
Introspective Insecurity

Tim Bayne

This paper examines the case for pessimism concerning the trustworthiness of introspection. I begin with a brief examination of two arguments for introspective optimism, before turning in more detail to Eric Schwitzgebel's case for the view that introspective access to one's own phenomenal states is highly insecure. I argue that there are a number of ways in which Schwitzgebel's argument falls short of its stated aims. The paper concludes with a speculative proposal about why some types of phenomenal states appear to be more introspectively elusive than others.

Keywords

Cognitive phenomenology | Emotion | Freestanding judgments | Imagery | Introspection | Introspection-reliant | Optimism | Pessimism | Scaffolded judgments | Schwitzgebel

Author

Tim Bayne

tim.bayne@manchester.ac.uk

The University of Manchester
Manchester, United Kingdom

Commentator

Maximilian H. Engel

M.H.Engel.1@student.rug.nl

Rijksuniversiteit Groningen
Groningen, Netherlands

Editors

Thomas Metzinger

metzinger@uni-mainz.de

Johannes Gutenberg-Universität
Mainz, Germany

Jennifer M. Windt

jennifer.windt@monash.edu

Monash University
Melbourne, Australia

1 Introduction

There is a curious ambivalence in current attitudes towards our epistemic relationship to consciousness. Some theorists hold an optimistic view of the powers of introspection, regarding judgments about one's current experiences as epistemically secure—perhaps some of the most secure judgments that we make. Optimists rarely claim that we have exhaustive and infallible access to consciousness, but they do hold the epistemic credentials of introspection in high regard, at least when introspection is directed towards the phenomenal character of consciousness. Those inclined to optimism don't doubt that it is possible to mis-remember or mis-report one's experiences, but they tend to assume that one has some

kind of epistemic access to one's experiences simply by having them.¹

Running alongside this vein of optimism is a rather more pessimistic strand of thought, according to which the epistemic credentials of introspection are chronically insecure. Far from regarding introspection as a light that illuminates every corner of consciousness, pessimists suspect that significant swathes of experience are accessible to introspection only with great difficulty if at all.²

¹ Theorists inclined towards optimism include Chalmers (2003), Gertler (2012), Goldman (2004), Horgan et al. (2006), Horgan & Kriegel (2007), Siewert (2007), and Smithies (2012).

² The contrast between "optimists" and "pessimists" is far from sharp, for optimists often grant that epistemic access to consciousness can be (very) challenging, and pessimists often allow that there are experiential domains with respect to which introspection is trust-

Glossary

Introspection	An unmediated judgment that has as its intentional object a current psychological or phenomenal state of one's own.
Discrimination	The capacity to attentively single the state out from amongst the other experiences that one has at the time in question.
Categorize	To categorize a phenomenal state is to locate it within a taxonomy of some kind.
Directly and indirectly introspective judgments	A direct introspective judgment concerns the phenomenal character/content of one's current phenomenal state(s) and is grounded in a single act of introspective attention, whereas an indirect introspective judgment concerns the general nature of one's conscious experience and is not grounded in a single act of introspective attention.
Scaffolded judgments	An introspective judgment is scaffolded if and only if it is accompanied by a disposition to make a first-order judgment (e.g., a perceptual judgment) whose content broadly corresponds to the judgment of the introspective judgment. For example, the judgment that one has a visual experience as of a red tomato in front of one is scaffolded insofar as it is accompanied by a disposition to make the perceptual judgment that there is a red tomato in front of one.
Freestanding judgments	An introspective judgment is freestanding if and only if it is not accompanied by a disposition to make a first-order judgment (e.g., a perceptual judgment) whose contents broadly corresponds to the judgment of the introspective judgment.

According to Dan Haybron, “[...]even the gross qualitative character of our conscious experience can elude our introspective capacities” (Haybron 2007, p. 415). Sounding a similar note, Maja Spener has argued that “philosophers and psychologists routinely overestimate the epistemic credentials of introspection in their theorizing” (Spener unpublished; see also Spener 2011a, 2011b, and 2013). But perhaps the most thoroughgoing pessimist is Eric Schwitzgebel:

Most people are poor introspectors of their own ongoing conscious experience. We fail not just in assessing the causes of our mental states or the processes underwriting them; and not just in our judgments about nonphenomenal mental states like traits, motives and skills, and not only when we are distracted, or passionate or inattentive or self-deceived, or pathologically deluded or when we're re-

worthy. Nonetheless, these terms are useful insofar as they capture the overarching attitude that the two groups of theorists express with regard to introspection.

flecting about minor matters, or about the past, or only for a moment, or when fine discrimination is required. We are both ignorant and prone to error. There are major lacunae in our self-knowledge that are not easily filled in, and we make gross, enduring mistakes about even the most basic features of our currently ongoing conscious experience (or “phenomenology”), even in favourable circumstances of careful reflection, with distressing regularity. (2008, p. 247)

Although Schwitzgebel's pessimism is tempered by moments of optimism, the dominant theme in his work is that introspection cannot be trusted to reveal anything other than the most mundane features of consciousness. Descartes, Schwitzgebel argues, “had it quite backwards when he said the mind—including especially current conscious experience—was better known than the outside world” (2008, p. 267).

I feel the pull of both optimism and pessimism. In my optimistic moments I find it hard

to take seriously the suggestion that I might be guilty of “gross and enduring mistakes” about the basic features of my current phenomenology. But the arguments for pessimism are powerful and not easily dismissed, and I worry that Schwitzgebel is right when he suggests that the allure of optimism might be due to nothing more than the fact that “no-one ever scolds us for getting it wrong” (2008, p. 260).

A central aim of this paper is to provide an overview of Schwitzgebel’s case for introspective pessimism, and to chart a number of ways in which the optimist might respond to it. But although this paper can be read as a defence of a kind of optimism, my central concern is not so much to take sides in this debate as to advance it by noting various complexities that have perhaps been overlooked. But before turning to the debate itself let me make a few comments about its importance. An account of the trustworthiness of introspection is likely to have a bearing on two important issues. Most obviously, it has implications for the use of introspection as a source of evidence regarding philosophical and scientific debates about consciousness. Whether or not introspection is our sole form of access to consciousness, there is no doubt that it is currently treated as a *central* form of such access, and thus doubts about the reliability of introspection engender doubts about the viability of the study of consciousness. A second issue on which the trustworthiness of introspection has an important bearing concerns debates about the nature of introspection, and in particular the relationship between introspection and consciousness. Some accounts of introspection take a person to be necessarily acquainted with his or her conscious states, where acquaintance is an epistemic relationship of a particularly intimate kind (Gertler 2012; Horgan et al. 2006; Smithies 2012). It is fair to say that such approaches are optimistic by nature, and although advocates of such accounts have attempted to accommodate the possibility of introspective ignorance and error (see e.g., Horgan 2012), the success of such attempts is very much an open question. Other accounts of “introspection”—such as those that deny that there are any distinctively first-per-

sonal modes of access to consciousness—can easily accommodate introspective ignorance and error, but they struggle to account for the epistemic security that often seems to characterize introspection. In short, an account of introspection’s epistemic profile would function as a useful constraint on accounts of its nature.

2 Motivating optimism

By “introspection” I mean an unmediated judgment that has as its intentional object a current psychological state of one’s own. Introspection can take as its object a wide variety of psychological states, but here I am concerned only with the introspection of *phenomenal states*—states that there is “something it is like” for the subject in question to be in. In principle one could have any number of reasons for self-ascribing a phenomenal state—for example, it is possible to self-ascribe pain on the basis of neural or behavioural evidence—but introspection involves the self-ascription of phenomenal states on the basis of seemingly “direct” contact with them.³

There are many aspects of consciousness with respect to which we clearly have little to no introspective access. For example, introspection is clearly not a source of information about the neural basis of consciousness or its functional role. But surely, one might think, introspection can provide trustworthy answers to such questions as, “Am I now in a conscious state with such-and-such a phenomenal character?” Roughly speaking, to regard introspection as able to reveal the phenomenal character of one’s conscious states is to have an optimistic attitude towards it. But there is more than one sense in which introspection might be said to reveal the character of consciousness, and thus more than one way to be an introspective optimist.

One way in which introspection can reveal a phenomenal state is by allowing one to *discriminate* it from its phenomenal neighbours. I take discrimination to be bound up with the ca-

³ Introspection may involve direct access to consciousness at a personal level and yet also be inferential and indirect at sub-personal levels of description.

capacity to single the state out from amongst the other experiences—e.g., thoughts, perceptual experiences, and bodily sensations—that happen to populate one’s field of consciousness. Discriminative access to an experience allows one to direct one’s attention towards it and to thus make it the potential target of demonstrative thought—“I wish that *this* experience would stop”. A second mode of introspective access to consciousness involves the deployment of categories. To categorize a phenomenal state is to locate it within a taxonomy of some kind. Categorical access to the experience of an itch, for example, involves recognizing it as a phenomenal state of a certain type—a state, perhaps, that has a certain intensity, bodily location, and relations to other experiences. Categorical access is a more sophisticated form of access than discriminative access. Just as it is possible to discriminate a bird from its surroundings without being able to recognize it as a bird—perhaps all one can do is bring it under the demonstrative, “that thing there in the sky”—so too it may be possible to discriminate a phenomenal state without being able to recognize it as the kind of phenomenal state it is. Mature human beings enjoy some degree of categorical and discriminative access to their phenomenal states, but many conscious creatures—non-linguistic animals and young children, for example—may enjoy only discriminative access to consciousness.⁴

With this in mind, we can distinguish two forms of introspective optimism. Moderate introspective optimism holds that being in a phenomenal state typically brings with it the capacity to discriminate that state from its phenomenal neighbours, while a more radical form of introspective optimism holds that being in a phenomenal state typically brings with it the capacity to both discriminate and accurately categorize it. By the same token, introspective

pessimism can be more or less radical depending on whether its scope is restricted to categorical access (moderate) or includes both categorical and discriminative access (radical). In what follows, I use the terms “introspective optimism” and “introspective pessimism” to refer to the moderate versions of these views unless noted otherwise.

2.1 The phenomenological argument

Although introspective optimism is often assumed rather than explicitly argued for, I think it is possible to discern two lines of argument for it in the literature. Neither argument is conclusive, but taken together they go some way towards justifying the widespread endorsement of introspective optimism.

The first argument is phenomenological: introspection seems to reveal itself as providing a trustworthy source of information about consciousness. In other words, the epistemic security of introspection seems to be something that is manifest in its very phenomenology. Consider [Brie Gertler’s](#) description of what it is like to attend to the experience that is generated by pinching oneself:

When I try this, I find it nearly impossible to doubt that my experience has a certain phenomenal quality—the phenomenal quality it epistemically seems to me to have, when I focus my attention on the experience. Since this is so difficult to doubt, my grasp of the phenomenal property seems not to derive from background assumptions that I could suspend: e.g., that the experience is caused by an act of pinching. It seems to derive entirely from the experience itself. If that is correct, my judgment registering the relevant aspect of how things epistemically seem to me (this phenomenal property is instantiated) is directly tied to the phenomenal reality that is its truthmaker. (2012, p. 111)

I suspect that Gertler’s comments will strike a chord with many readers—they certainly resonate with me. Introspection seems not merely to

⁴ This claim would need to be tempered if as seems plausible discriminative access requires a minimal form of categorical access. Consider again the case of discriminating a bird but failing to recognize it as a bird. This counts as a failure of categorical access insofar as one fails to bring it under the concept <bird> (or related concepts such as <robin>), but it is arguable that in order to discriminate it from its perceptual background one (or one’s visual system) must bind the various visual features together as the features of a single object, which may require a minimal form of categorical access to the object.

provide one with information about one's experiences, it seems also to "say" something about the quality of that information. This point can be illuminated by contrasting introspection with other forms of access to consciousness. Suppose that you believe that you have the phenomenology associated with anger because a friend has pointed out that you are behaving angrily. In cases like this, testimony provides one with a form of access to one's phenomenal states, but this access surely lacks the epistemic security that introspective access typically possesses—or at least seems to possess. It would be very odd to put more faith in "third-person" evidence concerning one's own conscious states than "first-person" evidence.

Now, one might think that even if the phenomenological consideration just surveyed can *explain* why optimism seems so compelling, it surely can't provide any *justification* for it. Appealing to introspection itself in order to establish its epistemic credentials would be as futile as attempting to pull oneself up by one's own shoelaces. If it's introspection itself that is in the dock, how could its own testimony exonerate it?

In considering this objection we need to distinguish two questions. One question is whether introspection makes claims about its own veracity. A second question is what to make of such claims should they exist—that is, whether to regard them as providing additional reasons for thinking that introspection is trustworthy. Beginning with the first question, it seems to me not implausible to suppose that introspection *could* bear witness to its own epistemic credentials. After all, perceptual experience often contains clues about its epistemic status. Vision doesn't just provide information about the objects and properties present in our immediate environment, it also contains information about the robustness of that information. Sometimes vision presents its take on the world as having only low-grade quality, as when objects are seen as blurry and indistinct or as surrounded by haze and fog. At other times visual experience represents itself as a highly trustworthy source of information about the world, such as when one takes oneself to have a clear

and unobstructed view of the objects before one. In short, it seems not implausible to suppose that vision—and perceptual experience more generally—often contains clues about its own evidential value. As far as I can see there is no reason to dismiss the possibility that what holds of visual experience might also hold true of introspection: acts of introspection might contain within themselves information about the degree to which their content ought to be trusted.

The foregoing addresses the first of the two questions identified above but not the second, for nothing in what I have said provides any reason to think that introspection is a reliable witness to its own veracity. It is one thing for introspection to represent its deliverances *as* trustworthy but it is another for those deliverances to *be* trustworthy. But this being noted, it seems to me not unreasonable to think that the claims introspection makes on its own behalf should be afforded *some* degree of warrant. In general, we regard perceptual testimony as innocent unless proven guilty, and even if introspection is not itself a form of perception it seems reasonable to apply that same rule here. (After all, it is not clear why we would have acquired a cognitive capacity if its deployment routinely led us astray.) The phenomenological argument certainly doesn't provide any kind of proof for introspective optimism, but it seems to me to do more than merely explain why optimism is so attractive: it also provides it with some degree of justification.

2.2 The conceptual argument

A rather different argument for optimism takes as its point of origin the very notion of a phenomenal state. By definition, a phenomenal state is a state that there is "something that it's like" for the subject in question to be in. Conscious creatures enjoy mental states of many kinds, but it is only phenomenal states that bring with them a subjective perspective. But—so the argument runs—if a phenomenal state is a state that there is something it is like to be in, then the subject of that state *must* have epistemic access to its phenomenal character. A

state to which the subject had no epistemic access could not make a constitutive contribution to what it was like for that subject to be the subject that it was, and thus it could not qualify as a phenomenal state. Call this the *conceptual argument*.⁵

How compelling is this argument? It seems to me that a lot depends on what is implied by the notion of “epistemic access”. There is little to recommend the conceptual argument if “epistemic access” is understood in terms of categorization, for it seems fairly clear that a subject need not possess the capacity to accurately categorize its phenomenal states in order for them to contribute to its phenomenal perspective. Of necessity any phenomenal state will fall under categories of various kinds, but the nature of these categories need not be transparent to the creature experiencing it.

But suppose that we construe epistemic access in terms of categorization, rather than identification. Might the conceptual argument justify a moderate form of optimism, according to which subjects must have discriminative access to their phenomenal states? To make this clearer, suppose that it is possible for phenomenal states to occur within the modules of early vision of the kind that are concerned with determining (say) texture or colour constancy. Such phenomenal states—assuming that they are possible—would be completely inaccessible to the subject in question. The creature in question would be unable to contrast the phenomenal character of these states with the phenomenal character of any of its other experiences; it would be unable to single such states out for attention, and it would be unable to make them the objects of demonstrative thought. As such, it seems to me that it is very plausible to hold that they couldn’t be genuinely ascribed to the subject in question, but could at best be ascribed only to one of the subject’s perceptual modules. The root of this intuition, I suspect, lies with the thought that a phenomenal state to which the subject has no discriminative access couldn’t be anything “to” the subject—

that in the relevant sense of the phrase there couldn’t be anything “that it’s like” for the subject to have the relevant experiences.

Although attractive, this argument is not without its problems. One challenge comes in the form of creatures that lack introspective capacities. A creature without introspective capacities might be able to use its conscious states to discriminate some features of the world from others, but it would not be able to make its conscious states themselves objects of its own discriminative activities. And yet—the objection runs—it would be implausible to hold that creatures that lack the capacity for introspective discrimination cannot have phenomenal states. Intuitively, having phenomenal states is one thing and being able to discriminate one’s phenomenal states for each other is another—and more sophisticated—thing. Thus—the argument runs—discriminative access to a phenomenal state cannot be a necessary condition for being in that state.

I certainly agree that it would be implausible to restrict phenomenal states to creatures that possess introspective capacities, but perhaps the objection can be met without making such a restriction. What we can say is that when a creature does acquire introspective capacities those capacities bring with them the ability to discriminate its phenomenal states from one another (at least under epistemically benign conditions). So, we can grant that being in a phenomenal state doesn’t require discriminative access to that state, but also hold that creatures with introspective capacities will be able to discriminate their phenomenal states from one another (again, at least when conditions are epistemically benign).

A second objection to the conceptual argument concerns states that occupy the “margins” of consciousness—such as the unnoticed hum of the refrigerator or the background phenomenology of mood experiences. It is arguable that in some cases experiences like this not only fail to fall within the scope of introspection but in fact cannot be brought within its scope, for to attend to them would be to bring them into the “centre” of consciousness and thus change their phenomenal character. Such states serve as po-

⁵ There are echoes here of the claim that phenomenal consciousness entails a certain kind of “access consciousness”. For some relevant discussion see Church (1997) and Clark (2000).

tential counter-examples to the claim that creatures with introspective capacities must be able to discriminate their phenomenal states from one another.

In response, one might grant that even if the phenomenal states that occur in the margins of consciousness cannot be singled out for introspective attention, there is still a sense in which they can be the objects of discrimination. Not only can they be discriminated from one another, they can also be discriminated from those phenomenal states that *do* fall within the scope of attention. Indeed, if such states cannot be discriminated from their phenomenal neighbours in any way then it is unclear what reason we could have for thinking of them as falling within the margins of *consciousness* at all, rather than being completely unconscious.

Where do these considerations leave us? I have suggested that the phenomenological argument provides some reason to take at least a moderate form of optimism seriously. It doesn't, of course, establish that our access to all kinds of phenomenal states is robust—indeed, one might even appeal to phenomenological considerations to motivate the idea that our epistemic access to significant regions of phenomenal space is very poor. (I return to this topic shortly.) The conceptual argument provides little reason to think that we will always be able to *categorize* our phenomenal states, but it does provide some motivation for the idea that being in a phenomenal state brings with it the ability to *discriminate* that phenomenal state, at least when it comes to creatures with introspective capacities. In short, optimism of at least a moderate form is not merely a holdover from Cartesianism but can be provided with some degree of support. With these considerations in mind let us turn now to the case for pessimism.

3 Motivating pessimism

Two distinctions will prove helpful in what follows. One distinction is between forms of pessimism that concern only our capacity to identify our phenomenal states and forms of pessimism that call into question our capacity to both discriminate and categorize our phenomenal states. A

second distinction concerns the *scope* of pessimism. At one end of the spectrum are local forms of pessimism that concern only a relatively circumscribed range of phenomenal states (say, imagery experiences), while at the other end of the spectrum are forms of pessimism that are unrestricted in scope. Perhaps no theorist has ever embraced a truly global form of pessimism—even Schwitzgebel grants that introspection is trustworthy with respect to certain aspects of consciousness—but some forms of pessimism are clearly wider in scope than others. These two distinctions are, of course, orthogonal to each other. One could be a moderate but global pessimist; alternatively, one could endorse a radical but highly local form of pessimism.

So much for the varieties of pessimism—how might one argue for the view? One influential line of argument for pessimism—or at least something very much like it—appeals to the alleged privacy of introspection. Because an individual's introspective judgments cannot be checked by anyone else, it follows—so the argument runs—that it would be inappropriate to trust them. This argument is often used to motivate the view that introspection is scientifically illegitimate, but it could also be used to motivate the view that one should adopt a sceptical attitude towards one's own introspective capacities.⁶ Although it has been influential, I will leave this argument to one side in order to focus on a trio of arguments that aim to establish not merely that there is no positive reason to trust introspection (as the argument just outlined attempts to do), but that there is positive reason *not* to trust it. My presentation of these arguments will draw heavily on Schwitzgebel's work, for he has done more than any other author to develop and defend them.⁷

But before I examine those arguments, I want to consider the overall structure of

⁶ For critical discussion of this argument see Goldman (1997, 2004) and Piccinini (2003, 2011). In my view the most plausible response to it involves denying that introspection is private in the sense required for the argument to go through. I touch briefly on this idea in section 4.

⁷ Schwitzgebel is clearly attracted to a fairly global form of introspective pessimism, but (to the best of my knowledge) he doesn't distinguish between discriminative and categorical access, and thus it is unclear whether his version of pessimism is radical or merely moderate. Generally, however, he seems to have something akin to radical scepticism in mind.

Schwitzgebel's case for global scepticism. As I read him, Schwitzgebel employs a two-step strategy (2008, p. 259). The first step involves attempting to establish a form of local pessimism via one (or more) of the three argumentative strategies to be explored below. The second step involves generalizing from the kinds of phenomenal states that are the targets of local pessimism to phenomenal states in general. The second step is clearly required, for without it we would have no reason to regard introspection *in general* as “faulty, untrustworthy, and misleading”—“not just possibly mistaken, but massively and pervasively” (Schwitzgebel 2008, p. 259).⁸

I will consider both steps in due course, but the crucial point to note for now is that, considered in the abstract, the second step of the argument looks somewhat suspect (Bayne & Spener 2010). Even if there are hard cases for introspection—that is, cases in which introspective access to phenomenology is insecure—there also easy cases—that is, cases in which introspective access to phenomenology is clearly secure. Indeed, Schwitzgebel himself grants that introspection “may admit obvious cases” and that some aspects of visual experience “are so obvious it would be difficult to go wrong about them” (Schwitzgebel 2008, p. 253). But if that's the case, then one might well ask why we shouldn't generalize from those cases rather than from the hard cases on which he focuses. Schwitzgebel complains that to generalize about introspection only on the basis of the easy cases “rigs the game”. That's true. But it's equally true that to generalize only on the basis of the hard cases—as Schwitzgebel seems to do—would *also* rig the game. In fact, it would seem pretty clear that any comprehensive account of the epistemic landscape of introspection must

take *both* the hard and easy cases into consideration. Arguably, generalizing beyond the obviously easy and hard cases requires an account of what makes the hard cases hard and the easy cases easy. Only once we've made some progress with that question will we be in a position to make warranted claims about introspective access to consciousness in general. What this suggests is that although there is a formal distinction between the two steps of Schwitzgebel's argument, the steps are not entirely independent of each other, for the fortunes of the second step rest in part on the case that can be made for the first step. With that thought in mind, let us now turn to the arguments for pessimism.

3.1 The argument from dumbfounding

One line of argument that features prominently in Schwitzgebel's work is what I call the argument from dumbfounding.⁹ Arguments of this form involve posing introspective questions that allegedly stump us—questions that we find ourselves unable to answer with any significant degree of confidence. Here's an example of such an argument:

Reflect on, introspect, your own ongoing emotional experience at this instant. Do you even have any? If you're in doubt, vividly recall some event that still riles you until you're sure enough that you're suffering from renewed emotion. Or maybe your boredom, anxiety, irritation, or whatever in reading this essay is enough. Now let me ask: Is it completely obvious to you what the character of that experience is? Does introspection reveal it to you as clearly as visual observation reveals the presence of the text before your eyes? Can you discern its gross and fine features through introspection as easily and as confidently as you can, through vision, discern the gross and fine features of nearby external objects? Can you trace its spatiality (or non-spatiality), its viscerality or cognitiveness, its involvement with conscious

⁸ Another reconstruction of Schwitzgebel's overarching argumentative strategy proceeds as follows. Although the arguments from dumbfounding, dissociation, and variation establish only local forms of introspective pessimism when considered on their own, when taken collectively they provide a good case for a relatively global form of pessimism given that each of the three arguments concerns distinct (albeit, perhaps, overlapping) domains of phenomenology. Thus understood, Schwitzgebel does not need to appeal to a generalization from the “hard cases” to introspection in general. Although this construal provides an alternative route to pessimism, I regard it as less promising than the one outlined in the text—both as a reading of Schwitzgebel's work and as an argument in its own right.

⁹ Following Hohwy (2011), Schwitzgebel (2011) calls this “the argument from uncertainty”.

imagery, thought, proprioception, or whatever, as sharply and infallibly as you can discern the shape, texture and color of your desk? (Or the difference between 3 and 27?) I cannot, of course, force a particular answer to these questions. I can only invite you to share my intuitive sense of uncertainty. (Schwitzgebel 2008, p. 251)

This argument does not appeal to independent evidence in order to motivate pessimism. Rather, it appeals to first-person considerations: introspection *itself* seems to suggest that there are aspects of our own conscious experience that elude our grasp. As Schwitzgebel puts it, “it’s not just language that fails us—most of us?—when we confront such questions [...] but introspection itself. [...] in the case of emotion, the very phenomenology itself—the qualitative character of our consciousness—is not entirely evident” (Schwitzgebel 2008, pp. 249–250).

Before examining the force of this argument, let us first consider what kind of pessimism it aims to establish. Does the above passage call into question our capacity to accurately *categorize* our emotional phenomenology, or is the claim rather that we lack even the capacity to *discriminate* our emotional experiences from one another and from the rest of our phenomenal states? Although Schwitzgebel’s concern seems to include questions of discriminative access—after all, the passage begins by asking if we can even tell whether or not we have any emotional phenomenology—I take his worries to centre on our capacity to accurately categorize our emotional phenomenology. As I read him, Schwitzgebel’s questions focus on our ability to determine how our emotional experience is structured, both internally and in terms of its relations to phenomenal states of other kinds.

I think that the questions Schwitzgebel raises *are* difficult to answer. However, it is not clear to me that this fact provides quite as much support for introspective pessimism as Schwitzgebel thinks it does. Lying behind the dumbfounding strategy is the assumption that the questions being posed have determinate answers—that they are appropriate questions to ask. However, I suspect that in an important

range of cases this assumption may be unjustified. With respect to the phenomenology of emotion it is natural to assume that the boundaries between the phenomenal states associated with emotion are as clean and sharp as the boundaries between our standard ways of categorizing emotional states. We regard boredom, anxiety, and irritation as distinct emotional states, and we also regard each of these states as associated with distinctive forms of phenomenology. On the basis of these two thoughts we assume that the phenomenal states associated with these categories can themselves be cleanly distinguished from one another. Thus, when one finds oneself at a loss to know whether one is in *the* phenomenal state associated with boredom, anxiety, or irritation one naturally assumes that the fault lies with one’s introspective capacities. But perhaps the mistake was to assume that the phenomenology of emotion can be cleanly demarcated into states that are uniquely associated with either boredom, anxiety, or irritation. Perhaps the phenomenal states associated with these emotional states overlap and interpenetrate each other. If this were the case, then although there might be certain contexts in which one’s emotional phenomenology is purely that of (say) boredom, there may also be other contexts in which one’s emotional phenomenology involves a complex mix of the phenomenal states associated with boredom, anxiety and irritation. And if one were in a context like this, one might be at something of a loss to know just how to categorize one’s emotional state. The only categories that might come to mind would be those associated with the folk psychology of emotion—*<boredom>*, *<anger>* and *<irritation>*—but these categories might fail to cut the phenomenology of emotion at its joints. In other words, emotional phenomenology may pose a particular introspective challenge not because introspection does a poor job of acquainting us with emotional phenomenology, but because the structure of the phenomenology of emotion fails to map onto the structure of our folk categories of emotions in a straightforward manner.

Other versions of the argument from dumbfounding raise a different set of challenges

for introspective optimism. Consider the question of introspective access to visual imagery. Schwitzgebel asks his readers to form a visual image of the front of his or her house, and to then consider the following questions:

How much of the scene are you able vividly to visualize at once? Can you keep the image of your chimney vividly in mind at the same time you vividly imagine (or “image”) your front door? Or does the image of your chimney fade as your attention shifts to the door? If there is a focal part of your image, how much detail does it have? How stable is it? Suppose that you are not able to image the entire front of your house with equal clarity at once, does your image gradually fade away towards the periphery, or does it do so abruptly? Is there any imagery at all outside the immediate region of focus? If the image fades gradually away toward the periphery, does one lose colours before shapes? Do the peripheral elements of the image have color at all before you think to assign color to them? Do any parts of the image? If some parts of the image have indeterminate colour before a colour is assigned, how is that indeterminacy experienced—as grey?—or is it not experienced at all? If images fade from the centre and it is not a matter of the color fading, what exactly are the half-faded images like? (Schwitzgebel 2002, pp. 38–39)

I think that this line of questioning poses one of the most significant challenges to optimism. Further, it is doubtful whether this challenge can be resisted in the way that the previous version of the dumbfounding challenge can, for these questions don’t seem to rest on any problematic assumptions. Schwitzgebel isn’t assuming that visual imagery must be pictorial in nature, or that it will always be fully detailed and determinate. Rather, one issue that he explicitly puts on the table is whether the phenomenology of visual imagery can be purely “generic” or “gisty”, or whether it must instead always be specific in some way or another.

But perhaps the dumbfounding challenge can be met in another way. As Jakob Hohwy (2011) has noted, one striking feature of visual imagery is its instability:

In the absence of specific goal parameters for simulations there will be much phenomenal variability because in such conditions subjects must themselves make up the purposes for which they imagine things, or engage in ‘simple’ free-wheeling imagery. For example, there is an indefinite number of purposes for which you can imagine the front of your house (walking up to it, standing close by, assessing its shape, its prettiness, flying around it, how the postman sees it, smelling it, repairing it, buying it, selling it etc), each of these purposes will constrain the imagery, and thus the introspected phenomenology, in different ways. This means that subjects probably do have *variable* phenomenology, and introspectively report so reliably. (2011, p. 279)

Hohwy’s comments are intended to explain the variability in the introspective reports that individuals give, but they also bear on the dumbfounding argument. Perhaps we are not sure how best to describe the phenomenology of imagery because it is so variable. Imagery experiences cannot be pinned down, but are constantly shifting in response to our own imagistic activity. Precisely how much of the scene we vividly visualize “all at once” depends on the goals that constrain the act of visualization. And, as Hohwy suggests, when we have no such goals our imagery may end up “freewheeling”, such that we move from one state to another. Hohwy grounds his analysis in a predictive-coding account of cognition, but his fundamental point is independent of that theoretical framework and should be fairly uncontentious: imagery surely *is* more labile than perceptual experience or bodily sensation. No wonder, then, that its phenomenal structure is that much more difficult to articulate.

I have suggested that the optimist has the resources to meet (or at least “problematize”)

two of the leading versions of the dumbfounding argument. But suppose that my responses are found wanting, and that the pessimist is able to show that our introspective access to both emotional and imagery phenomenology is insecure and impoverished. Even so, there would be a further question as to how such a finding would motivate *global* pessimism. It is certainly true that questions about the nature of certain kinds of experiences (e.g., emotional and imagery experiences) strike us as difficult to answer and may leave us flummoxed, but it is equally true that many introspective questions strike us as easily answered. Indeed, as the quotation from Gertler makes vivid, many of our introspective judgments appear to be accompanied by a sense of epistemic certainty. Why should we generalize from the first set of cases rather than the second? Without an account of why certain introspective questions leave us dumbfounded it is difficult to see why pessimism about a particular range of introspective questions should undermine the epistemic credentials of introspection more generally. So even if the threat posed by dumbfounding arguments were able to establish a form of local pessimism, that threat would appear to be easily quarantined.

3.2 Dissociation arguments

A very different case for introspective pessimism is provided by what I call dissociation arguments. Such arguments appeal to a lack of congruence between a subject's introspective judgments and their capacity to produce reliable first-order judgments—that is, judgments about the objects and properties in their environment. An example of this kind of argument is provided by Schwitzgebel's treatment of the so-called "grand illusion" (Noë 2002). Most people, Schwitzgebel claims, hold that a broad swathe of their environment—perhaps thirty or more degrees—is clearly presented within visual experience with its "shapes, colours, textures all sharply defined". Schwitzgebel argues that we have good reason to regard such claims as false. In making the case for this claim, he appeals to an example first popularized by Dennett (1991):

Draw a card from a normal deck without looking at it. Keeping your eyes fixed on some point in front of you, hold the card at arm's length just beyond your field of view. Without moving your eyes, slowly rotate the card toward the centre of your visual field. How close to the centre must you bring it before you can determine the colour of the card, its suit, and its value? Most people are quite surprised at the result of this little experiment. They substantially overestimate their visual acuity outside the central, foveal region. When they can't make out whether it's a Jack or a Queen though the card is nearly (but only nearly) dead centre, they laugh, they're astounded, dismayed. (Schwitzgebel 2008, pp. 254–255)

How might we explain the dissociation between subjects' introspective judgments and their first-order judgments? One explanation is that the subjects' introspective beliefs are false, and that people wrongly take themselves to have detailed visual phenomenology outside of the focus of attention. This is the explanation that Schwitzgebel endorses. But as Schwitzgebel (2008, p. 255) himself notes, it is possible to explain this dissociation by supposing that individuals are wrong not about which phenomenal states they are in but only about the *origin* of that state. With respect to the card trick example, the proposal is that subjects do indeed have detailed visual phenomenology outside of the origin of attention, but that this phenomenology derives from background expectation rather than environmental input—that is, it is "illusory".

Schwitzgebel's account of the dissociation may have more intuitive appeal than the account I have just outlined, but it is not clear how the data furnished by the dissociation argument allows us to choose between them. However, reasons to favour Schwitzgebel's account can be gleaned noting that the judgment on which we have focused—"thirty or more degrees of my visual field presents itself to me clearly in experience with its shapes, colours, textures all sharply defined"—is available to

introspection only indirectly. This judgment is not the direct reflection of any one introspective act, but is a belief about the nature of one's visual experience that one forms by tracking one's introspective capacities over time. Call such judgments *indirectly introspective*. Indirectly introspective judgments can be contrasted with *directly introspective* judgments—that is, judgments of the kind that one makes in the very context of the card trick experiment, such as “I am now experiencing the shape, colour, and texture of this card (which is presented to me slightly off centre) in sharp detail”. We can now see that although there is a dissociation between the first-order judgments that subjects make and their *indirect* introspective judgments, there is no such dissociation between their first-order judgments and their *direct* introspective judgments. Subjects in the card-trick experiment *don't* report experiencing the shape, colour, and texture of cards that are presented slightly off centre to them “in sharp detail”—rather, they claim to lack sharp and detailed experiences of such objects. Direct introspective judgments clearly have more warrant than indirect judgments, and thus there is good reason to prefer Schwitzgebel's explanation of the dissociation over the alternative account.

But although we have found reasons to support Schwitzgebel's analysis of the dissociation, we have seen that these very reasons undermine his pessimistic attitude to introspection in general, for the evidence in favour of Schwitzgebel's account involves an appeal to introspection. In other words, the pressure that the dissociation argument puts on indirect introspective judgments assumes that direct introspective judgments are trustworthy. The card trick case does indeed cast doubt on the epistemic security of our *background* beliefs about our own visual experience, but there is no reason to extend such doubts to include our direct introspective judgments; and it is surely direct introspective judgments that are at the heart of debates about the trustworthiness of introspection. (Indeed, indirect introspection judgments are not really a genuine form of introspection at all.)

Let us turn now to the second step of the dissociation argument: the inference from local pessimism to general pessimism. Suppose that we were to find a dissociation between a certain range of introspective judgments and the subject's capacity to make the corresponding first-order judgments. Suppose, furthermore, that one could show that this dissociation is best explained by assuming that the introspective judgments in question were false. Would one have any reason to think that introspection *in general* ought to be regarded with suspicion? Not as far as I can see. It seems to me that our faith in the robustness of introspective access to domains in which such dissociations are not to be found ought to remain completely untroubled by such a finding. In fact, one might even argue that coherence between first-order judgments and (direct) introspective judgments would provide evidence in favour of introspective optimism. If dissociations between a person's introspective capacities and their first-order capacities can *disconfirm* their introspective judgments (as the dissociation argument assumes), then *associations* between a person's introspective judgments and their first-order capacities ought to *confirm* them (Bayne & Spener 2010). In other words, the fact that a person's introspective judgments cohere with their capacity to produce reliable reports of their environment ought to provide us with positive reason to trust those judgments.¹⁰ And a great number of our introspective reports clearly *do* cohere with our first-order capacities. Although there are cases in which such coherence fails to obtain—for example, Schwitzgebel (2011, Ch. 3) provides a plausible case for the claim that introspective reports of visual imagery are only weakly correlated with the kinds of first-order cognitive capacities that one would expect visual imagery to subserve—such cases

¹⁰ This argument is closely related to an argument presented by Spener (2013) in defence of the idea that we can provide principled reasons for trusting introspection in certain contexts. Spener argues that certain everyday abilities, such as adjusting a pair of binoculars or ordering food in a restaurant, are introspection-reliant—that is, their successful execution requires that the subject have accurate introspective judgments. I find Spener's argument plausible, but, as Schwitzgebel (2013) notes, it is something of an open question just how many of our everyday abilities are reliant on introspection. At any rate, the argument I have given here makes no appeal to that notion.

are striking precisely because they stand out against the backdrop of coherence that characterizes the relationship between our normal introspective reports and our first-order perceptual capacities.

3.3 Arguments from introspective variation

Perhaps the strongest case for introspective pessimism derives from the phenomenon of introspective variation. Such arguments have as their starting point a disagreement about how best to describe some aspect of phenomenology. Pessimists then argue that the best explanation for the introspective dispute is that at least one of the two groups is mistaken about its own phenomenology, and thus that introspective access to the relevant phenomenal domain is insecure: despite their best efforts, at least one of the two parties to the dispute is wrong about its own phenomenology.

Schwitzgebel (2008) examines a number of arguments from introspective variation, but his central case study concerns a debate about the nature of conscious thought—the so-called “cognitive phenomenology” debate (Bayne & Montague 2011; Smithies 2013).¹¹ On one side of this dispute are those who deny that thought has a distinctive phenomenal character. Those who hold this view typically allow that conscious thought has a phenomenology of some kind, but they regard that phenomenology as purely sensory—as limited to the phenomenology of inner speech, visual imagery, and so on. We might call this the conservative account of conscious thought, for it treats phenomenal consciousness as limited to sensory aspects of the mind. On the other side of this dispute are those who adopt a liberal conception of conscious thought, according to which conscious thought is characterized by a range of non-sensory phenomenal states—states of “cognitive phenomenology”. It is tempting to conclude that at least one of these two sides is guilty of a

fairly radical introspective error: introspection either fails to inform conservatives of a wide range of phenomenal states that they enjoy on a regular basis, or it misleads liberals into thinking that they enjoy a wide range of phenomenal states that they don’t enjoy. Either way, introspection would seem to be untrustworthy with respect to what is clearly a central feature of phenomenology.¹²

But before we follow Schwitzgebel (and many others) in embracing this conclusion, we need to consider alternative explanations of the cognitive phenomenology dispute. One possible explanation appeals to group differences in phenomenology. Perhaps the descriptions of conscious thought that both liberals and conservatives give are right when applied to themselves but wrong when taken to describe conscious thought in general. In other words, perhaps both parties to the dispute are guilty of over-hasty generalization rather than introspective error.

Although an appeal to group differences might explain (away) some instances of introspective disagreement, it is unlikely to provide the best explanation of the cognitive phenomenology dispute. First, this account requires a degree of variation in phenomenology for which there are few (if any) parallels. This is not to say that phenomenal differences between individuals might not run much deeper than common-sense tends to assume—consider, for example, the phenomenal differences that characterize synaesthesia (Robertson & Sagiv 2005)—but the kinds of phenomenal differences that we already recognize are nowhere near as fundamental as the kinds of differences required by this explanation of the cognitive phenomenology debate, for liberals claim that conscious thought is characterized by a *sui generis* kind of phenomenology—a kind that is non-sensory in nature. Second, the group difference proposal predicts that there are cognitive and behavioural differences between the advocates of cognitive phenomenology and their detractors that simply don’t appear to obtain. In sum, it seems

¹¹ Other examples of recent introspective disagreement concern the apparent shape of the objects of visual experience (e.g., Siewert 2007; Schwitzgebel 2011, Ch. 2), the existence of high-level perceptual phenomenology (Siegel 2006; Bayne 2009), and the satisfaction conditions of the phenomenology of free will (e.g., Horgan 2012; Nahmias et al. 2004).

¹² The conservative view is also known as the “restrictive” (Prinz 2011) or “exclusivist” (Siewert 2011) view, while the liberal view is also known as the “expansionist” (Prinz 2011) or “inclusivist” (Siewert 2011) view.

highly unlikely that the debate about the existence of cognitive phenomenology can be explained by supposing that what it is like to be a liberal is different from what it is like to be a conservative.

But there is another deflationary explanation of the debate about cognitive phenomenology that cannot be so easily dismissed. Perhaps the parties to the debate are operating with very different conceptions of what it would take for thought to possess distinctive phenomenal character, and are thus talking passed each other (Bayne unpublished). On this proposal, liberals are willing to extend the notion of phenomenal consciousness beyond its sensory paradigms in a way that conservatives are not. If this account is right, then the dispute surrounding the existence of cognitive phenomenology is largely verbal. Rather than disagreeing about what introspection reveals, the two sides instead disagree about how the term “phenomenal consciousness” and its cognates ought to be employed.

Why take this proposal seriously? Well, one argument for it is that it would provide a good explanation of why there is such widespread disagreement about the nature of conscious thought—the very terms in which the debate are couched are contested. It is also widely acknowledged that there are different notions of “what it’s likeness” (see e.g., Tye 1996; Flanagan 1992; Georgalis 2005). Although this proposal clearly needs much more defence and development than I can give it here, I think it is not unreasonable to suppose that the disagreement surrounding the existence of cognitive phenomenology might turn out to be largely verbal. At any rate, it seems to me that this account provides at least as good an explanation of the dispute as that which is required by the argument from variation.¹³

¹³ Of course, the pessimist might argue that, even if the disagreement surrounding the phenomenology of thought is fundamentally semantic, it doesn’t follow that the optimist is off the hook. After all, using introspection to ground a science of consciousness doesn’t merely require the reliability of introspection, it also requires intersubjective agreement about its deliverances. And—the pessimist might continue—dispute about how to apply the term “phenomenal consciousness” and its cognates threatens to undermine intersubjective disagreement about what introspection reveals just as surely as introspective unreliability does. This is a

There are, of course, other introspective disagreements besides that concerning the phenomenology of thought, and nothing that I have said here goes any way towards showing that they too succumb to a deflationary analysis. Indeed, I suspect that certain introspective disputes—for example, those relating to the richness of visual imagery—may well be best explained by appeal to introspective error. But even if the argument from variation succeeds in establishing a local form of pessimism, it seems to me there is little reason to think that this pessimism generalizes. Indeed, domains that feature disagreement in introspective reports stand out against a general backdrop of introspective agreement. Arguably many domains of consciousness exhibit a great deal of uniformity with respect to introspective reports once individual differences and verbal disputes are taken into account. Now, although inter-subjective agreement doesn’t entail that the individuals in question are right, it does need to be explained, and it seems plausible to suppose that leading explanations of inter-subjective agreement will appeal to the trustworthiness of introspection.

4 Elusive phenomenology

In the previous section I argued that there are good reasons for resisting Schwitzgebel’s case for global pessimism. However, we also saw that there are domains in which our introspective access to phenomenal consciousness is rather less secure than we might have pre-theoretically assumed. In other words, we saw that there is reason to think that certain kinds of phenomenal states are introspectively elusive. In this final section I want to sketch an account of why certain types of phenomenal states are elusive and others are not.

Let me begin by distinguishing the form of phenomenal elusiveness with which I am concerned from another notion of phenomenal elusiveness that I want to set to one side. In a recent paper, Kriegel uses the label “elusive phe-

fair challenge, but in my view the prospects for securing a solution to the cognitive phenomenology dispute, should it turn out to be fundamentally semantic, are quite high. For further discussion of phenomenal disputes and introspective disagreement see Hohwy (2011) and Siewert (2007).

nomenology” to describe phenomenal states “whose very essence requires the absence of introspective attention” (2013, p. 1171). Among the examples that he gives of elusive phenomenology are the phenomenal states that occur at the fringes or margins of consciousness. As Kriegel notes, such states are elusive in that any attempt to make them the object of attentive introspection would change their nature. Although Kriegel’s notion of elusiveness is closely related to the one that I employ here, the two notions are not identical. (One way of seeing that they are distinct is that Kriegel’s elusiveness is primarily a matter of the phenomenology, whereas my elusiveness is a matter of one’s introspective access to the phenomenology.) Unlike Kriegel, I am interested in a type of elusiveness that is independent of attention. Consider again visual imagery. Although particular instances of visual imagery might be elusive in Kriegel’s sense because they happen to occupy the margins of consciousness, I am interested here in the fact that visual imagery *as such* appears to be introspectively elusive.¹⁴

Why might certain types of phenomenal states be elusive in a way that other types of phenomenal states are not? Broadly speaking, there are two places in which we might look for an answer to this question. On the one hand we might appeal to intrinsic features of the phenomenal states themselves. Perhaps there is something inherent in the very nature of certain kinds of phenomenal states that renders them relatively opaque to introspective access. Another possibility is that the elusiveness of certain types of phenomenal states has nothing to do with their intrinsic nature but instead reflects the structure of our introspective capacities. Just as our perceptual system is geared toward the identification of certain kinds of environmental states rather than others, so too it is possible that our introspective system is geared towards the identification of certain kinds of phenomenal states rather than others. On this view, the fact that our introspective access to some types of phenomenology is more secure

than it is to others tells us more about introspection than it tells us about phenomenal consciousness (as it were).

It is, I think, premature to speculate which of these two accounts might be the more plausible; indeed, it is possible that a full explanation of elusiveness will have to draw on both ideas. But rather than pursue that thought, I want instead to sketch one way in which the structural features of introspection might go some way towards explaining why certain types of introspective judgments are more secure than others. The account in question appeals to a distinction between two kinds of introspective judgments: *scaffolded judgments* and *freestanding judgments* (Bayne & Spener 2010). The distinction is perhaps best grasped by means of examples. Contrast an introspective judgment that is directed towards one’s visual experience of looking at a red tomato with an introspective judgment that is directed towards an experience of visual imagery involving a red tomato in front of one. In the former case, there is a perceptual judgment that one is disposed to make (“There is a red tomato in front of me”) whose content corresponds (broadly speaking) to the content of one’s introspective judgment (“I have an experience as of a red tomato in front of me”). In the latter case, however, there is no such first-order judgment that one is disposed to make whose content might correspond to the content of one’s introspective judgment. In a sense, the former judgment is “scaffolded” by a perceptual disposition in a way that the latter judgment is not.

I suggest that scaffolded judgments are typically more secure than freestanding ones precisely because they are scaffolded. At the very least, it is a striking fact that many of the most epistemically insecure introspective judgments appear to be freestanding. Further, one can tell an attractive story about *why* introspective scaffolding might contribute to epistemic security. In making scaffolded judgments, the subject is able to both exploit the resources that it has for making freestanding judgments and calibrate those resources by drawing on its dispositions to make first-order

¹⁴ Phenomenal domains that are at least somewhat elusive include the phenomenology of agency (Metzinger 2006; Bayne 2008; Horgan et al. 2006) and high-level perceptual phenomenology (Siegel 2006; Bayne 2009).

perceptual judgments.¹⁵ Just as beliefs that are derived from multiple (independent) sources are typically more secure than beliefs derived from just a single source, so too scaffolded introspective judgments might typically be more secure than their freestanding brethren.

5 Conclusion

This paper provides a partial response to Schwitzgebel's case for global pessimism with respect to introspection. I began by outlining two arguments for optimism; the first argument turned on an appeal to the phenomenology of introspection, while the second drew on a conceptual connection between the notions of introspective access and phenomenality. Neither argument comes close to being decisive, but taken together they provide some explanation for—and justification of—the widespread appeal of optimism. I then turned to a detailed examination of Schwitzgebel's case for pessimism, arguing that although his arguments go some way towards justifying local pessimism (particularly with respect to imagery), there is little reason to generalize that pessimistic attitude to introspection more generally.

But perhaps the central lesson of this paper is that the epistemic landscape of introspection is far from flat but contains peaks of security alongside troughs of insecurity. Rather than asking whether or not introspective access to the phenomenal character of consciousness is trustworthy, we should perhaps focus on the task of identifying how secure our introspective access to various kinds of phenomenal states is, and why our access to some kinds of phenomenal states appears to be more secure than our access to other kinds of phenomenal states. I have suggested that the notion of introspective scaffolding might play a role in answering this second question, but that that proposal is at

best only a very small part of a much larger account of introspective insecurity. There is certainly a lot more work to be done before we have a good grip on the epistemic structure of introspection.

Acknowledgements

An earlier version of this paper was presented at a graduate workshop on perception at the City University of New York, and I am grateful to the members of the audience on that occasion for their comments. I also gratefully acknowledge the support of European Research Council Grant *The Architecture of Consciousness* (R115798).

¹⁵ An influential account of introspection holds that introspection involves a semantic ascent routine in which one redeploys rather than represents one's introspective target (Byrne 2005; Evans 1982; Fernández 2013). Although I am not endorsing this account of introspection in general (or indeed of introspective access to perceptual phenomenology in particular), I am suggesting that such procedures might be implicated in introspective access to certain kinds of phenomenal states.

References

- Bayne, T. (2008). The phenomenology of agency. *Philosophy Compass*, 3 (1), 182-202. [10.1111/j.1747-9991.2007.00122.x](https://doi.org/10.1111/j.1747-9991.2007.00122.x)
- (2009). Perception and the reach of phenomenal content. *Philosophical Quarterly*, 59 (236), 385-404. [10.1111/j.1467-9213.2009.631.x](https://doi.org/10.1111/j.1467-9213.2009.631.x)
- Bayne, T. (unpublished). *The puzzle of cognitive phenomenology*.
- Bayne, T. & Montague, M. (2011). Cognitive phenomenology: An introduction. In T. Bayne & M. Montague (Eds.) *Cognitive phenomenology* (pp. 1-34). Oxford, UK: Oxford University.
- Bayne, T. & Spener, M. (2010). Introspective humility. *Philosophical Issues*, 20 (1), 1-22. [10.1111/j.1533-6077.2010.00176.x](https://doi.org/10.1111/j.1533-6077.2010.00176.x)
- Byrne, A. (2005). Introspection. *Philosophical Topics*, 33 (1), 79-104. [10.5840/philtopics20053312](https://doi.org/10.5840/philtopics20053312)
- Chalmers, D. J. (2003). The content and epistemology of phenomenal belief. In Q. Smith & A. Jokic (Eds.) *Consciousness: New philosophical perspectives* (pp. 220-272). Oxford, UK: Oxford University Press.
- Church, J. (1997). Fallacies or analyses? In N. Block, O. Flanagan & G. Güzeldere (Eds.) *The nature of consciousness* (pp. 425-426). Cambridge, MA: MIT Press.
- Clark, A. (2000). A case where access implies qualia? *Analysis*, 60 (1), 30-38. [10.1093/analys/60.1.30](https://doi.org/10.1093/analys/60.1.30)
- Dennett, D. (1991). *Consciousness explained*. Boston, MA: Little, Brown.
- Evans, G. (1982). *The varieties of reference*. Oxford, UK: Oxford University Press.
- Fernández, J. (2013). *Transparent minds: A study of self-knowledge*.
- Flanagan, O. (1992). *Consciousness reconsidered*. Cambridge, MA: MIT Press.
- Georgalis, N. (2005). *The primacy of the subjective*. Cambridge, MA: MIT Press.
- Gertler, B. (2012). Renewed acquaintance. In D. Smithies & D. Stoljar (Eds.) *Introspection and consciousness* (pp. 93-128). Oxford, UK: Oxford University Press.
- Goldman, A. (1997). Science, publicity, and consciousness. *Philosophy of Science*, 64 (4), 525-545. [10.1086/392570](https://doi.org/10.1086/392570)
- (2004). Epistemology and the evidential status of introspective reports. *Journal of Consciousness Studies*, 11 (7-8), 1-16.
- Haybron, D. (2007). Do we know how happy we are? On some limits of affective introspection and recall. *Noûs*, 41 (3), 394-428. [10.1111/j.1468-0068.2007.00653.x](https://doi.org/10.1111/j.1468-0068.2007.00653.x)
- Hohwy, J. (2011). Phenomenal variability and introspective reliability. *Mind & Language*, 26 (3), 261-286. [10.1111/j.1468-0017.2011.01418.x](https://doi.org/10.1111/j.1468-0017.2011.01418.x)
- Horgan, T. (2012). Introspection about phenomenal consciousness: Running the gamut from infallibility to impotence. In D. Smithies & D. Stoljar (Eds.) *Introspection and consciousness* (pp. 405-422). Oxford, UK: Oxford University Press.
- Horgan, T., Tienson, J. & Graham, G. (2006). Inner-world scepticism and the self-presentational nature of phenomenal consciousness. In U. Kriegel & K. Williford (Eds.) *Self-representational approaches to consciousness* (pp. 41-61). Cambridge, MA: MIT Press.
- Horgan, T. & Kriegel, U. (2007). What is phenomenal consciousness that we may know it so well? *Philosophical Issues*, 17 (1), 123-144. [10.1111/j.1533-6077.2007.00126.x](https://doi.org/10.1111/j.1533-6077.2007.00126.x)
- Kriegel, U. (2013). A hesitant defense of introspection. *Philosophical Studies*, 165 (3), 1165-1176. [10.1007/s11098-013-0148-0](https://doi.org/10.1007/s11098-013-0148-0)
- Metzinger, T. (2006). Conscious volition and mental representation: Towards a more fine-grained analysis. In N. Sebanz & W. Prinz (Eds.) *Disorders of volition* (pp. 19-48). Cambridge, MA: MIT Press.
- Nahmias, E., Morris, S., Nadelhoffer, T. & Turner, J. (2004). The phenomenology of free will. *Journal of Consciousness Studies*, 11 (7-8), 162-179.
- Noë, A. (2002). Is the visual world a grand illusion? *Journal of Consciousness Studies*, 9 (5-6), 1-12.
- Piccinini, G. (2003). First-person data, publicity, and self-measurement. *Philosophers' Imprint*, 9 (9), 1-16.
- (2011). Scientific methods must be public, and Descriptive Experience Sampling qualifies. *Journal of Consciousness Studies*, 18 (1), 102-117.
- Prinz, J. (2011). The sensory basis of cognitive phenomenology. In T. Bayne & M. Montague (Eds.) *Cognitive phenomenology* (pp. 174-196). Oxford, UK: Oxford University Press.
- Robertson, L. C. & Sagiv, N. (2005). *Synesthesia: Perspectives from cognitive neuroscience*. Oxford, UK: Oxford University Press.
- Schwitzgebel, E. (2002). How well do we know our own conscious experience? The case of visual imagery. *Journal of Consciousness Studies*, 9 (5-6), 35-53.
- (2008). The unreliability of naïve introspection. *Philosophical Review*, 117 (2), 245-273. [10.1215/00318108-2007-037](https://doi.org/10.1215/00318108-2007-037)
- (2011). *Perplexities of consciousness*. Cambridge, MA: MIT Press.

- (2013). Reply to Kriegel, Smithies, and Spener. *Philosophical Studies*, 165 (3), 1195-1206. [10.1007/s11098-013-0152-4](https://doi.org/10.1007/s11098-013-0152-4)
- Siegel, S. (2006). Which properties are represented in perception? In T. S. Gendler & J. Hawthorne (Eds.) *Perceptual experience* (pp. 481-503). Oxford, UK: Oxford University Press.
- Siewert, C. (2007). Who's afraid of phenomenological disputes? *Southern Journal of Philosophy*, 45 (S1), 1-21. [10.1111/j.2041-6962.2007.tb00107.x](https://doi.org/10.1111/j.2041-6962.2007.tb00107.x)
- (2011). Phenomenal thought. In T. Bayne & M. Montague (Eds.) *Cognitive phenomenology* (pp. 236-267). Oxford, UK: Oxford University Press.
- Smithies, D. (2012). A simple theory of introspection. In D. Smithies & D. Stoljar (Eds.) *Introspection and consciousness* (pp. 259-293). Oxford, UK: Oxford University Press.
- (2013). The nature of cognitive phenomenology. *Philosophy Compass*, 8 (8), 744-754. [10.1111/phc3.12053](https://doi.org/10.1111/phc3.12053)
- Spener, M. (2011a). Using first-person data about consciousness. *Journal of Consciousness Studies*, 18 (1), 165-179.
- (2011b). Disagreement about cognitive phenomenology. In T. Bayne & M. Montague (Eds.) *Cognitive phenomenology* (pp. 268-284). Oxford, UK: Oxford University Press.
- (2013). Moderate scepticism about introspection. *Philosophical Studies*, 165, 1187-1194.
- Spener, M. (Unpublished). *Calibrating introspection*.
- Tye, M. (1996). The function of consciousness. *Noûs*, 30 (3), 287-305. [10.2307/2216271](https://doi.org/10.2307/2216271)