

Experiential Pluralism and Mental Kinds

Maja Spener (University of Birmingham)

1. Introduction

We perceptually experience the world in various ways. Visual experiences are perceptual experiences which arise in exercises of one of our perceptual capacities, namely vision. These perceptual experiences essentially involve conscious visual presentation of the world: when undergoing visual experiences, it is visually for the subject as if things are presented to them. This paper is about the nature of visual experience and how to argue about it.¹ At its centre is a new argument for the view that there is kind of visual experience that is world-involving. The argument draws on the explanatory role of experience vis-à-vis certain aspects of our agency – specifically, certain types of abilities we routinely attribute to ourselves and others in ordinary situations. Before setting out the argument, let me first introduce some key notions and outline the background of the relevant debate in the philosophy of perception.

It is common to recognize the following distinction. Some visual experiences are part of seeing the world, in that they figure in a cognitively successful relation to the world where you perceive the world as it really is. Some of them fail to do so. Call the first sort of case ‘good case’, the other sort ‘bad case’. My visual experience as of the carob tree in the court yard is part of a good case if it is part of my seeing the carob tree and the court yard, but it would be a bad case if it looks to me that way when there are no carob trees within appropriate visual range. Mixed cases are cognitively successful in some respects but not in others.

Mere recognition that there are good and bad cases does not involve any commitment about whether the visual experiences at the centre of them are of different mental or psychological kinds in some substantial sense.² Both the following views are compatible with the distinction:

Experiential Monism: Visual experience has a single common psychological nature.

Experiential Pluralism: It is not the case that visual experience has a single common psychological nature.

Experiential monists hold that good and bad cases involve visual experiences which are of the same basic mental kind. According to them, what makes for a given case being good or bad resides in factors that are external to the visual experiences (typically in their causal aetiology and certain counterfactual conditions); there is nothing that distinguishes the visual experience as good *qua* significant mental kind. Experiential pluralists hold that good and bad cases involve visual experiences which differ significantly in their mental nature. On the basis of this they can argue that the cognitive success or failure of the good and bad cases, respectively, in which such experiences figure, is due to that difference. They hold that there is something that distinguishes the visual experience as good *qua* mental kind.³

¹ My paper focusses exclusively on visual experience, but I will drop the qualifier ‘visual’ at times.

² The category of mental kinds is very broad, including a variety of psychological aspects. Standard candidates are phenomenal character and various types of content.

³ The labels ‘experiential monism’ and ‘experiential pluralism’ are originally from (Snowdon, 2008). On my use experiential pluralism captures a great variety of positions. I am avoiding the term ‘disjunctivism’, since the

These positions come in different strengths. According to extreme versions of experiential pluralism, e.g. there is *nothing* mental, or nothing of any significance, in common between the good case and the bad case. More precisely, an extreme version involves denial that there is anything mental in common between the good and the bad case, other than what the experiences involved have in common with coarser varieties of mental states (e.g. being a visual experience as of a tree, etc.). Less extreme versions of experiential pluralism leave room for some mental properties in common between the good and the bad case, where these common properties are specific to these cases. Such moderate versions include the claim that the good case has at least one such mental property which is not shared by the bad case. This good-case-only-property is held to be significant in certain explanatory respects, and, on this basis, moderate versions judge the good case as involving an important mental kind.⁴

In this paper, I provide an argument for a moderate experiential pluralism. It is a novel argument because it does not lean on phenomenological or conceptual constraints on visual perception, or on epistemic constraints of the kind that extant arguments for this position do. Instead, the argument appeals to the explanatory role of experiences in the good case vis-à-vis certain aspects of our agency: possession of certain kinds of ordinary abilities.⁵

Further, I make a methodological point about how best to argue for moderate experiential pluralism. The latter takes off from the observation that claims about experiential pluralism versus experiential monism have little substance without further elucidation of the central notion of a mental kind involved. Moderate positions in particular ought to be clear on the matter. They allow for common mental properties between the good and the bad case, but nonetheless maintain that there is a good-case-only-property that grounds a significant mental kind. To understand this, we need some sense of why the mental kind in question is to be theoretically prioritized. In particular, the import of a moderate view depends on the conception of the wider explanatory significance of the good-case-only-property, and of the mental kind it forms in light of this explanatory role.

The methodological connection between wider explanatory significance and mental kindhood does not get enough attention in the philosophical debate about perceptual experience. It is often difficult to understand the broader importance of the claims about our psychological

latter is often used by proponents and critics of more specific positions on the nature of the experience all of which I include on the experiential pluralist side (such as naïve realism (e.g. Martin 1997, 2002, Campbell 2002) and content-based disjunctivism (e.g. McDowell, Tye 2007)). I include epistemic disjunctivism (e.g. McDowell 2010, 2013) so long as the view grounds the epistemic difference between the good and the bad case on a psychological difference between the two cases. Soteriou (2016) uses the label ‘metaphysical disjunctivism’ to cover views that postulate a psychological difference between the good and the bad case, and ‘epistemic disjunctivism’ to cover views that postulate an epistemic difference between the two cases (specifically, a difference in the kind of perceptual warrant for judgements accessible to the subject in the two cases). Epistemic disjunctivism does not automatically involve endorsing metaphysical disjunctivism. See (Soteriou, 2016, esp. 197-211) for discussion of different kinds of disjunctivism, terminologies and taxonomies and pp. 146-151 for discussion of the metaphysical commitments of McDowell’s epistemic disjunctivism).

⁴ Byrne and Logue (2008) use the label ‘moderate disjunctivism’ to cover such views and ‘metaphysical disjunctivism’ for what I am calling extreme experiential pluralist views. Their use of the latter term thus differs significantly from Soteriou’s use.

⁵For direct appeals to phenomenal adequacy conditions, see e.g. (Martin 1997, 2002). For appeals to the distinctive warranting role of perceptual experience, see e.g. (McDowell 2010, 2013). For appeals to conceptual constraints, see (Kalderon 2018). (Soteriou 2016) provides an excellent overview and detailed discussion of the variety of disjunctive positions and debate about them. My argument is perhaps closest to aspects of Campbell’s case for naïve realism, which includes appeals to the role of perceptual experience in explaining our capacities for perceptually-based, demonstrative thought about mind-independent objects and their features. But at bottom, Campbell’s case crucially rests on appeals to introspectively-acquired phenomenal adequacy conditions as well (for discussion, see Soteriou 2016, 83-114).

nature that are being put forward. As a consequence, it is often difficult to relate the conception of mind put forward by proponents of experiential pluralism to conceptions put forward in other areas, especially in those that take empirical psychology as a guide. By contrast, as I will explain, my argument for experiential pluralism is sensitive to these methodological considerations. The main purpose of the paper thus goes beyond merely defending a particular position on the nature of experience according to the terms of current debate. More ambitiously, it aims to show what a thorough case for such a view requires and to propose one route to meeting this demand. In turn, providing a successful route to meet the demand counts as additional support for my argument.

The plan for the paper is as follows. I briefly set out some preliminaries concerning the debate about visual experience in section 2. In section 3, I state the argument and defend it. In section 4, I explain the methodological point in relation to the current state of the debate, and show how my argument addresses it.

2. Visual experience in the good case

Experiential monists and pluralists agree on two conditions on what counts as a visual experience in a good case ('veridical perception'). The first condition is that such experience must consciously present or convey the world the way it actually is. In a very loose sense of content, such visual experiences must have 'veridical' content. An experience in the bad case, of course, can also have this feature: hallucinations can convey the world as being a certain way, when the world really happens to be that way. Moreover, hallucinations can be subjectively indistinguishable from veridical perceptions, when for instance, a subject hallucinating a carob tree in the court yard cannot tell from the inside ('by introspection') that they are not veridically perceiving the carob tree in the court yard. The second condition on which both views agree is that the good case must involve a certain kind of perceptual causal contact with the world. A subject in a good case is perceptually connected with the world such that it is *in virtue* of this connection that the experience is veridical. The perceptual-causal connection must rule out that the visual experience in the good case satisfies the first condition accidentally or as a result of a contrived artificial manipulation.

There are different ways to think about how these two conditions on visual experiences in the good case relate, i.e. how whatever satisfies the second condition, makes satisfying the first condition non-accidental and non-contrived. One suggestion – forming the focus of this paper – is to think of the two conditions as intimately related: visual experience in the good case presents the world the way it actually is, because it is itself in part constituted by the relevant aspects of the world. What is it for visual experience to be in part constituted by relevant aspects of the world? There are different options for how to spell this out, e.g. that the phenomenal character of the experience has a relational structure with worldly objects as *relata*, or that its representational content has object-dependent truth-conditions.⁶ But the core thought in common to these options is this: a visual experience *E* is constituted by a given object *o* exactly when the occurrence of *E* alone guarantees the appropriate presence of *o*. *E*'s occurrence conditions include *o*. Let us call such a visual experience a 'world-involving visual experience'.

⁶ There are two aspects to conscious visual presentation: content (in the loose sense) and phenomenal character, and a wide variety of views as to how these aspects relate to one another. As a result, there are different options concerning which of these aspects is the focus of world-involvingness (see (Haddock & MacPherson, 2008, p. 17ff especially). One might hold that the phenomenal character of visual experience in the good case is in part constituted by the worldly objects and properties presented in the experience, or that it is the content of such experience that is so constituted.

The idea of world-involving visual experience is of central importance in the debate between Experiential Monism and Experiential Pluralism: the latter typically appeal to it in defence of their view. However, as we saw earlier: there are two separate tasks involved in an adequate defence. If there are world-involving visual experiences *and* these types of experience form a substantive mental kind, then Experiential Pluralism is true (granting that hallucinations count as genuine experiences). Thus, an adequate argument for experiential pluralism should provide reason for both the claim that visual experience in the good case is world-involving, and the claim that such experience forms a substantive mental kind. Providing reasons for merely the first claim is not sufficient.⁷ The argument in the next section serves to support both claims.

3. The argument from abilities

Here is the argument:

- P1: We have situation-dependent abilities.
- P2: Possession of some of these situation-dependent abilities is explained by appeal to visual experiences in the good case.
- P3: If visual experiences in the good case play this explanatory role, then they are world-involving.
- C1: There are some world-involving visual experiences.
- C2: Experiential pluralism is true.

Let me now clarify and defend the premises.

3.1 *Situation-dependent abilities (P1)*

We are able to do many things: we are able to get to the museum, pick strawberries, play catch, etc. We attribute a wide array of abilities to ourselves and others all the time. It is an important part of how we make sense of one another, in particular, how we understand what kind of agents we are.

Abilities are properties of agents. Like other powers that agents have (e.g. competences, skills, potentialities), abilities can be had even when they are not exercised or manifested. I have the ability to walk even when I am standing still. This modal aspect of abilities admits of a distinction commonly made between what someone can do in a more general sense and what they can do with respect to more specific situations. For example, sitting at my desk I might say that right now I cannot swim but that nonetheless I can swim. The gist of this is clear: that in my current circumstances, which lack an adequate body of water nearby, I am not in a position to swim, but that, given such a body of water, the right temperature, etc. I would be in a position to swim. How these two claims about my overall swimming powers – one involving denial of ability, one involving attribution of ability – can be true at the same time is an interesting question. It is an instance of a more general issue that has been discussed extensively

⁷ Tyler Burge, for instance, has long advocated a view of certain mental states as object-dependent, but is staunchly against any version of experiential pluralism (see, for example (Burge, 2011; Burge, 2009)). Object-dependence, on his view, is an aspect of the context in which the thinker or perceiver has a certain type of mental state – spelled out in terms of his notion of application. It is not an aspect of the psychological type of the relevant state itself. For Burge, genuine psychological kinds are those that play a role in scientific psychology.

in debates on free will, abilities and dispositions. For instance, Mele distinguishes between general practical abilities and specific practical abilities:

Although I have not golfed for years, I am able to golf. I am not able to just now, however. I am in my office now, and it is too small to house a course. The ability to golf that I claimed I have may be termed a *general* practical ability. It is the kind of ability to A that we attribute to agent though we know they have no opportunity to A at the time of attribution and we have no specific occasion for their A-ing in mind. The ability to golf that I denied I have is a *specific* practical ability, an ability an agent has at a time to A then or to A on some specified later occasion. (Mele, 2003, p. 447)

Mele's distinction is routinely drawn by others using different terminology, and it is reflected in work on the semantics of abilitative senses of 'can'.⁸ I will use 'general abilities' and 'situation-dependent abilities' to mark it. An ability is usefully individuated by appeal to an action-type (in Mele's case, golfing, or in my example above, swimming) and a range of situations in which actions of this type are available, that determine the modal profile of the ability.⁹ Attribution of *general abilities* concerns what a person can do in some (suitably wide) range of situations. That is the sense in which I can swim as I sit at my desk, i.e. there is a range of situations where I am in a position to swim, though I am not in one then. Attribution of *situation-dependent abilities* concerns what a person can do indexed to a specific situation. That's the sense in which I cannot swim right now as I sit at my desk, i.e. in that situation, I am not in a position to swim. Often, we are interested in those situation-dependent abilities that an agent has in the actual situation they are in, but not always. I might consider, for instance, what I would be able to do if I were at Berlin central train station right now.

The two kinds of ability are connected. In fact, it is plausible that they are 'something like two modes of a single kind of power' (Maier, 2018a) and that one type is ontologically prior to the other. Work on free will, dispositions and powers, and on the semantics of 'can', includes accounts for one type of ability (or abilitative 'can') in terms of the other. Both directions of order of explanation have been explored. The connection is more complex than a simple type-token relationship, since there are cases where we can have a situation-dependent ability without a corresponding general ability, and vice versa (Whittle, 2010, pp. 3-4). While it is thus uncontroversial that there is a distinction between situation-dependent abilities and general abilities, the exact nature of that distinction is not. But my paper does not depend on this controversy being resolved in any particular direction.

Situation-dependent abilities constitute an important aspect of our agency and are critical to us in our everyday dealings with one another. We constantly refer to what we can and cannot do in certain situations. Assessment with respect to our general abilities is also significant in making practical sense of ourselves, but situation-dependent abilities have a central, structuring role in our daily planning and getting on with things. Having a good idea of someone's agency relative to a specific situation is crucial to predicting and assessing what they will do, and to determining how they and others around them should act on a given occasion. Call this aspect

⁸ (Whittle, 2010) distinguishes between 'local abilities' and 'global abilities', (Maier, 2015) between 'options' and 'general abilities', (Inwagen, 1983) between an ability as 'the power of an agent to act' and 'general ability', (Honoré, 1964) between 'can (particular)' and 'can (general)', (Berofsky, 2005) between 'token ability' and 'type ability'. See also (Bird, 1998) for a related distinction concerning the attribution of dispositions.

⁹ I do not assume any particular theory of abilities in this paper, only that attributions of abilities are associated with special restrictions on possibility. Relevant contemporary work on the semantics of 'can' predominantly relies on a framework proposed by (Kratzer, 1977), though see (Vetter, 2013). See also (Maier, 2018b) for critical development of the modal view of abilities. For an overview of theories of abilities, see (Maier, 2018a).

of our agency determined by situation-dependent abilities our ‘situational agency’, noting that it is not meant to prejudge overall conceptions of rational and non-rational agency.

(P1) is therefore uncontentious. There are many things that I am able to do, relative to the situation that I am in, or to a situation of interest. Given some specific situation, the class of my situation-dependent abilities determined in relation to it is typically much smaller than the class of what I can do more generally in this situation.

3.2 Explaining possession of situation-dependent abilities (P2)

How do we attribute abilities to ourselves and to others? At least one routine way of doing so exploits an explanatory link we commonly take for granted between having certain mental states or episodes and having certain abilities. Here, I focus on an explanatory link between having visual experiences in good cases and having certain kinds of situation-dependent abilities. Let me give some examples of the routine explanations at issue.

Finding Home: Suppose you are standing outside your office building, looking around. You are visually aware of the things around you, the street, the houses, etc. This puts you in a position to find your way home now. Of course, having the visual experience on its own does not put you in this position, there are other relevant aspects of you that must be in order and your environment has to remain relatively manageable, stable, etc. But your visual experiences of your surroundings as you are standing outside the department do play a key role in explaining how it is that you are able to find your way home right now.

Crossing River: Suppose you are standing at the bank of a stream full of boulders forming a natural path to hop across. You are visually aware of the things around you, the sloping river bank, the boulders, the water, the boggy patches etc. This puts you in a position to hop across the stream. Your visual experiences of your surroundings help explain how you are able to cross the stream.

Picking-up Cat: Suppose you are in my garden looking at my cat Pepperpot lounging in the sun. You are visually aware of Pepperpot sleeping on the bench. This puts you in a position to pick Pepperpot up and take him into the house. Your visual experiences of Pepperpot on the bench help explain how you are able to pick him up and carry him inside.

Each case focusses on attribution of a situation-dependent ability, exercise of which involves success at doing something with respect to a specific aspect of one’s environment. The explanations concern *possession* of these abilities, independently of whether they are actually exercised. The general explanatory pattern is that the potential for success is explained by cognitive access to relevant facts or things in one’s environment (e.g. landmarks, boulders, Pepperpot). It takes for granted that other factors must be in place as well (adequate use of limbs, etc.), but it singles out visual experiences as a key explanatory factor. Thus, according to the explanatory pattern, the cognitive access to the relevant environmental aspects underwriting the potential for success, is provided visually.

It is worth emphasising that the claim contained in the above explanatory pattern is merely that the situation-dependent abilities in question are typically, or perhaps even just sometimes, visually based. They could be had, and perhaps often are had, in virtue of other kinds of access to the relevant environmental aspects (e.g. hearing or touch).¹⁰ I will say more about this in a moment. But while the explanatory pattern I am drawing attention to does not allege

¹⁰See (Millikan, 2000) for interesting discussion of the different ways in which abilities can be had. Moreover, Finding Home also involves a longer process than the other two abilities and so requires continued cognitive access to aspects of the environment along the way.

exclusivity, it does assume that its application sometimes allows us to understand and explain that a subject possesses a certain range of situation-dependent abilities. In fact, the explanatory pattern is a firm part of common sense practice: we employ these sorts of explanations – attributing a range of situation-dependent abilities by appeal to subjects’ visual experiences – in our understanding of one another’s agency and our planning all the time. Note, though, that insofar as one does appeal to visual experience in the above three examples, only visual experiences in the good case can play the required explanatory role. The subject’s visual awareness must be such that it figures in genuinely seeing their environment because we want to explain their possession of an ability which turns on adequate access to certain parts in that environment.

One might object to the last point as follows. In Finding Home, suppose that instead of an experience in the good case, you are hallucinating as you come out of the department building and that your hallucinations happen to be veridical. Such veridical hallucinations would also allow us to explain plausibly how you could get home on this occasion. Similarly, in Crossing River and Picking-up Cat, appeals to veridical hallucinations of the boulders in the river, or of Peppercorn on the bench, would allow for plausible explanations of the respective situation-dependent abilities. So, according to the objection, certain visual experiences in the bad case can serve in these explanations as well.

The objection is misguided and to see why is helpful to understand (P2) better. I will focus on Finding Home, though the response generalizes. It is true that visual experiences in the bad case, specifically, veridical hallucinations, can explain how it is possible for you to find your way home on that occasion. Suppose, as you exit the building, you are unaware that you are hallucinating (scenario Oblivious). Since your hallucinations are veridical, you would get home by relying on this visual information. You would be lucky, it so happens that the environment actually is the way your hallucinations present it to be. Suppose, as you exit the building, you do know that you are hallucinating and veridically so (scenario Clued-up). Then it is also possible for you to find your way home and not just by luck. Since you know that you are hallucinating veridically, you can exploit the information visually presented to get home. Thus, in both Oblivious and Clued-up, it is possible that you find your way home and your veridical hallucinations seem to play a role in this. But, as I will argue now, neither scenario can make use of the explanatory pattern to which I am drawing attention.

In Oblivious, we are not explaining possession of an ability at all, but merely the possibility that I perform a certain action. It is at this point that my particular choice of explanatory target – situation-dependent abilities rather than perceptually-based behaviour or action – makes the important difference.¹¹ Contrast the sense in which it is possible for you to Φ in a given situation S, with the sense in which you have it in your power to Φ in S. It being in your power to Φ in S is a modal aspect of you and having it requires more than that Φ -ing is something the circumstances that you find yourself in allow (Inwagen, 1983, p. 9). It is possible that you pick the winning ticket from an opaque urn with thousands of tickets (you might be lucky) but it is not *in your power* to pick the winning ticket in this situation. You cannot choose to do so, for instance. As Millikan says ‘what one does successfully only by pure chance is not something

¹¹This choice of explanatory target – i.e. my focus on situation-dependent abilities and *not* on actions – distinguishes my approach from the older discussions of the role of object-dependent or externalist mental content in the psychological explanation of action (see, e.g. Segal (1989) and Noonan (1993); also Drayson 2018). There seems to be a certain argumentative stalemate in many of these debates, due to the fact that we can causally explain relevant behaviour and action by appeal to either object-dependent mental content, or a complex of an internal mental content and the obtaining of environmental conditions. As I argue here, situation-dependent abilities present a significantly different explanatory challenge.

one knows how to do. (...) Abilities [have] necessary involvement in the purposive and non-accidental order.’ (Millikan, 2000, p. 58).

The point is that mere performance of Φ -ing is not sufficient for possessing the ability to Φ . On wandering through a foreign city, where all the signs are in a foreign language I do not understand, I might happen to make my way from my hotel to the museum of modern art. But I do not thereby have the ability to find my way from the hotel to the museum. I might very well not be able to get there the next morning again, since it was quite by chance that I ended up where I did today. For it not to be a chance or one-off fluke performance, it must be the case that the (potential) performance is robust in some way: that I could have gotten from the hotel to the museum again in relevantly similar conditions. More generally, for me to exercise an ability, the successful performance constituting such an exercise must be repeatable in relevantly similar circumstances.

The point needs elaboration, for clearly, one’s Φ -ing entails that one can Φ and no-one will deny that. Austin, for instance, when discussing his famed golfing example, says that according to ordinary English merely doing something entails that one has the ability to do it (Austin, 1956). But in fact, something like the above condition on what counts as possessing an ability in the relevant sense, is standardly recognised in literatures where discussion of abilities plays a central role. Correspondingly, any notion of ability, or sense of ‘can’, that applies to a one-off performance is held to be significantly different. In this vein, Honoré talks about ‘the use of ‘can’ (particular) in relation to particular actions’ (Honoré, 1964, p. 464). Similarly, Mele introduces the distinction between simple abilities as opposed to intentional abilities. Simple abilities cover the sense in which one can Φ in virtue of a one-off Φ -ing (Mele’s own example is Ann’s being able to roll a six with a die in a chance game because she did so). Intentional abilities are abilities that manifest a sense of power or control that the agent has in possessing a given ability (Mele, 2003, pp. 447-8). It is the latter sense of ability that is at issue here and that admits of the distinction between situation-dependent and general abilities and that are subject to the condition of repeatability of performance. Thus, when clarifying her notion of local (i.e. my situation-dependent) ability, Whittle insists: ‘It isn’t enough that I just so happen to jump over a five-foot bar on one occasion. In addition, a reliable connection between being in a particular set of circumstances and the outcome is required’ (Whittle, 2010, p. 4).

The target of explanation in Oblivious is the possibility for performance of an action – finding your way home on that occasion. When you stand outside your department and hallucinate veridically while being unaware of this, it is easy to understand how these experiences, together with the fortuitously matching circumstances, make it possible that you get home on that occasion. But these experiences do not make it the case that you have it in your power to get home because it would be merely a happy accident if you did. There is no reliable connection between having a hallucination with a given content and performance, even if that content happens to be veridical on this occasion. Were you to be in relevantly similar circumstances – having an hallucinatory experience with the same content as the veridical hallucination you are having now – it is by no means clear that you could find your way home.¹² Having it in your power to get home – having the ability to get home, on the other hand, requires more than the circumstances conspiring in this manner. I will say more about this in the next section.

¹² Could the relevantly similar circumstances be stipulated to be those in which you have a veridical hallucination of your surroundings? Veridicality is, in the case of hallucinations, a matter of luck. This would amount to stipulating the relevantly similar situations as ones in which you are always lucky in that the content of your hallucinatory experiences correspond to the way your environment is. Any repeatability in successful performance based on such luck obtaining does not constitute a genuinely reliable connection between your visual experience and your performance: they are not anti-luck conditions.

In Clued-up, you are in a position to find your way home, and, that is a genuine power you have rather than merely the circumstances conspiring. Being in the hallucinatory state, together with knowing that you are veridically hallucinating, does provide for a reliable connection to the environment, making the performance robust. In relevantly similar circumstances, where you have an experience with the same content, and knowledge that this experience, though hallucinatory, does present your environment accurately, you will also be successful in finding your way home. However, attributing to you the ability to find your way home in Clued-up is based primarily – or at least to a large extent – on your higher-order knowledge about your visual experiences. What makes the connection between your particular circumstances and the outcome – the successful performance of finding your way home – reliable is that you could do so again in relevantly similar circumstances: circumstances in which you have the hallucinatory experience and knowledge that it is veridical. While the visual experience you are having (the veridical hallucination) does explain the possibility of you finding your way home on this occasion, the robustness of this potential for performance needs crucial appeal to your the higher-order knowledge about the relationship between your experience and the environment. It is the latter that provides the explanatory grounds for you having the ability to find your way home, *as opposed to* the fact that the possibility of finding your way home is something that is open to you in your specific situation.

Thus, in neither of the scenarios, are we explaining possession of the target ability by appeal to visual experiences in line with the explanatory pattern I set out above. In Oblivious, the target of explanation, the *explanandum*, is different: we are not explaining possession of an ability in the relevant sense. In Clued-up, the explaining factor, the *explanans*, is different: we are not explaining the relevant ability mainly by appeal to visual experiences, but by crucial appeal to higher-order knowledge about how the experience relates to the environment. The latter again brings out that abilities can be had in different ways. The explanatory pattern to which I am drawing attention in (P2) merely requires that certain situation-dependent abilities are sometimes had via normally-functioning vision, where this is consistent with them being had differently at other times. But when the pattern is at work – when we are appealing to visual experiences as the main factor to explain the possession of a given situation-dependent ability, only those in the good case can do the job.

3.3 The explanatory demand for world-involvingness (P3)

Turning now to (P3), in this section I argue that for visual experiences in the good case to do the explanatory job we said they do in the previous section, they are best conceived of as world-involving. According to the distinction between Φ -ings which are one-off fluke performances and Φ -ings which are exercises of the ability to Φ , the latter but not the former involve performances that satisfy some anti-luck or robustness condition. Such performances must be repeatable by the subject in sufficiently similar situations.

Recall, abilities are individuated in terms of an action-type Φ and a range of situations where the performance of Φ is available to the subject. In the case of situation-dependent abilities, the individuating range of situations is such that it includes

- (i) a particular situation with respect to which the agent's situation-dependent ability is fixed (the 'base situation'),
- (ii) more situations than the base situation,
- (iii) those situations which are relevantly similar to the base situation.

When we explain a subject's possession of a situation-dependent ability by attributing a visual experience in the good case to her, that experience must provide reason to pick out not only the

action-type that individuates the ability, but also to fix the ability's modal profile, such that the individuating range of situations includes (i), (ii), and (iii).

The base situation is identified as the one in which the agent has a given visual experience. Which situations count as being relevantly similar to it is also fixed by reference to the agent's experience: it is all those situations which differ minimally from the base situation, namely, from the situation in which the agent has the visual experience in question. That is, relevantly similar situations are minimally those in which the agent has a visual experience of the same type. What must be guaranteed is that these are indeed situations in which the action in question is available to the agent.

World-involving visual experiences can guarantee this. The base situation is identified in terms of the agent's having a given world-involving visual experience E_{w-i} . The base situation is thus one in which they have the right cognitive access to their environment for them to successfully perform Φ , where success in Φ -ing depends on such access. The range of relevantly similar situations, in which Φ is also available to the agent, is comprised of those situations in which the agent has experiences of type E_{w-i} (and nothing else relevant being amiss with the agent, i.e. their limbs are working, other cognitive faculties are in order, etc.). Since E_{w-i} is world-involving, that means that relevantly similar situations are those in which the required cognitive access obtains and thus things are as they are visually presented in relevant respects. If you are in a situation relevantly similar to the base one, where the latter is identified in terms of having a world-involving visual experience presenting Pepperpot as sitting on the bench, then the right cognitive access to the environment is guaranteed, such that it will put you in a position to pick him up. That is, relevant similarity requires that the situation is one in which you are having a world-involving visual experience presenting Pepperpot sitting on the bench over there. Such a situation is one in which you could easily have succeeded to pick him up if you tried.

Possession of the situation-dependent ability to Φ is thus explained by appeal to having a certain type of visual experience E_{w-i} : by having a certain kind of access to relevant aspects of one's environment provided by the visual experience. Importantly, this visual access explains how the potential for the performance of Φ is robust, how you could easily succeed to Φ in relevantly similar situations, *viz.* situations in which you have such cognitive access. The visual experience's being world-involving thus guarantees robustness.

So, conceiving of visual experience in the good case as world-involving provides a simple and elegant solution to satisfy the robustness requirement. That fact does not yet suffice for a defence of (P3), of course. Other conceptions of visual experience in the good case might also provide a decent solution. Indeed, given that the good case is characterized in part by the right kind of perceptual contact with relevant aspects of the environment, it may seem that *any* conception of visual experience in the good case satisfies the robustness requirement. A complex state $E+C_{ext}$ consisting of a veridical experiential mental component together with the right kind of external hook-up to the world guarantees robustness, too. Accordingly, one might object that the fact that world-involving visual experience guarantees robustness yields no support for (P3), other than perhaps by considerations of elegance.

Elegance considerations notwithstanding, the objection is wrong. There is a significant difference in how each conception of visual experience in the good case can guarantee robustness and this difference shows that there is, after all, an explanatory demand for world-involving experience providing support for (P3). If visual experience in the good case is world-involving, then conscious visual experience of the world, E_{w-i} does the explanatory heavy lifting with respect to both, fixing the relevant action-type individuating of the ability in question, as well as robustness. Visual experience itself – conscious visual awareness or presentation – explains possession of the situation-dependent ability under consideration

because the right perceptual hook-up to the world is a constituent part of it. This account makes best sense of the explanatory pattern expressed by (P2), according to which we sometimes explain possession of situation-dependent abilities by appeal to visual experiences.

By contrast, if visual experience in the good case is not world-involving, then the experiential element, the conscious visual presentation of the world, loses considerable force as an explanatory factor. On the $E+C_{\text{ext}}$ conception of visual experience in the good case, E can be had in a bad case, too. Just like in the case of Clued-up above, conscious visual experience E *per se*, does fix the action-type relevant to the ability, explaining how performance of this action is possible for the subject in situations where they have E . But, the range of situations fixed by E in this way do not amount to relevantly similar situations in which such performance is robust: E on its own will not fix the relevant similarity class of situations in which one could easily have succeeded to Φ . Instead, the explanatory burden of fixing the latter range is carried first and foremost by the additional external conditions C_{ext} which characterize cases of genuine seeing. In principle, these might obtain in conjunction with a variety of visual presentational states the agent is in, so the overall explanation needs to appeal to the complex $E+C_{\text{ext}}$. Although this kind of complex state can guarantee robustness, then, it diminishes the explanatory efficacy of conscious visual awareness itself considerably.

Our starting point is the ordinary explanatory pattern referred to in (P2), namely that there are occasions where we explain our possession of certain situation-dependent abilities by appeal to visual experience (accepting that other conditions, e.g. working limbs, etc. must be in place). A conception of world-involving visual experience does the best job of accounting for this explanatory role we accord to visual experience in ordinary practice vis-à-vis such abilities. This is because on this conception, conscious visual experience can be understood both to provide for the selection of the relevant action-type and to underwrite the robustness of potential performance of such action, necessary to turn mere performance into an exercise of an ability. World-involving visual experience yields the most appropriately straightforward *explanans* and in our explanatory pattern. The other conception of visual experience in the good case, $E+C_{\text{ext}}$, does not give conscious visual experience E the same comprehensive role, needing to appeal to external conditions C_{ext} to explain robustness. The explanation via world-involving visual experiences therefore more naturally reflects common sense explanatory practice and our understanding of our agency as often grounded in visually conscious access to our environment.

The explanatory pattern expressed by (P2) is reasonable: we ordinarily often do explain possession of such abilities by appeal to conscious visual experience. Furthermore, this common sense explanatory practice works: understanding our situational agency is at the core of our day-to-day planning and decision-making. It serves us well in understanding the kinds of agents we are *in situ* and guiding expectations and predictions about what we can and cannot do fairly well.

4. World-involving experiences as a mental kind

The methodological concern raised in the introduction focusses on the gap between establishing that there is world-involving experience and establishing that such experience forms a significant mental kind. Any moderate experiential pluralism of substance must close this gap by spelling out what the mental kindhood of world-involving experience consists in. The latter requires not only an argument for a good-case-only-property of experience, but also reasons for treating that property as constituting a mental kind.

Proponents of different versions of experiential pluralism often rely on terms such as ‘nature’, ‘kind’, and ‘fundamental’, as in ‘fundamental level of classification’ (McDowell, 2010, p. 252),

‘(ontologically fundamental) kind’ (Snowdon, 2005, p. 133), to state and clarify their key claims. However, they typically do not elucidate these terms much further, thus leaving it unclear what such appeals to fundamentality, nature, basic kind, etc. amount to in this context. Consider Martin’s suggestion that the fundamental kind at issue concerns the essence of the experience. The essence of a given entity determines what kind of entity it is, according to

“some privileged classifications of individuals, both concrete objects and events, and that our talk of what is essential to a given individual tracks our understanding of the kinds of thing it is. ... [E]ntities (both objects and events) can be classified by species and genus; for all such entities there is a most specific answer to the question, ‘What is it?’ In relation to the mental, and to perception in particular, I will assume that for mental episodes or states there is a unique answer to this question which gives its most specific kind; it tells us what essentially the event or episode is.” (2006, 361)

Suggestions along these lines seem to presuppose that we have some prior and shared common sense understanding of what constitutes the fundamental level, or the most specific kinds, etc. at issue.¹³ But fundamentality is an explanatory notion and claims about basic or most specific fundamental kinds must be situated within an explanatory context which in turn is tied to a given level of description of the world. This is not to say that everything goes, of course. Without wading into the general controversies about realism about natural kinds – specifically, whether such kinds are interest-relative, or conventional in some sense that undermines their objective reality – I think it is reasonable to assume that not all classifications for any given explanatory purpose are equal with respect to the world that they describe. We do want the fundamental kinds to carve the world’s joints in some significant manner, even if they are interest-relative to some degree. This is true of mental kinds as well. Both of these aspects of mental kindhood – their individuating connection to explanatory projects and levels of description, as well as their reality – need to be articulated adequately in order to address the methodological concern about the gap between an argument for a good-case-only-property of experience and reasons for treating that property as constituting a mental kind.

Claims about fundamental levels of classification and similar typically do not do so in my view. The assumption of a shared common sense understanding of the relevant notion of a fundamental level at which to locate the mental kinds in question, cannot be taken for granted. There is no firm grasp of a level of reality available, even in common sense, that is both independent of explanatory projects and uncontroversial enough to yield something like Martin’s unique, most specific answers in the case of experience. So, in the background of such an assumption must be something like the thought that there is an explanatory context, a level of description, which forms the shared understanding of common sense. One likely candidate for this is ‘the manifest image’ of mentality – the idea that there is a distinctive way in which we encounter our minds, including perceptual experience, from the person-level, subjective perspective. The manifest image of mentality is meant to provide a level of description that is at the heart of our basic mental concepts, and explanatory projects hitched to it are therefore held to be revealing of the nature of mentality itself.

But this background thought about the manifest image of mentality which putatively moors the shared understanding of perceptual experience providing for unique answers to the ‘What is it?’ question, is non-trivial. What is the source of claims reflecting our shared understanding and the manifest image of perceptual experience? It is fair to say that in the debate about the

¹³ Martin and McDowell, of course hold very different experiential pluralist views (see footnote 3) and this section is not meant to treat these views as the same. But given that such views typically use talk of the fundamental level, etc. this variety illustrates the need for more detail on the psychological kind invoked by such talk.

nature of perceptual experience, many of them are grounded in the deliverance of introspective reflection on the phenomenal character of experience. In addition to introspection-based claims about phenomenal character, there are considerations about the epistemic role of perceptual experience in sustaining various other manifest mental capacities (e.g. demonstrative thought, empirical propositional knowledge), as well as theoretical reflection on the nature and purpose of the capacity of perception itself. It is not at all clear to me that these sources (singly or jointly) offer the bases for the required shared understanding of perceptual experience. Insofar as the manifest image is meant to reflect *common sense* understanding, these sources are not obviously part of ordinary thinking about the mind, if by that we mean pre-theoretic, everyday thought and talk. Considerations about the epistemic role of perceptual experience in grounding perceptual demonstratives, or about the nature and purpose of the capacity of perception are theoretically motivated considerations about the mental phenomenon in question. One has to do some work to get the ordinary person to understand the questions being asked in the first place. Even introspective reflection on the phenomenal character of experience does not genuinely yield a ‘naïve’ take on experience, where that requires isolating (say) visual experience, from any background expectations, habits, other beliefs, and non-visual experiences in the vicinity. At the very least, *that* kind of naivety has to be practiced some before being available to an introspecting subject. But once we accept that the shared understanding requires some theoretical sophistication and, importantly, shared explanatory interests, none of the claims expressive of it offer particularly neutral dialectical starting points that opponents will easily agree upon.

Furthermore, the above suggestions typically have little to say about how exactly their explanatory projects relate to other explanatory projects concerned with the mind and perception in particular. According to Martin, the essence of experience as located at the fundamental manifest level is meant to be consistent with there being ‘a more fundamental level of reality’ on which the mental supervenes in some way (see his 2006: 360, footnote 9). But it is difficult to see how exactly the ‘manifest kinds’ at issue in their debates relate to, especially, scientific kinds in perceptual psychology.¹⁴ In the wake of a thriving science of perception, philosophers of perception are (rightly, in my view) expected to ensure that their views are appropriately empirically-informed. More generally, I think that the question about the reality of mental kinds demands some minimal integration of the picture of the mind emerging from philosophical investigation of the manifest image, with the picture of the mind emerging from scientific investigation in psychology. In some quarters, philosophical concerns and views about perceptual experience are held to be irrelevant and unmotivated unless they are contained in our best current scientific approaches to perception (Burge 2010, Rescorla, 2015). This strikes me as an unjustified restriction - based on insisting on the primacy and dominance of the explanatory project of contemporary vision science (e.g. Fish (this volume)). However, to insist in return on the availability of entirely different and autonomous explanatory projects is equally unsatisfactory. It leaves the question about the reality of mental kinds – as distinct from psychological kinds – hanging, undercutting any sense that the mental kinds of one such project and the psychological kinds of the other in the end contribute to an integrated picture of the mind.

My strategy – pursued in the argument from abilities - helps to refocus the issue away from insisting on the primacy or autonomy of a specific explanatory project, to the choice of appropriate explanatory targets for attribution of mental or psychological kinds more generally. According to my strategy, the explanatory targets are ordinary situation-dependent abilities.

¹⁴ The label ‘manifest kind’ is from Johnston (1997), discussed in Phillips (2018). For a related debate about the question of whether there is unconscious or subliminal perception, see, e.g. Block and Phillips (2016), Peters et al. (2017), Phillips (2018).

These are relatively uncontroversial candidates for counting as natural properties of the mind, from the scientific, philosophical, and common sense perspectives. They offer a genuinely neutral starting point for explanation and allow for different explanatory projects concerning them.

In addition, in focussing on ordinary situation-dependent abilities as explanatory targets, my strategy sticks closely to an innocuous rule of thumb concerning the reality of natural kinds in psychology: that mental states or episodes count as real psychological properties if they are causally trackable. Moreover, if the causal efficacy of such properties can be shown to be systematic and distinctive in certain ways, there is reason to hold that these properties form a significant psychological kind.¹⁵ The way in which such causal efficacy is often established is by showing that these mental states or episodes have an important – perhaps even essential or ineliminable – role in the causal explanation of behaviour (Drayson 2018).

Although my argument follows the basic outline of this approach, it does not fit the mould of a causal explanation of behaviour exactly. As we have seen, the argument turns on the explanatory role that world-involving experiences play and the explanatory targets in question are ordinary situation-dependent abilities. The latter, of course, are not instances of behaviour, if by that we understand performances of actions. But, as we also saw, such abilities are intimately connected with performances of action. Abilities encompass certain behavioural possibilities that characterize an agent. They determine an important part of their agency in setting a repertoire of actions that an agent has available to them in certain kinds of circumstances. So, while the explanatory targets in my argument are not instances of behaviour directly, they are properties that manifest in human behaviour.

Correlatively, the explanations at the heart of my argument are themselves not causal. They are constitutive, indicating that having world-involving experience endows the subject with certain other psychological properties. But the abilities which are my explanatory target form the basis for prediction of behaviour in ordinary planning and decision-making, where we need to know what an agent can and cannot do in specific situations. Abilities themselves thus become causally trackable in terms of their exercises being causally trackable. This is reflected in the fact that common sense explanatory practice involving attribution of situation-dependent abilities is empirically informed. Attributions of these abilities have been honed by observation of success and failure of performance of relevant actions. More importantly, abilities are open to empirical study. That means not only that abilities are scientifically recognised phenomena, but that common sense explanatory practice involving them is open to revision from science.¹⁶

In light of this, we can see that my argument follows the rule of thumb for recognition of a psychological kind mentioned above in spirit, if not in practice. It follows it enough to provide a firm first step towards recognizing a certain type of mental episode as a genuine psychological kind. World-involving visual experiences explain possession of some situation-dependent abilities. In turn, these abilities are a central ingredient in behaviour-related planning and decision-making and they are causally trackable via their causally efficacious behavioural manifestations. In light of their explanatory and predictive role and the fact that they are subject

¹⁵ This second application of the rule – underwriting natural psychological kindhood – is clearly less innocuous than the first, underwriting natural psychological properties. It can be filled out in different ways, depending on one's views about what constitutes kindhood more generally.

¹⁶ Much of psychology involves investigating which abilities, skills, capacities and expertise we have and how to explain them, and, moreover, how they contrast with our ordinary conceptions of them. Two give just two examples: studies on unethical amnesia and its effect on one's ability to regulate shame-induced social dysfunction (Kappes & Crockett, 2016) and studies on our ability to provide accurate eyewitness testimony (Wells & Loftus, 2003).

to empirical investigation and revision, situation-dependent abilities are genuine psychological properties of ours. In that case, mental states or episodes by appeal to which we attribute and explain possessing them, have a claim to be taken seriously as real psychological properties if not kinds as well. World-involving visual experience does have such an explanatory role and it is therefore a fitting candidate for counting as a psychological kind.

A detailed account of what constitutes natural kindhood, and especially natural psychological kindhood, is beyond the scope of this paper. The positive strategy I offer is meant as a first step, charting a plausible route for making a detailed case that the type of experience in question constitutes a natural mental kind in its own right, where such kindhood is not entirely disconnected from related psychological kinds recognised by empirical psychology. There may well be other good strategies for doing so (see, e.g. Phillips 2018 for a different proposal). The particular one I have advanced here lays the groundwork for a promising approach, if the wider aim is to facilitate engagement between scientific and philosophical theorising about perception.

References

- Austin, J. L., 1956. Ifs and Cans. *Proceedings of the British Academy*, Volume 42, pp. 109-132.
- Berofsky, B., 2005. Ifs, Cans, and Free Will: The Issues. In: R. Kane, ed. *The Oxford Handbook of Free Will*. Oxford: Oxford University Press, pp. 181-201.
- Bird, A., 1998. Dispositions and Antidotes. *The Philosophical Quarterly*, Volume 48, pp. 227-34.
- Block, N. and Phillips I., 2016. Unconscious seeing – a debate. In: B. Nanay, ed. *Current Controversies in Philosophy of Perception*. Routledge.
- Burge, T., 2009. Five Theses on De Re States and Attitudes. In: J. Almog & P. Leonardi, eds. *The Philosophy of David Kaplan*. Oxford: Oxford University Press, pp. 246-310.
- Burge, T., 2010. The Objectivity of Perception. OUP.
- Burge, T., 2011. Disjunctivism Again. *Philosophical Explorations*, 14(1), p. 43–80.
- Byrne, A. & Logue, H., 2008. Either/Or. In: *Disjunctivism: Perception, Action and Knowledge*. Oxford University Press, pp. 57-94.
- Campbell, J., 2002. *Reference and Consciousness*. Oxford: Oxford University Press.
- Drayson, Z., 2018. Extended minds and prime mental conditions: probing the parallels. In: J.A. Carter, A. Clark, J. Kallestrup, S.O. Palermos, and D. Pritchard, eds. *Extended Epistemology*. OUP.
- Haddock, A. & MacPherson, F., 2008. Introduction: Varieties of Disjunctivism. In: A. Haddock & F. MacPherson, eds. *Disjunctivism: Perception, Action and Knowledge*. Oxford: Oxford University Press, pp. 1-24.
- Honoré, A. M., 1964. Can and Can't. *Mind*, 73(292), pp. 463-479.
- Inwagen, P. v., 1983. *An Essay on Free Will*. Oxford: Clarendon Press.
- Johnston, M., 1997. Manifest Kinds. In *Journal of Philosophy* 94 (11):564-583.

- Kalderon, M., 2018. Experiential Pluralism and the Power of Perception. In John Collins and Tamara Dobler (eds.), *The Philosophy of Charles Travis: Language, Thought, and Perception*. Oxford: Oxford University Press, pp. 222-236.
- Kappes, A. & Crockett, M., 2016. The benefits and costs of a rose-colored hindsight. *Trends in Cognitive Sciences*.
- Kratzer, A., 1977. What 'Must' and 'Can' Must and Can Mean. *Linguistics and Philosophy*, Volume 1, p. 337–355.
- Maier, J., 2015. The agentic modalities. *Philosophy and Phenomenological Research*.
- Maier, J., 2018a. *Abilities*. [Online]
Available at: <https://plato.stanford.edu/archives/spr2018/entries/abilities/>
- Maier, J., 2018b. Ability, modality, and genericity. *Philosophical Studies*, pp. 411-428.
- Martin, M., 1997. The Reality of Appearances. In: M. Sainsbury, ed. *Thought and Ontology*. Milano, Italy: Franco Angeli, pp. 81-106.
- Martin, M., 2002. The transparency of experience. *Mind and Language*, 17:376-425.
- Martin, M., 2004. The Limits of Self-Awareness. *Philosophical Studies*, Volume 120, p. 37–89.
- Martin, M., 2006. On Being Alienated. In: T. Gendler & J. Hawthorne, eds. *Perceptual Experience*. Oxford: Oxford University Press, pp. 354-410.
- McDowell, J., 2010. Tyler Burge on Disjunctivism. *Philosophical Explorations*, 13 (3):243-255.
- McDowell, J., 2011. *Perception as a Capacity for Knowledge*. s.l.:Marquette University Press.
- McDowell, J., 2013. Tyler Burge on Disjunctivism (II). *Philosophical Explorations*, 16 (3): 259-279.
- Mele, A. R., 2003. Agents' Abilities. *Nous*, 37(3), pp. 447-470.
- Millikan, R., 2000. *On Clear and Confused Ideas*. Cambridge: Cambridge University Press.
- Millikan, R. G., 2000. *On Clear and Confused Ideas: an Essay on Substance Concepts*. s.l.:Cambridge University Press.
- Peters, M., Kentridge, R., Phillips, I. and Block, N., 2017. Does unconscious perception really exist? Continuing the ASSC20 debate. *Neuroscience of Consciousness*, 3 (1)
- Phillips, I., 2018. Unconscious perception reconsidered. *Analytic Philosophy*, 69 (4).
- Rescorla, M., 2015. Bayesian Perceptual Psychology. In: M. Matthen, ed. *Oxford Handbook of Perception*. Oxford: Oxford University Press.
- Siegel, S., Winter 2016. *The Contents of Perception*. [Online]
Available at: <https://plato.stanford.edu/archives/win2016/entries/perception-contents/>
- Snowdon, P., 2005. The Formulation of Disjunctivism: a Response to Fish. *Proceedings of the Aristotelian Society*, Volume 105, pp. 129-141.
- Snowdon, P., 2008. Hinton and the origins of disjunctivism . In: A. Macpherson & F. MacPherson, eds. *Disjunctivism: Perception, Action, Knowledge*. Oxford: Oxford University Press, pp. 35-56.
- Soteriou, M., 2016. *Disjunctivism*. Routledge.

Vetter, B., 2013. 'Can' Without Possible Worlds: Semantics for Anti-Humeans. *Philosophers' Imprint*, 13(16).

Wells, G. L. & Loftus, E. F., 2003. Eyewitness Memory for People and Events. In: A. Goldstein, ed. *In Handbook of Psychology: Forensic Psychology*. Wiley.

Whittle, A., 2010. Dispositional Abilities. *Philosophers' Imprint*, 10(12), pp. 1-23.