

# A Defence of the Resemblance Meaning of ‘What it’s like’

RICHARD GASKIN

*University of Liverpool*

5 *richard.gaskin@liverpool.ac.uk*

It is often held to be definitive of consciousness that there is something it is like to be in a conscious state. A consensus has arisen that ‘is like’ in relevant ‘what it is like’ locutions does not mean ‘resembles’. This paper argues that the consensus is mistaken. It is argued that a recently proposed ‘affective’ analysis of these locutions fails, but that a purported rival of the resemblance analysis, the property account, is in fact compatible with it. Some of the implications of this argument are briefly explored: it is suggested that the meaning of ‘what it’s like’ does not, in itself, have any special bearing on consciousness, and that the implications for the so-called ‘hard problem’ of consciousness are deflationary.

## 15 1. Introduction

Following Thomas Nagel (1974), the notion of ‘what it’s like’ (WIL) has in recent years acquired a central place in discussions of consciousness: indeed, it is often held to define what it is to be a conscious subject that it is ‘like something’ to be that subject and held to define what it is to be in a conscious state that it is ‘like something’ to be in that state.<sup>1</sup> As part of this development, a consensus has grown up that ‘is like’ in relevant WIL locutions does *not* mean ‘resembles’.<sup>2</sup> A typical sentence that is held to resist the resemblance gloss would be

(1) There is something it is like for me to see red.

25 In a recent paper in this journal, Daniel Stoljar offers an ‘affective’ account of the meaning of WIL statements: according to his story (2016, pp. 1173, 1176), a statement of the form

(2) There is something it is like for me to  $\psi$ ,

<sup>1</sup> See, for example, Chalmers (1996, pp. 4–6, 103–4; 2010, pp. 5, 104); Byrne (2004, §1); Kirk (2005, p. 58; 2017, pp. 73, 77, 134, 141–2, 153–4); Stoljar (2016, pp. 1188–91). See further Farrell (2016, pp. 51–2).

<sup>2</sup> See, for example, Nagel (1974, p. 170 n. 6); Lewis (1999, p. 265); Lormand (2004, pp. 311–22); Hellie (2004, pp. 352–6); Snowdon (2010, pp. 17–21); Farrell (2016, p. 52); Stoljar (2016, pp. 1184–5); McGrath (2017, p. 13); cf. Hacker (2002, p. 161); Janzen (2011, p. 279).

where ‘ $\psi$ ’ is a suitable psychological predicate, means roughly:

(3) There is some experiential way in which my  $\psi$ -ing affects me.

In the course of his argument, Stoljar rejects a number of rival construals of the WIL locution: one of these is the resemblance analysis; another reads (2) as saying that my  $\psi$ -ing has some property. My main aim in this paper is to show that the consensus on the meaning of WIL locutions is mistaken: ‘what it’s like’ *does* mean ‘what it resembles’. But I shall also argue that, while Stoljar’s affective analysis of WIL talk is unsatisfactory, the ‘property’ account can be accommodated. Because WIL talk is central to contemporary discussions of consciousness, one would expect that a correct understanding of this talk would have important implications in the philosophy of mind. I think that this is indeed the case. I will mention what I think those (deflationary) implications are in my final section.

*Prima facie*, the idea that WIL locutions do not introduce the idea of resemblance ought to seem odd. Consider the following *wh*-questions:

(4a) Whom is *s/he* like?

(4b) Which one is it like?

(4c) Where is it (this place) like?

(4d) When is it (this time) like?

These questions all ask after resemblances. So when one comes to

(5) What is it like?

one would naturally expect it to follow suit. Against this expectation, the consensus says that (5) patterns not with (4) but with

(6) How is it?

which seems to ask for a specification, not of a resemblance, but of a property or properties. Supporters of the consensus appeal to the *OED*, which imputes a ‘special idiomatic force’ to (5), and states (Like, *adj.*, A.I.b): ‘The question *What is he* (or *it*) *like*? means “What sort of man is he?”, “What sort of thing is it?”, the expected answer being a description, and not at all the mention of a resembling person or thing’. Stoljar agrees: the question

(7) What is Chicago like?

is, he tells us, 'asking for the way that Chicago is; it would be acceptable to say in response that it is cold or sports-crazy' (2016, p. 1171).<sup>3</sup>

But this is too simple as it stands. Certainly (7) *can* be answered by a specification of properties, but it can *also* be answered by a comparison, as in:

(8a) (It is like) London,

(8b) (It is like) no other city you've ever visited.

Here the parentheses indicate optional material. Whether or not this material is included, the availability of (8)-style answers to (7) shows that 'is like' in (7) at least *can* mean 'resembles': for otherwise the questioner in (7) could in turn respond to (8) with 'I didn't ask you what (other cities) Chicago *is like*'. But I assume that this would be an unacceptable response. So, the *OED*'s 'not at all' goes too far. How should we react to the availability of these different styles of response to (7)? It seems overhasty to conclude, as Stoljar does, that cases like (7) show that 'there is a non-comparative use of "like" in English' (2016, p. 1171), and leave matters there: we need first to settle the relation between the two styles of answer to (7). Perhaps the comparative style of answer gives the logical form of the other; or perhaps the two are in some sense equivalent. *Ab initio* we would expect a resemblance account to enjoy semantic priority. After all, the idiomatic use that the *OED* claims for (5), if it really does involve a distinct sense, is a *development* from the historically original sense of the adjectival 'like', which introduces resemblance.<sup>4</sup> And indeed several writers who reject a resemblance analysis of 'like' in (5) concede that its *literal* meaning in (5) is comparative.<sup>5</sup> So we need to be open to the possibility that the so-called idiomatic use of (5) is explicable in terms of that original, literal sense. Perhaps the idiom does not after all import semantic novelty.

If 'like' in (5) and (7) does not introduce resemblance, what is its formal semantics? Stoljar does not declare on this, but refers us (2016,

<sup>3</sup> So too Hellie (2004, pp. 336–41, 357–8); Lormand (2004, p. 309).

<sup>4</sup> According to *OED* the relevant form of 'like' is derived from OE *gelic*, itself a descendant of Proto-Germanic *galiko-*, which is composed of *ga-* (= with, corresponding to Latin *con-*/*com-*) + *liko-* (= body, form; cf. OE *lic* = corpse, Lyke-wake Dirge, lych-gate, German *Leiche* etc.), so is precisely analogous to Latin *conformis*, and yields *gleich* in modern German. The leading definition reads: 'Having the same qualities or characteristics as some other person or thing; of approximately identical shape, size, colour, character, etc., with something else; similar; resembling; analogous'.

<sup>5</sup> Lewis (1999, p. 265); Lormand (2004, pp. 318–19).

p. 1171 n. 15) to two accounts in the literature, implying that either of these will do: Eric Lormand's predicator-functor analysis, and Benj Hellie's propredicate analysis. So, to test the credentials of a supposed non-comparative 'like', we need to examine these. I take them in turn.

## 5 2. The predicate-functor analysis

Lormand's account (in his 2004) runs, in outline, as follows. He first notes what he styles a grammatical oddity. In relevant WIL sentences, such as

(9) *S*'s having *M* is like something for *S*,

where *S* is a subject and *M* a suitable mental state, Lormand tells us  
10 that "‘something’ is best specified by predicates, not terms', but  
'predicates typically cannot grammatically follow “is like”' (2004,  
p. 309). For in answer to the question 'What is it like?', we typically  
respond with, say, 'fatiguing', not 'fatigue'; but 'is like fatiguing' is not  
15 regarded as grammatically acceptable, though there are, Lormand  
notes, exceptions, such as 'is like new'. He then suggests (ibid.,  
p. 310) that this 'oddity' is best explained by analysing (9) as

(10) For some *F*, *S*'s having *M* is *F*-like for *S*.

In Lormand's story, (10) is then further analysed as

(11) For some *F*, *S*'s having *M* has the appearance of being *F* for *S*,

20 which in turn gets analysed as:

(12) *S*'s having *M* itself perceptually appears some way to *S*.

In (10), 'like' is a functor, taking a predicate as input and outputting  
another predicate. Lormand argues that a resemblance construal of  
'like' in this role fails.

25 Before considering Lormand's argument on this last point, we may  
note two difficulties with the alleged grammatical oddity. First, the  
linguistic data do not unequivocally support Lormand's assertion that  
'something' in (9) is best specified by predicates, not terms. Just as we  
can respond to (7) with (8), so we can extend (9) to a discourse that  
30 incorporates a term specification of 'something' in (9). Suppose that  
*M* is the experience of seeing a brilliant red object. We might then say  
that *S*'s having *M* is like something for *S*, namely (like) (hearing) the  
sound of a trumpet.<sup>6</sup> Again, the bracketed expressions indicate op-  
tional additions, and they are independent options: you can have

<sup>6</sup> Cf. Locke *Essay* III, 4, 11 (1975, p. 425).

either, both, or neither. If 'like' is omitted (and whether or not 'hearing' is included), we have specification by term (noun phrase). Secondly, Lormand suggests that a resemblance understanding of (9) which respects the allegedly exclusive specification of 'something' by predicates requires us to take it as generalizing over ways of being, and this, he says, 'yields logical gibberish' (ibid., p. 320) when factored into the resemblance analysis of (9), namely:

(13) *S*'s having *M* is similar to some way for *S*.

But, even by his own lights, Lormand has underexpressed (13): if predicates genuinely introduce ways of being, as he told us, then we should have not (13) but

(14) *S*'s having *M* is similar to some way of being for *S*,

which makes good sense. We can in turn take (14) to be equivalent to:

(15) For some *F*, *S*'s having *M* is similar to *S*'s being *F*.

Again, this argument indicates that Lormand's starting observation, that 'something' in (9) is best specified by predicates, not terms—that is, noun phrases—is incorrect: 'a way of being' and 'ways of being' are noun phrases.

Suppose, now, that we provisionally accept Lormand's analysis of (9) as (10): what is his case against a resemblance construal of 'like' in (10)? It is that this construal does not make good sense of '[adjective]-like' in (10) because *S*'s having *M* 'may appear (to have the property) *F* but is hardly similar to (the property) *F*' (2004, p. 311). There are two problems with this argument. First, it assumes the correctness of Lormand's analysis of (10) as (11) (and of (11) as (12)); but as Hellie has shown (2007a, pp. 443–4, 452–9) we should reject that analysis. Secondly, and more importantly, it presupposes too simple a reading of (10). We can approach this point by noting that Lormand concedes (2004, p. 311) that a resemblance construal is correct for '[noun]-like' phrases, such as 'electron-like particle', which refers to (the property of being) a particle that resembles an electron. But of course some '[noun]-like' expressions do not unpack in that neat way. A spider-like motion is a motion that resembles not a spider but a spider's motion; the 'cat-like tread' of Gilbert and Sullivan's Penzance pirates is a tread that resembles not a cat but a cat's tread; Shakespeare's 'saint-like sorrow', 'Christian-like conclusion', 'swan-like end', 'soldier-like phrase', 'giant-like rebellion', 'infant-like abilities', 'priest-like fasts' (and Keats's 'priest-like task') all unpack in a similarly indirect way. So too, in the case of '[adj]-like' phrases, a grim-like smile (an

example that Lormand takes from the *OED*) is not a smile that resembles grim—that collocation of words is not even grammatical—or grimness, but a smile that resembles a grim smile. These examples all illustrate a standard brachylogy or ellipse: both these terms denote the omission from surface syntax of semantically significant elements; so, for example, at the end of Horace’s Postumus ode (II, 14), the wine that you have carefully hoarded only for your heir to squander is said to be ‘richer than pontifical feasts’, by which is meant ‘richer than <the wine drunk at> pontifical feasts’. Such brachylogy is common in natural language. Consider this conversation involving the term of crucial interest to us:

Me: Did you know David wrote over a hundred comments on a recent student essay?

You: That is so like David.

Here you mean that writing over a hundred comments resembles not David but certain other things he does. In general, ‘It is just/so like *S* to *V*’ means that *V*-ing resembles not *S* but *S*’s typical (or most interesting) actions. A resemblance construal of (10) may appeal to this kind of brachylogy and analyse it as saying that *S*’s having *M* resembles, not *F* or *F*-ness, but *S*’s being *F*, for some *F*. That is to say, (10) gets analysed as (15), with a suitable restriction on the domain of the quantifier. (I shall return to this point directly.) And it is worth pointing out, given that Lormand relies so heavily on the authority of the *OED*, that it offers resemblance readings of all three of ‘[adj]-’, ‘[noun]-’, and ‘[adv]-like’ phrases.<sup>7</sup>

Lormand also considers a resemblance construal of (9) that does not follow his own favoured analysis of it as (10), namely:

(16) *S*’s having *M* is similar to something for *S*.

He rejects this analysis on the basis that it does not make sense of the ‘optional detachability’ of ‘for *S*’ (2004, p. 318). The suggestion is that while

(17) For some *F*, *S*’s having *M* is *F*-like

is easily heard as abbreviatory of (10), and so derivatively of (9),

(18) *S*’s having *M* is similar to something

is not easily heard as abbreviatory of (16), and so of (9). That, Lormand says, is because (18) is too trivial (*ibid.*, p. 319; cf. pp. 311,

<sup>7</sup> Cf. Hellie (2007a, pp. 453–7).

320): everything is similar, in some respect, to everything else. An obvious reply to this is to say that resemblance claims are typically made with tacit domain restrictions, to be gathered from the context. Lormand concedes this point, but he thinks that no such restrictions are in play with 'something' generalizations (ibid., p. 319). However, that is untrue. If I have an unusual experience and say 'Hmm. This is similar to something—what?', I do not expect my interlocutor to reply 'Well of course it's similar to *something*: everything is. Indeed, everything is similar, at some level of abstraction, to *everything*'. Depending on the context of my remark, some restriction on the domain of the quantifier will be operative. What I probably mean is that this is similar to some experience that I have already had. So, with suitable domain restrictions in place there is no difficulty hearing (18) as abbreviatory of (16), and so of (9), if (9) is being analysed as (16).

(9) and (16) talk about a mental state's being like, or similar to, something *for S*. The discussion so far has presupposed that we understand this relativization. But do we? Surely, one wants to protest, something is either like something else or it is not, period. What can it mean to say that *x* is like, or similar to, *y*—*for S*? That seems to make as little sense—at least, as little obvious literal sense—as the assertion that the cat is on the mat *for me*, or that the pressure and volume of a gas are inversely related *for you*. Well, one general sense that can be extracted from locutions of the form 'For *S*, *p*' is '*S* believes that *p*'. Stoljar notes (2016, p. 1187) that we say things like 'To/for me, Alice was the best player', where we mean 'I think that Alice was the best player'. And, concerning the relativization in (16), Lormand writes, on the resemblance theorist's behalf, that 'presumably, to ask whether [certain mental states] are similar to something "for me" is to ask whether they are similar to something *to me*, that is, whether I believe they are similar to something' (2014, p. 319).

But the resemblance theorist should reject the proffered assistance, for whether an experience of mine is like or similar to something (else)—and that is what, according to this theorist, really concerns us in (9) (and (16))—is a quite objective matter, not to be settled by whether I believe it to be like something. There is no entailment in either direction between my believing an experience of mine to be like something and its actually being so. Indeed, I suggested above that, assuming Lormand's analysis of (9) as (10), (10) in turn should be analysed by the resemblance theorist as (15); and in (15) the relativization 'for *S*' that is found in (9) and (10) (and (14)) has disappeared, which might suggest, assuming that (15) is a good analysis of (10) (and

(14)), that it was doing no useful work in (9) and (10) (or (14)).<sup>8</sup> One possibility is this: the sense that we are trying to extract from (9) and (10) (and (14)), and which according to the resemblance theorist finds adequate expression in (15), has no role for the operator ‘S believes that ...’; assuming that the ‘for S’ relativizations in (9) and (10) (and (14)) can only deliver that operator if they have a semantic role to play at all, and since we do not want that operator, we simply ignore the ‘for S’ relativization in giving our analysis of (10) (and (14)) in (15). Against that reductive strategy, one might consider an alternative construal of the ‘is like something for me’ locution, as applied to an experience, according to which it involves the empiricists’ idea of an ‘inner theatre’, that is, the idea that my experience is itself the object of an inner gaze.<sup>9</sup> Indeed, Lormand tells us that he favours this model of experience, and he glosses ‘for’ in the relevant relativizations as ‘in the presence or sight of’ (2004, pp. 311–12, 322–3). So it might look as though the relativizations of (9), (14), and (16) commit us to that model. But that seems as unsatisfactory an upshot as finding a role in our analysis of WIL locutions for the ‘S believes that ...’ operator.<sup>10</sup>

I shall leave the matter there for the moment, in order to complete my discussion of the semantics of WIL expressions; but we shall return to the status of the relativization in sentences like (9) and (16) later. Our conclusion from the discussion of this section is this. There is nothing in the analysis of ‘like’, in relevant WIL locutions, as a predicate functor, as least as Lormand develops that analysis, that need worry resemblance theorists. Lormand favoured an analysis of (9) as (10); alternatively, (9) might be analysed as either (14) or (16). The resemblance theorist can accept any of these moves: if (9) is analysed as either (10) or (14), these in turn will be glossed by that theorist as (15). (15) and (16), whichever of these we end up with, need to be understood as subject to suitable domain restrictions; but that is a standard and unproblematic feature of the use of quantifiers in natural language.

### 3. The propredicate analysis

Hellie defends a view according to which, in statements such as ‘*e* is like something’, “like” is used as a device that syntactically transforms

<sup>8</sup> Cf. Stoljar (2016, p. 1194).

<sup>9</sup> See Hume, *Treatise* I, 4, 6 (1978, p. 253).

<sup>10</sup> On both these points I am in agreement with Stoljar (2016, pp. 1186–8, 1193–5).



a pronoun into a propredicate' (2007a, p. 447), and he applies the analysis in particular to the phrases 'like this' and 'like that'. He rejects a resemblance construal of such phrases for reasons that I will explain, and rebut, below (§4). Rather, on his account, 'like this', though syntactically complex in the relevant uses, is semantically simple, being employed in an equivalent way to the semantically simple propredicates 'thus' and 'so'. Accordingly, Hellie regards 'Cheney is like that' as a variant of 'Cheney is thus', and just as we can subject the latter to quantification into the 'thus' position, getting 'Cheney is somehow', so we can do the same with the former, getting 'Cheney is like something' or the raised 'There is something that Cheney is like'.

But a problem immediately strikes us: as the quantificational examples show—and in fact the syndrome recurs under all the routine syntactic operations—'like', unlike 'thus' and 'so', both of which move without leaving a trace, does not move. So whereas in the case of 'thus' and 'so' 'Cheney is thus/so' is transformed to 'Cheney is something' or 'Cheney is somehow', where no trace of 'thus' or 'so' remains, we find that 'Cheney is like that' becomes 'Cheney is like something' or 'There is something that Cheney is like', where 'like' stays put. So, also, we have 'What is Cheney like?'; 'Cheney is like *what?*'; 'Whittington knows what Cheney is like' (ibid., pp. 447–8). Similarly, when 'like' combines with noun phrases, or when it is suffixed to an adjective (or noun or adverb), it again stays put under quantification (and other syntactic operations): so 'Chicago is like London' and 'His smile was grim-like' become 'Chicago is like something/somewhere' and 'His smile was like something' respectively. And in the rare cases where 'like' combines with adjective phrases, we observe the same phenomenon: so 'The car is like new' is transformed to 'The car is like something'. Hellie's response to these syntactic phenomena is despairing: 'I don't know how to explain this syntactic behaviour... syntax does weird things sometimes' (ibid., p. 448). But that is clearly unsatisfactory. The explanation is obvious, and the despair unnecessary: 'like this', 'like that', and 'like something' are not only *syntactically* but also *semantically* complex. And the natural reaction to *that* fact is to say, with the resemblance theorist, that 'like' introduces resemblance. As Hellie admits, "'like" is a real syntactic stick-in-the-mud!' (ibid., p. 448): the obvious explanation of the syntactic behaviour of 'like' is that it carries a proprietary semantic payload, and the obvious content of that payload is *resemblance*.

Once we have got clear on this point, we may then, in accordance with the position I have already started to develop in response to

Lormand, say that the resemblance theorist is not forced to analyse the discourse

(19) A: What is Cheney like?

B: Secretive

5 either as being ill-formed or as representing *B* as propounding the absurdity that Cheney resembles the property of being secretive; we can construe it as saying that Cheney resembles *a secretive person* (in respect of secretiveness, we must add for full explicitness). For, again, you can reply to the question what N.N. is like by saying that s/he is  
10 like no one you've ever met before. The comparative approach harmonizes the various syntactically admissible responses to such a question by discerning brachylogy in answers that deploy a simple adjective phrase. How the brachylogy is filled out will depend on context. Hellie cites from the *OED* the sentence 'He refused to keep  
15 his royal promise; kings are like that' (ibid., p. 451). Here the resemblance theorist will, having glossed 'are like' as 'resemble', fill out the second clause by supplying unvoiced elements, to yield 'kings resemble <people who behave in> that <way>'. Hellie objects to discerning semantic complexity in 'like that' that we will then find ourselves  
20 saying 'not that "Lear is like that" predicates of Lear that he *has* the indicated feature, but that he has some feature *related* to the indicated feature' (ibid.). But this objection begs the question: the resemblance theorist will say that 'like that' introduces semantic complexity and that 'like' refers to the resemblance relation, so that we indeed want a  
25 relational analysis of 'Lear is like that'.

This takes us to an important taxonomic point. Hellie makes the objection I have just quoted in the context of criticising Lormand's predicate-functor analysis, and his criticism is designed to show that 'like' in 'like that' is not a predicate functor, on the basis that 'like that'  
30 is semantically simple. As we have seen, his argument fails: 'like that' is semantically complex. And in fact resemblance theorists are free to adopt the predicate-functor analysis; they will then offer a supplementary story that, departing from Lormand, preserves the resemblance meaning of 'like'. But they can *also* accept from Hellie that 'like that'  
35 is a propredicate. So the predicate-functor and the propredicate analyses are not, as Hellie supposes, *competitors*. They are not competitors because they are analysing different things: the predicate-functor view analyses just the word 'like'; the propredicate view analyses the whole expression 'like that' (and similar). It is entirely compatible with the  
40 propredicate view that, *within* the complex expression 'like that', 'like'

serves as a predicate functor. In fact, Hellie accepts that 'like [PRED]' is a predicate functor (*ibid.*, pp. 452–3): he just wishes to treat 'like this' and its congeners as idiomatic exceptions to the rule. I am suggesting that there is no need to make any exceptions. The upshot of the discussion of this and the previous section is that both the predicate-functor and the propredicate analyses, when rightly understood and applied to their appropriate analysanda, are *correct*; but neither threatens the resemblance construal of 'is like' in the relevant WIL constructions.

#### 4. Rebuttal of further arguments against the resemblance analysis

I have already (§2) considered, and rejected, Lormand's arguments again a resemblance construal of relevant WIL phrases. In this section, I shall look at a number of further arguments against this construal offered by Hellie (2004) and Stoljar (2016). Hellie's first argument goes as follows. Suppose someone utters

(20) My couch is coloured like this,

while demonstrating a red object. Hellie says that, in this context, (20) 'may just mean the same as' the non-comparative

(21) My couch is coloured red.

Hence, 'intuitively, "like" does not mean "similar to"' (2004, p. 353). But (20) and (21) do *not* mean the same: they have quite different semantics, on anyone's view. In context, of course, they may at some suitably abstracted level of semantic analysis be equivalent; but that does not show that 'like' does not mean 'similar to'. (Actually, it should be 'similarly to': see the next point.) Many pairs of sentences that are semantically quite distinct are, at an appropriate level of abstraction, equivalent in a given context.

Hellie's second argument is that 'coloured' can be followed by an adjective, as in (21), whereas 'My couch is coloured resembling that' is unacceptable; 'by contrast, [(20)] does not suffer from this problem' (*ibid.*, p. 353). But here we must observe that 'like this' often functions adverbially ('She ran like this'), and that is how it works in (20); similarly, in (21), 'red' is an adverb. In fact, in older English an explicit *-ly* adverb was often used in predicative constructions ('My couch is coloured redly'). Whereas now we tend to say 'it sounds odd', formerly

one could say ‘it sounds oddly’,<sup>11</sup> and we still say ‘it reads oddly’ and not ‘it reads odd’. So the right resemblance analysis of (20) will be ‘my couch is coloured similarly to this’.

5 Hellie’s third argument is that ‘to understand the utterance of [(20)], one’s audience must merely determine which colour one is talking about. By contrast, to understand an utterance of (for instance) “my couch is similar in colour to that material body”, one’s audience must (i) determine which material body in the region of ostension one is talking about, and (ii) determine what its colour is’  
 10 (ibid., p. 353). But to understand (20) you need to grasp *which coloured object* is the relevant relatum. The speaker cannot simply gesture at a disembodied colour; you cannot directly demonstrate the *universal* redness—which, being an abstract object, is in any case (*pace* Plato) not red. Even if there is a pure colour sample to hand, say in a decorator’s chart, the speaker must still demonstrate a physical object,  
 15 perhaps a piece of paper that is coloured in a particular way. So there is no such disanalogy between the cases as Hellie attempts to establish. Ostending a red colour-sample is less liable to be misunderstood than ostending an object that is only partly red and has other salient prop-  
 20 erties, but, as Wittgenstein taught us, even the most favourable conditions for ostension presuppose background conventions of understanding and so can, in easily imaginable circumstances, misfire. The point here is that in both of Hellie’s cases the audience must carry out both tasks (i) and (ii), however automatic either or both of these  
 25 tasks may, in context, be.

Hellie’s fourth argument is that ‘The couches here are red’ can be continued by ‘couches that are coloured like that please me’, where ‘like that’ functions anaphorically, replacing ‘red’ (ibid., p. 354). But note here again that ‘like that’ is adverbial: the anaphora is enthymematic, depending on an implied ‘redly’ contained in the covert step  
 30 ‘so the couches here are coloured redly’. Here too the resemblance theorist will analyse ‘like that’ as ‘similarly to that’, where ‘that’ demonstrates a property instance. The instance is spread over several couches in the example, but that (*contra* a suggestion of Hellie’s:  
 35 ibid., p. 353) is unproblematic. (Parallel remarks apply to a further example that Hellie introduces in which ‘like that’ is a bound variable: ibid., p. 354.)

Stoljar offers three arguments against the resemblance analysis of relevant WIL phrases. The first appeals to examples like (7) and his

<sup>11</sup> See, for example, Shakespeare, *The Tempest* 5.1.200.

gloss on them. According to the second, the question 'What is it like to see red?' does not call for resemblances, but for 'various properties that seeing red might have or might be associated with'. Thirdly, 'one can know what an event resembles without knowing what it is like ... For example, someone who has never eaten a peach and does not know what it is like may still know (for instance because they have heard it on good authority) that eating a peach resembles something else, for example, eating a nectarine' (2016, p. 1184).<sup>12</sup>

Take the third point first. Here Stoljar's assertion is simply question-begging, and the illustration that follows the assertion does not supply an independent argument for it. Against Stoljar, the resemblance theorist will say that, if you have eaten nectarines but not peaches and you ask 'What are peaches like?' or 'What is the taste of a peach like?', it *is* informative to be told 'They are like nectarines' or 'It's like the taste of a nectarine', where 'like' introduces resemblance. At least, these answers tell you *something* about what the taste of peaches is like. But do they tell you what it is *really* like? In the background of Stoljar's third argument there lies, of course, Frank Jackson's Knowledge Argument, according to which someone who has not eaten a peach does not know what it is *really* like—does not know it 'from the inside', or something of the sort. The idea is that there is some piece of *knowledge*—concerning the experiential content of tasting a peach—which someone only acquires on first tasting a peach. We need not rule here on the right response to the Knowledge Argument. Suppose that there is something that you cannot know unless you have tasted a peach. Still, you can discover what the taste of a peach *is like* by eating a nectarine. Or, if sheer propositional knowledge (of the sort that Jackson's Mary can acquire in her black-and-white room) suffices for the relevant knowledge, then you can gain knowledge of what the taste of peaches is like by being reliably informed that it resembles the taste of nectarines. To maintain that the *only* way you can find out what the taste of peaches is like is by eating a peach would be, in effect, to introduce a special, technical sense of 'what it's like'; but Stoljar tells us that he is not doing this.<sup>13</sup> For, as far as ordinary language goes, you *can* find out what peaches taste like by ascertaining resemblances.

Ordinary language licenses such utterances as 'If you want to find out what peaches taste like and have none to hand, try a nectarine.

<sup>12</sup> Cf. Snowdon (2010, p. 20); McGrath (2017, p. 13).

<sup>13</sup> Stoljar (2016, pp. 1183–4). See further §7 below.

That will tell you what peaches taste like, because the tastes are very similar'. To insist that only a peach can *really* taste like a peach, or even that only *this* peach can really taste like *this* peach, would be to misuse language. A thing cannot be *like itself* except in a degenerate sense. That is because *likeness* normally implies (numerical) *difference*.<sup>14</sup> When, in *Hamlet*, Marcellus asks of the ghost of old Hamlet 'Is it not like the King?' and Horatio replies 'As thou art to thyself' (1.1.61–2), Stephen Greenblatt remarks on the strangeness of Horatio's response, for 'Marcellus is not "like" himself; he *is* himself' (2001, p. 211). To suggest that an experience, or its content, is like itself *par excellence* and only in a secondary sense like other things would get matters back to front: if 'like' is being used ordinarily, the primary sense in which it is like something to have an experience is that in which having that experience is like something *else*—perhaps, though not necessarily, having another experience.

Here it might be objected that being informed that peaches taste like nectarines can only tell you what it is like to taste a peach if, as the objector might put it, you already know what nectarines taste like: are we to apply the resemblance analysis again to this latter 'like'? Well, why not? The worry, I suppose, is that a subject might come to know a pair or even a whole network of resemblances without having had any of the relevant experiences, and so without knowing, as the objector wants to say, what any of the experiences *are like*. But here we need to recall just how much is presupposed to knowledge of empirical resemblances: in order to know that peaches taste like nectarines, or that violas look like violins but are slightly bigger,<sup>15</sup> one needs to grasp the relevant concepts, and it is plausible that only a subject who enjoys a considerable range of experiences—and also, plausibly (at least in culturally more sophisticated cases, such as the viola/violin one), who possesses language—will do so. Experience and language tie down what otherwise threatens to be a free-floating network of resemblances in such a way that we can be sure that the subject really does know what is going on. We should not forget how much is built into, say, the knowledge of what a violin is: no doubt it would be wrong to insist that anyone who can be credited with that knowledge must have seen a violin, or a picture of one, but, if not, many similar

<sup>14</sup> Poole (1987, p. 116). Frege famously made a similar point about correspondence (1918–19, p. 60).

<sup>15</sup> McGrath's example (2017, p. 13).

experiences and much contextual cultural knowledge will be requisite. Once all that empirical richness is in place—so once we can be sure that the subject really does know what a violin is—then it *is* informative for him or her to be told that violas look like (resemble) violins. That tells the subject *what a viola is like* (at least visually). So the resemblance theorist can insist that, given appropriate experiential and linguistic context, knowledge of resemblances *does* tell you what something is like. Putting it the other way around, the difficulty for the resemblance theorist's opponents is this: the more they set things up so that knowledge of resemblances does *not* suffice for knowing what something is like, the more they risk undermining the premisses of their case, because the less likely they make it that the subject satisfies the requisite conditions for grasping the concepts involved in knowing the resemblances in the first place.<sup>16</sup>

Stoljar's first two arguments are variants of one another: in effect they both say that the 'What is it like?' question asks after properties, not resemblances.<sup>17</sup> And it must be agreed that, as we have already noted, such questions are often answered by specifying properties. When I said above, in connection with the dialogue (19), that the answer given there is, in effect, an abbreviation of '<Cheney resembles a> secretive <person>', with the unvoiced elements here rendered explicit, someone might object that, though the reply *is* elliptical, the completion of the ellipse should be: '<Cheney *is* a> secretive <person>'. And this is the right place to mention a fact stressed by Stoljar and others, namely, that English WIL expressions are often translated into other languages as 'how' expressions, so in effect (5) comes out as (6): in German, for instance, 'What is it like?' is rendered 'Wie ist es?', which seems to ask after properties rather than resemblances.<sup>18</sup> How should the resemblance theorist tackle this point?

<sup>16</sup> As I have implied in this discussion of Stoljar's third argument, the resemblance theorist need take no particular line, for these purposes, in response to the Knowledge Argument. Whichever way the verdict falls on whether Jackson's Mary needs to *see red* things in order to grasp the concept *red* (see here Daly (1998); Kirk (2005, p. 65)), she can be told what red things *are like* by being given resemblances. (Of course, if we insist that in order to grasp the concept *red* Mary must see red things, then we will also say that in her black-and-white phase she cannot understand—or perhaps can only partially understand—these explanations.)

<sup>17</sup> So too Hellie (2004, pp. 337–8).

<sup>18</sup> Stoljar (2016, p. 1170); Snowdon (2010, p. 21).

## 5. The property and affective accounts

Despite his appeal to properties in his first two arguments against the resemblance theory, Stoljar rejects what he calls the ‘property’ account of WIL locutions, which analyses

5 (1) There is something it is like for me to see red

as

(22) For some *F*, my seeing red is *F*.

Stoljar objects (2016, p. 1185) that, while (1) entails (22), the reverse entailment does not hold: my seeing red can have all sorts of properties (including contextually salient ones) without its being *like* anything for me. This recalls Lormand’s objection to (16), and the response here can be the same: we operate with domain restrictions. That applies not only to (22) but also to what would be the resemblance theorist’s analysis of (1), namely:

15 (23) There is something it resembles for me to see red.

Stoljar considers this response (*ibid.*, pp. 1185–6) and replies that, whereas sentences like (1) are routinely used, both on their own and in inference, to report certain quite narrowly individuated facts about experience, (22)—and presumably he would extend the point to (23)—is not. He claims that, by contrast with the property account (and presumably also the resemblance account), his analysis explains these routine uses. But the explanation he offers (*ibid.*, p. 1178) is bogus: it is achieved by building into his account the point that WIL sentences are *stereotypically* used with a narrowly experiential content. (This is Stoljar’s hypothesis (H2): *ibid.*, p. 1176.) But this tactic simply shifts the burden of explanation. Neither Stoljar’s account nor the property and resemblance accounts explain *why* the experiential use of WIL sentences is stereotypical; they merely take it for granted. But then, should *any* of the accounts we are considering have to explain this fact of usage? The property and resemblance theorists can agree that ‘What it is like’ typically means ‘What it is experientially like’, and can say that the domains of the relevant properties and resemblance relations in (22) and (23) have to be restricted accordingly. It is just a fact of usage that ‘What it is like’ has come to be employed in this way, at least in some contexts (though not, by any means, in all),<sup>19</sup> and the property and resemblance theorists are no

<sup>19</sup> Snowdon (2010, pp. 23–4).



worse off in point of modelling that fact in their stories than is Stoljar with his affective account.<sup>20</sup>

So the property account is not dismissed by Stoljar's argument. Before we deal with the challenge that it poses to the resemblance theory, let us examine his affective account in more detail. I said (S<sub>1</sub>) that Stoljar analyses (2) as (3), but that construal in effect depends on taking 'for me' in (2) to do double duty, so that (2) is construed as:

(24) For me to  $\psi$  is like something for me.

Stoljar argues that such 'double indexing' is forced by the fact that the indices can be distinct: for you to have toothache can be like something for me—upsetting, say, or wearing.<sup>21</sup> We may label the two index positions in (24) 'primary' and 'secondary': the primary index is the subject of the psychological infinitive verb in the 'to  $\psi$ ' position (what linguists call 'PRO'); the secondary index specifies the subject for whom  $\psi$ -ing is like something. Now about the intelligibility of the primary index there can, I suggest, be no quibble: we need a subject of the psychological state which is being said to be like something. But how about the secondary index? Here, picking up a point that was anticipated at the end of S<sub>2</sub>, the resemblance theorist will interpose: what can it mean to say that (factoring in that theorist's analysis) it resembles something *for me*, for you (or me) to  $\psi$ ? Does it mean that *I believe that* it resembles something for you (or me) to  $\psi$ ? That was one possible meaning that (in effect) we extracted from the secondary index in the earlier discussion. But I also suggested that an allusion to the subject's beliefs in an analysis of WIL locutions is, at least in general, irrelevant. Stoljar's glosses are little help in making sense of the secondary index. At one point, as an analysis of 'There is something it is like to have a toothache', we are offered (2016, pp. 1165, 1171–2):

(25) There is some way such that for  $y$  to have a toothache is that way to  $x$ .

Here 'y' occupies primary position and 'x' secondary, and Stoljar distinguishes the indices prepositionally, using 'for' for primary and

<sup>20</sup> Stoljar asserts (2016, p. 1177) that the verb 'affect', deployed in his analysis of WIL sentences, is itself stereotypically used with a narrowly experiential content. Even if that were right, it would not favour his analysis, for the reason given in the text. But it is surely wrong: 'affect' is used in connection with causal transactions of all sorts, and in purely physical contexts it is generally (*contra* what Stoljar perhaps implies, *ibid.*) used without anthropomorphism.

<sup>21</sup> Stoljar (2016, pp. 1166–9); cf. Lormand (2004, p. 307).

‘to’ for secondary position. But what does ‘that way to  $x$ ’ mean? What is *a way to  $x$* ? Stoljar does not, of course, mean a route to  $x$ . But nor does he mean that  $x$  *believes* that  $y$ ’s having a toothache is such-and-such. In cases of co-indexing (so where we have, in (25) for example,  $x = y$ ) that would entail an analysis of WIL expressions in higher-order-consciousness terms, but Stoljar rejects such accounts (in my view rightly).<sup>22</sup>

An alternative reading of the ‘for me’ relativization in (24) in its secondary occurrence, briefly mentioned at the end of §2, is one that imports the empiricists’ ‘inner theatre’. This approach also finds higher-order consciousness present in WIL contexts, but in a different way. It is just as unacceptable, for reasons brought out by the later Wittgenstein, who inveighed against the temptation to ‘detach the colour-*impression* from the object, like a membrane’,<sup>23</sup> the point being that we should not think of a seen colour as lifting off from the worldly object like a peregrine ghostly film that becomes the object of a purported inner gaze. Despite Wittgenstein’s warning, the mistake is constantly encountered in contemporary discussions. Here is a typical example, from David Chalmers:

Mary looks at a red apple, and visually experiences its colour. The experience instantiates a phenomenal property  $R$ , which we might call phenomenal redness. It is natural to say that Mary is having a red experience, even though of course experiences are not red in the same sense in which apples are red. (2004, p. 270)

The concession made at the end of this passage does little to repair the damage done by the introduction of the spurious notion of ‘phenomenal redness’. For it is not just that experiences are not red in the sense in which apples are red: experiences are not red *at all*. That involves a category error.<sup>24</sup> If we nevertheless allow ourselves to speak of ‘red

<sup>22</sup> Stoljar (2016, pp. 1193–5); cf. Byrne (2004, §5); Kirk (2017, p. 165). Hellie also glosses the WIL locution as involving ‘a property which is “to” or “for” the subject of experience’ (2007b, p. 295); but it is clear that, unlike Stoljar, Hellie does intend this to be understood as importing higher-order consciousness (pp. 301–2); cf. Weisberg (2011, pp. 426–7). Stoljar’s attempt to supply other cases where ‘to’ means something like ‘affects’ (2016, p. 1174) do not help: ‘awful (to)’ and ‘destabilizing (to)’ are straightforward cases of expressions requiring relativization and make sense as such, whereas ‘a way to’ and ‘a property to’, in the sense in which these phrases are intended, look thoroughly opaque.

<sup>23</sup> Wittgenstein (1958, I, §276, emphasis in original); cf. Byrne (2002, p. 124); Kirk (2005, p. 55).

<sup>24</sup> Byrne (2002, pp. 125, 128).

experiences', we can *only* intelligibly mean by that experiences of red things in the world.

Chalmers' mistake is brought out by his (perfectly consequential, given his starting point) statement that there are possible worlds  
 5 where red objects do not cause red experiences (ibid., p. 271). Now, in an ordinary sense such worlds are indeed abundant: they include, for example, worlds containing red objects but no light, or no subjects of experience, or no subjects with the right perceptual equipment, and so in which red objects are never observed. But this is not what  
 10 Chalmers means: he aims to allow for the conceivability of a world where all the right conditions are in place, but where red objects cause not red but, say, green experiences. But this is contradictory. A red object is an object which is such as to look red to suitably equipped observers in suitable viewing conditions.<sup>25</sup> So you cannot have red  
 15 objects that do not look red when those conditions are met. Neither the possibility that red objects might not look red, under the appropriate conditions, but rather (say) green, nor the possibility that green objects might look red (under the same conditions) makes sense. Red objects might of course have looked green if they had been painted  
 20 green, but they could not have looked green while continuing to be red: if an object *looks* green (under the right conditions) then it *is* green. And contrariwise: if an object *is* green, then (under the right conditions) it *looks* green. It follows that the content of 'looks red', say, is fixed by ostending red objects in the world. If you ask me what  
 25 it is for something to look red, I respond not by inviting you to introspect, but by demonstrating (say) a British post box and saying 'It looks *that* colour' (that is, it looks the colour which that object *is*).<sup>26</sup> Chalmers is, in effect, detaching looking red from being red, which is precisely the membrane fallacy identified by Wittgenstein in the passage quoted above. Put another way, it is a version of the homunculus  
 30 fallacy that is often diagnosed in the old sense-datum theory of perception. So in the sense in which he intends the claim it is not, as Chalmers thinks, 'conceivable that when looking at red things, such as roses, one might have had the sort of colour experiences that one in  
 35 fact has when looking at blue things' (1996, p. 5). Note that merely to say that something looks red does not *itself* commit the membrane fallacy. Something *looks* red just if it seems to the subject, accessing the object in the usual visual way, as though it—the seen object—*is* red.

<sup>25</sup> McDowell (1998, p. 133).

<sup>26</sup> Cf. Wittgenstein (1958, I, §275).

There is no call to suppose, incoherently, that a ‘look’, itself coloured redly (despite being unable to reflect light), detaches itself from the red object, becomes the object of an inner gaze (with consequent looming regress),<sup>27</sup> and can vary independently of real redness.<sup>28</sup>

5 We have not so far managed to make good sense of the secondary index in statements formed on the model of (24): I have rejected analyses of it as introducing higher-order consciousness, whether in the form of a belief operator or an inner gaze. In Stoljar’s affective account of WIL locutions, we are told that the secondary index—‘to (x)’, in his terminology—means ‘affects (x)’ (2016, p. 1173). Here I think Stoljar faces a dilemma. There are two things one might mean by ‘affects’, as it figures in (3), neither satisfactory.<sup>29</sup> Sometimes Stoljar seems to mean that there is some *further* respect in which my  $\psi$ -ing affects me—that is, some respect beyond the sheer fact of my  $\psi$ -ing, something *caused* by the  $\psi$ -ing. Thus, he glosses ‘affect’ in the required sense as ‘influence’ and ‘bring about a change or condition in’: examples given include ‘Having the disease affects you by making you very tired’ (ibid., p. 1173). In this sense my having toothache might affect me by, for example, making me irritable. But that reading of (3) obviously renders it useless as an analysis of (2), because it changes the subject and introduces irrelevancy: an experience may be like something without affecting me in that ulterior sense. At other times Stoljar means that the affecting alluded to in (3) just is the experience of  $\psi$ -ing itself: you are affected by your  $\psi$ -ing simply by virtue of the fact that you are  $\psi$ -ing. This is what Stoljar has in mind when he comes to analyse a sentence such as (1): here we are told, in effect, that I am affected by my experience in the sense that I *feel* something when I see red, where ‘feel’ is interpreted broadly so as to apply to me, in this instance, by virtue of its *seeming* to me that things are a certain way (ibid., p. 1181).<sup>30</sup> But all there is to say in general about how things seem to me when I see red is that (it seems to me that) I see red. So we have been taken round in a circle: my seeing red affects me just in the sense that I see red. This is objectionable for a reason that is opposite to the reason that told against the first analysis: there the analysans

<sup>27</sup> Cf. Kirk (2017, pp. 121–2).

<sup>28</sup> As, for instance, in Chalmers’ ‘inverted Mary’ scenario (2004, p. 275).

<sup>29</sup> The ambiguity may be intentional on Stoljar’s part (cf. 2016, p. 1190 n. 31), but he does not see that it presents him with the dilemma I identify here.

<sup>30</sup> Note that this gloss appears to be inconsistent with Stoljar’s rejection of the ‘operator’ account of WIL locutions (2016, pp. 1186–8).

went off topic; here it travels no distance from the analysandum. (Alternatively, one might, on this horn of the dilemma, detect the membrane fallacy: when *I* see red, an inner *homunculus* sees red.)

## 6. Properties and resemblance

5 The affective account is unsatisfactory, then, but the property account is still standing as an apparent rival to the resemblance story that I have been canvassing hitherto in this paper: we might wonder how we should adjudicate between them. But here I propose that we need not choose: we can have both. The key to achieving this harmony is the  
10 reminder that properties just *are* resemblances.<sup>31</sup> My suggestion is this. If we ask what is meant by the statement that it is like something to see red, the answer, in the first instance, is that what is meant is that seeing red is similar to something, given a suitable restriction on the domain of the quantifier. That meaning is guaranteed, I have argued, by the  
15 fact that, however we may go on to use WIL expressions, we can always recover their original, literal meaning: the question 'What is it like?' can always be answered in the format 'It is like (= resembles) such and such'. But we can group together things that are similar in some respect, and when we do that we get a property, for a property is  
20 just the shared respect which collects similar things. In view of that, it is no surprise that WIL locutions may be heard as compendiously introducing properties; so we arrive at a developed position in which WIL questions, for example, can be answered by one-word property specifications. Recall (19): 'What is Cheney like?'; 'Secretive'. Here we may spell  
25 out the derivation of the answer as follows: Cheney *resembles* a secretive person in respect of secretiveness, and so *is* secretive. The one-word answer is convenient; but the analysis gives its semantic aetiology.

I do not see the fact that the resemblance theorist appeals to covert elements here as a serious cost: semanticists regularly adduce unvoiced  
30 elements in the analysis of ordinary discourse; and, however we analyse the answer in (19), we will have to discern covert elements. The property account must do so too. But if the property account were advanced as a *rival* to the resemblance account it would have to be rejected, since it does not explain why a resemblance meaning is  
35 always recoverable from the WIL locution. What I suggest, however, is that the property account is best seen not as a competitor, but as a supplement, to the resemblance theory. At a certain stage of semantic

<sup>31</sup> Quine (1969, pp. 116–7).

development, WIL locutions can be heard as asking after, or specifying, properties. The original resemblance meaning of the locution is there in the semantic background, is recoverable, and is in any case consistent with specification by property, given that properties are themselves just resemblances.

5 Where does that leave the secondary index in sentences like (24)? I have rejected a higher-order-consciousness gloss on this; but (24) appears to make sense, so what does the second ‘for me’ mean? I suggest that it can be read as restricting the range of the quantifier to resemblances, and so to properties, that concern me in the right way, whatever that way is. (The context will decide.) That implies the rejection of an idea mooted at the end of §2, namely, that the secondary index is semantically inert. We can still say, as we wanted to say there, that (15) is the right analysis of Lormand’s (10), itself an analysis of (9); but we will need to add that the ‘for *S*’ relativization in (10), though superficially disappearing in (15), is really present as a restriction on the domain of the quantifier. We noted in the earlier discussion that this restriction must in any case be presupposed in (15) and (16), and the point being added here is just that the ‘for *S*’ relativization, though not explicit in (15), in fact gets taken up into, and so is tacitly present in, that domain restriction. (Its explicit presence in (16), which I allowed as an alternative analysis of (9), should also be construed in this sense, namely, as a domain restrictor.) Lormand remarks that the pain in my leg is like something only for me—not for my leg or my medicine cabinet (2004, p. 311). What that comes to is that *I* (not my leg or medicine cabinet) have the experience and that the experience’s resemblances include ones that bear relevantly on and characterize—perhaps essentially characterize—my feeling. To make this clear, we can redraft (15) as: For some *relevant F*, *S*’s having *M* is similar to *S*’s being *F* (where ‘relevant’ codifies the need for suitable domain restriction on the quantifier).

## 7. Implications

‘Essentially characterize my feeling’: but not *intrinsically* characterize it—at least not if ‘intrinsic’ is being used as a stand-in for a construal of WIL talk that is resistant to an analysis such as I have given in resemblance (and property) terms. That is often how ‘intrinsic’ is used in the contemporary consciousness room: so, for example, Chalmers speaks of ‘the intrinsic nature of a sensation of red’ (2010, p. 22),<sup>32</sup>

<sup>32</sup> Cf. Weisberg (2011, p. 410).

where he means something much more ambitious than merely that, say, a sensation of red is more like a sensation of orange than it is like a sensation of green.<sup>33</sup> He means to advert to the 'hard problem of consciousness', according to which what experiences are like allegedly  
 5 eludes physical and functional description. But if what an experience is like is just a matter of what it (relevantly) resembles, and so, derivatively, of what (relevant) properties it enjoys, such ambitions are quite unwarranted. What we find time after time in discussions of consciousness is that the WIL locution plays a pivotal role in setting up  
 10 a supposed 'hard problem' of consciousness; but we are told next to nothing about what it means.<sup>34</sup> Many writers are happy to say that they are treating the WIL locution as a technical term.<sup>35</sup> I share Stoljar's aversion to this approach (§4).<sup>36</sup> After all, if 'what it's like' is a technical term, it needs to be properly defined. But such a definition is never given: glosses along the lines of 'the phenomenal content/intrinsic character of experience' take us nowhere, for they are themselves technical terms (indeed much more obviously technical than 'what it's like'), and simply re-pose the demand for a definition in plain terms. When we spell out in plain terms what 'what it's like'  
 20 actually means, as I have tried to do here, we find that there is nothing in that meaning to support the subjectivist and dualist superstructure that is often placed upon it: the 'further phenomenon'<sup>37</sup> that is supposedly inexplicable by ordinary science and is said to be grounded in the what-it's-likeness of experience presupposes an understanding of  
 25 WIL talk that is based on a simple linguistic error.

Stoljar begins his paper (2016, p. 1162) by noting how surprising it is that WIL talk has been commandeered for consciousness purposes, given that there is nothing on the surface of the phrase 'what it's like' to alert you to any such connection. I have in effect been arguing here  
 30 that there is nothing more to 'what it's like' than what appears on the surface. The phrase *looks* as though its meaning is comparative, and its meaning *is* comparative. It means *what it resembles*—no more than that. The phrase can be used in a context where consciousness is at issue, of course, as it can be used in many other contexts; but when we

<sup>33</sup> Cf. Byrne (2004, §5).

<sup>34</sup> For an illustration, see Chalmers (2010, Essay 1).

<sup>35</sup> See, for example, Lewis (1999, p. 326); Janzen (2011, p. 279).

<sup>36</sup> See also Hellie (2004, pp. 336–9); Farrell (2016).

<sup>37</sup> Chalmers (2010, p. 33).

have to do with what an experience is like, the burden of introducing consciousness will be borne by the explicit or implicit restriction of the resemblances (and so properties) that interest us to experiential ones, and not by anything in the semantics of the phrase ‘what it’s like’. In itself, WIL talk tells us nothing about consciousness, and so *a fortiori* nothing about a ‘hard problem’ of consciousness. When we restore its true comparative sense to ‘what it’s like’, we see that efforts to drag-oon WIL locutions into the service of a preterscientific solution to an alleged ‘hard problem’ of consciousness fail, and glosses of ‘what it’s like to see red’ in terms of ‘the intrinsic nature of a sensation of red’ simply risk tipping over—and in the case of some contemporary writers do tip over—into the membrane fallacy.

Again, when Chalmers writes that ‘there is nothing it is like to be a zombie’ (2010, p. 107), we can reply that in the ordinary and correct sense of ‘is like’ this is just false. Zombies, as they are characterized in the philosophical literature, are very similar to us indeed—in fact, they are identical to us in all respects except, supposedly, one.<sup>38</sup> Hence there are plenty of things it is like to be a zombie; and the one respect in which, allegedly, it is like nothing to be a zombie cannot, it would appear, be stated. At least, WIL locutions will not do the trick: for zombies are physically and functionally indistinguishable from normal human beings, which means that their experiences will boast the usual range of resemblances (properties), in fact, exactly the same ones as ours do. Moreover, zombies will be able to verbalize, just as we can, what those resemblances (properties) are: they will say things like ‘Ouch! This toothache is really agonizing’, ‘Look how blue the sky is!’, ‘Hmm, this feeling I’m getting is like something—what? Oh yes, it’s like that extraordinary feeling I remember from last year’, and so on.<sup>39</sup> There is

<sup>38</sup> Farrell (2016, p. 52).

<sup>39</sup> Although I do not have space to pursue the point here, it is worth noting that, as well as the hard problem of consciousness, another potential casualty of a correct understanding of WIL locutions is the popular idea that feelings are inexpressible. At any rate, the correct understanding of WIL locutions does not support that idea; as far as the resemblance theory is concerned, resemblances may all be perfectly expressible in words. Note that whether an experience can be put into words is a quite distinct issue from whether those words can only be understood by someone who has had similar experiences. (For the confusion, see Hellie (2004, p. 342).) The latter may hold of a range of statements—indeed, it is probably true that understanding most ordinary statements requires experience—but it does not follow that the relevant experiences cannot be verbalized. Putting experiences into words often involves demonstrating something in the world, but that is unsurprising given that experiences often have (as we noted in the case of colour experience) worldly things as their content: cf. McDowell (1996, pp. 56–7); McGrath (2017, p. 12). If I say ‘My experience is/was of *that*



nothing left for their experiences to be *like* that escapes the net of ordinary talk about ordinary resemblances, and if we try to force the WIL locution to express that supposed but non-existent remainder, we risk finding ourselves, once more, in the grip of the membrane fallacy and the illusion of a Cartesian inner theatre.<sup>40</sup>

## References

- Byrne, A. 2002: 'Something about Mary', *Grazer Philosophische Studien* 63, 123–40.
- 2004: 'What phenomenal consciousness is like', in R. Gennaro (ed.), *Higher-Order Theories of Consciousness* (Amsterdam: Benjamins), 203–25.
- Chalmers, D. 1996: *The conscious mind: in search of a fundamental theory* (Oxford: OUP).
- 2004: 'Phenomenal concepts and the Knowledge Argument', in P. Ludlow et al. (ed.), *There's Something About Mary: Essays on Phenomenal Consciousness and Frank Jackson's Knowledge Argument* (Cambridge, Mass.: MIT), 269–98.
- 2010: *The Character of Consciousness* (New York: OUP).
- Daly, C. 1998: 'Modality and acquaintance with properties', *Monist* 81, 44–68.
- Farrell, J. 2016: "'What it is like" talk is not technical talk', *Journal of Consciousness Studies* 23, 50–65.
- Frege, G. 1918–19: 'Der Gedanke: eine logische Untersuchung', *Beiträge zur Philosophie des deutschen Idealismus* 2, 58–77.
- Greenblatt, S. 2001: *Hamlet in Purgatory* (Princeton: Princeton University Press).
- Hacker, P. 2002: 'Is there anything it is like to be a bat?', *Philosophy* 77, 157–74.
- Hellie, B. 2004: 'Inexpressible truths and the allure of the Knowledge Argument', in P. Ludlow et al. (ed.), *There's Something About Mary: Essays on Phenomenal Consciousness and Frank Jackson's Knowledge Argument* (Cambridge, Mass.: MIT), 333–64.
- 2007a: "'There's something it's like" and the structure of consciousness', *Philosophical Review* 116, 441–63.

---

colour', demonstrating a colour sample to hand, I have verbalized my experience (with the help of the world).

<sup>40</sup> I am grateful to the journal's two readers for helpful and incisive comments on previous drafts of this paper.

- 2007b: ‘Higher-order consciousness and higher-order acquaintance’, *Philosophical Studies* 134, 289–324.
- Hume, D. 1978: *A Treatise of Human Nature*, ed. P. Nidditch (Oxford: Clarendon).
- 5 Janzen, G. 2011: ‘In defence of the what-it-is-likeness of experience’, *Southern Journal of Philosophy* 49, 271–93.
- Kirk, R. 2005: *Zombies and Consciousness* (Oxford: Clarendon).
- 2017: *Robots, Zombies, and Us* (London: Bloomsbury).
- Lewis, D. 1999: *Papers in Metaphysics and Epistemology* (Cambridge: CUP).
- 10 Locke, J. 1975: *An Essay Concerning Human Understanding*, ed. P. Nidditch (Oxford, Clarendon).
- Lormand, E. 2004: ‘The explanatory stopgap’, *Philosophical Review* 113, 303–57.
- 15 McDowell, J. 1996: *Mind and World* (Cambridge, Mass.: Harvard University Press).
- 1998: *Mind, Value, and Reality* (Cambridge, Mass.: Harvard University Press).
- McGrath, M. 2017: ‘Knowing what things look like’, *Philosophical Review* 126, 1–41.
- 20 Nagel, T. 1974: ‘What is it like to be a bat?’, in *Mortal Questions* (Cambridge: CUP, 1979), 165–80.
- Poole, A. 1987: *Tragedy: Shakespeare and the Greek Example* (Oxford: Blackwell).
- 25 Quine, W. V. 1969: *Ontological Relativity and Other Essays* (New York: Columbia University Press).
- Snowdon, P. 2010: ‘On the what-it-is-likeness of experience’, *Southern Journal of Philosophy* 48, 8–27.
- Stoljar, D. 2016: ‘The semantics of “what it’s like” and the nature of consciousness’, *Mind* 125, 1161–98.
- 30 Weisberg, J. 2011: ‘Misrepresenting consciousness’, *Philosophical Studies* 154, 409–33.
- Wittgenstein, L. 1958: *Philosophical Investigations*, tr. G. Anscombe (Oxford: Blackwell).