capacities, where those capacities are understood in ways that do not presuppose content. I think this means not helping herself to classificatory perceptual capacities in explaining content.

Rutgers University
New Brunswick, NJ 08901-1411
msm331@philosophy.rutgers.edu

Arché Philosophical Institute
University of St. Andrews
St Andrews Fife KY16 9AR
msm331@philosophy.rutgers.edu

References

Ducasse, C.J. 2015. *Truth, Knowledge and Causation*. New York: Routledge. Original publication date: 1969.

McGrath, M. 2016. Schellenberg on the epistemic force of experience. *Philosophical Studies* 173: 897–905.

Wright, L. 2013. Epilogue. In *Functions: Selection and Mechanisms*, ed. P. Huneman, 233–38. Kluwer: Dordrecht.

Accuracy Conditions, Functions, Perceptual Discrimination

BY SUSANNA SCHELLENBERG

I am deeply indebted to Alex Byrne, Jonathan Cohen and Matthew McGrath for their careful, constructive, and penetrating comments on *The Unity of Perception* and I am grateful for the opportunity to clarify my view further.

1. Reply to Byrne

1.1 The capacity to discriminate and single out

Byrne's first set of questions focuses on my central notion of perceptual capacities. He asks why I focus on capacities to discriminate and single out rather than, say, the capacity to know. In response, one reason against focusing on the capacity to know is that one goal of my account is to give an account of perceptual knowledge: it is not clear what explanatory progress would be made if knowledge were analysed in terms of the capacity to know. Such an account seems quite obviously circular. Similarly, it is unclear what

Analysis Reviews Vol 79 | Number 4 | October 2019 | pp. 739–754 doi:10.1093/analys/anz057

explanatory progress would be made if representational content is analysed in terms of capacities to represent.

My goal is to ground representational content, consciousness, evidence and knowledge in more basic and more primitive capacities, rather than to explain them in terms of, say, the capacity to know, the capacity to represent, the capacity to be in a conscious state or the capacity to have evidence. I argue that in employing perceptual capacities to discriminate and single out particulars one gains factive evidence (Ch. 7). And then, in a separate argument, I show that factive evidence is sufficient evidence for knowledge. In the good case, employing discriminatory, selective capacities yields knowledge (Ch. 9). So evidence and knowledge are grounded in non-normative, non-reliable and naturalistic properties of perception. This is an ambitious project and it is a project that takes a new stab at naturalizing epistemology and explaining epistemic features with more primitive properties of the mind. I argue, moreover, that perceptual content and consciousness are constituted by the perceptual capacities employed. Thus, the larger goal is to ground content, consciousness, evidence and knowledge in a basic property of perception.

But why focus on the capacity to discriminate and single out (rather than say the capacity to classify or the capacity to attribute properties to objects)? The reason is that discrimination is necessary for perception (whereas classification and attribution are not). In cognitive psychology and neuroscience, perception is fundamentally understood to be a matter of discriminating. In grounding content, consciousness, evidence and knowledge in the capacity to discriminate and single out, my account is grounded in research in cognitive psychology and neuroscience.[1]

Byrne asks, moreover, why I need the capacity to discriminate *and* single out rather than merely the capacity to discriminate. He suggests that singling out α simply amounts to attending to α and if that is right then capacitism would posit that one cannot perceive α without attending to α. In response, Byrne is right that if singling out α were to require attending to α, then the account developed would imply that one cannot perceive something without attending to it and that would be a problem for capacitism. But singling out α does not require attending to α and indeed I argue that attention should be understood in terms of perceptual discrimination (see Schellenberg forthcoming).

To explain why singling out is a necessary element of the account, consider Raul who is looking at a surface that is partly red and partly blue. He can see the scene in front of him by employing the perceptual capacity to discriminate and single out blue from other colours. Let's call this capacity C_{blue}. Alternatively, he can see the scene in front of him by employing his perceptual capacity to discriminate and single out red from other colours. Let's call

1 See my response to Cohen, for a discussion of why I focus on the capacity to discriminate and single out rather than other low-level capacities.

this capacity C_{red}. He could also see the scene by employing both C_{blue} and C_{red}. Regardless of whether he employs C_{blue} or C_{red}, he will be discriminating between an instance of red and an instance of blue. However, if he employs C_{blue} he will discriminate a blue-instance from a red-instance and single out the blue-instance. Likewise, if he employs C_{red} he will discriminate a red-instance from a blue-instance and single out the red-instance. If he sees the scene by employing both C_{blue} and C_{red}, he will single out both the blue and the red-instance. These differences generate a host of repercussions for his perceptual content, the phenomenal character of his perceptual state, and the evidence and knowledge he gains. For example, if he employs his capacity C_{blue} and thereby discriminates and singles out the blue patch from its surround, he will gain evidence of the presence of that blue patch (but not evidence of the presence of the red patch – assuming he did not also employ his capacity C_{red}).

1.2 Properties versus property-instances

Byrne's second set of question concerns my argument that we perceive property-instances rather than properties. As Byrne points out, there is a well-established view – established by two millennia of rich philosophical argument, culminating in work of Byrne himself – on which the only particulars we perceive are objects (see Byrne 2001). As Byrne puts it: 'We perceive the cup's shape (understood as a universal, something shared with other similar cups) because we are causally affected by a particular instance of that shape, the one qualifying the cup in front of us. But we do not *perceive* this particular property-instance: the only particular we perceive is the cup.' Byrne notes, moreover, that vision scientists talk of properties, rather than property-instances.

In response, it is no doubt true that vision scientists talk of properties. Typically, however, they talk of seeing features, and in doing so, they seem to assume that these features are mind-independent particulars (rather than abstract entities). When they do talk of properties, they similarly do so as if those properties were particulars (and so property-instances), rather than abstract entities. If this is true, then vision scientists' talk of properties is mere talk: what they really mean is property-instances.

What hinges on the matter? Properties are not spatio-temporally located and not causally efficacious. Due to this, it is mysterious how they could be the kind of things one perceives. On this basis, I argue that what we perceive are property-instances (for the details of this argument, see Schellenberg 2018: 145–50). It is unmotivated to argue that we discriminate objects, events and property-instances in our environment, but then only perceive objects – and not property-instances, despite discriminating them. Byrne goes on to argue that on my view there is an abundance of particulars in our environment that we discriminate and single out and thereby perceive. It is true that in the typical case of perception, there is according to capacitism a

plethora of particulars we discriminate and single out and thereby perceive. This is a good consequence of the view. After all, there is a plethora of objects, events and property-instances in our environment.

Byrne ends his discussion of my commitment to property-instances by pointing out that 'even if we do perceive property-instances, perceptual systems treat them very differently from particular objects and events'. This is true and I do not argue that the perceptual system treats objects and property-instances the same. We can all agree that there are deep differences between objects, events and property-instances. For example, objects and events can instantiate properties but properties cannot instantiate objects and events. While acknowledging these differences, my point is that there are two critical respects in which property-instances are on a par with objects and events: all three are (i) external and mind-independent particulars that we (ii) discriminate and single out in perception. Failing to acknowledge these two critical similarities has led to confusions and to the prevalence of over-intellectualizing attributional views over discriminatory views of perception. I contend that perception is fundamentally a matter of discrimination and not fundamentally a matter of attribution (for a dissenting view, see Burge 2010 and Block 2014).

1.3 Constitution and the particularity argument

Byrne argues that Premise 2 of the particularity argument is supported by a tautology. Premise 2 has it that if a subject *S* discriminates and singles out the particular α (as a consequence of being perceptually related to α), then *S*'s perceptual state *M* brought about by being perceptually related to α is constituted by discriminating and singling out α. I do not agree that this is a tautology and indeed, I believe, the best way for the generalist to avoid the conclusion of the particularity argument is to find a way to reject Premise 2. But let's assume for the sake of argument that the premise is supported by a tautology. It only strengthens my conclusion if it is derived from trivially true premises. So even if Byrne were right here (and I do not think he is) this would not negatively affect my argument. The important point is that the conclusion is not trivial and indeed rejected by generalism – one of the most widely held views among philosophers of mind. If this conclusion follows from premises that turn out to be trivially true, so much the better.

Byrne does not deny this, of course, and indeed he sympathizes with particularism. He points out, however, that the generalist could accept that perceiving α requires discriminating and singling it out but nonetheless avoid my particularist conclusion by denying that *S* is in perceptual state *M in virtue of* discriminating and singling out α. The generalist could insist that *S* is in *M* in virtue of *S*'s intrinsic state *B*.

In response, the problem for the generalist is, as Byrne acknowledges himself, that this involves denying that seeing α is a 'perceptual state' strictly speaking. So Byrne and I are in agreement on this issue. In addition to the

problem Byrne mentions, a further problem for the generalist is that *S*'s intrinsic state *B* is arguably constituted by discriminating and singling out particulars. So for this generalist strategy to work, the generalist would have to drive a wedge not only between discriminating α and S's perceptual state of seeing α but also between discriminating α and *S*'s intrinsic state *B*. By driving a wedge between perceptual discrimination and the nature of the perceptual state as well as between perceptual discrimination and the intrinsic state *B*, generalism seems now at best *ad hoc*.

1.4 Prosthetic eyes, flukes and luck

Byrne ends his comments by pointing out that in the visual case discriminating and singling out a cup amounts to seeing the cup and suggests that I simplify matters and analyse the capacity in question as a capacity to see rather than a capacity to discriminate and single out. In response, while the relevant capacity ultimately enables one to see and so could be described as a capacity to see, there is a more basic level at which to describe it: namely the level at which one discriminates and singles out particulars that constitutes seeing the cup. My point is that there is no such thing as brutely seeing. There is always something one does in virtue of which one sees, namely discriminate and single out mind-independent particulars in view. So it does not seem apt to stop the analysis at the level at which on has diagnosed the perceptual state as one of seeing the cup and the relevant capacity as the capacity to see.

To support his suggestion, Byrne considers Lewis's case of a prosthetic eye with a loose wire (Lewis 1980). It is a case in which seeing allegedly happens via a fluke without the relevant subject possessing the capacity to see objects. In response, I argue regarding such a case that in the moments that the wire is connected, the subject sees, has perceptual states with singular content, and gains factive evidence. I would even say she gains knowledge – leaving open that she might not have more high-level reflective knowledge. Similarly, I argue that Henry who sees the one and only barn in barn façade county (that is a county that is filled with barn façades rather than actual barns) gains knowledge of the barn he sees. This might be a surprising verdict on such cases, but arguably it is the right verdict. After all, Henry sees the barn and the subject with the prosthetic eye gains information about her environment when the wires are connected.

2. Reply to McGrath

McGrath is the kind of commentator who raises problems for my view and then gives an excellent solution to those problems.

2.1 The magic of perceptual capacities

McGrath's first and main question is 'what it is about perceptual capacities, unlike many other dispositions, which enables them to ground content,

phenomenology and epistemic force? Their functions. But how do their functions accomplish this?' The answer to this question is different for content, consciousness and evidence. The perceptual content argument answers McGrath's question regarding content. The phenomenal evidence and factive evidence arguments answer the question regarding evidence. The mental activism argument answers the question regarding consciousness.

Since McGrath's discussion zeros in on content and evidence, I will focus on those. Chapter 5 is devoted to arguing that employing perceptual capacities constitutes perceptual content. The argument is complex, but the pivotal idea is that the employment of perceptual capacities generates a perceptual state that is *repeatable* and has *accuracy conditions*. Being repeatable and having accuracy conditions are jointly key signatures of representational content. The very same perceptual capacity C_α can be employed to single out particular α_1 or to single out particular α_2, where α_1 and α_2 are both particulars of the type that the perceptual capacity functions to single out. As I argue in Chapter 1, if one singles out α_1 rather than α_2, one is in a distinct perceptual state, namely, a perceptual state that is constituted by α_1 (and not by α_2). This is the case even if α_1 and α_2 are qualitatively identical. So the same perceptual capacity can be employed in distinct environments and yield distinct perceptual states. Moreover, the same perceptual capacity can be employed to single out α_1 at time t_1 and at time t_2 and thus yield the same perceptual state at t_1 and t_2. If this is right, then there is a repeatable element that is constitutive of perceptual states, namely, the perceptual capacities employed and, moreover, employing perceptual capacities generates a perceptual state that has a repeatable element.

I agree with McGrath that the accuracy of a perceptual state cannot be a matter of whether one discriminates more or fewer property-instances in the scene, nor can it be a matter of whether one singles out particulars more or less correctly. Indeed, I want to argue that one cannot single out a particular incorrectly, rather one either singles out the particular or fails to do so. When one is perceptually related to a scene, one employs perceptual capacities which may or may not function to single out the particulars present. If I employ my capacity to discriminate and single out red from other colours in an environment in which there is no instance of red, the content of my experiential state will be inaccurate in that respect. Insofar as a perceptual capacity is repeatable and insofar as one either singles out the particular one purports to single out or one fails to do so, employing perceptual capacities generates a perceptual state that is repeatable and has accuracy conditions. So employing perceptual capacities yields perceptual states that exhibit key features of representational content: it yields something that is at least in part repeatable and that can be accurate or inaccurate. With a few plausible further assumptions, these considerations establish that employing perceptual capacities yields perceptual states with content. In Chapter 5, I defend this thesis in more detail.

The reason why employing perceptual capacities furnishes the perceptual states they constitute with epistemic force is due to a different aspect of perceptual capacities. Before explaining further it should be noted first that my concern is meta-epistemological. The question I am addressing is not the standard epistemological question of how perception justifies beliefs (assuming that perception has epistemic force), but rather the meta-epistemological question of why perception has epistemic force in the first place. Imaginations, beliefs, hopes and fears do not have epistemic force, but perception does. Why is that?

The key idea for why perception has epistemic force is that perceptual states are systematically linked to external, mind-independent particulars of the type that the perceptual state is of in the good case. Perceptual states are systematically linked to what they are of in the good case since they are constituted by perceptual capacities employed and the particulars thereby discriminated and singled out. So the successful employment of perceptual capacities relates perceivers to these particulars and thus the perceiver gains factive evidence of the relevant particulars. This aspect of my account is akin to a knowledge-first view, but it is an account on which the basic level of analysis is the capacities employed (not knowledge or any other such epistemic property). Moreover, in contrast to knowledge-first views, I argue that even when employed in the bad case perceptual capacities have the function of discriminating and singling out particulars, since the perceptual capacities employed in the bad case are explanatorily and metaphysically parasitic on their employment in the good case.

There is an *explanatory primacy* of the good over the bad case, since one can give an analysis of the perceptual capacities employed in the bad case only by appealing to their role in the good case. Licensing this explanatory primacy there is a *metaphysical primacy* of the good over the bad case. The metaphysical primacy is captured by the asymmetry condition on perceptual capacities that I develop in Chapter 2: the employment of a perceptual capacity C_α in cases in which C_α fulfils its function is metaphysically more basic than the employment of C_α in cases in which C_α fails to fulfil its function. Perceptual capacities function to single out particulars. They do not function to fail to single out particulars. Due to this function, I argue, we have phenomenal evidence (regardless of whether we are perceiving, hallucinating or suffering an illusion) in virtue of employing capacities with a certain function.

So the crucial difference between perception, on the one hand, and beliefs, hopes, fears and imaginations on the other, is that perceptual capacities function to single out the particulars to which we are perceptually related, whereas the capacities employed in those other mental states do not necessarily have this function. It should be noted that while I argue that we have additional factive evidence in perception (over the phenomenal evidence we have regardless of whether we are in the good or the bad case), I am not an

epistemological disjunctivist since on my view factive evidence is not reflectively accessible to the perceiver.

Now McGrath questions whether this is enough to infuse perceptual states with epistemic force (see also McGrath 2016). He notes that I neither appeal to any historical feature nor reliability features to help my case and suggests that if we cannot appeal to such features, then all we are left with is the idea that capacities are a certain kind of disposition. McGrath then argues convincingly that dispositionality is not up to the job.

In response, we do not need to accept this dichotomy of either appealing to history or to dispositions. I agree with McGrath that dispositionality is not up to the job, but capacities are not dispositions. To show why, it will help to take a closer look at the difference between dispositions and capacities. It should be noted first, however, that the notion of capacities is used in many different ways and I do not aim to police how we should use the term. There is a cluster of concepts that are closely related yet distinct: capacities, abilities, dispositions, competences, powers and skills. I develop a particular notion of capacities according to which they are repeatable, fallible and — in virtue of their function to discriminate and single out mind-independent particulars — they are systematically linked to those very environmental particulars. While dispositions are triggered passively, capacities are a kind of mental tool that can be employed more or less deliberately. Dispositions are properties picked out by predicates such as 'is fragile' or 'is soluble'. In contrast to dispositions, capacities are not merely triggered when the right conditions are met. Another way of putting the same point is to say that the manifestation conditions of capacities (as I understand them) are not the trigger-manifestation conditions of dispositions. There is more to the employment of a capacity than being triggered by the right stimuli. Thus, the manifestation condition of a capacity is not due merely to causal facts. What dispositions and capacities have in common is that they exist even when they are not manifested.

This brief analysis does not do justice to the subtle differences between dispositions and capacities. But for present purposes, this will have to suffice to justify why there are important differences between capacities and dispositions. The central point is that in contrast to dispositions, the manifestation condition of a capacity is not due merely to causal facts and due to this the function of perceptual capacities is up to the job of grounding the epistemic force of perception.

2.2 The accuracy conditions of perceptual content

McGrath's third set of questions concern how the function of perceptual capacities manages to determine contents with the right accuracy conditions in the case of illusions and specifically how gappy contents are bound in the right way.

In response, the content and accuracy conditions of mental states are often equated. This, I argue, is a mistake. Content determines how the world would have to be for the content to be accurate. So, the accuracy condition of the content 'That white cup is to my right' would according to capacitism be:

> The content 'That white cup is to my right' of a perceptual state brought about by being perceptually related to that white cup to my right is accurate if and only if that white cup is to my right.

More generally the accuracy conditions of content *c* can be specified as follows:

> The content *c* of a perceptual state brought about by being perceptually related to environment *E* is accurate if and only if *E* is the way *c* represents *E* to be.

The fact that a content is gappy implies that the content is necessarily inaccurate insofar as a gappy content could never make an accurate claim about the world. To motivate this, consider two ways in which a content can be inaccurate. One is for the content to make a claim about the environment that is not accurate. A second way is for it to fail to make an accurate claim about the environment. To illustrate this second sense of inaccuracy, suppose that I claim that Pegasus lives in my apartment. This claim is inaccurate. Given that 'Pegasus' does not refer, the inaccuracy in question is that I have failed to make an accurate claim about who lives in my apartment. If inaccuracy is understood in this second way, then an illusion or hallucination can have a gappy content and nonetheless be inaccurate.

Now McGrath raises a specific problem about how the content of a state in which one, say, sees α but misperceives it to be instantiating a property it does not in fact instantiate determines the right kind of accuracy conditions. I argue that the content of such a state will be:

$< \mathrm{MOP}_{r\alpha}(\alpha_1),\ \mathrm{MOP}_{r\pi}(\_\_) >$

where $\mathrm{MOP}_{r\alpha}(\alpha_1)$ is a singular mode of presentation of the cup α_1 that is the product of employing a perceptual capacity that functions to single out the kind of object under which α_1 falls and $\mathrm{MOP}_{r\pi}(\_\_)$ specifies the property π that this object would instantiate, where the experience a perception rather than an illusion. McGrath's question is how $\mathrm{MOP}_{r\pi}(\_\_)$ is bound with $\mathrm{MOP}_{r\alpha}(\alpha_1)$.

In response, $<\mathrm{MOP}_{r\alpha}(\alpha_1), \mathrm{MOP}_{r\pi}(\_\_)>$ does not specify which property-instance has to be present but only which kind of property-instance. It is crucial that it does not specify what specific property-instance has to be present. After all, any instance of the relevant property will do. Another way of expressing what is going on here is that we have a sense without reference. There is enough structure for the content to specify what kind of

property-instance would have to be present, but not enough to specify which specific property-instance has to be present.[2] Also contrary to what McGrath suggests $MOP_{r\pi(\_)}$ is not akin to an open sentence such as: '__ is π'. Gappy contents are not open sentences, they are rather token contents that are instances of the same content type as token singular contents.

The binding problem can be solved in many different ways and my view is compatible with a range of different solutions. McGrath considers the option of appealing to classification to solve the problem. I agree with McGrath that this is not a good solution. After all, classification is more high-level than discrimination and while it is plausible that discrimination is necessary for perception it is not plausible that classification is necessary. It should be noted that McGrath discusses classification in a way such that classification would occur on the basis of perceptual content. One does not need to understand classification in this way. But even if one understands classification such that it does not occur on the basis of perceptual content, it would not be a good solution to McGrath's worry. So McGrath and I are in agreement here.

How then should we solve the binding problem? One option is to insert markers in the content specifying that $MOP_{r\alpha}(\alpha_1)$ and $MOP_{r\pi(\_)}$ are bound. How this solution goes is more obvious if one considers a many-properties case of hallucination. Consider a subject who at t_1 hallucinates a white square cup and a red round cup and then at t_2 hallucinates a white round cup and red square cup. These hallucinatory states differ in content and one can mark their difference as follows:

$$(c - t_1) < [MOP_{r\alpha}(\_), MOP_{r\mathrm{WHITE}}(\_), MOP_{r\mathrm{SQUARE}}(\_)] \,\&\, [MOP_{r\mathrm{a}}(\_), MOP_{r\mathrm{RED}}(\_), MOP_{r\mathrm{ROUND}}(\_)] >$$

$$(c - t_2) < [MOP_{r\alpha}(\_), MOP_{r\mathrm{WHITE}}(\_), MOP_{r\mathrm{ROUND}}(\_)] \,\&\, [MOP_{r\mathrm{a}}(\_), MOP_{r\mathrm{RED}}(\_), MOP_{r\mathrm{SQUARE}}(\_)] >$$

But what correlates with these markers on the level of employing perceptual capacities? McGrath is right that co-location of the particulars singled out cannot do the job. After all, no particulars are singled out. What we need rather than co-location is co-directedness.

2 Similarly, the content of a hallucination

$<MOP_{r\alpha}(\_), MOP_{r\pi}(\_)>$

specifies the kind of object that would have to be present for the content to be accurate without specifying which particular of that kind; and it specifies the kind of property-instance that would have to be present for the content to be accurate again without specifying which particular of that kind.

3. *Reply to Cohen*

3.1 Hegemony of distal discrimination

Cohen questions why I focus on personal-level discriminatory, selective capacities rather than the many other capacities that play a role in perception and the mind at large.[3] There are a lot of different capacities at play in perception. My account zeros in on one kind: capacities to discriminate and single out mind-independent particulars in our environment. In doing so, I do not deny the existence of other kinds of capacities, nor do I argue that there is a discontinuity between discriminatory, selective capacities and those other kinds. Indeed, I favour analysing the mind in terms of capacities on multiple levels and domains. As Cohen notes, the reason I zero in on discriminatory, selective capacities is that these capacities are necessary for perception and because they play the pivotal role of grounding perceptual content, consciousness and evidence.

Cohen asks whether one can perceive an object one does not discriminate. In response: no, one cannot, at least not on the view I put forward. In the typical case of perception, there are lots of particulars that are present in one's visual field that one does not discriminate and single out and thus does not perceive. Consider Miriam, who is looking at a Richter painting in a museum. There is a speck of dust on the frame of the painting. The speck of dust is in her visual field and large enough that she could discriminate and single it out. But she is focused on the painting and thinking about a paper she read earlier that day. So despite the speck of dust being in her line of sight, she does not discriminate it and thus does not see it. There are many reasons why one may not discriminate and single out (and thus fail to see) a particular in one's field of vision. One is that one *does not* discriminate and single it out, even though one could. Another is that one *cannot* discriminate and single it out, perhaps because the particular blends in perfectly with its surround or because it is too small to be discernible to the perceiver. The important point for present purposes is that if one does not discriminate and single out a particular (whatever the reason for that may be), one does not perceive the particular.

Let's consider Cohen's case in which one has a chameleon in view, but does not notice the chameleon since it blends in perfectly with its surround. So one does not notice that one is looking at a chameleon despite seeing the coloured surface that (unbeknownst to one) happens to be the skin of the chameleon. In this case, capacitism posits that one perceives a property-instance that the chameleon instantiates, without perceiving *that* it is instantiated by the chameleon. One does not discriminate where the chameleon ends and where the surrounding leaves begin since this property-instance meshes perfectly with

3 For an explanation of why I do not focus on capacities to know or capacities to represent, see my response to Byrne's comments.

the colour of the surrounding leaves. Nonetheless, one perceives aspects of the chameleon – one just does not perceive them *as* aspects of the chameleon.

So contrary to what Cohen is suggesting, capacitism has it that in such a case one discriminates a coloured surface which (unbeknownst to one) happens to be at least in part the coloured skin of the chameleon. Cohen is right that capacitism has it that one does not see the object instantiating this colour – certainly one does not see it *as* that object. That is the right verdict on such a case. After all, having an object in view, is not sufficient for perceiving the object. Not perceiving an object is, however, compatible with perceiving some of its property-instances – properties instantiated by the object. So contrary to Cohen, my view does not imply that one does not perceive *tout court*. Indeed, cases like this are one reason why I treat perception of objects and property-instance as largely on a par and, at least to some extent, as independent of one another: one can perceive an object without perceiving all its property-instances; and one can perceive a property-instance without perceiving the object that instantiates the relevant property.

Capacitism can easily explain cases in which an object or a property-instance is in full view, but one nonetheless fails to perceive them. And the view can easily explain cases in which one perceives objects without perceiving at least some of its property instances as well as cases in which one perceives a property-instance without perceiving the object that instantiates the relevant property. Discrimination provides for a clear criterion by means of which one can separate the particulars one perceives from those that one does not.

3.2 Undergeneration and non-visual senses

Cohen raises the issue of discrimination of mind-independent particulars in non-visual cases where the particular is a property-instance (rather than an object), but the perception of this property-instance depends not just on a distal quality but also on something narcissistic or subject-involving: thermoreception, flavour perception, phenome perception, to name just a few.

In response, capacitism has the tools to account for such cases. After all, perceptual capacities contribute to the narcissistic or subject-involving element. Any case of perception involves two elements: a particular element (the mind-independent, external particular) and a general element (the perceptual capacities employed). Which perceptual capacity one employs to discriminate and single out any given mind-independent particular (or distal quality) will vary from situation to situation and will depend not just on the perceiving subject but also the larger surround. As I argue in Part III of the book, perceptual consciousness is determined entirely by the perceptual capacities employed and so by the narcissistic, subject-involving element with which Cohen is concerned.

So while there is a staunch externalism running throughout my project insofar as any episode of perception is constituted by the mind-independent

particulars perceived, there is – given the central role of perceptual capacities – no lack of subject-involving contributions.

3.3 Swampman and the function of perceptual capacities

Cohen asks what grounds the natural function of perceptual capacities, noting correctly that I argue that their ground lies neither in the intentional states of users, theorists or interpreters, nor in the phylogenetic or ontogenetic aetiology of users. In response, it is important to note first that as it so happens we have the capacities we have due to our phylogenetic or ontogenetic history. We can all agree on that. My point is that once any given capacity has the function it has, we can talk about that function without appeal to its history. A view of natural function allows that. The work functions do in my view does not in any way depend on the history of the function.

To explain, it is important to distinguish three distinct questions: What function does capacity C_α have? How did C_α come to have that function? How can we tell what function C_α has? Aetiologists seem to think that one cannot answer the first question without answering the second. But that is simply false. In answering the first question they are, moreover, often motivated by the third question. This third question is no doubt interesting but it is not a question that my project needs to answer.

For my purposes, we can say that any capacity has a function, remaining neutral on how it came to have that function (second question) and how we can tell what function it has (third question). The neutrality on these two questions is not due to laziness. It is due to the fact that the content, consciousness and evidence that are constituted by the capacities employed in no way depend on how those capacities came to have their function or how we can tell what functions they have.

To motivate this, consider mathematical functions. There is a story that could be told about how '+' came to have its function, but this story is not told when teaching addition. For good reasons. It is irrelevant. What matters is that '+' has a specific function. While there is no perfect analogy between the function of perceptual capacities and mathematical functions, in this respect the two are the same.

As Cohen notes, it is precisely because phylogenetic and ontogenetic history is irrelevant according to capacitism that I can say that swampman has the very same perceptual states, the very same evidence, and the very same phenomenal character as Davidson would have had (had he not been tragically killed by a bolt of lightning). That is right. And that is one of many advantages of my view of aetiological views.

3.4 The individuation condition of perceptual capacities

Cohen argues that according to capacitism 'one episode [of perception] can be counted as an instance of distinct, overlapping capacities', that one employs all capacities one could employ in any given episode of perception, and

that 'every episode that is a successful employment of a capacity C_1 is, additionally, a baseless employment of a distinct, overlapping capacity C_2'. In response, this is not my view and capacitism does not have the implications Cohen suggests it does. To show why, consider Dylan who sees a patch of scarlet. She could see this patch by employing her capacity to discriminate and single out scarlet. Or she could see the patch by employing her capacity to discriminate and single out red. Let's say that at time t_1, she employs her capacity to discriminate and single out scarlet (C_{scarlet}), and at time t_2, she employs her capacity to discriminate and single out red (C_{red}). Her perceptual state at t_1 and at t_2 are both accurate – after all, scarlet is a shade of red. They are simply more or less fine-grained.

While I argue that the same environmental particular can be discriminated and singled out by a range of capacities, the fact that Dylan could employ a range of different perceptual capacities to discriminate and single out the patch of scarlet, does not imply, as Cohen says that any specific episode of perception 'can be counted as an instance of distinct overlapping capacities'. After all, at any given moment, Dylan employs the specific capacities she is employing. At t_1, she employs C_{scarlet}; and at t_2, she employs C_{red}. Since she is employing different capacities at t_1 and t_2, her perceptual state at t_1 is distinct from her perceptual state at t_2.

So capacitism does not posit that an episode of perception is an instance of employing every capacity that could be employed in that episode. It posits that an episode of perception is an instance of the very capacities employed and no other. While the relevant perceiver could have employed other capacities, the content and phenomenal character of her perceptual state is constituted by the capacities she, in fact, employs (not the capacities she could have employed). Now there is a question as to how one can know which capacities a perceiver is employing. But as I argued in Section 3.3 above, that is a question that does not need to concern us here.

3.5 Is there a proliferation worry?

Cohen raises the problem that 'Snow is white' is true iff snow is white but also iff snow is white and $2 + 2 = 4$. Fregeans can circumvent this proliferation problem. After all, opaque contexts of 'snow is white' and 'snow is white and $2 + 2 = 4$' differ. Thus, the Fregean can block the proliferation. Capacitism is Fregean to its core. Perceptual capacities are the psychologistic flipside of Fregean modes of presentations. As Fregean modes of presentation are ways of grasping things, perceptual capacities are ways of grasping things in our environment. As Fregean modes of presentation have a cognitive significance that can differ even as the reference remains the same, employing perceptual capacities constitutes mental states that differ regarding their phenomenal character, content and evidence even as the particular singled out remains the same.

I argue that perceptual content is constituted by the perceptual capacities employed and the particulars thereby singled out. The contents are singular

de re modes of presentations. For the same reason as articulated in response to Cohen's worry about ubiquity of illusion, there is no sense in which perceivers employ infinitely many perceptual capacities. While for any typical scene of perception they could employ a range of different capacities to discriminate and single out the environmental particulars, at any given moment they employ specific perceptual capacities and their perceptual state is constituted by those (and not by the any others they could have employed). Thus, capacitism circumvents the proliferation worry.

3.6 The possession conditions of perceptual capacities

Cohen takes issue with my specification of the possession conditions of perceptual capacities. They are specified as follows:

> *Possession Condition:* A subject *S* possesses a perceptual capacity C_α if and only if the following counterfactual is true of *S*: *S* would be in a position to discriminate and single out a particular α_1, where α_1 is any particular of the type that C_α functions to discriminate and single out, if *S* were perceptually related to α_1, (i) assuming *S* is perceptually capable (awake, alert etc.), (ii) assuming no finking, masking, or other exotic case obtains, and (iii) where *S* being perceptually related to α_1 means that (a) the situational features are such that α_1 is perceivable by *S* (good lighting conditions etc.), (b) *S* has the relevant sensory apparatus that allows her to gain information about α_1, and (c) *S* is spatially and temporally related to α_1 such that *S* is in a position to gain information about α_1 via her sensory apparatus. (40)

Cohen's worry is that the right-hand side of this condition cannot specify the full range of possible cases. In response, the right-hand side specifies that a subject and the environment must be such that the subject can discriminate and single out a relevant particular. It rules out a range of extreme cases in which the subject or the environment is such that the subject cannot discriminate and single out a relevant particular. The first qualification (i) rules out extreme cases in which the subject is incapacitated and as a consequence is not, in that particular moment, in a position to employ the capacity she possesses (i). Similarly, the qualification (ii) that no finking, masking or other exotic cases obtain rules out extreme and recherché cases in which the subject mysteriously loses her capacity from one moment to the next (because, say, the world explodes whenever she is about to employ the relevant capacity such that she is never in a position to discriminate and single out a relevant particular).[4]

4 The inference from a claim about perceptual capacities to a counterfactual fails in such cases. However, all the standard ways of fixing the disposition-to-counterfactual inference can be exploited for the capacity-to-counterfactual inference (see Lewis 1997). Finding a formulation of the capacity-to-counterfactual inference that is indefeasible in light of all possible finking, masking, and similarly exotic cases would be a project of its own.

The rest of the condition (iii) specifies what it takes to count as being perceptually related to a particular. Here again the point is to rule out extreme cases such as there being no source of light in a visual case of perception (a), the subject failing to have the relevant sensory apparatus (b) or the particular α_1 being so far away that it is no longer perceivable by the subject (c). The parts in brackets are no more than examples and, as such, will naturally differ depending on whether one is seeing, hearing, touching, smelling, tasting or perceiving α_1 in some other sensory mode. Indeed, the details of what it takes for the situational features to be such that α_1 is perceivable by the subject will depend on α_1. But regardless of the nature of α_1 certain situational features must be met. For example, I cannot see the coffee cup on my table if there is no source of light whatsoever. I hope to have shown that the condition is general enough that it can accommodate the full range of possible cases.

Rutgers University
New Brunswick, NJ
08901-1411, USA
susanna.schellenberg@rutgers.edu

References

Block, N. 2014. Seeing-as in the light of vision science. *Philosophy and Phenomenological Research* 89: 560–72.

Burge, T. 2010. *Origins of Objectivity*. Oxford: Oxford University Press.

Byrne, A. 2001. Intentionalism defended. *Philosophical Review* 110: 199–240.

Johnston, M. 1992. How to speak of the colors. *Philosophical Studies* 68: 221–63.

Lewis, D. 1980. Veridical hallucination and prosthetic vision. *Australasian Journal of Philosophy* 58: 239–49.

Lewis, D. 1997. Finkish dispositions. *Philosophical Quarterly* 47: 143–58.

Martin, C.B. 1994. Dispositions and conditionals. *Philosophical Quarterly* 44: 1–8.

McGrath, M. 2016. Schellenberg on the epistemic force of experience. *Philosophical Studies* 173: 897–905.

Schellenberg, S. 2018. *The Unity of Perception: Content, Consciousness, Evidence*. Oxford: Oxford University Press.

Schellenberg, S. forthcoming. Attention and perceptual discrimination.

Therefore, I will here work on the independently plausible assumption that no such exotic cases obtain. For a discussion of masking, see Johnston 1992; for a discussion of finking, see Martin 1994.