Received: 13 January 2020 | Accepted: 13 July 2020

DOI: 10.1111/ejop.12582

ORIGINAL ARTICLE

European Journal of Philosophy WILEY

Thinking through illusion

Dominic Alford-Duguid

Faculty of Philosophy, University of Oxford, Oxford, United Kingdom

Correspondence
Dominic Alford-Duguid, Faculty of Philosophy, University of Oxford, Radcliffe Humanities Building, Oxford, OX1 2JD, United Kingdom.
Email: dominic.alford.duguid@gmail.com

Funding information
Leverhulme Trust, Grant/Award Number: ECF-2018-40

Abstract

Perception of a property (e.g., a colour, a shape, and a size) can enable thought about the property, while at the same time misleading the subject as to what the property is like. This long-overlooked claim parallels a more familiar observation concerning perception-based thought about objects, namely that perception can enable a subject to think about an object while at the same time misleading her as to what the object is like. I defend the overlooked claim, and then use it to generate a challenge for a standard way of thinking about the relationship between visual experience and rational belief formation. Put informally, that view holds that just as we can mislead others by saying something false, illusory experience misleads by misrepresenting how things stand in the world. I argue that we ought to abandon this view in favour of some radical alternative account of the relationship between visual experience and rational belief formation.

In this article, I defend an overlooked claim concerning perception-based thought about properties, and then use it to generate a challenge for a standard way of thinking about the relationship between visual experience and rational belief formation. The claim in question parallels a familiar observation concerning perception-based thought about objects, namely that perception can enable a subject to think about an object while at the same time misleading her as to what the object is like. The overlooked claim for the case of properties is that perception of a property (e.g., a colour, a shape, a size) can enable thought about it, while at the same time misleading the subject as to what it is like.

My eventual target is the view that, just as we can mislead others by saying something false, illusory experience misleads by misrepresenting how things stand in the world. For example, suppose you see a vase, and in response come to believe that the vase is oval. If the vase is round, and merely looks oval, then you have been misled. The target view says that in this case, you have been misled because your visual experience misrepresents the vase as oval.

What of the vase's actual shape (i.e., its roundness)? It seems to have disappeared. This apparent consequence becomes problematic when combined with my overlooked claim concerning perception-based thought about properties. For if you can still think about the vase's actual shape in such a case, and do so on the basis of your visual experience, the shape cannot in fact have disappeared. I turn this intuitive tension into a genuine challenge.

The paper has three parts. The first (Section 1) defends a precise version of the overlooked claim concerning perception-based thought about properties, and will be of independent interest to those working on perception-based thought. The second (Sections 2 and 3) unpacks the target view and develops my challenge. The final part consolidates the challenge: in Section 4 I establish a crucial premiss, and in Section 5 I show that if we reject the target view, we ought to pursue a radical alternative account of the relationship between visual experience and rational belief formation.

1 | THE TOLERANCE OF PERCEPTUAL ASCRIPTION

Suppose a ball of red yarn rolls into view, followed by your curious cat. This perceptual encounter puts you in a position both to think about the ball and to think about redness, each in a particular way. Call these ways of thinking (of the ball of yarn on the one hand, and redness on the other) "perceptual demonstration" and "perceptual ascription". These ways of thinking are made available only by a perceptual link with the object or property they pick out.[1] We standardly express thoughts that involve perceptual demonstration with sentences like "That object is red" or "That is round"; thoughts that involve perceptual ascription we standardly express with sentences like "The ball is that colour", "It's that shape", or simply "It's like this". Perceptual ascription of properties thus functions as the analogue for properties of perceptual demonstration of objects. (Some will insist that what I call "perceptual ascription" just is the perceptual demonstration of properties; I wish to not prejudge the exact relationship between these two perception-based ways of thinking.)

Philosophers have long supposed that a subject can perceptually demonstrate an object even when the beliefs she rationally forms about the object on the basis of her perceptual link with it are almost all false.[2] Call this mark of perceptual demonstration "tolerance". For instance, you might see a large red chair, but mistake it for something small. Yet these errors do not undermine your ability to think about the chair: beliefs you rationally form in response to perception, and express with sentences such as "That is small", will still involve successful perceptual demonstration of the chair.

What nobody seems to have recognised is that perceptual ascription is also tolerant.[3] A subject can be rationally misled about an observable property on the basis of her visual experience, yet nevertheless remain in a position to rationally form beliefs that perceptually ascribe the property. In more precise terms, the claim I shall defend runs as follows:

> TOLERANCE OF PERCEPTUAL ASCRIPTION:There are incompatible properties F and G (i.e. no object can have them at the same time in the same respect) such that cases with the following structure are possible: (a) It is rational for a subject S to believe on the basis of her visual experience of an object (which is in fact F) that the object is G (e.g., *it's oval*); and (b) at the same time, it is rational for S (on the basis of that same experience) to form a belief about the object that perceptually ascribes F (e.g., *it's that shape* [i.e., round]).

Cases with the required structure are "tolerance cases". In order to bring out a sample tolerance case, first consider the relationship between properties, classification, and misleading perception. We use properties to group or classify objects. For example, when presented with a collection of objects, and asked to group them by shape, we group like with like: the square objects in one group; the round objects in another; and so on. When perception

misleads us about an object, and we are asked to group the object with others like it, we misclassify it: if a round object looks oval, we group it with oval objects; if a red object looks orange, we group it with orange objects; and so on.

A parallel story applies to properties. We use *higher-order* properties to classify the properties of objects. For example, if we sort a collection of colours according to their hues (hue being a higher-order property of colours),[4] we group two colours together when they have the same hue. As a result, when perception misleads us about hue, we misclassify colours: if a red object looks orange, we classify its colour as a shade of orange rather than as a shade of red.

With the background connection between properties, classification, and misleading perception in place, here's our flagship tolerance case:

> ROUND VASE: Your visual experience misleads you about the shape of a round vase, so that you rationally but mistakenly take the vase to be oval. (I follow common usage, here and throughout, in assuming that being oval and being round are incompatible properties.) Hence when asked to classify the vase's shape, you group the vase with oval objects rather than other round objects.

Does your classificatory error undermine your capacity to perceptually ascribe the vase's actual shape? I think not. "That shape looks more oval than round", you might initially say, only to later revise your view when you learn that your perspective was strange, and insist "It actually wasn't the way it looked". So the initial "that shape"—which captures your original perceptual ascription—plausibly picks out the actual shape (i.e., round), not the merely apparent shape (i.e., oval).[5] And you could have used this capacity to perceptually ascribe the vase's actual shape when you rationally formed beliefs in response to what perception delivered. For instance, on the basis of your visual experience you might have rationally formed beliefs about the sort of behaviour that possession of *that shape* permits the vase to engage in (e.g., rolling down smooth inclines). Putting all this together: your visual experience of the vase misleads you into thinking a round vase is oval (satisfying [a] from TOLERANCE OF PERCEPTUAL ASCRIPTION), but at the same time permits you to rationally form beliefs such as *the vase is that shape* that perceptually ascribe the vase's actual shape (satisfying [b]).

Other cases like ROUND VASE are not difficult to construct. However, readers typically have one of three reactions to my claim that these are tolerance cases: many are immediately convinced; others mildly sceptical; and a few violently opposed. If you are in the first group, you can skip ahead to Section 1.3 where I introduce limits on the tolerance of perceptual ascription. If you belong to one of the latter groups, you worry that in ROUND VASE (for example) either the subject does not in fact perceive the vase's actual shape (and for that reason cannot perceptually ascribe the shape) or the subject cannot perceptually ascribe the vase's shape (despite somehow perceiving it). I shall use the next two subsections to address each worry in turn. My aim is avowedly modest: I seek to convince only the mildly sceptical reader. Yet even committed opponents of TOLERANCE OF PERCEPTUAL ASCRIPTION should admit, given what I say below, that their position requires more active defence than they frequently assume.

1.1 | No perceptual link?

Again, some opponents will object that any case that satisfies [a] from TOLERANCE OF PERCEPTUAL ASCRIPTION cannot also be one in which the subject maintains a perceptual link with the object's relevant property (i.e., F). If they are right, [a] and [b] cannot be jointly satisfied, since perceptual ascription requires a perceptual link with the ascribed property.

I cannot develop a decisive answer to the objection at this juncture.[6] But since my present target is the mildly sceptical reader, rather than the committed opponent, I can instead point to the work of others. Macpherson and Batty (2016) argue at length that cases of so-called "property illusion" are very much possible. In our terms, these are

cases in which a subject's perceptual experience of an object provides a perceptual link with one of the object's properties, and yet at the same time makes it rational for a subject to believe (on the basis of her experience) that the object possesses an incompatible property. Macpherson and Batty argue that if there is appropriate[7] counterfactual dependence between a subject's perceptual experience and a given property, and the subject's current experience puts her in a position to know facts about the property, then her perceptual experience provides what I have called a "perceptual link" with the property in question. Given this background, here is Macpherson and Batty's flagship case of property illusion:

> [Suppose you] experience objects and all of their visible properties accurately except for the fact that which colour you experience them as having is systematically skewed. You experience objects as the same amount slightly lighter, the same amount slightly more saturated … Again, suppose that, in these circumstances, you are looking at a blue car and experience it accurately except for the fact that you experience it as light blue when it is in fact dark blue. We think that this is a case of object perception, but property illusion. … [Y]ou are sensitive to the colours of things—just not in a way that gives you veridical experiences of their colours. Your experience is counterfactually dependent on the colours of the object before you, but it is systematically skewed, so that there is never a perfect match between your experience and the colours of objects in the world. Still, there is a close match. For, in addition to your counterfactual sensitivity, you can come to know some facts about the colours of objects solely based on your experience—such as which objects are lighter and darker… (p. 281; emphasis mine)

As the underlined portions make clear, Macpherson and Batty take their conditions for the existence of a perceptual link with the car's actual colour to be satisfied. I shall not go into more detail about their arguments. While ROUND VASE was not as complex as the cases Macpherson and Batty develop, it (and other putative tolerance cases) could be unpacked so as to meet Macpherson and Batty's conditions for property illusion. At this point, the mildly sceptical reader should abandon their worry that in ROUND VASE the subject does not perceive the object's actual shape, and committed opponents should recognise that pressing the worry requires an answer to Macpherson and Batty.

1.2 | Perception without perceptual ascription?

Most will find it simply obvious that it would be rational for the subject in ROUND VASE to believe (for example) *nothing else around here is that shape*. The second worry raised by mild sceptics and committed opponents alike is that this belief—and any other like it—fails to perceptually ascribe the vase's actual shape (even assuming that the subject in ROUND VASE maintains a perceptual link with the shape).

In response to this worry, I shall argue that the subject's perception-based beliefs in ROUND VASE behave just as we would expect if TOLERANCE OF PERCEPTUAL ASCRIPTION were true. When combined with the assumption that the subject maintains a perceptual link with her vase's shape, this result is strong evidence for the claim that in ROUND VASE (and analogous cases), a subject can perceptually ascribe the property about which he is misled. Again, I do not expect that this evidence will convince a committed opponent, but it should suffice to sway the mildly sceptical reader.

My argument relies upon a parallel between the perceptual demonstration of objects and the perceptual ascription of observable properties. The parallel concerns *when* aboutness-fixing works for perception-based thought. Consider an ordinary case of perceptual demonstration. You come across an object. Your perceptual attention focusses on it. This attentional focus results in an influx of perceptual information about the object—where it is, what properties it possesses, what changes it undergoes.[8] Much of this information you take at face value when forming beliefs.

This growing collection of beliefs evolves in response to changes in the nature and structure of the incoming information.

When do these "perceptual demonstrative beliefs"—which you would express using sentences like "That is red" or "That is round"—count as being about the object to which you attend? An old and attractive answer is that your perceptual link with the object must govern your dispositions to form and maintain these beliefs.[9] These dispositions must organise your beliefs in response to a perceptual link with the object they purport to be about. They must also preserve a degree of internal rational coherence (e.g., no obviously inconsistent beliefs permitted). Most importantly, the dispositions must treat the perceptual link as *authoritative*: incoming perceptual information will usually trump other putative sources of information about an object.

Note that the previous paragraph addresses only the question of *when* perceptual demonstrative beliefs count as being about an object. It leaves open the much harder question of *how* a perceptual link with an object does its aboutness-fixing work. Fortunately, nothing I want to say in this article requires taking a stand on how this harder question should be answered.[10]

I now turn once again from perception-based thought about objects to perception-based thought about properties. Whatever theorists eventually say about *how* aboutness-fixing works for these two types of thought, the stories about *when* it works plausibly remain almost exactly the same. Consider someone in a variant of the ROUND VASE case. His visual experience misleads him about the shape of a round vase, so that he mistakenly takes the vase to be oval. Yet his perception-based beliefs about the vase's shape will still be rationally formed in response to information delivered by his attentional link with that shape.[11] And the dispositions that lead him to form beliefs in this fashion not only treat his perceptual link to the shape as authoritative, but also seek to preserve a degree of internal coherence amongst the resulting beliefs about the vase's shape (these beliefs will be about the shape's higher order properties—a departure from the perceptual demonstrative case: those beliefs were about properties of the perceived object).

To see these dispositions at work, suppose our subject sees the vase roll, and as a result comes to (rationally) believe that *that shape* endows objects with the ability to roll (assume he already knows that an object's capacity to roll depends in part its shape). Now suppose that the vase stops rolling, and the situation evolves so that the subject would no longer classify the vase among the oval things but among the round things (so the perceptual experience ceases to be misleading). In that case the subject will retain his rationally formed belief that *that shape* endows objects with the ability to roll, though now he will revise his belief about which shape *that shape* happens to be (round, not oval). This is exactly the pattern of belief revision we would expect if our intuitive verdict about cases like ROUND VASE were correct. The pattern reflects an attempt to preserve the same kind of internal coherence that we try to maintain when forming beliefs about objects on the basis of a perceptual link. Consider the object case first. "I thought it was large and far away, but now I see that it is small and nearby" we say; "I thought the thing was red, given how it looked then, but now I see it is orange". In such cases of perceptual demonstration of objects, our beliefs evolve to maintain rationality in response to information delivered through the attentional perceptual link that is underpinning perceptual demonstrative aboutness. The case I rehearsed above manifests the same pattern for perceptual ascription. Suppose the subject is asked to make sense of his evolving beliefs about the object's shape. "I saw that objects with that shape roll," he will say. "I simply mistook that shape for oval (given how the shape looked to me before), whereas now I see it is round." The subject's beliefs have evolved to maintain rationality in response to the information that his perceptual link with the shape delivers.

1.3 | Limits of tolerance

This subsection rounds out my discussion of the tolerance of perceptual ascription by examining limits on this tolerance. These limits will play an important role in Section 4.

To get a feel for the limits on the tolerance of perceptual ascription, let us first consider limits on the tolerance of perceptual demonstration. Some theorists restrict this type of tolerance by endorsing a version of "sortalism" about perceptual demonstration. This is the view that a subject can perceptually demonstrate an object only if she successfully classifies the object according to its kind.[12] For instance, a sortalist about perceptual demonstration might insist that to think *that is tall* of a tree requires recognising it as a tree. Sortalism about perceptual demonstration entails restrictions on tolerance: a subject can perceptually demonstrate a perceived object only if she can correctly isolate properties of the object whose identification enables classification of the object according to its kind. The more specific the kind a subject must use when classifying an object for the purposes of perceptual demonstration, the fewer of an object's properties she can get wrong without undermining her capacity to think about it. Yet even a demanding version of sortalism will not make perceptual demonstration wholly intolerant.

Just as some endorse sortalism about perceptual demonstration, one could also endorse a version of sortalism about perceptual ascription. On this type of sortalism, to perceptually ascribe an observable property a subject must classify the property according to one of its kinds: it is a shade, a shape, a height.[13] Hence if sortalism about perceptual ascription is true, a thought you might express with "It's like this" must equally well be expressed by a sentence like "It's that colour" or "It's that shape" (compare how a traditional sortalist about perceptual demonstration might insist that "That is red" is always elliptical for a sentence such as "That table is red").

Some type of sortalism about perceptual ascription must be true. Perception cannot lead a subject to mistake a round object for something triangular, for example, yet at the same time permit him to perceptually ascribe the object's actual (round) shape.[14] But just as with sortalism about perceptual demonstration, sortalism about perceptual ascription is compatible with a degree of tolerance. To get a property's kind right—to correctly classify it as a shade, a shape, etc.—a subject need not know every observable higher-order property of the property. I therefore take it that there are non-trivial limits on the tolerance of perceptual ascription. Certain errors about the higher-order properties of observable properties—those errors that undermine our capacity to classify these properties according to their kinds—would prevent perceptual ascription of the properties. We can think about these limits in terms of determinates and determinables.[15] Given a determinate property (say scarlet) there may exist maximally dissimilar determinates such that if a perceptual experience were to lead a subject to mistake the former for any of the latter, the experience would still permit her to perceptually ascribe the original property. These determinates—along with every determinate in between—would fall under a common minimally specific determinable property. Limits on tolerance would correspond to such determinable properties.

2 | THE TARGET

In this section I unpack the target view that, just as we can mislead others by saying something false, illusory experience misleads by misrepresenting how things stand in the world. I introduce my challenge in Section 3.

2.1 | Common sense view

My target is an explanation of what I dub the "common sense view" of the relationship between visual experience and rational belief formation. The common sense view combines three points of widespread agreement over the relationship between visual experience and rational belief formation:

1. Visual experience of an object can lead us to rationally form beliefs about the object, and these beliefs may be true.

These are usually classified as cases of "veridical" experience. For instance, Sylvia might see an oval vase and in response come to believe *that is oval*. The vase is oval, so her belief is true.

2. Experience can also mislead us: the beliefs we rationally form about a perceived object may be false.

These kinds of cases are familiar, and are usually classified as cases of illusion. For instance, Psyche might see her round vase through a distorting medium, and in response come to believe *that is oval*. But the vase is round, not oval, so Psyche's belief is false.

3. Only some beliefs about an object formed in response to a visual experience of the object are rationally formed.

Strange are the cases that support 3. For instance, a subject might see a round vase, and the vase might *look* round to him, and yet somehow the subject might be caused by his experience to believe *that is oval*. But the vase is round, not oval, so his belief is false. Though the belief is caused by the experience, absent further exculpatory information one cannot make sense of it as a rational response to the experience. The subject sees a vase that both is round and looks round, so how can his belief that it is oval, formed in response to his experience, count as rational?[16]

There is also a fourth kind of case that I have not mentioned in the presentation of the common sense view. The view remains silent about the rational significance of hallucinatory experience—bizarre experiences that lack an object (e.g., it seems to Macbeth as though a dagger floats before him, but lo, there is in fact no dagger there). My own view is that common sense delivers no clear verdict about hallucinatory cases.[17]

The combination of 1–3 constitutes the common sense view: only some beliefs formed in response to experience are rationally formed, and these may be either true or false.

2.2 | Epistemic Representationalism

Philosophers find one explanation of the common sense view particularly tempting. Call this explanation 'Epistemic Representationalism'.[18] As gnomically stated earlier, Epistemic Representationalism holds that just as we can mislead others by saying something false, illusory experience misleads by misrepresenting how things stand in the world. The extant position that seems to most clearly entail Epistemic Representationalism is the dogmatism defended by Pryor (2000). Yet many non-dogmatists should also recognise their own positions as falling under the broad umbrella of Epistemic Representationalism, since the view turns out to be significantly weaker than dogmatism. In this subsection, I unpack Epistemic Representationalism and its relation to the common sense view.

Epistemic Representationalism uses a notion of perceptual content. A visual experience's content is what that experience conveys to a subject. Less loosely, an experience's content is the condition under which the experience would count as accurate.[19] For example, if a chair looks red to you, and the chair is red, then your visual experience of the chair is accurate; by contrast, if you could undergo the same experience, but in a situation in which the chair is actually yellow, your experience would be inaccurate.

Let us return to the cases that helped motivate the common sense view. Sylvia saw an oval vase and rightly believed it to be oval, whereas Psyche saw a round vase but was misled by her experience into believing it oval. We assume that Sylvia and Psyche share a normal epistemic background state: each takes herself to be perceiving properly, and not under some malign epistemic influence, etc. Why is it rational for them to form the same beliefs (and thus why is Psyche misled)?

Most agree that Psyche cannot discriminate her situation from Sylvia's whilst relying upon reflection alone. Epistemic Representationalism takes the extra step of insisting that this reflects a "phenomenal match" between Psyche and Sylvia: what it is like for Psyche to undergo her visual experience is the same as what it is like for Sylvia to undergo her experience (i.e., Psyche and Sylvia undergo experiences with the same "phenomenal character").

Epistemic Representationalism assumes that this sameness of phenomenal character helps explain Psyche and Sylvia's parallel rational positions, and thus endorses the plausible claim that what it is like to undergo an experience constrains what it is rational to believe in response to the experience.[20]

This connection between phenomenal character and rationality cannot serve alone as an explanation of the common sense view. We must still explain not only why Psyche and Sylvia form the same belief, but also why it is rational to form that specific belief rather than some other belief.

It is at this stage that Epistemic Representationalism puts the notion of perceptual content to work. In order to preserve the desired connection between phenomenology and rationality, Epistemic Representationalists insist that a visual experience's (rationally relevant) content capture the experience's phenomenal character. Sylvia and Psyche's respective beliefs are then *rational* because they each undergo experiences which would be accurate on the condition that their vase is oval (i.e., their experiences share a common content); Psyche's belief is *false*, despite being rational to form, because her experience's content is inaccurate (her vase is round, not oval).

This explanation of the common sense view by Epistemic Representationalism carries at least three important commitments:

> RATIONALITY AND CONTENT: What it is rational for a subject to believe in response to her visual experience is explained in part by which content that experience possesses (a full explanation of a subject's rational position will involve facts about a subject's background epistemic state, for example, whether she takes herself to be perceiving normally).
>
> CONTENT AND PHENOMENOLOGY: The rationally relevant content of a visual experience—the content which explains what it is rational for a subject to believe in response to the experience—must explain what it is like to undergo the experience (i.e. it must capture the experience's phenomenal character).
>
> COMMON CONTENT: In standard cases of misleading visual experience (roughly, those cases like Psyche's in which the subject presupposes that she is perceiving properly) the subject's visual experience must be inaccurate, and it must be possible for someone (in Psyche's case, Sylvia) to undergo an accurate visual experience with the same content.

The last two commitments together capture the phenomenal match that Epistemic Representationalists insist holds between subjects like Sylvia and Psyche. And RATIONALITY AND CONTENT explains why only some beliefs formed in response to a visual experience are rational to form: these beliefs count as rationally formed only if their contents stand in the right relations to the visual experience's content. For instance, Sylvia's belief that her vase is oval is rational because her experience is accurate on the condition that the vase is oval. But because it explains rational belief formation in terms of perceptual content, RATIONALITY AND CONTENT requires facts about which content a perceptual experience possesses to be *prior in order of explanation* to facts about which beliefs it would be rational to form in response to the experience.[21] CONTENT AND PHENOMENOLOGY secures this explanatory priority with a tight link between perceptual content and phenomenology: Psyche and Sylvia undergo experiences with the same content because they undergo experiences with the same phenomenal character. Finally, CONTENT AND PHENOMENOLOGY applies only to the rationally relevant content of visual experience (i.e., the content appealed to by RATIONALITY AND CONTENT). This leaves open the possibility that, for one explanatory purpose or another, we might posit perceptual content that does not explain what it is rational to believe in response to experience.

Epistemic Representationalism is significantly weaker than dogmatism because its commitments (and motivations) neither exhaust, nor are exclusive to, dogmatism. For instance, dogmatists and anti-dogmatists can endorse RATIONALITY AND CONTENT, yet disagree over which additional features help to explain what it is rational for a subject to believe in response to perception (e.g., they will disagree over whether a subject must already have grounds for

taking perception to be reliable).[22] Non-dogmatists also frequently explain visual phenomenology (including the veridical/illusory distinction) in terms of content that can sometimes be inaccurate.[23]

In Sections 3–4, I use TOLERANCE OF PERCEPTUAL ASCRIPTION to argue that Epistemic Representationalism is unstable, and that the view (along with any position that falls under its broad umbrella) ought to be abandoned. The instability arises because Epistemic Representationalism puts perceptual content to the same sort of work as testimonial intermediaries in the case of misleading testimony. The problem therefore lies not with any of Epistemic Representationalism's commitments taken in isolation, but with their combination in an explanation of the common sense view.

In order to consolidate Epistemic Representationalism, and bring out why so many have found the view tempting as an explanation of the common sense view, I wish to highlight a suggestive parallel—briefly noted in the last paragraph—with an attractive account of why beliefs formed in response to testimony can be false yet rational. When I testify, I use a sentence in a shared language to express something I wish to say. In a case of ordinary assertion, where I make my utterance with the aim of expressing a belief that I want my audience to share, successful testimony requires that my audience understand and believe what I say. If what I say is true, successful testimony will result in my audience forming a true belief. But if what I say is false, my audience will form a false belief. In both cases, what I say is the same, and my audience has no reason to doubt my word. These two facts explain why it is rational for my audience to form beliefs with the same contents in both cases, and thus why the subject in the second case is misled.[24] What I say when I testify also constrains the sorts of beliefs my audience can rationally form on the basis of my testimony. For example, if I say "London is cold tonight", my audience cannot, without substantial additional stage setting, rationally believe that London is *warm* tonight on the basis of my testimony.

In the testimony case, what I say acts as an intermediary between speaker and hearer. It explains my audience's false (or true) beliefs in a way that preserves their rational status, and it also constrains the class of beliefs that it would be rational for my audience to form in response to my testimony.

The seemingly tight parallel between what we are explaining in the testimony case, on the one hand, and the common sense view, on the other, suggests to many that an explanation of the latter must identify a feature of visual experience that accomplishes the same explanatory task as the testimonial intermediary (i.e., what I say when I testify). A misleading visual experience must share this feature with appropriate non-misleading visual experiences, since that feature must in turn explain why it is rational for a subject to form the same beliefs in response to the misleading and non-misleading experiences, and thus why the subject is rationally misled when he undergoes the misleading experience. Yet what feature of visual experience, if any, can play this explanatory role? Epistemic Representationalism insists that perceptual content should play this role.

3 | THE CHALLENGE

My challenge to Epistemic Representationalism takes the form of a trilemma that I dub the "Tolerance Fork". I begin, however, with some necessary terminology. The following definition isolates the perceptual conditions that sustain rational formation of beliefs that perceptually ascribe a property:

> Definition: A subject stands in "perceptual contact" with a property iff she has a perceptual link with the property that sustains rational formation of beliefs that perceptually ascribe it. (Similarly: a subject stands in "perceptual contact" with an object iff she has a perceptual link with the object that sustains rational formation of perceptual demonstrative beliefs about it).

The definition permits the reformulation of TOLERANCE OF PERCEPTUAL ASCRIPTION. This streamlined version says that a subject can be rationally misled about an observable property on the basis of her visual experience, yet at the same time remain in perceptual contact with the property. More formally:

TOLERANCE OF PERCEPTUAL ASCRIPTION: There are incompatible properties F and G such that cases with the following structure are possible: (a) it is rational for a subject to believe on the basis of her visual experience of an object (which is in fact F) that the object is G; and (b) at the same time, that visual experience puts her in perceptual contact with F.

3.1 | The path to the tolerance fork

I trace the path to the Tolerance Fork for the special case of a subject in ROUND VASE—call him "Gareth"—who sees a round vase but rationally believes it to be oval. Because Gareth's is a tolerance case, the following two claims are true (as instances of [a] and [b] from the streamlined TOLERANCE OF PERCEPTUAL ASCRIPTION):

1. It would be rational for Gareth to believe (in response to his experience) that his vase is oval.
2. Gareth's experience puts him in perceptual contact with roundness (i.e., his vase's actual shape).

The Tolerance Fork arises when we seek to accommodate **2** within an Epistemic Representationalist's explanation of **1**.

What must an Epistemic Representationalist say about **1**? Gareth's vase is round, and yet his experience makes it rational for him to believe that the vase is oval, so Epistemic Representationalism classifies Gareth's visual experience as misleading. And because they seek to explain the common sense view, Epistemic Representationalists must also hold that Gareth could have a non-misled "twin"—call her "Dina". Whereas Gareth sees a round vase and rationally believes it to be oval, Dina sees an oval vase and rationally believes it to be oval. Given RATIONALITY AND CONTENT, the contents of Gareth and Dina's respective experiences explain why these particular beliefs are rational to form (we are assuming for the moment that they share a common background epistemic state). Why is it rational for each of them to believe that their respective vase is oval? Epistemic Representationalists hold that Dina and Gareth undergo experiences with a common phenomenal character, and so use CONTENT AND PHENOMENOLOGY to conclude that Dina and Gareth's experiences possess the same rationally relevant content (at least with respect to the shape of their vases). This sameness of content connects the (supposed) common perceptual phenomenology with Gareth and Dina's respective rational positions. Yet given the common content, why is Gareth's belief false while Dina's true? COMMON CONTENT provides the answer: while Gareth and Dina undergo experiences with a shared content, the relevant content is inaccurate in Gareth's case (hence his false belief) and accurate in Dina's (hence her true belief).

The Tolerance Fork properly comes into view when Epistemic Representationalists recognise a substantial connection between perceptual contact and perceptual content. In its schematic form, this connection runs as follows:

> *Contact and Content*: if a visual experience of an object places a subject in perceptual contact with F, the experience depends for its accuracy upon that object being F.

I first show that Epistemic Representationalists must accept *Contact and Content*, and later consider how one might spell out the relevant dependence.

A positive and a negative argument together force Epistemic Representationalists to accept *Contact and Content*.

I begin with the positive argument. There are limits to the reach of perceptual ascription. For instance, we cannot perceptually ascribe microscopic properties (e.g., mass or charge). We also cannot perceptually ascribe colours or shapes or sizes in a pitch-black room (at least not on the basis of visual experience). These limits reflect a broader perceptual boundary: when restricted to visual experience, we can perceptually ascribe only those distal properties we can see.[26] Yet which properties count as visible depends in part upon visual phenomenology. CONTENT AND

PHENOMENOLOGY—and thus Epistemic Representationalism—requires that perceptual content capture visual phenomenology. When I see a red cup on my table, CONTENT AND PHENOMENOLOGY says that the cup will look red just if redness figures appropriately in the accuracy conditions of my experience (e.g., the experience is accurate on the condition that the cup is red). As a result, if Epistemic Representationalists wish to respect perceptual limits on perceptual ascription, they must accept that perceptual content constrains the reach of perceptual ascription. *Contact and Content* expresses this constraint in its most schematic and plausible form. Moreover, this line of thought generates an important restriction on how Epistemic Representationalists must understand the dependence invoked by *Contact and Content*: whatever specific dependence relations figure in a non-schematic version of *Contact and Content*, these relations must still permit the subject to see a property when they hold between that property and an experience's content.

Now for the negative argument: if Epistemic Representationalists were to deny *Contact and Content*, they would be unable to capture the perceptual ascription's tolerance. To see why, consider Gareth's case once again. Epistemic Representationalists hold that Gareth's belief that his vase is oval is rational because his experience's content depends for its accuracy on the vase being oval. If they deny *Contact and Content*, they cannot provide a parallel explanation of those beliefs of Gareth's which perceptually ascribe the actual shape of his vase (i.e., its roundness). A parallel explanation would require the accuracy of the content of the subject's experience to somehow depend upon the vase being round. Yet the rejection of *Contact and Content* excludes this sort of dependence. But if Epistemic Representationalism cannot explain why Gareth's perception-based beliefs count as rationally formed, it fails to accommodate perceptual ascription's tolerance.

While I have argued that Epistemic Representationalists must accept a schematic connection between perceptual content and perceptual contact, I have not yet suggested which non-schematic connection they should accept. Looking ahead, I argue in Section 4 that however Epistemic Representationalists understand *Contact and Content*, they cannot accommodate perceptual ascription's tolerance.

3.2 | Tolerance fork

We are finally in a position to state the Tolerance Fork. Since Gareth's is a tolerance case:

1. It would be rational for Gareth to believe (in response to his experience) that the vase is oval.
2. Gareth's experience puts him in perceptual contact with roundness (the vase's actual shape).

Now add *Contact and Content*:

3. *Contact and Content*: if a visual experience of an object places a subject in perceptual contact with F, then the experience depends for its accuracy upon that object being F.

But in Section 4 I establish **4**:

4. If **1**–**3** are true, Epistemic Representationalism is false.

The tension between Epistemic Representationalism and **1**–**3** arises in different forms, depending on which non-schematic version of *Contact and Content* one adopts. (Recall that 'depends for its accuracy' is a placeholder for a more specific notion of dependence.) I make the case in Section 4 that **1**–**3** and Epistemic Representationalism are jointly inconsistent under every candidate version of *Contact and Content*. While the argument for **4** is complicated, I can provide a brief intuitive motivation for the premiss. Given that Gareth's experience places him in perceptual contact with his vase's roundness (from **2**), an obvious way to satisfy *Contact and Content* is to hold that Gareth's

experience is accurate on the condition that his vase is round. Epistemic Representationalism must use this content to explain the rational status of Gareth's perception-based beliefs. Yet how could it be rational for Gareth to believe the vase is *oval* in response to an experience with that content (i.e., how could **1** be true)? We seem to have a mismatch between the perceptual content Epistemic Representationalists require for their explanation of Gareth's beliefs, and the perceptual content that perceptual ascription's tolerance forces them to recognise. In Section 4 I reveal how this tension reflects genuine inconsistency.

We now arrive at the Tolerance Fork. **1**–**4** entail that Epistemic Representationalism is false. Since Epistemic Representationalists must accept some version of *Contact and Content*, they face a hard choice: deny **1** or **2** (and thus TOLERANCE OF PERCEPTUAL ASCRIPTION); deny **4**; or abandon Epistemic Representationalism. This forced choice is the Tolerance Fork.

I have made the case both for TOLERANCE OF PERCEPTUAL ASCRIPTION, and for the claim that Epistemic Representationalism supports *Contact and Content*. So the Tolerance Fork generates only two questions: Why accept **4**? And if one accepts **4**, how far must one depart from Epistemic Representationalism to escape the Tolerance Fork? I address the first question in Section 4, and the second in Section 5.

4 | CONSOLIDATING THE TOLERANCE FORK

In Section 3.1 I argued that Epistemic Representationalists must accept *Contact and Content*:

> *Contact and Content*: if a visual experience of an object places a subject in perceptual contact with F, the experience depends for its accuracy upon that object being F.

We can assess the relationship between this principle, Epistemic Representationalism, and perceptual ascription's tolerance only once we replace 'depends for its accuracy' (a placeholder) with a specific dependence relation. I defend **4** in this section by arguing that however we understand the relevant dependence, Epistemic Representationalism cannot accommodate perceptual ascription's tolerance.

Section 4 has three parts: Section 4.1 canvasses candidate versions of *Contact and Content*; Section 4.2 shows that each prevents Epistemic Representationalism from accommodating perceptual ascription's tolerance; Section 4.3 closes a potential escape route from **4**.

4.1 | Variations on *Contact and Content*

A demanding version of *Contact and Content* fits most naturally with how we have been discussing perceptual content:

> *Contact and Content*$_{\Phi}$: An experience places a subject in perceptual contact with a property only if that experience is accurate only on the condition that something (e.g. a perceived object) possesses the property.

Contact and Content$_{\Phi}$ entails that if my experience of an object puts me in perceptual contact with red, and you undergo an experience with the same content as my experience, your experience will be accurate only on the condition that a given object is red. Evaluating this kind of non-schematic version of *Contact and Content* requires a more nuanced understanding of perceptual content than we have worked with so far. I avoided these nuances to highlight the generality of both Epistemic Representationalism and the Tolerance Fork. *Contact and Content*$_{\Phi}$ fits naturally with (but is not entailed by) extant views about the individuation of perceptual content. For example, so-called

"Russellians" identify perceptual contents with structured propositions and hold that two contents are the same only if they include the same properties as constituents.

Other versions of *Contact and Content* recognise a less direct relationship between perceptual contact and perceptual content than *Contact and Content*$_{\Phi}$. They allow that two experiences can share their content, and yet differ in which properties they put subjects in perceptual contact with. These versions must therefore hold that an experience puts us in perceptual contact with a given property only if that property satisfies an appropriate condition, where this condition may sometimes be such that it could have been satisfied by a different property. The versions differ in the relationship each recognises between this condition and the perceptual content of the corresponding experience. There are two options. *Contact and Content*$_{\Delta}$ includes the condition in the perceptual content:

> *Contact and Content*$_{\Delta}$ ("descriptivist"): An experience places a subject in perceptual contact with a property only if the experience's content includes an appropriate condition that the property satisfies (i.e. a condition such that the experience is accurate only if something possesses the property satisfying the condition).

Like *Contact and Content*$_{\Phi}$, *Contact and Content*$_{\Delta}$ fits naturally with (but is not entailed by) certain views of perceptual content.[27] Given our gloss of perceptual content as that which is conveyed to a subject by an experience, *Contact and Content*$_{\Delta}$ requires that the relevant condition reflect what is conveyed to the subject. For example, Chalmers (2004) suggests the following gloss on what a visual experience as of a red apple conveys: it presents the apple's colour "as *the normal cause of experiences of this sort*".

By contrast, *Contact and Content*$_{\Upsilon}$ does not require that perceptual content include the property-determining condition:

> *Contact and Content*$_{\Upsilon}$ ("demonstrative"): An experience places a subject in perceptual contact with a property only if the property satisfies an appropriate condition (a condition which is such that the experience is accurate only if something possesses the property satisfying the condition).

Which accounts of perceptual content might support this minimal thesis? The connection between the relevant condition and the experience's content cannot be too attenuated: satisfying the condition must be sufficient for a property to count as visible. One well-known type of account constructs perceptual contents from modes of presentation of properties and objects, where these modes of presentation are themselves fixed by the relevant property- or object-determining conditions.[28] Advocates use demonstratives to describe what an experience with this kind of content conveys: an experience of a red apple might present the apple's colour "as *that colour*". More generally, any account of perceptual content that supports *Contact and Content*$_{\Upsilon}$ must hold that our judgments about the content of visual experience adequately capture that content only if they sometimes involve perceptual ascription of perceived properties. For if the judgments were to only ever ascribe a perceived property as the satisfier of such-and-such a condition (e.g., *the colour which…*), the judgments would not respect the difference between accounts of content that support *Contact and Content*$_{\Delta}$ and accounts that support *Contact and Content*$_{\Upsilon}$.

This division between versions of *Contact and Content* is exhaustive. Each version can be understood as restricting the class of property-determining conditions: *Contact and Content*$_{\Upsilon}$ permits any condition that satisfies the minimal constraints derived from the original argument for *Contact and Content* (e.g., that the property must still count as visible); *Contact and Content*$_{\Delta}$ permits only those candidate conditions included in perceptual content; and *Contact and Content*$_{\Phi}$ can be understood as permitting only those candidate conditions that are both included in perceptual content and could not have been satisfied by a different property (e.g., the condition that a property be identical to F). The remainder of this section shows that we cannot accept one of these views and still maintain both Epistemic Representationalism and TOLERANCE OF PERCEPTUAL ASCRIPTION.

4.2 | Contact and Content$_{\Phi}$

TOLERANCE OF PERCEPTUAL ASCRIPTION entails that Gareth's experience places him in perceptual contact with the roundness of his vase. Applying *Contact and Content*$_{\Phi}$, we get the result that Gareth's experience is accurate on the condition that the vase is round. Yet this turns out to be inconsistent with Epistemic Representationalism.

There are two ways to bring out the inconsistency. The first targets Epistemic Representationalism's explanation of why Gareth is rationally misled by his experience. Even though Gareth's experience places him in perceptual contact with the roundness of the vase, it remains rational for him to believe (in response to his experience) that the vase is oval. Epistemic Representationalists must explain the latter fact by appeal to the content of Gareth's experience. The explanation in question plausibly requires that Gareth's experience is accurate on the condition that the vase is oval. We now have a violation of COMMON CONTENT (and thus Epistemic Representationalism). That thesis required that Gareth's experience could have been accurate; but the experience could never be accurate if its accuracy conditions include that the vase is both round and oval.

Epistemic Representationalists might respond by refusing to spot us the natural assumption that their explanation of Gareth's rational position requires that his experience be accurate on the condition that the vase is oval. This refusal would come at an immediate cost: they would owe an alternative explanation of why Gareth is rationally misled.

The natural assumption turns out to be unnecessary, however, since we can derive a contradiction without it. Epistemic Representationalism requires that Gareth share the content of his experience with a non-misled "twin"—Dina. Whereas Gareth sees a round vase and rationally believes it to be oval, Dina sees an oval vase and rationally believes it to be oval. According to Epistemic Representationalism, Dina's belief is true, while Gareth's is false, because Dina's experience is accurate and Gareth's inaccurate (i.e., veridical and illusory, respectively) with respect to the shapes of their vases. And Dina and Gareth rationally form beliefs with the same content, despite the difference in their perceived surroundings, because their experiences share a common content.

Yet if Gareth's experience is accurate on the condition that a given vase is round (as required by *Contact and Content*$_{\Phi}$), and Gareth and Dina undergo experiences with the same content, Dina's experience will also be accurate on the condition that a given vase is round. Dina's experience turns out to be inaccurate after all, since her vase is oval, not round. The resulting contradiction arises because of COMMON CONTENT. And what licences applying COMMON CONTENT to Dina and Gareth is Epistemic Representationalism's need to accommodate a supposed phenomenal match between Dina and Gareth's experiences.

4.3 | Against alternatives

If Epistemic Representationalism adopts either *Contact and Content*$_{\Delta}$ or *Contact and Content*$_{\Upsilon}$, it dodges the obstacles that prevent it from accepting *Contact and Content*$_{\Phi}$. Asked to describe the content shared between Gareth and Dina, someone who accepts *Contact and Content*$_{\Upsilon}$ [or *Contact and Content*$_{\Delta}$] could say that each subject undergoes an experience that is accurate on the condition that their respective vase is *that shape* [or: *the shape such that…*], and *that shape* [or: *the shape such that…*] is oval. Gareth's experience will be inaccurate, since it counts as accurate only if the actual roundness of his vase is identical with ovalness; Dina's experience will be accurate, since it is accurate only if the actual ovalness of her vase is identical with ovalness.

I argue in this subsection that neither *Contact and Content*$_{\Delta}$ nor *Contact and Content*$_{\Upsilon}$ permits Epistemic Representationalism to accommodate perceptual ascription's tolerance. The problem traces to the fact that accommodating perceptual ascription's tolerance requires accommodating the non-trivial limits on this tolerance introduced in Section 1.2. These limits force a division between 'tolerance' cases like Gareth's, and "intolerance" cases in which visual experience of an object that misleads a subject about the object's shape (for example) does not permit a subject to perceptually ascribe the object's actual shape.

Subjects in an intolerance case cannot perceive the properties that, due to exceeding the limits of tolerance, they fail to be in perceptual contact with. To see why, suppose for *reductio* that the claim is false. What immediately follows is a violation of a plausible condition on visual awareness. This condition holds that if a subject sees a property (and is not otherwise cognitively defective), she can perceptually ascribe it. Yet even if we reject this condition, we are left with an implausibly liberal construal of perceptual awareness. On this liberal construal, a colour might look like a shape; or (even if that possibility is independently ruled out) one might see an oval that looks like an equilateral triangle. Admitting these sorts of cases would destroy our grasp of what it is to be perceptually aware of a property.[29]

4.3.1 | Contact and Content$_{\Delta}$

Contact and Content$_{\Delta}$ assigns a condition two roles: it must both determine which property an experience puts a subject in perceptual contact with, and serve as a component of that experience's content. I argue that no condition can plausibly play both roles, given limits on perceptual ascription's tolerance.

Since the condition determines which property an experience puts a subject in perceptual contact with, it must be satisfied in a tolerance case, and unsatisfied in an intolerance case. But whether a given visual experience of an object constitutes an intolerance case depends upon—among other factors—both the degree to which the experience misleads a subject, and which property the experience misleads the subject about. A round object can look oval without generating an intolerance case, for example, but the same is not true of a square object that looks round.

Epistemic Representationalists now face a problem. If perceptual content includes property-determining conditions, these conditions must be sensitive to those factors that distinguish tolerance cases from intolerance cases. And given CONTENT AND PHENOMENOLOGY, Epistemic Representationalists must assume that this perceptual content reflects visual phenomenology. Yet visual phenomenology cannot justify introducing perceptual content that involves such complex conditions.[30] Such complexity is quite simply not conveyed to a subject who undergoes a visual experience. While we have some sense of what it is for experience to present us with one colour rather than another, for example, perceptual experience leaves it quite unclear where exactly the boundaries of tolerance lie (and, more importantly, why they lie where they do). The issue is not merely the addition of complexity, but the kind of complexity forced into visual phenomenology. Putting this together: Epistemic Representationalists cannot accept both perceptual ascription's tolerance and *Contact and Content*$_{\Delta}$, since that combination would force the abandonment of CONTENT AND PHENOMENOLOGY.

4.3.2 | Contact and Content$_{\Upsilon}$

By permitting the exclusion of the property-determining condition from perceptual content, *Contact and Content*$_{\Upsilon}$ avoids adding untenable complexity to that content, and so does not run afoul of CONTENT AND PHENOMENOLOGY. Nevertheless, I argue that an unavoidable tension still arises between Epistemic Representationalism and *Contact and Content*$_{\Upsilon}$.

In order to bring out this tension, I repurpose an argument from Soteriou (2005) for disjunctivism. For present purposes, 'disjunctivism' is the view that reflectively indiscriminable veridical and hallucinatory perceptual experiences possess different phenomenal characters. The view is compatible with Epistemic Representationalism, since the latter remains silent about hallucination. Soteriou derives disjunctivism from the following commitment[31]:

> **Particularity:** A subject's judgments about how things seem to her when undergoing a veridical visual experience (i.e. judgments about the experience's phenomenal character) adequately capture how

> things seem only if some of these judgments involve successful perceptual demonstration of the objects of perception.

Judgments that involve perceptual demonstration exhibit an asymmetry of defeat with respect to veridical and hallucinatory experience. If I come to know that the conditions for successful perceptual demonstration do not obtain (in our terms, that I am not in perceptual contact with an object), this knowledge will defeat those of my judgments that involve perceptual demonstration (e.g., "It seems to me that *that* blue table sits before me"). The knowledge need not derive from introspection: I may know on non-introspective grounds that I am hallucinating, and thus that there is no object with which I am in perceptual contact. By contrast, veridical visual experiences will not be subject to this kind of defeat, since they will always place a subject in perceptual contact with an object.

Soteriou exploits this asymmetry of defeat to argue that any account of perceptual content committed to Particularity must also accept disjunctivism. His motivates and applies the following test for disjunctivism:

> The disjunctive theory of perception is correct if the positive claims about the phenomenal characters of hallucinatory experiences that are warranted by introspective evidence are defeated by non-introspective evidence, whereas the positive claims about the phenomenal characters of veridical perceptions that are warranted by introspective evidence are not defeated by non-introspective evidence. (Soteriou, 2005, p. 187)

Given the asymmetry of defeat highlighted above, applying this test to Particularity yields disjunctivism.

As I argued earlier, introspective judgment will adequately capture the sort of perceptual content that supports *Contact and Content*$_{\Upsilon}$ only if it sometimes involves perceptual ascription of the perceived properties (e.g., "It's that colour, and that colour is red"). Hence, given that CONTENT AND PHENOMENOLOGY requires that perceptual content capture visual phenomenology, *Contact and Content*$_{\Upsilon}$ and Epistemic Representationalism together entail the following dual of Particularity:

> **Particularity*:** A subject's judgments about how things seem to her when undergoing a veridical visual experience (i.e. judgments about the experience's phenomenal character) adequately capture how things seem only if some of these judgments involve successful perceptual ascription of perceived properties.

Soteriou's original test for disjunctivism used Particularity to establish a difference in phenomenal character between indiscriminable veridical and hallucinatory perceptual experiences. Applying essentially the same[32] test to Particularity* promises to establish an even more radical thesis: that indiscriminable veridical and illusory perceptual experiences may differ in phenomenal character.

Of course, this "Radical Disjunctivism" follows only if an indiscriminable illusory perceptual experience could fail to support successful perceptual ascription (otherwise the judgments mentioned by Particularity* will remain immune to defeat by knowledge gleaned from non-introspective evidence). Many would no doubt attempt to resist Radical Disjunctivism at precisely this point. These opponents might grant Particularity* and Soteriou's test, but insist that no illusory perceptual experience can be both indiscriminable from a veridical experience and incapable of supporting successful perceptual ascription (i.e., incapable of placing a subject in perceptual contact with a perceived property).

We know better. Intolerance cases support the move (via Soteriou's test) from Particularity* to Radical Disjunctivism. These cases involve misleading illusory perceptual experiences that fail to place subjects in perceptual contact with perceived properties. And these perceptual experiences are also indiscriminable from veridical perceptual experiences; otherwise they would not be misleading. Hence, if we accept *Contact and Content*$_{\Upsilon}$ (and thus Particularity*), we get the following result: despite reflection being unable to discriminate a perceptual experience in an intolerance case from a veridical experience, these experiences do not share a common phenomenal character.

This result violates Epistemic Representationalism's tenet that misleading visual experiences mislead in part because they share their phenomenal character with veridical experiences. Epistemic Representationalists cannot easily dodge this problem. They must admit intolerance cases if they are to accommodate perceptual ascription's tolerance. Hence, their only options are to deny *Contact and Content*$_Y$ or undermine Soteriou's test.

I suspect that opponents can undermine Soteriou's test—and thus his argument from Particularity to disjunctivism—but I argue below that this move will not save Epistemic Representationalism.

Assume (with Soteriou, and many others) that positive claims about the phenomenal character of experience are warranted only if they are ultimately grounded in introspective evidence. Given this assumption, satisfying Soteriou's test entails that we cannot possess a warrant for the positive claim that a veridical experience and an indiscriminable hallucination share a common phenomenal character. Yet the absence of a warrant for believing p is not usually a warrant for believing not-p. So satisfying Soteriou's test does not obviously generate a warrant for the disjunctivist thesis that the experiences do not share a common phenomenal character. As a result, unless Soteriou can bridge this gap between the weaker thesis and disjunctivism, his test fails as a test for disjunctivism.

My use of Soteriou's test does not require bridging the gap. Epistemic Representationalism asserts that misleading illusory experiences and indiscriminable veridical experiences share a common phenomenal character. Showing that this positive claim lacks warrant—the analogue of the weaker thesis that clearly follows from satisfying Soteriou's original test—therefore suffices to undermine Epistemic Representationalism. So even if the original test does not establish disjunctivism, essentially the same test will combine with Particularity* to undermine Epistemic Representationalism.

4.4 | Tolerance cases are not exceptions

Have I established **4**? Not quite yet. I have taken it for granted that Gareth's case—and most other tolerance cases—are otherwise standard cases of misleading visual experience. It was this assumption that licenced my claim in §4.2 (on behalf of Epistemic Representationalism) that Gareth's experience was inaccurate, and thus that he must have a non-misled twin, Dina, who undergoes a veridical experience with the same content as Gareth's experience.

By contrast, a subject with a non-standard background epistemic state could be misled by an experience even when that experience is veridical (i.e., when that experience's content is accurate). For example:

> Dorothy glances at her desk and sees an oval vase. She also has independent support for the mistaken belief that (like Gareth) she's seeing the vase through a distorting medium. So in response to seeing the vase, she forms the false belief *that is round*.

A belief about her perceptual situation mediates Dorothy's response to seeing the vase. As a result, she is misled into believing that the vase is round. Yet if her experience possesses content, this content is accurate on the condition that the vase is oval (and so her experience counts as veridical).

Dorothy's "non-standard" case brings out the extent to which an account of how perceptual experience interfaces with belief must be sensitive to complications arising from the influence of a subject's background epistemic state.[33] Whether a given case of misleading visual experience counts as standard or non-standard depends on the sort of influence exerted by a subject's background epistemic state. In Dorothy's case, she mistakenly believes that her experience is illusory, and this belief imposes a rational constraint on which beliefs she forms in response to her experience.

Unfortunately, the complications Dorothy's case introduces cannot save Epistemic Representationalism. The existence of cases like Dorothy's—cases in which explicit mediating beliefs complicate matters—gives us no reason to think that in *all* cases the path from perception to belief is mediated by other explicit beliefs about one's

perceptual situation. In fact, it might be especially implausible to think this in cases like Gareth's, where someone is simply absorbing information delivered by a perceptual link.

Problems remain even if we shelve this initial worry. Insofar as it makes sense to talk about Gareth's (presumably highly tacit) beliefs about his own perceptual situation, his error lies in the other direction: he assumes that he is not a victim of illusion. Gareth shares this assumption not just with other subjects in tolerance cases, but with all subjects forming beliefs by uptake from a perceptual experience whose veridicality they have no reason to doubt. As a result, Gareth's assumption about his perceptual situation cannot make him an exception to the Epistemic Representationalist story. If it could, ordinary instances of belief formation by uptake from perception also count as exceptions. And in that case we lose our grip on the rationalising role of perceptual content (and thus RATIONALITY AND CONTENT). So the attempt to restrict Epistemic Representationalism undermines the view itself.

5 | BEYOND EPISTEMIC REPRESENTATIONALISM

Epistemic Representationalists will seek creative ways to avoid the Tolerance Fork. These may take the form of an account that avoids the arguments in Sections 4.1–4.3, or an alternative ground upon which to treat tolerance cases as exceptions. Yet while tortured modifications might defuse the threat posed by tolerance cases, the resulting baroque views are unlikely to enjoy the kind of deep theoretical or intuitive attraction that made Epistemic Representationalism so tempting in the first place.

What remains once we abandon Epistemic Representationalism in response to the Tolerance Fork? Not much—or so I argue below. We should instead seek a radical account of the relationship between visual experience and rational belief formation.

As we saw in Section 2, a pair of motivations drives Epistemic Representationalism. On the one hand, the view seeks to respect a deeply intuitive connection between rationality and perceptual phenomenology. What it is like to undergo an experience—its phenomenal character—constrains and (partially) explains what it is rational for a subject to believe in response to the experience. On the other hand, reflection often cannot discriminate misleading experiences from non-misleading experiences (recall Sylvia and Psyche). Epistemic Representationalists assume that the misleading experiences must therefore share their phenomenal character with the relevant non-misleading experiences. They seek to preserve this "phenomenal match".

Epistemic Representationalists satisfy these foundational motivations by introducing an explanatory role for perceptual content. CONTENT AND PHENOMENOLOGY and RATIONALITY AND CONTENT together explain the intuitive connection between rationality and perceptual phenomenology: perceptual content captures an experience's phenomenal character, and that content in turn constrains and (partially explains) what it is rational for a subject to believe in response to the experience. COMMON CONTENT permits those who accept CONTENT AND PHENOMENOLOGY to respect the supposed phenomenal match between misleading and non-misleading experiences. Given that (rationally relevant) perceptual content must capture an experience's phenomenal character, the phenomenal match between the experiences requires a common perceptual content.

My defence of the Tolerance Fork demonstrates that Epistemic Representationalists cannot jointly maintain their three central commitments. Something must go. Yet jettisoning any commitment risks abandoning a foundational motivation for Epistemic Representationalism. Epistemic Representationalism was so tempting in part because its foundational motivations—including its acceptance of the common sense view—possess enormous intuitive appeal. Those still gripped by this appeal must therefore pursue a radical alternative account of the common sense view. Such an account will discard Epistemic Representationalism's assumption that one can use a notion of perceptual content, taken as explanatorily prior to facts about the rational status of beliefs formed by uptake from experience, to explain patterns of rationality for these beliefs. The account will thus abandon both RATIONALITY AND CONTENT and CONTENT AND PHENOMENOLOGY.[34]

Others will see the Tolerance Fork as a reason to reject a foundational motivation for Epistemic Representationalism. The Tolerance Fork leaves these theorists two (non-exclusive) options: abandon COMMON CONTENT, and thus arguably the posited phenomenal match (Radical Disjunctivism threatens!)[35]; or abandon only one of CONTENT AND PHENOMENOLOGY or RATIONALITY AND CONTENT, and thus the intuitive connection between rationality and phenomenology.

I want to suggest in closing that a central lesson of the Tolerance Fork is that accounts of the relationship between visual experience and rational belief formation must take seriously a separation between two ways in which visual experience can "fall short" (a separation erased by Epistemic Representationalism). Visual experience can *mislead* us about whether a perceived object possesses a property. It can also *fail to put us in perceptual contact* with an object's properties, and thereby prevent us from perceptually ascribing them. Given *Contact and Content*$_{\Phi}$—and ignoring tolerance for a moment—Epistemic Representationalism entails that a visual experience misleads us about whether a perceived object is F just if it fails to put us in perceptual contact with F. Perceptual ascription's tolerance forces apart these two ways for experience to "fall short", and so exposes Epistemic Representationalism's confusion over the dual role of observable properties in visual experience.

Alternative accounts that accept the common sense view can still respect a tight connection between phenomenology and rationality. They need only recognise a division between two aspects of Gareth's visual phenomenology: one aspect captures the fact his vase looks oval; the other captures his perceptual contact with the vase's actual shape. The phenomenology characteristic of perceptual contact with a property is an analogue for properties of the phenomenology characteristic of an attentional "lock" or "focus" on an object. The phenomenology characteristic of attentional focus may be present even in the absence of an attentional link with an object. Likewise, visual experience can possess "contact" phenomenology even when the limits of tolerance are exceeded. Intolerance cases will therefore also sometimes involve 'contact' phenomenology.

These two aspects of perceptual phenomenology play different rationality-securing roles. In Gareth's case, how his vase looks makes it rational for him to believe it is oval; yet the presence of 'contact' phenomenology that arises from perceptual contact with the vase's actual shape helps make it rational for him to form beliefs that perceptually ascribe this shape. Abandoning Epistemic Representationalism ushers in a new project: to explain how and why these two aspects of perceptual phenomenology play their rationalising roles.

ACKNOWLEDGMENTS

This article would not exist in its present form without the generous feedback of others. Special thanks to Imogen Dickie, Bill Brewer, Matt Soteriou, Benj Hellie, Gurpreet Rattan, Phil Kremer, Ori Beck, Fatema Amijee, Tashi Alford-Duguid, Maarten Steenhagen, Boyd Millar, Zachary Irving, and Diana Raffman, all of whom read and commented on earlier drafts. Thanks also to Mike Martin, Rachel Goodman, Clayton Littlejohn, Alex Moran, Kathleen Akins, Matthew Parrott, Martin Hahn, Kristoffer Sundberg, Naomi Duguid, Ellen Fridland, David Papineau, James Davies, Thomas Raleigh, Ivan Ivanov, and audiences at King's College London, Simon Fraser University, the University of Manitoba, the National University of Singapore, the Philosophy of Language and Mind Conference at the University of Oslo, and the Society for Exact Philosophy.

ORCID


Dominic Alford-Duguid https://orcid.org/0000-0002-9213-4369

ENDNOTES

[1] I take no stand on many core features of perceptual links: they may be more or less fundamental than perceptual experience of a scene; and a perceptual link with a property F may not require that anything in the perceived scene look F.

[2] Cf. Evans (1982); Campbell (2002); Recanati (2012); Dickie (2015).

[3] Perhaps the closest has been Heck (2000, p. 495), who formulated the possibility of such tolerance, though did not endorse it, as part of an argument against McDowell (1994). Hellie (2011, Section 5.1) also comes close.

[4] Compare Hardin (1988, p. 25): "We may distinguish three dimensions of perceived color:hue, brightness, and saturation… the hue of a color is its [degree of] redness, or greenness, or yellowness, or blueness."

[5] I take no stand on whether one could perceptually ascribe the vase's merely apparent shape.

[6] I answer the objection in detail in "Perception-Based Thought about Properties Tolerates Perceptual Error" (ms), a paper wholly dedicated to defending TOLERANCE OF PERCEPTUAL ASCRIPTION.

[7] In spelling out the relevant counterfactual dependence, Macpherson and Batty follow Lewis (1980).

[8] For a textbook account of visual attention to objects, including a discussion of how attention responds to changes in the attended object, see Palmer (1999, pp. 554–6; 561–3)

[9] These dispositions constitute our perception-based governing or "controlling conception" of an object (cf. Evans, 1982, Ch. 5). The recent literature on mental files and singular thought descends from Evans's discussion, and includes specific proposals about the constraints that a governing conception must satisfy for a body of beliefs to be about a particular object (cf. Dickie, 2010; Recanati, 2012).

[10] The harder question remains a topic of ongoing debate. Cf. Evans (1982); Campbell (2002); Smithies (2011); Recanati (2012); Dickie (2015).

[11] As I understand "attention to the vase's shape", if the vase were to change its shape, so would the target of attention; the target of this attentional link is therefore a determinate shape rather than whatever determinable shape remains throughout changes in the vase's shape. Carrasco (2011) provides an overview of the empirical literature on visual attention—including a survey of work on attention to observable properties (what psychologists call "feature-based attention"). For a philosophical overview of recent work on attention, see Wu (2014).

[12] A sortalist must explain, inter alia, which kinds (of the many an object may fall under) a subject must use when classifying an object for the purpose of thinking about it. Sortalists include Strawson (1959), Quine (1960), Dummett (1973/1981), and Wiggins (1997). Dickie (2014) defends a less extreme form of sortalism. For arguments against sortalism, see Ayers (1974, 1997), Campbell (2002, 2006), and Goodman (2012).

[13] Cf. Peacocke (2008).

[14] Brewer (2011, p. 73) argues that there are limits to the kinds of ways an object can look—and thus limits to perceptual illusion (e.g., a round object cannot look triangular). Yet even within these sorts of boundaries on illusion, there plausibly remain cases that go beyond the limits of tolerance.

[15] Armstrong (1997, Section 4.12) provides a canonical statement of the distinction between determinates and determinables. He takes his lead from Johnston (1921). For a recent influential discussion of the distinction, see Funkhouser (2006).

[16] Cf. Hellie (2011, p. 130; 2014, Section 3); Campbell and Cassam (2014, pp. 86–7).

[17] In this respect I differ both from those like Matthen (2010) who insist that a subject can rationally form beliefs in response to hallucinatory experience, and from those like Campbell (2002) who deny that these beliefs can be rationally formed.

[18] Epistemic Representationalism should not be identified with any particular version of representationalism about perceptual phenomenology. The latter explains a perceptual experience's phenomenal character by appeal to its representational content. By contrast, Epistemic Representationalism concerns the relationship between perceptual experience and rational belief formation.

[19] Cf. Peacocke (1992, Ch. 3) for a seminal treatment of perceptual content in terms of accuracy conditions.

[20] While arguably intuitive, this connection is far from universally recognised (cf. reliabilists about perceptual justification).

[21] My arguments will therefore not apply to those who accept a link between perceptual content and rational belief formation, yet deny that facts about the former are explanatorily prior to facts about the latter. Compare McDowell (1994) and Hellie (2011, 2014).

[22] Non-dogmatist endorsements of the view that we must explain rational restrictions in terms of *some* notion of perceptual content can be found in Lewis (1994), McDowell (1994), and Hellie (2011, 2014).

[23] Cf. Tye (2000), Byrne (2000, 2009), Chalmers (2004), Siegel (2010), and Schellenberg (2018).

[24] Many deny that the rationality of testimonial uptake requires merely the absence of a positive reason to doubt a speaker. I'm not committed to the view here, but for defences see Burge (1993) and Dummett (1994).

[25] In the original formulation, (b) ran as follows: "At the same time, it is also rational for S (on the basis of that same visual experience) to form a belief about the object that perceptually ascribes F."

[26] While avoidance of potential counterexamples might require a more precise characterisation of the boundary, I need only this rough version.

[27] For example, the "Fregean Representationalism" defended by Chalmers (2004) on which perceptual contents are functions from scenarios to extensions. An experience then counts as presenting a given property in virtue of that property being the value of an appropriate function from scenarios to extensions.

[28] Burge (1991) develops a view of this sort; see Martin (2002) for a careful elucidation of this approach. Others might leave out modes of presentation and instead opt for a Russellian view that permits "gaps" in the structured propositions that serve as perceptual contents.

[29] For the parallel point about limits on our perceptual awareness of objects, see Brewer (2011, p. 73).

[30] Similar complaints have been made against Searle's (1983) descriptive-condition-heavy account of perceptual content, for instance by Burge (1991).

[31] Soteriou's proximate targets are Burge (1991) and Martin (2002): Burge endorses a theory of perceptual content that entails Particularity, and Martin argues that Burge's account does not lead to disjunctivism.

[32] We need only replace "hallucinatory experiences" with "illusory experiences". The motivation for the test permits the substitution.

[33] Cf. Armstrong (1955); Brewer (2011); Hellie (2011, 2014); Genone (2014). Hellie (2011, pp. 148–9) dubs Dorothy-type cases as "delusive good" (and later "infragood" in Section 4.1 of Hellie, 2014).

[34] Some extant accounts may serve as such radical alternatives (Cf. Brewer, 2011; Campbell, 2002; Dickie, 2015; Johnston, 2014; Kalderon, 2011).

[35] Cf. Hellie (2011, 2014).

REFERENCES

Armstrong, D. (1955). Illusions of sense. *Australasian Journal of Philosophy*, *33*, 88–106.

Armstrong, D. (1997). *A world of states of affairs*. Cambridge, UK: Cambridge University Press.

Ayers, M. (1974). Individuals without Sortals. *The Canadian Journal of Philosophy*, *4*, 113–148.

Ayers, M. (1997). Is "physical object" a sortal concept? A reply to Xu. *Mind and Language*, *12*, 393–405.

Brewer, B. (2011). *Perception and its objects*. Oxford, UK: Oxford University Press.

Burge, T. (1991). Vision and intentional content. In E. Lepore & R. van Gulick (Eds.), *John Searle and his critics*. Oxford, UK: Blackwell.

Burge, T. (1993). Content preservation. *Philosophical Review*, *102*(4), 457–488.

Byrne, A. (2000). Intentionalism Defended. *Philosophical Review*, *110*, 199–240.

Byrne, A. (2009). Experience and content. *The Philosophical Quarterly*, *59*, 429–451.

Campbell, J. (2002). *Reference and consciousness*. Oxford, UK: Oxford University Press.

Campbell, J. (2006). Does visual reference depend on sortal classification? Reply to Clark. *Philosophical Studies*, *127*, 221–237.

Campbell, J., & Cassam, Q. (2014). *Berkeley's Puzzle*. Oxford, UK: Oxford University Press.

Carrasco, M. (2011). Visual attention: The past 25 years. *Vision Research*, *51*, 1484–1525.

Chalmers, D. (2004). The representational character of experience. In B. Leiter (Ed.), *The future for philosophy*. Oxford, UK: Oxford University Press.

Dickie, I. (2010). We are acquainted with ordinary things. In *New essays on singular thought*. Oxford, UK: Oxford University Press.

Dickie, I. (2014). The Sortal dependence of demonstrative reference. *European Journal of Philosophy*, *22*, 34–60.

Dickie, I. (2015). *Fixing reference*. Oxford, UK: Oxford University Press.

Dummett, M. (1973/1981). *Frege: Philosophy of language* (2nd ed.). Cambridge, UK: Harvard University Press.

Dummett, M. (1994). Testimony and memory. In A. C. Matilal (Ed.), *Knowing from words*, New York, NY: Kluwer.

Evans, G. (1982). In J. McDowell (Ed.), *The varieties of reference*. Oxford, UK: Oxford University Press.

Funkhouser, E. (2006). The determinable-determinate relation. *Nous*, *40*, 548–569.

Genone, J. (2014). Appearance and illusion. *Mind*, *123*, 339–376.

Goodman, R. (2012). Why and how not to be a sortalist about thought. *Philosophical Perspectives*, *26*, 77–112.

Hardin, C. L. (1988). *Color for philosophers*. Indianapolis, IL: Hackett.

Heck, R. (2000). Nonconceptual content and the 'space of reasons. *The Philosophical Review*, *109*, 483–523.

Hellie, B. (2011). There it is. *Philosophical Issues*, *21*, 110–164.

Hellie, B. (2014). Love in the time of cholera. In B. Brogaard (Ed.), *Does perception have content?* Oxford, UK: Oxford University Press.

Johnston, M. (2014). The problem with the content view. In B. Brogaard (Ed.), *Does perception have content?* Oxford, UK: Oxford University Press.
Johnston, W. E. (1921). *Logic, part I*. Cambridge, UK: Cambridge University Press.
Kalderon, M. E. (2011). Color Illusion. *Nous*, *45*, 751–775.
Lewis, D. (1980). Veridical hallucination and prosthetic vision. *Australasian Journal of Philosophy*, *58*, 239–249.
Lewis, D. (1994). Reduction of mind. In S. Guttenplan (Ed.), *Companion to philosophy of mind*. Oxford, UK: Blackwell.
Macpherson, F., & Batty, C. (2016). Redefining illusion and hallucination in light of new cases. *Philosophical Issues*, *26*, 263–296.
Martin, M. (2002). Particular thoughts and singular thought. *Royal Institute of Philosophy Supplement*, *51*, 173–214.
Matthen, M. (2010). Color experience: A semantic theory. In M. Matthen & J. Cohen (Eds.), *Color ontology and color science*. Cambridge, UK: MIT Press.
McDowell, J. (1994). *Mind and world*. Cambridge, MA: Harvard University Press.
Palmer, S. (1999). *Vision science*. Cambridge, UK: MIT Press.
Peacocke, C. (1992). *A study of concepts*. Cambridge, UK: MIT Press.
Peacocke, C. (2008). *Truly understood*. Oxford, UK: Oxford University Press.
Pryor, J. (2000). The skeptic and the dogmatist. *Nous*, *34*, 517–549.
Quine, W. V. (1960). *Word and object*. Cambridge, UK: MIT Press.
Recanati, F. (2012). *Mental Files*. Oxford, UK: Oxford University Press.
Schellenberg, S. (2018). *The Unity of perception*. Oxford, UK: Oxford University Press.
Searle, J. (1983). *Intentionality: An essay in the philosophy of mind*. Cambridge, UK: Cambridge University Press.
Siegel, S. (2010). *The contents of visual experience*. Oxford, UK: Oxford University Press.
Smithies, D. (2011). What is the role of consciousness in demonstrative thought? *Journal of Philosophy*, *108*(1), 5–34.
Soteriou, M. (2005). The subjective view of experience and its objective commitments. *Proceedings of the Aristotelian Society*, *105*(1), 177–190.
Strawson, P. F. (1959). *Individuals*. London, UK: Methuen.
Tye, M. (2000). *Consciousness, color, and content*. Cambridge, UK: MIT Press.
Wiggins, D. (1997). Sortal concepts: A reply to Xu. *Mind and Language*, *12*, 413–421.
Wu, W. (2014). *Attention*. London, UK: Routledge.

How to cite this article: Alford-Duguid D. Thinking through illusion. *Eur J Philos*. 2020;1–22. https://doi.org/10.1111/ejop.12582