

The Multidisjunctive Conception of Hallucination*

Benj Hellie

Abstract

Direct realists think that we can't get a clear view the nature of *hallucinating a white picket fence*: is it *representing a white picket fence*? is it *sensing white-picket-fencily*? is it *being acquainted with a white 'picketed' sense-datum*? These are all epistemic possibilities for a single hallucination: after all, phenomenological reflection suggests that the nature of that hallucination is *being acquainted with a white picket fence*; but the suggestion is misleading, and we have no further evidence about this nature. But if these are all epistemic possibilities for a single hallucination, they are all metaphysical possibilities for the hallucinations which subjectively match it. Hallucination of a white picket fence itself is, therefore, a disjunctive or "multidisjunctive" category. While this undermines MGF Martin's widely discussed variant of the "causal argument from hallucination" for his "epistemic" conception of hallucination, Martin's epistemic category still serves as a "reference fixer" for my many disjuncts.

Hallucination, edited by Fiona Macpherson; MIT Press, forthcoming.

If one wishes to endorse a “direct realist” view of perception, on which in a case of seeing, seen objects or their particular features are somehow parts of the experience one undergoes, what other doctrines must one endorse? Michael G. F. Martin has modified the classic “causal argument from hallucination” to support the claim that direct realists must endorse the doctrine that “For certain visual experiences as of a white picket fence, namely causally matching hallucinations, there is no more to the phenomenal character of such experiences than that of being indiscriminable from corresponding visual perceptions of a white picket fence as what it is” (Martin 2006: 369); this doctrine has lately gone under the name “the epistemic conception of hallucination”, or, as I will prefer, “epistemic disjunctivism”. The epistemic disjunctivist’s prediction that hallucinations lack an “independent nature” has, in turn, come under assault from a number of directions (Siegel 2004, 2008; Hawthorne and Kovakovich 2006; Sturgeon 2008). Accordingly, both advocates and opponents of direct realism have reason to wonder whether Martin’s causal argument from direct realism to epistemic disjunctivism is successful.

This paper aims to promote, as a direct-realism friendly alternative to epistemic disjunctivism, the doctrine of “multidisjunctivism”: the view, roughly, that the phenomenal character of *this* hallucination of an apple consists in acquaintance with a sense-datum of a certain sort, while the phenomenal character

of *that* indiscriminable hallucination of an apple consists in the representation of a certain apple-relevant proposition. (To foreshadow, multidisjunctivism stands to epistemic disjunctivism roughly as “reference-fixing” functionalism stands to “sense-giving” functionalism.) By granting hallucinations an “independent nature”, multidisjunctivism frees direct realism from many of the anxieties provoked by epistemic disjunctivism.

To promote the multidisjunctive view, I work toward a negative aim and a positive aim. The negative aim is to rebut a central plank in Martin’s causal argument. The plank in question appeals to a probabilistic relation among properties known as “screening off”, which I will explain below. Considered very abstractly, the causal argument begins with the claim that (for certain groups of properties) *these* properties screen off *those* properties from *these third* properties; advances to a certain claim about the causal relations among instances of those properties; and then draws out epistemic disjunctivism as a consequence of this claim about causal relations. My central contention will be that the case for epistemic disjunctivism breaks down because the screening off relation fails to hold in the case at issue. As we will see, the reason for this failure is that the screening off relation does not hold if multidisjunctivism is coherent: but multidisjunctivism *is* coherent.

This brings us to the paper’s positive aim, which is to argue that multidisjunctivism is not only coherent but is also *more plausible* than its direct

realist competitors: it is superior not only to Martin's epistemic disjunctivism, but also to such older positions as those that Byrne and Logue (2008) have labeled "positive disjunctivism" and "the moderate view". Roughly and briefly, the reason for this is that multidisjunctivism alone respects the link between epistemic and metaphysical possibility, in light of the irremediable uncertainty about the nature of a hallucination that is at the core of direct realism. Accordingly, those friendly toward direct realism need not worry themselves over the status of epistemic disjunctivism: they can and should endorse multidisjunctivism instead.¹ In doing so, they would accept both (against Martin) that a hallucination has an independent nature and (with Martin) that there is little to be gained from speculating about what it is.

Three points about the negative aspects of my discussion deserve to be set into relief at this early stage. The first point is that I will grant Martin the assumption that "No instance of the specific kind of experience I have now, when seeing the white picket fence for what it is, could occur were I not to perceive such a mind-independent object as this" (Martin 2006: 357), which we shall understand to be the doctrine labeled by "direct realism". My aim will therefore be to take sides in an internal dispute among direct realists: I will not be assessing *whether* to adopt direct realism, but rather wondering how, once direct realism has been *adopted*, it should be *implemented*.

The second point about the negative aspects of the discussion is that I will not

be contesting the success of any of Martin's attacks on the positions he has considered. In my view, these may well all succeed (however, I will along the way highlight certain points at which these attacks are somewhat imperiled). Rather, the point of the negative aspects of the discussion will be that multidisjunctivism remains standing even if the positions under attack fall.

The third point is that the issues clustering around screening off, causation, explanation, and kind membership are intricate and delicate, and involve semi-technical considerations from the metaphysics of science. Accordingly, my discussion of these issues will proceed at a relatively slow pace, recognizing that some of the central points under discussion may be unfamiliar to a significant portion of my readership; the reader who finds the material more elementary may find this discussion skimmable.

Roadmap: section 1 contains a parable, intended to vividly illustrate the content of multidisjunctivism. Section 2 explains the notion of screening off. Section 3 establishes the facts about screening off in the parable. Section 4 describes in detail the various positions under discussion—direct realism, epistemic and nonepistemic disjunctivism, and so forth. Section 5 presents in detail Martin's argument from direct realism to epistemic disjunctivism. Section 6 highlights the point at which Martin's argument breaks down if multidisjunctivism is coherent. Section 7 addresses a certain misunderstanding of the notion of screening off, which might seem to undermine the force of my

rebuttal. Section 8 argues on grounds for the most part independent of the foregoing that multidisjunctivism is the best version of direct realism. Section 9 draws together my central contentions.

<1> A Parable

God had a few larks when She created the actual world. The platypus is of course well-known; what is less familiar is that She introduced certain odd bits of whimsy into the psychophysical laws, the laws that specify the nomically possible correlations between the forms of consciousness and the distributions of matter.

Of particular relevance to our current topic is a bit of fun God had with perceptual experience. There is a certain genetic structure such that it is a psychophysical law that when people with that structure have perceptual experiences, they are *acquainted* with *sense-data*: that is, they bear a certain relation of awareness to mental particulars that mimic or image seen situations in the external world. For instance, to use the familiar terminology, when someone with this genetic structure sees or dreams about a white picket fence, he is also acquainted with a sense-datum which is (in the familiar Peacockean jargon) “colored” “white” and “shaped” “picketed”. And there is a certain distinct genetic structure such that it is a psychophysical law that when people with this second structure have perceptual experiences, they *represent intentional contents*:

that is, they go into certain intrinsically normative states that target or are about or are correct in a certain special, representational, way just in the presence of seen situations in the external world. For instance, when someone with this second genetic structure sees or dreams about a white picket fence, he also is in a state which is correct in the special way just in the presence of a white picket fence.

As it happens, these days every resident of Sweden has the first, sense-datum, sort of genetic structure, while every resident of Italy has the second, intentional, sort of genetic structure. (Of course, this would change if someone moved from one country to another, but this is unlikely to happen due to the climactic differences, so we can ignore this possibility.)

Suppose then that Sven and Stefano are both seeing a white picket fence as what it is, while Dagmar and Daniella are both dreaming of a white picket fence as such; Sven and Dagmar being residents of Sweden and Stefano and Daniella being residents of Italy.

Sven and Stefano have something in common psychologically: both are seeing a white picket fence as what it is. Let us suppose that this psychological commonality is captured by there being a *property* that they share: *seeing a white picket fence as what it is*, or SEEING, for short. Neither Dagmar nor Daniella has SEEING.

Sven and Dagmar also have something in common psychologically: both are acquainted with a white', picketed' sense-datum. Let us suppose that this is also

captured by there being a property they share, namely *being acquainted with a white', picketed' sense-datum*, or DATUMING, for short. Neither Stefano nor Daniella has DATUMING.

And finally, Stefano and Daniella have something in common psychologically: both are in a state which is representationally correct just in the presence of a white picket fence. Let us suppose that this is also captured by there being a property they share, namely *being in a state which is representationally correct just in the presence of a white picket fence*, or REPRESENTING, for short. Neither Sven nor Dagmar has REPRESENTING.

Each of our subjects reflects on the character of his or her experience and, quite naturally, judges to him or herself “I am seeing a white picket fence as what it is” — Sven and Stefano correctly, and Dagmar and Daniella incorrectly, as it happens. We can introduce a fourth property, JUDGING: *judging that one is seeing a white picket fence as what it is*. All four of our subjects have JUDGING.

<2> Screening Off

The notion of “screening off” (Salmon 1971) concerns a certain probabilistic relation among properties.² The core idea is that of a first property rendering a second *probabilistically independent* of a third. That is to say, whenever the first property is around, whether the third is also around does not influence the

probability that the second is around. (As I will understand the notion, it doesn't matter whether, when the first *isn't* around, whether the second is influences the probability of whether the third is.)

In symbols, the criterion can be written: $P(C/A \wedge B) = P(C/A \wedge \neg B)$. ($P(X/Y)$ is the conditional probability of X given that Y , a value in all but extreme cases (Hájek 2003) equal to the value $P(X \wedge Y)/P(Y)$.) Here, A is understood as the proposition that the first property is “around”, and so on. Just what it means for a property to be “around” is rather vague, but the vagueness can be usefully exploited to pick out a relation that varies from context to context: when considering correlations among certain symptoms of disease and certain candidate explanations, being “around” might be understood as being instantiated in the same organism or in the same organ, as appropriate; when considering correlations among certain economic properties and certain legislative properties, being “around” might be understood as instantiation in single municipalities or states, or in communities of trading partners, or what have you.

<3> Screening Off in Our Parable

What screens what off from JUDGING in our parable? On the table are three properties, hence three distinct pairs of properties which might be assessed in answer to this question: (i) SEEING, DATUMING; (ii) SEEING, REPRESENTING;

(iii) DATUMING, REPRESENTING. Does either member of any of the three pairs screen the other off from JUDGING?

To make sense of the question, we need a notion of “being around”. Let’s consider this one: being instantiated in a single subject within a short interval of time. Then the question is whether, in the case of pair (i), SEEING’s being instantiated in a subject renders DATUMING’s being instantiated in that subject probabilistically independent of JUDGING being instantiated in that subject within a short time interval, or vice versa, or neither.

My answer will be: in cases (i) and (ii), we do not have enough information to say one way or another; while in case (iii), we can be confident that neither screens off the other.

Let’s begin with (iii). The resolution of the issues here matter little to our overall goal, so we may discuss the issue briefly and then move on to more important matters. The questions are: if a subject has DATUMING, does this render whether the subject has REPRESENTING probabilistically independent of whether the subject has JUDGING? and: if a subject has REPRESENTING, does this render whether the subject has DATUMING probabilistically independent of whether the subject has JUDGING?

Consider the first question. Assume that it is impossible for a subject to have both DATUMING and REPRESENTING. If so, then the left hand side of the equation has a value equal to zero. We know that the right hand side has a positive value,

as both Sven and Dagmar are among the subjects with DATUMING and JUDGING. If it is possible to have both DATUMING and REPRESENTING, then this reasoning does not go through. One possibility is that each property would screen the other off: having both contributes no extra phenomenological “oomph” that having only one lacks, in a way that would lead to the judgement. That has an aura of plausibility to it. Still, resolving whether it is correct doesn’t matter for our purposes, so let us push on to pair (i).

Here the questions are: if a subject has DATUMING, does this render whether the subject has SEEING probabilistically independent of whether the subject has JUDGING? and: if a subject has SEEING, does this render whether the subject has DATUMING probabilistically independent of whether the subject has JUDGING?

Consider the former question. In order for the screening off condition ($P(C/A \wedge B) = P(C/\neg A \wedge B)$) to be met, it needs to be the case that whether a subject has or lacks SEEING, it is equally likely that he or she has JUDGING, given that he or she has DATUMING. Is this so? On standard developments of the views under consideration, the answer is yes: one who sees a certain scene is no more (or less) likely to make a particular judgement about the character of his experience than is one who merely has a perfect hallucination of that scene, implemented in a sense-datum. So DATUMING screens off SEEING from JUDGING.

Now consider the latter question: in order for the criterion to be met, it needs to be the case that whether a subject has or lacks DATUMING, it is equally likely

that a subject has JUDGING, given that he has SEEING. Is this so? Once again, we may suppose, the answer is yes. If a subject has SEEING and lacks DATUMING, this is presumably because he or she instead has REPRESENTING. And on standard developments of these theories, whether a subject has DATUMING or REPRESENTING does not influence the chances that he or she has JUDGING, given that he or she has one or the other. So SEEING also screens off DATUMING from JUDGING.

Note that the answer to the latter question would be different if it were not possible to have SEEING while having *either* REPRESENTING or DATUMING—for instance, if anyone who had SEEING had DATUMING. After all, then the probability of having JUDGING given that one has SEEING but lacks DATUMING would be undefined: the condition in question would be an impossibility, and thus of probability zero. Accordingly it would not be the case that the probability of having JUDGING given that one has SEEING is independent of whether one has DATUMING: the probability is defined if one has DATUMING, but undefined if one lacks DATUMING. In this case, SEEING would *not* screen off DATUMING from JUDGING.

The reasoning is of course exactly parallel for pair (ii). Accordingly, we can make the following assessments:

- i. DATUMING screens off SEEING and SEEING screens off DATUMING;
- ii. REPRESENTING screens off SEEING and SEEING screens off REPRESENTING;

- iii. Either DATUMING and REPRESENTING each screen the other off or neither screens the other off (which answer is correct depends on the outcome of certain questions in metaphysics).

Note that the assumption that it is possible to have SEEING while having either REPRESENTING or DATUMING—the multiple narrow realizability of the broad property—plays a crucial role here. If SEEING *entailed* DATUMING (or REPRESENTING), our assessment of (i) (or (ii)) would be false. As we have seen, the screening off would not go in both directions: DATUMING (or REPRESENTING) would screen off SEEING, but not vice versa.

When two properties screen each other off from an effect, I will say that they *symmetrically* screen each other off from the effect; when one property screens off, but is not itself screened off by, another from an effect, I will say the former *asymmetrically* screens off the latter. Our assessment, then, is that in our parable, SEEING and DATUMING symmetrically screen each other off, as do SEEING and REPRESENTING. The status of DATUMING and REPRESENTING is up in the air, but it won't matter for our purposes; henceforth I will ignore this pair.

<4> What is “Epistemic Disjunctivism”?

Let's look away from our parable now, and focus in on Martin's case for

“epistemic disjunctivism”. I begin by explicating the nature of that position; in the next section, I will run through my interpretation of the argument for it.

Consider the cases of seeing a white picket fence as what it is, and all their “matching” hallucinations: the cases of dreaming of seeing a white picket fence as what it is, the cases of hallucinating seeing a white picket fence as what it is due to electrical stimulation, and so forth. (To clarify, I understand these experiences as involving seeing a white picket fence of a certain highly determinate look, from exactly the same perspective, and so forth—put in enough qualifiers so that the experiences all “match” one another.) In a roughly Martinesque framing of the landscape, we may distinguish three positions about the “fundamental kinds” of these experiences, those features which are most definitional of or essential to those experiences.

I frame the landscape of positions in terms of “fundamental kind” following Martin (2004: 39, 43). An alternative understanding concerns the *experiential properties* of the experiences, the natural determinates of the property *consciousness* that they instantiate (Hellie 2010). I am inclined toward the view that the latter understanding cuts closer to the real issues here; and, for reasons I will detail below, is also a better fit with Martin’s argumentation. The reader should of course feel free to mentally substitute the alternative “experiential” formulation for the official “kind” formulation if it is felt that this helps to clarify the issue.

First is the “Cartesian” or “common factor” view (terminology here and below due to Byrne and Logue: 2008, sec. 5), on which all the cases of seeing a white picket fence as what it is, and all their matching hallucinations, share all aspects of their fundamental kind, any respects in which they differ being entirely contained within aspects external to this fundamental kind. Certain classical implementations of the sense-datum and intentional views are Cartesian views, in this sense. For instance, the former would say that DATUMING exhausts the fundamental kind of the experiences, while the latter would say that REPRESENTING does.

Second is the “moderate” view, on which while they share certain aspects of their fundamental kind, certain respects of difference between the experiences are also characteristic of their respective fundamental kinds. For instance, one might think that DATUMING exhausts the fundamental kind of the hallucinations, while both SEEING and DATUMING are part of the fundamental kind of the cases of seeing. Or one might think that both DATUMING and *lacking* SEEING are part of the fundamental kind of the hallucinations. Or one might adopt the variant of either position that results by substituting REPRESENTING for DATUMING.

Third is the “disjunctive” view, on which hallucinations and cases of seeing do not share any aspect of their fundamental kind. For instance, on one version of the disjunctive approach, SEEING exhausts the fundamental kind of the cases of seeing, while DATUMING exhausts the fundamental kind of the cases of

hallucinating.

(The distinction between the moderate and disjunctive views is due to Byrne and Logue (2008). Martin himself does not explicitly recognize this distinction; still, as we shall see, doing so is not crucial to his line of argumentation.)

We can distinguish two varieties of the disjunctive view, differing in respect of the line they take on the hallucinatory disjunct. To begin with, note the attraction of the view that there is *some* sense in which a case of seeing and its matching dream are “introspectively indiscriminable”: this is, for instance, how Descartes gets on to his characterization of the notion of matching. Just how the notion of introspective indiscriminability is to be elucidated has received extensive discussion (Siegel 2004, 2008; Martin 2004, 2006; Hawthorne and Kovakovich 2006; Sturgeon 2008; Byrne and Logue 2008; Fish 2009; Hellie 2007, 2010), with the key inadequately understood notions being the force of the modality, the identity of the discriminator, and the notion of introspection. Since my purposes in this paper are to *minimize* the significance of introspective indiscriminability, resolving these issues is not important for my purposes. Let us introduce one last property, that of being “introspectively indiscriminable” (in whatever sense is necessary to elucidate the notion of match) from an instance of SEEING (of not being introspectively knowably not an instance of SEEING): call this property INDISCRIMINABILITY.

What is the fundamental kind of the hallucinatory disjunct, then? We can distinguish two roles of the property INDISCRIMINABILITY in answering this question, which are loosely analogous to Kripke's distinction between the "reference-fixing" and "sense-giving" roles of a description. The analogue to the former, reference-fixing role, would say that the fundamental kind of the hallucinatory disjunct is just that property, whatever it might be, that grounds or explains that experience's having INDISCRIMINABILITY. Here the property that fills the role might be DATUMING, or it might be REPRESENTING. I'm assuming the relation of "ground" is antisymmetric, so that INDISCRIMINABILITY can't ground itself. (This is the position I used as an initial illustration of the thesis of disjunctivism, above.) The analogue to the latter, sense-giving role, however, would say that the fundamental kind of the hallucinatory disjunct just *is* INDISCRIMINABILITY: that it lacks any more fundamental kind in virtue of which it has INDISCRIMINABILITY. Let us call the former, "reference-fixing" position *positive* disjunctivism, and the latter, "sense-giving" position *epistemic* (sometimes in the literature "negative") disjunctivism.

Positive disjunctivism in turn comes in two varieties, which I shall label *multidisjunctivism* and *unidisjunctivism*. Unidisjunctivism is the positive disjunctivist view that there is a single property *F* such that for every experience with INDISCRIMINABILITY, the property that grounds the experience's having INDISCRIMINABILITY is *F*; multidisjunctivism is the conjunction of positive

disjunctivism and the denial of undisjunctivism—in other words, the view that there are some experiences with INDISCRIMINABILITY for which differing properties ground their having of INDISCRIMINABILITY. Multidisjunctivism is true in our parable.

The aim of the next section, section 5, is to present Martin's case against positive disjunctivism. Since Martin overlooks multidisjunctivism, I will ignore the availability of that position in explaining his argument. This is to say that in the upcoming section I will equate positive disjunctivism with positive *undisjunctivism*. Multidisjunctivism will return to the stage in section 6.

<5> Martin's Case for Epistemic Disjunctivism

We have four views on the table: the common factor view, the moderate view, positive disjunctivism (ssh—by which I mean, recall, positive undisjunctivism), and epistemic disjunctivism. As I understand it, Martin's case for epistemic disjunctivism takes the form of an argument by elimination, proceeding in two stages:³

- A. Against the Cartesian view, it is argued that the cases of seeing have SEEING as part of their fundamental kind (see Martin 2004: 42, 2006: 354 ; Byrne and Logue 2008: sec. 7.1);

B. Against positive disjunctivism and the moderate view, it is argued that if there is some property (other than *INDISCRIMINABILITY*) which all the matching hallucinations have as their fundamental kind, the cases of seeing do not have *SEEING* as even part of their fundamental kind.

I am broadly friendly to stage (A) (at least when understood as concerning the *experiential* properties of the case of seeing, rather than its *fundamental kind*), and do not propose either to defend the Cartesian view or to expound upon the case against it. Rather, I hope to intervene in an internal debate among anti-Cartesians. And while I will raise difficulties concerning stage (B), it does have considerable force against its intended target. In my view, however, this does not succeed in establishing epistemic disjunctivism, as there is a further position, which is an alternative to epistemic disjunctivism, which neither the considerations in (A) nor the considerations in (B) undermine. To see this, however, it will be important to take a close look at stage (B).

As I understand the argument in stage (B), it can be cast as follows (a technical term to be explained below will be put in italics):

1. If the hallucinations all have *DATUMING* (for instance; *mutatis mutandis* for *REPRESENTING*) as their fundamental kind, the cases of seeing at least instantiate *DATUMING* (maybe not as part of their fundamental kind);⁴
2. If the cases of seeing and hallucinating all instantiate *DATUMING*, *DATUMING*

- asymmetrically screens off SEEING from JUDGING *without compensation*;⁵
3. If there is some property *F* which asymmetrically screens off *G* from *H* *without compensation*, *G* never causes/explains *H* when *F* is around;⁶
 4. If *G* never causes/explains JUDGING, *G* is not the fundamental kind of any experience, not even of the cases of seeing.⁷

From these premisses, it follows that if the hallucinations all have DATUMING as part of their fundamental kind, SEEING is not the fundamental kind of any experience, not even the cases of seeing; the same for REPRESENTING. Martin also argues that if the hallucinations have INDISCRIMINABILITY as part of their fundamental kind, then even though INDISCRIMINABILITY asymmetrically screens off SEEING, it does not do so *without compensation*; accordingly, the hallucinations having INDISCRIMINABILITY as part of their fundamental kind is compatible with the cases of seeing having SEEING as part of their fundamental kind. All this together amounts to the case for (B). Let me now briefly expand upon the reasoning behind each premiss.

Concerning (1), I understand the reasoning behind this premiss as follows. First, any property which gives the fundamental kind of the hallucinations as of seeing a white picket fence must be a narrowly supervenient property like DATUMING, rather than a broad property like *being a case of DATUMING with an abnormal causal chain not leading back to any white picket fence*: “hallucinations

are 'inner experiences'. We have the conception that the occurrence of such events imposes no additional condition on the world beyond the subject's putative state of awareness" (Martin 2004: 58). Second, for any case of seeing a white picket fence, there is some hallucination of this sort which exactly shares its narrowly supervenient properties.⁸ If so, for any case of seeing a white picket fence, it instantiates DATUMING.

The second assumption seems very plausible (Martin provides it with a somewhat intricate defense). By contrast, my intuitions about the fundamental kinds of experiences are not entirely in order, so I am not certain why we should accept the first assumption: the alternative would be that the fundamental kind of a hallucination takes in its external conditions, such as its failure to be caused in a normal way by external conditions. However, I think it is possible to repair the justification for this premiss on the alternative way of understanding the issue, as dealing with the experiential properties rather than the kinds of experiences. Then the assumption would then be "the experiential nature of a hallucination is narrowly supervenient". This looks a bit more plausible (if perhaps still contestable): after all, in a case of hallucination, intuitively one's experience is "cut off" from one's environment; we should not think of a hallucination as a case of consciously taking in the abnormal causation of one's internal condition.

However things may be with this principle, I will let it slide, as I believe that the overall argument can be criticized on more straightforward grounds.

Concerning (2): consider first the claim that results from ignoring the proviso “without compensation”. This proviso-free claim is correct, as we have seen toward the end of section 3: recall that the case in our parable for symmetric screening off relies on the possibilities *both* that one might have SEEING while having DATUMING *and* that one might have SEEING while having REPRESENTING. What about the proviso? The notion of “compensation” will be a bit easier to grasp after we have discussed premiss (3), so I will return to (2) complete with proviso in just a bit.

Concerning (3): this is clearly a principle in the metaphysics of causation, and assessing it requires taking leave of matters strictly perceptual. Accordingly, providing a final assessment of the idea in this context would require an immense digression. Still, we can at least sketch a motivation for the view. The idea is roughly that causation is associated with difference-making: that a factor should not in general be thought of as causal (with respect to some domain) unless switching it on or off would make a difference (to that domain). In principle (3), the relevant notion of difference-making with respect to an effect is understood as influencing the probability of that effect against a certain background; if a property screens another off from an effect, then the second property doesn't make a difference to the probability of the effect if the first property is around.

The principle (3) waters down the idea that causation requires difference-making in this sense in two ways. First, by requiring that screening off must be

asymmetric: surely if *A* and *B* each screen the other off, we don't want to deny that in virtue of this neither is a cause; simple cases of causal overdetermination (multiple shooters) are such cases, but surely the effects in such circumstances are not uncaused.

Second, that the asymmetric screening off must be *without compensation*. What do I mean here? The idea is that causation is intricate and complex, and our philosophical understanding of it is limited. Accordingly, no flatfooted slogan like "causation requires difference-making" is likely to capture anything like a necessary condition for causation. Rather, flatfooted slogans of this sort can be regarded at best as defeasible tests for the presence or absence of causation: the status of a case with respect to a flatfooted slogan can be regarded as giving *prima facie* motivation for judgements about causation, but is unlikely to ever count as the end of the story. So I want to regard "compensation" as a sort of placeholder, meaning roughly "except in the presence of factors which defeat the *prima facie* reason that asymmetric screening off gives for denying causal force".

Accordingly, my sense is that this premiss is plausible. It doesn't seem to be entirely trivial, as the defeaters shouldn't be regarded as mere "except when it doesn't" cases, but should rather strike us as sufficiently legitimate as overcoming the requirement of difference-making to count as causes.

The "without compensation" proviso is important for the current dialectic, as

without it, *no* view could sustain the causal powers of SEEING with respect to JUDGING. Martin's idea, gestured to above, is that if the fundamental kind of all hallucinations is INDISCRIMINABILITY, these causal powers can be sustained. In a nutshell, the idea is that although INDISCRIMINABILITY asymmetrically screens off SEEING from JUDGING (and from any other property, for that matter), it does not do so "without compensation". In the present case, the particular form of "compensation" amounts to INDISCRIMINABILITY having its causal powers "dependently": in Martin's view, our understanding of causal explanations involving INDISCRIMINABILITY requires a prior grasp of causal explanations involving SEEING. Accordingly, were we to learn that SEEING has no causal powers, we would have to conclude that the same was true of INDISCRIMINABILITY. Accordingly, INDISCRIMINABILITY could rob SEEING of its causal powers only by robbing itself of its own causal powers; but since *those* causal powers were supposed to deprive SEEING of its causal powers, the deprivation wouldn't occur after all. This is a sort of existence proof for a weakening of the asymmetric screening-off requirement; the natural conclusion is that such a case of causal dependency counts as a "compensation" (this dialectic can be found at Martin (2004: 70)).

(It seems that another defeater, in Martin's view, is a sort of *proportionality* constraint. Narrow properties that are asymmetrically necessitated by wide properties always screen the wide properties off from everything, even wide

causes. Still, Martin agrees with Williamson (2000: ch. 2) that ‘he picked up a glass because he experienced as of a glass’ is inferior as an explanation to ‘he picked up a glass because he saw a glass’ (Martin 2004: 62). The explanans in the former case can occur in a wider range of cases than that in the latter case; a large proportion of the expanded range of cases would be situations in which picking up a glass would not occur; accordingly the explanans and explanandum are more tightly coupled in the latter case than in the former. Plausibly, such tight coupling makes for causation in a way that overrides asymmetric screening off. Our present argument concerns JUDGING, however, which is a narrow effect, and is thus equally coupled with SEEING and DATUMING.)

Let’s return to (2), now that we have a better sense of the “without compensation” proviso. We can now see that (2) is problematic in a way that the unqualified principle is not. After all, the proviso is supposed to amount to “in the absence of defeaters”, where at least one such defeater is the sort of explanatory dependency allegedly involved in the relation between SEEING and INDISCRIMINABILITY. In order to conclusively establish (2), then, we would need to be assured of two things: first, that properties such as DATUMING and REPRESENTING are not explanatorily dependent on SEEING; and second, that there is no other defeater (in a potentially open-ended list) that might override asymmetric screening off in the present case and restore causal powers to SEEING. I regard the status of both of these claims as open questions. Anyway, let us move

on.

Finally, concerning (4). I regard the rationale for this principle as elusive. I agree with Byrne and Logue (2008: sec. 7.2) when they write “Evidently [“that the phenomenal aspects of the experience in the good case are best explained by [DATUMING] rather than [SEEING]”] is supposed to be highly problematic, but what is not clear to us is why Martin thinks this”.

Once again, however, understanding the argument in the alternative way, as concerning not the fundamental kinds of experiences but rather their experiential properties, restores considerable plausibility to the principle. After all, it is quite bizarre to regard a property as an aspect of one’s conscious experience when it has no influence whatever over one’s introspective judgements.

In conclusion, my assessment is that each of these premisses is somewhat problematic. As officially stated, the rationale for each of (1) and (4) is not easy to grasp, although the plausibility of each premiss can be enhanced by recasting it in the alternative form, as concerning experiential properties. Still, assessing the principles as recast requires a deep investigation of questions about the nature of consciousness and its connection to introspection. (2) is, as we have seen, quite plausible when written without its proviso (although the case for that principle in section 3 contained a number of escape clauses); still, with proviso, its status depends on a fuller understanding both of the sort of causation that the sense-

datum theorist and intentionalist hope to appeal to, and on large-brush questions about the metaphysics of causation. (3) seems to be in good standing; however, as with any defeasible principle, there is always a worry about triviality.

<6> Martin's Argument in Light of Our Parable

All these quibbles can be set aside, however, as the argument faces a more fundamental difficulty. In a nutshell: as I have cast the argument, it is an argument by elimination, but a position is missing. Stage (A) eliminates Cartesianism, and stage (B) eliminates the moderate view and positive disjunctivism. Allegedly, the last view standing is epistemic disjunctivism.

But now it is time to return the distinction between positive undisjunctivism and positive multidisjunctivism to the stage. We have assumed that the positive disjunctivist view under consideration is *undisjunctivism*; accordingly, a *fifth* view, positive *multidisjunctivism*, remains uneliminated.

Common to all the four views discussed is that *all hallucinatory experiences as of seeing a white picket fence as what it is share a fundamental kind*.

According to the sense-datum version of Cartesianism, the moderate view, and positive disjunctivism, it is DATUMING; according to the intentional version of these views, it is REPRESENTING; according to the epistemic disjunctivist it is

INDISCRIMINABILITY.

But in our parable, things do *not* shake out this way. Rather, some hallucinatory experiences of this sort—including those occurring in residents of Sweden—have DATUMING; others—including those occurring in residents of Italy—have REPRESENTING. And what reason would there be at this point to suppose these properties to be other than the *fundamental* kind of these experiences? Suppose that there is none; then there is no common fundamental kind to all hallucinatory experiences. If the story in our parable is coherent, so is *multidisjunctivism*, according to which none of the cases of seeing shares its fundamental kind with any hallucination, and some of the hallucinations fail to share fundamental kinds with other hallucinations: the cases of seeing all have SEEING as their fundamental kind, while some of the hallucinations have DATUMING as their fundamental kind, and still others have REPRESENTING as their fundamental kinds.

(A sixth position, the *multimoderate* view, agrees with multidisjunctivism about cases of hallucinating, but disagrees about cases of seeing: rather, such cases have dual fundamental kinds, such as SEEING WHILE DATUMING or SEEING WHILE INTENTIONALIZING. I believe that Martin's argument succeeds when directed at this view, as I will explain in section 7.)

Stage (A) of Martin's argument is in no position to rule out

multidisjunctivism, of course. Nor is stage (B): the antecedent of the conditional, that there is some property such that all the matching hallucinations have it as their fundamental kind, is just what the multidisjunctivist denies. Nor could the argument be modified to attack the multidisjunctivist. Premiss (2) is the key here. If some of the cases of seeing and hallucinating instantiate DATUMING and others instantiate REPRESENTING, it is not the case that *either* of those properties asymmetrically screens off SEEING from JUDGING. *Rather, the screening off is entirely symmetric: this is just the lesson of judgements concerning pairs (i) and (ii) in the case of our parable.*

My conclusion, then, is that Martin's case by elimination for epistemic disjunctivism fails, due to the availability of multidisjunctivism. I rely on two claims in drawing this conclusion. First, that if multidisjunctivism is coherent, it is immune to stage (B) of Martin's case for epistemic disjunctivism, the "screening-off" argument, so that Martin's argument by cases for epistemic disjunctivism fails. And second, that multidisjunctivism is coherent, or at least worth taking seriously as a metaphysical view of the nature of perceptual experience. I see objections arising to both claims, which I will now address serially in the next two sections.

<7> Objection: "Causally Matching" Experiences

On five independent occasions,⁹ colleagues to whom I have presented my parable and its import for Martin's screening-off argument have urged that I am ignoring a central aspect of Martin's discussion: namely, that the argument is intended to be restricted to "causally matching" hallucinations. Perhaps the reader, too, feels unease about this, so in this section I will address the relevance of this restriction.

Martin doesn't offer an explicit definition of "causal match", but the notion as used is not hard to interpret in context. Its first appearance is immediately after a discussion of a "complete causal condition [in a subject's body] just prior to [the time of an event of seeing] which determined the chance of this event of seeing occurring" and a situation of hallucinating occurring in a subject whose body is in exactly that condition "even if no candidate object of perception is present and conditions necessary for the occurrence of a perception are not met" (Martin 2004: 53). Clearly then for two experiences to be causally matching is for the organisms in which they occur to be intrinsic duplicates immediately prior to those experiences. Martin acknowledges that hallucinations that fail to causally match the experiences for which direct realism is true might have a more "robust" property such as *DATUMING* or *REPRESENTING* (Martin 2004: 71); the screening off argument is clearly intended to support epistemic disjunctivism, then, only for experiences which causally match the experiences for which naive realism is true.

Putting the point more bluntly, since *our* experiences are supposed to be the ones for which naive realism is true, epistemic disjunctivism only needs to be true for the hallucinatory experiences which causally match our veridical experiences.

I can see two worries to which this restriction might give rise. The first goes by rejecting the claim that our veridical experiences are relevantly similar to either the Italian or the Swedish experiences in our parable: in that case *none* of the hallucinations in the parable would be causally matching to any of our veridical experiences, so that it does not matter if epistemic disjunctivism is not true of them. The second allows that our veridical experiences are relevantly similar to—picking a side arbitrarily—the Swedish experiences (*mutatis mutandis* for the Italian experiences), but argues that the Italian experiences are not causally matching, and are therefore irrelevant.

Concerning the first worry, it would seem entirely question-begging in the present context to deny that our experiences could be like the Swedish (or the Italian) experiences. After all, the main point of the screening off argument is to rule out a view on which our hallucinatory experiences involve sense-data (or intentional content) as their fundamental kind. To rule out my parable as irrelevant to our case on the grounds that its hallucinations fail to causally match our veridical experiences would just amount to insisting on the falsity of positive undisjunctivism and the moderate view.

Let us then turn to the second concern: granting that we are the Swedes (for instance), how is the condition of the Italians relevant to us? To be a bit more concrete about the concern, my response to the screening off argument would be rebutted as follows: Sven and Dagmar are, let us suppose, intrinsic duplicates; as are Stefano and Daniella. Then we can run the screening off argument twice: once focusing only on the hallucinating subjects who intrinsically duplicate Sven—Dagmar—and once focusing only on the hallucinating subjects who intrinsically duplicate Stefano—Daniella. The first use of the argument would establish epistemic disjunctivism for Dagmar, and the second use would establish epistemic disjunctivism for Daniella. Aggregating, epistemic disjunctivism would be established for all hallucinating subjects.

In a bit more detail, consider the use of the argument to establish epistemic disjunctivism for Dagmar. It would run something like this: restricting our attention to Sven and Dagmar, we see that both have JUDGING and both have DATUMING, while only one has SEEING. Accordingly, toggling SEEING while keeping DATUMING constant has no effect on JUDGING, while no probability is assigned to JUDGING on the assumption of SEEING and not DATUMING. Thus, DATUMING asymmetrically screens off SEEING from JUDGING (presumably, without compensation). This restores premiss (2), and the argument runs on unimpeded.

This line of thinking is alluring, but fallacious. The fallacy is the assumption

that we can restrict our attention in this way without changing the subject: while insuring that we continue to make true claims about the same old properties, JUDGING, DATUMING, and SEEING, that we were interested in assessing probabilistic relations among in the first place. Properties are *universal*, so to assess the probabilistic relations among certain properties, we need to know their dependency relations in *all* cases in which they appear, not in a restricted set of such conditions. Otherwise we have not assessed how things are with *those properties*, but rather only with *restrictions* of those properties: conjunctions of those properties with the restricting conditions.

A toy example of statistical reasoning in epidemiology might help to drive the point home. Suppose we are interested in whether job stress causes heart disease. Suppose we survey a certain segment of the population, and determine that in that segment, job stress and heart disease are uncorrelated. Would we thereby be reasonable in concluding that job stress does not cause heart disease? Not without further information about the segment of the population we have studied. If our segment was not genuinely cross-sectional, but rather was restricted to wealthy middle-aged white men, it might be more appropriate to conclude that job stress *in wealthy middle-aged white men* does not cause heart disease (perhaps because members of the studied group are in a position to compensate for job stress by other means, such as taking frequent relaxing vacations).

So, similarly, when we restrict our attention to the residents of Sweden, we

have not assessed probabilistic relations among the properties SEEING, DATUMING, and JUDGING; rather, only among properties along the lines of SEEING WHILE DATUMING, DATUMING, and JUDGING WHILE DATUMING (or perhaps SEEING WITH THE DATUM-INDUCING GENETIC STRUCTURE, and so forth). Accordingly, we have not assessed whether DATUMING asymmetrically screens off SEEING from JUDGING, but only whether DATUMING screens off SEEING WHILE DATUMING from JUDGING WHILE DATUMING.

This might all be used in an argument that goes on to establish that, if SEEING WHILE DATUMING is the fundamental kind of certain cases of seeing a white picket fence, epistemic disjunctivism is true of the hallucinations that causally match those cases. But who cares? We have no argument for the antecedent of this conditional. Indeed, if Martin is correct, we seem to have on the table an argument for a position incompatible with this one, namely stage (A) in Martin's case: that SEEING is the fundamental kind of *all* cases of seeing a white picket fence. To yield a case from stage (A) to epistemic disjunctivism, we need a case for stage (B); and this, I am arguing, is entirely lacking.

(Well, someone should care, namely the friend of the multimoderate view: recall, the view that causally matching veridical and hallucinatory experiences share *some* but not *all* of their fundamental kind. On this view, Sven's experience has as its fundamental kind SEEING WHILE DATUMING and Dagmar's has as its fundamental kind DATUMING. The considerations we have just seen suggest that—

granting its premisses—the screening-off argument *does* cut against this view: it does seem to show that if SEEING WHILE DATUMING is the fundamental kind of Sven’s experience, INDISCRIMINABILITY is the fundamental kind of Dagmar’s.)

The reader may still feel as though my arguments have a somewhat bamboozling character. If so, perhaps the following line of thought is worrying the reader: “focus just on Sven and Dagmar. If they were the only subjects that existed, the screening off pattern would hold, and SEEING would be causally irrelevant to JUDGING in Sven. And how can bringing in people who are *utterly extrinsic* to how things are with Sven—Stefano and Daniella—make a difference to how things are causally *with Sven*? It can’t: hence the rebuttal to the screening-off argument fails.”

The fundamental difficulty with this line of thought is that it fails to account for the fact that we are assessing causal relationships among tokens of properties *by means of* assessing probabilistic relationships among the properties themselves. As far as assessing probabilistic relationships among properties goes, people “utterly extrinsic” to how things are with Sven matter insofar as they share or fail to share properties with him. At least, it would seem that if making appeal to such people is illegitimate, so is making appeal to *Dagmar* in assessing how things are with Sven. Her situation, too, is, from the point of view of properties, “utterly extrinsic” to how things are with Sven only if that of Stefano and Daniella is as well. We may have some intuition that Dagmar is somehow “closer” to Sven’s

situation than are Stefano and Daniella. But this intuition cannot be captured inside of the austere framework of relations of probabilistic dependence and independence among properties. Perhaps the intuition will at some point be shown to follow from some relatively precise and productive treatment of causation, but until then it does not seem (to me at least) to rise above the status of a vague unease; and from the point of view of blocking Martin's argument as it appears on the page, such a residual unease is entirely irrelevant.

Somewhat relatedly, the thought that we can consider probabilistic relationships among properties by focusing only on one world—the world without Stefano and Daniella—is the error that lies behind frequentist treatments of probability, with their many failures to capture the ordinary concept (Hájek 1997). Probability is immersed not just in actuality, but in possibility as well. Just which possibilities are relevant to probability is contested. But since properties spread out through modal space, it would be shocking if relationships among them as significant as probabilistic relations did not reflect this.

So—if our parable is coherent—I can see no principled reason to exclude reasoning about probabilities on the basis of it, if our aim is to assess causal relationships among particular situations by establishing probabilistic relationships among the properties instanced in those situations. There might still, of course, be some other reason to deny that SEEING can be part of the fundamental kind of Sven's experience if DATUMING is part of the fundamental

kind of Dagmar's. My aim is not to deny that there is any possible case for epistemic disjunctivism, but is rather the more moderate aim of rebutting the "probabilities-to-causes" strategy embedded in the screening off argument. If my parable is coherent, it seems to me, this probabilities-to-causes strategy must fail.

<8> On the Plausibility of Multidisjunctivism

This section argues that multidisjunctivism is the most plausible direct realist position on hallucination. I begin by addressing further concerns about my rebuttal of Martin's argument, stemming from the metaphysics embedded in our parable. The story is, one might think, a bit weird: and this weirdness might be thought to diminish the significance of my objection in one of two ways.

First, Martin frequently proclaims his acceptance of the doctrine of "Experiential Naturalism", a doctrine which rules out certain metaphysical possibilities as holding of the actual world. Accordingly, if our parable takes us out into one of the ruled out worlds, objecting to the case for epistemic disjunctivism on the grounds that the argument for the view fails at those worlds, my line of attack would not be to the point: the issue is whether epistemic disjunctivism holds at the actual world (and its near neighbors).

Experiential Naturalism, however, is the fairly anodyne doctrine that

our sense experiences, like other events or states within the natural world, are subject to the causal order, and in this case are thereby subject just to broadly physical causes (i.e. including neurophysiological causes and conditions) and psychological causes (if these are disjoint from physical causes). (Martin 2004: 39–40)

I don't think anything in my parable is at odds with this. The strongest dualist thesis my parable requires is that the psychophysical laws are contingent. I see no reason to deny that this assumption is compatible with the view that "our sense experiences [...] are [...] subject to broadly physical causes [...] and psychological causes".

A second concern about my rebuttal might stem from the fact that that the parable requires any amount of dualism at all. My rebuttal gets a purchase only if the sort of multiple realization by internal mental properties I have described is *possible*, and indeed possible for us (or at least beings like us). If dualism is true, I think we should grant the possibility of my case. After all, dualism is typically motivated by conceivability arguments, and the situation in the parable seems clearly coherent. However, while dualism is taken more seriously now than twenty years ago, still many philosophers might think that if the price of avoiding epistemic disjunctivism is dualism, they'd prefer to run with epistemic

disjunctivism.

To allay this concern, I should confess that the parable was chosen entirely because it is vivid and easy to understand: I don't think the aspects of the structure of the parable that are relevant to its place in the overall argument rely on the truth of dualism. To see this, consider a sort of simple "reference-fixing" version of physicalist functionalism, on which narrow mental properties are identical to brain properties (or disjunctions of brain properties), and whether a given term for a narrow psychological kind denotes a certain brain property is determined by whether the brain property is the entity which satisfies certain causal "platitudes" associated with the term.

Suppose that the terms 'DATUMING' and 'REPRESENTING' are associated with causal platitudes which are distinct, though overlapping in relation to matters concerning introspection. Then it might be that the brains of a certain population have structures which satisfy the former platitudes but not the latter, while those of a certain distinct population have structures which satisfy the latter platitudes but not the former. Finally, suppose that the platitudes associated with the broad term 'SEEING' fail to incorporate any of the platitudes outside of the region of overlap. Then it might be that members of either population could satisfy 'seeing' (when in the correct broad conditions) just when they satisfy the narrow term appropriate to their population.

Such a situation would be analogous in all relevant respects to the one in our parable. Accordingly, if it is coherent, direct realists can avoid epistemic disjunctivism without slipping into dualism. Given the utterly anodyne materials out of which the situation is constructed, I see no reason to doubt its coherence. But if it is coherent, I see no reason to deny that it is possible.

A somewhat more programmatic reason for regarding the scenario as possible can be sketched. A frequently encountered strand in direct realist thought is the elusiveness of the internal. Because we reflectively discriminate among visual experiences solely by attempting to assess the natures of their presented external objects—in accord with the “transparency of experience”—it is only their presenting objects of such natures that can be known to be among their fundamental kinds. In the case of a hallucinatory experience, which lacks this feature, the fundamental kind remains unknowable—*pace* the epistemic disjunctivist. But if it is unknowable, then there are a range of epistemic possibilities for what it might be. Both dualists and functionalists about the relevant fundamental kinds may at this point see an immediate path to a range of *metaphysical* possibilities for the fundamental kinds of indiscriminable hallucinatory experiences: such properties “reveal their own natures” under their canonical conceptions, so that there is no room for a range of perspectives on an underlying unary fundamental kind.

But if the story is possible, then multidisjunctivism may well be *true*.

Multidisjunctivism and its competitors are metaphysical theses about the natures of visual experiences: accordingly, they are intended to apply to *all possible* visual experiences. It does not matter for deciding among these views whether there is any of the sort “mixture” observed in our parable *around here*, or at the *actual world*; rather, these views should be assessed by the extent to which they are adequate to all possible experiences.

Seen in this light, we can provide an argument for multidisjunctivism: nondisjunctive or “Cartesian” views are eliminated by stage (A) of Martin’s argument and epistemic disjunctivism is eliminated by considerations discussed by the authors mentioned in the introduction. The unary moderate view and positive unidisjunctivism are eliminated because they are insufficiently general: if our parable is a possible scenario, then matching hallucinations must differ in their phenomenal kinds (unless epistemic disjunctivism is true). Eliminating these views does not require anything like Martin’s intricate probabilistic reflections. Such reflections are necessary only for the purposes of eliminating the multimoderate view (as discussed in the previous section). Multidisjunctivism is the only view left standing.

<9> Conclusion

I have argued for the following conclusions.

First: situations like the one described in our parable are possible: there is a pair of possible subjects undergoing subjectively indiscriminable hallucinatory visual experiences, such that the psychological properties of their visual experiences differ. I have argued that adherents to a range of positions about the metaphysics of mind should accept this thesis: in particular, dualists and reference-fixing physicalist functionalists should accept this, as should functionalist direct realists who are friendly to the transparency of experience.

Second: it follows from this possibility that a range of positions in the philosophy of perception are in trouble: positive unidisjunctivism and the unimoderate view, as well as “unary Cartesian” views are incompatible with the possibility. Only epistemic disjunctivism, positive multidisjunctivism, the multimoderate view, and “multi-Cartesian” views are compatible with the possibility. If we reject the Cartesian views (perhaps on the grounds that they are incompatible with direct realism), and we reject the epistemic conception of hallucination (perhaps on the grounds that a robust essence is needed to explain what we *can* know about hallucinations), the only views standing are positive multidisjunctivism and the multimoderate view.

Third: according to both multidisjunctivism and the multimoderate view, there

is what Kennedy (this volume) calls “mission creep”. Suppose that I see an apple and then hallucinate an apple due to direct artificial stimulation of the optic nerves, and no significant alterations to the constitution of my brain are made in the interim. Then whatever the fundamental kind of my hallucinatory experience is, my veridical experience shares that property: the multimoderate asserts, and the multidisjunctivist denies, that this property is among the fundamental kinds of my veridical experience.

Fourth: that in light of this mission creep, the multimoderate view entails, but multidisjunctivism does not entail, that the fundamental kind of my hallucinatory experience screens off the property of *seeing an apple* from my reflective judgements about what sort of visual experience I am enjoying.

Fifth: that Martin has provided us with a credible basis upon which to reject any view entailing this screening off thesis, and hence the multimoderate view.

Sixth: putting this all together, we have an argument by elimination of alternatives for multidisjunctivism.

So multidisjunctivism is the last view standing. Should we be happy about this? I am inclined to think that aside from the truth of some sort of direct realism, no claim about the metaphysics of perceptual experience has any distinctive pretheoretic allure.

At the same time, there is a certain respect in which multidisjunctivism

occupies a uniquely conciliatory position in the space of views we have canvassed. (And, accordingly, there is more to be said for it than that it remains uneliminated once a barrage of intricate argumentation has been brought against its competitors.) The guiding idea behind epistemic disjunctivism, it seems to me, is that it is a consequence of direct realism that the first-person perspective on hallucination cannot be expected to yield deep insight into the intrinsic essence of such experience. The epistemic disjunctivist takes from this the moral that there *is* no such intrinsic essence into which to acquire deep insight. The multidisjunctivist sees the lesson as being rather that the advantages of positing intrinsic essences can be preserved even if the first-person perspective is merely insensitive to, or “divides through” by, certain ranges of variation in this intrinsic essence.

Accordingly, multidisjunctivism represents a sort of compromise between epistemic disjunctivism and the range of competing theories postulating “robust” essences for hallucinations: the multidisjunctivist agrees with the robust theorist that a hallucination has a substantive nature of its own, while also agreeing with the epistemic disjunctivist that there is little profit in speculation about this nature.

Footnotes

*. Thanks are due to Alex Byrne, Jim John, Heather Logue, Matt Nudds, Ian Philips, Howard Robinson, and especially Susanna Siegel and Jessica Wilson.

1. My central complaint against Martin is in essence parallel to one of the two justifications Williamson (2000) provides for broad psychological explanation. Martin acknowledges one of them (Martin 2004: 64), that discussed in chapter 2 of Williamson's book: the widely recognized justification in terms of "proportionality" available even in cases of unique realizability, for broad explanation of broad explananda (more on this below) (Williamson 2000: ch. 2; see also Yablo 1997). Martin blocks appeal to this strategy by picking narrow explananda (Martin 2004: 64; see also Byrne and Logue 2008: sec. 7.2, 2(c)).

However, Williamson develops a *second* justification for broad psychological explanation in his *third* chapter, which Martin (along with the rest of the literature on this issue) has largely overlooked. This is available in cases of multiple realizability of a wide property by narrow properties, and is the essence of Williamson's "alpha-beta-gamma" cases (see also Williamson 1998). This strategy can provide for broad explanation even of narrow effects.

2. Perhaps there is also a more causal or explanatory sense to the expression, involving one property's depriving another of causal or explanatory force, but I intend it to be understood as having this solely probabilistic sense.

I suspect Martin is using the expression in the other sense. Consider the following passage: "Being red is more highly correlated with being sorted into the left hopper than being scarlet. So the property of being red here seems to screen off the property of being scarlet from having an explanatory role" (Martin 2004:

62). Here Martin seems to be suggesting (“so”) that the probabilistic relationship (“correlated”) provides a (nonconclusive?: “seems”) *reason* to believe in “screening”—which it could not do if it were just *analytically identical* to “screening”.

I don’t think, however, that we are talking at cross-purposes: merely using the expression ‘screening-off’ to denote differing elements of a common structure. Martin and I both distinguish the probabilistic relationship among properties from the causal/explanatory relation among their tokens; we agree that the probabilistic relationship would give reason to believe in the causal/explanatory relation, but (as we shall see below) disagree whether the probabilistic relationship holds. Beyond that, there is a difference of practice, in which Martin applies the terminology of ‘screening-off’ to the causal/explanatory relationship while I apply it to the probabilistic relationship.

Obviously there are certain advantages to granting Martin his use of the terminology; my view is that the advantages of using the terminology in the more established sense are greater.

3. My understanding of the argument has been greatly aided by Byrne and Logue (2008).

4. This seems to be a slightly but inessentially simplified version of Martin’s (6) (Martin 2004: 54). For discussion of the case for this principle see

pp. 54–8 of that paper; also Byrne and Logue (2008: sec. 7.2), where this stage goes under the rubric of “part 2(a)”.

5. For discussion see Martin (2004: 61–2, 64); also Byrne and Logue (2008: sec. 7.2), where this thesis is part of “part 2(b)”.

6. Here note that, as observed above in fn. 2 **Error! Reference source not found.**, Martin regards screening off as a defeasible test for causal/explanatory irrelevance. For discussion of defeaters, see Martin (2004: 63–4, 69–70); also “part 2(c)” in Byrne and Logue (2008: sec. 7.2).

7. “It would be a severe limitation on the disjunctivist’s commitment to Naïve Realism, if the Naïve realist aspects of perception could not themselves shape the contours of the subject’s conscious experience” (Martin 2004: 64).

It is not fully clear to me how Martin is understanding the notion of the “contours of conscious experience” here. The talk of “shaping” and the overall context of a discussion of screening off and explanation suggests a view according to which Naïve aspects are *distinct* from the contours of conscious experience, in particular from the “phenomenal aspects of experience” (59).

And yet outside of the context of the screening off argument, Martin refers repeatedly to “Naïve phenomenal properties” (49, 50, 53) and discusses the Common Kind Assumption on which matching veridical and hallucinatory experiences are of the same “phenomenal kind” as allegedly antithetical to direct realism (40). It is not entirely clear to me how Martin is understanding the notion

of a phenomenal property here, but according to the traditional use of this notion a phenomenal property *constitutes* the sort of experience one is undergoing. If Martin intends to speak with tradition here, this suggests that Naïve Realists think that some phenomenal properties *are* properties like *seeing a white picket fence*.

But then the claim that Naïve phenomenal properties are screened off from the shaping of the contours of conscious experience would then seem to make no sense: if they *are* the contours of conscious experience then they can't be screened off from themselves! It seems that on pain of rendering the case for epistemic disjunctivism nonsensical, we must reject the characterization of Naïve Realism as a doctrine about the natures of phenomenal properties, understanding it instead as a doctrine solely about the fundamental kinds of experiences. But this puts the epistemic disjunctivist in need of an argument against the positive (uni)disjunctivist who is willing to rest content with the initial characterization of Naïve Realism as concerned with the natures of phenomenal properties.

(It is noteworthy that throughout the discussion of the causal argument beginning on page 53 and reaching the conclusion that “there are certain mental events, at least those hallucinations brought about through causal conditions matching those of veridical perceptions, whose only positive mental characteristics are negative epistemological ones” on page 74, discussion of the “phenomenal” drops out—the incidental reference on page 59 aside.)

It is to bypass this exegetical and dialectical morass that I cast the argument in terms of JUDGING rather than in terms of some phenomenal notion.

It is a nice question whether the relationship between JUDGING and SEEING must in fact, as I am granting here, be *causal*. Kennedy (this volume) denies this. Put rather crudely, his view is that to count as an experiential property, SEEING must merely *rationalize* JUDGING: it must merely be the case that typically, self-ascriptions of SEEING formed in a way that exploits the fact that one instantiates SEEING count as knowledge; and that the sort of efficient causation the presence of which the screening off condition is designed to probe is not a necessary condition on this rationalization.

I agree with Kennedy's criticism of the principle. At the same time, the underlying question is rather deep, so that other readers may be less sympathetic. Such readers may welcome the more formal character of the criticism I offer. (Kennedy offers a version of positive undisjunctivism as an alternative to the epistemic approach. I am less sympathetic to the style of reasoning he brings on behalf of his alternative for reasons I discuss elsewhere: Hellie 2010, note 5; moreover, Kennedy's view is not compatible with the case for multidisjunctivism.)

8. I take it this is the import of premisses (1)–(3) at Martin (2004: 53).

9. Namely, exchanges with Heather Logue, Alex Byrne, an anonymous referee, Matt Nudds, and Ian Phillips.

References

- Byrne, A. and Logue, H. (2008) 'Either/Or', in A. Haddock and F. Macpherson (eds.) *Disjunctivism: Perception, Action, Knowledge*, Oxford: Oxford University Press.
- Fish, W. (2009) *Perception, Hallucination, and Illusion*, Oxford: Oxford University Press.
- Hájek, A. (1997) 'Mises Redux Redux: Fifteen Arguments Against Finite Frequentism', *Erkenntnis*, 45:209–27.
- _____ (2003) 'What Conditional Probability Could Not Be', *Synthese*, 137:273–323.
- Hawthorne, J. and Kovakovich, K. (2006) 'Disjunctivism', *Proceedings of the Aristotelian Society*, supplementary volume, 80:145–83.
- Hellie, B. (2007) 'Factive Phenomenal Characters', *Philosophical Perspectives*, 21: 259–306.
- _____ (2010) 'An Externalist's Guide to Inner Experience', in B. Nanay (ed.) *Perceiving the World*, Oxford: Oxford University Press.
- Kennedy, M. (2011) 'Explanation in Good and Bad Experiential Cases', this volume.

- Martin, M. G. F. (2004) 'The Limits of Self-Awareness: Disjunctivism and Indiscriminability', *Philosophical Studies*, 120:37–89.
- _____ (2006) 'On Being Alienated', in T. S. Gendler and J. Hawthorne (eds.) *Perceptual Experience*, Oxford: Oxford University Press.
- Salmon, W. (1971) *Statistical Explanation and Statistical Relevance*, Pittsburgh: University of Pittsburgh Press.
- Siegel, S. (2004) 'Indiscriminability and the Phenomenal', *Philosophical Studies*, 120:90–112.
- _____ (2008) 'The Epistemic Conception of Hallucination', in A. Haddock and F. Macpherson (eds.) *Disjunctivism: Perception, Action, Knowledge*, Oxford: Oxford University Press.
- Sturgeon, S. (2008) 'Disjunctivism about Visual Experience', in A. Haddock and F. Macpherson (eds.) *Disjunctivism: Perception, Action, Knowledge*, Oxford: Oxford University Press.
- Williamson, T. (1998) 'The Broadness of the Mental: Some Logical Considerations', in J. Tomberlin (ed.) *Language, Mind, and Ontology*, volume 12 of *Philosophical Perspectives*, Oxford: Blackwell.
- _____ (2000) *Knowledge and its Limits*. Oxford: Oxford University Press.
- Yablo, S. (1997) 'Wide Causation', in J. Tomberlin (ed.) *Mind, Causation, and World*, volume 11 of *Philosophical Perspectives*, Oxford: Blackwell.