Perceptual Content is Indexed to Attention

Adrienne Prettyman

(Penultimate Draft, June 2016)

Abstract: Attention seems to raise a problem for *pure representationalism*, the view that phenomenal content supervenes on representational content. The problem is that shifts of attention sometimes seem to bring about a change in phenomenal content without a change in representational content. I argue that the representationalist can meet this challenge, but that doing so requires a new view of the representational content of perception. On this new view, the representational content of perception is always relative to a way of attending. I call this this the *attention-indexed view* of perceptual content.

1. Introduction

It is widely accepted that attention affects perceptual experience. I argue that this effect of attention motivates a revision of traditional notions of perceptual content, and present a revised representationalist view that can account for attention's perceptual effects. On this new view, the content of perception is indexed to attention. The representationalist needs to expand beyond the traditional view that representational content is determined by accuracy conditions specifying states of the world. Instead, perceptual content is a function of not only objects or properties in the world, but also of the distribution of attention. I call this new view of perceptual content the *attention-indexed view*, since according to it the content of perception is indexed to the direction of attention.[1]

---

[1] For a preliminary discussion of this view, see Watzl (forthcoming) and Prettyman, *Shifts of Attention and the Content of Perception, t*alk presented at the Harvard-MIT Graduate Philosophy Conference, Harvard University, March 2011.

The evidence that attention affects perception has been thought to raise a challenge for representationalist accounts of the *phenomenal content of perception*: what it's like to perceive some object, property or event. In one form, representationalists hold that the phenomenal content of perception supervenes on representational content (e.g. Tye 1995; Bryne 2001; Crane 2007). I call this *pure representationalism* (following Chalmers 2004). Pure representationalists are committed to the claim that there can be no change in phenomenal content without a change in representational content.[2] The trouble is that shifting attention sometimes seems to bring about a change in phenomenal content without a change in representational content (Chalmers 2004; Nickel 2007; Speaks 2008; Block 2010; Wu 2011; Ganson & Bronner 2013). So, the effect of attention on perception seems to force us either to reject representationalism altogether, or to adopt an impure variant, on which phenomenal facts supervene on more than just facts about content. Either move is problematic, because in both cases we are left with a residual mystery: if not representational content, what determines the phenomenal content of a perceptual experience?

The attention-indexed view puts us in a position to respond to the challenge attention seems to raise for pure representationalism. In the next section, I will review the empirical motivation for the attention-indexed view, and present the view in more detail. I will then return to

---

[2] There are many variants of pure representationalism. Some pure representationalists hold that having a particular phenomenal content just is having a particular representational content. Others permit a distinction between phenomenal content and representational content, but maintain that representational content exhaustively determines phenomenal content (Chalmers 2004). These distinctions among different pure representationalist views won't make a difference to my arguments in this paper. Impure representationalism, in contrast, allows that phenomenal content might supervene on more than just representational content, like the mode of presentation. For the purposes of this paper, when I use the term 'representationalism' I mean pure representationalism, unless otherwise specified.

the problems attention raises for representationalism, and show that my view offers the best way forward for a representationalist account of attention's perceptual effects.

2. The Attention-Indexed View of Perceptual Content

According to the Attention-Indexed View (which I will be defending in this paper), many of the properties that we represent in visual perception are indexed to attention. The main motivation for this view is simply the pervasiveness of attention's effect on perception. In what follows, I will go into detail about one study providing evidence that attention affects how we perceive visual contrast. However, many additional properties represented in vision have been shown to vary across attentional contexts. In addition to contrast, these properties include shape (Gobell & Carrasco 2004), color saturation (Blaser, Sperling & Lu 1999; Fuller, Ling & Carrasco 2004), spatial frequency (Abrams, Burbot & Carrasco 2010) and the timing of events (Stelmach & Herdman 1991; Spence & Parisse 2009). There is also evidence that attention affects perception through learning and the development of expertise (Connolly 2014). Taken together, this literature supports the view that attention's effect on perception is widespread.

The effects of attention on perception are most evident when we look at the literature on covert shifts of attention, which show that attention affects what it's like to perceive some property or object even when the senses are held relatively fixed. During a covert shift of attention, attention selects an object without redirecting one's senses to that object, as when you keep your eyes still but attend to something in the periphery. In contrast, an *overt shift of attention* occurs when attention affects perception by redirecting the senses to a new object, as when I move my eyes in the direction of a bright flash of light. By studying covert shifts of visual attention, we

can determine whether attention affects what it's like to visually perceive objects or properties independent of attention's effect on eye movement. A growing body of evidence suggests that attention does affect what it's like to perceive properties, and therefore is likely to have a pervasive effect on ordinary perceptual experience.

Consider a case discussed by Chalmers (2004). Holding gaze fixed, a subject could switch attention between two red lights on a black background without moving her eyes. As she does, what it's like to perceive the lights changes in a subtle way. To try it for yourself, see **Figure 1**.

Figure 1
Keeping your eyes fixed on the cross at the center, attempt to shift attention first to the dot on the right, then to the dot on the left. While shifting attention involves a change in perceptual experience, it is not obvious that it involves a change in representational content.

In addition to cases introduced in the philosophical literature, there is a growing body of empirical research to support the claim that attention affects perceptual experience. In one study that has received considerable philosophical discussion (Carrasco et al. 2004; see also Block 2010; Speaks 2011; Stazicker 2011; Wu 2011; Ganson & Bronner 2013; Watzl 2014 & forthcoming), covert attention was shown to boost perceived contrast by 3 to 6 percent (see **Figure 2**). In Carrasco and colleagues' study, subjects were asked to fix their gaze on a central point, and were presented with two contrast patches (also called *Gabor patches*) on either side of their fixa-

tion point. They were cued either centrally (at the point of fixation) or peripherally (near the left or right patch), and instructed to indicate the orientation of the patch with greater contrast. The cues were timed to reduce eye movement, redirecting covert attention to the cued location. What researchers were really interested in was not the orientation judgment, but rather subjects' judgments of degree of contrast. More specifically, they were interested in determining the subjects' *point of subjective equality*, the point at which the right and left Gabor patches looked the same to the subject. They hypothesized that when cued neutrally, the point of subjective equality would align with the point of objective equality: the patches would look the same when they were set to the same objective contrast by the experimenter. But if attention boosts contrast, then the point of subjective equality for a peripherally cued task should occur when the patches were objectively unequal in contrast: for example, when the cued patch was lower than the uncued (see **Figure 2**).

This is precisely what they found. When the peripheral Gabor patch was cued, it would have to be objectively 3-6% lesser in contrast than the uncued patch for subjects to reach the point of subjective equivalence. They conclude that attention boosts perceived contrast. As they put it, "we found that when observers' transient attention was drawn to a stimulus location, observers reported that stimulus as being higher in contrast than it really was, thus indicating that attention changes appearance" (Carrasco et al. 2004 p. 309).

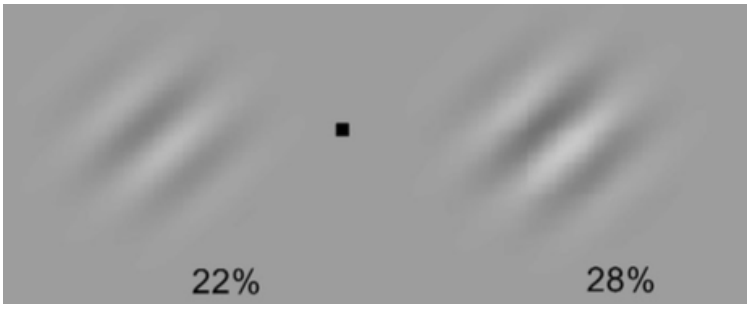

Figure 2
Fixing your gaze on the central point, covertly shift attention to the Gabor patch on the left. It should seem to you as though both Gabor patches are of the same contrast.

As I mentioned above, subsequent research has shown that in addition to contrast, attention affects perception of timing, spatial frequency, color saturation, and gap size. What is more controversial is how these and similar findings from psychology should be incorporated into traditional philosophical accounts of perception, like representationalism. If many visual properties vary across attentional contexts, as research suggests they do, this will raise problems for the traditional view of perceptual content, on which the function of perception is to represent the world as it objectively is. As I will argue, the attention-indexed account of perception puts us in a position to incorporate these attentional effects within a representationalist framework, and respond to the challenge that attention raises.

2.1 Vision is Narcissistic

Philosophers have long distinguished between properties that are *objective,* like shape or contrast, and properties that are perceiver-dependent, like taste or temperature. An objective property is a property of an object in the world. Objective properties are typically invariant across changes in the state of the subject. In contrast, properties like *hot* and *cold* are examples of what Kathleen Akins (1996) has called *narcissistic* properties: properties that are not determined by the world alone, but rather by the world relative to a subject's needs and interests.

According to traditional accounts of veridical perception, perceptual systems have the function of telling us about the objective properties of a mind-independent world. As Akins points out, the problem is that sensory systems often fail to present objective properties accurately. Rather than see perception as a deeply flawed veridical system, she suggests that sensory systems are narcissistic rather than veridical (Akins 1996 p. 350). They don't have the function of telling us how the world objectively is, but rather how the world relates to the perceiver. To make her case for the claim that sensory systems are narcissistic, Akins takes a detailed look at thermoreception (pp. 345-348). On the traditional view of verdiciality, thermoreceptors should serve as tiny thermometers, recording the temperature of the environment in a disinterested and objective way. On this view, we would expect the responses of those receptors to vary across environmental contexts but not across internal contexts or states of the subject. But thermoreceptors don't work like thermometers. On the contrary, the responses of thermoreceptors are complex, depending not only on the objective temperature of an object, but also on the perceiver's current skin temperature and the characteristics of the sensory system itself. For example, in regions of the body served by a great number of cold receptors, a cold area will be perceived as colder than it would in a region served by fewer receptors. As Akins points out, if we view thermoreception as a veridical system in the traditional sense, then it seems hopelessly flawed. But if we view it as a narcissistic system, then it is well-suited to its task. Thermoreception doesn't tell the subject what the environment is like objectively, but rather the relevance of the environment for the perceiver.

Generalizing from the case of thermoreception, Akins suggests that on a traditional view, each signal of the sensory system must correlate with an objective property or range of objective

properties. The function of perception, on this traditional view, is "servile" (Akins 350). It aims to represent accurately the properties, objects, and events in the world. Cases involving thermoreception — and, I will argue, attention's effects on perception — show that sensory systems are not servile. Unlike Akins, I think we can reconcile this fact with a traditional representationalist view of perceptual content on which content is specified by veridicality conditions. What we need is not to abandon veridicality, but rather to revise the notion to allow for narcissistic properties.

When we consider the effect of attention on visual experience, we find that — perhaps surprisingly — vision functions more like a narcissistic than a "veridical" system in the traditional sense. There are three features of thermal perception that lead Akins to posit that thermoreception is a narcissistic sensory system (pp. 351-352), and we can find analogs in the visual case due to the effect of attention. First, she points out that thermal sensations don't correlate with any particular objective temperature, since the same objective temperature can give rise to a variety of sensations depending on the state of the perceiver. Likewise, one and the same visual stimulus can elicit a variety of visual appearances, depending on the subject's direction of attention. Second, the thermal sensations don't reflect the "structure" of the thermal stimuli. That is, an environment may feel colder to us in some places and warmer in others, but this may reflect the structure of our own sensory system and not a structure of the environment. The same can be said for visual properties like contrast. For example, an object might look greater in contrast in some places and lesser and others, but this reflects our direction of attention and not a structure of the object. Third, thermoreception exaggerates or embellishes its representation of the world, so that properties that are relevant to its own interests and needs are emphasized. Similarly, when

a property is selected by visual attention, it is embellished in systematic ways. For example, it may appear great in contrast, more saturated, or to have greater spatial frequency.

The same reasoning that led Akins to posit that thermoreception is narcissistic can lead us to the the conclusion that the visual system is narcissistic. If the function of vision is to represent the world as it objectively is, then the features of visual attention outlined above would make it a flawed system, which does a poor job of representing the world in a servile manner. In contrast, if we view vision as a narcissistic system, then these features of visual attention are an asset. They serve to highlight an attended object in perceptual experience. This distinction gives us a way of understanding the philosophical moral to be drawn from the research on attention and perception. Many seemingly objective properties — like contrast — are relative to a subject's needs and interests, which guide the direction of attention. These properties are not about the world alone, but rather about the world as indexed to a particular direction of attention.

A narcissistic system utilizes what Akins calls "narcissistic encodings." That is, they encode information in a way that relates the object, property, or event to a relevant state of the subject. Which states of the subject are relevant to a particular narcissistic encoding is an empirical matter. In the case of temperature, thermal properties of objects are encoded in relation to a cluster of internal states of the subject, including the subjects' current skin temperature and the distribution of thermoreceptors. In the case of vision, some objective properties are encoded in relation to the direction of attention. For example, consider again the contrast patches from Carrasco and colleagues' 2004 experiments given in **Figure 2**. Each patch has an objective contrast, set by the experimenter. For ease of exposition let us suppose that if subjects represent degree of contrast, they do so by representing a specific numerical value, though it may in fact be a range of

values that is represented. Following Carrasco and colleagues, let us say that attention boosts perceived contrast by 3-6%. Under this assumption, we would expect the covertly attended 22% patch in **Figure 2** to be represented as 28% contrast. So, the reason that these patches look indistinguishable is that the subject represents both as 28%. When the objectively 22% patch is covertly attended, the narcissistic encoding would be 28%-attended-by-me (more succinctly, 28%-attended). In contrast, the unattended patch of objectively 28% contrast is 28%-not-attended-by-me (28%-unattended). So, there is also a difference in the narcissistic encodings, though this difference may not be reflected in the subject's perceptual experience. These numerical values are just by way of illustrating what a narcissistic encoding might look like in the case of perceived contrast. The general point is that we can account for the effect of attention on perceived contrast by specifying a narcissistic property, which in the case of visual contrast involves a contrast value (or range of values) indexed to attention. The attention-indexed view of perceptual content is just the view that a wide range of properties represented in vision — like contrast, saturation, and spatial frequency — are indexed to attention.

2.3. Attention-Indexed Properties and Veridicality

A narcissistic encoding is one that is defined relative to relevant states of a subject; we have evidence that a system is tracking narcissistic properties when it utilizes narcissistic encodings. Unlike Akins, however, I think we can make sense of narcissistic properties while retaining a revised notion of veridicality.

Consider a case involving temperature. A person who has just come in from the cold may perceive the temperature of a room to be warm, whereas a person who has been sitting in that

same room may perceive the temperature to be cold. In this case, neither is misrepresenting the temperature. Accurate representation of the temperature is indexed to a subject's sensitivity to hot or cold. The fact that the two people in our example have different sensitivities to hot or cold does not entail that at least one of them is misrepresenting the temperature of the room. Rather, there are many ways of correctly representing the same temperature. On the attention-indexed view, there are many more properties that can be handled in this fashion than previously thought. Perception research has revealed that many properties we previously thought were objective are in fact narcissistic. Accurately perceiving those properties depends on the way the world is relative to a subjects' needs and interests. Like temperature, accurate representation of visual contrast depends on the state of the subject, namely the direction of her attention, as well as the nature of the world. This shows that a seemingly objective property — contrast — is in fact a narcissistic property, and should be treated differently than objective properties with regards to veridicality conditions.

Akins considers this the "most promising" view for a proponent of the traditional account of perceptual content and veridicality. On this view, narcissistic systems do indeed detect properties, but they are not objective properties; rather, they are properties defined in relation to the subject. We can define a narcissistic property to correspond to each narcissistic encoding utilized by a sensory system, and thereby provide an account of how narcissistic systems are veridical in a revised sense. Unlike a traditional account, the veridicality conditions for a narcissistic property will specify more than just states of the world. They will also specify the relevant states of the subject, such as the direction of attention.

Akins' worry is that this attempt to reconcile narcissistic properties with veridicality is implausible because it is not demanded by a neurological description of narcissistic sensory systems. Rather, she thinks that what motivates the attempt to save the traditional view is a commitment to a view that she calls the *detection thesis.* The detection thesis states that *"each and every* sensory system functions to detect properties" (Akins 360, emphasis mine). But we don't need to accept a strong universal claim like the detection thesis to motivate a narcissistic view of properties. Instead, we can point to the need for an account of perceptual content that fits with a representationalist account of the mind, on which perceptual states have veridicality conditions. This is not to say that every sensory system detects properties. Rather, it is to acknowledge that the representationalist paradigm for theorizing about the mind has been tremendously fruitful; moreover, there is a way to reconcile the view that sensory systems are narcissistic with the representationalist's commitment to the view that perceptual systems are veridical. The evidence that the visual system is narcissistic rather than "veridical" in the traditional sense does not challenge this revised notion of veridicality, on which correctness conditions depend on states of the perceiver as well as objects and properties in the world.

3. A Defense of Representationalism

As I mentioned in the introduction, the evidence that attention affects appearance has received considerable philosophical discussion because it seems to raise a new counterexample to one of the dominant accounts of perception: representationalism. Recall that representationalism commits its proponents to the view that there can be no change in phenomenal content without a change in representational content. But it is controversial whether the representationalist can

identify a representational change to accompany each phenomenal change due to attention. The attention-indexed view helps the representationalist to respond to the challenge raised by attention.

The challenge from attention has taken a number of forms in the literature. Consider again the two lights example from Chalmers (2004; see also a version of this challenge posed by Nickel 2007). Though the phenomenal content of perceptual experience changes in a subtle way as attention shifts from a light on the left to a light on the right, it is not obvious that switching attention changes the representational content of that experience (see **Figure 1**). Of course, these examples raise a problem for representationalism only if there is no change in representational content that explains the change in phenomenal content. The reason for thinking that representational content does not change is that when I introspect on my experience of shifting attention, a change in representational content is not obvious to me. But there is a compelling representationalist response to this version of the challenge raised by attention: failing to notice a change in representational content provides a weak reason to think that the representational content of my experience in fact stays the same. Introspection is notoriously fallible, and some representational differences may be subtle and difficult to notice. When introspecting on experience, subjects may simply fail to notice a change in the representational content of their experience.

Although Chalmers (2004) thinks that shifts of attention provide the most compelling counterexample to the claim that phenomenal content supervenes on representational content, he agrees that the challenge as currently formulated is unsuccessful. Instead, he suggests that we should provisionally accept the view that all attentional differences are representational differences as a null hypothesis. That a subject is unable to identify a difference in representational

content through introspection is insufficient for generating a compelling challenge to representationalism.

A stronger counterexample to representationalism is raised by Block (2010) using Carrasco's experiments on perceived contrast. At first glance, a representationalist seems to be able to account for the results of Carrasco's experiment by maintaining that covertly shifting attention to a Gabor patch increases the contrast represented. For example, consider just the Gabor patch on the left back in **Figure 2**. The representationalist could say that when I covertly shift attention to the left Gabor patch, I represent it as 28% contrast, and when I covertly shift attention away, I represent it as 22%. Carrasco's study only challenges representationalism if it demonstrates a change in phenomenal content without a change in representational content; but the representationalist seems to be able to identify a change in representational content that accompanies this phenomenal change.

Block (2010) has given an argument for why this response won't work. He considers the view that a shift of attention changes the degree of contrast represented, and argues that a representationalist who embraces this strategy will have two options. She can either concede that at most one of the representations is an instance of veridical perception, or claim that the content of perception is a range of indeterminacy. Consider the first option. In Carrasco's experiment, subjects are misrepresenting the patch on the left either when they attend to it, or when they attend away. After all, though the representational content changes, the world itself stays the same. According to Block, the representationalist will run into trouble because there is no distribution of attention that engenders veridical perception. It is highly implausible that attending to a thing induces an illusion; but likewise, it's implausible that attending to something renders all back-

ground experiences illusory. According to Block, the problem is not merely epistemic –it is not that we can't know when we're representing correctly, and when we aren't. Rather, there is no particular way of attending that will make perception veridical.

According to Block, the second option for the representationalist is to hold that representational content is a range of indeterminacy that includes the actual contrast. Focusing just on the perception of the left Gabor patch, as attention shifts from the dot at the center to the Gabor patch on the left, representational content may change to a different range. Both represented ranges can be veridical, because both contain the actual contrast. This looks at first like a solution to the challenge attention raises to the supervenience of phenomenal content on representational content. But the representationalist isn't only interested in satisfying supervenience. Her primary aim is to provide an account of phenomenal content in representationalist terms. And for this purpose, Block thinks that this solution is inadequate. The reason is that the representational content "clashes with" the phenomenal content of perceptual experience. When I look at the Gabor patch, my visual experience is of a particular, determinate contrast, and not a range of indeterminacy. A range is indeterminate in a way that the content of my visual experience is not. Block thinks this shows that representational content can't be what explains why my perceptual experience has the particular phenomenal content that it does. More strongly, he claims that phenomenal content is actually incompatible with representational content: phenomenal content is determinate, but representational content is indeterminate.

Block's argument takes for granted two plausible assumptions: the first is that veridical representation must involve a match between representational content and the objective property represented. That is, if I veridically represent a patch that has an objective contrast of 22%, then

either I represent it as 22% or as a range that includes 22%. The second assumption is that phenomenal content is determinate. While there have been several attempts to defend representationalism against Block's challenge by rejecting his second assumption about determinacy (Nanay 2010, 2011; Stazicker 2011)[3] or by embracing widespread error (Watzl forthcoming), this paper challenges the first assumption: that veridicality requires a match between the property represented and the representing property. The attention-indexed view involves a new way of thinking about the conditions for veridical perception, which puts us in a position to respond to the challenge raised by perceived contrast.

3.1 The Attention-Indexing Response

I will first explain how the attention-indexed view can provide a response to the challenge raised by attention's effect on perceived contrast. In brief, the attention-indexed view would tell us that there are many ways of correctly representing the Gabor patch, depending on the direction of a subject's attention.[4] As attention shifts covertly toward and away from the Gabor patch on the left in **Figure 2**, representational content changes from a representation of (say) 28% contrast to a representation of 22% contrast. Although representational content changes, and

[3] A third option is to defend a version of impure representationalism. Ganson & Bronner (2013) have explored one version of impure representationalism which they call *quasi-representationalism*. A quasi-representationalist accounts for the effect of attention on appearance by introducing the notion of *prominence*. Prominence is a representational fact, but it is not a fact about representational content. My goal in this paper is to give an account of phenomenal change due to attention in terms of representational content, so I will not address their argument here.

[4] Block agrees that there are many ways of correctly representing the Gabor patch, but he means something different than what I mean. On Block's view, in order for two distinct representational contents of an objectively 22% patch to be veridical, the representational content on both occasions must include the actual contrast: 22%. On my view, distinct and contradictory contrast amounts can be veridical. The 22% patch could be accurately represented by 28%-attended and 22%-unattended, for example.

the world stays the same, the subject is not misrepresenting in either condition. For narcissistic properties, accurate representation depends on more than just the way the world is. It also depends on relevant states of the perceiver, in this case, the direction of attention. When unattended, I perceive the Gabor patch correctly when I represent it as an unattended-22% contrast, and I also perceive correctly when I represent it as attended-28% contrast.

The attention-indexed view doesn't deny that there are objective properties, but it does imply that perception isn't tracking them in the case of contrast. The Gabor patch has an objective contrast set by the experimenter, just as the room has an objective temperature set using the thermostat. In both cases there is a property that is had by an object, and a subject who represents that property relative to her internal states. In the temperature case, the subject represents the objective temperature (60F) relative to her internal state (the temperature of her skin having just stepped in from the cold, as well as the distribution of cold receptors). In the case of contrast, the subject represents the objective contrast (22%) relative to attention. For subjects that are attending in different ways, the objectively same contrast patch will be represented differently, just as for subjects with different internal states, the objectively same temperature will be represented differently. Accurately representing a 22% contrast patch involves representing the correct property for that contrast in a particular attentional context. As a result, the veridicality conditions for a representational content are indexed to attention.

On the attention-indexed view, we represent an object as having a different property when attended as compared to unattended. As Watzl (forthcoming) points out, a view like the attention-indexed view will face a challenge arising from perceptual constancies in vision. When we shift visual attention around a scene, it doesn't seem like the world around us is changing.

Unlike temperature, properties like contrast, size, shape, and color will look invariant across many changes in the environment and the subject, such as illumination and viewing distances (Georgeson & Sullivan 1975). Evidence for perceptual constancies is robust, and any adequate account of visual perception must take this evidence into account (Cohen 2015; Brainard et al. 2003; Bryne & Hilbert 2007). Since the properties represented in vision are invariant across different conditions of the subject, we have reason to think that they are indexed to the world and not to states of the subject. According to this objection to my view, the evidence for constancy in vision challenges my claim that many of the properties represented in vision are narcissistic rather than objective.

There is a simple response to the evidence for perceptual constancies (which Watzl acknowledges). While visual properties like contrast, size, shape, and color saturation may exhibit constancy across illumination and viewing distances, they don't exhibit constancy across attentional contexts. This is precisely what Carrasco's experiment shows for the perception of contrast: attention alters perception of contrast, even if illumination and viewing distance do not. From the fact that a property is invariant across illumination and viewing distance, we should not conclude that that property is objective rather than narcissistic. A narcissistic property may be invariant across certain sorts of changes in the subject — like their distance from the object — while nonetheless depending on other states of that subject — like the direction of their attention.

The attention-indexed view has consequences for how we understand perceptual constancy and its implications. It challenges the traditional distinction between properties that are narcissistic (like taste and temperature) and objective (like shape or contrast). If seemingly objective properties turn out to be indexed to states of the subject, this also raises new worries for how to

demarcate error from accuracy. Accurate perception will involve attributing seemingly incompatible properties to an object, like 22% and 28% contrast. These seemingly incompatible properties are in fact compatible once they are indexed to a direction of attention.

3.2 Widespread Error

In the cases that I have used to motivate the attention-indexed view, there is a simple alternative: subjects could be misrepresenting. According to this alternative, when two subjects disagree on the temperature of the room, at least one of them is misrepresenting. In Carrasco's experiment, subjects either misrepresent the Gabor patch when they attend to it, or when they don't attend to it, or both. On Watzl's version of this response, perception is a compromise between true representation and useful representation (Watzl forthcoming; see also Watzl 2014). He argues that perception and attention have distinct functions, which can conflict. The function of perception is to provide an accurate representation of the world, whereas the function of attention is to provide a usable representation. Attention serves to prioritize certain perceptual representations over others, and thereby organizes perceptual information in a way that makes the subject better able to exploit that information for use in guiding thought and action (for more on the notion of perceptual guidance, see Watzl 2014). But there are times when the function of perception is at odds with the function of attention. The most useful representation may not always be the most accurate representation. In cases where the function of perception and the function of attention are at odds, the subject will be in a state that is a compromise between the two. She will be in a perceptual state that is "close enough to being accurate while still being as useful as possible" (Watzl forthcoming p. 20).

Watzl's view has the implication that perception is very often inaccurate, a consequence that he accepts as a solution to the challenge to representationalism. As Watzl argues, in order to accurately represent a property — like the Gabor patch in Carrasco's experiment — we would have to attend in a way that optimizes the function of perception. This would be to ignore usability in favor of accuracy. Outside of highly idealized scenarios, usability matters. That is, the function of perception (accuracy) will almost always be balanced by the function of attention (usability), such that we rarely if ever attend in the way that would lead to accurate perception.

Although his view entails that we frequently misrepresent the world in perception, Watzl argues that this widespread misrepresentation is not problematic. One of the main concerns with accepting widespread misrepresentation is that it would seem to sacrifice the representationalist's account of accurate perception. Since Watzl is able to give such an account in terms of the function of perception, he thinks that the widespread misrepresentation predicted by his view is not problematic. Further, the sort of widespread misrepresentation that Watzl's view predicts involves relatively small distortions from accurate perception. For these two reasons, he thinks that a representationalist can accept widespread misrepresentation arising from a conflict between the function of attention and the function of perception.

Watzl's response may work for some versions of representationalism, but it will arguably cause problems for tracking versions of representationalism (as defended by Tye 2000; Dretske 1995; Millikan 1989 and Fodor 1987, among others). Tracking theories of representation hold that representational states carry information about the environment in virtue of correlating with objects or properties in that environment. This version of representationalism will meet with difficulty if misrepresentation turns out to be widespread and systematic, as it is on Watzl's view. To

see why, consider Mendelovici's (2012) distinction between reliability and veridicality. Reliable misrepresentation involves a similar response in similar circumstances, even though the world isn't really the way it is represented to be. Watzl's view seems to commits him to reliable misrepresentation: given the same distribution of attention and circumstances in the world, a subject will have to make the same compromise between the function of attention and the function of perception, and so they should misrepresent in the same way. Although a property is being tracked — a state reliably correlates with it — that state is nonetheless a misrepresentation. As Mendelovici puts it, "a reliable misrepresentation tracks *something*, but what it tracks is not what it misrepresents" (Mendelovici 2012, p. 4). This means that a representationalist who is committed to a tracking theory of representation would have trouble accepting Watzl's response to the challenge from attention. It is worth exploring whether Watzl's view can overcome these problems raised by reliable misrepresentation.

An alternative to widespread misrepresentation (and the problems for tracking theories of representation that come along with it) is to accept the attention-indexed view. On that view, we avoid the consequence that error is widespread, as well as related worries for tracking theories. Indexing veridicality conditions to attention allows us to account for the perceptual effects of attention without sacrificing the view that perception is ordinarily accurate.

3.3 Indeterminacy

The second alternative for responding to the challenge raised by attention is to question Block's assumption that indeterminacy is inconsistent with visual phenomenology. Several authors (including Nanay 2010, 2011 and Stazicker 2011) have argued for this response, defending

the view that indeterminacy is widespread in vision. Attention, on this view, enables us to represent an object in more detail or more determinately. I'll call the view that attention makes the content of perception more determinate the *determinacy view.* While the determinacy view offers a promising solution to some of the challenges raised by attention, other effects of attention on perception cannot be adequately characterized in terms of determinacy. This provides us with at least one reason to prefer the attention-indexed view.

The determinacy view takes several forms in the literature. Nanay has argued that the properties of an attended object are represented at a more determinate level than the properties of an unattended object (2010; 2011; see also Jagnow 2011). This view runs into a problem: we can and do sometimes attend to an object with the least determinate properties, or fail to attend to the object with the most determinate properties (as Wu 2011 has also pointed out). A thought experiment helps to illustrate this possibility. Suppose we were to construct a pair of glasses that mimic a spotlight of attention, with higher resolution in the center, fading to lower resolution in the periphery.[5] If you were to wear these glasses, you could certainly attend to the center, where properties are more detailed and determinate. But you could also shift attention to the periphery, where properties are more indeterminate. If you shift attention to the periphery, the properties of attended objects would be represented at a less determinate level than the properties of unattended objects. This shows that attended objects can sometimes be represented at a less determinate level than unattended objects. A stronger view, which may be closer to what Nanay has in mind, is that attention makes an object appear more determinate than it would appear if unattended.

---

[5] This example originates in Watzl (2013), though he uses it to illustrate a different point.

The problem for the determinacy view arises when accounting for the point of subjective equivalence in Carrasco's experiment, that is, that point at which the two stimuli appear to have the same degree of contrast. In some trials, subjects' performance on the task suggests that the attended stimulus appears indistinguishable from the unattended one. For example, when cued to covertly attend to the Gabor patch on the left in **Figure 2**, subjects behave as though the patches appear to have the same contrast. But if one property is represented at a determinable level and the other at a more determinate level, then this result is surprising. Consider an example involving color. Suppose that a thing could appear to you to be merely colored, without appearing any particular color. This would be perceptually distinguishable from something appearing *red.* It would be very surprising if a subject behaved as though a merely colored object and red object were indistinguishable with respect to color. Instead, it is plausible that properties represented at a more determinable level would be distinguishable from properties represented at a more determinate level. The evidence that subjects have a point of subjective equivalence, at which the two patches appear the same to them, gives us reason to think that the patches are represented at the same level of determinacy.

Attention has a wide variety of effects on perceptual experience, and some of these effects cannot be adequately explained in terms of determinacy. Even if attention sometimes has determinacy effects, the range of cases suggest that an appeal to determinacy will not suffice to explain attention's effects on perception. Following Carrasco, Ling & Read's (2004) study on perceived contrast, Gobell & Carrasco (2004) used a similar paradigm to investigate attention's effect on perceived gap size. They found that covertly attending to a figure increased the apparent size of a gap in that figure. In place of Gabor patches, subjects were presented with two

squares on either side of a central cross, each containing a gap. The size of the gap was measured by degree of visual angle, and was fixed by the experimenter. When cued to attend covertly to one of the gaps, the attended gap was objectively smaller than the unattended gap at the point of subjective equivalence. Gobell and Carrasco interpret this as showing that attention increases perceived gap size. This result is difficult for the determinacy view to explain. The attended gap did not appear more determinately sized. Subjects were not gaining more detail about the size of the gap. Instead, Carrasco and colleagues found that attending made the gap appear larger.

A revised version of the determinacy view can potentially handle these results. On a revised determinacy view, attending changes the range of properties represented, but it doesn't always involve a change to a more determinate level. Stazicker (2011, unpublished presentation) considers an example from Gobell & Carrasco's (2004) experiment in which a subject is presented with two stimuli: a peripherally cued test gap that is objectively 0.20 degrees, and an uncued gap of 0.23 degrees. On his view, attending to an object changes the range of values represented. For example, he suggests that both the covertly attended 0.20 degree gap and the uncued 0.23 degree gap might be represented as a range of indeterminacy between 0.19 and 0.27 degrees. Both are veridical because both ranges contain the objective contrast of their respective object; but we can nonetheless explain why subjects behave as though the two stimuli appear indistinguishable to them. As Stazicker notes, it is an open empirical question whether representational contents actually do span such a large range of indeterminacy. If the range is narrower than Stazicker suggests, he would lose the consequence that both contents are veridical. Stazicker's account would then leave us with the problem for veridicality raised by Block, while taking on the controversial assumption that vision is often indeterminate.

However, the larger point in this section is that the determinacy view struggles to accommodate the range of effects that attention has on perception. Given the wide range of attentional effects on perception, it is plausible to conclude that the effect of attention on the content of perception will be specific to the type of stimulus and sense modality involved. While we might be able to identify some representational changes by appealing to determinacy, many of these cases leave the phenomenal change unexplained or mysterious. In addition to contrast and gap size, attention has been shown to affect the perceived timing of stimulus presentation (Spence and Parisse 2009; see also Stelmach and Herdman 1991), and to make attended objects appear more saturated in color (Fuller, Ling & Carrasco 2004) or to have greater spatial frequency (Abrams, Burbot & Carrasco 2010). As Block summarizes, "attended items look bigger, faster, earlier, more saturated, stripier" (Block 2010, p. 41). There will likely be no simple answer to the question of how attention affects perception. The attention-indexed view allows for the possibility that attention has many different effects on the content of perception. Attention changes perceptual content in many diverse but specific ways, the nature of which is open to empirical investigation (see also Hill 2009). My view thus gives us a theoretical framework for making sense of a range of empirical findings.

4. Conclusion

I have argued that attention doesn't provide a counterexample to pure representationalism. Rather, the moral that we should draw from the research on attention and perception is that any adequate theory of perceptual content must make room for the role of attention. Accepting this moral motivates a revision of representationalist notions of content and veridicality. To accom-

modate the effects of attention on perception, we must revise the notion of veridicality to make room for the direction of a subject's attention, in addition to objects and properties in the world. The attention-indexed view of perceptual content that I have presented in this paper provides a framework for exploring how representationalism must change in response to the emerging empirical picture of attention's perceptual effects. This research points to a new view of perceptual content, and the pivotal role that attention must play in developing an adequate account of how we represent the world in perceptual experience.[6]

Works Cited

Abrams, J, Barbot, A. & Carrasco, M. (2010) "Voluntary Attention Increases Perceived Spatial Frequency." Atten Percept Psychophys. 2010 Aug; 72(6): 1510–1521.

Akins, K. (1996) "Of Sensory Systems and the 'Aboutness' of Mental States." *Journal of Philosophy.* 93: 337—372.

Armstrong (1978) *Nominalism and Realism, Vol. II.* New York: Cambridge University Press.

Blaser, E., Sperling, G., & Lu, Z.-L. (1999). Measuring the amplification of attention. Proceedings of the National Academy of Sciences, USA, 96, 11681-11686.

Block, N. (2010) "Attention and Mental Paint." *Philosophical Issues.* 20:23—63.

Block, R.A., George, E.J. and Reed, M.A. (1980) "A watched pot sometimes boils: A study of duration experience." *Acta Psychologica*. 46:81—94.

Brainard, D H, Kraft, J M and Longere, P. (2003). Color constancy:Developing empirical tests of computational models. In R. Mausfeld and D. Heyer, editors, Colour Perception: Mind and the Physical World, pages 307–328. Oxford University Press, New York.

Byrne, A. (2001) "Intentionalism Defended." *The Philosophical Review*. 110:199—240.

Byrne, A. and Hilbert, D. R., editors (1997). Readings on Color, Volume 2: The Science of Color. MIT Press, Cambridge, Massachusetts.

Carrasco, M., Ling, S., & Read, S. (2004) "Attention alters appearance." *Nature Neuroscience:* 7:308—313.

---

[6] My sincere thanks to Evan Thompson, whose support has been invaluable in developing the views in this paper. I would additionally like to thank the anonymous reviewers, as well as Ned Block, Kevin Connolly, Carolyn Dicey-Jennings, Todd Ganson, Benj Hellie, Mohan Matthen, and Bill Seager, whose comments and discussion were immensely helpful. I am indebted to the audiences at the 2011 Harvard-Columbia graduate philosophy conference, the 2011 NYU-Columbia graduate philosophy conference, the 2011 Pacific APA, the 2011 Eastern APA, the 2011 Interdisciplinary Conference on Consciousness at Boston University, and the 2012 Yale-UConn graduate conference. I owe special thanks to those people who served as commenters: Alex Bryne, Anya Farennikova, Bill Fish, Nemira Gasiunas, Aaron Norby, and Sebastian Watzl.

Chalmers (2004) "The Representational Character of Experience." in Leiter, B. (Ed) *The Future for Philosophy*: Oxford University Press.

Chalmers, D. (2004) "The representational character of experience." *The Future for Philosophy*. B. Leiter. Oxford, Oxford University Press: 153—181.

Cohen, (2015) "Perceptual Constancy," in Matthew (Ed.) The Oxford Handbook of Philosophy of Perception. Oxford University Press.

Connolly, K. (2014) "Perceptual Learning and the Contents of Perception." Erkenntnis 79: 1407-1418.

Dretske (1995) *Naturalizing the Mind*. Cambridge: MIT Press.

Fodor, J. (1987) *Psychosemantics.* Cambridge: MIT Press.

Fuller, S., Ling, S., Carrasco, M.(2004). Attention increases perceived saturation [Abstract]. Journal of Vision, 4(8): 329.

Gobell, J. & Carrasco, M. (2004) "Attention alters the appearance of spatial frequency and gap size." Psychological Science. 16(8): 644—51.

Ganson & Bronner (2013) "Visual Prominence and Representationalism" *Philosophical Studies* 164 (2):405-418.

Hicks, R.E., Miller, G.W., and Kinsbourne, M. (1976) "Propsective and retrospective judgments of time as a function of amount of information processed." *The American Journal of Psychology*. 89: 719—730.

Millikan, R. (1989) "Biosemantics" *Journal of Philosophy*. 86:281—297.

Nanay, B. (2010) "Attention and Perceptual Content." *Analysis*. 70: 263—270.

Nickel (2007). Against intentionalism. Philosophical Studies 136, 279-304.

Speaks, J. (2011) "Attention and Intentionalism" *The Philosophical Quarterly.* 60:325—342.

Spence, C. & Parise, C. (2009) "Prior-entry: A review." Consciousness and Cognition.

Stazicker, J. (2011) "Attention, Visual Consciousness, and Indeterminacy." *Mind & Language*. 26:2 156—184.

Stazicker, J. (2011) "Attention, Visual Knowledge and Psychophysics: Discriminating the Determinable" Talk given at the NYU Consciousness Project; draft available at https://philosophy.berkeley.edu/file/575/Attention__Visual_Knowledge___Psychophysics_web.pdf

Stelmach, L. B. & Herdman, C.M. (1991) "Directed Attention and Perception of Temporal Order." Journal of Experimental Psychology: Human Perception and Performance. 17: 539—550.

Tye (2000) *Consciousness, Color and Content*. Cambridge: MIT Press.

Watzl (forthcoming) "Can Intentionalism Explain how Attention Affects Appearance?" in Pautz, A. & Stoljar, D. (Eds.) *Themes from Block*: Cambridge, MIT Press. (http://folk.uio.no/sebaswat/materials/Watzl_Intentionalism_Appearances.pdf)

Watzl (2014) "Attentional Organization and the Unity of Consciousness." *Journal of Consciousness Studies*. 21:56—87.

Watzl (2014) "Perceptual Guidance" Ratio, 27(4): 369–505

Wu, W. (2011) "What is Conscious Attention?" *Philosophy and Phenomenological Research*. 82:93—120.