

Hallucinations, Dreams and Phenomenal Character

‘After all, we have no guarantee other than the [hallucinating] patient’s reports, which all seem subject to caution’ – Sartre, *The Imaginary*

Arguments from Hallucination are common philosophical tools which attempt to establish that ordinary, non-deceptive¹ perceptions have some feature in common with hallucinations and other experiences which can occur without the world containing any of the objects or features which they purport to reveal to their subject. Many traditional versions of such arguments attempt to show that hallucinations and non-deceptive experiences have the same objects: that is that the object of ordinary perception is just as much a mental entity, an ‘image’, as the object of hallucination. Such arguments have received sustained criticism on the grounds that we cannot infer from two things seeming the same, even from their being indiscriminable, that they have the same nature. However, even if that criticism is correct, we may think that if hallucinations and non-deceptive experiences seem the same, that is have the same phenomenal character, then the way they both seem, their phenomenal character, must be consistent with the possibility that all our experiences are hallucinations. This would have an impact on the theory of consciousness and also the epistemology of perceptual judgements: the former would be prevented from allowing that the external² objects³ of perception are even partial determinants of phenomenal character and the latter would be prevented from allowing that perceptual knowledge owes its epistemic status entirely to consciously accessible features of our experience. For which reasons, some philosophers (Austin, Hyman, Martin, Fish) have been inclined to deny that hallucinations and non-deceptive perceptions do in fact seem the same, if that is taken to mean that they have the same phenomenal character. However, these philosophers do allow that there is some sense – which they take pains to make precise – in which hallucinations can be indiscriminable from non-deceptive appearances.

In what follows I examine the argument from hallucination to the conclusion that non-deceptive experiences have a phenomenal character which is independent of the perceiver’s environment. I shall argue that such arguments fail because they are epistemically circular, for they require a

¹ I prefer ‘non-deceptive’ to ‘veridical’ because veridicality requires some, admittedly very weak, notion of correctness conditions for the experience itself whereas being non-deceptive can be cashed out entirely in terms of the truth of judgements the subject makes on the basis of it alone (in suitable idealized conditions).

² It is surprisingly difficult to define ‘external’. Many contemporary philosophers take it to be equivalent to ‘objective’ but either that is parasitic upon a definition of ‘subjective’, which is equally hard, or gets the extension wrong. Others, perhaps thinking this is the same as ‘objective’, take externality to be mind-independence. But apart from problems like perceiving other minds and the role of observation in quantum measurement, this seems to get the order of philosophical explanation wrong, making the cogency of the philosophy of perception depend upon the cogency of realist metaphysics. My preferred definition makes being an object of perception sufficient for being external. Externality can be extended to other objects by bridging principles. This has the consequence that if there are sense-data, these are external merely in virtue of being objects of perception. Thus externality does little metaphysical work: the questions of whether the objects of perception are physical and mind-independent have yet to be answered.

³ Being an object of perception turns out to be a rather different property on different theories, but there is an intuitive grasp of the property which it is the job of those theories to explicate.

premise about the phenomenal characters of non-deceptive experiences which we cannot know directly and can only know via indirect methods which are theory-mediated, specifically with respect to theories of phenomenal character. Thus the argument from hallucination cannot be used to decide between those theories.

PRELIMINARIES

When a subject is phenomenally conscious, she is conscious. But furthermore, she is conscious in a specific way which can vary over time and across subjects. The specific way she is conscious when she is phenomenally conscious can be called 'phenomenal character' (this definition is now common in the literature but appears to derive from Hellie, 2007, 261).

If we want to explain phenomenal consciousness and character, the most basic thing we can say is that when someone is phenomenally conscious, there is something it is like for her at that time, which it may not be like for her at other times and which it may not be like for other conscious subjects then or at other times. When there is something it is like for a subject, then there is some G such that it is G for her then. G is the phenomenal character.⁴

Not all conscious creatures are phenomenally conscious at all times. It may also be the case that there are times at which it is reasonable to say that there is something it is like to be that subject, and even to specify what that is, without thereby saying that the subject is phenomenally conscious. For example, it is plausible that there is something it is like to be an adult or to be 40 without thinking that being an adult or being 40 has a phenomenal character. What it is like to be an adult is not specified in terms of a special quality of 'adulthood'; rather what it is like to be an adult will be specified in terms of responsibility, having dependants, making substantial decisions, etc. These are all properties of conscious subjects and yet not ways of them being phenomenally conscious. Thus, while one can assume that all adult humans are conscious and, much of the time, phenomenally conscious, to say that there is something it is like for them to be an adult is not thereby to say that they are phenomenally conscious. One has only said that the creature is phenomenally conscious when the way it is like for her at that time is qualitative, when the replacement for 'G' in 'it is like G for her' is a predicate of a qualitative property.⁵

So defined, phenomenally consciousness is episodic: there are times when individuals are phenomenally conscious and times when they are not, and amongst those times when they are, their consciousness takes on some rather than other phenomenal characters. We can identify and individuate these episodes of consciousness by their phenomenal characters, subjects and durations. This might lead us to think that what is phenomenally conscious, what has phenomenal character, is a (presumably mental) episode, and thus we can talk of conscious mental episodes as experiences, asking about the properties of those experiences, their phenomenal characters, which distinguish them from unconscious mental episodes. It is worth noting that this inference from the episodic nature of phenomenal consciousness to phenomenal character being a property of an episode, while almost universal in the literature, is not obligatory (cf. Stoneham 2008), and if we do make it, we commit ourselves to the reification of experiences, the treating of experiences as kinds of entities which can have properties.

⁴ Thus having G is sufficient for being phenomenally conscious. Since G is a 'what it is like' property, I am here in disagreement with Snowdon (2010).

⁵ It may not be possible to say which properties are qualitative and which not without making at least an oblique reference to phenomenal consciousness. But the intuitive distinction between the qualitative and the non-qualitative can still help elucidate phenomenal consciousness by unpacking a little of its structure.

Personally I find the claim that experiences are entities which have properties hard to reconcile with direct realism, especially if those properties include phenomenal character. For present purposes, talk of experiences is harmless, but when arguing in propria persona, I will use the noun as a stylistically preferred nominalization of the present participle ‘experiencing’, examples of which are perceiving, seeing, tastings etc.

THE ARGUMENT AGAINST NAÏVE PROPERTIES

Benj Hellie (2007) offers a very subtle version of the Argument from Hallucination aimed at rejecting views which hold that non-deceptive experience has a phenomenal character which in some way depends upon the external world but hallucinations do not.⁶ For presentational reasons, Hellie attacks a particular version of that view, but he believes that his argument applies to many other versions as well (2007, 264-5) and my criticisms will be entirely independent of such details. Thus for present purposes, we can redefine Hellie’s term ‘naïve property’⁷ as a property which is both a phenomenal character and depends for its instantiation upon the external world being some specific way. Thus, for example, if we were to account for the phenomenal difference between tasting espresso and tasting cappuccino in part by reference to the bitterness of the espresso and the milkiness of the cappuccino, then we would be invoking naïve properties, for that difference requires the real, external world espresso and cappuccino to have certain properties.

Hellie’s argument has the following premises:⁸

Instantial Base

No hallucinatory experience *instantiates* any naïve property.⁹ (2007, 278)

Instantiation

If an experience has a property as a phenomenal character, it instantiates that property. (2007, 278)

Apparent Spreading

⁶ In fact, Hellie thinks that our ordinary concept of phenomenal character has that consequence and thus presents his argument as a paradox. His official view is that we need different ‘weakenings’ of the ordinary concept for different ‘branches of philosophy of mind’ (2007, 289). If my argument here is correct, the ordinary concept is in good shape, but certain ordinary beliefs about the phenomenal character of dreams and hallucinations are unjustified.

⁷ The property is ‘naïve’ in the sense that it is a property attributed to experiences by theories which try to show that the naïve realist understanding of experience is correct.

⁸ This is the full-dress version of the argument. The shorter version drops Infallibility and strengthens Apparent Spreading and Apparent Naïveté by dropping the modifier ‘phenomenally appears’. Phenomenal appearances are the prima facie, independent, conscious justifiers of beliefs made on the basis of phenomenological study about the phenomenal character of experiences (274). They are introduced to enable the theory to make epistemological claims about the phenomenal character of experiences which occur in contexts where phenomenological study is impossible. They may turn out to be identical to phenomenal characters.

⁹ If after-images are considered as hallucinations, this may not be fully general because the phenomenal character of an after-image may be naïve. But see Margolis (1966) for some discussion of the relation between after-images and their causes.

If some veridical experience phenomenally appears to have a naïve property, some hallucinatory experience phenomenally appears to have that same naïve property. (2007, 280)

Infallibility

If an experience phenomenally appears to have a certain property, it has that property as a phenomenal character (2007, 282)

Clearly these premises collectively contradict the claim:

Apparent Naïveté

Veridical visual experiences phenomenally appear to have naïve properties. (2007, 280)

But Apparent Naïveté is precisely the claim which a philosopher interested in finding a difference between the phenomenal characters of hallucinations and non-deceptive experiences will be trying to maintain, for it is this claim, in conjunction with Infallibility and Instantiation, which allows external objects to partially determine phenomenal character and thus for consciously available justifiers to give us perceptual knowledge of the external world. To avoid Hellie's conclusion, it seems that such philosophers will need to reject Apparent Spreading.

Hellie also considers those who might protect Apparent Naïveté by rejecting Infallibility. Now, an unqualified rejection of Infallibility would not help someone who wanted to claim that the phenomenal characters of non-deceptive experiences are naïve properties, but a partial rejection which preserved it for non-deceptive experiences would help. Hellie offers two arguments for unrestricted Infallibility. The first is that it is definitional of phenomenal character and the second is that there are no plausible counter-examples. Both these arguments, however, are question-begging against someone who holds that hallucinations have no phenomenal character at all (e.g. Fish, Hyman)). Consider the first argument, namely that 'phenomenal characters are just the phenomenally apparent features' (2007, 282). Someone who thinks that hallucinations have no phenomenal character may be inclined to become a disjunctivist about phenomenal appearances: phenomenal appearances are either the phenomenal characters, or, in the case of hallucinations where there are no phenomenal characters, misleading appearances of phenomenal character.¹⁰ This form of phenomenal disjunctivism would entail the falsity of Infallibility (though, interestingly, not Self-Intimation) and thus avoid the problematic conclusion. However, many who are attracted by Apparent Naïveté would not wish to go so far as to deny that hallucinations have phenomenal character at all. Furthermore, as we shall see, it is possible to object to Apparent Spreading by arguing not that it is false, but that it is unjustified, which has the advantage of rejecting Hellie's argument without committing oneself to any specific account of hallucinations. So it seems that the best way to challenge Hellie's argument will be to attack Apparent Spreading.

Hellie initially offers two arguments for Spreading.¹¹ The first involves an appeal to the local supervenience of phenomenal characters, or Phenomenal Internalism (2007, 271). He

¹⁰ Hellie in fact rules out such a position by stipulation when he tells us that phenomenal appearances are necessarily conscious (2007, 275), though one might think that he is really talking about access consciousness there.

¹¹ See footnote 4 for an explanation of the relation between Spreading and Apparent Spreading.

recognizes that this is controversial and it certainly is: anyone who thinks that phenomenal characters are naïve properties is thereby committed to denying local supervenience. The second is a raw appeal to intuitive plausibility propped up with a claim about what it is possible to dream. Again, anyone who is in the business of constructing an account of phenomenal character which accepts Apparent Naïveté is quite likely to have some nuanced things to say about dreaming which will cast doubt on Hellie's 'intuitive' (which in this context just means 'without philosophical reflection or empirical research') claim about what dreams are possible. However, for those who are rightly unconvinced at this point, Hellie also offers an appendix with a 'painstaking case for Spreading' (2007, 272).

This argument is designed to address a fundamental problem with appeals to the phenomenal character of hallucinations: almost by definition, when undergoing a hallucination, we are not in a position to engage in a phenomenological study of our current experience, so on what grounds is it possible to make a claim about what the phenomenal appearance is? We are left in the position of someone wanting to make a claim about the visual appearance of an object he cannot see and which those who can see cannot report upon. Hellie's argument develops an interesting epistemology for phenomenal appearances which are not directly accessible to investigation.

The first step of the argument is to persuade us that for any actual experience there is a 'Cartesian alternative', that is a possible experience the possibility of which could 'be reasonably used in the course of an attempt to raise sceptical doubts about' (292) a non-deceptive experience. Again¹² Hellie appeals to dreams to support this claim, though it is noticeably weaker than the intuition used to directly support Spreading.

Since we cannot get at the phenomenal appearance of these hallucinatory/dream experiences via the direct route of phenomenological study, we instead 'project' ourselves into them. That is, we assign them a 'projective appearance', which is how 'we envisage it to be when we project ourselves into having it, from the inside, imaginatively' (293).

Just as a phenomenal appearances justify introspective judgements of phenomenal character, it seems that projective appearances are at least candidates for justifying non-introspective judgements of phenomenal character, though clearly they are not infallible. However, Hellie goes for an even weaker claim, namely that what we can conclude from an experience having a given projective appearance is that it has a phenomenal appearance which is similar to the projective appearance, which in turn would justify the judgement that it has a similar phenomenal character. And then he argues that if it projectively appears to have a naïve property, then the phenomenal character it has would justify judging that it has a naïve property. I think these subtleties are unimportant for my response, so we can work as if the stronger, less subtle claim were being defended, namely that the projective appearance that some experience is F justifies the judgement that it phenomenally appears F, which entails that it is F (given Infallibility).

¹² The first appeal to dreaming came in his defence of Base (2007, 271). But that was unnecessary: we can just stipulate that for the purposes of the argument an hallucination has no naïve properties. Since Base is universally quantified, we don't need there to be any instances to accept it. The job of establishing that there can be such episodes comes in the defence of Spreading.

OUR KNOWLEDGE OF OUR DREAMS

Before we examine the proposed epistemology of projective appearances, we should note that another substantive epistemological claim has been smuggled in unexamined. How do we know, when a sceptic offers us a scenario, that it really is possible? Consider the specific example Hellie gives:

Last night as I slept, I dreamt. My dream was of writing this paper. In my dream, I saw my fingers moving over a keyboard and text appearing on the screen of a laptop sitting on a messy desk. It was all very realistic. (2007, 291)

This is a dream report. We are expected to think it reports on a genuine possibility. Why? Well, the obvious answer is to say 'Because we have all had dreams like that', i.e. from our own knowledge of our own dreams we have a grasp of what kinds of dream are possible and this is a possible dream. But that is only a good answer if we know that we have had dreams and we know what those dreams were like. Now that may seem like the merest commonsense, a claim that no reasonable philosopher could deny, but in fact when one sits down to work out the epistemology of dream reports, it is much less obvious that we have this knowledge, at least that we know as much about our dreams as we think we do. We might be tempted to think that the epistemology of dream reports is simply that of episodic memory and we have no particular reason to doubt our memories.¹³ But this simple answer faces two major problems. First, dream reports are not at all like other episodic memories. To give two examples: when we recollect waking experiences we can usually fill in more and more detail as we put more effort into the recollection, but when we recollect dreams, much of the detail fades away with just one or two striking details standing out; and normally if we recollect a waking experience once, this reinforces the memory and makes it easier to recollect later, but (unless we write them down or make a big show of the telling) we quickly forget even the explicitly recollected dreams. Secondly, all memory reports are particularly susceptible to the contexts in which they are recollected (Loftus...) with different prompts or cues eliciting very different 'memories' of the same scene. Dream recollections tend to be prompted in very specific ways which other research on memory suggests would result in considerable inaccuracies and we have no independent check on their accuracy. For example, research conducted in Kleitman and Aserinsky (1953) seemed to show that dreaming occurred predominantly during REM sleep, but when the experiments were repeated with a slight change in the question which the subjects were asked upon being woken (from "what did you dream of?" to "what were you thinking of when you woke up?") there was an 80% increase in reports of dreams during non-REM sleep (Foulkes, 1962). Similarly, there is the well-documented phenomenon that in the early days of dream research, in the 1940s and 50s, the majority of Americans reported dreaming in black and white, whereas now it is rare to report black and white dreams. Perhaps how we dream has changed in response to cultural influences such as cinema and television, but inaccuracy of reports may provide a better explanation (Schwitzgebel 2002). As Owen Flanagan summarized the situation:

[a]lmost every dream report contains either more *or* less than was experienced or *both* more and less than was experienced in the dream state. On [a]lmost every view, dream reports are putrid evidence about *what was experienced* while asleep, even if we accept

¹³ In fact, there is some physiological evidence that our memories of dreams may be impaired (Flanagan, ...).

that they provide reliable evidence that *something* was experienced while sleeping. (1996, 102)

Since subjects are generally unaware when they are repressing or embellishing their ordinary memories in response to contextual cues and often insist most strongly on their accuracy when they are most embellished, we should have similar caution about dream reports. Furthermore, in the case of dreaming, we have no independent means of telling the difference between an accurate recollection and one which is highly, though perhaps not deliberately, embellished, so we cannot even begin to assess their reliability.

Even dream researchers like Allan Hobson whose whole research programme relies upon taking dream reports at face value are ready to admit that 'memory for even very vivid dreams is evanescent and tends to fade quickly upon awakening unless special steps are taken to retain it' (Hobson, Pace-Schott and Stickgold, 2000, 795, quoted in Sutton 2009, 527). The special steps include, amongst other things, systematic dream diaries completed every day immediately upon wakening. They may also include many more obvious prompts which, a cynic might suggest, tend to encourage embellishment. For example, in a major study of children's dreaming in a home setting (Resnick et al. 1994), parents were asked not to pressure children to report in the morning but there was no control for such pressure in the pre-sleep period and even worse, 'on five of the study nights, children were told, according to the investigative team's own instructions, to repeat "I will remember my dreams" three times out loud just before going to sleep' (Foulkes, 1999, 22). This will strike most parents as a routine which would almost certainly guarantee that children will make up or significantly alter their dream reports in order to secure parental approval. It certainly will not give reliable evidence about what was dreamt. The influences which may operate on adult subjects asked to complete daily dream diaries may be a little less obvious, but it doesn't take much cynicism to find them.^{14,15}

David Foulkes (1999, ch.2) uses such arguments to defend his own practice of using sleep laboratories and professional interviewers to record dreams of children. Interestingly, he is initially less concerned about these problems with respect to adult dreams because 'there the question of veracity doesn't bother us so much' (1999, 24). However, as I have been arguing, it is not only insincerity and intentional exaggeration which should worry us with respect to dream reports, but also much more subtle influences from social and other contexts which will be harder, if not impossible, for even a sophisticated adult to detect as influencing their reports. It is quite clear that Foulkes is convinced – presumably from personal experience and conversations with other scientists – that dreams occur and that accurate reporting of them is possible, even if hard to achieve. However, the difficulties in ensuring accurate reporting and the extent of inaccuracy in most reports of dreams made outside the laboratory, lead him to an interesting conclusion: there are two quite different phenomena being called dreams. A-type

¹⁴ For example, a study of 285 subjects 'indicate[d] that keeping a dream diary increased dream recall in low and medium [frequency] dream recallers but decreased dream recall in high [frequency] dream recallers' (Schredl, 2002). Did the dream diary help the former to remember their dreams or incline them to 'find' more dreams to put in the diary?

¹⁵ In his excellent survey of the state of dream research across the disciplines, John Sutton asks the pertinent question: 'Do individually and culturally variable beliefs about dreaming only influence dream reports, or is the form of dreams themselves in certain respects also malleable?' (2009, 536). Given that the only access to the content of dreams is via dream reports, it is hard to see how we could ever answer that question. Any subject whose dreams could be influenced by those beliefs is also a subject whose dream reports will be so influenced.

dreams are ‘the involuntary conscious experience of mentation during sleep’ (1999, 35), whereas B-type dreams are ‘a person’s everyday account of an experience ... as having occurred during sleep’ (1999, 36). While B-type dreams may be a worthwhile object of anthropological and sociological study, their content is determined by many factors over and above what actually happened during sleep, so they can tell us little about A-type dreams. To find out about those we need the careful methods of controlled laboratory research. While one might argue that Foulkes has overestimated the efficacy of laboratory methods to remove all the extraneous influences going into B-type dreams and thus give us pure reports of A-type dreams, it should be striking that the argument we are considering for the possibility of Cartesian alternatives seems to depend entirely upon claims about B-type dreams. If Foulkes is arguing along the right lines here, philosophers’ casual appeal to dreams only tells us what is possible in B-type dreams and does not tell us what types of conscious experience during sleep do or can occur.

Thus before we can even accept the first step of the argument for Spreading, we have to have an account of how we know that certain types of dream which constitute Cartesian alternatives are genuinely possible. A careful examination of the epistemology of dream recollection might well conclude that we do not have good reason to trust much of what we report ourselves as having dreamt. In particular, we might think that the reports of ‘very realistic’ dreams which are in all their details just like waking experiences are embellished by our expectations of what dreaming is like, expectations which have become embedded in our culture and into which we are socialized as children. In which case an appeal to our experience of dreaming to establish that Cartesian alternatives are possible would fail to be persuasive.

However, it might be that there is an alternative story about how we know that the Cartesian alternative dreams are possible. Perhaps we do not know we have had sufficiently similar dreams ourselves and we may have no reason to think that it is possible for an actual human to have the Cartesian alternative dreams, and yet we could think that the possibility described is still a genuine possibility. For we could say that even if our ordinary dream reports are largely unreliable or at least subject to serious doubt, we can conclude from the large number of such reports and their common patterns that dreams do occur and that they contain elements such as seeing one’s fingers moving over a keyboard, seeing words appear on a screen, or seeing a messy desk. While no actual dream may ever combine enough of these elements for it to be a Cartesian alternative, the Principle of Recombination (Lewis, 1986, Efirid & Stoneham 2008) might lead us to think that there is a possible dream which combines enough elements of enough actual dreams to be a Cartesian alternative to my current perception. We would need to do much metaphysics of experience¹⁶ to make this a legitimate use of the Principle of Recombination, but at least it would give us an epistemically more secure route to the first step in the argument for Spreading.

PROJECTIVE APPEARANCES AND JUSTIFICATION

Merely establishing the possibility of Cartesian alternatives does not suffice to establish Apparent Spreading, because the potential of such possibilities to generate sceptical arguments turns upon the conditions for knowledge of the external world whereas what we need to establish now is that the same justifiers for introspective knowledge of experience, that is, the same phenomenal appearances, are present in the two cases. This is achieved via the notion of a

¹⁶ Some philosophers do seem to think experience has a mereological structure which would suit it for applications of the Principle of Recombination, e.g. Chuard (2010, ??).

projective appearance: the hallucination projectively appears a certain way from which Hellie infers that it phenomenally appears a similar way.

The problem here is not the subtle move from projective appearance to phenomenal appearance, nor even the general claim that projective appearances are an important tool in coming to know about phenomenal characters we do not experience, such as those enjoyed by other people (Hellie, 2007, 298). The problem is the epistemic status of applications of the method of projective appearances specifically to hallucinations and dreams. While we may think that projective appearances are a reliable guide to phenomenal appearances in cases where our imagination is constrained by closeness to experiences we have made a phenomenological study of, why should we think it is equally, or at all, reliable when we are projecting into experiences of types which we may not have had? And even if we have had experiences of such types when dreaming or taking drugs, we have never been in a position to make a phenomenological study of them, and thus to make justified judgements about their phenomenal characters, because the very circumstances of having them are such that we are not able to meet the conditions for ideal phenomenological study (cf. Hellie, 2007, 267) when we are dreaming or under the influence of hallucinogenic drugs or – as seems common amongst reports of hallucinations from academic colleagues – when extremely exhausted.¹⁷

We should be very puzzled about why Hellie is so confident that projecting is a good method for determining the properties of hallucinations or dreams. He goes so far as to make a very bold claim:

Think of projective appearance on analogy to phenomenal appearance: when I imagine an experience of seeing a white picket fence [his example of a dream], I am aware of the dream in a way that immediately justifies judging it to have certain properties. (294)

Now, phenomenal appearances are meant to work just like perceptual, mnemonic and logical appearances, so presumably projective appearances have analogies here as well. Suppose I have never tasted a dish on the menu, but I read the recipe and imaginatively project myself into the experience of tasting it. Does that give me immediate justification for judging it to taste a certain way? Well, no one would want to deny that it will be fallible justification, for the flavours may interact in an unexpected manner, or the cooking process may affect the chemistry in a way I do not imagine. Also, some people are much better at doing this than others, and experience helps. By experience here, I mean, experience of the gustatory results of certain recipes: if I am to go from a recipe to a taste via imagination, I need to have learnt how types of recipe result in certain tastes. This is clearly the model Hellie has in mind for projective appearances and we certainly do have plenty of experience of projecting ourselves into experiences on the basis of a story and then later experiencing the same thing for ourselves, confirming or confuting the projection and thereby improving future projections.

Unfortunately, the situation with dreams – which are clearly Hellie's paradigm hallucinations – is very different indeed, for when I am dreaming I am not in a position to know what my

¹⁷ If we accepted the possibility of short-term envatted brains, we might think that being in such a circumstance would not harm one's ability to conduct a careful phenomenological study. But the case for the possibility of Cartesian alternatives for an envatted brain seems to rest on some internalist principle of local supervenience. It is possible that the argument in Martin (2004) from Experiential Naturalism could be used here, but it turns out that even that principle is not neutral on theories of phenomenal character (vide Stoneham 'Objectualism ...')

experience is like, I cannot conduct a phenomenological study of my experience. The same is true of drug or exhaustion induced hallucinations. What would be the equivalent here for imagined gustatory appearances? I guess a similar situation might be eating or drinking while engaged in an emotionally-charged argument – I think we all know that in such circumstances one is in no position to make an epistemically responsible judgement about the taste of the food or wine, for one often doesn't even notice what one is eating or drinking, even though it presumably does taste much the same as it would were one not distracted. Now suppose that there is some particular dish which one only ever eats while arguing like this (we might need the help of a philosophically mischievous dinner party host here), and what is distinctive about it is the method of preparation rather than the ingredients. Then one is given the recipe and asked to imagine what it tastes like – I take it to be fairly certain that one would be as unreliable here as one would be had one never tasted a dish prepared like that. Neither the fact that you have tasted it before nor an apparent memory of that taste make it any more likely that you will get it right, given the circumstances in which that tasting occurred. It is cases like this which provide the correct analogy for projection into the experience of dreams and hallucinations: one is not significantly more reliable if one has dreamt or hallucinated something similar to the sceptic's scenario than if one has not. Projecting into such scenarios is pretty close to guessing.

Thus the fact that, when someone projects themselves into the story the sceptic is telling, they find nothing in the projective appearance which grounds the judgement that the phenomenal character is different from the waking experience gives us little or no reason to think that the phenomenal character of the dream or hallucination is in fact similar to that of the waking experience. So we are left with no reason at all to accept Apparent Spreading.

THREE RESPONSES

One response might be that if projection into dreams and hallucinations was so unreliable, we would not expect so much agreement that the sceptic's scenario is indiscriminable from a waking experience. Surely if we were doing something close to guessing, as we would be in the gustatory case I constructed, then we would expect wide divergence in judgements of the phenomenal character of those experiences. To answer this we need to distinguish dreams from drug-induced hallucinations: in both cases our projections are mediated, but in different ways.

When we imagine ourselves into a dream of a white picket fence, what are we really imagining? Well, we are imagining a dream of a white picket fence. But how do we know about what it is like to dream? Not by the phenomenological study of dreams, for that cannot happen, but by dreaming and recollecting those dreams. So we are in fact imagining dreams through the filter of our dream reports. And the problems I raised above for our knowledge by recollection of our dreams shows that we cannot treat dream recollection as epistemically transparent. Hence, we can explain a great deal of the consensus on projection into sceptical scenarios by there being common factors influencing dream recollection, most notably a common culture of describing and discussing dreams, which is reinforced by literary, cinematic and philosophical traditions widely disseminated through our society which are likely to vary across cultures and even social groups. This is a point which is widely recognised in the social sciences, for example the 'new anthropology of dreaming' (Tedlock, 1991) which produces 'highly-nuanced, linguistically informed analyses of dream narration and interpretation as psychodynamic intercultural social processes'. Of course, what we are talking about here are Foulkes' B-type dreams, but since few

of us have been through the rigorous processes design to get accurate reports of A-type dreams, whatever reliability-implying consensus there is will be based on reports of B-type dreams.¹⁸

Few have experienced drug-induced hallucinations, so our projections here will not be mediated via a non-transparent recollective process. But common culture is still important, since the rarity of drug-induced hallucinations means that most of us, when projecting into such experiences, are drawing upon the same set of reports of those experiences, many of which are embellished for literary purposes or gathered by poor empirical methods. Furthermore, those who do have the experiences and report them often do so within a culture of expectations about what those hallucinations will be like. This suggests a clear sociological explanation of why there is so much consensus in our projections into drug-induced hallucinations.

What I have been arguing so far is that we lack any account of how we can know that there are hallucinations in the philosophers' sense, that is, possible experiences which are (a) Cartesian alternatives and (b) possessed of very similar or even exactly the same phenomenal character as some ordinary, non-deceptive experiences.¹⁹

A second significant line of response to this would be to argue that the evidence of generations of philosophers who have accepted the possibility of hallucinations should convince us that this is an epistemic possibility, and since there is no good reason to think that we have a case here where epistemic and metaphysical possibility come apart, we should accept the metaphysical possibility. Though rarely articulated like this, the argument from tradition to epistemic possibility to metaphysical possibility seems to underlie much of the incredulous reaction to the suggestion that hallucinations, as philosophers understand them, may not be possible.

The first thing to note about this line of thought is that not all apparent possibilities are epistemic possibilities. An apparent possibility may not even be an epistemic possibility for two reasons: it may be internally inconsistent in a subtle way which has not been noticed, or it may be inconsistent with some necessary truth that is known. A lot of science fiction represents objects and events which are not epistemically possible in one or both of these ways, but does it in such a plausible manner that many readers may take them to be possible. A clear example is the Babel fish in Douglas Adams' *Hitchhiker's Guide to the Galaxy*, which is meant to be able to translate from every possible language without a learning period. This apparent possibility is inconsistent with what we know about the nature of languages, namely that the connection between the physical form of words and their meanings is 'arbitrary'; words do not have their meanings essentially (even though it may feel like that to us) so there is no way of working out what a utterances of a previously unencountered language mean without some learning period.

An epistemic possibility fails to be a metaphysical possibility when it is inconsistent with some necessary truth that we do not (yet) know. These are difficult cases to describe correctly. Suppose that identity is necessary and someone in the 18th century does not know that heat is mean molecular kinetic energy. They conceive of a possible world in which heat is something else, perhaps the presence of caloric. Because what they have conceived of is internally

¹⁸ One may also be in general sceptical about how well results which require laboratory conditions to be measured may be indicative of extra-laboratory experience.

¹⁹ It is probably worth mentioning again that sceptical arguments turn on the claim that if you were in a Cartesian alternative, you would not be able to know that you were. This is weaker than the claim that the Cartesian alternative is phenomenologically similar to the non-deceptive experience. So the arguments given here will not engage with scepticism.

consistent and consistent with all the known necessary truths, we are not inclined to say that it is merely an apparent possibility. Rather, following Kripke's lead, we can say that they have conceived of a genuine possible world but mis-described it: the caloric world is one where there is no heat but there is another phenomenon which feels very much like heat and is confused with heat in the description of the epistemic possibility.

So the argument from tradition to epistemic possibility to metaphysical possibility must first establish that the tradition is not a merely apparent possibility, by establishing that it is internally consistent and also consistent with known necessary truths. Suppose we grant that, then the move from epistemic possibility to metaphysical possibility may be fine at the moment, that is, is consistent with all we know about the phenomenal character of dreams and hallucinations, we must still accept that it is mortgaged against future discoveries of necessary truths about these episodes. In particular, a metaphysics of consciousness might reveal that hallucinations and dreams with similar phenomenal character to non-deceptive experiences are in fact impossible and that what the tradition has been discussing is a mis-described possibility. Were someone to deny that hallucinations as Cartesian alternatives are metaphysically possible, they would still have to accept the epistemic possibility in which there are reports or intuitions of hallucinations matching non-deceptive experiences, much like one who denies that heat is caloric does not deny that there are sensations of warmth in the caloric world, but they would not have to take those reports or intuitions at face value. Crucially for present purposes, if the inference from epistemic possibility to metaphysical possibility is mortgaged to the correct theory of the nature of the phenomenal character of hallucinations, we cannot use the possibility of hallucinations as a constraint upon the metaphysics of phenomenal consciousness, that is we cannot say that there is an argument from hallucination which shows certain theoretical positions about phenomenal consciousness to be incorrect.

A third response might hold that we were too quick in claiming that dreams and hallucinations cannot be known by careful phenomenological study and are only ever known on the basis of later memories and reports. There are two sorts of case we should consider here: thought-experiments and certain clinical conditions.

Mark Johnston (2004, 122) gives an example of a philosophical thought-experiment designed to show the in principle possibility of hallucinations being available to phenomenological study:

You are undergoing an operation for an aneurysm in your occipital lobe. The surgeon wants feedback during the operation as to the effects of the procedure on the functioning of your visual cortex. He reduces all significant discomfort with local anaesthetic while he opens your skull.

So far so good, since no claims are being made yet about phenomenal character in such circumstances. We need to add more to the example:

He then darkens the operating theater, takes off your blindfold, and applies electrical stimulation to a well-chosen point on your visual cortex. As a result, you hallucinate dimly illuminated spotlights in a ceiling.

The point of stimulation is 'well-chosen' because Johnston wants to develop the example into a case of veridical hallucination. Notice that if the surgeon had very little knowledge about the detailed workings of the visual cortex, mine in particular for individuals may differ significantly

at this fineness of grain of brain function, and arbitrarily applied electrical stimulation, we could then legitimately wonder whether the effect was to produce a hallucination which was then available for phenomenological study or to produce a rather more complex phenomenon which resulted in the subject sincerely, but not necessarily accurately, reporting a hallucination due to a variety of factors.

So Johnston's thought-experiment appeals to the thought that if a visual experience is caused by the stimulation of your visual cortex which is in turn caused by the stimulation of the optic nerve and that by the stimulation of the retina, then we can 'cut out the middle man' and produce a matching hallucination by directly stimulating the visual cortex, and if we can do this in a way which doesn't impair cognitive functions, that hallucination is available for phenomenological study. Now, I have just noted that the antecedent of the second conditional, that 'the effects on one's brain [of the surgeon's actions] will be very similar in respects relevant to the causation of experience' (2004, 123), is based on the thought that the surgeon can have very detailed knowledge which allows him to intervene to produce exactly the effect on the cortex which would be produced by light hitting the retina, and nothing else. We certainly don't have that knowledge, and it is an empirical question whether, given the distributed nature of brain functions and the likelihood that specific regions are engaged in more than one function at once, it is possible to achieve that separation of effects.

Setting that question aside, however, we should consider the antecedent of the first conditional: a non-deceptive visual experience is caused by the stimulation of your visual cortex (which is in turn caused by ...). If that is false, then the surgeon could not know how to bring about a specific hallucinatory experience by directly stimulating a visual cortex. I don't need to argue that it is false (though as a matter of fact I do think it is), but merely note that it relies upon an assumption about the metaphysics of visual experiences, and thus about the phenomenal characters which are the properties of those experiences. The assumption is that visual experiences are episodes or events which sit somewhere in the causal nexus from light source, to reflective surface, to retina, to optic nerve, to visual cortex, to frontal lobe, to the primary motor cortex to muscle contraction, etc. One way of denying this would be dualistic parallelism or epiphenomenalism about experiences, but that is not what I had in mind, for the parallelist seems to share a characterization of experiences as being (immateriality aside) the kind of thing which could fit in a causal nexus. Instead we should note that there is a tradition in philosophical thought, a tradition which I have elsewhere called 'Pre-Kantian Innocence' (Stoneham 2006), which denies that to have a visual experience is to have your mind affected in some way, so the assumption Johnston needs is a sort of category mistake. If a visual experience is not part of the causal nexus but is a relation between subject and object which is not identical to but constituted by a set of underlying causal relations, then when the surgeon recreates a subset of that constituting basis, we cannot conclude anything about what will happen experientially. A rough analogy might help to explain this. A statue is constituted by a lump of bronze. Recreating one part of that lump may not even create a statue let alone one similar to the original.

The point here is not that Johnston's assumption is wrong, but that the example only motivates the thought that hallucinatory experiences can be open to phenomenological study if we grant a non-obligatory assumption about the nature of experience. Without that assumption, we have no reason to believe that the knowledge the surgeon needs in Johnston's thought experiment is even coherent.

Moving away from troublesome thought experiments, we come to troubling clinical conditions. Patients who suffer from congenital sensorineural deafness report experiences as if there was ‘a radio playing or a group speaking’ (personal communication). Crucially, these experiences do not appear to be associated with over-tiredness or excess of caffeine or alcohol, so we have no independent reason to think that the subject is unable to engage in careful phenomenological study.

There are two important questions here. One is whether such experiences are correctly characterized as hallucinations for the purpose of philosophical argument (as opposed to for clinical purposes), that is, whether they are hallucinations in the philosophical sense. The other is whether they do really occur in situations where there is no bar to careful phenomenological study. Since we cannot answer the second without conducting further research, let us address the first question only.

Consider tinnitus, which is a not uncommon experience. Should we classify it as hallucination? Given that DSM-IV defines hallucination as ‘false perception’, there is no doubt about the clinical classification and that clinical category clearly serves a good purpose. But for the purposes of a philosophical argument aiming to show we can draw conclusions about the character of non-deceptive experiences from the character of hallucinations, we need to be a little more careful. It seems that tinnitus is more like a visual after-image than a hallucination, for it is non-directional and does not vary with the normal ways of changing auditory perception (moving around, closing doors, putting fingers in ears etc.). While it is possible, on first experiencing tinnitus or even an after-image to make a false judgement about the world, it is quickly²⁰ apparent that the false judgement was not supported by the experiential data – we were rash. In fact, careful phenomenological study of after-images and tinnitus reveals that they have a fair bit in common with bodily sensations.²¹ However we should correctly categorise tinnitus, and there will be a range of alternatives here, it is not simply obvious that it is a hallucination in the philosophical sense. In particular, these phenomena do not provide evidence for either of two claims needed by Hellie, namely (1) that there are Cartesian alternatives to actual experiences, and (2) that either the phenomenal appearance or the projective appearance are of naïve properties.

So now we are in an interesting dialectical situation. Suppose we grant that these experiences had by the patient with sensorineural deafness are available for careful phenomenological study. Then we have to decide – presumably on the basis of that study – whether to classify them with hallucinations or with after-images or with something else entirely. This will be on the basis of their phenomenal similarity to those other types of experience. So if we are to say they are hallucinations, it will be because they are more phenomenally similar to hallucinations than to after-images. And now we have a problem. This sort of experience is only evidence of the phenomenal character of hallucinations if the experience is a hallucinatory one, but we can only classify it as that if we have knowledge of the phenomenal character of hallucinations. Thus these experiences cannot break us out of the epistemic prison and give us the access point we need for knowledge of the character of hallucinations. At the very most we can say that these experiences are of a category distinct from after-images, day-dreams etc. Whether we can then construct an argument from their phenomenal character to conclusions about non-deceptive

²⁰Perhaps not quite so quickly. If we start from a position of certainty about the cause of the experience (‘someone must have left a computer on’), we may be subject to confirmation bias about its being a perception rather than something entirely different.

²¹ An interesting intermediate category of experience are ‘floaters’ – the experiences caused by debris in the vitreous gel. While they do not fit the standard model of bodily sensations because those are normally associated with a position on the pleasure-pain spectrum, they do seem to be perceptions of the body qua perceptual organ.

experiences remains to be seen, but it is far from obvious that Apparent Spreading must be true for experiences like this.

A theme is emerging in these responses, namely that there is a special philosophical conception of hallucinations and dreams, which may be culturally quite widespread, but does not appear to be based in any solid evidence. What the evidence supports as actual or possible is rather different and not well-suited to the purposes of philosophical argument about phenomenal character.

THEORY-MEDIATED KNOWLEDGE

Could the argument from hallucination be improved by a better epistemology of dreams and other experiences which we cannot directly access via phenomenological study? Perhaps, but I want to argue that the argument so far is best interpreted as showing that knowledge of phenomenal characters which are not accessible by phenomenological study is theory-mediated. If that is correct, no argument which has a claim about the phenomenal character of hallucinations or dreams as a premise could ever be used to decide between theories of phenomenal consciousness. We can illustrate the problem schematically. Suppose person A claims to remember having a dream/hallucination with phenomenal character C. Person B might challenge this as evidence that dreams/hallucination can have C by asserting that what A in fact remembers is dreaming/hallucinating having an experience with C, but since dreaming that p does not entail that p, we cannot take this as evidence that A actually had an experience with C, only that she dreamt/hallucinated that she did. What further data could resolve the dispute between A and B?

Consider again the case of dream recollection. If we take this at face value, it looks like ordinary memory knowledge and, if any knowledge is not theory-mediated, it will be perception and direct remembering (Wiggins 1992). But what I was arguing above is that we should not assimilate dream reports to standard cases of recollection, for they seem in many ways quite different. Which is not to deny that we have an incredibly strong subjective impression that we are directly recollecting dreams, but the question is whether that is supported by even the subjectively available evidence, let alone the empirical data. Some subjects with déjà vu also have incredibly strong subjective impressions of remembering (Moulin, 2009?), but we are happy to discount those (rather than accept pre-cognition). If the subjective impressions of directly recalling dreams are not supported by the facts, then whatever the epistemology of dream reports turns out to be, it will not be a simple case of direct remembering and consequently we have plenty of reason to think that it is theory-infected, if not theory-mediated.

Now let's consider the case of projective appearances of Cartesian alternatives. When we consider the phenomenal character of the experiences had by our friends and acquaintances, we certainly can imaginatively project ourselves into their experience. But in so doing, our imagination is constrained in two ways: by the evidence of the external objects they are perceiving and by their own reports. In fact, our imagination is not merely constrained in these ways but it needs that kind of evidence before it can even get going. Suppose I tell you: John is situated 10 miles due north of you; you don't know whether he is indoors or outdoors, what is visible, audible, tangible etc. in his environment; you don't know what he is saying; and you cannot ask him any questions. Then I ask you to imaginatively project yourself into his experience. Of course you cannot do it. This is because the imagination is not an autonomous source of knowledge; what it allows us to do is to fill in gaps and draw out non-deductive

consequences of limited evidence. And by and large the way the sensory imagination achieves this is by relating what we are trying to imagine to similar experiences.²² That is why we cannot imagine the totally new (e.g. the taste of a new fruit) and we can, pace Hume, imagine the missing shade of blue.

So how does this imaginative projection work when we are considering experiences occurring in Cartesian alternatives which, by definition, are not related to the perceivable environment and are not being reported on by the subject. Are we any better off than in the case of John, 10 miles north of here? What we do have to go on is the sceptic's account of the subject's mental states. In our analogy, if we had added to the story of John that he believed himself to be walking through a pine forest on a sunny August afternoon, then that would give the imagination something to go on. For if that is what he believes, we can reasonably expect him to have perceptions which would (*ceteris paribus*) justify such beliefs, and that those would have a phenomenal character similar to my perceptions of such things. By analogy, if the sceptic's Cartesian scenario has the subject (falsely) believing that she is typing and watching words appear on the screen, then surely we can also expect suitably related experiences are in place to justify her beliefs?

Sticking to the dreaming case for simplicity, though the points apply equally well to drug-induced hallucinations, we quickly realize that the correct description of the mental states of someone who is dreaming is a substantive matter for philosophical theory. It is controversial whether someone who has a dream about seeing themselves typing actually has a visual experience or just imagines that she does, and equally whether she actually believes, while dreaming, that she is typing or merely imagines that she believes that (see Ichikawa 2009 for a summary of the positions and an argument that we merely imagine ourselves to be seeing and believing; Sartre is a good example of one who also thinks that hallucinations and dreams are imaginings (2004, 148-59)). Since seeing something and imagining seeing it are phenomenally different,²³ which philosophical theory of the nature of mental states had while dreaming is correct will influence what we can conclude about the phenomenal character of the episode, and even whether we need to postulate phenomenally conscious experiences to explain the data about belief that the sceptic is offering us.

A philosopher who has not thought about the issue of the nature of the conscious mental states involved in dreaming or hallucination may well not be aware that the sceptic's claim that in the sceptical scenario the subject has certain beliefs based on certain visual experiences depends upon a substantive and controversial theoretical claim. But the existence of rival theories, and the fact that those states are not available to direct, introspective observation makes it undeniable that imaginative projection into Cartesian alternatives is a heavily theory-mediated activity, and consequently the results of that projection cannot be used to settle disagreements between the theories which mediate it.

Owen Flanagan (1996) has tried to get around this problem by his 'natural method' which allows us to draw upon psychology and neuroscience to resolve disputes in phenomenology. His example is the question of whether dreams are experiences which occur during sleep, as traditional view holds, or whether we are simply disposed to make them up upon waking, which

²² The creative imagination may function rather differently, but it does not result in judgements about the phenomenal character of experiences.

²³ Also, merely imagined beliefs may not need even imagined experiences to justify them.

sceptics suggest (e.g. Malcolm and Dennett). It looks like appeals to dream reports cannot settle this (though some have thought that precognitive dreams and lucid dreams are relevant to the issue). Flanagan argues that if we find a neural correlate of waking consciousness, say 40Hz general oscillations, and that occurs in sleep which is followed by dream reports, then that gives us reason to challenge the sceptic. But clearly it is not a theory-independent reason, for much theorising has gone into finding the neural correlate. Furthermore, it makes theoretically loaded assumptions about the structure of causation, namely that the correlation between 40Hz oscillations and conscious experience in a waking subject is the same as that in a sleeping subject, where other variables are quite different. So it seems that, in our schematic example, the only evidence which could decide between A and B's claims about the phenomenal character of the dream/hallucination would be evidence that supported their general theory of those states and their phenomenal character.²⁴

CONCLUSION

While this discussion of arguments from hallucination began in a rather ad hominem vein as a critique of the sophisticated and nuanced version put forward by Benj Hellie (2007), the conclusion is entirely general. Consider, for example, a recent ingenious argument by Adam Pautz (2010). Pautz argues from hallucinations to the conclusion that phenomenal characters are intentional properties, which he takes to be relations to propositions. His conclusion entails that phenomenal characters are not naïve properties in Hellie's sense but the argument appears to have a very different structure, because it does not take a claim about the phenomenal character of hallucinations as a premise.²⁵ Pautz first argues that hallucinations have a property that only an intentional theory can explain, then that if we have an intentional theory of hallucinations, it is preferable to also have an intentional theory of non-deceptive experiences rather than any candidate disjunctive theory which treats the two entirely differently.

Setting aside the important question of whether Pautz's analysis of possible theories of phenomenal character is in fact exhaustive, let us consider the key premise that hallucinations have a certain property which needs explaining and which, it will turn out, only intentionalism can explain. The property is what he calls the 'grounding property' (2010, 268):²⁶ someone undergoing a hallucination has a 'visual experience property' (2010, 256) H which has the second-order property 'being a property such that, if a believer has it, then he thereby has the capacity to have [false (2010, 267)] beliefs involving [suitable external world properties]'. Thus hallucinating Macbeth has a property H which can ground his belief that there is a dagger before him, but not that there are pink rats before him. Grounding is weaker than justification but

²⁴ We can see this with simpler, everyday examples such as crying out in one's sleep. Suppose this was found to be correlated with a certain type of dream report, maybe one in which the dreamer is frightened, we could ask whether we should posit a frightening experience which causes both the crying and the report, or whether we should simply take the crying as a cause of the report, a 'combined confabulation-misremembering' (Flanagan 1996, 101) to explain the nocturnal experience of crying out. The preference for the first option expresses a belief that the usual explanation of waking crying, namely a distressing experience, is also the best explanation of sleeping crying.

²⁵ In fact, some of Pautz's sub-arguments do rest on such claims, but I concentrate here on the 'master argument', which appears different from other arguments from hallucination in just this respect. Also, his definition of 'visual experience property' simply assumes that in hallucinations 'your experience is, as we would put it, exactly the same' (2010, 256). However, the argument from grounding does not rest on that aspect of the definition.

²⁶ Pautz also argues from some other properties, but those arguments have more limited scope. The argument from the grounding intuition applies to all non-intentional theories.

stronger than causation, and in particular, it depends upon the phenomenal character of the experience. Thus we can assess how well a given theory of phenomenal character can explain the grounding property.

Pautz doesn't argue for the grounding property, merely asserting it is based on a 'grounding intuition' and assumes his readers will share that. It is very likely that the overwhelming majority of his readers will concur, but agreement is not enough for his argument, since the grounding intuition must be *pre-theoretical* for his argument to work. Even if some have the intuition whose theories are inconsistent with it, that does not show it to be pre-theoretical, for they may have a tacit commitment to a theory which is inconsistent with their explicit commitments. Widespread acceptance of a theoretically controversial intuition merely shows widespread (if tacit) acceptance of those theoretical presuppositions, not their correctness.

When we examine the intuition, we can see that it does not have the status Pautz needs it to have. Firstly, given that hallucinations are not available for phenomenological study, are not something we can have direct knowledge of, it is doubtful that any knowledge about hallucinations is entirely theory-neutral. Secondly, when we examine the particular claim about grounding, we see that it involves easily disputable theoretical commitments. Outside hallucination, it is only beliefs which are grounded and only by perceptual experiences, so the grounding relation holds in normal, non-hallucinatory cases between a perceptual experience or experiences and a belief. Thus, the grounding intuition is committed to someone undergoing an hallucination having both perceptual experiences and beliefs. If they do not have perceptual experiences but merely imagine they do, as some theories postulate of dreaming, then grounding does not hold, since imaginings do not ground beliefs. Equally, if they do not have beliefs but merely imagine they do, then again there is no grounding relation. Which is to say, Pautz's argument assumes that Macbeth has a grounded belief about a dagger, infers that the property H is a suitable perceptual experience, then asks how we can explain that experience. But when we ask why we should believe that Macbeth has a grounded belief about a dagger at all, we face exactly the issues I have raised against Hellie's argument.

Thus arguments from hallucination simply cannot be used to settle theoretical disputes about the nature of phenomenal character because the arguments themselves have to draw upon contested theories about the nature of phenomenally conscious mental states in order to establish their premises. Even if valid, they are epistemically circular.