

Too much ado about belief

Jérôme Dokic & Elisabeth Pacherie

Institut Jean-Nicod
CNRS-EHESS-ENS
1 bis, avenue de Lowendal
75007 Paris
dokic@ehess.fr
pacherie@ens.fr

Three commitments at least appear to be guiding Dennett's approach to the study of consciousness. First, an ontological commitment to materialist monism. Second, a methodological commitment to what he calls 'heterophenomenology'. Third, a 'doxological' commitment that can be expressed as the view that there is no room for a distinction between a subject's beliefs about how things seem to her and what things actually seem to her, or, to put it otherwise, as the view that there is no room for a reality/appearance distinction for consciousness.

Our main aim will be to investigate how Dennett's third doxological commitment relates to his first two commitments and whether its acceptance should be seen as a mere logical consequence of acceptance of the first two. We will argue that this is not the case, that Dennett's doxological commitment is in need of independent motivation, and that this independent motivation is not forthcoming. More specifically, we will argue that there is more to the conscious content of perceptual experience than what transpires in a subject's phenomenal beliefs¹. We will develop two lines of arguments. First, as Marcel (2003) has shown, a distinction of levels of awareness is needed to account for both the relations and distinctions between various pathologies of perceptual consciousness. Second, there is a sense

¹ Here and in the remainder of this paper, we use the phrase 'phenomenal beliefs' to refer to the subjective beliefs a subject has about how things seem to her.

in which perceptual experience has nonconceptual content and this content cannot be fully captured in terms of beliefs the subject has or could have. We will argue that positing the existence of this nonconceptual content is perfectly compatible with a methodological commitment to heterophenomenology, that is, the presence of this non-conceptual content is assessable using heterophenomenological methods. Finally, we will show that one can acknowledge the existence of nonconceptual content without fear of thereby smuggling dualism back into the picture.

In the first section, we will say more about what Dennett's three commitments involve and how they relate. In the second section, we will discuss Marcel's distinction between visual phenomenology and awareness thereof and show its relevance to an evaluation of Dennett's doxological commitment. In the third section we will try to motivate a distinction between conceptual and nonconceptual contents and to show that this distinction leads in turn to a distinction between experience and beliefs about experience. Some of the arguments that have been adduced in favor of the distinction between conceptual and nonconceptual content may not pass muster when subjected to the strictures of the heterophenomenological method. Others, however, do, or so we will argue. Similarly, some ways of conceiving of the contents of perceptual experience may indeed smack of dualism, but others are unobjectionable in that respect.

1. Dennett's three commitments

Dennett's commitment to materialism is a commitment to the view that there is only one sort of stuff, namely matter, or, to put it otherwise, whatever entities and processes are studied by physics, chemistry and biology, and that the mind is no exception and is ultimately nothing but a physical phenomenon. This commitment is shared by many philosophers and scientists

today, at least up to a point. Some people are happy to consider that most mental processes are indeed physical processes but have qualms when it comes to consciousness. More precisely, using Block's distinction between access consciousness and phenomenal consciousness, they tend to agree that access consciousness falls within the purview of materialism, but they stop short at phenomenal consciousness. This reluctance takes several forms. One can accept, as Nagel does, that physicalism about conscious experience might be true, but claim that so far "we do not have the beginning of concept of how it might be true" (Nagel, 1974: 447). One can pay even more perfunctory lip-service to materialism by arguing that, although consciousness is part of the natural realm, our chances of ever understanding how it can be are about the same as the chances of monkeys coming to grasp the principles of quantum physics (McGinn, 1989). Still others think that conscious experience cannot in principle be accounted for in materialistic terms and see this as a vindication of dualism (Chalmers, 1996).²

Dennett will have none of this. His is a die-hard materialism meant to embrace all aspects of the mental, including consciousness in all its forms. Indeed for him, if dualism is the best we can do, then any hope we may have of understanding consciousness is forlorn, but as he says, "such defeatism, today, in the midst of a cornucopia of scientific advances ready to be exploited, strikes me as ludicrous, even pathetic" (1991: 40). Although we have reservations about some aspects of Dennett's own theory of consciousness, we certainly agree with him that it is more than premature to forfeit materialism about consciousness. Thus, nothing we will say in defense of the existence of perceptual experience as distinct from phenomenal beliefs should conflict with a commitment to materialism.

² Note though that this dualism is more limited in scope than classical Cartesian dualism as it applies only to the phenomenal aspects of mental life.

Dennett's second commitment is a commitment to the method of heterophenomenology.

Heterophenomenology is intended as "the neutral path leading from objective physical science and its insistence on the third-person point of view, to a method of phenomenological description that can (in principle) do justice to the most private and ineffable subjective experiences, while never abandoning the methodological principles of science" (1991: 72).

The two main tenets of this methodology are the adoption of a third-person point of view and the adoption of a neutral stance towards the assertions of the subjects. It is, Dennett insists, nothing more than the standard third-person scientific method applied to the particular realm of phenomena of consciousness. This method involves three main steps. The first step involves recording raw data: verbal reports, button-pushings, eye blinks, as well as all other manifestations of belief, conviction, expectation, emotion, etc, including behavioral responses, brain activities, visceral reactions, hormonal reactions, and other changes in internal states that are detectable by objective means. The second step involves moving from raw data to interpreted data. In the case of verbal reports this means moving from vocal sounds to speech acts, which, assuming the heterophenomenologist has no reason to doubt the sincerity of the subjects, are treated as expressions of their beliefs about their experiences.

The neutrality or agnosticism of heterophenomenology resides in the fact that "the heterophenomenological method neither challenges nor accepts as entirely true the assertions of subjects, but rather maintains a constructive and sympathetic neutrality, in the hopes of compiling a *definitive* description of the world according to the subjects" (1991: 83).

This definitive description constitutes a description of the heterophenomenological world of the subject.

This initial bracketing for neutrality of the subject's phenomenal beliefs is required, according to Dennett (2001a), because of two failures of overlap, which he labels false positives and false negatives. On the one hand, some beliefs that subjects have about their own conscious states are provably false. For instance, many people think, wrongly, that their visual experience presents them with a detailed picture-like representation of the world, with high-resolution and uniformly rich detail from the center out to the periphery.³ On the other hand, "some psychological things that happen in people [...] are unsuspected by those people". For instance, on priming experiments using masking, subjects may be unable to report a masked prime and yet it can be demonstrated that their performance was influenced by the prime. The neutral stance is, however, only a temporary tactic. The third and final step in the inquiry consists in investigating whether the subject's beliefs can be confirmed or not. According to Dennett, there are three main types of results this enquiry may yield:

Often, indeed typically or normally, the existence of a belief is explained by confirming that it is a *true* belief provoked by the normal operation of the relevant sensory, perceptual, or introspective systems. Less often, beliefs can be seen to be true only under some arguable metaphorical interpretation — the subject claims to have manipulated a mental image, and we've found a quasi-imagistic process in his brain that can support that claim, if it is interpreted metaphorically. Less often still, the existence of beliefs is explainable by showing how they are illusory byproducts of the brain's activities: it only *seems* to subjects that they are reliving an experience they've experienced before (*déjà vu*). (2001a: 2)

³ However, see Noë (2004: Ch. 2) for reservations about whether people really have such a belief.

Thus, a phenomenal belief may be true, metaphorically true, or plainly false or illusory. One natural way to construe these three possibilities is as follows. A given phenomenal belief is: (1) true, if the subject's has or had the experience that he believes he had and the experience constitutes the basis for the belief; (2) metaphorically true, if the subject had an experience that, under some metaphorical interpretation, fits the content of the belief; or (3) false, when instead of being based on an experience with a corresponding content, it is an accidental by-product of the brain's activities. It is what Dennett seems to be saying in the above quoted passage. It is also, we think, what he should say. But is it what he really means to say? Turning now, to his doxological commitment, we will see that this is far from obvious.

Towards the end of chapter 4 of *Consciousness Explained*, when giving a summary of the heterophenomenological method he spent the chapter explaining, Dennett writes: "we are giving you [the subject] total, dictatorial authority over the account of how it seems to you, *about what it is like to be you*" (1991: 96). Here, the subject's beliefs about what it is like to be her are suddenly equated with what it is actually like to be her. On first reading this passage, a charitable reader would suppose it is a slip of the pen, for how could our phenomenal beliefs ever be false, as Dennett insists they can be, if they were, so to speak, constitutive of the reality they are supposed to be about? Yet, this strange assertion can't be blamed on a momentary lapse of attention. For one thing, Dennett is happy to quote this passage in several of his later writings.⁴ For another, in chapter 5 of *Consciousness Explained*, Dennett makes it crystal clear that this is what he actually means, saying for instance:

Postulating a "real seeming" in addition to the judging or "taking" expressed in the subject's report is multiplying entities beyond necessity. Worse, it is multiplying entities

⁴ For instance, Dennett 2001a, 2003, 2005.

beyond possibility; the sort of inner presentation in which real seemings happen is a hopeless metaphysical dodge, a way of trying to have your cake and eat it too, especially since those who are inclined to talk this way are eager to insist that this inner presentation does not occur in some mysterious, dualistic sort of space perfused with Cartesian ghost-ether. (1991: 134)

What we call Dennett's doxological commitment is indeed a commitment to the claim that there is no room for a distinction between what you believe it is like to be you and what it is really like to be you – or more generally that there is no such thing as an appearance/reality distinction for consciousness. As Schwitzgebel (this issue) forcefully argues, there is, to put it mildly, a tension between this claim and Dennett's other claim, used to motivate the neutrality of the heterophenomenological method, that we can be mistaken about our conscious experience.

The best way to bring out the tension between these claims is to draw a threefold distinction between:

- (1) A belief that one has an experience of type ϕ
- (2) A seeming to have an experience of type ϕ
- (3) A conscious experience of type ϕ

Here, type ϕ can include both aspects of the content of the experience and other (perhaps qualitative) aspects of the experience itself. Both (2) and (3) are *seemings*. Whereas (3) typically describes an experience whose content is entirely about the world, (2) describes a second-order experience whose content includes in addition the fact that one has the first-order experience. Now Dennett could be interpreted as meaning that there's no room for a distinction between (1) and (2): seeming to have an experience is just having the belief that

one has this experience. This would be his doxological commitment. On the other hand, keeping (1)-(2) distinct from (3) would allow him to maintain that we can be mistaken about our conscious experience. It is debatable whether (1) and (2) should be identified, but at least under this charitable interpretation, Dennett's position could be exonerated from the charge of downright self-contradiction.

Unfortunately, there's no evidence that this is the view he holds and there is some evidence that he would not accept this interpretation. For keeping (1)-(2) distinct from (3) seems to re-introduce some appearance/reality distinction for consciousness, something Dennett rejects. The passage quoted above suggests that for Dennett accepting an appearance/reality distinction for consciousness would ineluctably reinstate a form of dualism. So it would go against his ontological commitment to materialism. But is it so? Dennett seems to assume that accepting in one's mental ontology not just beliefs about experiences but also experiences involves committing oneself to the existence of such dubious entities as "sense data", "sensations", "raw feels", etc. But this is unwarranted. There have been several attempts by philosophers strongly committed to naturalism such as Tye (1995) and Dretske (1995) to develop purely representational theories of experiential states. In particular, Dretske's distinction between epistemic and non-epistemic seeing can readily be interpreted as a form of the distinction between conscious experience and belief about conscious experience. Whatever objections one may have against these representational theories — and most of them have been voiced by qualophiles, not exactly Dennett's bedfellows —, it is difficult to accuse them of succumbing to dualism.

Another reason Dennett advances against an appearance/reality distinction for consciousness is that this would open the possibility that some of your conscious experiences occur

unbeknownst to you, reality without appearance so to speak. However, Dennett points out, such conscious experiences would be inaccessible both from the first-person point of view and from the third person point of view. But then, "putative facts about consciousness that swim out of reach of both "outside" and "inside" observers are strange facts indeed" (1991: 133). Whence Dennett's 'first-person operationalism':

We might classify the Multiple Draft model, then, as *first-person operationalism*, for it brusquely denies the possibility in principle of consciousness of a stimulus in the absence of the subject's belief in that consciousness. (1991: 132)

In a nutshell, if a purportedly conscious state is neither accessible from the inside nor from the outside, how are we to distinguish it from an unconscious state? This is indeed the brunt of Dennett's objection to Dretske's talk of non-epistemic seeing. Dretske needs a difference that makes a difference between non-epistemic seeing which counts as visual experience and unconscious visual information which doesn't. According to Dennett, Dretske doesn't provide it, and he can't provide it because it doesn't exist.

On our reading, Dennett's doxological commitment equates what it is like to be us with what we believe it is like. Perhaps Dennett might prefer weaker versions of the commitment. For instance, one might claim that one cannot consciously experience a stimulus without believing that one is so conscious (first-order operationalism), but deny that merely believing that one has certain conscious experience implies having the experience itself. This would give one limited authority: in the presence of a conscious experience, one's beliefs about it cannot be but true, although other beliefs about what one experiences can be false.

On another version, ordinary beliefs about experience are fallible only in the sense that they embody a mistaken conception about what experience is. These beliefs picture experience as a mental state distinct from belief, which it is not. This is compatible with first-person operationalism and with some form of authority, for instance over the *contents* of one's first-order beliefs. Yet another version restricts dictatorial authority to what our *second-order* beliefs about our experience are, independently of whether they are true or false, thus leaving room for the claim that introspection is fallible.⁵

We cannot here argue in detail against each of these versions. Let us say that we doubt that any of them is a faithful representation of what Dennett means. Briefly, the first version does not motivate the required *epistemological* notion of authority, since the subject would still be unable to distinguish, among her second-order beliefs, those that are associated with real first-order experiences. The second and the third versions considerably weaken Dennett's radical claims about, respectively, fallibilism about introspection and first-order operationalism. Fallibilism about mode does not seem to capture what Dennett has in mind when he says that our introspective beliefs about our experience are highly fallible. When I mistakenly think that my visual experience presents me with a detailed picture-like representation of the world, it seems I am wrong about the content of my visual experience not just its mode. On the third version we considered, first-person operationalism would be restricted to second-order beliefs. But why should we have dictatorial authority over what our second-order beliefs are but not over our first-order beliefs?

In what follows, we offer two different lines of arguments which directly attack Dennett's doxological commitment, by motivating a distinction between conscious experience and

⁵ Here we are grateful to an anonymous referee for pointing out these last two readings of Dennett's position.

belief about such experience. These arguments show that what it is like cannot be (or at least cannot be fully) constituted by what the subject believes it is like.

2. Pathologies of perceptual consciousness

Not all cognitive scientists — not to mention philosophers — are of Dennett's conviction that a difference that makes a difference between unconscious states, conscious experiences and phenomenal beliefs is out of reach of third-person scientific methodology. First, as Dennett rightly emphasizes, heterophenomenology gives us much more data than just the subject's verbal reports; behavioral responses as well as all the covert, internal reactions that can be detected are part of the data. Second, inference to the best explanation is part of the scientific toolkit. Black holes are not among the data of science, yet we have good reasons to infer their existence.

Marcel (2003) provides an elegant illustration of how this could apply to the study of consciousness. Marcel argues in favor of a distinction between phenomenal experience and awareness of phenomenal experience (i.e., phenomenal belief in the terminology used here). Among his arguments for this distinction is the following, which is worth quoting at length:

This distinction is also illustrated by and explains the separate occurrence of blindsight (loss of conscious visual experience with preserved nonconscious vision), Anton's syndrome (unawareness of blindness), and Anton's syndrome with bilateral blindsight [...]. Blindsight patients have normal access to their visual experience, which in their case is absent in the scotomic area, which is why they deny seeing; yet they are shown to have nonconscious vision by accuracy of guessing [...]. In Anton's syndrome unawareness of blindsight is plausibly due to lack of access to absent visual experience.

In the third case, a patient assessed as totally bilaterally blind after traumatic impaction of the occipital poles was unaware of his blindness for 7 months. Yet during this time his ability to point to light targets and guess luminance levels correctly and other indirect tests all indicated blindsight, i.e. nonconscious vision without visual phenomenology. Apparently, he lacked access to the latter's absence. The difference between blindsight and unawareness of blindness appears to be in one's awareness of visual phenomenology or its loss. The presence of these three conditions and normal vision suggests a dissociation between nonconscious vision, visual phenomenology, and awareness of the latter. (Marcel, 2003: 173).

If Marcel is right, we should abandon Dennett's doxological commitment. For instance, the phenomenal belief that one does not see is not constitutive of the fact that one does not see. Patients with Anton's syndrome form the belief that they do not see, but this belief is not based on genuine access to a first-order (lack of) experience. It is disconnected from the first-order level and thus cannot be constitutive of it.

Marcel's data are impeccably heterophenomenological, consisting of the verbal reports of the patients, their behavioral responses to various tests, and data on their brain conditions and activities. The distinction between nonconscious vision, visual experience and awareness of visual experience is the product of an inference to the best explanation. We need this distinction, Marcel argues, in order to make sense of the difference between the three conditions that are blindsight, Anton's syndrome, and Anton's syndrome with blindsight. This distinction also does useful work in accounting for other phenomena, such as the split awareness of sensation that can be observed in anosognosia for hemiplegia, in the 'hidden observer' phenomenon in hypnosis or in general anesthesia induced by centrally acting

analgesics. Thus, the distinction has substantial explanatory power. In other words, what we have here is a difference arrived at using data and a mode of reasoning that are perfectly in line with the dictates of the heterophenomenological method. Moreover, it is a difference that makes a difference. What more could Dennett want?

Of course, Dennett could argue that since it is arrived at through a method of inference to the best explanation, it might one day be superseded by another better and more parsimonious explanation. Granted, but given that the distinction is arrived at by scientifically sound means and that so far we have no better alternative explanation, surely it should be provisionally accepted in the same way that, barring new evidence or a better explanation not yet in, astrophysicists accept the existence of black holes. Since Marcel's form of theorizing does not infringe the principles of heterophenomenology, Dennett's fallback line of resistance may be that his theorizing commits Marcel to the existence of entities of dubious metaphysical standing and thus runs against an ontological commitment to materialism. But once again, this is unwarranted. Nothing in what Marcel says commits him to the existence of sense data and the like.

So, how does Dennett's doxological commitment relate to his ontological and methodological commitments? It seems pretty obvious that it is not a direct consequence of either commitment. As we have seen, it certainly doesn't square well with the motivation Dennett gives for the neutrality of heterophenomenology, namely that we can be mistaken about our conscious experience and often are. In addition, as our discussion of Marcel's arguments for a distinction between phenomenal experience and awareness of phenomenal experience shows, one can be faithful to both materialism and heterophenomenology while eschewing Dennett's doxological commitment. Rather, it seems that Dennett's doxological commitment is tied to acceptance of

the Multiple Draft model of consciousness. Indeed, in one of the quotes given above, Dennett admits as much when he says that the *Multiple Draft model* is a form of first-person operationalism that denies the possibility in principle of consciousness of a stimulus in the absence of the subject's belief in that consciousness. But this model is intended as an empirical theory of consciousness, a theory whose value lies in part in its empirical adequacy. If the model cannot make sense of differences that make a difference, then surely there is something wrong with its empirical adequacy.

3. Belief and experience

In an exchange with Dennett (reported in Dennett, 2001a), Goldman summarizes one of their main points of disagreement as follows:

I realize that a 'doxological' (or representational) reductionist like yourself will want to reduce feeling states to dispositions-to-believe. A resistor like myself need not deny, of course, that feeling states do have a tendency to produce beliefs. The question is whether there are 'categorical' features of feeling states in virtue of which they have that tendency, or whether they are just pure doxological tendency and nothing else. I find the former view more compelling, and don't think that representational reductionism will work across the board. (Goldman, quoted in Dennett, 2001a: 10)

Along with Goldman, we think it unlikely that feeling states (i.e. conscious experiences) reduce to dispositions-to-believe. Unlike Goldman, however, we would not equate 'doxological' reductionism with representational reductionism, for we think one can be a representational reductionist without being a doxological reductionist. Indeed, our aim in this section is to motivate a distinction between experience and phenomenal belief that hinges on

differences in types of content. More specifically, we will argue that the content of perceptual experiences, although representational, is not fully conceptual or conceptualizable; instead it is non-conceptual. Furthermore, we will argue that some of the empirical arguments that have been adduced in favor of the idea that perception has non-conceptual content are perfectly kosher with respect to the methodological precepts of heterophenomenology, relying not on questionable intuitions but on differences between recognitional and discriminative capacities that are detectable by objective means.

For a start, let us state some rather obvious facts about the links between beliefs and concepts. Although hardcore realists about propositional attitudes and moderate realists like Dennett have serious disagreements about their ontological status, they at least share the view that it only makes sense to attribute propositional states, including beliefs, to a system insofar as by so doing we can explain and predict its behavior. In other words, if it didn't have some explanatory or predictive leverage, the practice of attribution would be purely gratuitous. Now the practice of attribution can only have such explanatory or predictive leverage if the transitions between attitudes are rational transitions and transitions between propositional attitudes are rational insofar as they are content-dependent.

What needs explaining is how these transitions can be content-dependent. This is where concepts step in: the content of a propositional state is a function of its conceptual constituents and the rational transitions between propositional contents are in turn a function of their conceptual constituents. As Crane (1992) aptly puts it, concepts are the inferentially relevant constituents of intentional states, they account for the inferential relations among beliefs and more generally, among propositional attitudes. To borrow a famous example from Stich (1983), someone who believes that President McKinley was assassinated is expected to

believe that President McKinley is dead. This rational transition between the two beliefs is itself a function of the inferential link between the concept of assassination that figures as a constituent of the content of the first belief and the concept of death that is a constituent of the content of the second belief. Indeed, Mrs. T. who claims to remember that President McKinley was assassinated, but denies knowledge of whether he is now dead or alive can be diagnosed as lacking mastery of the concept of assassination.

In short, the argument runs as follows: (1) it is an essential feature of propositional attitudes that they play a certain inferential role; (2) concepts are the constituents of the content of these states that are responsible for their playing this role; hence (3) propositional attitudes in general, and beliefs in particular, have conceptual content. But if the content of beliefs is by definition conceptual and we can show that perceptual experience has content that is not fully conceptualizable, then we will also have shown that perceptual experiences do not reduce to dispositions to belief. To this task we now turn.

According to the conceptualist approach in the philosophy of perception, perceptual content is fully conceptual. Against this view, nonconceptualists have argued that perceptual experiences have a fineness of grain that far outstrips what can be captured in terms of concepts possessed by the perceiver. The argument from fineness of grain should not be confused with the argument from richness.⁶ The argument from richness of content refers to the putative fact that in perceptual or at least visual experience we enjoy a richly detailed picture-like representation of the world around us. Fineness of grain relates to the fact that there are many dimensions and properties— hue, shape, spatial magnitudes, loudness, pitch,

⁶ One reason this distinction is often overlooked is that the same phrase, namely ‘the analog character of perceptual content’, is sometimes used to refer to the fineness of grain of perceptual content, sometimes to its richness. Thus, while Peacocke (1992) uses the phrase to refer to fineness of grain, Dretske (1981) uses it to refer to richness.

sourness, acidity— such that any value on that dimension may enter the fine-grained content of an experience.

To the argument from richness it has been objected that the putative phenomenological fact on which it is based is actually a ‘grand illusion’ (Noë *et al.*, 2000). Recent work in psychology on change blindness, inattention blindness and related phenomena seems to call into question the idea that we really enjoy perceptual experiences which represent the environment in rich detail. Whether and in what sense we should really be skeptical that we really enjoy such richly detailed visual experiences is at present a hotly debated question.⁷ However, we won’t go into this debate here, but instead will focus on the argument from fineness of grain. The argument from fineness of grain and the argument from richness are independent and the former may hold even if the latter is rejected.

The argument from fineness of grain exploits the idea that we can normally make perceptual discriminations that are much more fine-grained than the discriminations for which the subject has corresponding concepts. In other words, perceptual representations have a fineness of grain that far exceeds what can be captured by using concepts possessed by the perceiver. One often used illustration of this phenomenon is that of color discrimination. Normal subjects can perceptually discriminate many more shades of color than they have concepts for such shades. The same remark applies to many other properties such as shapes, spatial magnitudes, sounds, and smells. For instance, to borrow an example from Peacocke: “If you are looking at a range of mountains, it may be correct to say that you see some as rounded, some as jagged. But the content of your visual experience with respect to the shape of the mountains is far more specific than that description indicates. The description involving the

⁷ See Noë (2002) for an in-depth investigation of this issue.

concepts round and jagged would cover many different fine-grained contents that your experience could have, contents that are discriminably different from one another” (1992: 67-68). The conclusion the nonconceptualists have wanted to draw is that since the range of discriminable contents perceptual experiences can have goes far beyond the ranges picked out by concepts in the perceiver’s repertoire, these perceptual contents must be nonconceptual.

The conceptualist reply to this argument, offered by both McDowell (1994) and Brewer (1999), is to acknowledge the fine-grainedness of perceptual experience and yet contend that it can perfectly be captured by appropriate demonstrative concepts. According to them, the nonconceptualists’ argument rests on an unduly narrow view of concepts, which restricts the concepts available to capture perceptual discriminations to those associated with verbal expression and having context-independent norms of application. Once this restriction is lifted, it becomes possible to deny that a perceiver capable of a perceptual discrimination between, say, two shades of red for which he lacks correspondingly different context-independent color concepts has an experience with nonconceptual content. It is perfectly possible for her to capture this difference in her perceptual experience of the two shades in terms of demonstrative concepts like *colored thus* or *this shade*, that exploit the presence of samples of the shade in question.

For this strategy to be successful, the conceptualists should demonstrate that the purported demonstrative observational concepts (DOCs for short) they claim can capture the fine-grained content of color experience satisfy central criteria of concepthood and hence really qualify as concepts. Yet, as we will show, their DOCs fail to meet highly plausible criteria for concepthood that form part of the conceptualist’s own conception of concepts.⁸

⁸ Some of the arguments in what follows are adapted from Dokic and Pacherie (2001).

Let us start with some empirical data. There is a wealth of evidence from both psychophysics and perceptual psychology that our capacity for perceptual discrimination far surpasses our capacity for perceptual recognition. (Burns and Ward, 1977; Halsey and Chapanis, 1951; Hardin, 1988; Hurvich; 1981). For instance, Hardin (1988) points out that although we can discriminate something on the order of ten millions colors, when the task is to identify a color in the absence of a reference standard, even for fairly good observers positive identification is of the order of about thirty. Quite a drop! To put it vividly, although you may be very good at discriminating two very similar shades of red when they presented simultaneously, you are more than likely to be at chance when they are presented in succession and separated by a mask. The same difference is also found in pitch perception. Although we can make very fine-grained pitch and interval discriminations, even trained listeners are not, in general, able to recognize reliably intervals to a finer degree than chromatic semitones (Raffman, 1993).⁹

We can now voice our concerns with the conceptualist claim that the content of perceptual experience can be fully captured in terms of DOCs. First, it seems plausible to suppose that mastery of an observational concept implies a capacity to identify and reidentify perceptual instances of the concept. This means that to have a concept of a perceptually accessible property, one must be able to recognize instances of the property over time, not just to discriminate between something that instantiate that property and something that doesn't. Obviously enough, one can recognize only what one can remember. But empirical evidence makes it clear that perceptual memory is limited and its grain much coarser than our

⁹ Persons with absolute pitch may constitute an exception. Yet, it is estimated about 3% of the general population in Europe have absolute pitch and this figure only increases to about 8% in professional or semi-professional musicians. Furthermore, it seems that the ability exists either from early childhood or not at all. Although with considerable effort and extensive training an adult can acquire perform at a level that is close to that of persons with absolute pitch, this painfully acquired skill turns out to be quite fragile, i.e., it gets rapidly lost when practicing on tone recognition is terminated (Terhardt, 2000).

perceptual discrimination thresholds. It follows that if concept possession requires a certain recognitional capacity, the maximal fineness of grain of our perceptual concepts will correspond to the maximal fineness of grain of perceptual memory encoding. It is overwhelmingly unlikely that the conceptualist's DOCs meet this constraint and hence are associated with a recognitional capacity.

Of course, the conceptualists could deny that a recognitional requirement should be imposed on DOCs. Yet there is a price to pay for this move and one may doubt whether it is worth paying. One important reason for insisting that concept possession requires an associated identification or recognition capacity has to do with the inferential role of concepts. As we argued earlier, concepts are plausibly thought to be what accounts for the inferential powers of our beliefs and other propositional attitudes. In most cases, what explains the validity of an inference is the presence of certain conceptual constituents as parts of the content of the premises. To borrow an example from Crane (1992), a thinker who believes that *a is F* and *b is F* and that *a is not b* will be disposed to believe that *at least two things are F*. What is essential for the inference to be valid is that the same part of the content should occur in the first two premises, in other words they should both contain the concept *F*. But in the absence of any genuine recognitional capacity associated with DOCs, how are we to ensure that the same DOC figures in both premises? The only case in which such insurance can be given is when the two premises are simultaneously available to the thinker in perceptual experience in such a way that he can attend to both at once. As soon as the two premises are obtained separately, the warrant disappears, since, for lack of recognitional capacity, the thinker will not be in a position to ascertain whether the DOCs involved in the two premises are the same or different. The inferential potential of such concepts is therefore extremely restricted. These

concepts have neither past nor future and their use in reasoning is confined to the here and now of perceptual experience.

So far we have argued that lack of diachronic recognition for DOCS severely limits the inferential potential of demonstrative judgements involving them and seriously threatens the conceptualists' claim that they are genuine concepts. But the worse is still to come for even synchronic recognition is problematic.

According to a plausible Criterion of Difference, two concepts (considered as the senses of predicates) are different if the subject who grasps them at the same time can rationally adopt different epistemic attitudes toward thoughts containing them (which are otherwise identical). For instance, if it is possible for a rational subject to believe (in a particular context) that everything which is thus_a is thus_a , where both tokens of ' thus_a ' express the same DOC grounded on the perception of a , while not believing that everything which is thus_a is thus_b , where the token ' thus_b ' expresses a DOC grounded on the perception of b , then the concept expressed by ' thus_a ' is not the same as the concept expressed as ' thus_b '. The Criterion of Difference for concepts implies that if a subject grasps two concepts at the same time, she is at least sometimes in a position to know that the concepts are different. That is, if she can rationally adopt different epistemic attitudes toward certain thoughts containing them, she knows that they are different.

It is a well-known fact that perceptual indiscriminability is non-transitive. For instance, color samples (or pitches) a , b , and c can be found such that, for a normal observer, b is indiscriminable from a and c is indiscriminable from b and yet c is discriminable from a . Now suppose that a subject, who is aware of this empirical fact, perceives two samples of the same

color shade – say, a_1 and a_2 . It is natural to suppose that the DOCs grounded on the perception of these samples are the same. After all, not only do they *seem* to be colored exactly alike, but they really have the same color shade. However, it seems always possible for the subject to believe that everything which is thus_{a_1} is thus_{a_1} , while doubting whether everything which is thus_{a_1} is thus_{a_2} . For all she knows on the basis of perceptual appearance, there might be a sample that is indiscriminable from a_2 but discriminable from a_1 . Such a sample would fall under thus_{a_2} but not under thus_{a_1} . This is a coherent epistemic possibility, which implies, according to the Criterion of Difference, that ‘ thus_{a_1} ’ and ‘ thus_{a_2} ’ express different concepts (in this context).

If we accept the Criterion of Difference for concepts, then, it seems impossible to grasp the same DOC through the perception of numerically distinct objects that look exactly the same as far as their most specific color is concerned (or, *mutatis mutandis*, through the perception of distinct auditory events that sound exactly the same as far as their pitch is concerned). In fact, there is an infinite number of DOCs for a given shade or pitch, since DOCs are necessarily tied to particular samples. So either conceptualists admit that there are conceptual differences in the contents of perception which do not correspond to any phenomenological differences, or they slice the phenomenal world *too* finely. Beside the fact that neither option is intrinsically plausible, the normal experience of perceiving internal relations between different shades presented at a given time (for instance, perceiving that two samples have exactly the same shade or two sounds the same pitch) has not been provided for.

Thus, DOCs have implausible conditions of individuation and violate an intuitive Criterion of Difference for concepts. Moreover, they are not associated with recognitional capacities and their inferential potential is severely restricted. We conclude we are left with no reason to admit that DOCs qualify as concepts. Let us make clear, however, that we do not deny the

coherence of context-dependent classifications. On the contrary, a judgement expressed by 'That is thus' in a given context can have a fully conceptual content. What we claim is that although such a content is based on the perception of things and their qualities, it cannot substitute for the content of perceptual experience itself. The modes of presentation expressed by demonstrative predicates like 'is thus' cannot both satisfy the essential constraints on the individuation of concepts and capture the phenomenology of perceptual experience. So the content of perceptual experience cannot be fully captured by concepts, even if we allow for conceptual, context-dependent classifications of things.

The relevance of the foregoing discussion to the evaluation of Dennett's doxological commitment is rather obvious. The fineness of grain of our perceptual experiences makes an essential contribution to what it is like to enjoy them. Our point about diachronic recognition is that phenomenal beliefs cannot constitute the *whole* of what it is like to have perceptual experiences. There are phenomenal differences which cannot be captured at the level of beliefs, at least if beliefs are partly defined in terms of an inferential network of beliefs. Moreover, these differences can be manifested in controlled behaviour, which makes them accessible in principle using the methods of heterophenomenology.

The point about diachronic recognition is *prima facie* compatible with the claim that our phenomenal beliefs cannot be false. What it shows is just that they cannot provide a complete description of what it is like to have experiences. Our discussion of synchronic recognition goes a step further. Blind adherence to the doxological commitment runs the risk of slicing the phenomenal world too finely. So if our phenomenal beliefs were really constitutive of what they describe, our phenomenal world would be very different from the one we actually live in. This comes close to being a *reductio* of Dennett's commitment.

Some final considerations. One might object that our criteria of concepthood are too stringent. Noë (2004) has argued that the sensori-motor skills that enable us to experience sensory qualities should already be considered as conceptual (or "proto-conceptual"). In our view, sensori-motor skills cannot ground by themselves concepts that are sophisticated enough to play a substantial role in inferences. However, perhaps the issue is at least partly terminological. Suppose Noë is right and the content of perceptual experience is conceptual in a loose sense of the term. Then our claim should be rephrased as the claim that some conceptual contents cannot be the contents of beliefs and other high-level propositional attitudes. Perception would have a conceptual content, but one which is not substantive enough to support full-blown inferential interactions with the contents of genuine propositional attitudes. The gist of our criticism of Dennett's doxological commitment would be unaffected.

Alternatively, one might object that our criteria of beliefhood are too stringent. Dennett might allow that perceptual discrimination involves a rather cognitively primitive form of belief, one which is not necessarily associated with recognitional and reidentification capacities. Once again, the issue would seem to be purely terminological. But some terminologies are more perspicuous than others. The term 'belief' would label mental states with radically different functional roles and would lose most of its theoretical usefulness.

4. Conclusion

In this paper, we have put forward two arguments which cast doubt on Dennett's claim that our phenomenal beliefs are somehow constitutive of what they are about, including what it is like to have conscious perceptual experiences of the world. First, a coherent account of consciousness can motivate a threefold theoretical distinction between unconscious

experiences, conscious experiences, and awareness of conscious experiences. Second, conscious perceptual experiences have contents which cannot be fully captured at the level of beliefs. Both arguments are compatible with materialism, and neither introduces dubious mental intermediaries between mind and world. Our opinion is that they put Dennett's insights about the relevance of heterophenomenology to the study of consciousness in a much better light.

References

- Brewer, B. 1999. *Perception and Reason*. Oxford: Oxford University Press.
- Burns, E. M. and W. D. Ward. 1977. Categorical perception — phenomenon or epiphenomenon: Evidence from experiments in the perception of musical intervals. *Journal of the Acoustical Society of America*, 63: 456-68.
- Chalmers, D. J. 1996. *The conscious mind*. New-York: Oxford University Press.
- Crane, T. 1992. The non-conceptual content of experience. In *The Contents of Experience*, ed. T. Crane, 136-157. Cambridge: Cambridge University Press.
- Dennett, D. C. 1991. *Consciousness explained*. Boston: Little, Brown & Co.
- Dennett, D. C., 2001a. The Fantasy of First-Person Science.
<http://ase.tufts.edu/cogstud/papers/chalmersdeb3dft.htm>.
- Dennett, D. C. 2001b. Are we explaining consciousness yet? *Cognition*, 79: 221-237.
- Dennett, D. C. 2003. Who's on first? *Journal of Consciousness Studies*, 10, 9-10: 19-30.
- Dennett, D. C. 2005. *Sweet dreams. Philosophical obstacles to a science of consciousness*. Cambridge, Mass.: MIT Press.
- Dokic, J. and Pacherie E. 2001. Shades and concepts. *Analysis*, 61/3: 193-201.
- Dretske, F. 1995. *Naturalizing the Mind — The Jean Nicod Lectures - 1994*. Paris: Editions du CNRS & Cambridge, MA: MIT Press.
- Goodman, N. 1951. *The Structure of Appearance*. Indianapolis: Bobbs-Merrill.
- Halsey, R. M. and A. Chapanis. 1951. Number of absolute identifiable hues. *Journal of the Optical Society of America*, 41: 1057-8.
- Hardin, C. L. 1988. *Color for philosophers*. Indianapolis: Hackett.
- Hurvich, L. M. 1981. *Color vision*. Sunderland, MA: Sinauer Associates.
- Marcel, A. J. 2003. Introspective report. *Journal of Consciousness Studies*, 10, 9-10: 167-86.

- McDowell, J. 1994 (second edition 1996). *Mind and World*. Cambridge, MA: Harvard University Press.
- McGinn, C. 1989. Can we solve the mind-body problem? *Mind* 98:349-66.
- Nagel, T. 1974. What is it like to be a bat? *Philosophical Review* 4:435-50.
- Noë, A., Pessoa, L., & Thompson, E. 2000. Beyond the grand illusion: what change blindness really teaches us about vision. *Visual Cognition*, 7: 93-106.
- Peacocke, C. 1992. *A Study of Concepts*. Cambridge, MA: MIT Press.
- Raffman, D. 1995. On the persistence of phenomenology. In *Conscious Experience*, ed. T. Metzinger, 292-308. Thorverton: Schöningh /Imprint Academic.
- Stich, S. 1983. *From Folk Psychology to Cognitive Science*. Cambridge, MA: MIT Press.
- Terhardt, E., 2000. Absolute pitch. <http://www.mmk.e-technik.tu-muenchen.de/persons/ter/top/absolute.html>.
- Tye, M. 1995. *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind* Cambridge, Mass: MIT Press.