

35 Can Representationism Explain How Attention Affects Appearances?

Sebastian Watzl

Gustav Fechner and William James, both famous early psychologists as well as philosophers, disagreed on whether attention affects appearances. Fechner thought that “the pendulum-beat of a clock [appears to us] no louder, no matter how much we increase the strain of our attention upon [it]”¹ William James, by contrast, believed that “in listening for certain notes in a chord, the one we attend to sounds probably a little more loud.”²

At least for the case of visual appearances, the debate between Fechner and James has recently been resolved. James was right. A Gabor patch, for example, looks to have a higher contrast with attention than without attention (Carrasco, Ling, and Read 2004). Similarly, attention affects apparent brightness (Tse 2005), apparent color saturation (Fuller and Carrasco 2006), apparent size (Anton-Erxleben, Henrich, and Treue 2007), apparent spatial distance (Liverence and Scholl 2011), and apparent duration (Tse et al. 1997). While the effects of attention on appearances are small (and so Fechner and James couldn’t resolve the issue on the basis of introspection alone), they are statistically significant and reproducible.³ Carrasco and colleagues sum up their findings as follows:

These studies confirm the intuition of William James..., who proposed that attention could intensify the sensory impression of a stimulus. Attention not only affects how we perform in a visual task, it also affects what we see and experience. (Carrasco, Fuller, and Ling 2008, 1163)

What is the significance of these findings concerning how attention affects appearances?

In an influential publication, Ned Block (2010) has argued that Carrasco and colleagues’ findings show that a widely held view about the relationship between the phenomenal character of perceptual experience and its representational content, a view Block calls representationism, must be mistaken. Block says,

[These] empirical facts concerning [how attention affects appearances] ... conflict with forms of representationism.... According to which the phenomenal character of perception is or supervenes on its representational content (how it represents the world to be). (Block 2010, 25)

—1
—0
—+1

Block's conclusion is surprising. If correct, it would show that representationism cannot even capture those aspects of phenomenology that make it most appealing—namely, those captured by how the world looks.

Block's conclusion is opposed to an extremely natural interpretation of Carrasco and colleagues' findings—one indeed originally suggested by them—namely, that attention *exactly* affects the way perceptual experience represents the environment, or “[the properties] we see and experience” (Carrasco, Fuller, and Ling 2008, 1163). In particular, Carrasco and her colleagues propose that their studies show that attention may lead to distortions of perceptual representation. They say,

Attention augments perception...by emphasizing relevant details at the expense of a faithful representation of the sensory input....Attentional enhancement of neural firing is interpreted as if the stimulus had a higher contrast. (Carrasco, Fuller, and Ling 2008, 1162–1163)

What, then, is Block's argument against representationism? Why does he think that Carrasco's findings cannot be interpreted in terms of effects on representational content?

The first goal of this chapter is to provide a reconstruction of Block's argument. I show that the argument is best viewed as presenting us with a new instance of what I call the *problem of arbitrary phenomenal variation* (or the arbitrariness problem). The most famous instance of this problem, discussed by Block in some earlier writings (1990, 1999, 2007), appeals to the conceivability of spectral inversion, or the idea that the same object might look red to one person but green to another (see also MacPherson 2005; Thompson 2008). I provide a statement of the arbitrariness problem in terms of several plausible but jointly inconsistent claims (one of them is representationism).

The second goal of the chapter is to consider responses to Block's argument. Once the structure of Block's argument is revealed as an instance of the arbitrariness problem, we can get a clearer view of the space of possible options. I argue that many of these responses, including some that have been proposed in the recent literature, have serious drawbacks. In particular, some responses that might be quite appealing when considering spectral inversion are much more problematic when considering Block's new attention argument. I also suggest, though, that the view originally proposed by Carrasco and colleagues emerges surprisingly untouched; it is the view according to which attention distorts perceptual representations because veridicality gets sacrificed for usability. Attention creates useful perceptual fictions. I end my discussion by showing how to develop that view.

1 The Target of Block's Argument

We can state the target of Block's discussion as follows:

Representationism. How things phenomenally look to the subject in an experience is fully determined by the representational content of that experience.

-1—
0—
+1—

This presentation of Block's target differs from the presentation Block himself provides in the quote above.

First, Block speaks of "phenomenal character," whereas my formulation uses "how things phenomenally look to the subject." This change is intended as a restriction on the aspect of phenomenal character that is determined by representational content, namely one that amounts to a difference in appearance ("how things look"). Block is interested in how attention affects features like apparent contrast or apparent size. The view Block is attacking thus is compatible with the claim that other aspects of visual phenomenology, such as the salience of some part of visual experience or what is central and what is peripheral in an experience, cannot be explained in terms of content, as long as these aspects of phenomenology do not affect how things look to the subject.⁴ Since the explanation of phenomenal character in terms of representational content is most appealing for those aspects of the character that *do* make a difference to how things look (see, e.g., Byrne 2001; Siegel 2010), Block thus attacks representationism where it might have seemed strongest.

Second, by speaking of things phenomenally *looking* a certain way to the subject, I have explicitly restricted representationism to the visual modality. The relevant form of representationism is thus compatible with views on which phenomenal differences between the sensory modalities are not explained by differences in content. Representationism thus would not be undermined by pointing to a phenomenal difference between seeing a certain shape and touching that shape that gets explained in terms of a difference in mode or attitude instead of the content of the relevant states (Crane 2003; Chalmers 2005).⁵

Third, I have formulated representationism not just in terms of supervenience or modal entailment but in terms of what *determines* how things look. Block himself clarifies his target in this way when he says that his "case against representationism depends on a version of it in which phenomenal character is determined by or flows from representational content rather than merely supervening on it" (Block 2010, 27). I take it as understood that the relevant version of representationism thus rules out accounts of the contents of visual experience that specify them in irreducibly phenomenal terms.

Fourth, representationism makes no assumptions about what, if anything, explains why experiences have the contents they do. The view is supposed to be compatible with both naturalistically reductivist accounts that attempt to reductively explain phenomenal intentionality in naturalistic terms (e.g., Dretske 1995; Tye 2000) and antireductivist accounts that take the relevant intentional notions as primitive (e.g., Pautz 2006, 2010a, 2010b). Relatedly, the target of Block's argument is compatible with both externalist accounts, on which the content of experience depends on factors about the subject's environment (Dretske 1995; Byrne and Tye 2006), and internalist accounts, on which the content of experience depends only on factors roughly within the head of the subject (Pautz 2006).

—1
—0
—+1

Representationism is a widely held position, and the literature contains a number of arguments in its favor.⁶ It thus is a worthwhile target for philosophical debate, and if Block's argument were successful, we would have to rethink much about the connection between phenomenal character and representational content as well as consider consequences for other areas such as the epistemology of perception.⁷

2 The Problem of Arbitrary Phenomenal Variation

2.1 The Problem Schema

Let us start with a rough statement of Block's argument.

Carrasco and colleagues' findings show that without attention an object looks to have one contrast, while with attention it looks to have a different (higher) contrast. Now we can ask the following crucial question: Which of these two experiences is an *accurate* representation of the object? It seems that they cannot both be accurate or both be inaccurate. So exactly one of them must be accurate. But, we can continue to argue, it would be completely arbitrary to say of either one of these experiences that *it* is the accurate one, while the other is inaccurate. But then it seems that having an experience in which something looks to have a certain contrast does not settle *conditions of accuracy*. And *that* seems to be incompatible with the idea there is a representational content that determines how things look—that is, with representationism.

We can get a better grip on the structure of Block's argument by comparing it to cases of so-called *spectrum inversion*. In these cases, we are asked to conceive of a person, call her Invert, whose visual system is wired differently from ours so that when some object looks red to one of us, call her Vert, it looks green to Invert (and vice versa). Now we can ask a question similar to Block's: Which of these two experiences (Vert's or Invert's) is an accurate representation of the object? Again, it seems that they cannot both be accurate or both be inaccurate. But, the argument continues, it would be completely arbitrary to say of either one of these experiences that *it* is the accurate one, while the other is inaccurate. And again, with the same argumentative steps as above, we seem to have put pressure on representationism.

These rough ideas can be made precise in terms of the following argument schema, a schema that is exemplified by both the attention argument and the spectrum inversion argument. The schema consists of a number of jointly inconsistent claims. One of these claims is representationism. If the other claims can be shown to have more weight than our attachment to representationism, the schema can thus be used as a *reductio ad absurdum* against representationism. Here then is the *problem of arbitrary phenomenal variation*.⁸

1. How things phenomenally look to the subject in an experience is fully determined by the representational content of that experience (representationism).

-1—
0—
+1—

2. The representational content of an experience consists in a type of accuracy condition C for that experience (content as accuracy conditions).
3. Necessarily, if an experience has accuracy conditions C, then for each possible scenario given C either the experience is accurate with respect to that scenario or it is inaccurate with respect to that scenario (accuracy is determinate).
4. Necessarily, if an experience is accurate or inaccurate with respect to a scenario, then it is not arbitrary to suppose that this experience is accurate or inaccurate with respect to that scenario (accuracy rules out arbitrariness).
5. Possibly there exists a scenario S and two experiences E1 and E2 such that
 - a. E1 and E2 differ with respect to how things phenomenally look (different looks).
 - b. If how things phenomenally look in E1 and E2 is fully determined by accuracy conditions C, then given C
 - i. E1 and E2 are not both accurate with respect to S (incompatibility).
 - ii. E1 and E2 are not both inaccurate with respect to S (no error).
 - iii. It is arbitrary to suppose that (with respect to S) E1 is accurate and E2 is not, and it is also arbitrary to suppose that E2 is accurate and E1 is not (arbitrariness).

Claims 1–5 together are inconsistent: From 1 and 2 we get that how things phenomenally look to the subject in an experience is determined by a type of accuracy condition. With this, we can discharge the antecedent of 5b. From 5bi, 5bii, 5biii, and 4 we get that, given the accuracy condition implied by claim 1 and 2, E1 and E2 are neither accurate nor inaccurate with respect to the scenario at hand. But claim 3 says that given any accuracy condition for an experience that experience must be either accurate or inaccurate. And so we have a contradiction.

If the arbitrariness problem is to be used as an argument against representationism, the other claims have to be plausible (i.e., more plausible than representationism). Conversely, a proponent of representationism has to show that one of them can plausibly be rejected.

Let us start to look at these other claims. They are of two kinds. Claims 2–4 are general principles about what it takes to have representational content or accuracy conditions and about what is necessarily true of such accuracy conditions. Their plausibility will have to be assessed on general grounds. Claim 5, by contrast, is about the possibility of a certain type of scenario. Its plausibility will have to be assessed by considering specific scenarios of that type.

The general claims come with a high degree of initial plausibility (I look more closely at them in the second half of this chapter).

Content as accuracy conditions states a thesis about what content is. Content consists in a kind of accuracy condition (see Siegel 2010). Its intuitive force is that if two experiences differ in content then they “say” something different. But if they say something

—1
—0
—+1

different, then it seems that the situations in which what they say is accurate are different, and so they have different accuracy conditions.

Accuracy is determinate states that the answer to the question Is this experience accurate with respect to some situation S? can never be it's both, or it's neither, or it's accurate to some degree but not fully. The intuitive idea here is that "being accurate" (like "being true") is not a vague predicate. When we say that something is approximately accurate or partially accurate, the only thing we mean is that in some respects it is accurate and in others it is not. We don't mean that it is genuinely vague whether that experience is accurate with respect to a specific situation.

For the motivation of *accuracy rules out arbitrariness* consider the following. Suppose that it is arbitrary to suppose that an experience is accurate or inaccurate. Consider, for example, that on a certain occasion I feel nauseous. Now suppose, not too implausibly, that it would be completely arbitrary to suppose that my feeling is either accurate or inaccurate with respect to my situation. If one believes that there is this arbitrariness, it seems to follow that there is simply no fact of the matter as to whether that feeling was accurate. Consequently, someone who wants to defend that feelings of nausea *have accuracy* conditions would probably argue that it is not arbitrary to suppose that in one circumstance the feeling is accurate (e.g., when your stomach is upset) while being inaccurate in another (e.g., when it is not upset).

The claims comprising 5, as I said, get supported by appeal to a specific scenario (call this an *arbitrary scenario*). The details of description of this scenario are supposed to make plausible that it is a possible scenario in which two experiences differ in how things phenomenally look to the subject (*different looks*). But given the accuracy conditions that are supposed to determine how things look to the subject, it is not the case that both are accurate (*incompatibility*) or that neither is accurate (*no error*). But it would also be completely arbitrary to suppose that it is one but not the other experience (*arbitrariness*).

2.2 Two Instances of the Schema

Let us then look at two instances of the schema provided by the problem of arbitrary phenomenal variation, the spectrum inversion case and the attention case.

Here is the spectrum inversion case again.

The Spectrum Inversion Case. Consider a normal human being called Vert. When Vert looks at a tomato she has an experience E1 in which the tomato looks phenomenally red to her. But now we can—on the basis of what we know about color processing in the brain—conceive of a color system that is wired differently than Vert's color system. Suppose that Invert has such a different color system and that consequently when Invert looks at a tomato she has an experience E2 in which the tomato looks phenomenally green to her.

The inverted spectrum case describes a scenario with two people, Vert and Invert, who when faced with qualitatively identical objects have experiences in which things

-1—
0—
+1—

look different. How do we know that such a scenario is possible? The scenario seems to be conceivable, and its conceivability lends at least some support to its possibility.⁹ Furthermore, what we know about color processing in the brain seems to support the claim that color experience depends on contingent facts about that processing and is not fixed (at least not fixed completely) by the properties of the objects we are perceiving.¹⁰ Since these facts about color processing are contingent, it seems that another creature could process color differently and hence have different color experiences. Considerations like these, pertaining to empirically constrained conceivability, thus seem to support this instance of *different looks*.

What about the other three claims?

First, it seems that E1 and E2 cannot both be accurate in the relevant situation. For E1 (in which the tomato looks red) would seem to be accurate just if the surface of the tomato is red, while E2 (in which the tomato looks green) would seem to be accurate just if the surface of the tomato is green. But no surface, it seems, can be both red and green. And so we get *incompatibility*.

Second, it also seems that E1 and E2 cannot both be inaccurate. For if they were, by generalization, all (or almost all) our color experience would be inaccurate. And such a massive error theory, one might think, should not be accepted. And so we get *no error*.

Third, it seems that the case can be described in a way that it would be completely arbitrary to pick either Vert's or Invert's experience as the one that is accurate in the relevant situation: suppose, for example, that half a population is like Vert and the other half is like Invert. There seem to be no grounds on which to suppose that Vert's experience is accurate, for example, since the situation of Vert and Invert is fully symmetrical. And so we get *arbitrariness*.

There are many ways to resist this description of the scenario. Some people have found it less than obvious that when we consider two people we can actually say of the phenomenal character of their experience that it is the same or different.¹¹ Furthermore, it might be said that maybe both of their experiences *can* be accurate in this situation. Maybe the tomato is red-for-creatures-like-Vert and green-for-creatures-like-Invert? That is, one might say that experienced color properties are, in some way, relative to a kind of perceiver.¹² Or maybe it would not be arbitrary to suppose that Vert, who is like all of us, is the one who perceives color accurately and not the merely imagined Invert.¹³ Accepting the spectrum inversion scenario thus requires (at least) accepting comparison of qualia across persons, rejecting color relativism, and accepting a certain kind of symmetry between actual and merely possible people. All of these might plausibly be doubted.

Let me now get to Block's attention argument. Here we have a different instance of an arbitrary scenario. Focusing on how attention affects apparent contrast (Carrasco, Ling, and Read 2004), we can present the scenario as follows:

—1
—0
—+1

The Attention Case. If a subject focuses her attention on a Gabor patch *g*, then she will have an experience *E1* in which *g* looks to have an apparent contrast *C1* (28 percent). If the same subject focuses her attention at a fixation cross (2° visual angle away from *g*), she will have an experience *E2* in which *g* looks to have a different and lower apparent contrast *C2* (22 percent).

As in the spectrum inversion case we have two experiences of qualitatively identical objects such that in the respective experiences the object phenomenally looks different to the subject. We can immediately see that this scenario is an improvement over the spectrum inversion case in several respects.

First, there is no need to appeal to any potentially controversial move from conceivability to possibility. We know that the scenario is *actual*, and hence possible, on the basis of firmly corroborated scientific research. The claim that there would be a phenomenal difference between *E1* and *E2* is not left to intuitive judgments: it is a statistically significant result that subjects judge that things look phenomenally different in *E1* and *E2* (hence the rejection of *different looks* seems to be no serious option).

Second, we also don't need to compare qualia across persons. The Gabor patch looks different to one and the same person in one and the same scenario depending on whether she attends to it or not.

Third, it seems that we cannot avoid incompatibility by suggesting relativism about the relevant property, since one and the same person has a different contrast experience depending on her focus of attention.

Why should we also accept the other three claims about this scenario?

The case for *incompatibility* seems to be the easiest. It seems clear that *E1* and *E2* could not both be accurate. For how could both the experience in which *g* looks to be 28 percent contrast and that in which it looks to be 22 percent be accurate representations of its contrast? For both experiences to be accurate, it seems that *g* would have to be both 22 percent contrast *and* 28 percent contrast, and that seems to be impossible.

To support *no error* and *arbitrariness* we need to consider another detail of Block's argument:

Block considers empirical evidence that suggests that "the distribution of attentional resources" (2010, 45) is a graded phenomenon: we can ask about how *much* attention is allocated to *g* (for simplicity, let's assume that attention comes in degrees between zero and one). Plausibly, the gradation in the allocation of attentional resources is reflected in gradation of the effect of attention on apparent contrast.¹⁴ The more attentional resources are allocated to *g*, the higher the apparent contrast it looks to have. The 28 percent look and the 22 percent look thus plausibly are just two points in a relatively smooth spectrum. With an intermediary amount of attention allocated to *g* it would look, for example, to have 25 percent.

With this further assumption concerning the gradation of the attention effect, we can now support both *no error* and *arbitrariness*. The crucial question becomes, How many attentional resources are needed for accuracy?

-1—
0—
+1—

Suppose we say *all* attentional resources (degree one). Arguably, it is almost never (maybe never) the case that all attentional resources are devoted to a single item in the visual field. In almost all circumstances some attentional resources are devoted to something else. If that is right, then in almost all circumstances our experience of contrast would be inaccurate. As Block says, “If the subject were speaking on a cell phone, nothing would be veridically perceived!” (2010, 45). We seem to be forced into the uncomfortable position of a massive error theory about contrast perception.¹⁵

We get the same result if we say that *no* attentional resources must be directed at *g* to represent its contrast accurately (degree zero). For research suggests that most consciously perceived items receive at least some degree of attention (Cohen, Alvarez, and Nakayama 2011).¹⁶ It might also seem implausible that to perceive the contrast of *g* correctly the best thing to do is to focus as hard as you can on a difficult practical or theoretical problem that absorbs (almost) all your attentional resources. Block’s argument for *no error* then is that otherwise we would have to accept that we almost never perceive contrast accurately.

So it seems that accuracy cannot go with *all* or *no* attentional resources. Either accuracy must be achieved with some specific intermediary amount of resources or it must be achieved if the resources are above or below some threshold. Now we get to *arbitrariness*. To pick any one specific amount of resources as the amount that yields accuracy of contrast perception seems completely arbitrary. Why should a level of attention of .5 make contrast perception accurate instead of a level of .3? The same arbitrariness results from picking some threshold θ (either more or less than θ of attentional resources directed at *g*). “An arbitrary cutoff—which some have proposed for baldness—would make no sense for veridicality,” Block says (2010, 46). Furthermore, the problem would recur within the range below or above which we are supposed to get accuracy since *g*’s apparent contrast will be different depending on whether subjects direct, for example, $\theta + .1$ or $\theta + .2$ resources at *g*. Again, we would need to pick either as accurate (*incompatibility* seems still in place), and that would be arbitrary. And so we have *arbitrariness*.

This concludes my reconstruction of Block’s attention argument. I have shown that it can be represented as an instance of an arbitrariness problem. In my view Block succeeds in making the claims that comprise that problem at least initially appealing so that there is at least some pressure exerted on representationism. I have compared the attention case to the inverted spectrum case. Given the amount of scientific support for various aspects of the attention case in contrast to the better-known inversion case, it seems to me that we should accept at least the following. If the conceivability of spectral inversion lowered our confidence in representationism at least a little, then our confidence in representationism should be lowered quite a bit more by considering the attention case.

—1
—0
—+1

3 The Space of Responses and Why Most of Them Fail

The formalization of Block's argument presented in the preceding section enables us to systematically consider responses to the argument. It also enables us to compare responses to the attention argument to responses to the parallel inverted spectrum argument. In this section I consider some such responses. I show that several of the responses that seem quite plausible in the inversion case are much less plausible in the attention case. This implausibility extends to some of the responses that have entered the recent literature. This section thus in part serves to map the logical space of options and in part serves to strengthen Block's case against representationism.

In section 4 I suggest that Carrasco and colleagues' original suggestion that attention distorts perception might still be correct. When properly assessed this leads to a plausible way of rejecting *no error*; that is, both E1 and E2 are inaccurate. As we have seen, Block's case against this depends on the idea that a rejection of *no error* would imply that almost no contrast experience is accurate. I suggest a proposal about the interplay between attention and perception that makes such an idea palatable.

3.1 Deny Arbitrariness

The most direct response to the arbitrariness problem is to provide reasons to deny that it would be arbitrary to suppose that exactly one of the relevant two experiences is accurate. As I have suggested, in the inverted spectrum case there seems to be room for this response, mostly based on one of the experiences typically being *ours*.

In the attention case one might defend an empirically oriented account on which whether attention distorts or helps the accuracy of perceptual experience is strongly context and situation dependent (Hill, unpublished manuscript).¹⁷ On this view, the accuracy conditions of the experience of a specific perceptual property are not to be discovered a priori, or from the armchair. Rather, it is a possibly complex empirical matter that determines those accuracy conditions. The discussion of which degree of attention (or which threshold) yields correct representations of contrast has to occur within this complex empirical discussion. It is quite possible that attention sometimes has one effect (making experience more illusory) and sometimes the other (making experience more accurate). While from the *armchair* exactly what degree of attention yields accuracy in a specific case might seem arbitrary, we should expect the feeling of arbitrariness to go away after the science is in.

Given its reliance on detail, the plausibility of this response depends on what these details are. Let us look at what Christopher Hill says. He proposes several options. The "forms" of attention that promote veridicality might be those that "count as defaults," or "the ones that seem particularly relevant to the more basic informational functions of attention and vision," or "the ones that best serve some weighted average of the various functions of attention and vision" (Hill, unpublished manuscript, 7).

-1—
0—
+1—

None of this, in my view, is responsive to the best version of the attention case. In this version (which is the one presented above) we compare two different degrees of attention strictly between 0 and 1. For example, consider that with a degree of attention of .8 subjects experience a 28 percent contrast, while with a degree of attention of .6 they experience a 25 percent contrast, and with a degree of .4 they experience a 22 percent contrast. If this turns out to be the case, it is hard to see why any particular level of attention could count as default, be most relevant to the most basic functions of vision, or best serve the weighted average of those functions. Different levels of attention are not different forms of attention but are more like aspects of a single form, and arguably every visual experience is infused with all levels of attention. It is hard to see how scientific discovery could single out one specific level as the one that makes vision veridical.

I believe that there is something important and correct in Christopher Hill's response. As I argue in section 4, we can indeed turn to considerations about the functions of attention and perception in order to delineate conditions for accuracy. Yet, as I will show, once that view gets developed, we see that it makes both E1 and E2 inaccurate.

3.2 Deny Incompatibility

Compatibility responses attempt to make it plausible that both E1 and E2 can be accurate in the relevant situation. How could such compatibility arise? There are several options.

3.2.1 Indeterminacy A first view would be that the contents represented by E1 and E2 are *indeterminate* and hence overlap in being accurate with respect to the situation at hand. Block discusses this form of the compatibility response and rejects it because of the vagueness in phenomenal character it is supposed to imply. How might the view be defended against Block's charges? Stazicker (2011a, 2011b) provides a defense. He argues for the general claim that experience represents merely determinable properties and not absolutely determinate ones (e.g., being red instead of being some absolutely specific shade of red). For the case of contrast properties, Stazicker claims that experience represents merely determinable contrast properties such that, for example, the determinate contrasts between 22 percent and 34 percent are all determinates of such a determinable.

It is easy to see how on such a view both E1 and E2 might be accurate (Stazicker 2011a, 2011b). Suppose, for example, that E1 represents 21–29 percent (which is a determinable contrast property such that all contrasts between 21 percent and 29 percent are determinates of that determinable), while E2 represents 15–23 percent. In this case, it would be true that attention increases represented contrast by about 6 percentage points, and yet both E1 and E2 would be accurate, since the actual contrast of the patch (22 percent) is a determinate of both (distinct) determinables represented by the two experiences. We have a *compatible* determinables scenario.

—1
—0
—+1

Block raises a number of objections to such a response.

First, he suggests that the indeterminacy view cannot account for the Gabor patches looking to have a specific contrast, since “there is no single ‘look’ that something has if it is 22 percent plus or minus 6 percent in contrast” (Block 2010, 52, my stylistic changes).

While it is somewhat unclear what having a “single look” comes to, Stazicker has a plausible reply to this objection.¹⁸ Some faraway object, for example, might have the look of something one to two meters tall, without looking to be any more specific size. Examples like this suggest that on a natural understanding of a “single look” the experience of merely determinable properties can present a single look to the subject.

Second, Block suggests that this view would predict fuzziness in the experience of the relevant properties that is not phenomenally apparent and so the view couldn’t *explain* or *determine* the phenomenology (as representationism claims).

Yet the consideration that answered the last objection applies here too. Stazicker (2011a) appeals to a faraway object, visually represented as being one to two meters tall and not looking fuzzy, and so it is not generally true that the visual representation of merely determinable properties is associated with any form of fuzziness.¹⁹

Block’s objections to the indeterminacy proposal thus can be answered.

Yet my reconstruction of Block’s argument shows the limits of the indeterminacy response.

Consider that, for both E1 and E2 to be accurate representations of the Gabor patch’s contrast, the indeterminacy has to be just large enough (or correspondingly the shift in represented contrast due to attention has to be just small enough). The size of the indeterminacy and the size of the shift in contrast have to be coordinated in just the right way for the case to be a compatible determinables scenario (as in the example given above). But to point to the possibility of these scenarios does nothing to exclude *other* possible scenarios in which the shifts in represented contrast due to attention are a little larger (or the indeterminacy is a little smaller) so that one of E1 or E2 must be inaccurate. Suppose, for example, that somewhat more determinate contrast properties are represented so that E1 represents 24–28 percent while E2 represents 18–22 percent. This scenario, like the one suggested by Stazicker, is compatible with the empirically observed shift in apparent contrast by 6 percentage points, and yet the determinable properties represented by E1 and E2 do not overlap, and so at least one of them must be inaccurate. We can call this an *incompatible* determinables scenario.

The mere possibility of an incompatible determinables scenario is enough to generate the arbitrariness problem, since in that scenario it would still be arbitrary to say that it is E1 rather than E2 (or the other way around) that is accurate. Indeed, there is no current evidence to decide whether the *actual* scenario in Carrasco and colleagues’ experiments is a compatible determinables scenario rather than an *incompatible* determinables

scenario (though such evidence could probably be gathered). The arbitrariness problem makes clear that the possibility of *one* scenario with arbitrary phenomenal variation is enough, and so the indeterminacy response remains unsatisfactory.

3.2.2 Phenomenal-response dependence The indeterminacy proposal is just one way to implement the idea that the properties represented by E1 and E2 are compatible so that the Gabor patch could have both of them. Another option (one that has played an important role in discussions of the inverted spectrum version of the arbitrariness problem) suggests that these properties are *response dependent*.

There are many ways to spell out the idea of a response dependent property.

One suggestion is that the relevant properties are either relations to certain phenomenal characters or dispositions to elicit certain phenomenal characters in certain situations (Shoemaker 1994, 2006b). We can call this the phenomenal-response-dependence view.

This form of compatibility response appeals to phenomenal character in its account of content. Therefore, it is incompatible with a form of representationism according to which contents determine or explain phenomenal character. Since the target of Block's article is such a form of representationism, the phenomenal-response-dependence view thus falls outside Block's target zone, and I don't discuss it here further.

3.2.3 Indexing The phenomenal-response-dependence view was unexplanatory because it appealed to phenomenal character in order to specify the relevant represented properties. Yet one can avoid this feature. The properties represented by E1 and E2 could be indexed to or be relative to a subject's mental state in some way that doesn't mention phenomenal character. I call this the *indexing response*.

In the case of the inverted spectrum the indexing response might appeal to color *relativism* (Cohen 2009) or a (selectionist) color *pluralism* (Kalderon 2007).²⁰ On either view, color properties are indexed to the way an organism responds to and processes spectral information (the pluralist and the relativist disagree on the nature of that indexing). The indexing view thus helps with the inverted spectrum case since Vert and Invert process spectral information differently. On the resulting view, Vert represents the object as red_{Vert} , while Invert represents the object as $\text{green}_{\text{Invert}}$. Since red_{Vert} and $\text{green}_{\text{Invert}}$ are compatible properties, the object could have both of them, and so both Vert's and Invert's experience would be accurate.

Could these ideas be applied to the attention case? Here E1 and E2 are experiences of one and the same organism, and so indexing to the way a specific type of organism processes perceptual information won't do the job. In the attention case the relevant experiences would have to represent a contrast property that is *indexed to a specific state of attention* (Prettyman 2016). Represented would be something like 28-percent-contrast-with-attention or 22-percent-contrast-without-attention. These would be compatible

—1
—0
—+1

properties so that a Gabor patch could be both 28-percent-contrast-with-attention *and* 22-percent-contrast-without-attention.

The attention-indexed pluralist (or relativist) view of experienced contrast, in my view, is problematic.

One argument against the view might appeal to what is known about contrast *constancy* (e.g., Georgeson and Sullivan 1975): since experienced contrast is invariant under different illumination conditions, viewing distances, and the like, one might argue that it must be an objective property and not one that is indexed to a specific state of the organism.²¹ Carrasco and colleagues' results, of course, seem to show that perceived contrast is *not* invariant across different distributions of attention, and so it is not entirely obvious how appeal to perceptual constancies gets a grip here. Nevertheless, its rejection of the claim that perceptual experience represents objective contrast properties reveals the radicalness of the application of the indexing view to the attention case.

There is an even clearer argument against the indexing view. Subjects are able to phenomenally compare contrast with one level of attention to contrast with a different level of attention (which is just what the subjects of Carrasco's experiments do): an unattended item might look to the subject to be the same contrast as an attended item. On the indexing view it is hard to see on what basis subjects would make that phenomenal comparison; contrast with one level of attention, after all, is a different property than contrast with a different level of attention. The attention-indexed view of contrast properties thus has to reject the highly plausible claim that if two items look to have the same contrast, then there is some contrast such that both items look to have that contrast.

Overall, the indexing view thus is much less plausible for the *intrapersonal* attention case than for the *interpersonal* inverted spectrum case. It is hard to deny that E1 and E2 indeed represent incompatible properties, and so only one of them can be accurate.

3.3 Deny Different Looks

So far we have assumed that we can compare the phenomenal character of the two experiences that generate the arbitrariness problem. But maybe the phenomenal character of our two target experiences is incommensurable so that it can be said neither that they differ with respect to how things phenomenally look in them nor that they are the same in this respect.

The incommensurability response initially seems puzzling. Why should we *not* be able to compare the phenomenal character of two experiences at least to say whether in having them things look either the same or different? Yet in the case of the inverted spectrum this response has been motivated by appeal to the claim that the phenomenal character of experiences cannot be compared between subjects, a response that has come to be known as the Frege-Schlick view (Stalnaker 1999; discussed in Shoemaker 2006a).

-1—
0—
+1—

Whatever the virtues of the Frege-Schlick view for cases of *interpersonal* comparisons of phenomenal character (or even for cases of *diachronous* intrapersonal comparisons), the view won't apply to *synchronous intrapersonal* applications of the arbitrariness problem like Block's attention argument. One and the same subject straightforwardly can and does compare the phenomenal character of (the possibly simultaneous) experiences E1 and E2 in this case. The incommensurability response thus doesn't get off the ground.

3.4 Deny Content Determines Accuracy Conditions (or Rather, Point Out an Equivocation)

So far, I have discussed several of the claims about the specific scenario at hand. I have not yet discussed the idea that *both* experiences might be inaccurate. In the end this is the view that I believe to have the highest plausibility. But before I get there, I discuss whether one might respond to Block's argument by denying one of the general principles embodied in claims 2–4 in the statement of the arbitrariness problem.

How about *content determines accuracy conditions*? The most plausible response in this area, in my view, would appeal to a Fregean theory of content. On this view the content of visual experience is composed of modes of presentation of properties and not of properties themselves. Two visual experiences may differ in content, even though *in one sense* they are accurate under the same conditions, because they differ in their phenomenal modes of presentation (Thompson 2008; Chalmers 2005).²²

A Fregean view of phenomenal representation has been motivated mostly on the basis of inversion arguments (Thompson 2008; Chalmers 2005). In this case the idea is that Vert's experience might represent the same physical property as Invert's phenomenally different experience because Vert's and Invert's experiences represent that physical property under different phenomenal modes of presentation (roughly, a reddish mode of presentation vs. a greenish mode of presentation).

The Fregean response is in principle also open as a response to Block's attention argument. Yet because it does not aim at an *explanation* of phenomenal character in terms of content (since the contents explicitly mention phenomenal characters), it is like the phenomenal-response-dependence view outside the target zone of Block's argument. We can thus put it to the side.

3.5 Deny Accuracy Rules Out Arbitrariness

How might one defend that a certain experience might be accurate in a certain situation, even though it would be arbitrary to suppose that it is accurate in that situation?

To make progress let us look a little closer at what "it would be arbitrary to suppose" here comes to.

Let me start with two epistemic readings.

First, consider that I might say that it would be arbitrary to suppose that exactly 5,326 grains of rice are in a certain jar. This statement is naturally understood as

—1
—0
—+1

expressing the thought that I currently have *no reason to believe* that this is the number of grains. The arbitrariness is epistemic and hence perfectly compatible with there being a determinate fact of the matter as to how many grains of rice there are in the jar, as well as with the existence of a perfectly obvious and accessible procedure for coming to know the number of grains. The present understanding of arbitrariness would seem to undermine *accuracy rules out arbitrariness*. It is implausible to hold that the fact that we currently don't have any reason to believe that an experience is accurate in a certain scenario implies that it isn't accurate in that scenario.

Second, one might suggest a stronger epistemic reading. On this reading it is arbitrary to suppose that *p* just if there is no accessible procedure for coming to know that *p*. Would this reading make *accuracy rules out arbitrariness* plausible? The answer seems to be no. Only a proponent of a form of verificationism that ties accuracy to the existence of an accessible epistemic procedure would accept it.

On the epistemic readings *accuracy rules out arbitrariness*, then, seems to have very limited plausibility. If the epistemic readings were the only ones, then a defender of representationism would have an easy response to Block's argument.

Yet the epistemic readings are not the only ones. There also seems to be a *metaphysical reading*. On this metaphysical reading the claim that it is arbitrary to suppose that *p* entails that *nothing objectively explains or grounds the fact that p*. So the claim that it is arbitrary to suppose that an experience is accurate in a certain situation would entail that there is no objective explanation for why that experience is accurate in that situation.

This metaphysical reading, I believe, is the intended understanding of arbitrariness in the arbitrariness problem: the intuition behind *accuracy rules out arbitrariness* seems to be that if an experience is accurate in a certain range of situations, then there must be something that explains why it is accurate in these situations (rather than others). It seems to me that this is the reading of arbitrariness that is also intended in the other parts of the arbitrariness problem. Consider *arbitrariness*. Block is not just claiming that we could not know which level of attention yields accuracy, which would seem compatible with the claim that there is a certain specific level that yields accuracy. Rather, I take it, Block's claim is that objectively there isn't anything that could explain why an experience of contrast would be accurate at one level of attention rather than some other level.

One way to reject Block's argument then would be to hold that facts about the accuracy conditions that are determined by how things look are *fundamental* facts that have no further explanation. We can call this *primitivism about phenomenal representation* (see Pautz 2010a, 2010b). Primitivism about phenomenal representation would undermine *accuracy rules out arbitrariness* on what I take to be the intended and most plausible reading. A proponent of this way of resisting Block's rejection of representationism

-1—
0—
+1—

would thus claim that it is a fundamental fact that a certain level of attention yields an accurate representation of contrast, a fact without any further explanation. A discussion of primitivism about phenomenal representation is beyond the scope of this chapter (for discussion see Pautz 2010a, 2010b). Most who are attracted to any form of physicalism or naturalism would probably find it unattractive (since such primitivism explicitly denies that representational facts have a physical or naturalistic explanation). I now sketch a response to the attention version of the arbitrariness problem that won't depend on it.

3.6 Summary

So far I have considered a variety of responses to Block's argument. I have shown that once we see that his argument has the shape of an arbitrariness problem, several of the responses against Block's argument in the literature are much less plausible. Furthermore, I have shown that many of the strategies that seem at least somewhat appealing when considering cases of spectral inversion are much less appealing when considering the attention case. Those attracted to explaining how things phenomenally look to the subject of an experience in terms of the representational content of that experience seem to be led to a form of primitivism about the relevant form of representation that many proponents of representationism would probably find hard to accept.

4 A Better Response: Deny *No Error*

One route, though, we have not yet considered. This is the rejection of *no error*, or the claim that both E1 (in which the subject attends to the Gabor patch) and E2 (in which she does not attend to it) are inaccurate. I will now show that this response is the one naturally to take for someone who develops Carrasco, Fuller, and Ling's (2008) initial response to their own findings; namely, in these experimental settings attention distorts the accuracy of perceptual representation.

4.1 The Diverging Functions of Attention and Perception

My starting point is an observation about the central functions of perception and perceptual attention. Obviously, conscious perception and perceptual attention are intimately connected: most of what we consciously experience receives at least some attention. Yet perception and perceptual attention being closely connected does not entail that the capacity for perception and the capacity for perceptual attention are one and the same capacity. Indeed, it is highly plausible that they are different capacities. Their central functions diverge: while the central function of the perceptual systems is to *enable accurate* personal-level *representations* of the organism's environment, the central function of perceptual attention is to *make* such personal-level *representations usable*.

—1
—0
—+1

Consider first the function of visual perception. At least since Marr's (1982) celebrated book on vision, it has been widely accepted that vision serves a representational function. It enables the representation of a certain range of features of the environment. The claim that perceptual capacities should indeed be thought of as enabling personal-level (or rather, organism level) representation of properties of the organism's distal environment has received a detailed defense in Burge (2010). The perceptual representation of such distal properties is connected to the existence of perceptual constancies: perceptual representations of distal properties remain invariant under a variety of distinct viewing conditions. Important for present purposes is that perceptual constancies manifest in invariant appearances: an object will *look* to the subject to have the same contrast, size, shape, or color independently of viewing distance, lighting conditions, and so on.

What about the function of perceptual attention? The idea that perceptual attention serves to *prioritize* certain stimuli over others seems to be the one most consistent with current research.²³ On this conception attention serves to organize perceptual information so that the organism can better exploit that information for further processing, bodily action, reasoning, deliberation, and so forth. We can summarize this idea by saying that the function of attention is to make perceptual representations *usable*. The view of attention as enabling usable representations can be developed in a number of different ways, but in some form or other this conception of attention has been explicitly endorsed in many discussions (Allport 1987; Neumann 1987; Mole 2010; Smithies 2011; Jennings 2012; Wu 2011; Watzl 2011, 2017).

Suppose then that perceptual capacities are indeed identified in terms of their function to deliver accurate representations of distal properties, while the capacity for perceptual attention is identified in terms of its role in making information usable. These two functions are distinct but closely related. They are distinct in that the most usable representation might not be the most accurate one and the most accurate one might not be the most usable one. But they are closely related in that a wildly inaccurate representation will most often not be very usable (think of representing an object as twice its actual size). Given that an organism has both perceptual capacities and the capacity for perceptual attention, the organism's perceptual engagement with its environment will thus involve a compromise between getting accurate representations of its distal environment and making information about its environment usable. The organism will end up in perceptual states that are close enough to being accurate while still being as useful as possible.

Let us then return to Carrasco and colleagues' initial interpretation of their data. They suggest, remember, that "attention augments perception...by emphasizing relevant details at the expense of a faithful representation of the sensory input" (Carrasco, Fuller, and Ling 2008, 1162). Their suggestion thus is that the effects of attention on appearances can be understood in exactly the terms just outlined. Ways the representation of an object might be made more useful are, for example, to make it look somewhat

-1—
0—
+1—

higher in contrast, larger in size, and to occur earlier. The appearance of perceptual items thus should be viewed as a compromise between an accurate representation and a useful representation.

4.2 How These Considerations Make Both Experiences Inaccurate

What then should we say about the accuracy of experiences E1 and E2? In section 2.2 I have suggested that Block's claims about these experiences depend on the plausible assumption that the effects of attention on apparent contrast are gradable: the larger the amount of attentional resources devoted to processing the Gabor patch, the higher its apparent contrast. The amount of attention that would yield an accurate representation then would be the one that would help optimize the function of perception. But that function would be optimized, it seems, if usability were no concern but only accuracy was. And this would happen in an idealized scenario in which the subject either did not have the capacity for attention or in which her attention was evenly distributed, thus avoiding any potential distortion of accuracy in the service of usability. Since these conditions are not in place in Carrasco's experiments, we thus get the result that both E1 and E2 are inaccurate.

It is important to avoid misunderstanding this suggestion. The suggestion is *not* that veridical perception of the contrast of some item is achieved by attending to it as little as possible. There is nothing to rule out that peripheral perception does not itself distort the representation of the properties of the peripheral item. The suggestion is rather to suppose that optimizing the function of attention is simply no concern. It is unclear whether this is empirically possible, but its closest approximation would be to distribute attention evenly and not to withdraw it. We should thus reject answering Block's question concerning which amount of attentional resources delivers accurate representations of contrast. There is no such amount and no threshold. The scenario that yields accuracy instead would be a scenario in which optimizing the function of attention played no role.

Considerations about the diverging functions of perception and perceptual attention thus deliver a plausible account of how *no error* could be false.

4.3 Objections and Responses

How might one object to this response?

Objection 1. One might again raise concerns about massive illusions. Since the condition in which concerns about optimizing the function of attention, or optimizing usability, never or almost never obtains, wouldn't all or almost all appearances of contrast, size, color saturation, and so on, still be illusory? And wouldn't accepting such massive illusion be too hard a bullet to bite?

Response. Consider as a start that what we need for veridicality is neutralizing concerns about usability. It is possible that such neutralization is approximated by even

—1
—0
—+1

distributions of attention and thus might be actualized in many actual cases. More importantly though, the relevant massive illusion, if it obtains, should not be viewed as problematic. First, the relevant distortion of perceptual representation will never, as a matter of principle, become too large. If it did, neither the function of perceptual attention nor the function of perception would be even close to optimized. Because of that, and because perceptual representations of contrast, size, and so on are much more fine-grained than the content of most perceptual beliefs and judgments, the relevant distortions will almost never result in false perceptual beliefs or judgments (except in the rare cases, as in Carrasco's experiments, in which subjects are asked to make very fine-grained discriminatory judgments). Second, consider that what seems unacceptable about the attribution of massive illusion is that we lose our grip concerning the conditions under which perceptual experience *would* be accurate. This concern though has been taken into account. Relying on the considerations about the function of perception and attention, I have provided an account of the conditions under which perceptual representations would be accurate.

Objection 2. One might call into doubt my claims about the diverging functions of perception and perceptual attention. How, if at all, can these claims be squared with other plausible claims about attention, such as the claim that it *enables* conscious perception (see, e.g., Prinz 2012) or that perceptual attention to some object *improves* our epistemic situation with respect to it (see, e.g., Stazicker 2011a)?

Response. The topic of whether attention enables conscious perception and hence is sufficient for it is a large one. Several recent lines of evidence suggest that there might be unconscious forms of attention and that attention therefore is not sufficient for consciousness.²⁴ Yet we can sidestep these debates. Even if attention were sufficient for conscious perception, its function could still be to enable usable representations and not accurate ones. With respect to the epistemic role of attention, it is important to note that even if attention might lead to distortions of perceptual representation, it could still have epistemically beneficial upshots: *beliefs* formed about the properties of attended objects might be more stable, and indeed more reliable, even if (and indeed because), in the relevant *experience*, properties such as contrast, size, or color saturation are slightly exaggerated (in part because attention often enables the discrimination of finer spatial detail).²⁵

Objection 3. One might remark, fine, maybe the preceding response answers the attention version of the arbitrariness problem, but nothing in the response seems to generalize to other cases of arbitrary phenomenal variation such as the inverted spectrum.

Response. First, and most importantly, let me note that it is not obvious that the same response should be given to all versions of the arbitrariness puzzle. As we have seen in the discussion above, several responses are implausible for the attention case, even though they have some plausibility for the inverted spectrum case. Given the different

-1—
0—
+1—

concerns raised by various versions of the problem, a more piecemeal approach might be desirable. Second, though, there *is* a general lesson. The claim that both experiences in an arbitrary scenario are inaccurate should not be viewed as worrisome in itself, even if it leads to the claim that with respect to the relevant properties *most* experiences are, strictly speaking, inaccurate. Massive perceptual inaccuracy is not in itself problematic if it can be attributed to interference with optimal perceptual functioning and if the attribution of massively false beliefs can be avoided.

5 Conclusion

Block's recent work on attention and consciousness exemplifies his approach to the philosophy of mind that smoothly integrates close engagement with empirical work and detailed philosophical argument. He has shown a whole generation of philosophers that this is how philosophical work on consciousness *should* be done and that this is how it *can* be done. In this chapter, I have attempted to both bring out the strengths of Block's argumentation and provide a response to it that aspires to embody some of the virtues of Block's approach to philosophy: once we take seriously empirical work on attention and recognize that the capacities for conscious perception and perceptual attention diverge, representationism can explain how attention affects appearances. Whether representationism is *true* is a further question. Other arguments will have to be made.

Notes

1. Fechner 1889, 452–453; quoted in James 1981, 425.
2. James 1981, 425.
3. For some doubts about this empirical evidence, consider research by Schneider (2006, 2011), Schneider and Komlos (2008), and Valsecchi, Vescovi, and Turatto (2010). I believe that the replies by Anton-Erxleben, Abrams, and Carrasco (2010, 2011) and Block (2010) are successful, and I thus take the empirical evidence at face value.
4. I have defended a view on which there is such an aspect of visual experience—structure (Watzl 2011). While Block and I thus agree that some aspects of how attention affects perceptual phenomenology cannot be captured in terms of representational content, my view is compatible with the version of representationism that Block attacks.
5. The relevant form of representationism similarly is also compatible with the claim that there are nonintentional qualitative states outside the visual modality such as feelings of nausea or moods like sadness whose phenomenal character cannot be explained in terms of representational content (e.g., Block 1996; Farkas 2009).
6. See the long list of defendants provided by Block (2010, 25).

—1
—0
—+1

7. See Siegel and Silins (2015) for an overview of the relevant debates. In particular, suppose that there is no representational content that determines how things look to the subject. For example, suppose that *x* looks to *S* to have a higher contrast than *y*, and yet *S*'s experience does not represent *x* to have a higher contrast than *y* (a possibility Block's position seems to entail). Does the subject's experience provide her with justification that *x* is a higher contrast than *y*? If justification goes with appearances, the answer would be yes. If justification goes with content, the answer would be no.
8. There are many similar schemata. Kalderon's (2007) location problem, for example, has an extremely similar form (as Kalderon notes, this problem has a distinguished history, peaking with Berkeley).
9. Thompson (2008) argues that the conceivability of spectral inversion is enough to generate a problem for a view like representationism.
10. For an overview of some of the issues about spectral inversion, see Byrne (2010).
11. Stalnaker 1999.
12. See, for example, Cohen 2009.
13. See Byrne 2010.
14. In fact, while—as Block notes—a lot of scientific evidence supports the claim that the allocation of visual attention is graded, no direct evidence supports the claim (which is needed at this point of the argument) that the effect of attention on apparent contrast itself is graded. But I follow Block in accepting this plausible empirical conjecture.
15. The same, of course, would hold for all other appearances that are similarly affected by attention such as size and color saturation.
16. This is plausible even if one accepts with Koch and Tsuchiya (2007) and against Cohen, Alvarez, and Nakayama (2011) that there are some conscious experiences that receive no attention at all.
17. Note that Hill (2016) later came to hold a different view, which is more in line with the one of this paper. Based on the empirical results by Carrasco and colleagues discussed by Block, Hill argues in this new paper that “attention-caused representations attribute ... properties that objects do not really possess” (199).
18. Stazicker 2011a, 65.
19. Ibid.
20. Stazicker (2011a, 143–147.) also responds to problems Block raises about the subjects' ability to discriminate a 28 percent patch from a 22 percent patch.
21. In fact, we don't need to take these as views about *color*. It is enough that such pluralism holds for whatever properties are represented by those experiences that generate the arbitrariness problem. Whether these should be identified with *colors* is a further question that need not concern us for present purposes.

-1—
0—
+1—

22. Such invariantly represented properties thus are not narcissistic properties in Akins's (1996) sense. See Burge (2010) for arguments that go from perceptual constancies to the representation of objective properties (i.e., properties that, if instantiated at all, are instantiated by objects in the subject's environment).
23. Alternatively, one might think of these as *centering features* (see Egan 2006).
24. It has even been argued that some such form of prioritization might be *necessary* for coordinated behavior and action (Allport 1987; Neumann 1987; Wu 2011).
25. See e.g. the review by Kentridge (2011). Block (2013) himself endorses this evidence.
26. Shiu and Pashler 1995; Yeshurun and Carrasco 1998; Montagna, Pestilli, and Carrasco 2009.

References

- Akins, K. 1996. Of sensory systems and the "aboutness" of mental states. *Journal of Philosophy* 93: 337–372.
- Allport, D. A. 1987. Selection for action. In *Perspectives on Perception and Action*, H. Heuer and H. F. Sanders, 395–419. Hillsdale, NJ: Lawrence Erlbaum.
- Anton-Erxleben, K., J. Abrams, and M. Carrasco. 2010. Evaluating comparative and equality judgments in contrast perception: Attention alters appearance. *Journal of Vision* 10: 1–22.
- Anton-Erxleben, K., J. Abrams, and M. Carrasco. 2011. Equality judgments cannot distinguish between attention effects on appearance and criterion: A reply to Schneider. *Journal of Vision* 11: 1–8.
- Anton-Erxleben, K., C. Henrich, and S. Treue, S. 2007. Attention changes perceived size of moving visual patterns. *Journal of Vision* 7: 1–9.
- Block, N. 1990. Inverted Earth. *Philosophical Perspectives* 4: 53–79.
- Block, N. 1996. Mental paint and mental latex. *Philosophical Issues* 7: 19–49.
- Block, N. 1999. Sexism, ageism, racism, and the nature of consciousness. *Philosophical Topics* 26: 39–70.
- Block, N. 2007. Wittgenstein and qualia. *Philosophical Perspectives* 21: 73–115.
- Block, N. 2010. Attention and mental paint. *Philosophical Issues* 20: 23–63.
- Block, N. 2013. The grain of vision and the grain of attention. *Thought: A Journal of Philosophy* 1: 170–184.
- Burge, T. 2010. *Origins of Objectivity*. Oxford: Oxford University Press.
- Byrne, A. 2001. Intentionalism defended. *Philosophical Review* 110: 199–240.
- Byrne, A., and M. Tye. 2006. Qualia ain't in the head. *Noûs* 40: 241–255.

—1
—0
—+1

- Byrne, A. 2010. Inverted qualia. In *Stanford Encyclopedia of Philosophy*, E. Zalta. Stanford University. Article published November 10, 2004; last modified January 20, 2010. <http://plato.stanford.edu/archives/spr2010/entries/qualia-inverted/>.
- Carrasco, M. 2009. Attention: Psychophysical approaches. In *The Oxford Companion to Consciousness*, T. Bayne, A. Cleeremans, and P. Wilken, 78–84. Oxford: Oxford University Press.
- Carrasco, M., S. Fuller, and S. Ling. 2008. Transient attention does increase perceived contrast of suprathreshold stimuli: A reply to Prinzmetal, Long, and Leonhardt (2008). *Perception and Psychophysics* 70(7): 1151–1164.
- Carrasco, M., S. Ling, and S. Read. 2004. Attention alters appearance. *Nature Neuroscience* 7: 308–313.
- Chalmers, D. J. 2005. The representational character of experience. In *The Future for Philosophy*, B. Leiter, 153–181. Oxford: Oxford University Press.
- Cohen, J. 2009. *The Red and the Real: An Essay on Color Ontology*. Oxford: Oxford University Press.
- Cohen, M. A., G. A. Alvarez, and K. Nakayama. 2011. Natural-scene perception requires attention. *Psychological Science* 22(9): 1165–1172.
- Cohen, M. A., P. Cavanagh, M. M. Chun, and K. Nakayama. 2012. The attentional requirements of consciousness. *Trends in Cognitive Sciences* 16: 411–417.
- Crane, T. 2003. The intentional structure of consciousness. In *Consciousness: New Philosophical Perspectives*, Q. Smith and A. Jokic, 33–56. Oxford: Oxford University Press.
- Dretske, F. I. 1995. *Naturalizing the Mind*. Cambridge, MA: MIT Press.
- Egan, A. 2006. Appearance properties? *Noûs* 40: 495–521.
- Farkas, K. 2009. Not every feeling is intentional. *European Journal of Analytic Philosophy* 5: 39–52.
- Fechner, G. T. 1889. *Elemente der Psychophysik*. 2nd ed. Vol. 2. Leipzig: Breitkopf und Härtel.
- Fuller, S., and M. Carrasco. 2006. Exogenous attention and color perception: Performance and appearance of saturation and hue. *Vision Research* 46: 4032–4047.
- Georgeson, M. A., and G. D. Sullivan. 1975. Contrast constancy: Deblurring in human vision by spatial frequency channels. *Journal of Physiology* 252: 627–656.
- Hill, C. S. Unpublished manuscript. Reply to Ned Block. Presented at the Eighty-Fifth Annual Meeting of the American Philosophical Association (Pacific Division) on April 21, 2011.
- Hill, C. S. (2016). Perceptual relativity. *Philosophical Topics* 44(2): 179–200.
- James, W. 1981. *The Principles of Psychology*. Cambridge, MA: Harvard University Press. First published 1890.
- Jennings, C. D. 2012. The subject of attention. *Synthese* 189: 535–554.

- Kalderon, M. E. 2007. Color pluralism. *Philosophical Review* 116: 563–601.
- Kalderon, M. E. 2011. The multiply qualitative. *Mind* 120: 239–262.
- Kentridge, R. 2011. Attention without awareness: A brief review. In *Attention: Philosophical and Psychological Essays*, C. Mole, D. Smithies, and W. Wu, 228–246. New York: Oxford University Press.
- Koch, C., and N. Tsuchiya. 2007. Attention and consciousness: Two distinct brain processes. *Trends in Cognitive Sciences* 111: 16–22.
- Liverence, B. M., and B. J. Scholl. 2011. Selective attention warps spatial representation: Parallel but opposing effects on attended versus inhibited objects. *Psychological Science* 22: 1600–1608.
- Macpherson, F. 2005. Colour inversion problems for representationalism. *Philosophy and Phenomenological Research* 70: 127–152.
- Marr, D. 1982. *Vision. A Computational Investigation into the Human Representation and Processing of Visual Information*. San Francisco: W. H. Freeman.
- Mole, C. 2010. *Attention Is Cognitive Unison: An Essay in Philosophical Psychology*. Oxford: Oxford University Press.
- Montagna, B., F. Pestilli, and M. Carrasco. 2009. Attention trades off spatial acuity. *Vision Research* 49: 735–745.
- Neumann, O. 1987. Beyond capacity: A functional view of attention. In *Perspectives on Perception and Action*, H. Heuer and A. F. Sanders, 361–394. Hillsdale, NJ: Lawrence Erlbaum.
- Norman, L. J., C. A. Heywood, and R. W. Kentridge. 2013. Object-based attention without awareness. *Psychological Science* 24: 836–843.
- Pautz, A. 2006. Sensory awareness is not a wide physical relation: An empirical argument against externalist intentionalism. *Noûs* 40: 205–240.
- Pautz, A. 2010a. Why explain visual experience in terms of content? In *Perceiving the World*, B. Nanay, 254–309. New York: Oxford University Press.
- Pautz, A. 2010b. A simple view of consciousness. In *The Waning of Materialism*, G. Bealer and R. Koons, 25–66. Oxford: Oxford University Press.
- Prettyman, A. 2011. Shifts of attention and the content of perception. Paper presented at the Harvard-MIT Graduate Philosophy Conference, Harvard University, March.
- Prinz, J. 2012. *The Conscious Brain: How Attention Engenders Experience*. Oxford: Oxford University Press.
- Schneider, K. A. 2006. Does attention alter appearance? *Perception and Psychophysics* 68: 800–814.
- Schneider, K. A. 2011. Attention alters decision criteria but not appearance: A reanalysis of Anton-Erxleben, Abrams, and Carrasco 2010. *Journal of Vision* 1113: 1–8.

—1
—0
—+1

- Schneider, K. A., and M. Komlos. 2008. Attention biases decisions but does not alter appearance. *Journal of Vision* 8: 1–10.
- Shiu, L.-P., and H. Pashler. 1995. Spatial attention and Vernier acuity. *Vision Research* 35: 337–343.
- Shoemaker, S. 1994. Phenomenal character. *Noûs* 28: 21–38.
- Shoemaker, S. 2006a. The Frege-Schlick view. In *Content and Modality: Themes from the Philosophy of Robert Stalnaker*, J. J. Thomson, 18–33. Oxford: Clarendon Press.
- Shoemaker, S. 2006b. On the way things appear. In *Perceptual Experience*, T. S. Gendler and J. Hawthorne, 461–480. Oxford: Oxford University Press.
- Siegel, S. 2010. *The Contents of Visual Experience*. Oxford: Oxford University Press.
- Siegel, S., and N. Silins. 2015. The epistemology of perception. In *Oxford Handbook of Philosophy of Perception*, M. Matthen, 781–811. Oxford: Oxford University Press.
- Smithies, D. 2011. Attention is rational-access consciousness. In *Attention: Philosophical and Psychological Essays*, C. Mole, D. Smithies, and W. Wu, 247–273. Oxford: Oxford University Press.
- Stalnaker, R. 1999. Comparing qualia across persons. *Philosophical Topics* 26: 385–406.
- Stazicker, J. 2011a. Attention and the indeterminacy of visual experience. PhD diss., University of California, Berkeley.
- Stazicker, J. 2011b. Attention, visual consciousness and indeterminacy. *Mind and Language* 26: 156–184.
- Thompson, B. 2008. Representationalism and the conceivability of inverted spectra. *Synthese* 160: 203–213.
- Thompson, B. 2009. Senses for senses. *Australasian Journal of Philosophy* 87: 99–117.
- Tse, P. U. 2005. Voluntary attention modulates the brightness of overlapping transparent surfaces. *Vision Research* 45: 1095–1098.
- Tse, P. U., P. Cavanagh, J. Intriligator, and J. Rivest. 1997. Attention distorts the perception of time. *Investigative Ophthalmology and Visual Science* 384: S1151.
- Tye, M. 1995. *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*. Cambridge, MA: MIT Press.
- Tye, M. 2000. *Consciousness, Color, and Content*. Cambridge, MA: MIT Press.
- Tye, M. 2002. Representationalism and the transparency of experience. *Noûs* 361: 137–151.
- Valsecchi, M., M. Vescovi, and M. Turatto. 2010. Are the effects of attention on speed judgments genuinely perceptual? *Attention, Perception, and Psychophysics* 72: 637–650.

Watzl, S. 2011. Attention as structuring of the stream of consciousness. In *Attention: Philosophical and Psychological Essays*, C. Mole, D. Smithies, and W. Wu, 145–173. Oxford: Oxford University Press.

Watzl, S. 2017. *Structuring Mind*. Oxford: Oxford University Press.

Wu, W. 2011. Attention as selection for action. In *Attention: Philosophical and Psychological Essays*, C. Mole, D. Smithies, and W. Wu, 97–116. Oxford: Oxford University Press.

Yeshurun, Y., and M. Carrasco. 1998. Attention improves or impairs visual performance by enhancing spatial resolution. *Nature* 396: 72–75.

-1—
0—
+1—