

Introspective Humility

Tim Bayne and Maja Spener
University of Oxford and St. Catherine's College
Manor Road
Oxford OX1 3UJ
United Kingdom
tim.bayne@gmail.com

Published in *Philosophical Issues*, 20: 1-22 (2010). This version of the paper is only a **draft**. For purposes of quotation please consult the published version.

1. Introduction

Viewed from a certain perspective, nothing can seem more secure than introspection. Consider an ordinary conscious episode—say, your current visual experience of the colour of this page. You can judge, when reflecting on this experience, that you have a visual experience as of something white with black marks before you. Does it seem reasonable to doubt this introspective judgement? Surely not—such doubt would seem utterly fanciful. The trustworthiness of introspection is not only assumed by common-sense, it is also taken for granted by many of theorists about the mind. Within both philosophy and the science of consciousness it is widely held that introspection is generally reliable, at least with respect to the question of one's current (or immediately prior) conscious states. Without this assumption, we could not make sense of theorists' widespread use of introspection, both in support of their own position and to undermine that of their opponents.¹

Despite all this, the assumption that introspection is trustworthy is not unproblematic. One line of argument for introspective scepticism takes starts with the thought that introspection is unscientific because it delivers purely private data. This worry animated

¹ In this paper we use 'introspection' and its cognates solely to refer to a mode of first-personal access to one's current (or immediately prior) conscious states, episodes or processes, and we leave to one side uses of the term that apply to non-occurrent and non-conscious mental states (background motivations, dispositional states, etc).

behaviourism, and in contemporary discussion it lies behind much of Dennett's (1991) criticism of introspective methodology in the study of consciousness. Another line of argument for introspective scepticism appeals to the fact we are often uncertain about the exact characterization of our own conscious states; when pressed, we find it surprisingly difficult to describe what (say) our emotional experiences or our experiences of mental imagery are like (e.g. Schwitzgebel 2008). These challenges are serious but they are not our concern here. Instead, our focus is on the challenge from introspective disagreement: the fact that apparently sincere and competent subjects often disagree about the deliverances of introspection puts pressure on the assumption that introspection is trustworthy.

In responding to this challenge we are guided by two thoughts. On the one hand, the challenge cannot be dismissed with a Moorean wave of the hand. Introspective disagreement poses a genuine threat to the trustworthiness of introspection, and it must be given a proper response. On the other hand, radical scepticism about introspection is untenable: denying that we have any first-person knowledge of consciousness is not a serious possibility.

Our attempt to do justice to both of these constraints proceeds as follows. We start by providing three examples of introspective disputes (§2). We then sketch an argument from these disputes for 'radical introspective scepticism'—the claim that introspection is rarely, if ever, to be trusted (§3). Having motivated the threat of radical scepticism we spend the remainder of the paper defusing it. We diagnose each of the three introspective disputes (§4). With a better grip on how introspection might go awry, we then show that there is a range of introspective judgments whose trustworthiness is not threatened by sceptical worries generated from introspective disagreement (§5).

2. Introspective disagreement

Introspective disputes have troubled the study of consciousness since its inception. In fact, it was precisely an introspective dispute over the existence of non-imagistic thought that brought introspectionist psychologist into disrepute in the early years of the 20th century (Boring 1929/50; Lyons 1986; Woodworth 1906). The science of consciousness has recently been resurrected, and with its revival has come a renewed appreciation of the depth and range of disagreement that introspection engenders.

The introspective disputes that are of interest to us here share the following structure. Some theorists say that introspection supports the claim that P. Other theorists say either that introspection supports the claim that not-P, or that it supports claims that directly entail or at least strongly support not-P. Advocates of each position appear to have

engaged in careful and extensive introspection with respect to the question at issue, yet neither of the two camps has been able to win over its opponents. We turn now to the first of our three examples of such disputes.²

2.1 The nature of perceptual experience

One of the major fault-lines running through treatments of perceptual experience divides those who hold that phenomenal character is exclusively world-directed from those who hold that the phenomenal character involves aspects that are not world-directed. In contemporary discussion this division fuels the debate between pure representationalists, who hold that the phenomenal properties of a given perceptual experience can be adequately specified in terms of objects and properties represented by experience and nothing more, and impure representationalist, who hold that an adequate specification of the phenomenal properties of a given experience involves appeal to the experience itself (in addition to objects and properties represented by experience).

In motivating their views, both pure and impure representationalists appeal to how things seem to them when reflecting on their own experience. Consider, for example, the way in which Levine motivates a non-representational conception of colour phenomenology:

[T]his is the intuitive starting point. It seems plausible to think of visual experience as having an intrinsic qualitative character. The reddishness of my visual experience of the diskette case seems to be a property of my experience. ... We needn't expect *a priori* to know what sort of property it is, but that my experience has it seems apparent. (I am not appealing to privileged access, or any sort of hyper-certainty. I'm just saying let's not doubt the obvious until we have to.) (Levine 1995, 277-8)

By contrast, Tye holds that introspection supports a strongly representationalist conception of colour phenomenology:

Standing on the beach in Santa Barbara a couple of summers ago on a bright sunny day, I found myself transfixed by the intense blue of the Pacific Ocean. Was I not here delighting in the phenomenal aspects of my visual experience? And if I was, doesn't this show that there are visual qualia? I am not convinced. It seems to me that what I found so pleasing in the above instance, what I was focussing on, as it were, were a certain shade and intensity of the colour blue. I experienced blue as a property of the ocean not as a property of my experience. My experience itself certainly wasn't blue. Rather it was an experience

² For further discussion of introspective disputes see Kriegel 2007, Siewert 2007 and Schwitzgebel 2008.

that represented the ocean as blue. What I was really delighting in, then, were specific aspects of the content of my experience. It was the content, not anything else, that was immediately accessible to my consciousness and that had aspects I found so pleasing. (Tye 1992: 160)

Appealing in this way to the deliverances of introspection is standard practice amongst the participants in this dispute. Theorists agree on the following constraint on their theorizing: an account of the nature of experience must be phenomenally adequate, i.e. it must neither ignore nor distort the phenomenal character of experience. They also assume that we have access to the phenomenal character of experience via introspection. Specifically, they assume that we can formulate particular phenomenal adequacy conditions based on introspective evidence. These introspection-based phenomenal adequacy conditions are then used to support a favoured view and to reject rival views. But the introspective reports that theorists produce are often incompatible, as in the example provided above.³

2.3 The phenomenology of thought

Our second case of introspective disagreement concerns the nature of conscious thought. Although it is (almost) universally granted that thoughts can be conscious, there is little agreement concerning the phenomenal character of conscious thought. Roughly speaking, we can distinguish two approaches: a liberal (or 'inclusivist') position and a conservative (or 'exclusivist') position. Liberals holds that the reach of phenomenal consciousness is not restricted to sensation, perception and affect, but extends to thought as well. According to liberals, some thoughts have a *sui generis* phenomenal character. For example, they typically hold that there is a distinctive phenomenology involved in (say) judging that the tree in one's backyard ought to be cut down, hoping that the bus will soon come, or deciding to referee a paper. Conservatives, by contrast, deny that thoughts have a distinctive phenomenal character. Conservatives generally admit that thinking brings with it phenomenal states of various kinds, but they hold that the phenomenology of

³ For other (introspection-based) defences of pure representationalism see Harman (1990) and Dretske (1995); and of impure representationalism see Peacocke (1980) and Block (2003). This debate is part of a larger debate about the nature of perceptual experience that involves further theories, notably sense datum theory and disjunctivism, the proponents of which also take themselves to be bound by introspectively-based constraints of phenomenal adequacy. For discussion see e.g. Martin (2002) and Spener (ms).

thought is merely sensory or imagistic—that it can be exhaustively captured by appeal to the phenomenal character of inner speech or motor imagery.⁴

Although liberals do provide non-introspective arguments for the existence of a *sui generis* cognitive phenomenology (see e.g. Pitt 2004), appeals to introspection—sometimes more direct, sometimes less direct—constitute the central line of argument for their view. For example, Horgan and Tienson write:

Intentional states have a phenomenal character, and this phenomenal character is precisely the what-it’s-like of experiencing a specific propositional-attitude type vis-à-vis a specific intentional content.... Attentive introspection reveals that both the phenomenology of intentional content and the phenomenology of attitude type are phenomenal aspects of experience, aspects that you cannot miss if you simply pay attention. (Horgan & Tienson 2002: 522-23)

Conservatives are not persuaded. Here, for example, is Wilson:

In the spirit of Horgan and Tienson’s appeal for a reader to “pay attention to your own experience”, I have just done the decisive experiment: I thought first that George Bush is President of the United States, and had CNN-mediated auditory and visual phenomenology that focused on one of his speeches. I then took a short break, doodled a little, wandered around the room, and then had a thought with that very same content and...nothing. (Wilson 2003: 417)

So here we have another example of an introspective dispute. Whereas the debate about perceptual experience concerns the structure of phenomenal consciousness, this dispute concerns its scope.

2.4 The richness of consciousness

Our third example concerns the dispute about how rich the contents of consciousness are. According to some theorists, the stream of consciousness consists of very little beyond one's immediate focus of attention at a given time. On this view, conscious awareness is sparse: at a given time it is typically confined to a small range of objects and details, in

⁴ For defence of conservatism see e.g. Tye (1995); Nelkin (1989); Robinson (2005); Lormand (1996); for defence of liberalism see e.g. Horgan & Tienson (2002); Pitt (2004); Siewert (1998); Strawson (1994). For an overview of the debate see Bayne & Montague (2011).

While there are points of contact between this debate and the debate about the existence of non-imagistic thought in the early decades of the 19th Century, the two should not be conflated.

perhaps only one modality. Suppose that while reading a novel you suddenly notice the scratchy collar of your shirt. The sparse view holds that although the collar may have been scratching you for hours, you were not consciously aware of the scratchiness until your attention was drawn to it. Other theorists, by contrast, hold that the stream of consciousness is typically rich with content. On this view, we often enjoy experiences in different modalities simultaneously. A single conscious state might include (say) auditory, visual, affective and pain content. According to this view, conscious awareness is abundant in that it significantly exceeds the (narrow) focus of attention at a given time.

Proponents of each of these two views invoke introspection in support of their position. Here, for example, is Searle's defence of the abundant position:

Within the field of consciousness, we need to distinguish between those things that are at the centre of our attention and those that are at the periphery. We are conscious of a very large number of things that we are not attending to or focusing our attention upon. For example, up to this moment I have been focusing my attention on the philosophical problem of describing consciousness, and I have not been paying any attention to the feeling of the chair against my back, the tightness of my shoes, or the slight headache I have from drinking too much wine last night. Nonetheless, all of these phenomena are part of my conscious awareness. (Searle 1992: 137f.)

Although some advocates of the sparse conception of consciousness grant that introspection seems to support the rich conception of consciousness (Blackmore et al. (1995) and O'Regan (1992)), others take introspection to support the sparse view. For example, Zimmerman claims that 'personal experience' supports the view that consciousness can be directed to the products of only one sense organ at a time (1989: 172).

A variant of the debate between sparse and rich theorists occurs even with respect to the analysis of visual experience. Rich theorists hold that visual experience itself presents us with a detailed representation of the world, whereas austere theorists hold that at any one point in time visual experience is limited in the number of objects and properties that it represents. Here too theorists draw on introspective evidence in support of their views. Consider for example the debate about visual content in the context of briefly-presented visual displays. In the Sperling paradigm subjects are shown a grid of twelve alphanumeric figures for a brief period (say, 250 ms), and are required to report the identity of each figure (Sperling 1960). Typically, subjects are able to reliably identify only four of the twelve figures. Advocates of a rich conception of visual experience sometimes

appeal to the fact that subjects of these experiments often insist that they have seen more than they can report (Sperling 1960: 1; see also Block 2007: 531; Tye 2006: 513). Such reports have not gone unchallenged, with other theorists denying that there is any introspective evidence for thinking that the content of visual experience ‘overflows’ what is reportable.

So we have two manifestations of the debate between rich and sparse theorists, one that concerns the question of whether consciousness draws on multiple modalities at a time and another that concerns the structure of visual experience. Participants in each of these debates appeal to introspective evidence in support of their position.

3. The threat of radical scepticism

In each of the three domains just surveyed, theorists disagree about the deliverances of introspection despite taking care to introspect as carefully as they can. Subjects are typically confident in the accuracy of their own judgments, but their judgments are at odds with those of others’.

The argument from introspective disagreement involves two steps. The first step is to argue that the introspective judgements involved in these disputes are unreliable, and that one (or perhaps even both) parties to the dispute are getting things wrong. If this step can be secured, this motivates scepticism concerning the introspective judgements involved in the disagreement. The second step extends this sceptical stance, taking it to undermine the *general* trustworthiness of introspection. If we have reason to think that introspection is untrustworthy even when theorists are doing their very best (and take themselves to be getting things right), it is little comfort that sometimes people agree (or at least fail to disagree) about the deliverance of introspection. The damage is already done: deep introspective disputes show that introspection should not be trusted even when deliberately and carefully conducted. At the very least, the burden of proof would seem to be on advocates of introspection to think that we ever get it right. In short, introspective disagreement of the sort we set out appears to motivate *radical introspective scepticism*.

This scepticism is radical in two ways. Firstly, it is radical in *scope*. The idea is not merely that we ought to distrust introspection in contexts of introspective disagreement or subjective uncertainty, but even in contexts of introspective agreement and high subjective confidence. Secondly, it is radical in *significance*. The claim is not that introspection is open to merely ‘philosophical’ forms of scepticism that can be safely ignored in ordinary contexts of inquiry, but that we have reason to doubt whether introspection meets even ordinary, everyday standards of trustworthiness.

The form of the argument from disagreement to radical scepticism is a familiar one. In ethics it is often argued that persistent and apparently intractable disagreement on most foundational issues in moral theory shows that we have no moral knowledge (see e.g. Sinnott-Armstrong 2006). In the philosophy of religion it is commonplace for theorists to argue that the protracted disagreement about religious experience undermines the claim that such states might function as a source of religious knowledge (see e.g. Martin 1990). Similarly, deep introspective disagreements pose a radical threat to ordinary introspective knowledge.⁵

Radical introspective scepticism is rarely (if ever) endorsed, but some theorists certainly feel its pull. In a recent paper Eric Schwitzgebel writes, 'the introspection of current conscious experience, far from being secure, nearly infallible, is faulty, untrustworthy, and misleading—not just *possibly* mistaken, but massively and pervasively' (2008: 259; emphasis in original).⁶ Although Schwitzgebel draws on a number of considerations in making his case, the argument from introspective disputes clearly plays a central role in drawing him towards radical scepticism. He says, 'even the meekest and most tentative reflections about experience are bound soon to conflict with what others have said, so widespread and fundamental are the disagreements in consciousness studies' (Hurlburt & Schwitzgebel 2007: 250).

But despite his introspective pessimism Schwitzgebel stops short of whole-heartedly embracing radical scepticism:

People must have at least some inkling of what's going on in their own present and immediately past conscious experience. That inkling is, I think, surprisingly poor and unstable ... but it would be a radical skepticism indeed to suppose that we have no clue whatsoever about the ongoing flow of our experience. Asking people about their present or immediately past experience is not entirely pointless. (Hurlburt & Schwitzgebel 2007: 227)

⁵ There is also renewed interest in the epistemology of disagreement (for an overview see Christensen 2009). Spener (2010) discusses the epistemic significance of introspective disagreement about cognitive phenomenology.

⁶ Dennett (1991) appears to endorse radical introspective scepticism in places, although there are other passages in which he places a great deal of trust in introspection. See Dokic and Pacherie (2007), Hurlburt & Schwitzgebel (2007) and Schwitzgebel (2007a) for illuminating discussion of Dennett's position.

We agree. Focus once again on your current experience of the colour of this page. Is it reasonable to doubt the deliverance of introspection in this instance, i.e. that it visually appears to you as if there is a white surface with black marks before you? Surely not: this introspective judgment might not be infallible, but it is one in which a high degree of confidence would be entirely warranted. What holds here holds of countless other introspective judgements one might make: that it is as if one is in pain; hears a high-pitched sound; tastes something tangy; is tired, etc. These are just some of the many introspective judgments concerning which serious doubt seems perverse. Let us call instances of introspection that seem secure in this way 'trustworthy cases.'

Our notion of a trustworthy case is not particularly demanding. We are not assuming that trustworthy cases are infallible. Perhaps certain types of introspective judgments are infallible in a way that goes beyond mere self-reference; if so, that is not our concern here.⁷ Nor are we assuming that trustworthy introspective judgments are epistemically superior to their perceptual counterparts. Our claim is only that there is a class of ordinary introspective judgments that are trustworthy in that we are justified in being highly confident in them.

This then is the challenge that we face: although introspective disagreement provides strong motivation for radical introspective scepticism, that position seems untenable. But unless we have positive reasons for rejecting radical scepticism its dismissal amounts to nothing more than wishful thinking. What we require is an account of introspective disagreement that makes it *reasonable* for us to retain our pre-theoretical conviction that introspection is often trustworthy.

At this point it might be suggested that introspective disputes arise because one (or even both) parties to these disputes are failing to introspect properly, and that when conducted in the right way introspection would produce inter-subjective agreement. Ordinary introspection might not be trustworthy (even when conducted by theoretically naïve subjects), but *suitably educated* introspection can be trusted to deliver reliable judgments.

We are not persuaded. For one thing, it is very much an open question whether introspection can be significantly improved. As a matter of historical fact, introspective training did not lead to any kind of convergence on the question of imageless thought, for subjects tended to produce introspective reports consistent with the prevailing ideologies of the lab in which they were trained (Boring 1929/50; Woodworth 1906). But even if it

⁷ For accounts of (a narrow class of) infallible introspective judgments see e.g. Papineau (2000), Chalmers (2003), Horgan & Kriegel (2007).

were true that introspective training of some kind could lead to introspective convergence, we would still be faced with the challenge of radical introspective scepticism, for we do not typically think that any regimen of introspective training is *required* to make trustworthy introspective judgments. The challenge provided by introspective disagreement would remain even if it were the case—which we doubt—that introspective disagreements could be overcome by training.

4. Diagnosing the disputes

Our response to the challenge of introspective disputes is to ‘quarantine’ the sceptical threat that they motivate. As we show, the best explanation of the disputes may indeed show that the introspective judgements involved are untrustworthy, but this does not undermine introspection’s ability to deliver trustworthy judgements elsewhere. We begin by sketching four potential sources of introspective disagreement, and then turn to the question of how this framework might account for the three introspective disagreements set out in §2.

One potential source of introspective disagreement concerns *individual differences*.

Although introspection delivers a first-person present-tense judgment such as ‘It looks as if there is something red before me’ or ‘It sounds as if there is a loud clanging,’ introspective disagreement strictly speaking concerns generalizations about experience that are made on the basis of such judgements. The underlying assumption is of course that humans are relevantly similar with respect to the subject matter of the introspective judgment. But this assumption may be mistaken. Introspective disputes might arise not because any of those who disagree are failing to introspect correctly, but because they are over-generalizing from their own case.

A second potential source of introspective disagreement is *terminological variation*. Perhaps participants in certain debates are using different terms to express the same introspectively available facts. This would explain introspective disputes without calling into question the introspective judgments of the participants. The source of disagreement would be located in the expression of introspective judgments rather than in the contents of introspective states *per se*.

A third potential source of introspective disagreement arises from the *influence of background belief and expectation*. Introspective judgments might be driven by expectation in much the same way as perceptual judgments can be. On this view, introspective disagreement results from the fact that the deliverances of introspection have been significantly modulated by theory. Since introspective judgements themselves do not provide reliable clues as to their aetiology, theorists will usually be unaware of the degree

to which their claims are expectation-driven. Note that this explanatory strategy concerns the possibility of expectation influencing introspective judgement; it is a further question whether expectation might affect the subject's conscious states themselves.

A fourth explanatory strategy focuses on *operational constraints*. Perceptual capacities generate accurate perceptions only when they are functioning properly, the environmental conditions under which are deployed meet certain specifications, and they are directed at appropriate kinds of things. For example, accurate visual experience requires not only that the machinery of vision be intact but that it also be directed towards appropriate properties and objects under suitable conditions of illumination and proximity. Similar constraints apply to introspection: in order to generate trustworthy judgments, the introspective capacity and whatever machinery underlies it must not be impaired, and it must be deployed under appropriate conditions. According to this diagnosis, introspective disputes arise because the introspective judgements involved are made under conditions that fall outside of introspection's range of reliable operation.

Note two things about these four explanatory strategies. Firstly, there is a fundamental contrast between the first two strategies and the second two strategies: whereas the second two strategies attempt to explain introspective disagreement in ways that leave introspection itself untouched, the second two strategies directly impugn the reliability of introspection. Secondly, these potential sources of introspective disagreement are not mutually exclusive, and a full account of each of these disputes might have reason to draw on more than one strategy.

With these four strategies in hand let us return to the introspective disputes set out in §2.

4.1 Perceptual experience revisited

Why might some theorists claim introspection reveals that the phenomenal character of perceptual experience is exhausted by its representational content whereas other theorists deny this?

Appeal to individual differences does not yield a satisfactory account of this debate. Of course, there is nothing to rule out the possibility that there is far more variation in perceptual experience than we tend to assume, but the individual differences that would need to be posited in this case would be of a particularly striking kind. It is hard to believe that the perceptual phenomenology of some of us is 'pure' whilst that of the rest of us is 'impure.'

Similarly unconvincing is a diagnosis based on terminological variation. The introspective descriptions at the heart of this debate often conflict in a straightforward manner,

centring, as in our sample quotations for instance, on the presence or absence of intrinsic, non-representational properties of experience. It seems unlikely that the parties to this dispute are merely using the different terms to describe the same phenomenal aspect in this case.

How about operational constraints? At first glance, a diagnosis in terms of operational constraints seems unattractive. After all, debates about the nature of phenomenal experience appear to involve introspective judgments produced by introspection operating under optimal—or at least perfectly adequate—conditions. Theorists make genuine efforts to introspect carefully, repeatedly and under conditions conducive for reliable introspection, and their judgments are not made under conditions of high cognitive load, time pressure, or fatigue.

Nonetheless, we are inclined to think that the appeal to operational constraints can do some work here. The key idea is not that these reports are produced under conditions in which introspection is unreliable, but that the questions being asked of introspection are not ones that it is well-equipped to answer—they fall outside of introspection's domain of expertise. Although we don't have an account of that domain to offer here, but we do think that whatever such account is offered must be tied to the basic function of introspection in our mental economy. The further one moves away from that basic function, the less reason there is to think that introspection will be able to address the questions that are put to it.⁸

As an analogy consider the constraints on visual perception. Vision enables us to perceive ordinary objects in our environment, but we do not expect it to give us access to all of the properties of such objects. Some properties are of the wrong scale, being either too small or too large to be detected by the human visual system. Other properties—such as being made in France, being rare, being expensive—are just not the kinds of properties that our visual system has the function of detecting, and thus we wouldn't expect these properties to be given in the contents of perception. Similar, introspection has also a function, and the domain within which introspection is reliable is likely to be closely tied to that function.

Of course, providing an account of the function of introspection is by no means uncontroversial. The spectrum of views on offer ranges from the view that introspection has no (useful or serious) function at all to the view that it provides an incorrigible foundation for empirical knowledge. Leaving the wide questions to one side, we suggest

⁸ See Goldman (2003) for a related point.

that at least part of the basic function of introspection is to facilitate explicit rational and practical assessment of one's situation in the light of one's current experiential circumstances. It is often useful to be able to reflect upon how things consciously strike one, and in what manner or mode one perceives or thinks about one's surrounding at a given moment. Awareness of one's current conscious experiences and thoughts facilitates the rational appraisal of theoretical and practical deliberation. According to this picture, the basic function of introspection is attuned to the basic function of perception: introspection reinforces and improves the subject's capacity to successfully navigate and manipulate her physical surrounding. If this is even roughly right, then the normal operational conditions of introspection are limited to relatively banal information concerning the mode and contents of experience. Information about the metaphysical structure of perceptual experience, by contrast, is not likely to improve the subject's ability to deal with her environment, and so is unlikely to be part of how the mind is available to first-person reflection.⁹

The foregoing might explain why introspective judgements at the centre of disputes about the nature of perceptual experience are unreliable, but it does not explain why introspection gives rise to *divergent* reports here. Appeal to the influence of background belief supplies the missing part. Arguably, the introspective judgments of some (or even all) of the participants in this debate are significantly modulated by their theoretical commitments (Spener ms). Theorists who endorse pure representationalism might be driven in part by a general commitment to a physicalism and the worry that any other view about experience makes defence of physicalism problematic. Some pure representationalists might also be committed to the view that theorizing about conscious experience ought to preserve as much as possible of common sense belief. This background commitment might lead some theorists to the claim that introspective attention to perceptual experience presents one with only external, mind-independent objects and properties.

Impure representationalism might be motivated by inverted spectrum thought experiments. The latter are taken to show the character of perceptual experience must

⁹The notion of introspective operational constraints does not presuppose a broad perceptual model of introspection (see Shoemaker 1996). It is compatible with our view that the rational and practical appraisal of one's situation (i) can only be achieved by a subject with the capacity to introspect and (ii) that it is constitutive of having conscious experiences that they could potentially figure in such appraisal.

have non-representational features. Theorists who accept the existence of such features might then be inclined to think that the presence of these features is introspectable. Or they might focus on the idea that perceptual experience is suffused with subjectivity (e.g. Levine 2001: 6-7) or 'me-ishness' (Block 1995). Such theorists might believe that experience involves some essential reference to the subject of experience itself and that this is manifest in its phenomenal character. Jointly or separately, background expectations such as these might drive theorists' introspection-based judgements that some features are experienced *as* mind-dependent.

Of course, these accounts are highly speculative; we have no direct evidence that these introspective judgements have been shaped by background commitments in precisely these ways. However, we do hold that a story along these lines provides a plausible diagnosis of this debate.

4.2 Cognitive phenomenology revisited

We turn now to the debate over cognitive phenomenology. Here, too, we can set aside the individual differences strategy as unpersuasive. There is good reason to assume that the presence or absence of cognitive phenomenology would manifest in cognitive and behavioural differences, but as far as we can tell no such differences are to be found—apart, that is, from the fact that liberals endorse, whereas conservatives reject, the claim that there is cognitive phenomenology.

What about the suggestion that this debate requires introspection to operate outside of its reliable bounds of operation? Again, on first glance this analysis of this dispute is not attractive. Both liberals and conservatives appear to produce their introspective judgments under optimal conditions. Neither party is restricted to situations involving high cognitive load, time pressure, etc. But appeal to operational constraints in the explanation of this debate might nonetheless be instructive. This is because the debate about cognitive phenomenology is not a debate about whether there *is* a phenomenology of thought. Rather, it is a debate about whether the phenomenology of thought includes phenomenal states that are *sui generis*, i.e. whether the phenomenology of thought is exhausted by the phenomenology of the sensory or imagistic accompaniments of thought, or whether it contains a non-sensory and non-imagistic component. And this question may not be one that falls within introspection's domain.

Again, this strategy cannot explain why introspection gives rise to divergent reports, but here too we might appeal to the influence of expectation (as, indeed, the participants in this debate often do when explaining what has gone wrong with their opponents, see e.g. Siewert (2011); Strawson (2011); Tye (2011)). Motivations on each side can be found easily.

Those who are attracted to the thought that intentional content is grounded in phenomenal character have reason to think that there must be a proprietary and distinctive phenomenology of thought. More generally, those attracted to anti-reductionist accounts of phenomenal consciousness might also be attracted to the view that thoughts have a cognitive phenomenology. By contrast, those attracted to certain reductive accounts of phenomenal consciousness have reason to deny that thought has a proprietary and distinctive phenomenal character, for it would pose a further obstacle to their reductive vision. Conservatism might also be motivated by the desire not to multiply phenomenal properties beyond necessity, together with the conviction that the phenomenology of thought can be accounted for in terms of more familiar kinds of phenomenal properties.

We also want to make room for a diagnosis of this debate in terms of terminological variation. According to this explanation, liberals and conservatives alike are aware of the same introspectively accessible facts but they describe those facts differently because they understand the key terms in different ways. Arguably the main culprit here is the term 'phenomenal' and its cognates. Conservatives, one might think, implicitly equate phenomenal character with qualitative or sensory character, a usage according to which thoughts couldn't possess proprietary and distinctive phenomenal character. Liberals, on the other hand, use the term 'phenomenal' in a wider sense—a sense that is not restricted to the qualitative or sensory. Of course, both parties to this dispute are (at times) willing to unpack the notion of phenomenal character in terms of 'what it's likeness', but there is reason to doubt whether this phrase is itself employed univocally. Indeed, the participants in this debate sometimes suggest their use of 'what it's likeness' departs from that of their opponents (see Tye 1996: 302, n. 3; Georgalis 2006: 69; Flanagan 1992: 67). There is good reason to suspect that conservatives and liberals might simply be talking past each other. In short, there are two plausible accounts of the debate about cognitive phenomenology. One account appeals to the combined influence of operational limitations and expectation; the other account appeals to terminological differences.

4.3 The richness of consciousness revisited

Let us now turn to the dispute about the richness of consciousness. As we noted, there are at least two manifestations of this debate: one concerns the question of whether or not consciousness includes multiple modalities at a single time; another concerns the amount of detail presented in visual experience.

Could individual differences provide an explanation of the different introspective reports associated with rich and austere theorists? Might those who produce introspective reports

that support the rich view have a richer stream of consciousness than those who produce introspective reports in support of the austere view? That seems unlikely. Here, too, one would expect such differences to manifest themselves in cognitive and behavioural capacities closely associated with consciousness. For example, one would expect abundant theorists to be able to report more of the items in a Sperling paradigm than sparse theorists are able to report. Although it would be premature to rule out the possibility of such differences, we are not aware of any evidence for them and would be rather surprised if such evidence were forthcoming.

What about a terminological analysis of this debate? Perhaps advocates of these two positions actually enjoy the same (or at least very similar) states, but report those states in rather different terms? There are a number of reasons to take this suggestion seriously. The two positions are usually presented impressionistically rather than in careful and quantitative detail. Theorists claim that the contents of consciousness are ‘sparse’ or ‘abundant’, ‘poor’ or ‘rich’, but these terms are rarely given precise analyses. Even those theorists who do specify their position with some exactness—say, by claiming that visual experience is (or is not) restricted to attention—might be talking past one another, for what one subject regards as falling outside of attention might not be so-regarded by another subject.

Further, subjects in the Sperling paradigm often say that they are aware of each—or at least most—of the alphanumeric figures in the display, but such claims are crucially ambiguous between two readings. On the one hand, subjects could be reporting only ‘generic phenomenology’—that is, an experience of each of the items as alphanumeric figures (see Grush (2007); Byrne et al. (2007)). Alternatively, subjects could be reporting ‘non-generic phenomenology’—an awareness of the determinate identity of each of the letters (say, that the letter in the top left hand corner was a ‘9’). The former interpretation of subjects’ reports is consistent with a relatively sparse view of visual experience, whereas the latter reading would support a richer conception of the contents of visual experience.

Terminological factors might figure in an analysis of the debate about the richness of experience, but it is likely that the lion’s share in a full explanation of this dispute will invoke operational constraints on introspection. Claims about the richness of consciousness will need to engage with the question of whether there is consciousness outside of—or at least at the margins of—attention, and this clearly requires introspection to operate at the very limits of its capacity to produce trustworthy judgments.

Of course, here too we face the question of why subjects go on to produce different reports. If (say) our perceptual experience is 'moderately rich', why do some theorists under-report its content whilst others over-report its content? Part of an answer may involve the fact that subjects have different thresholds (or criteria) for deciding whether or not a representation is conscious. Those with an undemanding criterion might be inclined to think that they enjoy quite a lot of conscious content outside of attention, whereas those with a more demanding criterion will be more likely to think that attention limns the limits of consciousness. Furthermore, any noise in the system will also be exploited by subjects' theoretical expectations and biases. Indeed, this is the analysis that most of those engaged in this debate tend to give of what has gone wrong with their opponents. It is common for advocates of the sparse conception of experience to argue that the introspective judgments of their opponents are expectation-driven: those who report a rich experience of the world mistakenly assume that their experience is detailed because they have experiential access to a richly detailed world. Advocates of the rich conception of experience argue that their opponents produce the introspective reports that they do only because they are antecedently committed to a sparse conception of experience. There is, then, sufficient reason to think that the introspective reports of both sparse theorists and abundant theorists could be driven by expectation along these lines.

5. Quarantining the sceptical threat

Consider again some typical trustworthy cases. Suppose that you are in what you take to be introspectively good conditions and judge that you are in pain, or that you are seeing something blue before you, or that you are hungry. In the light of sceptical worries arising in the context of the three debates, should you now be worried about the trustworthiness of these introspective judgements? On reporting them, should we be worried?

Although there are common elements to the diagnoses we offered in the previous sections, no single explanatory strategy provided a full account of all three disputes. There is thus no simple generalizing move to be made, extending scepticism about introspection from one area to all. In light of this, we need to consider whether the possible sources of disagreement might, as a group, make mischief for putatively trustworthy cases.

The first two of the four sources of introspective disagreement—namely individual differences and terminological disagreement—can be set to one side immediately, for explanations in terms of terminological variation or individual differences do not lead to scepticism about introspection as such. Of course, if these factors were pervasive then they might generate worries about the utility of *introspective methodology*—in other words, they

might undermine the use of introspection to theorize about the mind. But there is little reason to think that either individual variation or terminological differences *are* prevalent in the context of apparently trustworthy cases, for such cases are characterized by introspective agreement. Two (normally functioning) subjects gazing at the sky would agree that they have blue experiences; if they both stub their toes and wince, they would both agree that they are in a state of pain. Agreement in introspective judgment, in turn, suggests agreement in both terminology and introspective states.

Let us turn to the question of operational constraints. We have seen that subjects are sometimes inclined to make introspective judgments under conditions in which introspection is not reliable, either because the conditions in question are not hospitable or because introspection is required to address questions that fall outside of its domain. Might this also occur in the context of apparently trustworthy cases?

The first thing to ask is just how radical the radical sceptic is prepared to be. If the sceptic is willing to deny that introspection has any genuine function in our mental economy whatsoever then we have nothing to say at this point, other than that this position seems counter-intuitive in the extreme and is not motivated by introspective disagreement. So let us suppose that the sceptic is willing to grant that introspection has a function. In that case, she ought also to grant that there will be a range of conditions under which introspection will generate reliable judgments, namely those conditions that are tied to its basic function. Moreover, we have some reason for thinking that there will be a relatively close match between those conditions in which we intuitively take introspection to be trustworthy and those conditions in which it actually is trustworthy. Recall our proposal of the basic function of introspection as outlined in §4.1, according to which introspection is in the business of facilitating rational and practical assessment of one's situation on the basis of current experience. If this proposal is on the right path, then we should expect—as, indeed, we do—that ordinary, coarse-grained introspective judgments about the subject-matter of perception and thought are trustworthy.

The most forceful case for radical scepticism involves an appeal to the influence of expectation. The sceptic might argue that inter-subjective agreement in so-called trustworthy cases results from nothing more than a shared set of expectations about the kinds of states that are the targets of introspection. But, the thought goes, this is to put introspection at the mercy of expectation, for the accuracy of introspection would derive from that of the expectations that govern it. And this, it might be said, would be to give in to a radical introspective scepticism of sorts. Introspection as such would not merit our trust. What should we say about this argument?

The first point to be made is that introspection could be trustworthy *even if* it were subject to the influence of expectation on a grand scale. The influence of expectation will lead to introspective *error* only when the expectations involved are mistaken. Being expectation-driven to this degree might make introspective judgments vulnerable to error, but being vulnerable to error is not to be confused with actually being unreliable. The sceptic would need to show that the expectations governing a class of apparently trustworthy judgments were not only influenced by expectations but that those expectations were largely mistaken. Of course, the sceptic might reply that introspection could not generate knowledge if it were vulnerable to the influence of expectation, for knowledge requires a degree of ‘modal robustness’ that might be undercut by such vulnerability. We grant that there is a challenge here, but the onus is on the sceptic to show that intuitively trustworthy cases are vulnerable in a way that is at odds with the possibility of introspective knowledge.¹⁰

But—and this is the second point—we should not think of introspection as invariably subject to the influence of expectation on a massive scale. In the same way that perceptual judgment is more or less expectation-driven depending on context, so too we should expect that the degree to which introspective judgments are expectation-driven will also be context-dependent. Expectation will have a greater influence on judgment in contexts in which the signal from the target state is weak; correlatively, the stronger the signal from the target state, the less room there will be for expectation to impact on the subject’s judgment. Consider the influence of expectation on perception: the less stimulus-driven information the subject has, the more influence expectation will have on perceptual judgment. The upshot of these considerations is that it is reasonable to suppose that the influence of expectation on introspective judgment will be markedly weaker in contexts of apparently trustworthy cases than it is in the context of introspective disputes.

In sum, none of the sources of disagreement that figure in our diagnoses of the three introspective disputes threatens the trustworthy cases. We do have good reason to doubt the reliability of introspective judgements in the context of the disputes considered, but these reasons do not support any form of radical introspective scepticism.

Thus far our defence of introspection has been a negative one in that we have sought to quarantine the threat of disagreement-based scepticism. In the remainder of this section we sketch a positive case in favour of the trustworthiness of certain introspective judgments. Begin by distinguishing between two sorts of introspective judgments:

¹⁰ We thank Ian Phillips for prompting us to say more here.

scaffolded judgments and *freestanding* judgments. A scaffolded judgment is an introspective judgment about a conscious state P the content of which matches closely the content of a different judgment the subject would be disposed to make when endorsing the content of P. A freestanding judgement is an introspective judgement that is not scaffolded. For example, where P is a conscious perceptual experience, the content of a scaffolded judgement matches closely the content of a perceptual judgement one is disposed to make strictly on the basis of what P presents to one. Suppose you see a red tomato and you introspectively judge that it looks to you as if there is a red tomato before you. The content of that introspective judgement closely matches the content of a perceptual judgment you would be disposed to make on the basis of what your visual experience presents to you: that there is a red tomato before you. Contrast this with another introspective judgment you might make: that you are currently enjoying a perceptual image of a red tomato. There is no close content match between that judgement and a perceptual judgement you are disposed to make when endorsing what the experience presents. The second introspective judgment is freestanding whilst the first is scaffolded. Note that introspective judgments can be more or less scaffolded depending on the degree to which the content of the introspective judgment matches that of the relevant first-order judgement.¹¹

Many of the introspective judgments we are naturally inclined to regard as trustworthy are strongly scaffolded. Introspective judgements such as ‘There appears to be something blue before me’ or ‘I am thinking about how delicious fresh blueberries are’ have contents that correspond to perceptual judgements and thoughts that the subject will be disposed to endorse. By contrast, the introspective judgments that occur in the context of introspective disputes—risky cases—will often be freestanding (or at most weakly scaffolded). For example, there are no first-order judgements whose contents match those of the introspective judgments that one finds in the cognitive phenomenology debate. The same point holds for the debate about the nature of perceptual experience. Neither the introspective judgments made by pure representationalists nor those made by impure representationalists are scaffolded.

Matters are slightly more complicated when it comes to the debate about the richness of experience. Advocates of the sparse conception point to the fact that subjects can only reliably produce perceptual reports of only a very limited range of objects. Although this

¹¹ There are obvious points of contact between scaffolding and the discussion of the so-called transparency or diaphanousness of consciousness. However, exploration of these points of contact will have to wait for another occasion.

fact does provide some support for the sparse view, there is no perceptual report here that qualifies as scaffolding the introspective reports of sparse theorists. What about the introspective reports of those who endorse a rich conception of experience? These reports do appear to be scaffolded, for subjects are typically disposed to produce the first-order judgment that they are aware of a great deal of their environment. But although this introspective judgment is scaffolded, it is arguable that *both* the perceptual judgment and its introspective counterpart owe more to top-down expectation than to perceptual experience. (This is one area in which the veridicality of the subject's perceptual judgments bears on the trustworthiness of their introspective judgments: we would have more reason to think that the subject's introspective judgments were correct if their first-order perceptual judgments were veridical than we do given that they are not veridical.) In this way, the contrast between scaffolded and freestanding judgments provides us with a positive characterization of the difference between trustworthy and risky cases: the former tend to be strongly scaffolded, the latter tend to be freestanding (or weakly scaffolded).

We might also expect that weakly scaffolded judgments will be less secure than strongly scaffolded judgments but more secure than freestanding judgments. Introspective judgements about affective experience provide some evidence in favour of this proposal. Such judgements are weakly scaffolded insofar as the first-order judgments which the subject might be disposed to make when endorsing the content of that emotional experience will not typically bear a very tight match to the content of the introspective judgment. (In part, this is because it is unclear what the precise content of the first-order judgements would be.) This suggests that introspective access to affective experience is likely to be problematic, and indeed there is a good amount of empirical evidence for this view (Haybron 2007; also Schwitzgebel 2008).

An important upshot of the contrast between scaffolded and freestanding judgements is this: it reveals a source of evidence for the reliability of scaffolded judgments that is not available for freestanding judgments. Consider the scaffolded introspective judgment "I see a frog in front of me." We can 'check' this judgment in that we can ask whether its subject is willing to endorse the perceptual judgment "There is a frog in front of me." In this way, facts about which first-order judgments the subject is prepared to endorse can be used as evidence for claims about whether or not their scaffolded judgments are trustworthy. Such facts, however, are not available to us in the context of free-standing judgments. Consider the debate about the nature of cognitive phenomenology: because the introspective judgments that give rise to this debate are freestanding, we have no independent way of evaluating rivals introspective judgements about the nature of

thought. Given that many, perhaps even most, of the introspective judgments that intuitively strike us as most trustworthy are scaffolded, we will generally have an independent source of evidence in favour of the reliability of intuitively trustworthy judgments that is lacking in the context of freestanding judgments.¹²

6. Conclusion

Our aim in this paper has been to dismantle the disagreement-based case for radical introspective scepticism. We began with three introspective disputes: a dispute about the nature of perceptual experience, a dispute about the nature of cognitive phenomenology, and a dispute about the richness of experience. We argued that it is possible to provide plausible diagnoses of each dispute by appealing to one or more of four possible sources of introspective disagreement: individual differences, terminological variation, operational constraints and influence of background commitment and expectation. A central component of our diagnoses is that introspective judgements in these debates are unreliable because they require that introspection be deployed outside its normal range of operation. This is at least one way in which our diagnoses make room for scepticism about introspection as used by theorists in these debates. However, we argued that the sceptical threat can be quarantined, and that it does not infect the use of introspection in the context of intuitively trustworthy cases. Moreover, in some of the trustworthy cases there is positive reason to think that introspection is reliable insofar as the judgments in question will be suitably scaffolded by perceptual judgments. When it comes to introspection we should be neither radical sceptics nor uncritical fideists. Instead, we should adopt an attitude of introspective humility. Introspection is not to be rejected outright, but its use in the study of consciousness should be accompanied by a great deal more caution than is currently the case.¹³

¹² Of course, this evidence is not decisive. One could make the judgment that there is a frog in front of one without having a perceptual experience as of a frog, and one could have a perceptual experience as of a frog without judging (or being prepared to judge) that there is a frog in front of one. But given the intimate links between perceptual experience and perceptual judgment we will often be in a position to use the subject's (reports of) perceptual judgments as evidence of their perceptual states.

¹³ We are very grateful to Jakob Hohwy, Ian Phillips, and Eric Schwitzgebel for their valuable comments on an earlier draft of this paper.

References

- Bayne, T. and Montague, M. 2011. Cognitive phenomenology: an introduction. In T. Bayne & M. Montague (eds.) *Cognitive Phenomenology*. Oxford: Oxford University Press.
- Blackmore, S.J., Brelstaff, G., Nelson, K., Troscianko, T. 1995. Is the richness of our visual world an illusion? Transsaccadic memory for complex scenes. *Perception* 24: 1075-81.
- Block, N. 1995. On a confusion about the function of consciousness, *Behavioral and Brain Sciences* 18/2: 227-287.
- Block, N. 2003. Mental paint. In M. Hahn and B. Ramberg (eds.) *Reflections and Replies: Essays on the Philosophy of Tyler Burge*, Cambridge, MA: MIT Press, pp. 165-200.
- Block, N. 2007. Consciousness, accessibility, and the mesh between psychology and neuroscience, *Behavioral and Brain Sciences*, 30: 481-548.
- Boring, E.G. 1929/1950. *A History of Experimental Psychology* (2nd ed). New York: Century.
- Byrne, A., Hilbert, D.R. & Siegel, S. 2007. Do we see more than we can access? *Behavioral and Brain Sciences*, 30/5-6: 201-502.
- Chalmers, D. 2003. The content and epistemology of phenomenal belief. In Q. Smith and A. Jokic (eds.) *Consciousness: New Philosophical Perspectives*. Oxford and New York: Oxford University Press.
- Christensen, D. 2009. Disagreement as Evidence: The Epistemology of Controversy, *Philosophy Compass* 4.
- Dennett, D. 1991. *Consciousness Explained*. Boston: Brown and Little.
- Dokic, J. & Pacherie, E. 2007. Too much ado about belief. *Phenomenology and the Cognitive Sciences*, 6: 185-200.
- Dretske, F. 1995. *Naturalizing the Mind*. Cambridge, MA: MIT Press.
- Flanagan, O. 1992. *Consciousness Reconsidered*. Cambridge, MA: MIT Press.
- Georgalis, N. 2005. *The Primacy of the Subjective*. Cambridge, MA: MIT Press.
- Goldman, A. 2003. Epistemology and the evidential status of introspective reports. In A. Jack and A. Roepstorff (eds) *Trusting the Subject*, Vol 2. Imprint Academic, pp. 1-16.

- Grush, R. 2007. A plug for generic phenomenology. *Behavioral and Brain Sciences*, 30: 504-5.
- Haybron, D. 2007. Do we know how happy we are? On some limits of affective introspection and recall. *Noûs*, 41(3): 394-428.
- Horgan, T. & Tienson, J. 2002. The intentionality of phenomenology and the phenomenology of intentionality. In David J. Chalmers (ed.) *Philosophy of Mind: Classical and Contemporary Readings*. Oxford: OUP.
- Horgan, T. & Kriegel, U. 2007. Phenomenal epistemology: what is consciousness that we may know it so well? *Philosophical Issues* 17 (1): 123-144.
- Hurlburt, R. & Schwitzgebel, E. 2007. *Describing Inner Experience?* Cambridge: MIT Press
- Kriegel, U. 2007. The phenomenologically manifest, *Phenomenology and the Cognitive Sciences*, 6:115-36.
- Levine, J. 2001. *Purple Haze*. Oxford: Oxford University Press.
- Levine, J. 1995. Qualia: intrinsic, relational or what? In *Conscious Experience*, ed. by T. Metzinger. Paderborn: Schoenigh.
- Lormand, E. 1996. Nonphenomenal consciousness. *Noûs*, 30: 242-61.
- Lyons, W.E. 1986. *The Disappearance of Introspection*. Cambridge, MA: MIT Press.
- Martin, M. 1990. *Atheism*. Philadelphia: Temple University Press.
- Martin, M.G.F. 2002. The transparency of experience. *Mind and Language*, 17/4: 376-425.
- Nelkin, N. 1989. Propositional attitudes and consciousness, *Philosophy and Phenomenological Research*, 49 (3): 413-30.
- O'Regan, K. 1992. Solving the 'real' mysteries of visual perception: The world as an outside memory, *Canadian Journal of Psychology*, 46: 461-88.
- Peacocke, C. 1983. *Sense and Content*. Oxford: Clarendon Press.
- Pitt, D. 2004. The phenomenology of cognition Or What is it like to think that P? *Philosophy and Phenomenological Research*, 69: 1-36.
- Robinson, W.S. 2005. Thoughts without distinctive non-imagistic phenomenology. *Philosophy and Phenomenological Research*, LXX: 534-61.
- Shoemaker, S. 1994. Self-knowledge and 'inner sense'. Lecture II: The broad perceptual model, *Philosophy and Phenomenological Research*, 54: 271-290.
- Searle, J. 1992. *The Rediscovery of the Mind*. Cambridge, MA: MIT Press.

- Siewert, C. 1998. *The Significance of Consciousness*. Princeton, NJ: Princeton University Press.
- Siewert, C. 2007. Who's afraid of phenomenological disputes?, *Southern Journal of Philosophy*, XLV: 1-21.
- Siewert, C. 2011. Phenomenal thought. In T. Bayne and M. Montague (eds) *Cognitive Phenomenology*. Oxford: Oxford University Press.
- Schwitzgebel, E. 2007a. No unchallengeable epistemic authority, of any sort, regarding our own conscious experience – Contra Dennett? *Phenomenology and the Cognitive Sciences*, 6: 107-113.
- Schwitzgebel, E. 2007b. Do you have constant tactile experience of your feet in your shoes? Or is experience limited to what's in attention?, *Journal of Consciousness Studies*, 14/3: 5-35.
- Schwitzgebel, E. 2008. The unreliability of naïve introspection, *The Philosophical Review*, 117/2: 245-73.
- Spener, M. Manuscript. Phenomenal adequacy and introspective evidence.
- Spener, M. 2011. Disagreement about cognitive phenomenology. In T. Bayne and M. Montague (eds) *Cognitive Phenomenology*. Oxford: Oxford University Press.
- Sperling, G. 1960. The information available in brief visual presentations, *Psychological Monographs*, 74, no. 11.
- Strawson, G. 1994. *Mental Reality*. Cambridge, MA: MIT Press.
- Strawson, G. 2011. Cognitive phenomenology: Real life. In T. Bayne and M. Montague (eds) *Cognitive Phenomenology*. Oxford: Oxford University Press.
- Tye, M. 1992. Visual qualia and visual content. In T. Crane (ed.) *The Contents of Experience*. Cambridge: Cambridge University Press.
- Tye, M. 1995. *Ten Problems of Consciousness*. Cambridge, MA: MIT Press.
- Tye, M. 1996. The function of consciousness, *Noûs*, 30: 287-305.
- Tye, M. 2006. Content, richness, and fineness of grain. In T.S. Gendler and J. Hawthorne (eds) *Perceptual Experience*. Oxford: Oxford University Press.
- Tye, M. 2011. Is there a phenomenology of thought? In T. Bayne and M. Montague (eds) *Cognitive Phenomenology*. Oxford: Oxford University Press.
- Wilson, R. 2003. Intentionality and phenomenology. *Pacific Philosophical Quarterly*, 84 (4): 413-431.

Woodworth, R.S. 1906. Imageless thought. *The Journal of Philosophy, Psychology and Scientific Methods*, 3/26: 701-08.

Zimmerman, M. 1989. The nervous system in the context of information theory. In R. F. Schmidt & G. Thews (eds.), *Human Physiology* (2nd Ed.) Berlin: Springer-Verlag, pp. 166-173.