

CHAPTER 6

Dual-Content Theory: the Explanatory Advantages

Over the course of the last three chapters various different arguments have been sketched in support of my dual-content theory of phenomenal consciousness, as well as against a variety of opponents (first-order representational theories, inner-sense theories, and actualist higher-order thought theories, in particular). Likewise in Carruthers (2000) multiple lines of argument were presented, both in support of my own views, and against a number of different opponents. Even in the chapter of that book where the main argument against first-order theories was laid out (chapter 6), the lines of reasoning were complex, with different real or imaginary opponents being in my sights at various different points throughout it, and frequently requiring reference to the arguments of earlier and later chapters as well (especially chapters 4, 7, and 9). Within this welter of different arguments it isn't easy to keep one's eye on the ball. What I propose to do in the present chapter is to lay out the *main* line of reasoning supporting dual-content theory as against what I take to be the strongest, most plausible, form of first-order theory. I shall also take up the question that has now begun to be broached in Chapter 5 of the present volume, which is whether a first-order theorist can *also* accept the higher-order analog contents postulated by dual-content theory, and can thereby obtain all of the advantages of the latter while retaining their distinctively first-order account of phenomenal consciousness itself.

1. SOME BACKGROUND

One of my arguments against first-order theories of phenomenal consciousness (of the sort espoused by Dreske, 1995, and Tye, 1995, for example; see Chapters 2 and 3 for explanation and discussion) has been that such theories can't account adequately for the distinction between conscious and non-conscious perceptual states. (Such a distinction is warranted by the two-visual-systems hypothesis of Milner and Goodale, 1995; see also Jacob and Jeannerod, 2003; see Chapters 4 and 11 of the present volume for elaboration.) But there is a version

of first-order theory that can easily *seem* invulnerable to that argument.¹ Let me explain.

I have in mind someone who accepts the dual-systems theory of vision, and who accepts that while the perceptual states produced by the ventral/temporal system are phenomenally conscious ones, the perceptual states produced by the dorsal/parietal 'how-to' system aren't. But it might be felt that this is just the result that someone who endorses a version of Tye's (1995) first-order theory of phenomenal consciousness might predict. For according to that account, phenomenally conscious experiences are those that are poised to have an impact on the belief-forming and conceptual-reasoning systems of the agent in question. And the outputs of the dorsal/parietal system *aren't* so poised (rather, they are made available for on-line fine-grained control of movement). Only the outputs of the ventral/temporal system are. So such a first-order theory will predict that the dorsal/parietal perceptual states won't be phenomenally conscious, whereas the ventral/temporal states will be phenomenally conscious—just as we seem to find to be the case.

All this is perfectly true; and it is a possible combination of views that I addressed at some length in Carruthers (2000). My objection is that it isn't *explanatory*. Each of the dorsal/parietal and ventral/temporal visual systems produces analog intentional contents, representing distal properties of the environment (motion, shape, color, and so forth). So it is mysterious why the availability of such contents to conceptual thought and reasoning should transform them into states with the distinctive properties of phenomenal consciousness, whereas the availability of such contents to guide movement shouldn't.

To this it can be remarked that consciousness is mysterious in any case; hence the mystery in question is only to be expected; hence there is no objection to a first-order theory here. But this would be to miss the *whole* point, which is that representational theories of phenomenal consciousness are in the business of proffering reductive explanations of their target. They are in the business of *explaining* the distinctive features of phenomenal consciousness, not just in finding natural properties that happen to be co-extensive with those features. Nor is it adequate for them merely to postulate a brute identity between phenomenal consciousness, on the one hand, and the properties that figure in the theory (availability to first-order conceptual thought, as it might be), on the other. For an identity, even if true, can only count as a successful reduction when we can use the properties proposed in order to explain the distinctive properties of the target—just as we can use the properties of H₂O in explaining why water boils at 100° C, why it is such a good solvent, and so on (see Chapters 1 and 2).

¹ I am grateful to Shriver and Allen (2005) for impressing this upon me.

In fact my ultimate argument against first-order theories is that they are incapable of explaining the distinctive, puzzling, to-be-explained features of phenomenal consciousness; whereas dispositional higher-order thought theory *can* explain those features. (Seen from this perspective, the contrast between conscious and non-conscious experience isn't even strictly necessary to the argument; although it does help greatly to make the point more vivid.) I shall spend the remainder of this chapter elaborating this claim.

2. THE PHENOMENA TO BE EXPLAINED

What are the main *desiderata* for a successful reductive theory of phenomenal consciousness? I shall lay out six of them.²

First and foremost, a reductive account of phenomenal consciousness should explain why phenomenally conscious states have a *subjective aspect* to them, meaning that it is *like something* to undergo them, and that they should each have a distinctive subjective *feel*. This is more or less the defining feature of phenomenal consciousness, and any adequate explanation needs to be able to account for it.

Second, a successful theory should explain why there should seem to be such a pervasive *explanatory gap* between all physical, functional, and intentional facts, on the one hand, and the facts of phenomenal consciousness, on the other. (See Chapter 2.) At a minimum, it seems that one can take any proposed explanatory story—including one framed in terms of dispositional higher-order thought theory, it should be stressed—and think, 'All of that might be true, and still *this* [type of experience] might be different or absent.' So a successful explanation needs at the same time to explain *why* such thoughts will always remain thinkable.

² It might easily seem that a seventh should also be listed. This is that a successful theory should be able to explain the difference between conscious and non-conscious experiences. But in the context of the present debate this isn't really a distinct *desideratum*. For both sides (first-order representationalism, on the one hand, and my dual-content theory, on the other) are agreed in *predicting* that the difference between phenomenally conscious and non-phenomenally conscious experience should fall where it seems to, in the outputs of the ventral/temporal and dorsal/parietal visual systems respectively. What is at issue is whether those theories can successfully *explain* the distinctive properties of our phenomenally conscious states; and this is what my six *desiderata* are intended to capture. Moreover, once we widen the scope of the debate to include other versions of first-order theory, some of which claim that dorsal/parietal percepts *are* (or might be) phenomenally conscious despite not being *access-conscious*, then it is obvious that explaining why the conscious–non-conscious distinction should fall where it seems to is no longer an appropriate constraint on a theory of *phenomenal* consciousness. This is because such theorists will deny the reliability of our intuitive beliefs about which of our experiences fail to be phenomenally conscious. For they claim that these beliefs may merely reflect the inaccessibility of those states to conceptual thought and verbal report. (Against this alternative form of first-order theory my objection is essentially the same: it still can't satisfy the *desiderata* for a successful theory of phenomenal consciousness.)

Now in fact there is an emerging consensus amongst would-be naturalizers of consciousness that the key to explaining (and to defusing) the explanatory gap lies in the existence of purely recognitional concepts of experience—sometimes called ‘phenomenal concepts’—which have no conceptual connections to physical, causal-role, or intentional-content concepts (Loar, 1997; Tye, 1999; Carruthers, 2000 and Chapter 5 of the present volume; Sturgeon, 2000; Papineau, 2002). In which case, this second *desideratum* amounts to saying that a successful theory needs, *inter alia*, to be capable of explaining the existence of purely recognitional concepts of experience.

Third, a successful theory ought to be able to explain why people should have such a persistent intuition that the properties of their phenomenally conscious states are *intrinsic* ones, being non-relationally individuated. When we reflect on the distinctive qualities of an experience of red, for example, it seems to us that those properties don’t depend upon the existence of the worldly property *red*, nor upon the experience in question occupying any particular sort of functional role within our mental lives (Kripke, 1972). Rather, those properties seem to us to be intrinsic to the experience. This is the *qualia-intuition*, which all representational theories of consciousness (whether first-order or higher-order) are committed to denying, and to explaining *away*.³

Fourth, a successful theory of phenomenal consciousness ought to be capable of explaining why phenomenally conscious experiences should seem to their possessors to be peculiarly *ineffable* (indescribable and incommunicable). When I reflect on the distinctive qualities of any experience that I am currently undergoing—say a perception of a warm vibrant shade of red—I seem to be aware of qualities that slip through the net of any descriptive scheme that I attempt to impose on them. I can say, ‘It is an experience of a warm vibrant red’, of course, or ‘It is an experience of a bright scarlet.’ But such descriptions feel wholly inadequate to the task of expressing what I am aware of. And I can *exhibit* the red object to another person, of course, saying, ‘It is the experience of *this* color.’ But then I am aware that I am *neither* describing *nor* exhibiting the experience itself, and am forced to rely upon the presumed similarities between our respective perceptual systems in order to communicate my meaning.

Fifth, and relatedly, a successful reductive theory ought be able to explain why the properties of phenomenally conscious experiences should seem in some way *private* to their possessors. We are strongly tempted to say that only we ourselves can truly know what our own experiences are *like*, and that others can only approximate to that knowledge by means of a more or less shaky inference. (For a recent

³ Some writers use the term ‘qualia’ in a weak sense, to refer just to whatever the properties are that render our experiences phenomenally conscious. See Sturgeon (2000), for example. Throughout this chapter I adopt the more standard philosophical usage, in which ‘qualia’ is used in a strong sense, to mean the alleged intrinsic, non-relational, properties of phenomenally conscious experience.

elaboration on this sort of view, see Block, 2002.) Yet if any representational theory of consciousness is correct, of course, then the relevant properties *aren't* private ones. On the contrary, they are in the public domain, knowable in principle by others as well as by oneself. So again there is a thesis here that needs to be denied, and a persistent tendency that needs to be explained away.

Finally, and again relatedly, a successful theory ought to be able to explain why it should seem to people that they have *infallible* (and not just *privileged*) knowledge of the properties of their phenomenally conscious experiences. There is a strong temptation to insist that we can be completely certain of the qualities of our own conscious states, in a way that we couldn't be certain of any physical, causal-role, or intentional-content property. This, too, is a thesis that will need to be resisted by any representational theory, and our sense of infallibility will have to be explained away.

3. HOW MUCH CAN A FIRST-ORDER THEORY EXPLAIN?

How much progress can a first-order theory of the sort espoused by Tye (1995, 2000) and others make with these six *desiderata*?

First, it is plain that the core property of possessing a *subjective aspect* can't get explained by anything distinctive about the *contents* of the first-order perceptual experiences in question. Granted, the percepts produced by the ventral/temporal system must involve a subjective *take* on some aspects of the environment and not others, determined by the properties of the subject's perceptual apparatus.⁴ But the same is equally true of the percepts produced by the dorsal/parietal system. And granted, the percepts produced by the ventral/temporal system have a distinctive sort of analog and/or non-conceptual intentional content, different from the digital/conceptual contents of belief and judgment.⁵ But the same is also true of the non-conscious percepts produced by the dorsal/parietal system. So there aren't the resources, here, with which to explain the distinctive subjectivity—the 'what-it-is-likeness'—of the former set of percepts.

⁴ In Ch. 3 of the present volume I refer to this as the 'subjectivity of the world', contrasting it with the 'subjectivity of experience' that is truly distinctive of phenomenally conscious experience. All perceptual states have *worldly* subjectivity, in that they possess a distinctive representational profile, determined by the properties of the perceptual apparatus in question. (Only some wavelengths of light are detected by the visual system, for example, and some wavelengths are perceived as being more similar to, or as being the opposites of, others.) But only phenomenally conscious percepts have *experiential* subjectivity; only they possess, themselves, a subjective dimension or *feel*.

⁵ Intentional contents are *analog* when they have a finer grain than any concepts that the subject can possess and recall. Intentional contents are *non-conceptual* when their existence is independent of any concepts that the subject possesses. The latter implies the former, but not vice versa. I therefore work with the weaker notion of analog content throughout this chapter.

The only option for a first-order theorist, then, is to attempt to explain the subjectivity of ventral/temporal percepts in terms of their distinctive functional role. In particular, it might be said that such percepts are, by virtue of their availability to conceptual thought and judgment, *available to the subject* in a way that the dorsal/parietal percepts aren't. Now, one might wonder with what right *the subject* should get identified with the concept-deploying sub-systems of the mind rather than with the sensorimotor ones. Why shouldn't percepts that are available to guide the actions of a subject count as 'available to the subject' just as much as percepts that are available to inform belief and judgment? For each is in fact available to less than the whole person, but rather to sub-systems that nevertheless form essential parts of the whole person. Although this is a powerful criticism, it isn't really the main point, as I shall now try to explain.

From the perspective of a first-order theorist, the sense in which ventral/temporal percepts are available to the subject can't be that the subject is capable of introspecting and identifying those percepts as such. For this would then be a higher-order theory, not a first-order one. Rather, it is just that the percepts in question are apt to give rise to beliefs and judgments about the worldly (and bodily) items that those percepts concern. So my perception of a red tomato is apt to give rise to the belief that I am confronted by a red tomato; my perception of a moving light causes me to judge that something is in motion; and so on. And it just isn't clear why this sort of availability should give rise to the subjective *feel* that is distinctive of phenomenally conscious states. Why *should* the functional role of causing first-order judgments give rise to states that it is *like* something to undergo, when the functional role of guiding action doesn't? The connection is far from transparent, to say the least.

In fact at this point first-order theorists are reduced to postulating a brute *identity* between phenomenal consciousness, on the one hand, and percepts that are available to first-order thought, on the other. But this is to give up on the goal of seeking a reductive *explanation* of phenomenal consciousness. We want to know *why* it should be *like* something to undergo perceptual states that are available to first-order judgment and decision-making. Whereas the best that a first-order theorist can do, is to tell us that this is just what it *is* for a state to be phenomenally conscious. And that isn't good enough.⁶

What, then, of the second *desideratum*: that a good theory should be capable of explaining the existence of purely recognitional concepts of experience, necessary to explain (and explain away) the explanatory gap? It should be obvious, I think, that first-order theories lack the resources to explain the existence of such concepts. I have argued this at length elsewhere (see Chapters 4 and 5). But here

⁶ See Chs. 1 and 2 of the present volume for arguments that identities need to be backed up by reductive explanations if they are to count as successful reductions.

let me point out, first, that a recognitional concept of *experience* is inevitably higher-order rather than first-order in character; and second, that there is nothing in the content of a first-order percept that could serve to ground the application of such a higher-order concept. A perceptual state with the first-order analog content *red* is of the right sort to ground a first-order recognitional judgment of *red*; but not of the kind required to ground a higher-order recognitional judgment of *this* (where the object of recognition is a state of seeming red, or an experience of red).⁷

Might a recognitional judgment of experience be grounded, not in the first-order *content* of the experience, but rather in the *occurrence* of that experience? A first-order theorist might try claiming that the higher-order purely recognitional judgment *this* [experience of red] is grounded, not in the analog content *red* (nor in any higher-order analog content, either, as dual-content theory claims—see below), but rather in the *state* of perceiving red. Such claims can't do justice to the distinctive character of our purely recognitional concepts of experience, however. In particular, they can't explain the manner in which we have awareness of that which grounds our judgments. For when making a judgment of *this* I am aware of a fineness of grain in the object of my judgment that seems to slip through the mesh of my conceptual net. Indeed, the proposal on offer would seem to be a species of 'brute causal' account, of the sort that is criticized extensively in Chapter 5.

First-order theories can seem a bit more successful with the third *desideratum*, which is to explain the *qualia-intuition*. But for the explanation to work, two distinct sources of confusion have to be postulated. Recall that the qualia-intuition is the belief that our phenomenally conscious perceptual states possess non-relational (non-intentional, non-functionally individuated) properties. A first-order theorist *might* be able to explain why we have the intuition that there are properties of the world (or of our own bodies) that are intrinsic ones—for example, the property of being red. For a percept of red doesn't *represent* the relational character of the property that it represents. (It certainly doesn't represent it *as* represented, for example.) Rather, it just represents the surface in question as covered with a distinctive recognizable property. So if (as seems plausible) we have a tendency to confuse an absence of representation with a representation of absence, we might come to believe that the property of being red isn't a relational one.

⁷ Does the displaced-perception account of introspective awareness proposed by Dretske (1995) serve as a counter-example to this claim? On this view, when I make an introspective judgment of *this* [experience of red] what I am really doing is judging that I am experiencing *this* [redness]. In which case the experience of redness *is* of the right kind to ground an introspective judgment, contrary to what I claim. However, on this account the recognitional concept in question turns out not to be *purely* recognitional, after all. On the contrary, it has embedded within it a concept of experience. And this is then sufficient to preclude Dretske from explaining all of the intuitions that can seem to give rise to an explanatory gap. See Ch. 5 of this volume.

But even if this were convincing (and I will return in a moment to argue that it isn't), it doesn't yet explain the intuition that the experience of red itself (and not just the redness experienced) possesses non-relational properties. For that, we have to postulate a second source of confusion: between the properties presented in or represented by our experiences, and the properties of those experiences. So we have to suppose that people have a tendency to move from thinking that there is a fine-grained intrinsic property (*redness*) represented by an experience, to thinking that this is really a fine-grained intrinsic property of the experience itself.⁸

There may well be some such tendency. (Thus Dennett, 1991, for example, argues that we mistake the richness and unlimited detail of the world perceived for an equivalent richness in our own perceptions of the world.) But I believe that this explanation of the *qualia-intuition* fails at the first hurdle. For we don't actually possess the intuition that redness is an intrinsic property of external surfaces. The relational nature of color is a deeply ingrained part of common-sense belief. For we know that the colors that objects appear to have will vary with the lighting conditions, and with the state of the perceiver. So even if our perceptions of color don't represent the relational character of the colors represented, most of us nevertheless believe firmly in their relational nature.

As for the fourth *desideratum*, first-order theories can make *some* progress with the *ineffability intuition*. This is because of the fine-grained (analog or non-conceptual) character of the first-order perceptual contents in question. Our color-percepts represent subtle variations in shade that far outstrip our capacities for classification and description, for example. (Similar things are true of perceptible properties in each of the various sense modalities.) So when I am forced to think about or describe the color of the object that I am seeing in comparatively coarse-grained conceptual terms, it is well-nigh inevitable that I should feel that the property in question is indescribable (even though it *can* be *exhibited* to others).

But this is all about the ineffability of *color*, not about the ineffability of our experience of color. The latter is an intrinsically higher-order phenomenon. (To say that an experience of color is ineffable is to say that there is something about the experience that outruns our capacity for higher-order conceptual thought and/or our capacity for higher-order linguistic description.) And while the fineness of grain of our first-order percepts can explain why we can't give a satisfying higher-order *description* of those experiences—I am forced to say, 'It is an experience of *red*', while being aware that there is much more detail than that present in the objects of my experience—it *can't* explain why we should feel that we can't even *exhibit* the relevant properties of those experiences to another

⁸ This would be the inverse of the tendency famously postulated by Hume (that the mind has a tendency to 'spread itself onto' the world); it is a tendency of the world to spread itself onto the mind. In contrast with Hume's *projectivism*, then, this tendency might be labeled 'back-projectivism'.

person. (So it can't explain why the experience should seem fully *incommunicable* and not just *indescribable*.) For if all that were present were a first-order representation of a fine-grained property of a surface of the object, then I *should* be able to exhibit the experience to others just as I can exhibit the relevant property, by saying, 'It is the experience of *this*', while demonstrating the object that I am perceiving.⁹

For related reasons a first-order theory can't, as such, explain the intuition that our conscious experiences are *private*, and unknowable to others. For knowledge of experience is necessarily a higher-order phenomenon. If you are *aware* of your experience of red, or know *that* you are undergoing an experience of red, then you must be *representing* that experience; and this is, of course, a higher-order representation. Moreover, there is nothing about the content of a first-order experience that should suggest privacy. While I can indeed wonder whether other people are really tuned to the same external properties as myself when I perceive things (and hence while I can know that they can be subject to the very same doubts concerning myself), it ought to be equally obvious that this will be relatively easy for others to test. It is just a matter of determining whether or not our discrimination-profiles are sufficiently similar.

Likewise, then, with the intuition that we have *infallible* knowledge of our own phenomenally conscious experiences: since knowledge *of* experience is intrinsically higher-order, infallible knowledge would have to be so too. And there is nothing in the story that a first-order theorist tells us that can explain why we should be so tempted to believe that our higher-order knowledge of the character of our experiences should be infallible. Certainly there is nothing in the content of a first-order experience of red, say, that should suggest any sort of infallibility.

4. THE EXPLANATORY POWER OF DUAL-CONTENT THEORY

In contrast with first-order theories, dispositional higher-order thought theory *can* meet each of our six *desiderata* in a satisfying way. According to this account, the percepts produced by the ventral/temporal visual system acquire *dual* analog intentional contents, because they are available to higher-order thought,

⁹ Admittedly, with a little bit of background theory it will be possible to wonder whether other people have sense organs that tune them to the very same external properties (e.g. the reflective properties of surfaces) as do mine. So in exhibiting a red surface to another person, I can doubt whether I have succeeded in exhibiting the *relevant* property (the property that forms the object of my own perception). But this should give me no reason to feel that the property in question *cannot* in principle be communicated to others. On the contrary, it will be fully *effable* to other people on condition that their perceptual systems are relevantly like my own.

and by virtue of the truth of some or other form of consumer semantics (either teleosemantics, or causal or inferential role semantics). Those experiences become both world-representing (e.g. with the analog content *red*) and experience-representing (e.g. with the analog content *seems red* or *experience of red*) at the same time. In contrast, the percepts produced by the dorsal/parietal system retain analog intentional contents that are merely first-order in character, representing properties of the world (e.g. red), but not properties of our experience of the world. For the latter percepts *aren't* available to higher-order thought.

This immediately gives us an account of the key subjectivity, or 'what-it-is-likeness', of phenomenally conscious experience. For by virtue of possessing a dual analog content, those experiences will acquire a subjective aspect. This is their higher-order analog content. Hence they come to present *themselves* to us, as well as presenting properties of the world (or of the body) represented.¹⁰ Moreover, recall that by virtue of having the first-order analog content *red*, there is something that *the world* is *like* to one who has that experience. But now by virtue of having, in addition, the higher-order analog content *seems red*, there is something that the experience itself is *like* for the one who has it. So dual-content theory ascribes to our conscious experiences both *worldly* subjectivity and *experiential* subjectivity, with the contents of the latter being parasitic upon, and exactly mirroring, the contents of the former. And the 'what-it-is-likeness' of phenomenally conscious experiences is thereby explained.

In similar manner, dual-content theory gives us a satisfying explanation of the possibility of purely recognitional concepts of experience ('phenomenal concepts'). I have explained this at length elsewhere (see Chapter 5). But in brief, the idea is that the higher-order analog contents of experience enable us to develop purely recognitional concepts of experience, in just the same sort of way that the first-order contents of experience enable us to develop recognitional concepts of secondary qualities like redness and smoothness. The first-order analog content *red* (of some specific shade) serves to ground, and to guide the application of, a recognitional concept of red. Likewise, according to dual-content theory, the higher-order analog content *seems red* or *experience of red* (again, where the *seeming* is of some specific shade) can serve to ground, and to guide the application of, a purely recognitional concept *this* [experience of red].

¹⁰ See Ch. 8 of the present volume for further discussion of this idea. Note that I ought really to say, strictly speaking, that it is by virtue of having higher-order analog contents that our experiences come to present themselves *to their consumer systems*, rather than to *us*, if I am to avoid the objection that I leveled earlier against first-order theories. (With what right, it might be objected, do I assume that by virtue of their availability to the sub-systems responsible for higher-order thought and for linguistic judgment, our dual-content experiences are thereby available *to the subject*?) This then leaves open as a (remote) conceptual possibility that there might be phenomenally conscious (= dual content) experiences that are *inaccessible* to the systems that govern practical reasoning and linguistic report, as I duly acknowledge in Ch. 4 of the present volume.

Such a recognitional concept need have no conceptual connections with any physical, functional-role, or intentional-content concepts. (This is what it means to say that it is *purely* recognitional.)¹¹ And hence it is possible for someone deploying such a concept to entertain thoughts like the following: 'It might be possible for someone to be undergoing a perceptual state with the dual analog contents *red* and *experience of red*, while nevertheless lacking *this* [experience of red].' But of course this needn't prevent us from accepting that the property picked out by the recognitional concept *this* just *is* a perceptual state with the analog content *red*, recognized via its higher-order analog content *seeming red* (supposing that we were convinced of the virtues of dual-content theory overall). So the so-called 'explanatory gap' between physical, functional-role, and intentional facts, on the one hand, and the facts of phenomenal consciousness, on the other, is both explained and explained away.

What, then, of the *qualia-intuition*? How does dual-content theory explain the persistent tendency for people to believe that phenomenally conscious experiences possess properties that are *intrinsic* and non-relational in character? The explanation comes in two parts. The first part is conceptual. Since the concept *this* [experience of red] need have no conceptual connections with functional-role and/or intentional-content concepts, the concept *this* doesn't represent the property that it picks out *as* a relational one. And then the second part of the explanation is perceptual. The dual analog content of my phenomenally conscious experience of red neither represents the first-order property in question (redness) by its relations with anything else, nor does it represent the second-order property (experienced redness) relationally, either. Rather, just as the first-order contents of the experience seem to present to us an intrinsic property of the surface of the object perceived, so the second-order content of the experience seems to present to us an intrinsic (but introspectible) property of the experience itself.

But in contrast with first-order perception, we have no common-sense theory of introspection that might lead us to believe that the properties introspected are really relational in character after all. Although the first-order content of an experience of red seems to present us with an intrinsic, non-relational, property of the surface perceived, we have enough theoretical knowledge of the nature of perception to believe that redness is, really, a relational property of objects. In contrast, there is nothing to stop us from taking the higher-order content of a phenomenally conscious experience of red at face value, hence coming to *believe*

¹¹ In contrast, the recognitional concept *red* isn't a *purely* recognitional one, since it is embedded in a proto-theory of vision; whereas the recognitional concept *this* [experience of red] can be independent of any sort of proto-theory of introspection. And recall from Ch. 5 that it is vital that we should be able to explain how we can have a concept here that is *purely* recognitional, or we shan't be able to explain the intelligibility of some of the philosophically problematic thoughts, such as, 'Someone might have *this* occupying some quite different role within their cognitive system' (e.g. without it even being an *experience*, considered as a state that is functionally individuated).

that the experience possesses an intrinsic property, just as its higher-order analog content seems to present it as having.

It is only natural, then, since our experiences seem to present us with intrinsic properties of our experiences themselves, and since we aren't required to conceptualize those properties as *non*-intrinsic, that we should come to believe that our phenomenally conscious experiences possess properties that are intrinsic and non-relationally individuated (*qualia*, in the strong sense). But for all that, of course it is possible for us to believe, with dual-content theory, that there *are no* intrinsic properties of experience qua experience, since all of the relevant properties are in fact representational ones.¹²

The *ineffability-intuition*, too, is neatly explained by dual-content theory. Recall that both the first-order and higher-order contents in question are analog, or fine-grained, in character. So not only will the distinctive properties of the objects perceived (e.g. a particular shade of redness) seem to slip through the mesh of any conceptual net that we use to describe them, but so too will the distinctive properties of the perceptions themselves (e.g. a *seeming* of a particular shade of redness), presented to us by their higher-order analog contents. So we will have the persistent sense that our phenomenally conscious experiences cannot be adequately described to others. But then nor, of course, can they be *exhibited*. While I can exhibit the object of my first-order experience (e.g. a red surface), I can't exhibit the object of my higher-order experience (e.g. an experience of redness). So by virtue of the kind of awareness that I have of my own experiences, I shall have the persistent feeling that they are wholly *incommunicable* to others—for I can neither describe, nor exhibit to others, the object of my awareness.

Although the ineffability-intuition is powerful, and is explained by dual-content theory, it is nevertheless misguided if the explaining theory is correct. For if my phenomenally conscious experience of a particular shade of red is nothing other than a dual-content perceptual state with the analog contents *red* and *experience of red*, then it will be possible, in principle, to describe that experience fully to another person. But actually providing such a description may have to wait on an adequate theory of intentional content. For only when that theory is in hand will we be able to specify fully the distinctive properties of the experience in question, in virtue of which it has the particular analog contents that it has (representing both a specific *shade* of red and a *seeming* of that very shade).

The *privacy-intuition* is likewise both explained and undermined by dual-content theory. For if the properties of our phenomenally conscious experiences are apt to seem ineffable, such that we can neither describe nor exhibit them to

¹² It is vital to this explanation, of course, that one should keep straight the distinction between *concepts* (mental representations), on the one hand, and *properties* (features of the world or the mind represented), on the other. And one should have a robust, 'thick', conception of the latter, allowing that one and the same property can be picked out by a number of distinct concepts. See Carruthers, 2000.

others, then it is quite natural to conclude that they cannot be known to other people, either. One will be apt to think something like the following: ‘How can another person ever know that I have *this* [experience of a specific shade of red], given that I can’t adequately describe it to them, and given that I can’t exhibit it to them for their inspection, either?’ But if the properties in question are dual-content ones, then plainly it should be possible, in principle, for others to know of them—provided that the intentional properties in question can be known to others, of course. (Here, too, detailed knowledge of the kind required may have to wait on a worked-out theory of intentional content.)

Finally, what of the *infallibility-intuition*? Can dual-content theory explain why people should be so tempted to believe that they have infallible knowledge of the properties of their own conscious experiences, when they have them? Indeed, it can. For notice that the higher-order analog contents postulated by the theory are entirely parasitic upon their first-order counterparts, only coming to exist by virtue of the availability of the first-order analog contents in question to a faculty of higher-order thought. So, in contrast with so-called ‘inner sense’ theories of higher-order perceptual content (Armstrong, 1968; Lycan, 1996),¹³ there is no possibility of the higher-order contents getting out of line with the first-order ones. On the contrary, whenever a percept has the first-order analog content *F*, that state will also possess the higher-order analog content *seems F* or *experience of F*—at least, provided that the state in question is available to a higher-order thought faculty of the right sort, and given the truth of some or other kind of consumer semantics.

According to the dispositional higher-order thought version of dual-content theory, our higher-order experiences really do track our first-order ones in a way that is infallible. This is a consequence of the fact that the former are parasitic upon the latter. So whenever we have a higher-order analog experience with the content *seems F*, we must be undergoing a first-order experience with the content *F*. In which case there is barely any room for error when we make immediate recognitional judgments about our own phenomenally conscious experiences, while we are having them.¹⁴ Since there is no scope for *misperceiving* our own experiences (in the way that we *can*, of course, misperceive properties of the world) the only remaining room for error lies in various forms of mental disorder, which might in one way or another interfere with our conceptual abilities.

¹³ Such theories postulate that we have inner sense-organs, designed to deliver higher-order perceptions of (some of) our perceptual states, just as we have outer sense-organs designed to deliver perceptions of our environments and bodies. It will therefore be possible for such sense organs to malfunction, yielding higher-order *mis*-perceptions of our own first-order experiences. See Ch. 4 of the present volume.

¹⁴ When it comes to *remembering* our phenomenally conscious experiences, of course, our judgments will be just as liable to failure as are any other judgments about the past. Equally, we can of course be mistaken about the *causes* of our experiences, and in judgments that require some sort of mental operation to be performed on the immediate contents of experience, such as counting the sides of a mentally imaged figure, or making comparative judgments between one experience and another.

5. WHAT CAN A FIRST-ORDER THEORIST GAIN BY EMBRACING DUAL CONTENT?

I have argued that dual-content theory does a fine job of meeting the various *desiderata* that we laid down in section 2 for a successful theory of phenomenal consciousness, whereas first-order theories don't. So far, then, dual-content theory is winning 'hands down'. But it might seem that there is a further option open to first-order theorists, at this point. For couldn't they *also* believe in the existence of dual-analog contents? And couldn't they thereby get all of the explanatory advantages of dual-content theory while remaining first-order theorists about phenomenal consciousness as such?¹⁵ I shall elaborate on this idea first, before turning to evaluate it.

It is plain that first-order theorists can accept that some of our first-order perceptual states (specifically, in the case of vision, those produced by the ventral/temporal system) are made available to a faculty of higher-order thought, which is capable of forming partly recognitional concepts that apply to the first-order states in question.¹⁶ And it is also plain that first-order theorists can accept some or other form of consumer semantics. In which case first-order theorists might believe that our phenomenally conscious experiences are *also* dual-content states, having both first-order and higher-order analog intentional contents. But they will insist that it isn't in virtue of possessing dual analog contents that the states in question are phenomenally conscious ones, and thus come to have *experiential* as well as *worldly* subjectivity. Rather, their phenomenally conscious status will derive from the fact that they have *first-order* analog contents available to conceptual thought and conceptual judgment.

Now, if first-order theorists can think that phenomenally conscious states are also dual-content ones, then they can surely endorse many of the explanations offered in section 4 above, hence fulfilling the relevant *desiderata* of a theory of consciousness. For example, they can explain the existence of purely recognitional concepts of experience as being grounded in the higher-order analog contents of our perceptual states, in something like the way that first-order recognitional concepts of color are grounded in the first-order analog contents of

¹⁵ Tye (1995, 2000) sometimes seems to come close to endorsing this combination of views. I am grateful to Keith Frankish for pushing this idea on my attention.

¹⁶ At this initial stage the concepts, although recognitional, will still be theoretically embedded. Subjects who understand what an experience is will know that, whenever they are inclined to make a judgment of 'red', they are also entitled to make a judgment of 'seems red' or 'experience of red'. The component concepts of *seeming* and *experience* here are theoretical ones. It is only when the perceptual states in question thereby acquire higher-order analog contents that the subject can acquire *purely* recognitional concepts of experience, making judgments of 'this again' grounded in the higher-order perceptual contents in question.

our experiences. But they can do this while continuing to claim that all and only those mental states that possess first-order analog intentional contents that are available to conceptual thought are phenomenally conscious ones.

One might be tempted to respond to this suggestion by saying that phenomenal consciousness is, surely, whatever it is that gives rise to the 'hard problem' of consciousness (Chalmers, 1996). And it is the various features laid out in our six *desiderata* that make the problem of explaining phenomenal consciousness seem particularly hard. (This is especially obvious in respect of the second feature, which is precisely that phenomenal consciousness seems to give rise to a persistent 'explanatory gap'.) But what has now emerged from the discussion in sections 3 and 4 above, is that it is the dual-analog content of our perceptual states, rather than their first-order analog contents, that gives rise to the appearance of a 'hard problem'. In which case phenomenal consciousness itself should be identified with states possessing a dual-analog content, rather than more generally with first-order analog states that are available to conceptual thought.

There is some danger that the disagreement between first-order theorists and higher-order theorists may become merely verbal at this point, however. Each would be agreeing that the various puzzling features of phenomenal consciousness are to be explained by appeal to dual-analog contents. But the first-order theorist would be insisting on identifying phenomenal consciousness itself with the *object* of those puzzles—with what we are thinking *about* when we worry about the intrinsic properties of our experience or the ineffability of a sensation of red, say. And a higher-order theorist, too, should agree that what we are thinking about in such cases is a first-order analog perceptual state (albeit one that is presented to us via a higher-order analog representation of itself). But the higher-order theorist would be insisting, on the contrary, that phenomenal consciousness itself should be identified with that which *gives rise to* the various puzzles. And this is the higher-order analog content of our conscious perceptual states, as we have seen.

It is unclear, at this point, whether there would be anything substantive remaining in this disagreement. For it is unclear whether anything of theoretical importance turns on the question whether phenomenal consciousness should be identified with what we are thinking *about* when we are puzzled by the 'hard problem', or rather whether it should be identified with that which *gives rise to* the puzzles that constitute the 'hard problem'.

I believe that the dispute is *not* merely verbal, however. Recall from section 2 that the distinctive subjectivity, or 'what-it-is-likeness', of our phenomenally conscious experiences was said to be the defining feature of them. A phenomenally conscious state just *is* a state possessing the relevant kind of subjectivity. But we also said that first-order theory as such lacks the resources to explain it; whereas dual-content theory provides the needed explanation. If these claims are correct, then first-order theory (even when combined with a belief in dual analog contents)

isn't an acceptable theory of phenomenal consciousness. For it is the higher-order analog content, not the first-order content, that constitutes the defining feature of such consciousness (its subjective *feel* or 'what-it-is-likeness').

Does this reply commit me to any kind of objectionable *qualia*-realism? By insisting that the subjectivity of our phenomenally conscious states is a real (and defining) property of them, standing in need of explanation, have I let in *qualia* by the back door? Not at all. For there is, indeed, a real but naturalistically acceptable property here: it is the higher-order analog content possessed by the experiences in question. This gives those experiences a subjective dimension, making them available to their subjects in something like the way that first-order analog contents make available to us properties of our environments (and bodies).

Moreover, the higher-order analog content of the content-bearing states with which my sort of higher-order theorist identifies phenomenal consciousness serves to mark a real and substantial difference between the perceptual states produced by the ventral/temporal visual system (which possess such content) and those forming the output of the dorsal/parietal system (which don't). So we have identified a real difference between the states of the two kinds, which marks out the former as phenomenally conscious and the latter as not, in a way that is thoroughly relevant to their phenomenally conscious status. A first-order theorist, in contrast (even one who accepts that dual analog contents are differentially possessed by the ventral/temporal states) is committed to *denying* that it is the presence of the higher-order analog content that renders these states phenomenally conscious. On the contrary, it will have to be something about the wider functional role of those states (e.g. their availability to first-order thought and planning) that constitutes them as phenomenally conscious. And we have already seen the inadequacies of this approach (see section 3 above).

I conclude, then, that although first-order theorists can accept the existence of dual analog contents, and can hence reap the benefits of explaining many of the puzzling facts about phenomenal consciousness, they still can't provide an adequate account of what phenomenal consciousness itself *is*. For they can't explain the key subjectivity, or *what-it-is-likeness* of such consciousness. For this, they would have to accept that phenomenal consciousness should be *identified* with states possessing dual analog contents. But that, of course, would be a higher-order theory and not a first-order theory. Indeed, it is my own theory.

6. CONCLUSION

In summary, then, the case against first-order theories of phenomenal consciousness is that they are incapable of providing a reductive explanation of all of the various distinctive and problematic features of phenomenal consciousness.

114 | *Dual-Content Theory*

In contrast, the kind of higher-order dual-content theory that arises from combining dispositional higher-order thought theory with some suitable form of consumer semantics *can* explain those features. This provides us with powerful reasons for rejecting first-order theories, and for embracing dual-content theory instead.^{17,18,19}

¹⁷ The case against other varieties of higher-order theory (specifically, inner-sense theory and actualist higher-order thought theory) needs to be made rather differently. It should be obvious that inner-sense theory can meet the first five of our six *desiderata*, since it, too, postulates the existence of higher-order analog contents targeted on (some of) our first-order experiences. But the theory predicts that our inner sense-organs should sometimes malfunction, just as our outer senses do. In which case there should be instances where people are inclined to judge, spontaneously and without theoretical influence, 'The surface is red' while at the same time judging, 'The surface *seems* orange.' This seems barely coherent. There are, moreover, no good evolutionary reasons why an organ of inner sense should have evolved. See Carruthers, 2000, and Ch. 4 of the present volume for further discussion. Actualist higher-order thought theory, in contrast, can't meet our *desiderata* as such, since it fails to postulate higher-order analog contents. It *can* make appeal to such contents by accepting some version of consumer semantics; but then the motivation for being actualist rather than dispositionalist will evaporate, since all forms of consumer semantics are dispositionalist in form. In addition, actualist higher-order thought theory faces essentially the same evolution-based objections as does inner-sense theory, since it is mysterious why evolution should have ensured the generation of so many higher-order thoughts on a regular basis. See Carruthers, 2000, and Ch. 4 of the present volume for further discussion.

¹⁸ Note that if it also follows (when dual-content theory is supplemented by suitable premises concerning the cognitive powers of non-human animals) that non-human animals aren't subject to phenomenally conscious mental states, then that is a conclusion with which we are somehow going to have to learn to live. This topic is taken up at length in Chs. 9, 10, and 11 of the present volume.

¹⁹ I am grateful to Keith Frankish, Robert Lurz, Benedicte Veillet, and an anonymous referee for Oxford University Press for their thoughtful comments on a draft of the present chapter.