

Phenomenal Contrast: A Critique*

Ole Koksvik

December 2011

1 Introduction

Phenomenal experience is important to us. Each of us has a multitude of experiences each day: gustatory, olfactory and visual experiences; moods, emotions, and bodily sensations. Each has a phenomenal character; there is something it is like to have it. There is also a global phenomenal character, something it is like overall to be a particular conscious being at a particular time. We care greatly both about the characters of individual experiences, and about the character of our overall phenomenal experience.

Phenomenal experience is also important in our attempts to understand the world and our place in it. In psychology, subjects' reports of phenomenal experiences are used as data both in experimental and clinical settings.

Phenomenal experience is also important in philosophy. Some argue, for instance, that the phenomenal character of experience allows perceptual beliefs to be immediately justified (Pryor 2000, 2004); others that it seeming to someone that such-and-such is the case (partly a phenomenal notion) is the *only* thing that can justify a belief, and that 'seemings' can justify beliefs whether they be sensory, mnemonic, intellectual or introspective (Huemer 2005, 2007). Some claim that phenomenal experience is neces-

*Under review. Please do not cite without permission.

sary for us to have immediate knowledge of the content of our own thoughts (Pitt 2004); others that justification supervenes on phenomenology (Smithies 2006); that phenomenal experience is part of what makes certain beliefs infallible (Chalmers 2003); or that it is what makes the content of thought and speech determinate (Horgan and Graham Forthcoming). And many argue that the character of perceptual experience is closely related to its content (see Chalmers 2004 for an overview).

The problem is that we lack good and authoritative methods for answering questions about the character of phenomenal experience. Widespread stalemate therefore threatens: two parties might agree on what follows from phenomenal experience having a particular character, but disagree about whether or not it does.¹

In recent years, however, an optimistic consensus has begun to arise, according to which use of *phenomenal contrast* can provide answers about the phenomenal character of experience.² However, while I share the yearning for a method that can determine the character of phenomenal experiences, I do not share the optimism. Against the growing consensus I argue that important facts about human mental lives systematically block a large class of uses of phenomenal contrast from achieving their aim, and that these *minimal pair* arguments (as I will dub them) therefore fail, quite generally (§4). Not all uses of phenomenal contrast amount to minimal pair arguments, but many do (§3). If I am right, and if significant objections can be rebutted (§5), the growing optimism is to a large extent unwarranted. At the very least, the case here presented constitutes a significant challenge, to which

¹This well-known example can serve to highlight the problem:

[T]here is always to be remembered that famous session of the Society of Experimental Psychologists in which Titchener, after hot debate with Holt, exclaimed: “You can see that green is neither yellowish nor bluish!” And Holt replied: “On the contrary, it is obvious that a green is that yellow-blue which is just exactly as blue as it is yellow (Boring 1946, 176).

²For explicit discussion see Siegel (2006a), Kriegel (2007) and Kriegel and Horgan (Forthcoming). Others signal their agreement by using such arguments; numerous examples follow.

proponents of minimal pair arguments must respond.

2 The Structure of Minimal Pair Arguments

Any argument in which a significant role is played by the claim that two or more situations differ from each other with respect to phenomenology could be said to employ phenomenal contrast. The class of minimal pair arguments, on the other hand, is quite restricted: it is unified both in aim and method.

The aim of such arguments is to rationally persuade us that a particular mental feature M contributes to the character of experience.³ For instance, pain is widely acknowledged to be a contributor, but other mental features are more controversial. Minimal pair arguments have often been used in an attempt to establish that thinking that p makes a different contribution than thinking that q : that there is *content-specific cognitive phenomenology*.⁴ As we shall see, however, such arguments have been advanced for many other conclusions as well.

As for the *method*, minimal pair arguments proceed by describing a pair of situations in which a person might find herself. There are three desiderata for this description. First, the situations should of course differ from each other with respect to the crucial mental feature M , the one which the proponent

³An intuitive understanding of ‘contributing to experience’ suffices for our purposes. Peacocke (1998) also uses this locution.

⁴As the terms are used here, to say that there is something it is like to see something red—to say that seeing something red has a phenomenal character—is to say that seeing something red makes a difference to the character of the overall experience of the person who is seeing. However, we can distinguish between two different ways such contributions can be determined. On the one hand, the contribution a mental state makes can depend on the *content* of that mental state. Suppose that if I perceive something red, this makes a different contribution to the character of my overall phenomenal experience than perceiving something green does. In that case, perception has *content-specific phenomenology* (and if not, it does not). On the other hand, it may be that *whatever* I visually perceive, the fact that I am having a visual perceptual experience makes a contribution to the character of my overall phenomenal experience: perhaps it contributes a certain ‘visualness’ (Grice 1962/1989). In that case we may say that perception has *attitude-specific phenomenology*. Of course, perception may have, and indeed actually does have, both attitude-specific and content-specific phenomenology.

of the argument seeks to establish is a contributor. Unless the proponent blunders badly this desideratum is always fulfilled, usually by stipulation. Second, the situations should approximate a *truly minimal pair*: situations which differ *only* in **M**. Finally, the description should produce a clear a reaction that what it would be like to be in one of the described situations is different from what it would be like to be in the other: the character of a person's overall phenomenal experience would differ in the two cases.

In the ideal case the audience has no doubt about this: the person's overall phenomenal experience would be different in the two cases. Because the pair is truly minimal, no explanation of this fact can be given in terms of differences in *acknowledged contributors*, mental features which both sides take to be contributors. This gives rise to pressure to acknowledge a new contributor. Since **M** is the only difference between the two situations, **M** must be it.

Actual instances of minimal pair arguments are not ideal. In particular, the second and third desiderata pull in opposite directions, and proponents must seek a compromise between the two. The closer we come to a truly minimal pair, the less certain we are that there really would be a difference in the character of the person's overall experience. The more certain we are that there would be, the further we are from a truly minimal pair, and the less pressure there is, consequently, to conclude that **M** really is responsible for the difference. I claim that no compromise position leaves minimal pair arguments rationally persuasive. But before defending this claim, I want to present a more formal version of the (non-ideal) argument:

- (1) If a person were to find herself in the two situations described, the character of her overall phenomenal experience would be different in the two cases
- (2) The person would also differ with respect to **M**
- (3) That **M** contributes to phenomenal experience is the best explanation for the difference in overall experience

-
- (4) So, we have good reason to believe that **M** contributes to phenomenal experience

Here the tension is between premises (1) and (3): as the plausibility of (1) increases the plausibility of (3) diminishes, and *vice versa*.

3 Examples

In addition to being used to argue that there is content-specific cognitive phenomenology, minimal pair arguments have been used to argue that there is *not* content-specific cognitive phenomenology, that high-level properties (being the face of a parent, agency, membership in natural kinds, subject independence, perspectival connectedness, causation, ...) are represented in perceptual experience, that understanding has phenomenology, and more besides.

Here I discuss four examples from the literature, and show how they fit the structure I have outlined. In so doing I impose a particular interpretation on the texts. In particular, in each case I take the explanandum to be *that there is a difference* in the character of overall phenomenal experience. This is important. To foreshadow the criticism presented in §4, my objection will be that many other equally good explanations are available for this fact.

While other interpretations of the texts are possible, in this section I try to show that it is reasonable to interpret the texts in the way that I do, and in particular that it is reasonable to take the explanandum to be that there is a difference in the character of overall phenomenology. I wish to emphasise that a simple dismissal of that claim must be regarded as an inadequate response to the challenge. If that is not the explanandum, then what is? A definite answer is required, along with an explanation of how this changes the outcome. In §5.4 below I consider four suggestions for other explananda,

and argue that none of them serves the proponent of minimal pair arguments any better. If no better candidate can be found, and if the argument in §4 succeeds, we are forced to accept that minimal pair arguments fail.

COGNITIVE PHENOMENOLOGY. In *The Significance of Consciousness*, Charles Siewert argues that there is content-specific cognitive phenomenology:

[O]n some occasions someone utters a sentence, and you momentarily understand it one way . . . and then are struck by the realization that the speaker meant something else altogether. . . . [O]ne can note a difference in the way it seems to understand it, depending on which way one takes the story. And this is so even if one does not picture anything differently, or picture anything at all, as one interprets it differently.⁵ (Siewert 1998, 278–9)

Here is how I understand Siewert. A person could find herself in two very similar situations, such that in one she understands a recent utterance in one way, and in the other she understands it a different way. The content of her thought (**M**) would then be different in the two situations. What it would be like to be her would also be different; the character of her overall phenomenology would differ (so premise 1 is true). This would be so even if she did not visually imagine different things in the two situations. The best explanation for this is that thinking a thought with one content makes a different contribution to the character of overall phenomenology than thinking

⁵I take the content of experiences and thoughts to be their accuracy conditions, and assume that they have content in this sense; see Siegel (2005/2010) and Siewert (1998, §§6.2 and 8.4). Minimal pair arguments play an important role in Siewert’s argument (see especially his §8.3), but do not exhaust it. Pitt (2004) might also appear to argue in this way for content-specific cognitive phenomenology, but his intention is to draw attention to the type of phenomenology he claims exists (see §5.1). Horgan and Tienson (2002), Horgan and Graham (Forthcoming), Kriegel (2003) and Peacocke (1998) can also be understood to use minimal pair arguments to support the claim that there is content-specific cognitive phenomenology. In contrast, Jacob (1998) argues that there is *not* cognitive phenomenology; Tolhurst (1998) that there is a distinction between a ‘mere desire’ and a ‘felt demand’, and Kriegel (2007) that there is non-sensory phenomenology (see also Kriegel 2003).

a thought with a different content. So, we have good reason to believe that this is so.

RECOGNITIONAL CAPACITIES. In ‘Which Properties are Represented in Perception?’, Susanna Siegel argues that exercising one’s ability to recognise a tree of a certain kind contributes to the character of one’s overall experience:

Suppose you have never seen a pine tree before, and are hired to cut down all the pine trees in a grove containing trees of many different sorts. . . . [Y]our disposition to distinguish the pine trees from the others [gradually] improves. Eventually, you can spot the pine trees immediately. . . . Gaining this recognitional disposition is reflected in a phenomenological difference between the visual experiences you had before and after the recognitional disposition was fully developed. (Siegel 2006b, 491)

Here is how I understand Siegel. A person might find herself in two very similar situations, such that in one but not the other a capacity to immediately recognise pine trees is exercised (M). If she did, the character of her overall phenomenology would differ in the two cases (premise 1). The best explanation for this is that exercising a capacity to immediately recognise pine trees contributes to phenomenal experience (premise 3). So we have good reason to believe that this is so.⁶

PHENOMENOLOGY OF AGENCY. In ‘The Phenomenology of First Person Agency’, Horgan, Tienson and Graham argue as follows:

Suppose that you deliberately perform an action—say, holding up your right hand and closing your fingers into a fist. . . . [The experience of doing that] is certainly not like this: first experiencing an occurrent wish for your right hand to rise and your fingers to move into clenched position, and then passively experiencing your hand and fingers moving in just that way. (Horgan, Tienson, and Graham 2003, 327–8)

Here is how I understand these authors. It is possible for a person to find

⁶Kriegel (2007) argues that ‘being mommy’s face’ is represented in experience; Siegel (2006a) that subject independence, perspectival connectedness and causation are.

herself in two very similar situations, such that in the first but not the second she takes herself to perform certain actions voluntarily (**M**). What it would be like to be her would be different—the characters of her overall phenomenologies would differ (premise 1). The best explanation for this is that taking oneself to perform an action voluntarily contributes to phenomenal experience (premise 3). So we have good reason to believe that this is so.⁷

UNDERSTANDING. In *Mental Reality*, Galen Strawson argues that there is something it is like to understand a sentence:

[D]oes the difference between Jacques (a monoglot Frenchman) and Jack (a monoglot Englishman), as they listen to the news in French, really consist in the Frenchman's having a different experience? ... It is certainly true that Jacques's experience when listening to the news is very different from Jack's. And the difference between the two can be expressed by saying that Jacques, when exposed to the stream of sound, has what one may perfectly well call ... 'an understanding-experience', while Jack does not. (Strawson 1994/2010, 5-6)

Here is how I understand Strawson. It is possible for a person to find herself in two very similar situations, such that in the first she hears a sentence without understanding it, whereas in the second she does understand it (**M**). If she did, what it would be like to be her would be different in the two cases—the character of her overall phenomenology would differ (premise 1). The best explanation for this is that understanding a sentence contributes to overall phenomenology (premise 3). So we have good reason to believe that this is so.⁸

⁷The authors say they wish to 'bring into focus' the aspect of phenomenology they are interested in, perhaps indicating an ostensive use of phenomenal contrast (§5.1), but the passage is also naturally understood as an argument. See also Searle (1983, Chapter 3).

⁸That Strawson uses two different people in his presentation is of course immaterial. I interpret Strawson as taking as his explanandum the existence of a difference in overall phenomenology, but it is possible to understand him otherwise. In particular, the surrounding text might lead one to interpret him as claiming that a certain *kind* of expe-

*

I have given an informal explanation of minimal pair arguments, explained how a tension arises between two desiderata for such arguments, given a more formal version of the argument, and shown how some prominent examples of phenomenal contrast can be understood as minimal pair arguments. I next offer my critique of such arguments.

4 A Minimal Pair Argument

Suppose you encounter someone who doubts what I presume seems obviously true to you: that perceiving something red makes a different contribution to experience than does perceiving something green. How might you rationally persuade her? Suppose you tried the following ‘Cinema Screen Argument’:

Imagine that you are sitting in a comfortable seat in the middle of a dark movie theatre. You are not in pain, you are not hungry or thirsty, and you are sitting still. The screen turns a uniform green. You stare at the screen, concentrating on it. This is the first situation. You relax, closing your eyes. When you open them again, the screen is a uniform red. You stare at it, concentrating on it.

rience cannot be explained by other contributors. I discuss this possibility in §5.4 below. It would be natural to understand Strawson as defending attitude-specific phenomenology here (see n. 5). But Strawson does not seem to think that there is attitude-specific phenomenology *or* content-specific phenomenology of understanding. He seems to think that understanding makes a difference to the character of a person’s overall phenomenal experience, but not necessarily the same one each time:

To talk of understanding-experience, then, is not to commit oneself to the implausible view that there is some single qualitative type of experience that anyone who has understanding-experience must have. It is not to commit oneself to the view that particular qualitative experiences invariably go with understanding particular sentences (1994/2010, 7).

The Jack–Jacques passage is sometimes understood as an argument for content-specific phenomenology (see e.g. [Horgan and Graham Forthcoming](#), n. 6), but it seems clear that this is not the author’s intention.

This is the second situation. Clearly, the character of your overall phenomenal experience would be different in the two situations. The best explanation for this is that perceiving something green makes a different contribution to experience than does perceiving something red.

I do not deny that this argument *seems* convincing. But I claim that it is in fact not rationally persuasive: we should resist the pull and not be convinced. To see why, we need to note two important truths about our mental lives:

Richness

At most times there is a lot going on in our mental lives: several remembered, occurrent and imagined bodily sensations, moods and emotions usually occur at the same time (or near enough), and many thoughts go through a person's head. Our mental lives are *rich* with activity.

Flux

Many (or most) of these goings-on are evanescent; a remembered bodily sensation may last only a fraction of a second. Attention changes around often. Our mental lives are in constant *flux*.

An argument may be rationally persuasive given agreement on one set of facts (actual agreement, or merely for the sake of argument), but not given another. Let us suppose that you and your interlocutor agree that *occurrent*, *remembered* and *imagined* bodily sensations, moods and emotions all contribute to the character of a person's overall phenomenology. You also agree that there is content-specific cognitive phenomenology (thinking that *p* makes a different contribution than thinking that *q*) and that attention makes a difference *to the difference* each of these contributors makes.⁹ Together, these claims constitute your common ground, call it CG₁.¹⁰

⁹One could classify attention as itself a contributor, but it seems more natural to regard it as something which modifies the contribution other contributors make.

¹⁰In §5.5 below I consider what happens if we choose a different common ground.

Since human mental lives are rich with activity, a large number of the contributors acknowledged by CG_1 will obtain at any given time. Given Flux, the set will be different at any other time, barring a massive coincidence. As understood here, the explanandum in a minimal pair argument is the existence of a difference in overall experience. (Different candidate explananda are considered in §5.4 below.) But all these differences in contributors between the two situations provide possible explanations of that datum. Each of them is as good as all the others, and each explains the datum as well as the hypothesis that M is a contributor. That hypothesis is not the best explanation of the datum. Premise (3) is false.

There is no denying that the existence of a difference in the character of overall phenomenal experience really is the explanandum in the Cinema Screen Argument. As I have shown, the examples given above are also reasonably understood in this way. The critique against the Cinema Screen Argument therefore applies with equal force to all those arguments. Richness and Flux, these two important facts about our mental lives, show that none of them can rationally persuade us of their conclusions.

4.1 An invitation to be resisted

The Cinema Screen Argument specifies that the subject is not in pain, not hungry or thirsty, and is sitting still. The purpose of this is of course to exclude pain, hunger, thirst and the feeling of the seat against the subject's body as possible explanations for the difference overall phenomenology.

Similar specifications are common in minimal pair arguments; we saw some examples above. They are rhetorically important. With differences in pain, hunger, thirst and the feeling of the chair against the body excluded, there is still a strong reaction that there would be a difference in overall phenomenology. As we have seen, we can only be rationally persuaded that M is a contributor if all the other potential explanations of the difference in overall phenomenology are out of the picture. Therefore, by listing *some*

exclusions, the proponent must be understood to invite us to think that the list could be extended until all acknowledged contributors are excluded, with no difference in the result.

That invitation should be resisted. The list of contributors acknowledged by CG_1 is long, and the listed exclusions gives us no reason to believe that none of the other contributors differ between the situations. Occurrent bodily sensations (apart from pain, hunger and thirst), occurrent moods and emotions, remembered and imagined bodily sensations, moods and emotions; thought, and attention, all these are in play. Our mental lives are characterised by Richness and Flux, so we have every reason to believe that many of these will differ between the two situations. Each difference makes possible an explanation as good as every other. Again: premise (3) is false. The minimal pair argument fails.

5 Objections

In this section I consider and reply to a number of objections.

5.1 Ostension

One can distinguish two quite different ways of using phenomenal contrast. On the one hand, phenomenal contrast can be used as a way of ‘pointing’; on the other it can be used in an effort to rationally persuade someone. Call these ‘ostensive’ and ‘argumentative’ uses, respectively.

It is natural to understand ostensive uses as targeting a particular audience, namely those who are in at least approximate agreement with the proponent with respect to contributors. Such uses can be understood to have a particular aim, namely to assist the audience in focusing attention on certain features of phenomenology, or to make those features easier to notice and appreciate.

For argumentative uses, the audience is those who disagree with the proponent about what kinds of phenomenology there are, or at least those who are undecided on the matter. The goal is to rationally persuade the audience that a claim about phenomenal experience is true.

The Cinema Screen Argument is most naturally understood as, precisely, an argument, and §4 proceeds from this perspective. But the case could also be used ostensively; perhaps with the aim of acquainting the audience with a difference in phenomenology that the author believes would obtain. And it is in general possible to understand at least some uses of phenomenal contrast as ostensive.¹¹ However, as we have seen (§3), many uses of phenomenal contrast are reasonably understood as presenting arguments properly speaking.

It is possible that my argument should also give one pause about ostensive uses. But nothing here hinges on this. If someone responds to my argument by saying that their use of phenomenal contrast was never intended to rationally persuade, then they do not contest my conclusion. I therefore set aside ostensive uses in what follows.¹²

5.2 Stipulation

One way to attempt to resist the argument in §4 is to claim that there is a difference in overall phenomenology *even though* none of the other contrib-

¹¹It is significant, for example, that many of those who use phenomenal contrasts to argue for content-specific cognitive phenomenology also seem to hold the view that a difference in the content of thought would *metaphysically necessitate* a difference in phenomenology. Perhaps most of the argumentative weight is placed on the considerations presented in favour of that view. See [Horgan and Tienson \(2002\)](#) and [Pitt \(2004\)](#). That claim of necessitation is usually restricted to narrow content, and Siewert seems to hold that most though not all differences in narrow content necessitate a difference in phenomenology, see his (1998), §§8.4–8.5, esp. 287–9. See also [Kriegel and Horgan \(Forthcoming\)](#). For a recent argument that this restriction is unnecessary, see [Farkas \(2008\)](#).

¹² [Name omitted] has suggested that *showing* a person that a certain feature contributes to overall phenomenology might count as a legitimate way to convince her of this. [She or He] suggests we might call this ‘experiential persuasion’, and mentions smacking someone who does not believe in pain phenomenology as an example. However attractive this option may at times seem, I do not count it as a method of rational persuasion.

utors differ between the two situations. On this line, Richness and Flux do not bar one from showing that something is left unexplained by acknowledged contributors. There are two main ways to go here: via stipulation or memory. I deal with these in turn.

Suppose someone said the following:

I do not need to *invite* my opponent to believe that there are no other relevant differences between the first and the second situation; I simply *stipulate* that there isn't. I add to the description that there is no difference in occurrent bodily sensations, in remembered or imagined bodily sensation, mood or emotion, and no difference in thought or attention. This allows me to establish my conclusion.

If indeed there were no differences in *any* of the contributors acknowledged by CG_1 but there clearly would be a difference in the character of overall experience, the cinema screen argument would give non-negligible support to the view that there is content-specific phenomenology of perception. Still, I think we can easily see that this strategy does not get one far.

First, it is very far from clear that the intuition that there would be a difference in the character of a person's overall phenomenology between the two situations actually survives this stipulation. I doubt that it does. When I really try to imagine that there is *no* difference between the two situations other than the difference in the colour projected on the screen, attempting to exclude one by one all the other contributors, the situation soon becomes too alien. My mental life is characterised by Richness and Flux, and in the situation I am trying to imagine these deep facts about me no longer hold. The result is that I fail to have any intuition about overall phenomenology at all.

That the cinema screen argument has a plausible conclusion must not be allowed to muddle our judgement. I agree that there is content-specific

phenomenology of perception; indeed I take myself to know that there is. And it is of course legitimate for anyone pondering the case to *apply* her existing knowledge to it. If I know that there is content-specific phenomenology of perception and that the situations have different perceptual content, I can come to know that the situations differ in phenomenal experience. But this knowledge does not *arise* from the minimal pair case, and such considerations are of no help when the aim is to rationally persuade someone that lacks this knowledge. What is needed is a clear intuition *about the case as described*. That is what I am reporting that I lack.

Secondly, if someone were to report such an intuition, I think we would have little reason to trust it. Our intuitions about what it would be like to be in a certain situation are likely to be heavily influenced by our previous experiences. Since our mental lives are characterised by Richness and Flux, it is overwhelmingly likely that nearly all our experiences are from situations that differ from each other in *many* contributors—only a truly astronomical coincidence could bar this.¹³ Therefore, even if we try to respond to the case as described, it is likely that our experience will overwhelm our attempts, so that our judgements reflect our experience rather than the case. So even if we had intuitions about such cases, we lack good reason to give them weight.¹⁴ So stipulation offers no help.

We can begin to understand why minimal pair arguments *seem* persuasive by noting that, when we consider them, we naturally consider counterfactual conditionals. And this conditional is, of course, *true*:

If a human were to find herself in situations such as those originally described (before extra stipulation was brought in), there

¹³I take it to be possible, but extremely unlikely, to find oneself in two situations differing in only one contributor. This possibility does no damage to my case. It can still be true, of course, that the vast majority of our experiences are from situations which differ from each other in a range of contributors.

¹⁴This does not lead to wholesale scepticism about intuition, since there is a special reason here to suppose that our intuition is simply an upshot of our experience.

would be a difference between the first and the second situation in the character of the person's overall phenomenology

In close possible worlds human beings are very much like us, so their mental lives are substantially characterised by Richness and Flux. Between any two situations there are for them, just as there are for us, differences in many contributors. So the counterfactual is true.

The claim is that we in effect respond to the above counterfactual, even if we try to consider a case where all the acknowledged contributors are excluded by stipulation. And in most cases, of course, one does not even make an effort to exclude all acknowledged contributors, but contents oneself with thinking about a case where a few of them are excluded.

5.3 Memory

One might instead claim that one can *remember* situations that differ only in what is visually perceived and in overall phenomenology. To see why this response fails we need to make a further note about our mental lives:

Poor Identification and Remembrance

A large proportion of the episodes that contribute to the richness of our mental lives are of short duration, and are typically not paid much notice. Partly as a consequence of not paying attention, but also because our introspective abilities are just not that acute, our mental goings-on are often poorly identified. A mental goings-on not identified at the time of occurrence is unlikely to be correctly remembered later. Many of those that are correctly identified at the time will also not be committed to memory. So our mental goings-on are usually poorly remembered later.¹⁵

¹⁵For example: earlier in the day I was irritated after having spent, I thought, far too long in a queue. (This happens a lot.) Various things can now remind me of this emotion, without thereby making me irritated again. When I remember this emotion,

If Poor Identification and Remembrance is true, that there are pairs of situations *about which* one can only remember the difference in, say, perceived colour and overall phenomenology, does not provide one with good evidence for the belief that those *really were* the only differences between the two situations. Given Richness and Flux, we have every reason to believe that we experience no such pairs, and so do not remember any. So this response fails.¹⁶

Finally, it is worth noting that, while we also often fail to notice and remember changes in our overall phenomenology, we are comparatively better at recognising such changes than we are at determining which contributor(s) is (are) responsible for them. This, of course, is precisely what makes arguments which start with a difference in overall phenomenology attractive to begin with.

this changes the character of my overall phenomenology. I may well fail to recognise this episode of remembering for what it is. I might, for instance, easily misidentify the episode as frustration with my current progress in writing. (We might instead say that I do get irritated again when I am reminded of my former irritation, but only very mildly. This difference does not matter for our purposes, and may well be terminological.)

¹⁶Some might worry that Richness, Flux, and Poor Identification and Remembrance are jointly self-undermining. If we are so bad at identifying and remembering our mental goings-on, how can we know that our mental lives are characterised by Richness and Flux? Here is how.

We do not usually focus our attention on our phenomenology, but we can. The Richness of the experience is then revealed to us, and its fluctuating character likewise. Even just within tactile phenomenology, for example, there is a lot going on, and there are rapid changes. In one moment I am strongly aware of the feel of the soles of my feet, a moment later my typing fingertips typing dominate the tactile phenomenology, then a slight back-pain, the feeling of the arm-rest under my arm, the feeling of my body making contact with the chair, and so on. As I shift my attention around, what is in focus comes to make a more significant contribution to the character of my experience.

One can easily see that the same goes for other types of perceptual phenomenology by switching one's attention to, say, auditory phenomenology. If contributors do not contribute to overall phenomenology only when in focal attention—which I assume without argument—we also come to know that there are contributions that we do not recognise at the time, and so do not remember. So many contributors go unnoticed and without being remembered, not just one or a few. So Richness, Flux, and Poor Identification and Remembrance do not undermine each other.

5.4 Thicker explananda

The two above replies claim that there is a difference in overall phenomenology even when no acknowledged contributor differs. A different strategy is to accept that there will always be differences in acknowledged contributors, but to claim that the variance in **M** is nevertheless needed for adequate explanation. What needs explaining, it might be claimed, is not merely *that there is* a difference in overall phenomenology, but something more. I consider four such responses.

The first claims that the *magnitude of difference* in phenomenal experience is greater than can be accounted for by differences in other contributors. To explain the entire difference in experience we need to make reference to variance in **M**. Call this ‘the magnitude reply’.

A second response argues that the difference in overall phenomenology is partly *of a kind* such that the differences in other contributors cannot explain (a certain aspect of) it. There are differences in other contributors, but they leave a particular kind of difference unexplained. Call this ‘the kind reply’.

A third response claims that if we iterate the cinema screen experiment with different colours, the person will be aware that the differences are not the same in each case. Call this ‘the different difference reply’.

Finally, one might claim that if we repeat the original cinema screen setup numerous times while systematically varying other features of the situations, the person will be able to recognise an aspect of the overall difference each time. That aspect is the real explanandum. Call this ‘the persistent difference reply’.

*

The first reply fails because it relies on implausibly denying Poor Identification and Remembrance, and on claiming that we can accurately estimate magnitudes of change in overall phenomenology and magnitudes attributable to various contributors. If Poor Identification and Remembrance is true,

there is no way to support the claim that the amount of difference overall cannot be accounted for by acknowledged contributors, since in each instance there will be contributors that were not correctly identified initially, or that were since forgotten.

The magnitude reply also implausibly relies on abilities to accurately estimate magnitudes of difference in our overall phenomenology, to accurately estimate the magnitude of difference attributable to a certain specific contributor, and to accurately ‘add together’ several instances of the latter. If we cannot do these things, how could we tell that the difference in phenomenal experience is larger than what can be accounted for by differences in acknowledged contributors? The very least that would be required are abilities to classify an overall change as too large to be accounted for by changes in certain other contributors. But there is no reason to think that we have such abilities. So the magnitude reply fails.¹⁷

The kind reply, on the other hand, begs the question. The proponent wishes to show that the acknowledged contributors cannot explain the difference in phenomenology, and that reference to **M** is therefore needed. If she takes as her *datum* that there is phenomenology of a kind that the other contributors cannot possibly explain, then she has assumed what she set out to demonstrate.

It should be clear that taking a different difference as one’s starting point does not help. Our minds are characterised by Richness and Flux, so this datum is easily explained without reference to variance in **M**.

In the persistent difference reply, the alleged datum is an aspect recognised across iterations of the setup. Can the existence of this datum can be adequately supported? To say that we *remember* a common aspect though various pairs of situations is to place a strong bet on our abilities to identify, separate out, and remember detailed aspects of our phenomenology. We have good reasons to doubt that we have such abilities: things often taste, feel,

¹⁷Of course, I need not and do not claim that we have *no* such abilities whatever.

sound and look very different from what we seem to remember.

Can we perhaps *intuit* that an aspect of what changed from one situation to the next would be invariant across iterations of the setup? I could not reply to *that* claim as before, by saying that our judgements are overwhelmed by previous experience and not responsive to the situation described. Our experience is characterised by Richness and Flux, so there is no persistent difference in our experience to ‘overwhelm’ our judgment of the case. So my account predicts that the intuition should not obtain.

Fortunately, there is no intuition of sufficient strength and clarity to create a problem.¹⁸ We can, as before, apply previous knowledge to the cases, and in this way come to know that there would be a persistent difference. But we must distinguish this from an intuitive reaction to the case as described. The latter would support the minimal pair argument, but the former cannot.

Would the proponent have what she needed if we bracketed this concern, and granted the purported datum? Not obviously. In Figure 1, p_1 and p_2 are two different characters of overall phenomenology, a and b are perceptions with different content (a green screen and a red screen, for example), a_1 and a_2 are associated states, and $a_1 \neq a_2$. The proponent wishes to establish is that what is perceived *directly* contributes to the character of a person’s overall phenomenology; it makes such a difference *by itself*. She thus wishes to endorse the picture on the left but reject the picture on the right. But remembering a persistent difference is not enough for this.

If associated states intervene between the content of perception and the character of overall phenomenology, the question of whether the difference in overall phenomenology is attributable to perception in the correct way depends on what those states are. If the intervening states themselves are acknowledged contributors, the conclusion cannot be established. Moreover, it is plausible that the common ground must *rule out* the intervening states as contributors for the conclusion to be established. If it is an open ques-

¹⁸This is even clearer in the case of cognitive phenomenology, discussed further in §6.1.




Figure 1

tion whether the associated state is a contributor, how could the argument establish that it is M which explains the difference? Let a_1 and a_2 be slight emotional reactions. A green colour may be stably associated with one emotional reaction (a_1) and a red colour with another (a_2). On CG_1 emotions are contributors, so the datum could be explained without appeal to M .¹⁹

For the persistent difference reply to work, then, two significant hurdles must be overcome. First, the existence of the datum must be established. Second, an intervening state that is itself a contributor must be ruled out. But there is no reason to think this can be accomplished.²⁰

Could we make the datum even thicker? Could we combine the alleged memory of a persisting difference with a memory of associated states *not* obtaining? Proponents of minimal pair arguments sometimes appear to claim that this can be done. For example, in arguing that there is content-specific cognitive phenomenology, Siewert claims that we can exclude differences in

¹⁹These considerations also apply if the purported datum is a magnitude of phenomenal difference or a difference of a kind that cannot otherwise be accounted for.

²⁰It is crucial to keep in mind the distinction between an argument, capable of rationally persuading through conscious deliberation, and a process that happens at the sub-personal level. If I know that there is content-specific phenomenology of perception I must have come to know that somehow. One possibility is that this comes about through sub-personal abstraction from a large number of experiences. Keeping the door open for this possibility does nothing to undermine the point made here.

imagined or remembered visual perception: “clearly [the difference in overall phenomenology] does not have to do with whether or not it seems to us as it does to image different things. For there may be no difference of that kind” (1998, 276).

This much is true: we can sometimes check for imagination of some particular thing; we can tell whether we are currently visually imagining a pink rhinoceros, for example. But that does not show that we can exclude a contributor wholesale. I think we cannot do that. Suppose we agree, for instance, that the person in the cinema can be certain that she feels no hint of anger. To exclude the picture on the right (in Figure 1) she would have to exclude *any* other emotional reaction, as well any other acknowledged (or even: not ruled out) contributor being stably related to the perceptual states. It is not plausible that this can be done.

5.5 A different common ground

One might complain that CG_1 is implausibly comprehensive, and that a minimal pair argument would do better against a thinner common ground.

The common ground for the Cinema Screen Argument was not arbitrarily chosen, for it mirrors the common ground operative in the cognitive phenomenology debate. What is at issue there, is here part of the common ground, and *vice versa*, otherwise the two are identical. The failure of the Cinema Screen Argument therefore rather directly shows that minimal pair arguments fail in the cognitive phenomenology debate, too.

Secondly, slimming down the common ground will not help. If *no* mental feature is acknowledged as a contributor by both parties, it will, of course, be easy to show that a particular difference in overall phenomenology is not explained by an acknowledged contributor. But that is of no help if the goal is to establish a particular feature as a contributor.

Suppose we aim for something in between, and subtract, say, cognitive phenomenology from CG_1 . Then, not only does the proponent have to con-

vince her audience that there is no difference in occurrent, remembered or imagined bodily sensations, moods, emotions and attention, but still a difference in the character of overall phenomenology. Should she manage that formidable task there would now be (at least) two candidate explanations for the difference in overall phenomenology, and she would have to make the case that her favourite is the most plausible.

5.6 Proving too much?

Besides begging the question, another problem with the kind reply is that the interlocutor will not grant the existence of the explanandum. She will deny that there is a kind of difference in phenomenal experience that the other contributors are not suited to explain. But is there not something wrong with such a person? And am I not committed to denying this?

I do deny that her *failing to be convinced* by the minimal pair argument shows that something is wrong. She is right to not be convinced by that argument. This is of course entirely compatible with the possibility that *the fact that she does not already know* shows there is something wrong with her. That may be.

Some worry that the Cinema Screen Argument is *obviously* correct, so there must be something wrong with my argument. I answer that the argument merely *seems* rationally persuasive, probably because we already know that the conclusion is true, and because applying the knowledge to the case is so quick and easy. But applying knowledge is not the same as that knowledge arising from the case in question. The argument in §4 shows that it does not.

If a medical doctor presses my chest twice and gets an angry “Ow!” the second time, is she not entitled to conclude that I experienced pain? Am I not committed to denying this? She is so entitled, but I am not thus committed. That bodily sensations contribute to the character of experience is part of the common ground which the doctor and I share. The issues I raise do not

arise.

Am I committed to rejecting all manner of apparently good explanations, because there is always the possibility of confounding factors? No. I have not merely presented the possibility of a confounding factor; I have argued that factors will almost certainly be present which can explain everything that needs explaining. If that situation is found elsewhere the conclusion should be the same, even if the arguments initially seem worthwhile. So my argument does not prove too much, nor does it commit me to denying the obvious.

5.7 Lessons Learned

The discussion has concerned a particular example and a particular common ground. But the conclusion is general: minimal pair arguments cannot rationally persuade us.

Because our mental lives are characterised by Richness and Flux, on any reasonable common ground there will be many differences in acknowledged contributors between two situations. Each of these can explain a difference in overall phenomenology. We cannot stipulate that there are no differences in these contributors, because if we do, the intuition that there is a difference in the overall phenomenal experience is likely to vanish, and is in any case not to be trusted. Claiming to remember such situations is of no help, for a situation about which one can only remember the difference in a certain feature does not provide good evidence that that actually was the only difference.

We cannot salvage the argument by claiming phenomenology of a special kind: that begs the question. We cannot claim that the magnitude of difference in overall phenomenal experience is greater than what the acknowledged contributors can explain, for we cannot detect magnitudes of phenomenal differences to the required level of accuracy. We get no help from the different difference reply: that the difference should vary is precisely what the argument predicts. And we cannot adduce a persistent difference across it-

erations, for this alleged datum cannot be established either by memory or by intuition, and would anyway not rule out that associated states distinct from **M** are really responsible.

The conclusion stands.

6 Applications

I have defended my critique against a range of objections. I now apply it to two cases.

6.1 Cognitive Phenomenology

Many philosophers advance minimal pair arguments for content-specific phenomenology of cognition. We already saw a case of this in §3; let us now consider some further examples.²¹

AMBIGUITY

Many sentences in English are ambiguous. Examples include ‘Wake up, time flies!’ (motivational speech or reprimand to an official at the insect races); ‘Don’t be mad!’ (plea not to be angry, or crazy); and ‘This curry is hot!’ (exclamation about degrees centigrade, or chilli-content). Suppose I record myself uttering these sentences, and play each recording to you twice. I give you a paper with the sentences on it, and ask you to read along. Suppose you are aware of one meaning the first time, and the other the other. There is a difference in the content of your thought (**M**), and a difference in your overall experience. The best explanation for the latter is that **M** contributes to experience. So we have good reason to believe that this is so.²²

²¹I have modified some examples slightly to remove unimportant problems.

²²The examples are from [Siewert \(1998, 278-82\)](#) and [Horgan and Tienson \(2002, 523\)](#).

HIDDEN MEANING

Some sentences in English are difficult to parse. Examples include ‘The boat sailed down the river sank’ (the boat, which someone sailed down the river, sank); ‘The boy the man the girl saw chased fled’ (the boy, who was chased by the man who the girl saw, fled); ‘Dogs dogs dog dog dogs’ (Dogs which are such that other dogs dog (chase) them, dog (chase) yet other dogs). I present such sentences to you as before. You are unable to parse each the first time, but able to do so the second. There is a difference in the content of your thought (**M**), and a difference in your overall experience. The best explanation for the latter is that **M** contributes to experience. So we have good reason to believe that this is so.²³

SUDDEN THOUGHT

Suppose you are driving toward a green traffic light. Suddenly you become aware that the light may be about to change, and you wonder whether you should speed up. You do not say, out loud or to yourself, ‘The light may be about to change’, nor do you visualise the change. There is a difference in the content of your thought (**M**), and a difference in your overall experience. The best explanation for the latter is that **M** contributes to experience. So we have good reason to believe that this is so.²⁴

Most discussants agree that occurrent, remembered or imagined perceptions, bodily sensations, moods and emotions contribute to the character of overall phenomenology, so I take this to be the relevant common ground. Call it **CG**₂.

Just as in the cinema screen case, these examples are designed to exclude a small number of contributors that easily come to mind. **AMBIGUITY**, for

²³These examples are from Pitt (2004, 27-8) (but see fn. 5).

²⁴The example is from Siewert (1998, 277).

example, is obviously built to exclude auditory phenomenology as responsible for the difference in the character of overall phenomenology. It seems that you can become aware of the different meanings without there being anything different about the way the sentence *sounds*. (This can be disputed; I grant it here.) You also read the sentence, so the visual contribution to your overall phenomenal experience appears to be the same. Parallel considerations hold for HIDDEN MEANING. SUDDEN THOUGHT is superficially different, since contributors that stay the same are not expressly described. But the description is designed to rule out plausible candidates for differences in them (overt or inner speech, visualisation).

As before, we are invited to think that the list of exclusions can be extended to all the acknowledged contributors, and we should resist that invitation. But the examples can be resisted at an even earlier stage.

It is easy to accept that the *dominant* part of your auditory phenomenology in the first two examples will be the sound of my voice as it is played back to you. But this does not show that this will be all your auditory phenomenology consists in. The sounds in the background will likely differ between the two cases, and even if not, your attention to background sounds probably will. These are candidate explanations. Similarly, the dominant visual phenomenology stemming from seeing the written sentence does not imply that that is all your visual phenomenology consists in. If you attend slightly to something in your peripheral vision in one case but not in the other, this may explain the difference in experience. In the last example, what the person hears and sees changes as the car moves. This may well account for what difference there is between the cases. So, even before moving beyond visual and auditory phenomenology we have found reason to doubt these arguments.

There are, of course, many other contributors of the occurrent variety that could also account for the difference. Occurrent perception in the other modalities is likely to change. There will likely be occurrent emotional re-

actions triggered by coming to think a content that was not thought before. For example, understanding ‘Don’t be mad!’ as being about anger seems perfectly likely to trigger an emotional reaction.

In his discussion of *SUDDEN THOUGHT*, Siewert notes that the traffic light thought need not be associated with any great emotional highs or lows. True. But emotions contribute to overall experience even if they are short-lived and not of great intensity. So emotion may well explain what difference there is in overall phenomenology. If you have experienced severe peak-hour traffic, the thought may be associated with at least a weak annoyance or frustration for you. And so on.

We have considered occurrent versions, but contributors also include the remembered and imagined varieties. In *AMBIGUITY*, when becoming aware of the reading that corresponds to a reprimand at the insect races, you may briefly imagine the way buzzing insects sound on a summer day, or faintly remember an experience of that sound. Maybe the starter’s gun at a different type of race comes to mind, via memory or imagination. You might imagine or remember the feeling of the sun warming your skin on a summer day, the faint smell of honey on the sandwich that attracted the buzzing bees, or some other perception altogether. And there are countless other possibilities, just among perceptual experiences. In addition, you might remember or imagine bodily sensations (a bee sting, feeling over-full after too many honey sandwiches, . . .), or moods or emotional reactions (fear of bees, delight of impending honey-consumption, . . .).

What is needed, then, is not only that the difference in experience cannot be accounted for by differences in occurrent perceptions, bodily sensations and moods and emotional states, or differences in the way the subject focuses on one or more of these. The friend of cognitive phenomenology also needs to establish that there is likely to be no difference in *remembered* or *imagined* perceptions, bodily sensations, moods and emotional states, nor (arguably) a difference in the vivacity with which any of these are remem-

bered or imagined, nor a difference in the attention the subject affords either of these things. I hope it is becoming apparent just how monumental a task this would be.

Some might be tempted to apply one of the manoeuvres considered in §5 above to the case of cognitive phenomenology. But they fail, for reasons exactly parallel to those outlined there. So minimal pair arguments cannot rationally persuade us in the case of cognitive phenomenology.

6.2 Properties Represented in Perception

In ‘Which Properties are Represented in Perception?’ Susanna Siegel argues that some properties which are not standardly thought to be represented in experience (‘K properties’), in fact are. She presents the following argument:

- (0) The overall experience of which E1 is a part differs from the overall phenomenology of which E2 is a part.
- (1) If the overall experience of which E1 is a part differs from the overall phenomenology of which E2 is a part, then there is a phenomenological difference between the sensory experiences E1 and E2.
- (2) If there is a phenomenological difference between the sensory experiences E1 and E2, then E1 and E2 differ in content.
- (3) If there is a difference in content between E1 and E2, it is a difference with respect to K properties represented in E1 and E2 (2006b, 491).

E1 is “the sensory experience had by a subject S who is seeing the pine trees before learning to recognize them”, and E2 is “the sensory experience had by S when S sees the pine trees after learning to recognize them” (491).

The case is a little underspecified, but we can add to the description in a way that is, I take it, consistent with the author’s intention. Let the overall experience of which E1 is a part be the experience had by the subject when

she stands at some distance to the grove, looking at it, before starting the job for the first time, at time t_1 . At this time, she has no recognitional capacity for pines. She is not so far away as to be unable to distinguish features of individual trees, but far enough away to see a fair portion of the grove. The overall experience of which E2 is a part is that which the same person has standing in the same spot a few weeks into the job (at t_2), looking at it. She has now acquired the capacity to immediately recognise pine trees.

The intuition with which Siegel starts is that the characters of the overall phenomenologies in these two cases differ from each other. Granting this, Siegel's argument can be resisted along the same lines as before.

If we adopt the reasonable assumption that the relevant common ground is CG_2 , it is clear that many acknowledged contributors are likely to differ between the situations. Occurrent differences in other sensory modalities are likely to be present, as are differences in attention. At t_1 , the person is likely to experience some anxiety, to imagine what a pine tree looks like, what the work will feel like, smells and sounds, bodily sensations, and so on. At t_2 it is unlikely that any of these imaginings will take place. Instead, a number of remembered perceptions, bodily sensations, moods and emotions are likely to appear, involving bodily sensation associated with the work, remembered feelings of satisfaction or annoyance, sights, sounds, smells. There will probably be no anxiety. So there are numerous candidate explanations that do not involve accepting Siegel's premise (1).

Siegel says that resisting premise (1) means denying that something perceptual is responsible for the difference in experience. That seems wrong. The difference in experience can be imagined or remembered visual perception, occurrent perception in a different modality, or differing attention in the visual modality. All of these are perceptual, but none of them support premise (1).

7 Concluding Remarks

I have argued that minimal pair arguments systematically fail to rationally persuade us of their conclusions. I have also said that I think we know, among other things, that perceiving something blue makes a different contribution than does perceiving something red. How do we know this? This might be thought to pose a particular challenge for me, since someone might claim that minimal pair type reasoning is what teaches us that there is content-specific perceptual phenomenology. But, as we have seen, that argument does not work. So there can be no particular challenge here; how we know what we do about experience is, as the saying goes, everyone's problem.

I suspect the solution most likely lies in 'abstractable aspects' of overall experience (Horgan and Tienson 2002). It seems that we somehow manage to 'distill' out some information about specific contributors from the many experiences we have, a process which I suspect takes place at the sub-personal level. Of course, this raises further questions. How does the abstraction take place? What are good conditions with respect to it? And what is the epistemic status of claims that rely on it? These are likely to provide fruitful avenues for further research.

That I cannot answer these questions here does not detract from the challenge I have presented. Perhaps the challenge can be met, so that minimal pair arguments can rationally persuade us after all. If so, the proponents of these arguments should show us how.

References

- Boring, Edwin G. 1946. “Mind and Mechanism”. *The American Journal of Psychology* 59: 173–192.
- Chalmers, David J. 2003. “The Content and Epistemology of Phenomenal Belief”. In *Consciousness: New Philosophical Perspectives*. Oxford University Press.
- . 2004. “The Representational Character of Experience”. In *The Future for Philosophy*, edited by Brian Leiter. 153–181.
- Farkas, Katalin. 2008. “Phenomenal Intentionality Without Compromise”. *The Monist* 91: 273–293.
- Grice, H. P. 1962/1989. “Some Remarks about the Senses”. In *Studies in the Way of Words*. Harvard University Press, 248–268.
- Horgan, Terrence, and John Tienson. 2002. “The Intentionality of Phenomenology and the Phenomenology of Intentionality”. In *Philosophy of Mind: Classical and Contemporary Readings*, edited by David Chalmers. Oxford University Press.
- Horgan, Terry, and George Graham. Forthcoming. “Phenomenal Intentionality and Content Determinacy”. In *Prospects for Meaning*, edited by Richard Schantz. De Gruyter.
- Horgan, Terry, John Tienson, and George Graham. 2003. “The Phenomenology of First Person Agency”. In *Physicalism and Mental Causation: the Metaphysics of Mind and Action*, edited by Sven Walter and Heinz-Dieter Heckmann. Imprint Academic, 323–340.
- Huemer, Michael. 2005. *Ethical Intuitionism*. Palgrave Macmillan.

- . 2007. “Compassionate Phenomenal Conservatism”. *Philosophy and Phenomenological Research* 74: 30–55.
- Jacob, Pierre. 1998. “What Is the Phenomenology of Thought?” *Philosophy and Phenomenological Research* 63: 443–448.
- Kriegel, Uriah. 2003. “Consciousness as Sensory Quality and as Implicit Self-Awareness”. *Phenomenology and the Cognitive Sciences* 2: 1–26.
- . 2007. “The phenomenologically manifest”. *Phenomenology and the Cognitive Sciences* 6: 115–136.
- Kriegel, Uriah, and Terry Horgan. Forthcoming. “The Phenomenal Intentionality Research Program”. In *Phenomenal Intentionality: New Essays*, edited by Terry Horgan and Uriah Kriegel. Oxford University Press.
- Peacocke, Christopher. 1998. “Conscious Attitudes, Attention, and Self-Knowledge”. In *Knowing Our Own Minds*, edited by Crispin Wright, Barry C. Smith, and Cynthia Macdonald. Oxford University Press, 63–98.
- Pitt, David. 2004. “The Phenomenology of Cognition, or, What is It Like to Think That P?” *Philosophy and Phenomenological Research* 69: 1–36.
- Pryor, James. 2000. “The Skeptic and the Dogmatist”. *Nos* 34: 517–549.
- . 2004. “What’s Wrong with Moore’s Argument?” *Philosophical Issues* 14: 349–378.
- Searle, John R. 1983. *Intentionality*. Cambridge University Press.
- Siegel, Susanna. 2005/2010. “The Contents of Perception”. In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. (Available from plato.stanford.edu).

- . 2006a. “Subject and Object in the Contents of Visual Experience”. *Philosophical Review* 115: 355–88.
- . 2006b. “Which Properties are Represented in Perception?” In *Perceptual Experience*, edited by Tamar Szab Gendler and John Hawthorne. Oxford University Press.
- Siewert, Charles P. 1998. *The Significance of Consciousness*. Princeton University Press.
- Smithies, Declan. 2006. “Rationality and the Subject’s Point of View”. Ph.D. thesis, New York University. (Available from <https://sites.google.com/site/declansmithies/>).
- Strawson, Galen. 1994/2010. *Mental Reality*. Second edn. The MIT Press.
- Tolhurst, William. 1998. “Seemings”. *American Philosophical Quarterly* 35: 293–302.