
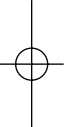




8

The Epistemic Conception of Hallucination

Susanna Siegel





Disjunctivism about perception is a view about the relation between veridical experiences and hallucinations. Suppose that I see a green cube, and my experience is veridical—no illusion or hallucination is involved.¹ The veridical experience, according to disjunctivists, includes as constituents the bit of the world that is perceived and the perceptual relation that the subject bears to it, and these constituents characterize what it is like to see the green cube. The phenomenal character of my veridical experience, they hold, is thus constituted by facts that could not obtain if I were hallucinating. Disjunctivists typically grant that there is a corresponding hallucination in which I could not know just by introspection that I was not veridically seeing a green cube. But they hold that the most basic characterization of what it is like to have such a hallucination—its most basic phenomenal character—must differ from that of the veridical green-cube experience. The main opponents of disjunctivism are *common-kind theorists*, who hold that the most basic phenomenal character in both experiences is the same.

Although the primary commitment of disjunctivism concerns veridical perception, rather than hallucination, disjunctivists aim to respect two central assumptions about hallucination: that they can ‘match’ veridical perceptions, and that they can differ phenomenally from one another. Both assumptions are commonplaces in their own right, and the first is presupposed by the debate between disjunctivism and the common-kind theory. Common-kind theories accommodate both assumptions straightforwardly. ‘Matching’ hallucinations and non-hallucinations share the same most basic sort of phenomenal character, and variation among hallucinations is variation in that same sort of phenomenal character. Disjunctivists, in contrast, cannot

For helpful discussion and criticism, many thanks to Matthew Boyle, Alex Byrne, Bill Fish, Adrian Haddock, Benj Hellie, Heather Logue, Fiona Macpherson, M. G. F. Martin, Ram Neta, Bernhard Nickel, and Stephen White. Special thanks to David Chalmers, Matt Soteriou, and Scott Sturgeon for so many helpful conversations and for comments on multiple drafts.

¹ Throughout this essay, I’ll use ‘veridical experience’ to pick out instances of perception that are not hallucinations or illusions. So there will be no such thing as a veridical hallucination in this sense of ‘veridical’.





accommodate the assumptions in this way, since they deny that matching experiences share the same basic phenomenal character. To respect these assumptions consistently with disjunctivism, some other conception of hallucination is needed.

Recently, some disjunctivists, including M. G. F. Martin and William Fish, have developed an epistemic conception of hallucination. According to this conception, there is nothing more to the phenomenal character of hallucination of a green cube, besides the fact that the hallucination is indiscriminable from a veridical perception of a green cube as such. No further mental properties underlie its being so indiscriminable. Fish (this volume) describes what he takes to be the essence of hallucination, quoting Martin (2004: 72) at the end:

So the essence of hallucination—what distinguishes hallucinations as a class from other mental states—lies in their being indistinguishable from veridical perceptions, not in some antecedently identifiable feature of the state. This is why, “when it comes to a mental characterization of the hallucinatory experience, nothing more can be said than the relational and epistemological claim that it is indiscriminable from the perception”.²


Martin and Fish offer more detailed treatment of hallucination than earlier defences of disjunctivism offered. The main idea, however, that ‘matching’ should be given a purely epistemic construal is central to disjunctivism itself. Disjunctivism takes facts about matching experiences to be facts about how those experiences seem to the subject, rather than taking those facts to reflect any deeper similarity between the matching experiences themselves. We can thus assess the plausibility of disjunctivism by assessing the plausibility of the conception of hallucination on which it depends.


In this essay, I argue against the epistemic conception of hallucination. In section 1, I state some central theses about hallucination put forward by Martin, who has done the most to develop the epistemic conception. In section 2 I introduce a notion of indiscriminability that figures in the epistemic conception. In section 3, I argue that the epistemic conception falters with its treatment of cognitively unsophisticated hallucinators. In section 4, I introduce Fish’s version of the epistemic conception and raise two objections to it. In section 5, I argue that neither version of the epistemic conception has a promising account of what hallucinators can know on the basis of introspection. I conclude that the prospects for a plausible disjunctivist theory of hallucination are not promising. Although much of the discussion focuses on commitments made explicitly by Martin and Fish, my main aim throughout is to explore the space of possible disjunctivist proposals about hallucination and their pitfalls, rather than to locate Martin’s or Fish’s exact positions in it.

1 DISJUNCTIVISM

Return to the case in which I veridically see the green cube. My visual experience in this case is indiscriminable from a veridical perception of a green cube. This follows

² See also Martin (2006).





trivially, because the experience is a veridical perception of the green cube. On this much disjunctivists of Martin's and Fish's stripe (from now on I'll simply call them 'disjunctivists')³ and their opponents can agree. They can also agree that a hallucinatory experience can be indiscriminable from a veridical perception of a green cube.


The central theses of disjunctivism concern the status of specific indiscriminability properties of experiences, such as the property of being indiscriminable from a veridical perception of a green cube. For our purposes, the most important theses concern hallucinations. According to disjunctivism, my hallucination of the green cube is indiscriminable from a veridical perception of a green cube, even though there is no robust property the hallucination has in virtue of which it is so indiscriminable. This is Martin's line on *causally matching hallucinations* generally: hallucinations with the same proximate causal antecedents as veridical perceptions. (So let us retroactively suppose that my hallucination of a green cube causally matched a normal veridical perception of a green cube.) The most basic mental kind that hallucinations belong to, disjunctivists hold, is: being indiscriminable from a veridical perception. Veridical perceptions also belong to this kind, they think, but it is not their most basic mental kind. That kind is: being veridical perceptions; whereas the most basic mental kind of causally matching hallucinations is: being indiscriminable from a veridical perception. This is the most basic mental kind (or in Martin's terminology, the fundamental kind) of hallucinations in the sense that there is no further mental property hallucinations have, in virtue of which they are so indiscriminable.


Veridical perceptions are, trivially, indiscriminable from veridical perceptions; but according to disjunctivists further properties underlie their indiscriminability. For instance, my veridical perception of the green cube, they think, is indiscriminable from a veridical perception of the green cube in virtue of the perceptual relation that holds between the perceiver (me), on the one hand, and the cube and the properties of it that appear to me, on the other. These further properties, on their view, constitute the most basic phenomenal character of the veridical experience. This phenomenal character is thus robust in the sense that it underlies the indiscriminability.

So disjunctivists agree with the common-kind theorist that there is a common element between causally matching hallucinations and the veridical perceptions that they causally match;⁴ but they disagree about the depth and significance of the

³ This use of 'disjunctivism' may be sociologically misleading since many theorists doing business under that label say very little about the nature of hallucination, and do not even invoke the notion of indiscriminability. (Cf. McDowell (1982) and Putnam (1999). Campbell (2002) repudiates the label but holds that the natures of non-hallucinations and hallucinations fundamentally differ.) To uphold their views, however, these theorists need some conception of hallucination that is consistent with disjunctivism, and they could sign on to the epistemic conception if they wanted to. So the discussion here bears on possible developments of their views.

⁴ Martin grants the conclusion of the Causal Argument from Hallucination (see section 4), which he formulates as follows: "whatever kind of experience does occur in situations like *h* [hallucination], it is possible that such a kind of experience occurs when one is veridically perceiving" (Martin 2004: 38). Of course, "whatever kind..." is restricted to exclude the kind: hallucination.





common element. For disjunctivists, it goes no deeper than the indiscriminability property, and it does not constitute the fundamental kind to which both experiences belong.⁵

We can now draw out the central claim that disjunctivists defend about hallucinations. Hallucinations, according to them, are events or states⁶ whose phenomenal character consists exclusively in the having of indiscriminability properties. Let us formulate the thesis for a specific case and call it (H):

(H) If S hallucinates a sausage as such, then the hallucination is indiscriminable from a veridical perception of a sausage as such, and there is no further mental property that constitutes the phenomenal character of the hallucination.

I am going to let ‘S hallucinates (/veridically perceives) a sausage’ abbreviate ‘S hallucinates (/veridically perceives) a sausage as a sausage’ (and similarly for things other than sausages). The additional clause (‘as a sausage’) makes a difference because in principle one could veridically perceive a sausage that looked nothing like a sausage (for example, if it was disguised as a carrot), yet the experience could still be veridical: the experience might take a stand on only the colour and shape of the disguise (if these were the only properties ‘manifested’ to the subject, and these were properties the disguise actually had). This would be a case in which a sausage was veridically perceived, but wasn’t perceived to be a sausage.⁷ For the sake of brevity I’ll sometimes leave out the extra clause, but the result should always be taken as an abbreviation.

The indiscriminability property that figures in (H) is both epistemic and purely negative. If your experience has it, then a certain fact about your experience is not introspectively knowable by you—namely, the fact that it is not a veridical perception of a sausage. Since a hallucination is a kind of visual experience, and since nothing is a visual experience if it lacks phenomenal character, disjunctivists who accept the generalization of (H) are also committed to the view that an event or state’s having the relevant sort of indiscriminability property suffices for it to be a visual experience. Let us call this claim $I \rightarrow E$:


$I \rightarrow E$: If a mental state or event has the property of being indiscriminable from a veridical perception of an F as an F, then it is an experience.


The conditional $I \rightarrow E$ is a schema. Once a more specific indiscriminability property is filled in, such as the property of being indiscriminable from a veridical perception of a sausage as such, the result is a claim that having that property suffices for having

⁵ “...while the perceptual event is of a fundamental kind which could not occur when hallucinating, nonetheless this very same event is also of some other psychological kind or kinds which a causally matching hallucinatory event... belongs to” (Martin 2004: 60).

⁶ Martin speaks of events rather than states. My discussion is indifferent to whether disjunctivism is formulated in terms of events or states.

⁷ Some might deny that anything can be visually perceived as a sausage, on the grounds that only colours, shapes, and other ‘low-level’ properties can be visually perceived. Nothing turns on whether this denial is correct. The sausage hallucination example comes from Martin (2006), but the green cube hallucination would work just as well throughout.





a correspondingly specific experience, such as an experience as of a sausage. If $I \rightarrow E$ is true, then there are also strong links between the notions of experience, on the one hand, and veridical perception, on the other. In Siegel (2004) I criticize the claim that the notion of experience is as closely tied to the notion of veridicality as $I \rightarrow E$ suggests, but I will set that issue aside here.

Let us now consider more closely what the notion of indiscriminability is that figures in H and $I \rightarrow E$.

2 INDISCRIMINABILITY

So far our discussion has allowed that a particular, unrepeatable experience could be indiscriminable from another (perhaps repeatable) event. It's not entirely clear what this means. The relevant notion of indiscriminability can't be a statistical notion, since the event said to be indiscriminable is unrepeatable. Nor can the notion be reasonably understood by considering what would happen if the subject had two simultaneous experiences, compared them, and found that they were the same in the relevant respect (as one might be able to do with two physical objects). Perhaps one could think of how the subject would regard the pair of experiences, if she had them successively; but that introduces complications about memory that seem extraneous. Nevertheless, there is some intuitive sense in which certain pairs of experiences seem the same to the subject. The disjunctivist (and anyone else applying the notion of indiscriminability to experiences) thus faces the task of specifying what it is for one experience to be indiscriminable from another. And the disjunctivist alone faces the task of specifying what indiscriminability is, without bringing in robust phenomenal character.

As Martin thinks of it, indiscriminability is a notion defined in terms of judgement. "To discriminate two things" Martin writes, "is to judge them non-identical" (2004: 62). This suggests that when A and B are *indiscriminable* for a subject, the subject cannot tell them apart *in judgement*. This suggests the following notion of indiscriminability:


X is indiscriminable from Y by a subject S at time t iff S cannot know at time t by introspection alone that X is not Y.⁸


Here indiscriminability is a relation that holds between tokens. According to a closely related notion, it is a relation that holds between a token and a type:

X is indiscriminable from Fs by a subject S at time t iff S cannot know at time t by introspection alone that X is not an F.

Of the two relations, the latter is closer to the one endorsed by Martin. According to him, if a subject S hallucinates a green cube, and if her hallucination—call it h—is indiscriminable from a veridical perception of a green cube, then applying this notion

⁸ This formulation of indiscriminability comes from Williamson (1990).





of indiscriminability yields the result that S cannot know by introspection alone that h is not a veridical perception of a green cube.⁹

We are now in a position to see what role in (H) is played by the notion of being indiscriminable from a veridical perception of a sausage, when it is cashed out in terms of such knowability. If (H) is true, then in the case of hallucination, a fact about what can't be known introspectively about the hallucination is supposed to characterize how things are visually from the subject's perspective. For instance, if you are hallucinating a sausage, then by Martin's lights, the impossibility of knowing introspectively that you are not veridically perceiving a sausage is all there is to the phenomenal character of your experience.

3 COGNITIVELY UNSOPHISTICATED HALLUCINATORS

When (H) and its generalization are combined with the interpretation of indiscriminability discussed above, they seem to have no application to cognitively unsophisticated creatures such as dogs, even though it seems plain that dogs can hallucinate.¹⁰ Consider a dog that lacks the cognitive equipment needed to form judgements, including ones that count as knowledge. For such a dog, there will be no perceptual experience such that he can know that it is distinct from a veridical perception. So all of the dog's perceptual experiences will count as indiscriminable from every kind of veridical perception. If the fact that two experiences share a specific indiscriminability property suffices for them to seem the same to the subject, then by accepting this interpretation of $I \rightarrow E$, disjunctivists predict that all the dog's experiences will seem the same to the dog. That seems plainly wrong.


This consequence hits hardest in the case of hallucinations. In the case of dog perceptions, the disjunctivist thinks there are robust world-involving properties that constitute the experiences, and that the dog's perception of a sausage will differ phenomenally from his perception of a carrot because of the difference between the corresponding robust properties. Although these experiences also differ in their indiscriminability properties, the disjunctivist need not rely upon these properties to make it the case that the experiences differ from the dog's point of view. In contrast, in the case of hallucinations, there is supposed to be nothing to the experiences (by disjunctivists' lights) besides the having of indiscriminability properties.


Martin considers and responds to this worry. His response introduces the idea that the relevant notion of discriminability should be impersonal:

when we turn to the experiences of sentient but unselfconscious creatures, to the extent that we do have a positive grip on the kinds of experience that they can have, and which can differ

⁹ It is possible to define other notions of indiscriminability. Some of these are discussed and criticized in Siegel (2004) and Farkas (2006). The most extended discussion of the notion is in Williamson (1990).

¹⁰ For brief discussion of empirical evidence of hallucination on the part of non-human animals, see Fish, this volume.





one from another, we also have a grip on how such experience would be discriminable through reflection or not . . . a dog might fail to discriminate one experience from another, making no judgment about them as identical or distinct at all, [but] that is not to say that we cannot judge, in ascribing them such experience, that there is an event which would or would not be judgably different from another experience. (2004: 54)



[W]e are interested in the impersonal notion of inability or incapability here. That is we are interested in the claim that John [or the dog] is in a situation for which it is impossible *simpliciter* and not just impossible for John [or the dog] to tell apart through introspective reflection from a veridical perception [of a sausage]. (2006: 381)


In considering this worry, Martin here is discussing John, an inattentive, hasty subject who treats samples of scarlet and vermillion indifferently. Intuitively, his experiences of each should count as distinct; yet, as Martin points out, it seems appropriate to say that inattentive John can't discriminate scarlet from vermillion. Martin takes the moral to be that the disjunctivist should adopt an 'impersonal' notion of indiscriminability, where this impersonal notion applies equally to the dog, to inattentive John, and to perfectly attentive cognitively sophisticated subjects.

Let's say that an indiscriminability property I of an experience is *impersonal* just in case there are no cognitive abilities that the subject must actually have, beyond any that may be needed to have conscious experiences, in order for her to have an experience with I. Martin proposes that we should understand these properties by comparison with the ordinary notion of invisibility. We can make sense of the idea that a bit of mending on a jacket is invisible, by considering whether someone ideally placed to see things—an ideal visual perceiver in ideal circumstances—could see the mending or not. If not, then the mending is invisible (Martin 2006: 381). By analogy, he suggests, we can make sense of the idea that the dog's experience is objectively indiscriminable from a veridical perception of a sausage while being objectively discriminable from a veridical perception of carrots, by considering whether by some means not employable by the dog himself, someone could know that the dog's experience is not a veridical perception of a sausage. If she could not know this, and if she could know that it's not a veridical perception of a carrot, then the experience is (non-trivially) indiscriminable from a veridical perception of a sausage.

The crucial question for the disjunctivist is what it could be for a dog's experience to have such an impersonal indiscriminability property. The disjunctivist must deny that an experience's having it involves its having underlying robust properties, in virtue of which the experience is indiscriminable from a sausage-perception. If the experience had such robust properties, that would violate the thesis that the hallucinator's experience consists merely in its having a negative epistemic property.

The fact that on pain of violating (H), impersonal indiscriminability properties cannot involve underlying robust properties suggests that they should be understood as some sort of counterfactual about the kind of knowledge that figures in the epistemic notion of indiscriminability endorsed by Martin and Fish. Here are three





counterfactuals that might be proposed as making explicit what it consists in, for the case of the dog:

- (1) If the dog could ideally reflect on his situation, he would not be able to know by reflection that he was not veridically perceiving a sausage.
- (2) If I could ideally reflect on the dog's situation, I would not be able to know by reflection that he was not veridically perceiving a sausage.
- (3) If an ideal introspector were in the dog's situation, such an introspector would not be able to know by reflection that she was not veridically perceiving a sausage.

If any of these explications is a way to understand what impersonal indiscriminability properties are, then given the claim that $I \rightarrow E$, it must support this conclusion:


C. The dog has an experience as of a sausage.

These counterfactuals all mention the dog's 'situation'. What is it to be in the [hallucinating] dog's situation? To be in this situation can't be to have a robust phenomenal character, by disjunctivists' lights. It cannot be a subpersonal physical or functional state that was identified with the experience, since that too would introduce an underlying common phenomenal character. It cannot be a subpersonal physical or functional state that normally is causally antecedent to the experience, because then the account of what the hallucinating dog's experience consists in will undergeneralize if there are other causal routes to having the experience, and overgeneralize in cases where the subpersonal state does not cause any experience.¹¹ And it can't be simply to have the property of being impersonally indiscriminable from a veridical perception of a sausage—or else 3, for example, would mean the same as the trivially true 3*:

(3*) If an ideal introspector couldn't know by reflection that she wasn't veridically perceiving a sausage, then she couldn't know by reflection that she wasn't veridically perceiving a sausage.

A similar objection applies to counterfactuals 1 and 2. 'The dog's situation' in each case cannot pick out a robust phenomenal character or a subpersonal state that is or is causally antecedent to the experience. An appeal to an indiscriminability property moves us in a circle, since that is the property that the counterfactuals are attempts to explain. Finally, it seems ill-advised for the disjunctivist to get rid of mention of a situation altogether in 1, so that the antecedent read 'If the dog could ideally

¹¹ To elaborate the counterfactual is supposed to provide an account of what it is for the hallucinating dog to have an experience as of a sausage. It accounts for this only if the cases where the counterfactual holds are all and only cases where the dog is in the situation mentioned in the antecedent. If that situation is characterized in terms of the normal causal antecedents of the experience, then the account will be adequate only if those causal antecedents occur when and only when the dog has the relevant kind of experience. There will thus be a problem of undergeneralization if there are other causal routes to the same effect (namely, the experience), and a problem of overgeneralization if there are cases where the normal causal antecedents don't produce the effect (say because something intervenes, or because of a chance occurrence). Since effects typically can have more than one cause, and since causal antecedents are not guaranteed to have their normal effects, both problems arise.



reflect . . .'.¹² Substituting this for I in $I \rightarrow E$ would have the consequence that the dog was always having an experience, assuming that the dog could ideally reflect.

Disjunctivists must find a way to talk about the dog's situation—that is, his being in a state or event that $I \rightarrow E$ says shall count as an experience. Here all three counterfactual proposals face the same problem. If this situation is not a specifically experiential situation, then there is no reason to think that the counterfactual in which it appears explicates what it is to have an experience. Once the situation in the antecedent is an experiential situation, however, the account becomes trivial, since the disjunctivist's view is precisely that there is nothing else to the experience besides the having of the indiscriminability property itself. These doubts are general, and seem to apply to any attempt to explicate what impersonal indiscriminability properties could be. Given the assumption that some creatures who can have experiences are nonetheless not capable of having the kind of knowledge invoked by the impersonal indiscriminability properties, the prospects for developing the modalized epistemic notion of impersonal indiscriminability consistently with disjunctivism thus seem dim.¹³


Disjunctivists might try responding that dogs do not after all lack the kind of knowledge involved in the indiscriminability properties. To follow this strategy, the disjunctivist would first have to identify a general cognitive ability that dogs as well as cognitive sophisticates actually have, so that the notion of indiscriminability could be formulated using that notion. A starting point might be Sosa's notion of animal knowledge in his (1991). So long as there could be a perceiver who lacked the ability to have the suitably undemanding kind of knowledge, however, the dog problem will return. To follow this strategy, disjunctivists would thus have to argue that *any* creature capable of having hallucinatory experiences is also capable of having this kind of knowledge.

This last key claim seems difficult to establish. Suppose that the animal knowledge in question is a reliable connection of some sort between putative perceptual states and the environment (as suggested by Sosa's notion of animal knowledge). Then the key claim says that any creature capable of hallucinating is also capable of being in states that are reliably connected to the environment. It seems easy to imagine a creature with separate bits of mental machinery dedicated to producing conscious experience on the one hand and to receiving inputs from the environment on the other, where the two bits could operate independently. Suppose that in such a creature, the machinery dedicated to receiving inputs from the environment was so defective that no correlation ever existed between states of the environment and states of the machinery. It seems coherent to suppose that the other bit of machinery might produce conscious experiences perfectly well, including those phenomenally like the ones had by normal animals when they perceive. The result would be a creature capable of conscious hallucinatory experiences, yet incapable of animal knowledge, on

¹² This was suggested by a referee.

¹³ Some of these difficulties are acknowledged by Martin (2006: 383 and fn. 44) where he disowns some of these ways of developing the notion of impersonal indiscriminability. Martin retains a notion of impersonal indiscriminability. His preferred version of impersonal indiscriminability seems to be implicit in his view of introspection, which is discussed further in section 5.





this construal of it. More generally, many ways of developing the notion of animal knowledge seem to leave it an open possibility that the features enabling animals to have that kind of knowledge need not be the very features that enable them to hallucinate. The disjunctivist would have to show that despite appearances, these two capacities (for animal knowledge on the one hand, and for hallucination on the other) could not each operate separately.

There is a notion of animal knowledge, however, on which these two capacities could not operate separately. Suppose that the kind of knowledge that animals could have is a sort that one is automatically capable of having as a result of having conscious experiences of the right sort, where the right sort includes hallucinations. Then it would not be possible for a creature to hallucinate without the requisite kind of knowability facts obtaining. In section 5 (under the heading 'brute fact proposal') I criticize a general version of this proposal, one that is indifferent to whether the kind of knowledge in question is available to cognitively unsophisticated creatures or not.

Sympathizers with disjunctivism might try to dismiss the entire worry about cognitively unsophisticated creatures, by responding that the view as a whole should not stand or fall with its treatment of hallucinating animals. This dismissal is inconsistent with a commitment of Martin's: that being indiscriminable from veridical perception "is the most inclusive conception we have of what sensory experience is" (Martin 2004: 56). If disjunctivists propounding (H) take on this commitment, then they are stymied by the dog objection. If they deny this commitment, opting out of providing a unified account of perceptual experience as such, then the dog objection will not arise, but other objections will still arise. Some of these objections target Fish's version of epistemic conception, which will be discussed next. Others target both Fish and Martin's versions of the epistemic conception, and they are presented in section 5.


4 FISH'S EFFECTS-BASED VERSION OF THE EPISTEMIC CONCEPTION

So far we've been discussing Martin's knowability-based notion of indiscriminability. In contrast to Martin, Fish develops a notion of indiscriminability of hallucinations from veridical perception that fixes on their effects. In this section I argue against Fish's version of the epistemic conception.

Fish takes it as analytic that hallucinations have certain effects:

[W]e can say that the indistinguishability of the two states [hallucination and veridical perception] is manifested by the hallucination having sufficiently similar effects to those that the veridical perception would have had. But the demand for a more substantial intrinsic characterization of the hallucinatory mental state is misguided. For a similar reason, it would also be misguided to demand any further explanation of why hallucinations, considered as a mental kind, have similar effects to veridical perceptions. As a mental state only qualifies as a hallucination inasmuch as it has the same kinds of effects as a certain kind of veridical perception, there is therefore something wrong with asking why *hallucinations* have these kinds





of effects—it would be akin to asking what bachelors have in common in virtue of which none of them are married. (Fish, this volume: section 7)

Which effects are sufficiently similar? According to Fish, the relevant effects in creatures like us are beliefs, including higher-order beliefs, such as might be expressed by saying ‘I am veridically perceiving a sausage’:

Say I veridically perceive a pink ice cube: what kind of effects would such a mental state have? . . . [T]here would be a vast array of effects such a state might have—it might make me want a cold drink, or remember the last occasion on which I saw such an ice cube, or think of Wilfrid Sellars, and so on. But there would also be more standard effects than these. In particular, where subjects such as ourselves are concerned, someone who veridically perceives a pink ice cube would, unless something very curious was going on, acquire a range of fairly predictable beliefs. We would expect these beliefs to include, for example, the belief that a pink ice cube exists and the belief that she sees the pink ice cube. (Fish, this volume: section 5)


So according to Fish, the property of being indiscriminable from a veridical perception of a sausage is constituted by the property of having sufficiently similar effects as a veridical perception of a sausage.


On the face of it, Fish’s effects-based notion of indiscriminability makes it easier for a proponent of (H) to solve the dog problem. According to Fish, in the case of creatures who are not sophisticated enough to have higher-order beliefs, hallucinations will not include among their effects the belief on the part of the subject that she is veridically perceiving. Instead, the effects that constitute their hallucinations will be limited to behavioural effects:

[W]hen it comes to a creature—a cat, for example—which lacks the conceptual sophistication required to entertain higher-order beliefs, a veridical perception of, say, a butterfly will still have certain kinds of effects. Of course, given that we are assuming that the cat is lacking in higher conceptual capacities, these effects will not include cognitive effects such as higher-order beliefs. Instead, they will be primarily behavioural, possibly including trying to strike the butterfly with its paw, and so on. But then our evidence for animal hallucinations is essentially a set of behavioural observations too. If we were to find a cat “striking at imaginary objects in the air” (Slade and Bentall 1990: 12)—in other words, behaving as though it perceived a butterfly when there was not one there—then we might have good reason to say that, in such a case, the animal is not perceiving a butterfly but is nevertheless in a mental state which has the same effects as such a perception. As the two states have sufficiently similar effects . . . we can therefore allow that the two states are indistinguishable for the cat, and that the animal is having a hallucination of a butterfly . . . (Fish, this volume: section 6)

With the effects-based notion of indiscriminability in place, worries about how the cat could non-trivially fail to know that she is not veridically perceiving a butterfly disappear. If effects-based notion is correct, then it avoids the dog problem.

There are problems facing this notion, however. The theory does not ensure that hallucinations have any felt reality from the point of view of the hallucinator—not even a deflationary sort of felt reality that Fish himself invokes. Fish considers his theory to be an irrealist account of hallucination that “denies that hallucinations have phenomenal character whilst agreeing wholeheartedly that they certainly seem to”.





Their seeming to have phenomenal character, he says, consists in part in the subject's believing that they are veridically perceiving.¹⁴

[I]f a mental state which was *not* a veridical perception of a pink ice cube somehow led the subject to form beliefs of this kind, then these beliefs could be appealed to in explaining why the subject takes herself to be having a veridical perception of a pink ice cube even though she is not. So long as I believe that I am seeing something, I will take myself to be enjoying a certain kind of veridical perception If a mental state which was not a veridical perception were somehow to come to have effects which included such higher-order beliefs, then this would look to be sufficient to explain why that mental state was mistaken for a veridical perception of that kind. (Fish, this volume: section 5)


The appeal to actual higher-order beliefs cannot account for what makes all hallucinators take themselves to be enjoying a veridical perception when they are hallucinating, because some of these hallucinators will not actually form the higher-order beliefs. Cognitively unsophisticated subjects may be one example. We can find other examples if we unpack the effects-based theory a bit more.

Fish's position seems best construed as a disjunctivist version of analytical functionalism. On this view, hallucinations are defined using an a priori theory such as folk psychology (or some analogue for animals). Hallucinatory experiences are whatever states have the same effects as veridical perception, with the stipulation that they are not themselves veridical perceptions.¹⁵ Veridical perceptual experiences, in contrast, would have a functional role that included these effects but also included worldly causes as well.

Suppose, with Fish, that there is a set of effects that veridical perceptions of a butterfly standardly have—call this set E. These effects are not effects of the veridical perception considered in isolation, but rather they are effects of the veridical perception together with other mental states. As Geach (1957: 8) pointed out in criticizing logical behaviourism, mental states do not each have their own individual fund of behavioural effects. A cat veridically perceiving a butterfly standardly paws at it partly because she wants to make contact with it. If the cat is lethargic from illness and so lacks that desire, even if she veridically perceives the butterfly, she won't paw at it or even be disposed to paw at it. In identifying a set E of standard effects of veridically perceiving a butterfly, we are identifying effects of the veridical perception together with other mental states. Fish does not present his view this way, but Geach's point seems correct, so let us proceed by considering a Fish-inspired theory of hallucinations according to which they are states that have the certain effects, given the other mental states of the creature. What necessary conditions does (H) impose on an event or state's being a hallucination, given an effects-based notion of indiscriminability?

¹⁴ Fish offers separate treatment of 'resisted hallucinations', which are hallucinations in which the subject is 'aware of the non-veridicality' of the experience. His treatment of these hallucinations appeals to different higher-order beliefs. See Fish (this volume: section 10).

¹⁵ According to Fish's definition of hallucination, "a hallucination is a mental state which, whilst not being a veridical perception of a particular kind, has effects which are sufficiently similar to those a veridical perception of that kind would have had" (section 5). This definition includes illusion too—though to Fish this is not an unwelcome result (Fish MS).



Where E is still the set of effects standardly brought about by veridically perceiving a butterfly, the condition is this:


(*) An event or state e is a hallucination of a butterfly only if: if the subject has e and certain other mental states, then e will ground a disposition to produce E.¹⁶


By itself, the necessary condition on hallucinating butterflies laid down by (*) does not reflect any specifically disjunctivist treatment of hallucination. For all (*) says, e could ground a disposition to produce E in virtue of its phenomenal character, and that phenomenal character could be shared with a veridical perception of a butterfly. Condition (*), however, is compatible with the conditions that (a) butterfly-hallucinations lack robust phenomenal character altogether, and in particular (b) butterfly-hallucinations lack robust phenomenal character that veridical perceptions of butterflies have. In the conjunction of (*), (a) and (b), we can see the structural similarity between this treatment of hallucination and Martin's treatment of them. Both hold that an event or state is a hallucination only if it is indiscriminable from a veridical perception. And both hold that there is no common phenomenal character underlying the indiscriminability.

Let us now return to the objection that Fish's theory of hallucination does not ensure that they have any felt reality. According to Fish, having a higher-order belief that you are veridically perceiving a butterfly is supposed to account for why you take yourself to be having an experience with phenomenal character, even though (by his lights) you are not. But even some hallucinators will not form them—even if they are not 'resisting' their hallucinations (that is, even if they do not believe that they are hallucinating—Fish offers a separate treatment of 'resisted' hallucinations). We can leave aside the question whether failing to form these higher-order beliefs would be irrational—rationality is not a prerequisite of hallucinating. For subjects who fail to form these higher-order beliefs, the effects-based theory lacks a crucial resource for accounting for the felt reality of the hallucination. Neither the hallucination's actual effects nor any phenomenal basis of them is available to Fish to account for the felt reality of the hallucination. It is hard to see what else could play this role.

In response to this problem, the effects-based theory could be reformulated to make it analytic that hallucinations actually have the effects in set E, rather than making it analytic that they have those effects given appropriate background conditions. But this refinement would have implausible consequences. It should be possible to have a hallucination that is indiscriminable from a veridical perception of a butterfly, even if the standard effects E don't actually come about. And given a veridical perception of a butterfly that itself lacks the standard effects E, it should be possible for there to be a hallucination that is indiscriminable from it. A subject might suddenly expire just after veridically perceiving or hallucinating a butterfly, in which case the veridical perception or hallucination would not have E. So to ensure that such hallucinations are possible, the effects-based theory should not formulate the definition of hallucination in terms of actual effects.

¹⁶ The Fish-inspired theory would arguably increase in plausibility by adding "if the subject is rational . . ." to the necessary condition in (*). But adding this would not change the dialectic.





5 GROUNDING EPISTEMIC FACTS

My next objection to (H) focuses on what is introspectively knowable in cases of hallucination. First, I will argue that in addition to countenancing negative epistemic facts about the indiscriminability from other states, proponents of (H) must countenance certain positive epistemic facts about the discriminability of hallucinations from other states. Second, I will suggest, by considering and rejecting four proposals, that the disjunctivist cannot give a good account of the epistemic ground of these positive facts about discriminability.

Earlier, we mentioned two assumptions about hallucinations that the disjunctivist aims to respect. First, there can be ‘matching’ hallucinations; second, hallucinations can vary in their phenomenal character (the ‘difference datum’). According to the difference datum, the character of hallucination of a sausage differs from that of a hallucination of a pyramid or of a butterfly. So for any hallucination, there will be a set of possible experiences with contrasting characters. Let a butterfly-experience be an experience with the phenomenal character of veridically perceiving a butterfly as such, or with the phenomenal character of hallucinating a butterfly as such. (Disjunctivists will cash out the latter kind of phenomenal character in terms of indiscriminability, whereas common-kind theorists won’t see the need for the disjunction in the definition).


Given the contrast in character between a sausage-hallucination and the butterfly-experiences, it seems plausible to suppose that when you hallucinate a sausage, you can know introspectively that you’re not veridically perceiving a butterfly as such. Generalizing from this, it is equally plausible to suppose that if you have a butterfly-experience, you can know introspectively that you’re not veridically perceiving a sausage as such. For such facts to be knowable, it need not be the case that every detail of the character of the hallucination is introspectively accessible. Introspection just has to be able to be sensitive to the phenomenal contrast between sausage-hallucinations and q-experiences. We can abbreviate the last claim that I said was plausible to accept as claim (i) below, where $K\neg VP$ (sausage) is an abbreviation for the claim that it is introspectively knowable that you are not veridically perceiving a sausage as such. (So note that ‘K’ has a modal element built into it.)

(i) S has a butterfly-experience $\rightarrow K\neg VP$ (sausage).

The consequent of (i) is a fact (some may consider it a merely putative fact) about what is introspectively knowable. Let’s call this the Positive Epistemic Fact. (It is capitalized because it is specific to sausages.) My objection to (H) concerns what epistemic ground disjunctivists can offer for the Positive Epistemic Fact and others analogous to it.

Someone might challenge (i), on the grounds that it overestimates introspection’s actual epistemic capacities, inappropriately crediting it with being a source of knowledge about what one is not veridically perceiving. Could a disjunctivist avoid discussing the objection by denying (i) on these grounds?





The position on introspection is not open to the proponent of (H) who construes indiscriminability as unknowability. According to this construal of (H), when you hallucinate a sausage, you can't know by introspection that you're not veridically perceiving a sausage, and this introspective unknowability is all there is to the phenomenal character of your hallucination. The difference in character between your sausage-hallucination and, say, a hallucination of a butterfly is supposed to be reflected in the fact that each hallucination consists in a different fact about introspective unknowability. The fact that you can't know by introspection that you're not veridically perceiving a sausage is thus supposed to suffice to make it the case that you are not having a butterfly-experience: for example, you are not hallucinating a butterfly, and you are not having an experience with the character of veridically perceiving a butterfly. We can abbreviate this claim as claim (ii):



(ii) $\neg K\neg VP(\text{sausage}) \rightarrow S$ is not having a butterfly-experience.

The proponent of thesis (H) is thus committed to (ii), given that the unknowability fact in its antecedent is supposed to respect the datum that hallucinations can differ in their character from other hallucinations (as well as from other non-hallucinations). But since claim (ii) is the contrapositive of claim (i), the proponent of thesis (H) is committed to claim (i) as well. So there is no escaping a discussion of what grounds the Positive Epistemic Fact, if one accepts thesis (H) with the unknowability construal of indiscriminability.

What grounds the fact, illustrated by the consequent of (i), about what is introspectively knowable in the case of hallucination? It is of course not open to the proponent of (H) to appeal to the robust phenomenal character of the hallucination here. So let us consider the other options that are open to them and see whether they are plausible.

Is the epistemic fact grounded by the indiscriminability property? According to the first proposal, the knowability in the Positive Epistemic Fact can be derived from the unknowability that constitutes the sausage-hallucination—that is, the unknowability by introspection that one is not veridically perceiving a sausage. Call this unknowability the Negative Epistemic Fact. This proposal may appeal to a proponent of (H) because such unknowability is the most basic phenomenal character of hallucination. So the proposal preserves a sense in which the experience itself grounds the knowability.

However, it is not possible to derive the positive epistemic fact from the negative epistemic fact in this way. Suppose that I cannot know by introspection that I am not veridically perceiving a sausage (as such). Then it is an open possibility for me, for all introspection tells me, that I am veridically perceiving a sausage (as such). But this open epistemic possibility for me may be just one among many. For all introspection says, I may alternatively be veridically perceiving a pyramid as such, or a bridge as such, or carrots as such. All that follows from the Negative Epistemic Fact is that it is an open possibility that I am veridically perceiving a sausage as such, but it is compatible with the openness of that possibility that I am instead veridically perceiving (say) a pyramid as such.




In contrast, these other possibilities (that I am veridically perceiving a pyramid as such, and so on) are ruled out if the Positive Epistemic Fact obtains. According to the Positive Epistemic Fact, it is knowable to me introspectively that I am not veridically perceiving a pyramid. So it is not possible to derive the Positive Epistemic Fact from the Negative Epistemic Fact. To make that transition, something else is needed to rule out these other possibilities. The first proposal thus does not pin-point what makes the proposition that I am not veridically perceiving a non-sausage introspectively knowable, in the case where I am hallucinating a sausage (as such).

Is the epistemic fact grounded by introspection construed as a procedure? Suppose that there is some feature of introspection in virtue of which the Positive Epistemic Fact (and others like it) obtains. Versions of this proposal will vary depending on what the feature of introspection is taken to be. For example, according to a reliabilist version, the proposition that I'm not veridically perceiving a pyramid is introspectively knowable to me when I'm hallucinating a sausage, in virtue of the fact that I could employ a type of introspective process that reliably generates true beliefs that I'm not veridically perceiving a pyramid. Reliabilism is just one example. Let us focus more generally on the idea that there are dedicated mechanisms or processes of introspection whose employment is an optional addition to experience itself, as opposed to being constitutively linked to experience. Call this the procedural model of introspection.

In our discussion so far we have been focusing on cases in which the Negative and Positive facts obtain, and in which by hypothesis I am hallucinating a sausage. But the proponent of (H) holds that these facts suffice for me to have an experience. They accept the conditional: $I \rightarrow E$. Once we fix on the procedural model of introspection, it is coherent to suppose that a procedure or mechanism of introspection could operate in a way that resulted in the subject's knowing that she was not veridically perceiving a pyramid, even if she was not having any experience at all. For instance, suppose I have an introspective process or mechanism that reliably generates true beliefs to the effect that I am not veridically perceiving a pyramid. Suppose it operates and generates such a belief. By reliabilist lights, I will know introspectively that I am not veridically perceiving a pyramid. The Positive Epistemic Fact will thus obtain. Suppose further that the mechanism cannot in this situation generate the belief that I'm not veridically perceiving a sausage. Then I will not be able to know introspectively that I am not veridically perceiving a sausage. So the Negative Epistemic Fact will thus obtain. But it is compatible with the case as it has been described that I might lack experience altogether. If the introspective procedure is not constitutively linked to my having an experience, then it could operate in the absence of any experience. So the procedural model itself undermines $I \rightarrow E$, because of the lack of a constitutive connection between the operation of introspection and the experience itself. So if the proponent of (H) appealed to some feature of introspection, procedurally construed, to ground the Positive Epistemic Fact, she would run afoul of another disjunctivist commitment, namely $I \rightarrow E$.

Is the epistemic fact a brute fact? I'll call the third proposal for what grounds the Positive Epistemic Fact the brute-fact proposal, because it gives the Positive Epistemic Fact the status of brute fact. Giving the Positive Epistemic Fact this status may seem attractive for disjunctivists. Prima facie, the natural place to look for the aetiology of



such knowledge or for its ground would be the specific character of the experience. The epistemic conception of hallucination replaces talk of robust phenomenal character of hallucinations with facts like the Negative Epistemic Fact. The whole point of the epistemic approach to hallucination is that epistemic facts suffice to reflect that specific character. In effect, they are substitutes in the theory of hallucination for robust phenomenal character. But as we have seen, such facts cannot ground the Positive Epistemic Fact.



According to the brute-fact proposal, when I form the belief (while hallucinating a sausage) that I am not veridically perceiving a non-sausage, the status of my belief as knowledge, and the introspective knowability of what I believe, are both just brute epistemic facts, with no further explanation. In cases where I actually know introspectively that I am not veridically perceiving a non-sausage, the proposal says that there is no feature of introspection that makes my belief count as knowledge. And in cases where I have no non-introspective basis for this belief, the proposal says there is nothing at all that makes the belief count as knowledge. Its status as knowledge will be a brute fact.

The brute-fact proposal puts pressure on the idea that the beliefs in question can have the status of knowledge at all. If there is no feature of introspection that gives the belief its status as knowledge, then introspection contributes nothing to giving the belief its putative epistemic status as knowledge. In cases where there is no other basis for the belief, there will be nothing else that gives it this status either. So if the belief had the status of knowledge, it would be knowledge without any grounding or basis, evidential or otherwise. At that point is hard to see what makes the belief deserve its status as knowledge, as opposed to being merely true. Yet it seems plain that we can know introspectively, while hallucinating a sausage, that we are not veridically perceiving a pyramid—and moreover, proponents of (H) are committed to this. The brute-fact proposal is not clearly compatible with this.

Even if the brute-fact proposal could allow that the belief has the status of knowledge, it closes off the possibility of explaining a fact for which we would like an explanation anyway. Independently of assessing disjunctivism, we would like to know how it is possible for us to know what we do from a first-person perspective about our experiences. According to the brute-fact proposal, we will be barred from learning anything about this, because there is nothing to learn.

We can distinguish between an extreme and a less extreme version of the brute-fact proposal. The extreme version embodies a general view about introspection, holding that there is never any such thing as introspective evidence for knowledge claims, and no other feature of introspection that can make it the case that a belief formed on the basis of introspection counts as knowledge (when it does). A less extreme version holds this only for the relevant introspectively anchored knowledge claims in the case of hallucination (for example, the knowledge claim that can be derived from the Positive Epistemic Fact).

The extreme claim is pleasingly uniform, but not all disjunctivists can accept it. Suppose that introspecting a veridical experience of seeing a sausage as such gives us access to some aspect of its nature. If the extreme claim is true, such access will not have the status of knowledge. But some disjunctivists, including Martin, motivate



disjunctivism on the grounds that its treatment of veridical experience best accords with the way introspection presents those experiences as being.¹⁷ They thus seem to treat introspection as a source of knowledge, or at least justified belief, about the nature of experience. This claim seems weakened if it turns out that there are simply no answers to the question of what introspection is such that it provides such knowledge, or what in the nature of experience enables us know about it introspectively. So disjunctivists opting for the brute-fact proposal will have to adopt a version of it according to which introspection sometimes contributes to making a belief count as knowledge, and sometimes does not.


Are the epistemic facts grounded in the experiential perspective? According to Martin, the fact (as he sees it) that I cannot know by introspection that I'm not veridically perceiving a sausage constitutes the phenomenal character of my hallucination. If so, whatever introspection can tell me about my situation, it will not include that I am not veridically perceiving a sausage. The Negative Epistemic Fact that constitutes my sausage-hallucination thus gives me a 'perspective' on my situation. (Let us set aside the tricky question of what exactly 'my situation' is supposed to be—this point arose in our discussion of the dog objection.) Martin anticipates that someone might challenge (H) by asking: why think that any perspective on my situation delivered by introspection is accurate? Why couldn't the deliverances of introspection be inaccurate about the specific character of my experience—or even be wrong in supposing that I am having any experience at all? (Compare the putative case of zombies who falsely self-ascribe experiences.)

Martin responds to this challenge by holding that introspection's deliverances determine the existence and specific character of experiences, including hallucinations. More specifically, they constitute experiences—so 'deliverances' may be misleading, to the extent that the word suggests deliverances of verdicts about experiences. Martin's proposal links the perspective of introspection constitutively to experience itself:

[T]he perspective we have on our own phenomenal consciousness cannot be grounded in some specific mode or source of knowing about something independent of that perspective. If it is true of someone that it seems to them as if things seem a certain way, as if they are having a certain sense experience, then they are thereby having that experience. Our reflective standpoint on our own experience cannot stand outside of it. (Martin 2006: 397)

If the suggestion about the link between introspection and hallucination applies to the entirety of introspection's perspective, so that it is the entire perspective that constitutively determines the facts about one's experience, then there does not seem to be anything special about the Negative Epistemic Fact as opposed to the Positive Epistemic Fact. When the sausage-hallucinator can know by introspection that she is not veridically perceiving a non-sausage, it seems, this deliverance of introspection should also be constitutively linked to the hallucination. This in itself suggests a revision of

¹⁷ For example, Martin writes: "the disjunctivist advocates Naive Realism [the disjunctivist's view about the nature of veridical experience] because they think that this position best articulates how sensory experience seems to us to be just through reflection" (2006: 354).



thesis (H), according to which both Epistemic Facts constitute the sausage-hallucination, not just the Negative Epistemic Fact.


Suppose the disjunctivist accepts this revision of thesis (H). Can this position on the status of introspection help account for what grounds the Positive Epistemic Fact? It does not seem to help. We cannot look to the experience as the ground, since the experience is identified with introspective perspective. Alternatively, a proponent of identifying experience with the introspective perspective might try to dismiss the question as senseless, on the grounds that experience (unlike belief) has no ground, hence no question of the ground of introspective perspective can arise. But this response will not work either, since the proponent of (H) who relies on the knowledge-based notion of indiscriminability—and Martin is just such a proponent—is committed to cashing out the introspective perspective in terms of knowability claims. Such claims have grounds, whatever identity claims they may figure in. So the approach that identifies introspection's perspective with hallucination does not seem to do any better at accounting for the Positive Epistemic Fact.

Martin addresses something like the objection just made in a footnote, where he considers a comparison of his proposal about introspection with theories that posit constitutive relations between self-ascriptions of thoughts and the thoughts self-ascribed. He notes that some philosophers object to these theories that such self-ascriptions could have no rational ground. But he says no analogous objection applies to “the constitutive connection between the subject’s perspective on his or her own mind, how it seems to be, and how his or her mind then is” because the perspective “need not be identified with judgments he or she actually makes” (Martin 2006: n. 42). But even if the perspective is identified with judgements that the subject could make if she were rational, the question of what grounds them will still arise. And given the notion of indiscriminability Martin invokes, it is not in the cards to identify the perspective with something other than judgements. So the ground of the Positive Epistemic Fact still seems unaccounted for.

Given the difficulties with the four proposals for grounding the positive epistemic facts, I conclude that the chances of developing a plausible account of this grounding consistently with disjunctivism are low. When combined with the difficulties generated by the case of cognitively unsophisticated hallucinators, this result makes the prospects seem dim for the epistemic conception of hallucination.

REFERENCES

- Campbell, J. (2002) *Reference and Consciousness*, Oxford: Oxford University Press.
Farkas, K. (2006) ‘Indiscriminability and Sameness of Appearance’, *Proceedings of the Aristotelian Society*, 106: 205–25.
Fish, W. (this volume) ‘Disjunctivism, Indistinguishability and the Nature of Hallucination’.
——— (MS) *Perception, Hallucination and Illusion*.
Geach, P. T. (1957) *Mental Acts*, London: Routledge and Kegan Paul.



- McDowell, J. (1982) 'Criteria, Defeasibility, and Knowledge', *Proceedings of the British Academy*, 68: 455–79.
- Martin, M. G. F. (2004) 'The Limits of Self-Awareness', *Philosophical Studies*, 120: 37–89.
- (2006) 'On Being Alienated', in T. S. Gendler and J. Hawthorne (eds.), *Perceptual Experience*, Oxford: Oxford University Press.
- Putnam, H. (1999) *The Threefold Cord*, New York: Columbia University Press.
- Siegel, S. (2004) 'Indiscriminability and the Phenomenal', *Philosophical Studies*, 120: 91–112.
- Sosa, E. (1991) 'Intellectual Virtue in Perspective', in his *Knowledge in Perspective: Selected Essays in Epistemology*, Cambridge: Cambridge University Press.
- Williamson, T. (1990) *Identity and Discrimination*, Oxford: Blackwell.