

Thomas W. Clark

Killing the Observer

‘Nor is it any longer clear how to understand the notion of our grasping the “simple facts of consciousness” from the perspective of the first person.’¹

Abstract: *Phenomenal consciousness is often thought to involve a first-person perspective or point of view which makes available to the subject categorically private, first-person facts about experience, facts that are irreducible to third-person physical, functional, or representational facts. This paper seeks to show that on a representational account of consciousness, we don’t have an observational perspective on experience that gives access to such facts, although our representational limitations and the phenomenal structure of consciousness make it strongly seem that we do. Qualia seem intrinsic and functionally arbitrary, and thus categorically private, because they are first-order sensory representations that are not themselves directly represented. Further, the representational architecture that on this account instantiates conscious subjectivity helps to generate the intuition of observerhood, since the phenomenal subject may be construed as outside, not within, experience. Once the seemings of private phenomenal facts and the observing subject are discounted, we can understand consciousness as a certain variety of neurally instantiated, behaviour controlling content, that constituted by an integrated representation of the organism in the world. Neuroscientific research suggests that consciousness and its characteristic behavioural capacities are supported by widely distributed but highly integrated neural processes involving communication between multiple functional sub-systems in the brain. This ‘global workspace’ may be the brain’s physical realization of the representational architecture that constitutes consciousness.*

1: Introduction

In characterizing consciousness, it is often said that there exists a first-person perspective or point of view associated with having phenomenal experience. On some construals of this perspective, the subject gains knowledge of,

Correspondence: Email: twc@naturalism.org

[1] Thomas Metzinger (1995a) p. 20.

acquaintance with, or access to certain categorically private first-person facts, the phenomenal ‘what it is like’ of experience, or qualia (Nagel, 1974; Jackson, 1982; Chalmers, 1995a). It is supposed that such facts about experience (the redness of red, the painfulness of pain) are not reducible or explicable in terms of third-person, objective facts about brains, neurons, patterns of excitation, and other researchable aspects of cognitive states. Nothing about such physical, functional, or representational states of affairs implies that qualia should feel precisely as they do to a particular subject, or that representational states should feel any way at all, in which case the particular qualitative looks and feels of sensory experience certainly seem a realm apart from what science can predict and explain. This difficulty is what David Chalmers dubbed the ‘hard problem’ of consciousness.

The challenge for those who seek unification of the apparently disparate realms of qualitative consciousness and scientific objectivity is to show that, despite appearances to the contrary, the phenomenal is entailed by the functional-representational (if in fact it is), and that all actual facts about experience are third-person facts. In what follows, I will pursue such unification by suggesting, taking a page (or several) from Daniel Dennett, that what seem to be non-functional, categorically private facts about experience are indeed explicable as *seemings*, not facts, seemings generated by the way in which consciousness comes to be. The key to all this, I will argue, is that as subjects we don’t have a first-person perspective on experience, even though as persons we most certainly have a unique cognitive *agent* perspective on the world and a unique cognitive connection to our bodies. To understand consciousness, we must extirpate any lingering notion that we witness experience, or to put it somewhat melodramatically, we must kill the observer.

2: The Observational First-Person Perspective

The notion of a first-person perspective, when construed in a certain sense, arguably helps to perpetuate the intuition that experience includes categorically private facts, facts that are inaccessible and unsubsumable by any sort of shared, objective, third-person understanding.² This construal is that in undergoing experiences, the person is in a literally perceptual relation to experience itself. The person somehow witnesses or observes experience such that it becomes a private presentation, involving a set of categorically private phenomenal facts, namely the qualitative facts about what experience is like. On this picture, phenomenal feels or qualia have a non-functional, non-representational, and by implication, non-physical aspect involving private facts about experience that only the individual can access.

Such an observational perspective is suggested by locutions that crop up in the literature which place the experiencer in a privileged position of seeing,

[2] For a fact to be irreducibly first-person and private bars it from scientific categorizations or descriptions that are prototypically objective. As Dennett (2001b) suggests, there can’t be a first-person *science* of consciousness.

accessing, or directly apprehending facts about experience. For instance, Thomas Nagel remarks that ‘For if the facts of experience – facts about what it is like *for* the experiencing organism – are accessible only from one point of view, then it’s a mystery how the true character of experiences could be revealed in the physical operation of that organism’ (Nagel, 1979, p.172 original emphasis). Similarly, Nagel says that ‘It is difficult to understand what could be meant by the *objective* character of an experience, apart from the particular point of view from which its subject apprehends it’ (Nagel, 1979, p.173, original emphasis). Now, to *apprehend* an experience, or have special access to facts about it, might suggest that the subject is in a unique observational or perceptual relation to it, which in turn implies that certain facts about that experience might be privy to the subject alone. As John Biro (1993, p. 180) puts it, ‘A point of view, it is claimed, gives its owner access to a special kind of fact that is different from, and irreducible to any other fact or set of facts equally available to others.’

A more recent expression of the observational perspective is found in Max Velmans’ idea that first- and third- person perspectives offer two complementary *views* of experience, views that depend on what he calls ‘observational arrangements’ (Velmans, 2002, p. 11). The subject is in a position to observe experience in a way that no outside observer can, so that, for instance ‘*Other* people’s experience might be hypothetical constructs, as we cannot observe their experiences in the direct way that we can observe our own...’ (p. 22, original emphasis). The difference in observational arrangements produces two different sorts of facts or information about experience, the first-person information about phenomenology and the third-person information about the brain’s representational mechanisms (p. 15).³ These views or perspectives and the facts they support are, Velmans claims, mutually irreducible – no reduction of the phenomenal to the physical or functional is possible. Similarly, Steven Lehar (2004) proposes what seems a modified sense data theory of consciousness (see Dretske 1995, pp. 128–9), in which experience, from the first-person point of view, is the perception of what he calls ‘internal effigies’ or ‘percepts’. He says, for instance, that ‘We cannot ... in principle have direct experience of objects in the world itself, but only of the internal effigies of those objects generated by mental processes’, and that ‘consciousness is indeed observable...because objects of experience are first and foremost the product or “output” of consciousness, and only in a secondary fashion are they also representative of objects in the external world’. Lastly, Antonio Damasio speaks, undoubtedly metaphorically, of experience as a ‘movie in the brain’ which is composed of mental images generated by neural processes (how neural processes do this, he admits, isn’t clear). But he slips from metaphor to literalism in supposing that we indeed observe such images, not the world, when we have experience: ‘The image *we see* is based on changes that occurred in our

[3] Similarly, in an earlier paper (1996) Velmans says that ‘Once one accepts that consciousness and its contents (viewed from a first-person perspective) provide valuable psychological data, one can get on with the business of working out how given conscious states *relate to* given forms of processing in minds or brains (viewed from a third-person perspective) in a way that does not prejudge either the ontological nature of consciousness, or its causal status’ (original emphasis).

organisms...’ (Damasio, 2003, pp. 198–9, my emphasis). William Lycan calls the supposition that experience is a perceived *object* ‘the banana peel’ since, as he puts it, ‘anti-materialists typically...slip on it into the Movie Theater Model of the mind’ (1987, p. 17), as do, it seems, some unwary materialists.

This notion of having an observational perspective on experience sets up a realm of private phenomenal facts that one observes about one’s experience, facts that can never become shared, objective knowledge as specified by science. But, if we consider what it means to observe and to be in possession of facts about things, and if we consider our phenomenological situation, including both the qualia of experience and the person who has them, I think it will become clear that we observe and know the *world*, not experience.

One long-standing approach to the issue of phenomenal facts has been via Jackson’s ‘knowledge argument’ (Jackson, 1982), in which anti-reductionists hold that someone conversant with all the neurophysical facts that correlate with an experience of, say, red (the philosopher’s archetypical quale) nevertheless learns a new, non-physical fact when she first *experiences* red. There are good replies to Jackson’s original argument in the literature (e.g., Van Gulick, 1993; Levine, 1993; Biro, 1993; Tye, 1995) and Jackson himself has abandoned it in favor of representationalism (Jackson, 2001), so I’ll for the most part avoid these well-worn paths and instead undertake a deliberate consideration of qualia and the subject to whom they might appear.

3: Informational Characteristics of Qualia

What precisely *are* qualia — the looks and feels of sensory experience — and why might they be supposed to incorporate private, non-representational, and non-functional facts? This is to take seriously the question, for instance, ‘what *is* it like to see red?’. In specifying ‘what it’s like’ we must pinpoint what *about* qualia is supposed to resist explanation in terms of representational, informational functions, for it turns out that there’s a good deal which is either explanatorily unproblematic, or at least unproblematically *conceivable* on a representationalist account.

Most, if not all, qualia are occurrently determinate, stable, non-conceptual *values* within the various modes of conscious sensation. By this I mean simply that, for instance, each colour quale has a particular place in the phenomenal structure of my colour experience as defined by its relations to other hues. The phenomenal character of the blue of my mouse pad is one of a huge number of *particular* hue values found in colour experience whose specific place lies somewhere on the continuum of distinguishable blues.

Let us provisionally adopt, on the basis of the representationalist paradigm, what I will call for convenience the *informational hypothesis* — that sensory qualia are representational contents embodied by world-responsive, neurally-instantiated, multi-dimensional representational phase or state spaces that feed information to higher-level perceptual representations (Churchland, P.M., 1989, Ch 9; Churchland, P.S., 1988, pp. 455–7; Tye, 1995, pp. 101–3, 119, 138). On

this assumption, a particular phenomenal hue corresponds to the cluster of specific values of each component dimension of hue state space, and its relations to its cousins are fixed by the relative proximity or remoteness to other hues along these dimensions. As a non-conceptual, primary sensory representation, it constitutes information fed to the sorts of categorizing, higher-level belief systems that constitute concept-based cognition (Tye, 1995, p. 104).

A singular, particular quale, or particular quality of experience, is just that which you can discriminate in a particular situation as admitting of no further detail; it's *non-decomposable* and homogeneous, having no further phenomenal structure. Whatever mode of sensation we choose to explore, we can find singular, monadic, and homogeneous elements of experience, for instance the uniform cream colour of a particular portion of my door. On the informational hypothesis, this quale is non-decomposable since our colour sensing systems have a finite resolution, such that (very crudely) when wavelengths within a particular frequency band impinge on a certain portion of the retina, the same set of colour state-space values is assigned to all points within a corresponding portion of the internal visual map that contributes to an object representation and that in turn partially constitutes my experience. This account of qualitative homogeneity helps explain the seeming *intrinsicity* of qualia, often thought to be a barrier to their reduction in terms of relational processes and states (Clark, 1995; Feser, 2001).

Significant regarding sensory qualities is that, although we can think of them as basic, non-conceptual and non-decomposable bits of phenomenology, they are usually experienced as belonging *to* or associated *with* phenomenally complex objects and events. Qualia are ordinarily experienced not as features of experience or representations, but as characteristics *of* things that we perceive or sense such as the colour of the door, the pain in my back, the taste of the coffee, the sound of a car going by (Van Gulick, 1993, p. 149; Metzinger 1995a, p. 11–12). They are essential in forming the often unexpressed, but expressible *conceptual* judgments that something is the case, e.g., that the door is cream-coloured. The informational hypothesis helps to explain this fact about qualia in terms of higher level binding processes: lower level, sensory state-space values contribute specific information — intentional content — to higher-level, occurrently bound, and behaviourally useful perceptual object representations. At the highest level, these representations are amalgamated into the experience of objects within a single, more or less coherent and predictable world with the subject at its centre (Metzinger, 2000a). So qualia are all assigned, more or less, to the larger perceptual ensembles that figure in moment-to-moment experience.

4: The Essential Characteristic of a Quality

Thus far there is nothing in this account of qualia that doesn't seem potentially explicable on the informational hypothesis, although of course nothing in the sketch above counts as a serious attempt at specifying the mechanisms that could instantiate these informational characteristics of phenomenal qualities (for a bit

more on such explanations see section 9). But this picture of qualia has seemingly left out what anti-reductionists contend is the characteristic which most resists functionalization: the ‘way it is for me’ phenomenal *feel* of a quality. True, informational functions might specify world-tracking intentional content that contributes to representations which are necessary to guide behaviour, but an explanation of why such content assumes a categorically *phenomenal* aspect seems nowhere to be found on the informational hypothesis. Saying what this categorically phenomenal aspect *is*, precisely, is notoriously difficult of course, but the mark of basic phenomenal particulars is more or less that they involve something simply and essentially qualitative, *something that itself defies further qualification*. It’s something quite specific, but having no internal structure, it can’t be further described and therefore seems *arbitrary* with respect to any informational function qualia might serve (Levine, 1983, p. 359; Van Gulick, 1993, p. 143–4). It’s not just that blue differs from red in consistent ways, thus permitting reliable discriminations and categorizations of objects, but that it has its own unique, seemingly intrinsic qualitative look that, apparently, could have been different while serving the same discriminative role. In fact, every discriminable quality in my experience (and there are thousands of such qualities) appears or feels to me a certain ineffable, non-relational way. Nothing about the informational hypothesis seems capable of specifying this rather large set of facts, facts ostensibly about my particular experiences, not about the world that experience might represent. Nor does it seem capable of saying why represented content must appear any way *at all* to a perceptually and cognitively adept system.

The difficulty for representationalism, then, is what I will call the *essential characteristic* of a quality (as have others, e.g., Nagel, 1979, p. 175 [footnote]; Metzinger, 1995, p. 15): being a particular, ineffable, seemingly intrinsic and functionally arbitrary *way* for a particular subject. But I want to reiterate the fact that qualia are particular values within sensory modalities, and that this particularity is a straightforwardly third-person fact, since as conscious creatures we all agree that my blue mouse pad (could you observe it) looks a definite distinguishable way. So it’s not particularity *in general* that’s the difficulty but the particular subjective look of blue to each of us as separate conscious individuals. This, it seems, is a fact not about my mouse pad, but about each subject’s particular *experience*. The question before us is whether or not the essential characteristic of a quality — that it is like *this* for *me* — involves categorically private first-person facts about experience over and above third-person facts about representation.

5: The Unspecifiability of Private Facts

The complaint often lodged against reductionist accounts of consciousness is that no third-person description of a system that subserves consciousness, however elaborate, can *capture* these facts (Jackson, 1982; Nagel, 1986, p. 15; Flanagan, 1992, p. 117). As Chalmers (1999) puts it, ‘I also take it that

first-person data can't be expressed wholly in terms of third-person data about brain processes and the like... That's to say, no purely third-person description of brain processes and behavior will express precisely the data we want to explain...' (quoted in Dennett, 2001a). But, given the ineffability of qualia, it's important to see that *no first-person description does any better*. Like you, I can reliably recognize and report qualities as they occur in experience and place them in relations to their qualitative cousins, but I can't descriptively specify what it is *about* a particular blue that makes it appear precisely the way it does to me. Likewise for pain: I can't specify what it is *about* pain that makes it painful.

To see the difference between ostensibly private qualitative facts and public facts, imagine we are standing before a blue chair. That a blue chair is in front of us counts as a third-person fact, since we can agree on a description: it's a chair, it's blue, and it's here. We've achieved inter-subjective consensus and now share collective knowledge about the world, so *that* chair is the *same* chair for both of us. But on the first-person-fact understanding of qualia, the what-it's-like-to-you of your blue might be different than the what-it's-like-to-me of my blue. Since neither of us can offer a description of our blues which could confirm or disconfirm a difference, they remain seemingly private, first-person facts incapable of third-person verification.

However, the unspecifiability of the essential characteristics of qualia, though it may drive the intuition of privacy, is at least somewhat troublesome for their status as private *facts*. That the chair is blue is an indisputably informative fact about the chair. That my private blue is like *this*, and possibly unlike your blue, is a claim that sounds as if it had informational content, but as we've seen, apart from the story about the relations of a particular quale to its cousins (relations which are the same for all of us with similar representational capacities), there is no further informational story to be told about the essential characteristic of a particular quale, even to ourselves as the presumptive cognizers of such facts. If a purported fact ('my blue is like *this*') delivers no informational content to its possessor, then one wonders if it's a fact at all. Is a fact that delivers no content and no knowledge indeed a fact? Remember, casting aspersions on what-it's-like-for-me blue's ambition to be a private fact is not to deny that blue looks a particular way to me, or that I can reliably recognize and name it on the basis of such a look, or that it plays an essential informational role in my cognitive economy. It's only, and strictly, to suggest that the way it looks to me does not involve a private fact, and (generalizing to all qualia) that there is no categorically private first-person information, available to the subject alone, to be gleaned about consciousness. Put another way, it's to say that the demonstrative 'this', when referring to the way blue looks to me, refers strictly to informational content *about the world, not my experience*, content that gets misconstrued as a private fact accessible only to the subject.

6: The Limits of Representation

To flesh out the story of why qualia might not involve private facts, it will be helpful to discuss briefly our situation as cognitive creatures. As organisms (or more generally, intelligent systems), we observe, in a straightforward and unequivocal sense, the world around us. (I'll stick with the visual modality in what follows, although we could tell the same story about observation in terms of hearing, touch, echolocation, or any world-responsive sensory modality that participates in perception of the world outside the body.) Light reflects off objects, enters the eye, and produces a cascade of neural events that ends up contributing to a freshly updated, neurally-instantiated representation of those objects which helps us get around successfully in our environment.⁴ This is what constitutes *directly* seeing the world, being in a *direct* perceptual or observational relation to it, and, ultimately, possessing facts and knowledge about it: the creation of informational, intentional content by incoming stimulation which helps determine values within the various representational state spaces that make up the visual system, values which (crucially) get integrated with other content into higher-order, suitably bound object representations that subservise behaviour. To be in possession of facts is for a cognitive system to consist of (partially) and have access to such behaviourally useful representations, useful because they reflect both the way the world is in some respect, or regularities it manifests, *and* the needs of the organism.

Observing the world and coming to know facts about it are a matter of getting the intentional content of one's representational system to be constrained (mostly) *by* the world, although of course it's still a highly selective representation. Thus constrained, the content ends up being *about* the world, and less about how you'd like it to be, always a good strategy for survival. Such content ranges from first-order, conceptually indeterminate content of sensory experience (e.g., my door is *that* colour) all the way up to higher-order, abstract conceptual content (e.g., knowledge is justified true belief). But whether we're considering non-conceptual or conceptual knowledge, the story about the organism observing the world and knowing facts about it ends here, in that there is no further inner observer of representational content. Although there might *seem* to be such an observer (of which more below in section 8), to take Dennett's line, representations don't have to be literally *witnessed or appreciated* by anyone or anything outside the network of representations to be efficacious (Dennett, 1991, 2001a).

This means that when it comes to the content delivery systems of the various modes of sensation, these too are unwitnessed and unappreciated. We don't have, for instance, a built-in representational apparatus that represents the *process* of colour-coding which then delivers content *about* that process to help guide behaviour with *respect* to that process. Although meta-representation of representational processes often occurs in various higher-level cognitive

[4] Although they get short shrift in my discussion, I do not mean to minimize the role of internal schemas that provide default representational assumptions for the system, nor am I making a claim about the moment-to-moment completeness or detail of internal representations as discussed, for instance, in Noe (2002).

contexts (as in applying concepts when reporting experience, e.g., ‘I’m in pain’, or as exemplified by this paper) we simply don’t have a behavioural need for information about how we accomplish the basic sensory work of representing the world. We just need sense modality-encoded information about the world, not information about how we encode such information.

This entails that we are *representationally blind* to the various multi-dimensional state spaces that constitute our vehicles of sensory representation, and *a fortiori* we are representationally blind to the dimensional aspect of the content they code for. We are not in a position to directly, non-conceptually represent the fact that a particular bit of sense-derived content is instantiated by such and such a set of dimensional values in a state space. Now, proceeding on the informational hypothesis, if that content were, for instance, the blue of my mouse pad as it phenomenally appears to me, this would *explain* the fact that I can’t directly see or grasp why — in virtue of what further set of facts — my blue looks a particular way to me, even though, on the informational hypothesis, it *is* a particular way by virtue of it’s being a certain set of dimensional values. I’m simply not in a position, *vis à vis* the experienced content *blue*, to know anything *about* the way it looks in and of itself because I don’t observe the process of sensory representation, rather I *consist* of it (along with many other representational and non-representational states and processes) as an organism that observes the world. Since I don’t have an observational perspective on the informational goings-on of experience, my blue will *necessarily* seem arbitrary with respect to its informational function and it will *necessarily* seem intrinsic, that is, non-relational and *sui generis*. It will seem as if some other apparently intrinsic phenomenal hue could have served as the particular color of my mouse pad, in which case it will seem eminently conceivable that someone with the same representational set-up could be experiencing at least a slightly different phenomenal hue from mine under the same perceptual conditions.⁵

I may recognize and report blue as a feature of my experiences but such recognition shouldn’t be confused with cognition or observation of blue as an *object* of my experience. I’m simply not in a position to experience my experience, that is, to take experience itself, in its basic qualitative particulars, as an observed, represented object about which I possess facts (Dretske, 1995, pp. 100–3). Apart from being able to specify its relations to its hue cousins, I don’t perceptually cognize the particular look of blue as a specifiable private fact since I don’t have the representational capacity to do so.⁶ Its ineffability, unspecifiability, and seeming

[5] Because the structure of subjective colour space is asymmetric, it’s plausible that simple or wholesale inversions of colour qualia are ruled out, and indeed that carefully imagined intersubjective qualia differences are very tightly constrained (Van Gulick 1993, pp. 144–5). Chalmers has argued that significant differences in qualitative experience are nomologically impossible for functionally identical systems. However, he allows that this ‘does not refute the possibility of a very mild spectrum inversion’ (Chalmers 1995b, p. 325).

[6] Another way to see this is that we can’t change our perspective on qualia, only on represented objects in experience. As I move around the blue chair, my represented perspective on it changes, and so my experienced qualia change. But I can’t achieve a perspectival, observational relationship with (that is, ‘move around’) any of the qualia involved in the experience (Clark, 1995).

intrinsicity as being *that* color are entailments of my representational limitations.

Since I can't assume a perceptual observer relation to the basic informational contents of sensory perception or the representational vehicles that encode them, sensory qualia don't include privately given or directly observed facts *about my experience*, rather they are represented facts *about what is represented*, what I as an organism directly observe, namely the world. Instead of cognizing (representing) blue, I cognize (represent) the world about which that particular blue is properly a specifiable fact: it's the specific colour of my mouse pad *as represented by the sensory system that I am*, partially, as a cognitive creature. Although they are instantiated by me as a representationally adept organism, and so are informational properties of representational states, qualia aren't *about* a special subjective world that only I have access to, rather they are informational content that represents the world as being certain ways, content that (as I will argue below in section 9) participates in those higher-level, integrative, complex behaviour-guiding functions that it turns out are specific only to consciousness. The essential characteristic of a quality, that it feels or looks like *this* to me, is thus a fact about the world-as-represented-by-me, not something I observe about my experience. There is, therefore, nothing factually available to me about experience over and above its intentional content, i.e., the relational, extrinsic, informational characteristics of qualia outlined in section 3 above (plus other non-sensory content) that I can successfully report, as can others. Consequently, the essential characteristic of an experience doesn't involve a further, private fact, something distinct and separate from such content. It *seems* to involve such a fact only because we're not in a position to directly, non-conceptually represent the fact that the values in various sensory state space dimensions constitute the look and feel of experience.⁷ The absence of categorically private first-person qualitative facts entails that qualities given in sensory experience all represent third-person, *objective* facts about the world as represented by a particular cognitive system.

There isn't, it turns out, anything metaphysically essential about the essential characteristic of a quality; instead, the 'mark of the qualitative' for basic sensory qualities is just that they can't be further described or specified *because* they are first-order, non-conceptual discriminative representations whose dimensional structure is not itself represented by the system, and therefore cannot be a fact *for* the system. The concrete, unanalyzable qualitiveness of sensory particulars is simply the fact that, as Thomas Metzinger puts it, the state space dimensions of qualia are 'impenetrable to cognitive operations'.⁸ All that's directly available to us is the *content* delivered by sensory representations, not the fact that such

[7] As Jackson (2001) puts it, 'I know only too well the residual feeling that *redness* could not be got out of the physical picture alone, but that is nothing more than a hangover from the conflation of instantiated property with intensional property [the property of representing *that* something is red]. That "redness" is not a feature one is acquainted with, but instead is a matter of how things are represented to be' (original emphasis).

[8] Metzinger asks rhetorically, implying an answer in the affirmative: 'Does the extremely high-dimensional form of a bodily state, which I do not recognize *as* such, lie concealed behind Yves Klein's

content is being represented or how it's represented. Another way to express this is that sensory representations are 'transparent', that is, cognitively invisible to us (Metzinger, 2003, p. 387), even though they are, of course, neurally instantiated in our heads.

7: The Contingency of Private Facts

The absence of categorically private first-person qualitative facts entails that qualities given in sensory experience all represent third-person, *objective* facts about the world as represented via a particular *type* of consciousness (human consciousness in the present case). To the extent to which others who share the same type of consciousness assume more or less the same perceptual, observer relation to the world, or the extent to which they can imagine being in such a relation, such qualities are shared, non-conceptual knowledge about the world. Phenomenal qualities pick out (represent) the same properties and regularities in the world in service to more or less the same set of cognitive needs.

My first-person point of view of the world as a cognitive agent — what we might call my *agent* perspective — consists of the unique set of third-person, objective facts (conceptual and non-conceptual) that I represent about the world as a particular instance of human consciousness. But my view of the world is also conditioned by the rather special, but nevertheless third-person fact that only *I* am constituted by this body, so I represent facts about this body in a way that no other consciousness can. On the view I'm recommending, the most private, subjective experience, e.g., the current feel of my toothache, represents a third-person, objective fact that only *I* *happen* to know by virtue of the fact that only *my* brain is recursively hooked up to itself and my body (Flanagan, 1992, p. 94). My toothache isn't private by virtue of incorporating categorically private phenomenal facts *about my experience*, although experience indeed consists of phenomenally represented facts about the world, in this case my tooth. Were a yet-to-be-devised (but perhaps not so far off) brain/body scanner trained to recognize the neuro-physical correlates of my toothaches, we'd see that the privacy of such facts is contingent, not categorical or metaphysical.

This is to say that only *I*, as *this* particular representational set-up, am now representing the objective state of my tooth (assuming that the tooth is in fact the problem) in this particular non-conceptual, affectively-laden, behaviour-guiding *manner*, which is simply to say it's my experience, not anyone else's. This constitutes a philosophically unproblematic sort of first-person perspective, that of my undergoing my particular token experiences. It is this individually tokened, non-conceptually based *mode of representing* particular objective facts about my body that others can't share by virtue of not being me — only *I* instantiate this mode *vis à vis* third-person facts about my tooth, a mode which, on the informational hypothesis, *is* what it's like to have a toothache. As Tye says, to fully understand or know what it is to experience a toothache, one has to have

"dimensionless depth" of the subjective sensory quality International Klein Blue?" (1995b, p. 449, original emphasis).

exercised similar non-conceptual representational capacities such that one can imagine what it's like to be in that representational state (Tye, 1995, pp. 54-6, 169-70), which is to say that one must in this respect instantiate more or less the same type of consciousness (human, or more broadly, toothed mammalian). While only I instantiate this particular non-conceptual representation, you, as someone who likely has undergone similar experiences, know what it's like to instantiate representations of analogous third-person facts about yourself. But neither of us is privy to subject-specific, first-person, unshareable facts about our *experiences* of a toothache.⁹ The privacy of sensory experience that follows from this unproblematic sense of a first-person perspective is not, therefore, a matter of access to special, categorically private data about experiences as objects, but just to be the only person who now instantiates this *instance* of representing a particular fact about the world in this particular non-conceptual mode.

Regarding Jackson's knowledge argument, my thesis suggests that in experiencing red for the first time, Mary doesn't come into possession of a new private fact about experience. Rather, she experiences (and learns to recognize and report) what it's like to instantiate a non-conceptual representation of a certain complex property of normal human environments; it's a new (for her) representational take on the *world*, on a set of third-person facts that she, a well-trained neuroscientist, already knew via high-level conceptual representations. Crucially, there *is* no private phenomenal fact *about red* to be entailed by third-person scientific facts, which is why Mary couldn't know in advance about that (non-existent) fact.

8. The Observing Self

Despite such considerations, it might still seem intuitively the case that we *are* in a privileged, observer-like relation to first-person facts about our experience, and indeed I think that our cognitive situation can provoke and maintain the illusion that experience is literally a presentation, not a representation. To help explain this illusion we must consider not only the purportedly private facts about experience but the observing subject that supposedly has access to them, since they stand or fall together.

As an organism (directly) observing the world by means of representing it, crucial objective facts for me to represent about the world are that I am in a particular position in it and that I'm interacting with it. As Antonio Damasio (1999, 2000), Thomas Metzinger (2000a), and others have pointed out, to survive I have to have a continuously updated egocentric *reality model* of my overall situation, a situation that includes me-the-organism contained in a larger environment that, from the organism's survival-driven perspective, is not-me. The most basic, essential distinction to be drawn within this model is between self and non-self. Thus, a more or less veridical and functional *self-model* (Van Gulick, 1993,

[9] As Dretske puts it, 'What we are conscious of when we feel pain (hunger, thirst, etc.) are not the internal representations of bodily states (the pains), but the bodily states that these representations (pain) represent. Though we can be — and most often are — aware that we are in pain, pains, like visual experiences, are awarenesses of objects, not objects of which we are aware' (1995, p. 103).

p. 150; Metzinger, 2000a, p. 289) must participate in a larger global representation that represents the self as being in, but other than, a world that is not-self. At all times the organism must have a sufficiently robust representation of its bodily boundaries, its location and trajectory, and its plans and purposes for its self-interested agenda to bear fruit. This constitutes a functionally essential, self-in-the-world-modeling *representational architecture*.

As a further elaboration of the reality model, the system must also represent itself as interacting with — being affected by and affecting — the world; it must model what Metzinger calls the ‘intentionality relation’. As he puts it, it’s functionally essential that the system have an ‘ongoing, dynamical representation of the system as *currently interacting with an object-component*’ (2000a, p. 296, original emphasis). On the informational hypothesis, this represented object-component consists of the intentional content about objects in the not-self world as delivered by perceptual processes fed by sensory representations (and modulated by top-down processes), while the self-model is built from internal representations of homeostatic functions that preserve bodily integrity (Damasio, 1999, 2000). Although I won’t defend it here, Metzinger’s explanation of how the *phenomenal* self arises from a representational architecture involves somewhat the same (but far more elaborate) considerations I adduced with respect to qualia in sections 4–6 above: by virtue of the fact that it can’t represent the self-model *as* a model (which is to say the self-modelling process becomes *transparent* or invisible to the system), the system inevitably falls into naïve realism with respect to the model, which therefore becomes an experienced *reality* of centred subjectivity (2000a, pp. 298–301). Such subjectivity involves the untranscendable feeling that we are seeing and perceiving the world around us as it’s *presented to us* by our senses. This feeling of self-as-observer to whom the world is presented thus models, in experience, the objective situation of being an organism in a perceptual relationship to its environment. The way in which our experience of the world changes as we experience ourselves acting with respect to it is the more or less reliably co-varying and behavior-controlling phenomenal analog of the organism’s interaction with the actual world. In describing this situation, Metzinger and others, such as Revonsuo (2000, pp. 65–6), have suggested that experience constitutes, in effect, a *virtual world* with the virtual subject at its center.¹⁰

Being a subject to whom the world is presented — the subjectivity of experience — is thus on Metzinger’s account a construction *within* experience. This phenomenal first-person perspective is analyzed and explained in representational terms that are themselves neither perspectival nor phenomenal (again, although I find Metzinger’s account persuasive, I won’t attempt to defend it

[10] Of course this world need not, as when we’re dreaming, correspond in real time to anything in the actual world, although when we’re awake it does so sufficiently to keep us out of trouble. The vertiginous strangeness of thinking of consciousness as a virtual world is evoked when Metzinger observes: ‘...a fruitful way of looking at the human brain, therefore, is as a system which, even in ordinary waking states, constantly hallucinates at the world, as a system that constantly lets its internal autonomous simulational dynamics collide with the ongoing flow of sensory input, vigorously dreaming at the world and thereby generating the content of phenomenal experience’ (2003, p. 52).

here). But despite this type of theoretical reduction, the strong, perhaps untranscendable *feeling* of being an observer or witness to the world can, I think, mislead some into supposing that experience itself is witnessed. That is, some, such as Velmans, Lehar, and Damasio (see section 2), will make the mistake of modeling their analysis of experience on the phenomenal sense of observerhood conferred by our representational architecture.

To see this, first note that the felt sense of being a self can vary from person to person, and over time for the same individual. Some persons (me, for instance) routinely experience the phenomenal subject as a more or less unitary, point-like self, nearly or totally devoid of any intrinsic character of its own. This self looks out at the not-self world, with its felt location more or less between and behind the eyes, most likely a function of the fact we are such visually oriented creatures. On this version of the phenomenal subject — call it the *bare* phenomenal self — the body is most certainly mine, but perhaps not essentially *me*. By contrast, others (who don't live 'in their heads' as much as I do) may have a more distributed sense of the essential self, constituted by some combination of bodily sensations, occurrent feelings, and thoughts, such that the phenomenal subject isn't bare, but consists of a set of characteristics. And of course it's possible and likely that the felt sense of self can vary within individuals over time, depending on their circumstances, for instance when meditating or under the influence of neurological disorders affecting the self-model (Metzinger, 2000a, pp. 295–6).

What I want to suggest is that the experience of being a self presented with the world, particularly when the self is experienced as the bare phenomenal subject, can create the impression that not just the world is presented to us, but that *experience* is presented to us. The bare phenomenal self is experienced as being essentially other than what is presented to it, and for such a subject even bodily sensations, occurrent thoughts, and emotions (internal content), along with the external world, are experienced as presentations *to* the essential me. This phenomenal situation involves, then, what non-conceptually feels to be a quintessentially witnessing subject, something that stands apart from and observes all of what is presented in experience. I'm suggesting that it's this extreme separation of the self from other contents of experience (for the phenomenal self is *itself* such content) that helps generate the problematic intuition of having an observational perspective on experience. For example, Baars (1996) says, '*You* are the perceiver, the actor and narrator of your experience' (original emphasis), and Deikman (1996) identifies the 'I' — what he describes as bare, contentless awareness — as the observer of the contents of consciousness. Now, when this notion of observerhood is combined with the conceptual (and veridical) understanding that the world is represented to us partially via experience (and partially via unconscious processes), the temptation might be to conclude that the representations themselves, the 'images' of experience as Damasio and Lehar call them, are observed from some sort of vantage point or perspective. Thus, instead of properly construing the phenomenal subject as constructed *within* experience, the subject might be construed as the witness *to* experience, so that experience is thought of, finally, as a literal presentation to an observing

self, not what it truly is, a virtual world-creating representation that *includes* the self.

In the grip of such an observational picture, the conviction might arise that the intentional contents of sensory experience are witnessed as private, first-person facts about *experience*, e.g., as the look *of my experience* of blue, the feel *of my experience* of pain. However, as we saw in section 6, we are not in a position to directly represent basic representational processes and thus know facts about particular, considered-in-themselves sensory representations that on the informational hypothesis constitute qualia. While the *phenomenal* situation of our virtual worlds involves the functionally useful sense that the experienced self is being presented with the world and the body via sensory experience, no one and no thing is literally presented with experience; as Dennett says, experience is ultimately and literally *unwitnessed and unappreciated*.

This, then, is my psychological-cognitive diagnosis of the philosophical mistake underlying the modern descendants of sense data theories such as Lehar's, which duplicate, in the purportedly direct perceptual relation of the person to experience, the actual direct perceptual relation between the person and environment. Naïve realists, since they don't suppose representations are involved in witnessing the world (the world is simply presented) can't fall prey to this mistake. It's only those sophisticated enough to realize that the world is known via experience who might be seduced by the strong (and functionally useful) sense of *phenomenal* observerhood into supposing that experience itself is a presentation involving a first-person observational perspective on private facts about experience. But, contra Velmans, no observational or representational arrangements exist which could provide this sort of perspective or give access to such facts.

9: Neuroscientific Accounts of Qualia

Although my approach so far has been largely negative, seeking to undermine intuitions that qualia involve categorically private phenomenal facts, a positive account of sensory consciousness as informational states is emerging from neuroscience and neurophilosophy (see for instance Dehaene, 2002; Metzinger, 2000; 2003).¹¹ Defined methodologically, consciously available information is just that embodied in representations that participate in functions subserving the empirically discovered *capacities* conferred by conscious states as opposed to unconscious states (Baars, 1999). For instance, conscious states have the capacity to make information available over extended time periods in the absence of continued stimulation; they permit novel, non-automatized behaviour; and they allow spontaneous generation of intentional, goal-directed behaviour with respect to perceived objects (Dehaene & Naccache, 2001). Studies of neural activity which contrast conscious and unconscious capacities indicate that

[11] This section seeks to implement Nicholas Humphrey's suggestion (2000), that in explaining consciousness we must work from both the philosophical, conceptual side *and* from the neuroscientific, empirical side, bringing them into mutual accommodation.

phenomenal experience is associated with widely distributed but highly integrated neural processes involving communication between multiple functional sub-systems in the brain, each of which plays a more or less specialized role in representing features of the world and body (Kanwisher, 2001; Dehaene & Naccache, 2001; Jack & Shallice, 2001; Parvizi & Damasio, 2001, Crick & Koch, 2003). Such processes, it is hypothesized, constitute a distributed, ever-changing, but functionally integrated ‘global workspace’ (Baars, 1988; Dehaene & Naccache, 2001).

The recruitment of sub-systems into the global workspace suggests that the neural correlates of phenomenal states integrate represented features into exactly those sorts of bound, coherent, object-level representations which dominate in subjective awareness and that seem necessary for most high-level, flexible behaviour and cognition. The informational aspects of qualia covered in section 3 may well correspond to various lower-level representational processes that intercommunicate during conscious episodes and thus are representationally linked to constitute conscious contents. But it’s critical to see that we aren’t normally aware of any of these aspects independently or alone, but always in a represented world of self, objects, and processes (relations and interactions of objects over time), all in more or less stable arrangements (Metzinger, 1995b, p. 448; Metzinger & Walde, 2000).¹² It takes a myriad of qualia to constitute a phenomenal, virtual world, and if we could ‘experience’ just one quale in isolation, it wouldn’t be experience at all; it would simply be the unconscious existence (part of no represented object and present to no represented subject) of one representational component of what *together* constitutes the phenomenal. Consciousness constitutes (and is constituted by) a represented holistic *context* within which discriminable elements are embedded (Kanwisher, 2001, pp. 107–8). The neural sub-processes responsible for each component of conscious representations, when acting in an unbound fashion, are perforce unconscious, but when they participate in the global workspace that unifies them into sensory percepts within such a context, they contribute their informational content to consciousness: the amalgamated, contextualized, object- and self-creating content that controls complex, flexible behavior. Although sections 3 through 6 concentrated on explaining qualia — the basic sensory elements of experience — the twist in the plot (flagrantly anticipated in section 8) is that phenomenology nearly always instantiates a virtual *world*, in which qualia participate as bound elements in globally representing the self in its environment of objects and events. To be phenomenal is (at least) for representations to be cognitively impenetrable and thus phenomenally transparent on the level of sensory elements (see section 6), *and* for these components to participate in an integrated, behaviour-controlling, and (on Metzinger’s account) transparent self-in-the-world model.

[12] Metzinger & Walde: ‘Not only is it impossible to experience hue without saturation or brightness, but it is also impossible to experience hue plus saturation plus brightness without an integrated percept — typically segregated from a background. Conscious experience seems to start on the object level, and elementary states in the true sense of the word do not exist’ (2000, p. 6).

What Dennett (2001a) suggests we should conclude from the study of the neural correlates of consciousness, and I concur, is that experience just *is* that behaviour-controlling information represented in the brain which is sufficiently globally available, i.e., not segregated in a lower level schema or in one sensory modality (see also Jack & Shallice, 2001, pp. 185–7). Shared and utilized by multiple functionally integrated sub-systems, this information largely ends up driving behaviour (largely but not exclusively, since unconscious representations may have their effects). On Dennett’s ‘fame in the brain’ gloss on the global workspace model, conscious representations are just those that dominate in the workspace by winning out over competing sets of less processed sensory information and incompletely bound higher-level cognitive content, and they win the competition via a staggeringly complex function of how the system’s drive states, represented goals, and self-model interact with perceptual input (Dennett, 1991 and 2001a; Desimone & Duncan, 1995; Cooper & Shallice, 2000; Franklin, 2000). We know such information has won out since, after all, it’s what we largely act on and can remember, report, and include when formulating intentions and plans. Again, it’s not that unconscious representations, e.g., masked stimuli and post-hypnotic suggestions, don’t have their influence on behaviour, but typically it’s conscious representations that rule as we go about our cognitively flexible business.

Once we discount the seeming observation of private phenomenal facts, we can begin to see that specifically phenomenal aspects of information – the highly integrated, self-and-objects-in-a-world aspect, and the resistance-to-further-representation aspect — involve nothing beyond the functioning of the global workspace in which representations of the self and environment dominate in controlling behaviour. Such integrated representations dominate because we need a global, integrated, reality model to behave effectively, and, for reasons of cognitive efficiency, we *don’t* need to immediately know — that is, non-conceptually represent — facts about the representational processes that contribute sensory content and build the integrated model.¹³ Crucially, consciousness isn’t something extra over and above the ever-fluctuating processes that instantiate such representations; it doesn’t ‘arise’ out of them or get ‘generated’ by them, rather it’s a property *of* those processes (Clark, 1995; Churchland, 1999; Dennett, 2001a). The highly integrated nature of conscious content is both phenomenally apparent (as phenomenally integrated selves, we experience coherent objects in a coherent space-time context) and neurally realized (in the high degree of sub-system communication manifested in the global workspace). Likewise, the cognitive impenetrability of qualia is phenomenally apparent (indeed, it’s the mark of qualitative particulars) and an empirical fact about the limits of representation.

[13] Indeed, to realize and believe in our gut that the reality model was ‘merely’ a model would undermine its effectiveness; naïve realism is arguably a necessary illusion for survival. As Metzinger puts it, ‘The evolutionary advantage of the underlying dynamical process of constantly confusing yourself with your own self-model is obvious: It makes a selfless biological system *egotistic* by generating a very robust self-illusion’ (2000a, p. 301, original emphasis).

Ongoing research will eventually describe in detail the characteristics of representational states that endow them with phenomenally conscious content: they participate in as yet obscure higher-level binding processes that instantiate a global reality model, a model which in turn supports abilities whose hallmarks are behavioural flexibility and the use of robust, multi-modal information about the world, in contrast to the more discrete, modular capacities subserved by unconscious processes.¹⁴ This picture of what constitutes the phenomenal contrasts markedly with Block's (2001), in which qualitative consciousness is something categorically other than informational representation, integration and dominance in behaviour control. According to Block, what this something is remains deeply mysterious. My Dennettian claim is that there may *seem* to be some mysterious, extra, non-informational, categorically private phenomenal component attached to qualitative experience, but there isn't. Once we've subtracted the seemings of private phenomenal facts, and accounted for cognitive impenetrability and informational integration, we're well on our way to capturing and explaining the phenomenal.¹⁵

10: So, What's Missing?

Despite the foregoing analysis, the lingering suspicion for many, perhaps most, readers is likely that the subject is nevertheless in possession of qualitative facts about experience, and that what's *essentially* qualitative or phenomenal is still being left out on this representationalist account. But what *is* essentially phenomenal or qualitative about consciousness? If it's the character of particular qualia, that's been accounted for as the informational content of multi-dimensional sensory representations that contribute information to higher level, bound, object representations. If it's the seemingly intrinsic, ineffable, not further specifiable *thusness* of qualia, the 'essential characteristic', that's been accounted for by our representational limitations: we can't directly grasp the dimensional aspect of sensory content. If it's subjectivity, that's been accounted for as the construction of the subject/world distinction *within* experience — a functionally essential, robust piece of non-conceptual representational content. If it's that consciousness gives us an entire, immediately present, coherent world, with the self at its centre, that's accounted for in terms of the representational architecture of the reality model which is neurally instantiated by the global workspace.¹⁶

[14] For general theoretical constraints on the notion of content see Van Gulick (1980), and for a well-articulated, comprehensive theory of phenomenal content, see Tye (1995). More recently Jackson (2001) has suggested that phenomenal representations have distinctive features of conveying contextual richness, immediacy, and causal origination, while playing a special functional role in determining beliefs.

[15] On our way, but with far to go, since there are further aspects of consciousness to be accounted for, e.g., other global properties of experience adduced by Metzinger ('presence', 'dynamicity', and 'convolved holism', 2000a, p. 286), and the sorts of non-sensory phenomenal content mentioned by Mangan (2001).

[16] Not fully accounted for here, of course, but I think by Metzinger (1995b; 2000a; 2003).

But, you persist, what about conscious experience *itself*? Well, what *about* experience do you have in mind? If you feel with regards to qualia that there's still something this account doesn't capture, but you can't quite specify it, that's to be expected. That, after all, is the *mark of the qualitative*: it's something specific, but you can't say what it is, and we've explained why it's both specific and yet unspecifiable. Some might object that what's been left out on this account is, for instance, the vivid, concrete, subjective interiority of experience (Chalmers, 1996; Nagel, 1974; 1986; 1998). This objection employs what are in fact adjectival descriptors that pick out variation *among* experiences and suggests, wrongly, that they are essential aspects of *all* experience which a representational account can't capture. In fact, the informational hypothesis *explains* why some (but, importantly, not all) sensory experiences are vivid, that is, *relatively* more intense; why the elements of sensory experience (but not experience as a composite phenomenal whole) are phenomenally concrete and monadic, i.e., *seemingly* intrinsic; and why experience (usually, but not necessarily always) seems an internal or inward subjective phenomenon, i.e., presented to a phenomenal subject felt to be here *inside* behind my eyes.

But how can 'mere' information *be* the marvellous, multi-faceted, completely engrossing world of qualitative subjective consciousness? The answer is that for each of us, this isn't *mere* information, it's personally *crucial* behaviour-guiding information without which we wouldn't last for long. It's information that's *represented* as inescapably significant and that, precisely because it's neurally instantiated, has ineluctable effects on behaviour. It is alternately lovely, terrible, exhilarating, saddening, and of course imbued with extraordinary detail and variety far beyond what we can express in mere concepts (not to disparage these very useful and recent cognitive innovations). Since we can't transcend this representation and its connections to behaviour — since we consciously and cognitively *consist* of it — there's no escaping the motivational hegemony of experience, for good or for ill.

Such considerations might help convince those like Kanwisher, who still wonder 'why perceptual awareness *feels* like anything at all' (2001, p. 90, original emphasis), that the 'hard problem' has been considerably attenuated as a philosophical conundrum. To feel is to non-conceptually, sensorily represent the world in the context of a sufficiently articulated and ramified self-world model. Both Nagel (1998) and Block (2001, pp.198, 212) suppose that some major empirical discovery or conceptual revolution perhaps in the far-flung future (Nagel: 'long after we are all dead') is needed to solve the problem of the fundamental nature of the phenomenal, but I think the revolution is well underway, driven by current neuroscientific research. If, as I and other functionalists and representationalists suppose, there is no extra-functional or non-representational private aspect of the phenomenal to account for, then as Dennett has long maintained, we're home free, but, I would add, without having quined — that is, eliminated — qualia (Dennett, 1990). Once we've quined the *seemings* of private first-person facts about experience presented to a non-experiential observer, there's no reason that qualia, conceived of as the reported basic particulars of

experience, can't persist in our theories as non-conceptual representations in good standing, with a relatively clear-cut neural, functional, and behavioral basis.

Given deeply entrenched folk-psychological intuitions about consciousness, e.g., that it simply can't be instantiated by physical processes, it is predictably counter-intuitive that sensory qualia might be nothing over and above informational contributions to integrated representations that dominate in the control of higher-level, flexible behaviour. If you suppose there's still something left out on this account, then my guess is you're thinking of consciousness in terms of presentations of private facts about experience, not the world, given to a non-experiential subject that has observational access to experience. But if, as I hope to have made plausible, no such animals exist, and we simply *consist of* experience as represented subjects-in-a-world, then there is no theoretical or empirical obstacle to reconfiguring our concept of the phenomenal to exclude the notion of private facts, and thus to fully naturalize qualia and consciousness.

Acknowledgements

My thanks to the anonymous reviewers for their insightful comments on an earlier draft of this paper.

References

- Baars, B. (1988), *A Cognitive Theory of Consciousness* (Cambridge, MA: Cambridge University Press).
- Baars, B. (1996), 'Understanding subjectivity: global workspace theory and the resurrection of the observing self', *Journal of Consciousness Studies*, **3** (3), pp. 211–16.
- Baars, B. (1999), 'Contrastive phenomenology: a thoroughly empirical approach to consciousness,' in Block *et al.* (1999).
- Biro, J. (1993), 'Consciousness and objectivity', in Davies and Humphreys (1993), pp. 178–196.
- Block, N. (2001), 'Paradox and cross-purposes in recent work on consciousness', *Cognition*, **79**, pp. 197–219.
- Block, N., Flanagan, O., Guzeldere, G. (ed. 1999), *The Nature of Consciousness: Philosophical Debates* (Cambridge, MA: MIT Press).
- Chalmers, D.J. (1995a), 'Facing up to the problem of consciousness', *Journal of Consciousness Studies*, **2** (3), pp. 200–219.
- Chalmers, D.J. (1995b), 'Absent qualia, fading qualia, dancing qualia', in Metzinger (1995).
- Chalmers, D.J. (1996), *The Conscious Mind* (New York: Oxford University Press).
- Chalmers, D.J. (1999), 'First-person methods in the science of consciousness', *Consciousness Bulletin*, Fall, pp. 8–11.
- Churchland, P.M. (1989), *A Neurocomputational Perspective* (Cambridge, MA: MIT Press).
- Churchland, P.S. (1988), *Neurophilosophy* (Cambridge, MA: MIT Press).
- Churchland, P.S. (1999), 'Can neurobiology teach us anything about consciousness?', in Block *et al.* (1999).
- Clark, T.W. (1995), 'Function and phenomenology: closing the explanatory gap', *Journal of Consciousness Studies*, **2** (3), pp. 241–54, online at <http://www.naturalism.org/consciou.htm>.
- Cooper, R. and Shallice, T. (2000), 'Contention scheduling and the control of routine activities', *Cognitive Neuropsychology*, **17**, pp. 297–338.
- Crick, F. and Koch, K. (2003), 'A framework for consciousness', *Nature Neuroscience*, **6** (2), pp. 119–26.
- Damasio, A.R. (1999), *The Feeling of What Happens: Body and Emotion in the Making of Consciousness* (New York: Harcourt Brace).

- Damasio, A.R., Grabowski, T., Bechara, A., Damasio, H., Ponto, L.L.B., Parvizi, J., Hichwa, R.D. (2000) 'Distinctive patterns of subcortical and cortical brain activation associated with self-generated emotions and feelings', *Nature Neuroscience*, **3** (10), pp. 1049–56.
- Damasio, A.R. (2003), *Looking for Spinoza* (New York: Harcourt).
- Davies, M. and Humphreys, G.W. eds. (1993), *Consciousness* (Oxford: Blackwell).
- Deikman, A. (1996), "'I" = awareness', *Journal of Consciousness Studies*, **3** (4), pp. 350–6.
- Dehaene, S. and Naccache L. (2001), 'Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework', *Cognition*, **79**, pp. 1–37.
- Dehaene, S. (ed. 2002), *The Cognitive Neuroscience of Consciousness* (Cambridge, MA: MIT Press).
- Dennett, D. (1990), 'Quining Qualia', in Marcel and E. Bisiach, eds (1988), reprinted in Lycan (1990).
- Dennett, D. (1991), *Consciousness Explained* (Boston, MA: Little, Brown and Company).
- Dennett, D. (2001a), 'Are we explaining consciousness yet?', *Cognition* **79**, pp. 221–37.
- Dennett, D. (2001b), 'The fantasy of first-person science', debate with David Chalmers, <http://ase.tufts.edu/cogstud/papers/chalmersdeb3dft.htm>.
- Desimone, R. and Duncan, J. (1995), 'Neural mechanisms of selective visual attention', *Annual Review of Neuroscience*, **18**, pp. 193–222.
- Dretske, F. (1995), *Naturalizing the Mind* (Cambridge, MA: MIT Press).
- Feser, E. (2001), 'Qualia: irreducibly subjective but not intrinsic', *Journal of Consciousness Studies*, **8** (8), pp. 3–20.
- Flanagan, O. (1992), *Consciousness Reconsidered* (Cambridge, MA: MIT Press).
- Franklin, S. (2000), 'Deliberation and voluntary action in "conscious" software agents', *Neural Network World*, **10**, pp. 505–21.
- Humphrey, N. (2000), 'How to solve the mind-body problem', *Journal of Consciousness Studies*, **7** (4), pp. 5–20.
- Jack, A.I., Shallice, T. (2001), 'Introspective physicalism as an approach to the science of consciousness', *Cognition*, **79**, pp. 161–96.
- Jackson, F. (1982), 'Epiphenomenal Qualia', *Philosophical Quarterly*, **32**, pp. 127–36, reprinted in Lycan (1990).
- Jackson, F. (2001), 'Mind and illusion', Royal Institute of Philosophy Lecture, 'Minds and Persons' series, 2000–2001, forthcoming in *Philosophy*, supp. vol. 53, pp. 253–73. Reprinted in *Minds & Persons*, ed. A. O'Hear (Cambridge: Cambridge University Press, 2003).
- Kanwisher, N. (2001), 'Neural events and perceptual awareness', *Cognition*, **79**, pp. 89–113.
- Lehar, S., (2004), 'Gestalt isomorphism and the primacy of subjective conscious experience: A gestalt bubble model', *Behavioral and Brain Sciences*, **26** (4), online at <http://cns-alumni.bu.edu/~slehar/webstuff/bubw3/bubw3.html>.
- Levine, J. (1983), 'Materialism and qualia: the explanatory gap', *Pacific Philosophical Quarterly*, **64**, pp. 354–61.
- Levine, J. (1993), 'On leaving out what it's like', in Davies and Humphreys (1993).
- Lycan, W.G. (1987), *Consciousness*, (Cambridge, MA: MIT Press).
- Lycan, W.G. (ed. 1990), *Mind and Cognition* (Cambridge, MA: Blackwell).
- Mangan, B. (2001), 'Sensation's ghost: the non-sensory "fringe" of consciousness', *Psyche*, **7**(18), online at <http://psyche.cs.monash.edu.au/v7/psyche-7-18-mangan.html>.
- Marcel, A. J., Bisiach, E., eds. (1988), *Consciousness and Contemporary Science* (Oxford: Oxford University Press).
- Metzinger, T. (ed. 1995), *Conscious Experience* (Thorverton, UK: Imprint Academic; Paderborn: Mentis).
- Metzinger, T. (1995a), 'The problem of consciousness', in Metzinger (1995).
- Metzinger, T. (1995b), 'Faster than thought', in Metzinger (1995).
- Metzinger, T. (ed. 2000), *Neural Correlates of Consciousness: Empirical and Conceptual Questions* (Cambridge, MA: MIT Press).
- Metzinger, T. (2000a), 'The subjectivity of subjective experience: a representationalist analysis of the first-person perspective', in Metzinger (2000), pp. 285–306.
- Metzinger, T., Walde, B. (2000), 'Commentary on Jakab's "Ineffability of Qualia"', *Consciousness and Cognition*, **7**, 353–62.
- Metzinger, T. (2003), *Being No One: The Self-model Theory of Subjectivity* (Cambridge, MA: MIT Press).

- Nagel, T. (1974), 'What is it like to be a bat?', *Philosophical Review*, **83**, pp. 435–51, reprinted in Nagel (1979).
- Nagel, T. (1979), *Mortal Questions* (Cambridge: Cambridge University Press).
- Nagel, T. (1986) *The View From Nowhere*, (New York: Oxford University Press).
- Nagel, T. (1998), 'Conceiving the impossible and the mind-body problem', *Philosophy*, **73** (285), pp. 337–52.
- Noe, A. (2002), 'Is the visual world a grand illusion?', *Journal of Consciousness Studies*, **9** (5–6), pp. 1–12.
- Parvizi, J. and Damasio, A. (2001), 'Consciousness and the brainstem', *Cognition*, **79**, pp. 135–59.
- Revonsuo, A. (2000), 'Prospects for a scientific research program on consciousness', in Metzinger (2000).
- Tye, M. (1995), *Ten Problems of Consciousness: A representational theory of the phenomenal mind*, (Cambridge, MA: MIT Press).
- Van Gulick, R. (1980), 'Functionalism, information, and content', *Nature and System*, **2**, reprinted in Lycan (1990), pp. 107–29.
- Van Gulick, R. (1993), 'Understanding the phenomenal mind: are we all just armadillos?', in Davies and Humphreys (1993), pp. 137–54.
- Velmans, M. (1996), 'Consciousness and the "causal paradox"', *Behavioral and Brain Sciences*, **19** (3), pp. 538–42.
- Velmans, M. (2002), 'How could conscious experiences affect brains?', *Journal of Consciousness Studies*, **9** (11), pp. 3–29.

Paper received March 2003, revised August 2004

CONFERENCE ANNOUNCEMENT

Phenomenology and Psychiatry for the 21st Century

5th & 6th September, 2005

Venue: Institute Of Psychiatry, London, UK

Interest in Phenomenology ebbed in the 1980s with the growth of the neurosciences, but in recent years there has been a swing of interest back towards phenomenology and how it relates to the new findings from neuroscience and epidemiology. This conference represents an exciting opportunity to reinvigorate and advance an important cross-disciplinary, clinical, research, and conceptual debate for 21st century psychiatry.

Speakers

Nancy Andreasen, Paul Bebbington, German Berrios, Peter Chadwick, John Cutting, Anthony David, Thomas Fuchs, Bill Fulford, Nassir Ghaemi, Arthur Kleinman, Elizabeth Kuipers, Paul Mullen, Robin Murray, Josef Parnas, Matthew Ratcliffe, Steven Rose, Louis Sass, Sean Spence, Giovanni Stanghellini, Jim Van Os and Kai Vogely.

For Further Information

Dr Gareth Owen and Dr Robert Harland
Email: phenomenology@iop.kcl.ac.uk

Or visit our website: <http://www.iop.kcl.ac.uk/iopweb/events>