

What is the concept of perceptual experience?

Questions about the metaphysical nature of perceptual experiences are of obvious interest to theorists of perception. And if you want to know about the metaphysical nature of something, then one useful avenue of approach is to first understand the concept of that thing, so that you can see what fits the bill. That is one good reason to ask the titular question of this paper.

In response to this titular question, one suggestion that was made half a century ago is that the concept of a perceptual experience is roughly the concept of something that seems to be a veridical perception.¹ More recently, this suggestion has been refined into:

The *indiscriminability view*: The concept of a perceptual experience is just the concept of something that is impersonally introspectively indiscriminable from a veridical perception.²

The indiscriminability view is said to have the great virtue of being an *epistemically modest* account of the concept of perceptual experience (§1). However, it faces a number of serious problems. One familiar problem is that the view is extensionally inadequate. In addition, I will raise three new problems for the indiscriminability view: it must appeal to a special-purpose notion of impersonal introspective indiscriminability, it is semantically immodest, and it is semantically incongruous (§2).

¹ Something like this idea appears in Hinton (1967).

² See Martin (2004, pp. 75-6) and (2006, §5).

In light of these problems, my goal in this paper is to argue for an alternative view – which I will call the *sampling view* – of the concept of perceptual experience. The sampling view is broadly Kripkean in spirit. The core idea is that we acquire the concept of perceptual experience in much the same way that we acquire the concepts of water, gold, or cats: by becoming familiar with samples of the natural kind in question. My suggestion will be that the sampling view avoids the four problems that threaten the indiscriminability view while being just as epistemically modest (§3). I will then propose a refinement of the sampling view in order to block an important objection (§4) before wrapping up (§5).

1. The indiscriminability view

The central argument for the indiscriminability view begins with an observation about one common way of picking up the concept of perceptual experience.³

Suppose that I am currently seeing a salsa dance performance. And suppose that you ask me to reflect more generally on situations that are introspectively indiscriminable for me from my actual perceptual situation. In some of those situations, I am really seeing a salsa dance performance. In others of those situations, I am not; I am *merely* in a situation that is introspectively indiscriminable for me from my situation when I am really seeing a salsa dance performance. Thinking about all of these situations together seems to be all that it takes for me to pick up the concept of perceptual experience, and thinking about just the latter situations seems to be all that it takes for me to pick up the concept of a hallucination.

From this observation about how we commonly pick up the concept of perceptual experience, some naïve realists would draw a defeasible conclusion about what the concept of

³ The argument described below is due to Martin (2004, pp. 47-52).

perceptual experience might be. These naïve realists start by outlining two alternative theories. According to the first view, which these naïve realists want to reject, we get our concept of a perceptual experience of a salsa dance performance *first* by introspecting on a veridical perception of a salsa dance performance and *then* by detecting some positive property *P* of the veridical perception that can be instantiated even if we are not veridically perceiving a salsa dance performance. Thus we have the following schematic proposal:

The *positive view*: The concept of a perceptual experience is the concept of something that has property *P*.

This can be contrasted with a second view, which these naïve realists favor. According to this second view, thinking about perceptual experiences is nothing more than thinking about situations that are introspectively indiscriminable from a veridical perception.

Of course, not just any kind of introspective indiscriminability will do in this characterization. After all, my powers of introspection might be compromised because I am tired, drugged, or insane. Moreover, it is possible for creatures such as dogs to have hallucinations, even though dogs are not cognitively sophisticated enough to discriminate any experiences from one another. But we might try to solve these problems by appealing to an *impersonal* form of introspective discrimination – a form of discrimination that has to do with what can be discriminated via introspection in general, rather than by any particular subject.⁴

Thus we have:

⁴ For a detailed characterization of the notion of impersonal introspective indiscriminability, see Martin (2004, pp. 74-81) and (2006, pp. 379-96).

The *indiscriminability view*: The concept of a perceptual experience is just the concept of something that is impersonally introspectively indiscriminable from a genuine perception.⁵

Why do these naïve realists favor the indiscriminability view over the positive view? Because they think that the indiscriminability view is the more *epistemically modest* view. In particular, these naïve realists say, surely the indiscriminability view is extensionally adequate – that is, surely something is a perceptual experience just in case it is impersonally introspectively indiscriminable from a veridical perception. So suppose, for the sake of reductio, that the positive view is correct – suppose that the concept of a perceptual experience is the concept of anything that has some positive, introspectively detectible property *P*. Then introspection must be a faculty that is *perfect* at detecting the presence or absence of *P*. After all, if introspection were not perfect in this way, then it would not be true that something is a perceptual experience – that something has property *P* – just in case it is impersonally introspectively indiscriminable from a veridical perception. This yields an immodest view of the epistemic powers of our introspective faculties.⁶

By contrast, the indiscriminability view does not posit any immodest epistemic powers, and indeed it need not treat introspection as a faculty of detection at all. At the same time, the indiscriminability view agrees with the positive view about the extension of the concept of perceptual experience. So – these naïve realists provisionally conclude – it is the better view.

⁵ Again, see Martin (2004, pp. 75-6) and (2006, §5). Martin takes inspiration from Hinton (1967).

⁶ See Martin (2004, pp. 49-51).

2. Four problems for the indiscriminability view

I agree that its epistemic modesty is an attractive feature of the indiscriminability view. But I think that the view faces four important problems, the first of which is familiar and the remainder of which will be first introduced here.

The familiar problem of extensional inadequacy. An immediate problem for the indiscriminability view is that it is subtly extensionally inadequate. For suppose that I am hallucinating, and suppose in particular that it seems to me that I am seeing an impossible Escher staircase. My hallucination is obviously a perceptual experience. But there is no veridical perception that it is impersonally introspectively indiscriminable from. After all, Escher staircases are impossible, so it is impossible to have a veridical perception of one.⁷

In response to this problem, indiscriminability theorists have suggested that any hallucination of the impossible can be exhaustively decomposed into various parts, each of which is impersonally introspectively indiscriminable from some veridical perception.⁸ But this suggestion does not seem to handle all of the cases. For example, it is possible to experience *supersaturated redness*, a shade of redness that seems to be more saturated than any shade of red could possibly be. In the most familiar actual cases such an experience would just be an illusion, but surely it is possible for a subject to have an experience of supersaturated redness that is a hallucination. And the problem is that such a hallucination does not seem to be exhaustively decomposable into parts that are each impersonally introspectively indiscriminable from some

⁷ This worry is due to Siegel (2004, p. 92).

⁸ This is the response of Martin (2004, pp. 80-1).

veridical perception. That is, hallucinating the impossible property of supersaturated redness does not seem to be a matter of hallucinating two or more possible properties.⁹

Perhaps the indiscriminability theorist will disagree with me on this point.¹⁰ Even so, I think that she faces a problem here. For she is offering a characterization of our shared *concept* of perceptual experience; she is saying that our concept of perceptual experience is the concept of something that is exhaustively decomposable into parts that are each impersonally introspectively indiscriminable from some veridical perception. So the indiscriminability theorist makes the following prediction: as long as I believe, perhaps falsely, that an experience of supersaturated redness is not decomposable into such parts, I should also believe, perhaps falsely, that such an experience cannot be perceptual. I hereby report that this is not what happens in my own case.

No doubt, there are sophisticated moves that the indiscriminability theorist could make here. Still, I think that it would be an advantage to have a theory that could give a smoother treatment of hallucinations of the impossible.

The new problem of special indiscriminability. The indiscriminability theorist appeals to an *impersonal* notion of introspective indiscriminability. How might she explicate this notion?

Well, there is a standard strategy for explaining what it is for two things to be impersonally indiscriminable via some method. The idea is that we first determine what kinds of properties the method generally can tell apart. We then say that what it is for two things to be

⁹ This point might even undercut the argument against the positive view. For the argument against the positive view rested on the premise that something is a perceptual experience iff it is impersonally introspectively indiscriminable from a veridical perception, and this premise (and nearby variants of it) is now coming under pressure.

¹⁰ Interestingly, Martin does mention the example of supersaturated redness (2004, p. 80), even though – as I have just mentioned – his response does not seem to me to handle that example aptly.

impersonally indiscriminable via the method is for those two things to be identical with respect to all of the kinds of properties that the method can tell apart. For example, we might say that two things are impersonally indiscriminable via the method of seeing just in case those two things are identical in all of their visible properties: color, shape, texture, etc.

Notice, however, that indiscriminability theorists tend also to accept the following view:

Negative view of hallucination: Any hallucination has no mental properties except for the property of being exhaustively decomposable into parts that are each impersonally introspectively indiscriminable from some veridical perception (and any properties entailed by that).¹¹

But if we accept this view, then we run into trouble when we try to explain the notion of impersonal introspective indiscriminability in the standard way. Why? Well, we want to explain the notion of impersonal introspective indiscriminability by identifying the mental properties of hallucinations and veridical perceptions that introspection can generally tell apart. But if we accept the negative view of hallucination, then the only mental property that a hallucination has is an impersonal introspective indiscriminability property (and properties that follow from that).

And, of course, we cannot appeal to an impersonal introspective indiscriminability property to

¹¹ Martin expresses a localized version of this principle: “for ... causally matching hallucinations, the only mental properties that such events possess are those of being indiscriminable from veridical perceptions and any properties which follow from their being so indiscriminable” (2004, p. 71). A similar principle appears in Brewer (2011, p. 109). Note that a “causally matching hallucination” is a hallucination that has exactly the same type of proximal cause as some veridical perception. Now, Martin’s restriction to causally matching hallucinations lets us avoid counterexamples from impossible hallucinations, but the restriction also seems to be *ad hoc*. I prefer to work with the principle in the text because it likewise avoids counterexamples from impossible hallucinations while nevertheless being fully general.

explain the notion of impersonal introspective indiscriminability – that explanation would just be circular.

In response to this problem, indiscriminability theorists have offered new accounts of what impersonal indiscriminability amounts to in the case of introspection. One possibility is that for two mental situations to be impersonally introspectively indiscriminable is for them to be alike in certain sub-personal respects, such as neurological or functional respects.¹² Another suggestion is that for two mental situations to be impersonally introspectively indiscriminable is for them to have the same range of effects of a certain type – e.g., the same cognitive effects.¹³

Now, it has been argued that these accounts are inadequate in various respects.¹⁴ But these arguments are complex, and I propose to set them aside so that I can focus on a different problem that is easier to appreciate.

I want to focus on the fact that it is always a cost to add a special-purpose notion to your theory. And that is precisely what the indiscriminability theorist is doing here: she is developing a special notion of impersonal introspective indiscriminability that does not have any obvious application to other cases. (At least, I have never seen any proposals for further applications from indiscriminability theorists.)

¹² This is Martin's suggestion: "What it takes for a creature so to satisfy this condition [of being in a situation that is impersonally introspectively indiscriminable from a veridical perception] may well involve ... the same neurological organization and functioning [as a veridical perceiver] ..." (2006, p. 395).

¹³ This is Fish's suggestion: "My definition of the hallucinatory state captures this insight [about the indiscriminability of hallucinations from veridical perceptions] by requiring a hallucination to produce the same cognitive effects that a veridical perception would have produced" (2009, p. 97).

¹⁴ See Siegel (2008).

That is by no means a devastating objection to the indiscriminability view. Still, it would be an advantage to have a view that did not require us to construct any special-purpose notions of this sort.

The new problem of semantic immodesty. Those who accept the indiscriminability view tend also to accept the following view:

Metaphysical disjunctivism: All that it is to be a hallucination, metaphysically speaking, is to be something that is exhaustively decomposable into parts that are each impersonally introspectively indiscriminable from some veridical perception.¹⁵

This view is obviously closely related to the negative view of hallucinations: if you accept one of these two views, then there is at least some pressure to accept the other. But these two views are logically independent, so I think that it is worth distinguishing them.

Now, there is something notable about the theorist who combines the indiscriminability view with metaphysical disjunctivism. For on this view, our concept of perceptual experience captures the very metaphysical nature of *some* perceptual experiences – namely, of hallucinatory experiences.¹⁶

That should come as a surprise. After all, since Kripke, philosophers have arrived at a strong consensus that *almost none* of our concepts capture the metaphysical natures of *any* of

¹⁵ See Martin (2004, p. 72) and (2006, p. 393). Again, I have fully generalized the principle that Martin accepts for reasons described in fn. 11. A similar principle is endorsed in Brewer (2011, p. 109).

¹⁶ As Martin acknowledges: “For these events [namely, causally matching hallucinations – see fn. 11] alone, there is nothing more to characterize them than what can be said about the concept of perceptual experience in general” (2004, p. 71).

their referents. Think of the concept of water: it is now almost universally agreed that we refer to water via its superficial features. Perhaps we pick out water as the stuff that is transparent and odorless, that fills the rivers and the lakes, that falls from the sky as rain, and so on. Still, there is no sample of water whose *essence* is to be transparent and odorless, to fill the rivers and the lakes, and so on. The same point holds for our concepts of gold, cats, money, food, and so on: none of these concepts reveals the metaphysical natures of any of their referents.

There might be a few exceptions to this rule. Maybe the metaphysical nature of the referent is captured by the concept of a bachelor or by the concept of a circle. But these are exceptional cases. The naïve realist who combines the indiscriminability view with metaphysical disjunctivism says that the concept of perceptual experience is another (partly) exceptional case: this concept reveals the metaphysical nature of some (though not all) of its referents. That amounts to an immodest view of our *semantic* powers. Such semantic immodesty is a significant cost.

The new problem of semantic immodesty. There is a related problem. Again, everyone agrees that there are many concepts that do not reveal the metaphysical natures of any of their referents. And, again, quite a few people also think that there are a handful of concepts – again, like the concept of a bachelor or the concept of a circle – that reveal the metaphysical natures of all of their referents. But, naïve realists aside, I cannot think of anyone who has ever held that there were any concepts that revealed the metaphysical natures of *some but not all* of their referents. That would be a strange view: the idea would be that a concept locks onto its referents in virtue of their having certain features, where some of those referents have these features accidentally and others of those referents have the same features essentially. Yet this strange

view is precisely the one that indiscriminability theorists hold when it comes to the concept of perceptual experience.

Now, there is nothing incoherent about this view. But it is not attractive to say that the concept of perceptual experience is the lone semantic dangler. This is the problem of *semantic incongruence*.

Conclusion. None of the four problems that I have mentioned is decisive. But each one is serious, and taken together they inspire me to look for an alternative view. Ideally, this view would solve all four of the problems that arise for the indiscriminability view while still being epistemically modest.

How might we begin our search?

3. An alternative view

Here is a natural thought. Since the Kripkean revolution in semantics,¹⁷ there is a certain standard way that philosophers have usually thought about natural kind concepts. So let us see what happens if we apply this standard way of thinking to the concept of perceptual experience.

In outline, the now-standard view is that we normally acquire concepts of natural kinds in virtue of encountering some paradigmatic instances of the kind. We see one or two samples of water and think about *that stuff*; that is all that it takes, or at least most of what it takes, to acquire the concept of water. There are, however, two importantly different ways of filling in the details of this view.

On the one hand, some theorists think that our encounter with a sample of water simply *fixes the reference* of our concept of water; we need not have any personal-level understanding,

¹⁷ Due, of course, to Kripke (1972).

even an implicit personal-level understanding, of how facts about causation, natural kindhood, etc. guide our application of the concept of water. The idea might be that H₂O just is the kind of thing that causes our water-beliefs, and in virtue of that our concept of water refers to H₂O. In fact, this was the view of Kripke himself, and I will call it the *reference-fixing view*.

By contrast, other theorists think that to have a concept of water, we do need to have a personal-level understanding, at least implicitly, of how facts about causation, natural kindhood, etc. guide our application of the concept of water. That is, these theorists think that we must understand at the personal level, at least implicitly, that *if* there is a single kind of thing that is transparent and odorless, that fills the rivers and the lakes, that falls from the sky as rain, and so on, *then* water is the thing of that kind. I will call this the *neo-descriptivist view*.¹⁸

For the sake of neutrality, let's not try to decide between the reference-fixing and neo-descriptivist views. Instead, let's just focus on what these theorists agree on: they agree that a natural-kind concept has its referent in virtue of the fact that this referent is what causes the subject to apply the concept in certain privileged conditions. For example, they agree that the concept of water refers to H₂O in virtue of the fact that H₂O is what causes the subject to apply the concept of water in certain privileged conditions – perhaps baptismal conditions, evolutionarily ideal conditions, or conditions of full background information of a certain sort.

What happens if we think of the concept of perceptual experience as an analog of the concept of water? Well, then we get something like this:

The *sampling view*:

¹⁸ Prominent neo-descriptivist works include Lewis (1984), Jackson (1998), and Chalmers (2012).

- (a) The concept of a perceptual experience is the concept of anything of the kind whose paradigms include assorted samples of veridical perceptions, illusions, and hallucinations.
- (b) We can conceive of particular samples of veridical perceptions and illusions because we have actually had such experiences.
- (c) We can conceive of particular samples of hallucinations, even if we have never had any hallucinations, primarily by relying on our knowledge of veridical perceptions.

I will suggest that the sampling view can avoid all of the problems that we have identified so far, and as I make my case I will simultaneously clarify the sampling view.

The problem of epistemic immodesty revisited. Recall the schematic view that some indiscriminability theorists have used as a foil for their own approach:

The *positive view*: The concept of a perceptual experience is the concept of something that has property *P*.

Indiscriminability theorists worry that since any perceptual experience is impersonally introspectively indiscriminable from a veridical perception, the positive theorist must say that the faculty of introspection is perfect at detecting the presence of property *P*. That is an epistemically immodest view, and such epistemic immodesty is costly. Is the sampling view epistemically immodest?

Well, the key issue is precisely how the sampling theorist can explain our ability to conceive of particular samples of hallucination. I see several options here. Perhaps we think of *some* hallucinations by introspectively identifying a certain positive property *P* that we have introspectively identified in *some* of our veridical perceptions. Or perhaps we think of *some* hallucinations simply as mental entities that are introspectively indiscriminable from veridical perceptions but that do not involve the perception of any ordinary objects.¹⁹ Either way, however, the sampling theorist does not have to suggest that we can precisely identify any (metaphysically interesting) property shared by *all* hallucinations, whether by introspection or by any other means. We just have to be able to identify a *few* hallucinations for purposes of sampling. And surely all theorists can agree that we can do that much, somehow, by extrapolating from veridical perceptions.

Thus, the sampling view is extremely modest – indeed, virtually minimalistic – in its claims about the epistemic powers of introspection.

The problem of extensional inadequacy revisited. We can conceive of hallucinations that are introspectively discriminable from any genuine perception, such as hallucinations of Escher-like staircases. The indiscriminability view has trouble explaining how this is possible. For, recall, the indiscriminability theorist might respond that such hallucinations are exhaustively decomposable into parts which are each introspectively indiscriminable from genuine perceptions. But some hallucinations, such as a hallucination of supersaturated redness, seem not to be decomposable into such parts. Can the sampling view explain how we can conceive of such hallucinations?

¹⁹ For what it is worth, my own view is that [brief identifying remark and reference to Author's Work A removed].

I think that it can. The explanation would be that we notice some important similarities between hallucinations of the impossible and paradigmatic cases of perceptual experiences, which leads us to believe that hallucinations of the impossible belong to the same relevant kind as the paradigmatic perceptual experiences. The sampling theorist does not need to take a stand on which similarities those are, exactly, but I find it immediately plausible that there are important similarities here. For example, just as there is an obvious, if hard to articulate, difference between thinking abstractly about the color crimson and hallucinating it, there is also an obvious, if hard to articulate, difference between thinking abstractly about supersaturated redness and hallucinating it.

More generally, I think that the sampling view gives us an especially promising approach to drawing the boundary around perceptual experiences. One problem with the positive view and the indiscriminability view is that they both try to specify in advance some reasonably simple, precise description that distinguishes veridical perceptions, illusions, and hallucinations from everything else. But we all know from the history of philosophy how difficult a feat it is to find such simple, precise descriptions in advance. The sampling theorist wisely refrains from trying to give such a simple, precise description in advance: she just collects samples and points, imprecisely, to *whatever* relevant kind those samples belong to.

But what if there is no single natural kind here – what if there is no single natural kind to which all and only the perceptual experiences belong? A fair question, and one that I will return to later.

The problem of special indiscriminability revisited. The indiscriminability theorist appeals to a notion of impersonal indiscriminability that is special-purpose: the notion seems not

to apply to anything besides introspection. Does the sampling theorist deploy this or any other special-purpose notion?

She does not. To be sure, she might want to use the notion of introspective indiscriminability, *simpliciter*, to explain how we can conceive of particular samples of hallucinations. In particular, the sampling theorist might want to say that we can conceive of particular samples of hallucinations, even if we have never had any hallucinations, simply as mental entities that are introspectively indiscriminable from veridical perceptions but that do not involve the perception of any ordinary objects.

But the sampling theorist does not rely on the notion of *impersonal* introspective indiscriminability. She can just use the *ordinary* notion of introspective indiscriminability, for on her view we just need the notion of introspective indiscriminability to help us collect a few samples of hallucinations; we are not trying to collect every possible sample of hallucination. And the ordinary notion of introspective indiscriminability is one that all parties to the debate can agree to use.

The only other remotely unusual notion that the sampling theorist uses is the notion of a natural kind. But this is certainly not a special-purpose notion: it is used by contemporary theorists for a whole slew of purposes: to solve Quine's problem of semantic indeterminacy, to solve Goodman's problem of gruesome inductions, to provide an analysis of qualitative duplicates, to characterize nomological laws, and much, much more.²⁰

The problems of semantic immodesty and semantic incongruence revisited. According to the indiscriminability theorist who also accepts metaphysical disjunctivism, the concept of perceptual experience characterizes the very essence of *some* perceptual experiences. In

²⁰ See, e.g., Lewis (1983). See also Sider (2011) on "structure."

particular, it characterizes the essence of hallucinations: the *concept* of a hallucination is the concept of something that is exhaustively decomposable into parts that are impersonally introspectively indiscriminable from a veridical perception, and that is also *what it is* to be a hallucination. This suggestion is semantically immodest: it says that our concept tracks the metaphysical natures of some (though not all) of its referents. It is also semantically incongruent, since there are no other obvious examples of concepts that work in this way.

How well does the sampling view handle these problems?

It handles the problem of semantic immodesty very well, I think. That should not come as a surprise: the sampling view is an instance of a broader contemporary approach to semantics that is intended to avoid such semantic immodesty. One idea that guides the sampling view, and contemporary approaches more broadly, is that we do not have to know the essence of a natural kind in order to form a concept of that kind. We can just collect some paradigmatic samples of the kind and use those to vacuum in all remaining members of the kind, whether by a sub-personal process of reference-fixing or by a personal-level but implicit understanding of the relationship between the paradigmatic samples and the other members. So the problem of semantic immodesty – the complaint that our very concepts reveal the essences of some of their referents – does not arise on the sampling view.

Similarly, the problem of semantic incongruence never arises on the sampling view. From the start, the sampling view is just an attempt to extend an utterly standard account of concepts to one more case. So rather than treating the concept of perceptual experience as a lone dangler, it treats it as utterly ordinary.

Conclusion. In light of these considerations, I believe that the sampling view has the same major advantage – the advantage of epistemic modesty – as the indiscriminability view.

But the four problems for the indiscriminability view either are easily solved or do not even come up in the first place. All in all, I think that makes for a strong *prima facie* case for the sampling view.

Still, it is worth checking whether the sampling view generates any major new problems of its own. Does it?

4. An objection

If the sampling view is true, then it is a conceptual truth that if there are any perceptual experiences, then there is a natural kind that includes all and only the perceptual experiences. You might worry, however, that this is not a conceptual truth. It is perfectly possible to discover that jade exists but is not in fact a chemical kind, but rather comprises the quite different chemical kinds *jadeite* and *nephrite*. And similarly, you might worry, it is perfectly possible to discover that there is no single mental kind that includes all perceptual experiences, but only several quite different mental kinds. For instance, perhaps there is just one kind that comprises veridical perceptions and illusions, and a very different kind that comprises hallucinations. Or perhaps things are even worse – perhaps hallucinations fragment into many kinds, or fail to belong to any kinds at all. These are live philosophical views!²¹ So if the sampling view says that these views are conceptual impossibilities, then so much the worse for the sampling view.

²¹ For theorists who sympathize with some such disjunctivist theory, see Hinton (1967); Campbell (2002); Martin (2004) and (2006); Snowdon (2005); Fish (2009); Nudds (2009); Brewer (2011); Logue (2012). For a bracingly clear overview of different forms of disjunctivism, see Soteriou (2016).

I think that this objection is a good one. But I suggest that it shows only that the sampling view that I have advanced is an oversimplification. And there are well-known options for avoiding the objection.

One option, which will likely appeal to reference-fixing theorists, is to say that *the world picks up the slack*: if there is not a single kind consisting of all and only the veridical perceptions, illusions, and hallucinations, then the concept of perceptual experience refers to the multiple natural kinds that the sample perceptual experiences belong to. Or, if there are *no* natural kinds in the vicinity, then the concept of perceptual experience just refers to anything that has enough of the superficial features shared by the sample perceptual experiences.

Another option, which might appeal to neo-descriptivists, is to say that the concept of perceptual experience is a conditional concept. The idea would be that it is *not* a conceptual truth that there is a single natural kind comprising all and only the veridical perceptions, illusions, and hallucinations. Rather, it is a conceptual truth that *if* there is a single natural kind comprising all and only the veridical perceptions, illusions, and hallucinations, then perceptual experiences are just the members of that kind; but if not, then perceptual experiences are something else. For example, it might also be a conceptual truth that *if* there is not a single natural kind comprising all and only the veridical perceptions, illusions, and hallucinations, but if there *is* a relatively small number of natural kinds comprising all and only the veridical perceptions, illusions, and hallucinations, then perceptual experiences are just the members of those kinds.²²

Notice that our theory of reference already needs some such refinement. *Everyone* needs to explain how the concept of jade gets its reference. And, I suggest, once we add the refinement

²² On conditional concepts, see Braddon-Mitchell (2003).

elsewhere, we might as well add it here. So the objection will be blocked, and at no additional cost.

5. Conclusion

It is metaphorical but not hyperbolic to say that *Naming and Necessity* has revolutionized the way that philosophers think about reference. And once Kripke's insights have been digested, the sampling view becomes a very natural hypothesis. Here I have argued that this hypothesis has many advantages: it is epistemically modest, it is extensionally adequate, it does not require special-purpose notions, it is semantically modest, and it is semantically congruent. I therefore suggest that it is the best going hypothesis about the concept of perceptual experience.

REFERENCES

- [Author's Work A].
- Braddon-Mitchell, D. (2003). "Qualia and analytical conditionals." *Journal of Philosophy* 100 (3): 111-135.
- Brewer, B. (2011). *Perception and Its Objects*. Oxford: Oxford University Press.
- Campbell, J. (2002). *Reference and Consciousness*. Oxford: Oxford University Press.
- Chalmers, D. (2012). *Constructing the World*. New York: Oxford University Press.
- Fish, W. (2009). *Perception, Hallucination, and Illusion*. Oxford: Oxford University Press.
- Hinton, J. (1967). "Visual experiences." *Mind* 76: 217-227
- Jackson, F. (1998). *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Clarendon Press.
- Kripke, S. (1972). *Naming and Necessity*. Cambridge, Massachusetts: Harvard University Press.
- Lewis, D. (1983). "New work for a theory of universals." *Australasian Journal of Philosophy* 61 (4): 343-377.
- Lewis, D. (1984). "Putnam's paradox." *Australasian Journal of Philosophy* 62 (3): 221-236.
- Logue, H. (2012). "What should the naïve realist say about total hallucinations?" *Philosophical Perspectives* 26: 173-199.
- Martin, M. (2004). "The limits of self-awareness." *Philosophical Studies* 120: 37-89.
- Martin, M. (2006). "On being alienated." In *Perceptual Experience*, eds. T. Gendler and J. Hawthorne (Oxford University Press), pp. 354-410.
- Nudds, M. (2009). "Recent work in perception: Naïve realism and its opponents." *Analysis Reviews* 69 (2): 334-346.
- Sider, T. (2011). *Writing the Book of the World*. Oxford: Oxford University Press.
- Siegel, S. (2008). "The epistemic conception of hallucination." In *Disjunctivism*, eds. A. Haddock and F. Macpherson (Oxford: Oxford University Press), pp. 205-224.
- Snowdon, P. (2005). "The formulation of disjunctivism: a response to Fish." *Proceedings of the Aristotelian Society* 105 (1): 129-141.