Philosophers' Imprint

VOLUME 25, NO. 48
DECEMBER 2025

NON-INFERENTIAL KNOWLEDGE OF PERCEPTION

Jonathan Brink Morgan

Montclair State University

doi.org/10.3998/phimp.4265

1. Introduction

Suppose that I see a lizard clinging to a wall. In doing so, I am aware of the lizard. With relative ease, I can shift my awareness from the lizard's head to its toes, or even to the texture of the wall. But I cannot in the same way shift my awareness to my seeing *itself*. When I try to do this, no new item "shows up"—despite my best efforts, there is only the lizard and the scene before my eyes. In this sense, my seeing is *transparent* to me.[1] Of course, I *know* that I see and, more specifically, that I see a lizard. This sort of knowledge is utterly mundane. But how exactly do I achieve it if visual perception is transparent? How do I know that I see without being aware *of my seeing*? This is *the puzzle of transparency.*[2] In what follows, I'll motivate and develop a novel solution to this puzzle.

To appreciate this puzzle's force, note that the puzzle of transparency is structurally similar to, but something of an inversion of, traditional puzzles about knowledge of the material world. As an illustration, consider the fact that some (like Jackson 1977) think that the only things we directly perceive are mental proxies for the material world—sensations, sense data, ideas, etc. At most, these theorists urge, we perceive the material world *indirectly i.e.* we perceive it only in virtue of perceiving whatever mental proxies "stand in" for material objects. But now we have a puzzle: How can I be sure that the mental proxies stand in for *any* material objects at all? That is, if I do not directly perceive the material world, then how do I know that there is such a world? This is similar to the problem posed by the puzzle of transparency: If I am not directly aware of my own seeing, then how

1. For a sampling of perspectives on transparency, see Dretske 2003a, Harman 1990, Martin 2002, Pautz 2007, Siewert 2004, and Tye 1995: 30-31, 2002. Remarks about transparency date back at least to G.E. Moore (1903: 446-451). There are traces of the idea in Gilbert Ryle's work as well. For example, Ryle writes that "[i]f I descry a hawk, I find the hawk but I do not find my seeing of the hawk. My seeing of the hawk seems to be a queerly transparent sort of process. . . " (1949/2009: 134).
2. The puzzle is well-known. See Byrne 2018: 74-98 and Dretske 1995: 39-63, 1996, 1999, 2003b.

do I know that I see?

One way out of the puzzle of transparency is to deny the relevant claim about awareness. That is, one could deny that visual perception is transparent and hold that one is aware of her own seeing. But this isn't the avenue I take here. I accept that visual perception is transparent—it doesn't *merely seem* transparent.[3] Moreover, I accept a maximally strong version of transparency on which seeing as well as other perceptual relations are *never* objects of awareness.[4] So, I won't attempt get out of the puzzle by rejecting or weakening transparency. Of course, as an alternative, I could simply deny that one has knowledge of her own seeing. However, I regard knowledge of one's own seeing as a Moorean fact—"[d]eny it, and the most credible explanation of your denial is that you are in the grip of some philosophical (or scientific) error" (Lewis 1997: 325). My assumptions are, therefore, that visual perception is transparent and that, nonetheless, we regularly have knowledge of our own seeing. The puzzle is how both could be true.

Extant solutions to the puzzle of transparency tend to be *inferentialist*. They propose that one infers that she sees thus-and-such from a certain proposition about her environment in the sense that one transitions from a *belief* about her environment to the belief that she sees thus-and-such, and that this inference is knowledge-generating. Alex Byrne's (2018: 128-155) solution to the puzzle of transparency clearly fits this description. Fred Dretske's (1995: 39-64) solution resembles Byrne's in many respects and is generally thought of as inferentialist, although the inference might be located in a slightly different place.[5] However, inferences of this sort have been criticized as either not characterizing what we actually do when gaining self-knowledge (Bar-On 2004) or as "mad" (Boyle 2011). But transparency advocates seem to have little choice but to defend inferentialism. Given visual transparency, it would appear that one's knowledge of seeing must somehow involve an inference.

However, friends of transparency need not accept inferentialism—there is, it turns out, a *non*-inferential alternative. The alternative becomes visible once we see that inferentialism is an instance of a more general kind of solution to the puzzle of transparency—what I call a *rule-following solution* (Sect. 2). Why seek a non-inferential rule-following solution? The above-mentioned objections (from Bar-On and Boyle) provide some reason to do so. But my primary reason is that extant inferentialist accounts suggest something striking: *the contents of vision cannot be expressed in language* (Sect. 3). Quite literally: you cannot say what you see. However, there are both intuitive and theoretical reasons to reject this kind of inexpressibility. And, in light of this, I present a non-inferential alternative: *the experiential rule-following account* (Sect. 4). On this account, there is no transition from a *belief* about one's environment to the belief that one sees thus-and-such. No such belief-to-belief transition is necessary. Instead, one simply transi-

3. Gow (2016) explains how to accept the latter without accepting the former. An interesting exercise would be to see how the puzzle of transparency fares if one denies that visual perception is transparent but accepts that visual perception *seems* transparent.
4. See Kind 2003 for an overview of different strengths of transparency.
5. Dretske seems slightly uncomfortable saying that an inference is involved on his account (1995: 60-62), but is ultimately willing to grant that it is. Here's how. Dretske thinks that *displaced perception*—the general sort of "perception" by which we come to know about our own mental states—requires a transition between beliefs. An example of displaced perception: You come to know something about your weight by seeing a particular reading on a bathroom scale. Dretske writes that "[t]hose who do not come to the perceptual situation with a connecting belief (as I will call it) will perceive the same object (bathroom scale) but not the same facts (that they weigh so many pounds)" (Dretske 1994: 265). In this scenario, forming the belief that you are a certain weight requires *transitioning from another belief*, namely, a belief about the connection between scale readings and weight. Since knowledge of one's seeing is also an instance of displaced perception for Dretske, there is a similar sort of connecting belief that must be involved in coming to know that one sees (Dretske 1995: 61). In a way, the epistemic rules that Byrne uses to form his account must, for Dretske, be *believed* in order to attain knowledge of seeing.

tions from *having* (but not being aware of!) a certain visual experience to the belief that she sees. Properly understood, this transition turns out to be perfectly reliable in that it always issues in a true belief. And given this reliability, the experiential account can offer a non-inferential explanation of how one knows that she sees—an explanation that is nonetheless compatible with the transparency of visual perception.

2. Rule-Following Solutions

2.1 The puzzle of transparency within a rule-following framework

To begin, we need to recognize that inferentialist solutions to the puzzle of transparency are instances of a more general kind of solution. Specifically, there is a family of solutions to the puzzle of transparency whose members we might call *rule-following solutions*. (These solutions also allow us to give a more precise formulation of the puzzle. More on this shortly.)

The most general feature of a rule-following solution is that it aims to solve the puzzle of transparency by means of characterizing certain belief-forming or *epistemic* rules. The relevant rules are material conditionals of the following form:

If p, then believe that q.

The rule tells a subject to believe that q in light of p, where p is a proposition mentally represented by the subject following the rule.

There are various ways of filling out this sort of schema—if there is sufficiently strong evidence for p, then believe that p; if the ground is wet, then believe that it has rained; if there are two people in the room, then believe that there is at least one person in the room; if there are ducks in the pond, then believe that the sun will explode tomorrow, etc. Most of these rules are in some sense "good". Others are not. Still, they are all epistemic rules.

To be clear, one does not need to *know* that she is following an epistemic-rule in order to follow it. If I believe that p and then come to believe that q as a result, I have followed the rule *if p, then believe that q*. I can do all this without even so much as knowing that this epistemic rule exists let alone that I am following it.

Epistemic rules can be *inferential* or *non-inferential*. The precise nature of inference is subject to debate (see Boghossian 2014 and Quilty-Dunn and Mandelbaum 2018 for details) and different conceptions might be necessary for different theoretical needs. However, within the context of the puzzle of transparency, inference is understood to be "a broadly causal process of transitions between belief states" (Byrne 2018: 125). Thus, one infers q from p if and only if she believes p and then, through the appropriate causal process, comes to believe q. For this reason, on inferentialist rule-following accounts, following a rule like *if p, then believe that q* requires, at a minimum, believing the antecedent of the rule and, as a result, forming the belief specified by the consequent. A non-inferential rule is thus defined *negatively*—it is a rule that can be followed without performing an inference. So, for example, if we imagine that a rule like *if p, then believe that q* is non-inferential, then this simply means that one needn't *believe* that p in order to follow the rule. Of course, one must mentally represent p in some way in order to follow the rule. It's just that the relevant representation need not be a belief.

The rule-following framework is useful partly because it allows us to sharpen up the puzzle of transparency.[6] The challenge presented by the puzzle: Find an epistemic rule that has two features. First, the rule must be such that it enables one to gain knowledge of her own seeing. In fact, it must do slightly more than this. As I said in the introduction, knowledge of seeing is easy to achieve and we have it regularly. The relevant epistemic rule must reflect this by being such that one does, in fact, regularly follow it and, in doing so gains knowledge that she sees. Second, the rule must be compatible with the transparency of visual perception. In other words, we need a rule that captures the fact that when one gains knowledge of her own seeing, one somehow tran-

6. Or, at the very least, it provides *one way* of sharpening up the puzzle. There may be other, legitimate ways.

sitions from an awareness of wholly non-mental items to knowledge that she sees these items. We need what I'll call a *transparent epistemic rule i.e.*, an epistemic rule whose antecedent concerns ordinary visible objects, events, and features but does not concern visual experiences or seeing itself.

Within the rule-following framework, the question of whether the puzzle of transparency can be solved becomes the question of whether there is a transparent epistemic rule that we regularly follow and is knowledge-generating with respect to one's own seeing. And in light of this, the puzzle becomes especially acute. For there seems no transparent epistemic rule that yields knowledge of one's seeing. To illustrate, consider a simple rule:

If there is an F in front of you, then believe that you see an F.

And for concreteness, let's focus on an instance of this rule:

If there is a lizard in front of you, then believe that you see a lizard.

This instance of the rule is transparent, for to follow it, I need only recognize that there is a lizard in front of me and then form the belief that I see a lizard. But the problem is that even if this transparent rule does to some extent describe what one does when she comes to believe that she sees, it is not enough to solve the puzzle. Following the rule wouldn't generate *knowledge* of seeing. At the very least, the fact that a lizard is in front of you is not a good reason to believe that you *see* a lizard. Even if it provides some probabilistic evidence that you see a lizard, this evidence is rather weak and seems insufficient to ground the relevant knowledge.[7]

The problem is that it is not immediately clear how to revise such a rule such that it is knowledge-generating *and* remains transparent. For example, one might suggest the following rule:

If you have a visual experience of a lizard in front of you, then believe that you see a lizard.

Maybe this rule would be knowledge-generating if followed. But it would not be transparent. The reason is that the antecedent contains information about something that, by the lights of transparency, one is not aware of, namely her own visual experience. Though tempting, simple solutions of this kind simply aren't available to the proponent of transparency. We must find some other solution.

3. Inferential solutions

Alex Byrne (2018) thinks he has one. His proposal depends in part on the idea that whatever makes the relevant epistemic rule knowledge-generating will be somehow more *externalist* than internalist. In short, while it is true that the antecedent doesn't provide adequate evidence for believing the consequent, there are other epistemically-relevant properties that can explain why following a transparent rule is knowledge-generating (2018: 122-124). The natural candidate is some kind of *reliability*, broadly construed. That is, the rule must be such that following makes it sufficiently likely that one will form a true belief to the effect that one sees thus-and-such. And, in the present context, this means the rule must be such that mentally representing the antecedent (whether doxastically or in some other mode) makes it sufficiently likely that that one in fact *sees*. Let's provisionally assume that an epistemic rule is defeasibly knowledge-generating if (but not only if) the rule is sufficiently reliable in producing true beliefs in this way. In other words, if one follows a sufficiently reliable epistemic rule and thereby forms a true belief, then absent sufficiently strong defeaters, she possesses knowledge.[8]

7. See Connee and Feldman 1985 for a discussion of evidentialism and Barnett 2016 for the claim that inferentialist accounts fail for roughly this reason. See Aydede 2003 for related criticism.

8. I emphasize that I am *assuming* this and not arguing for it. One might, of course, reject this assumption—especially if she is inclined to reject all forms of reliabilism. For worries about even *perfect* reliability and knowledge, see Hawthorne's (2002) discussion of *hyperreliable* belief-forming methods. For a contrasting view, see Williamson 1986: 114.

On Byrne's rule-following account, the reliability of the relevant rule depends in part on what it takes to follow it. For Byrne, this requires an *inference*. To follow the relevant epistemic rule according to Byrne (ibid.: 101-102), one must

(1) recognize (and thus believe) the antecedent;
(2) form the belief specified by the consequent; and
(3) form the belief specified by the consequent *because* one recognizes (and thus believes) the antecedent.

There is thus a kind of causal transition from belief to belief *i.e.* an inference. But what exactly is *recognition*? Recognition, according to Byrne, is a kind of propositional attitude (ibid. 101). It is propositional in that one recognizes *that* thus-and-such is so rather than recognizing some object (say, one's car or favorite hat). Moreover, Byrne claims that recognition implies knowledge in that if one recognizes that p, then one knows that p (ibid.). And given that knowing that p implies believing it, recognizing that p implies believing that p. Now, my suspicion is that Byrne uses 'recognition' as a convenient way of marking off a kind of occurrent or conscious state of knowing. To recognize that p just is to occurrently or consciously know that p. For this reason, I will occasionally drop talk of recognition in favor of talk of knowledge when describing Byrne's account.

What we need, then, is a sufficiently reliable rule that allows one to transition from knowing some fact about her environment to the belief that she sees. In other words, our question is whether there is a transparent rule whose following reliably generates true beliefs that one sees thus-and-such. It is not, for reasons already mentioned, immediately clear what such a rule would look like. But Byrne thinks he has found a rule that is *almost* perfectly reliable—successfully following it nearly always results in a true belief.[9] I'll focus on a simplified version of this rule.

SEE: If $[\ldots o \ldots]_V$, then believe that you see o.[10]

Now, $[\ldots o \ldots]_V$ is what Byrne calls a *v-proposition* or, equivalently, a *v-content*. V-contents are meant to capture what is *seen* by a perceiver. They thus concern ordinary objects and visually sensible qualities of those objects like color, illumination, shape, size, motion, and so on (ibid. 137).[11] Because they feature particular, ordinary objects, they are typically *singular.* We may contrast singular propositions like *o is F* with existential propositions like *there is an x such that x is F*. The former concerns a specific individual o. The latter concerns whatever individual (in the relevant domain) happens to be F. V-contents fall into the former category. So, for example, when I see the lizard l, the v-content that characterizes what I see might be something like $[\ldots l$ is green and lizard-shaped $\ldots]_V$. This content concerns a specific lizard l, attributes sensible qualities to this lizard, and, in general, captures the visual appearance of the lizard upon this occasion of perception. When a singular content concerns a particular object o, I will say that the content is *o-involving*.

SEE is a transparent rule in the sense described above: its antecedent makes no mention of seeing as such and concerns only visible objects and their features. And it seems to describe (at least approximately) what one does in coming to believe that she sees. The pressing question, therefore, is whether SEE is sufficiently reliable *i.e.* whether it is good enough at producing true beliefs.

However, this solution faces an immediate and familiar challenge.

9. Throughout Byrne 2018, a perfectly reliable rule of this sort is called *self-verifying*. Some rules are also *strongly* self-verifying for Byrne in that if one merely *tries* to follow them, one is guaranteed a true belief (ibid. 107). I do not pursue strongly self-verifying rules here.

10. Byrne's version of the rule reads: If $[\ldots o \ldots]_V$ and o is F, then believe that you see an F (2018: 139). This is a rule for gaining what I would call *sophisticated knowledge* of seeing where one brings the thing seen under some general, usually non-demonstrative concept. I focus on *simple knowledge* of seeing that does not involve this kind of conceptualization since it harmlessly simplifies the discussion of rule-following solutions.

11. To be clear, I am therefore assuming that v-contents are Russellian propositions *i.e.* structures of worldly entities like objects, properties, and relations (as opposed to, say, modes of presentation of these things).

It isn't clear how merely recognizing an *o*-involving v-content could serve a reliable basis for truly believing that one sees *o*. For whether an *o*-involving v-content is true is just a matter of whether *o* has the sensible qualities attributed to it by the v-content. But even if it is true that lizard *l* is (say) green, that seems largely irrelevant to whether I see *l*. And even setting that problem aside, See seems to be a rule that one can follow even if she sees *nothing*. Imagine someone—Ally—who is congenitally blind. Ally can quite easily come to recognize that lizard *l* is green and lizard-shaped by means of testimony. And if they do so, then they too know the relevant v-content, and are thereby in a position to follow See and come to form the false belief that they see lizard *l*. But this would mean that See isn't sufficiently reliable. Following it would quite easily lead to false beliefs.

Byrne is well-aware of this. And his solution is clever. His suggests that "in all ordinary situations, one knows that $[\ldots x \ldots]_V$ only if one *sees* x" (ibid. 140). In other words, recognizing an *o*-involving v-content is sufficient for seeing *o* in ordinary circumstances. Why exactly? Because, Byrne claims, "[v]ision is, at least in creatures like ourselves, an exclusive conduit for v-facts [*i.e.* true v-contents]. Hence [if one knows a v-content] one's information source must be vision, not audition, olfaction, testimony, or anything else" (ibid). So, if you've managed to recognize and hence know $[\ldots o \ldots]_V$, the source of your knowledge is vision itself. You must be visually presented with this content. But being visually presented with $[\ldots o \ldots]_V$ is sufficient for seeing *o*. In fact, in the current context, it's reasonable to think that being visually presented with $[\ldots o \ldots]_V$ *constitutes* seeing *o*. For this reason, the very fact that you know an *o*-involving v-content is, it would seem, sufficient for your seeing *o* in ordinary contexts. And this, in turn, is how Byrne gets around worries about those who see *nothing* being able to follow See. In short, someone who sees nothing *can't follow* See *at all*. Why? Following See requires knowing a v-content. But visual perception is "an exclusive conduit" (ibid.) for knowledge of v-contents. Since those who are congenitally blind lack visual perception, they cannot know v-contents. Hence, they cannot follow See at all and the challenge to See's reliability is met.

At this point, we can piece together Byrne's transparency-friendly, inferentialist story about how one comes to know that she sees an object *o*. First, one comes to recognize (and thus believe) a v-content $[\ldots o \ldots]_V$. Then, on the basis of this recognition, one comes to believe that she sees *o*. Importantly, forming the belief in this way is sufficiently reliable because if one has recognized $[\ldots o \ldots]_V$, then she must actually see *o* since the source of $[\ldots o \ldots]_V$ can be none other than visual perception itself (in ordinary circumstances). So, one has a true belief that one sees *o* on the basis of following a sufficiently reliable rule. And given a modest version of reliabilism, it will follow that following such a rule enables one to know that she sees *o*. This process is perfectly consistent with the transparency of visual perception. It also seems to describe what we do when we come to believe that we see. It is therefore a strong contender for solving the puzzle of transparency.

3.1 Are the Contents of Vision Inexpressible?

Despite its promise, I am going to argue that Byrne's inferentialism has a bizarre consequence: the contents of vision are linguistically *inexpressible*.[12] Now, I think Byrne likely realizes that he is committed to something like this, but he isn't terribly explicit about this commitment. For example, he writes that one might learn that a v-content is true by "reading it in the—as-yet-unwritten—language of vision" (2018: 140). This strongly suggests that Byrne thinks that we do not currently possess a language in which v-contents are expressible. Regardless of what Byrne thinks, it is important to see why he is in fact committed to the inexpressibility of v-contents. Here is the argument in outline.

P1. If Byrne's inferentialism solves the puzzle of transparency, then v-contents cannot be known by means other than vision—otherwise, See becomes unreliable.

12. See Hofweber 2017 for a general discussion of inexpressibility and ineffability.

P2. If v-contents cannot be known by means other than vision, then they cannot be known via testimony.

P3. If v-contents cannot be known via testimony, then they are linguistically inexpressible.

C. Thus, Byrne's inferentialism solves the puzzle of transparency only if v-contents are linguistically inexpressible. (from P1-P3)

The consequent of this conclusion is hard to justify. It is immensely plausible that v-contents *are* linguistically expressible. And given this, we should likely reject Byrne's inferentialism as a solution to the puzzle of transparency—it simply isn't worth the cost.

Let's unpack this argument premise by premise. Begin with the first: Byrne's inferentialism solves the puzzle of transparency only if v-contents cannot be known by means other than vision. To see why this is true, consider the contrapositive: If v-contents *can* be known by means other than vision, then Byrne's inferentialism *does not* solve the puzzle of transparency. Byrne explains the knowledge-underwriting reliability of See by appealing to the idea that "in all ordinary situations, one knows that $[\ldots x \ldots]_V$ only if one *sees* x" (2018: 140). Or, put differently, See is reliable precisely because successfully following it requires sight. However, if one can come to know that $[\ldots x \ldots]_V$ by means other than sight, then it is possible to follow See even when one does not in fact see x. We would once again find that See is not sufficiently unreliable—as illustrated in the above case of Ally—and following it cannot underwrite our knowledge of seeing. Thus if v-contents are knowable by non-visual means, Byrne's inferentialism does not solve the puzzle of transparency. Contrapositively, Byrne's inferentialism solves the puzzle of transparency only if v-contents cannot be known by means other than vision.

The second premise can be defended rather quickly. From the fact that v-contents cannot be known by any means other than vision, it follows that v-contents cannot be known by *testimony*—specifically, the testimony of others as delivered through auditory perception.

The third premise is, however, in need of explanation. It says that if v-contents cannot be known by testimony, then they are not linguistically expressible. Again, I'll defend the contrapositive: If v-contents *are* linguistically expressible, then they *are also knowable* by means other than vision.

Suppose that v-contents are expressible in natural language *i.e.* there is at least one grammatical expression in an extant natural language that has a v-content as its semantic value. If they are thus expressible, then they are quite likely knowable via testimony. This is for two reasons. The first is general: Under normal conditions, if you hear someone say (*i.e.* express) that p, then you are thereby in a position to know that p via testimony. The details of the story about how exactly this works may be complex and are by no means uncontroversial. For example, perhaps you need to have reasons to accept the speaker's epistemic credibility (a view sometimes attributed to David Hume). Or perhaps things are less demanding, and you simply have a presumptive right to accept p when it is asserted by another (a view sometimes attributed to Thomas Reid). But the details do not matter for present purposes. It should be relatively uncontroversial that *some* story vindicates the fact that under at least some conditions, if you hear someone say that p, then you are thereby in a position to know that p via their say-so. If there isn't such a story, then a general kind of testimonial skepticism follows.

The second reason is specific to v-contents: If v-contents are expressible, then reflection on cases suggests that they can be known via testimony. Let's focus on a concrete example to illustrate.

> Leadbelly: You and your friend are taking a break from your hike. As you stare at the river beside you—the Leadbelly—you tell your friend, "What a day. The Leadbelly is really blue." Suppose that part of what you express here is the v-content $[\text{Leadbelly is blue}]_V$. Your friend is currently digging through your hiking packs. Hearing you, but not looking up, he replies, "That sounds right." Your friend is affirming what you said. And his linguistic affirmation is a reflection of a cognitive affirmation: he

> accepts the very v-content that your utterance expressed, namely, [Leadbelly is blue]$_V$. So, there is a v-content that you linguistically express and your friend entertains that v-content as a result of encountering this linguistic expression. On the basis of this, your friend comes to know that [Leadbelly is blue]$_V$.

If we like, we may further stipulate that your friend has no reason to doubt your sincerity, doesn't think that you have ill-intentions, and, in general, regards you as reliable (which you in fact are). Thus, if you have actually said that [Leadbelly is blue]$_V$, then your friend is in a position to know that [Leadbelly is blue]$_V$ on the basis of your testimony. And, generalizing, if v-contents are linguistically expressible in this way, then they are knowable by means other than vision.

Taken jointly, these reasons provide two different, albeit related, avenues of support for the claim that if v-contents are linguistically expressible, then they are knowable by means other than vision. Likewise, they provide support for the contrapositive: if v-contents are not knowable by means other than vision, then they are linguistically inexpressible. Thus, we have the third premise of the above argument. And given these three premises, it follows that the viability of Byrne's inferentialism depends on the inexpressibility of v-contents.

The problem is that it is immensely plausible that v-contents *are* linguistically expressible and there seem no sufficiently strong reasons to think otherwise. Consider the fact that a v-content captures *what you see*. If you cannot express this sort of content, then you cannot say what you see. But saying what we see is one of the most basic communicative acts in which we engage. *Toddlers* can do it. Worse yet, presumably the rule-following account isn't just supposed to apply to vision. It should apply to other sense modalities as well: hearing, taste, touch, smell, etc. The result will be that if I cannot say what I see, then I also cannot say what I hear, smell, taste, touch, etc. Nothing *at all* about the perceptual world will be expressible. For some, this might be acceptable. For adherents of common sense, it will require rejecting Byrne's inferentialism.

Common sense can only go so far, of course. Another reason to reject the idea that v-contents are inexpressible is that there is no plausible story as to why they *can't* be expressed in language. One thought is that v-contents don't concern public properties like colors and shapes. Instead, they concern something like private, experiential analogues of these properties. So when I say that the apple is red, the redness I attribute to the apple is *not*, strictly speaking, a property that I see. The content of my utterance is entirely distinct from the content of my current visual episode. This is, of course, a controversial view, though it is not without proponents (*e.g.* Thau 2002: 22 and Strawson 1989).[13] But setting aside controversiality (and motivation for that matter), this view simply does not line up with the communicative intentions of ordinary speakers. If, while looking at the apple, I tell Ally that the apple is red, I intend to tell them about a feature of the apple that I see. And if we assume that there is even a loose relationship between the contents of utterances and the communicative intentions of speakers, the idea that the properties expressed in language are not properties that we see is implausible.

Of course, there are various ways in which Byrne might try to soften the blow of this inexpressibility. For example, he might say that (1) only *complete* v-contents cannot be expressed in language and (2) following See requires knowledge of a complete v-content. A complete v-content captures the totality of what one sees whereas a *partial* v-content captures elements of what one sees, but not the totality. If so, we could allow that partial v-contents are expressible, but complete v-contents are not. To illustrate, recall the case of Leadbelly. Even if partial v-contents are knowable via testimony, this would not enable your friend in Leadbelly to follow See since he would need to be in possession of a complete v-content to do so. We could thereby mitigate (to some extent) concerns about expressibility all while maintaining

13. It's worth noting that Strawson (1989) holds a nuanced version of the view on which no *specific* experienced property is expressed by natural language terms like 'red'.

that the only way to know a complete v-content is via vision.

The follow-up question, of course, is this: Why is it that only partial, but not complete, v-contents are expressible? The obvious answer is that complete v-contents are tremendously complex and highly determinate, and, for this reason, aren't typically expressed in language. Consider the complexity and determinacy of the viewed scene before you. Now, imagine trying to express this in language. While it may be possible, we certainly don't linguistically express contents like this on a regular basis.

There is something to this idea. But the problem is that if it is right, and if following See requires believing a complete v-content, *then we likely never actually follow* See. If v-contents are too complex/determinate to be ordinarily expressed in language, then, for similar reasons, they are likely too complex/determinate to be ordinarily entertained in thought. The richness and complexity of the scene I presently see outstrips my *general*, conceptual, representational capacities. I do not have concepts for all the colors, shapes, and qualities that I see.[14] Maybe, if given enough time, I could conceptualize the full richness of this scene via demonstrative concepts. But that is not something that I or anyone else ordinarily does when they come to believe that they see. In short, I am (at least normally, but perhaps not in principle) incapable of entertaining a complete v-content in thought. And this makes it implausible that I or *anyone* ever actually believes, knows, or recognizes complete v-contents. We don't normally have the conceptual resources or time to do this. And, if so, and following See requires recognizing complete v-contents, this would make See virtually unusable by creatures like us.

What one *wants* to do here is somehow modify See such that one can use it to know an *o*-involving v-content only if one in fact sees *o*. But—to echo a point made earlier—we need to remember the puzzle we are trying to solve. We require that See be a *transparent* epistemic rule. We cannot, for example, modify See to say something like "If [. . . *o* . . .]$_V$ is conveyed to you by means of vision, then believe that you see *o*". This would, of course, make See impossible to follow for the unsighted and thus potentially secure reliability. But it would do so at the cost of rendering the rule non-transparent since the modified rule's antecedent concerns things *other than* ordinary visible objects and their visible features. It would also make the rule non-explanatory. Successfully following See is meant to explain how one knows that she sees. But if following the rule involves recognizing and thus knowing *that a content is conveyed to you by means of vision*, then following it more-or-less requires that one *already knows that she sees*. The rule could not, therefore, offer an illuminating explanation of how this knowledge is achieved.

Where does this leave us? One could simply bite the bullet here and accept that v-contents are linguistically inexpressible. I myself cannot find sufficient reason to do this. Byrne's account is, I think, far less plausible than the claim that we can say what we see, and trading the latter for the former seems like philosophical overreach. However, I also suspect that *some* account within the rule-following family is broadly-speaking correct. So, what's the alternative to Byrne's account?

3.2 *A Non-Inferential Account*

At this point, the rule-following framework developed earlier becomes especially important. For it allows us to see that Byrne's inferentialism isn't the *only* kind of rule-following solution to the puzzle of transparency. There are possible non-inferential solutions. The difference between the two sorts is, at bottom, in how they claim the antecedent of a rule like See is to be mentally represented in order for the rule to be successfully followed. If the antecedent is to be represented doxastically, then following See will require a transition from belief to belief, and thus we will wind up with an inferential account. However, if the antecedent is to be represented *non-doxastically*, then following See will require a transition from a non-doxastic state to belief, and thus we will wind up with a non-inferential account.

14. See Speaks 2005 for a related discussion about conceptual content.

In this section, I construct a non-inferential account that fits this description and also avoids the problem of inexpressibility. But this is not all my account does. What makes my non-inferential account interesting—and what enables it to solve the problem of inexpressibility—is that the non-inferential transition that one makes in coming to know that she sees is *perfectly* reliable. If one successfully follows the relevant rules—which I claim she does on a regular basis—then she is guaranteed to form a true belief to the effect that she sees thus-and-such. We are thus giving a plausible explanation of why knowledge of one's own seeing strikes us as especially secure. Along the way, I'll also explain how my account solves or avoids other problems that Byrne raises for his own solution.

Let's begin by reflecting on an ordinary case of coming to know that one sees something. I am going to use this case to suggest that (1) we do ordinarily follow a rule like See when we form beliefs about our own seeing but (2) the *way* in which we follow this rule is non-inferential.

> Clock Tower: If you ask me what I presently see through my office window, I can tell you with ease: I see a clock tower. Phenomenologically speaking, my coming to know this isn't a complex process requiring much cognitive effort. I simply have a visual experience and focus on the clock tower that is part of my experience's content. As a result, I come to believe that I see a clock tower—or even just that I see *that*.

I take Clock Tower to be representative of a wide range of ordinary cases where one comes to have knowledge that she sees.[15] And in this scenario, I claim that I am following See (or some rule that is not importantly different) since the basis for my belief that I see a clock tower is a v-content concerning some particular object, namely, the clock tower itself. However, I am not following See by transitioning from a *belief* about the external world to a belief about what I see. Instead, I simply *have a visual experience* and transition from this experience to the belief that I see a clock tower.

Now, if Clock Tower is indeed representative of a wide range of ordinary cases, as I claim, then it suggests a generalization about how one comes to know that she sees.

> The Experiential Rule-Following Account: Typically, when one knows that she sees object *o*, it is because she follows See; and a subject S follows See if and only if
>
> (1) S has a visual experience with content $[\ldots o \ldots]_V$ (where *o* is some particular object, event, or fact);
> (2) S attends to $[\ldots o \ldots]_V$ (or some aspect thereof); and
> (3) S forms the belief that she sees *o* as a (non-deviant) causal consequence of having the experience in condition (1) and attending to the content of this experience in the way specified

15. As an anonymous referee points out, this is crucial. For I placed a constraint on any adequate rule-following solution to the puzzle of transparency, namely, that it specifies a rule that we *do* regularly follow. I take the description of Clock Tower to aptly characterize *what one ordinarily does* when she comes to believe that she sees something. Importantly, and as Clock Tower suggests, what one ordinarily does in coming to know that she sees involves, among other things, following See. What I am suggesting, then, is reflection on these sorts of familiar cases supports a *reasonable empirical hypothesis*: we regularly follow See.

by condition (2).[16]

What is distinctive about this account is that it says that *following* See has an experiential aspect—it requires having a certain visual experience. Furthermore, the experiential account says nothing about coming to *know* or even *believe* the full content of one's visual experience. This seems right. Whatever the v-content of my experience in the clock tower example, I do not need to know or even believe that it is true when forming the belief that I see a clock tower. Instead, I simply *have* the experience, focus on its content, and *then* come to believe that I see a certain object. For this reason, the experiential account is a non-inferential account. To be sure, the account claims that a transition between representational states is required to follow rule See. But the transition one undergoes is a non-inferential one. It is a transition from an experience to a belief rather than a transition from a belief to a belief (given the plausible assumption that experiences aren't beliefs). In short, since a belief-to-belief transition is absent, we do not have an inference.[17]

The experiential account is consistent with the transparency of experience.[18] Notice that the experiential account *does not* say that one must become aware of her visual experience itself to follow See. She only needs to *have* the experience. What she is aware of is the *content* of this experience. And, again, the relationship between having the experience and forming the belief is purely causal. In other words, the fact that S has a visual experience with a certain content—a content to which she attentive—(non-deviantly) *causes* her to believe that she sees thus-and-such. This is part of what it takes to follow See. And one needn't know that she is following See in order to generate knowledge of her own seeing. One can successfully follow a knowledge-generating rule without knowing that one follows it. Indeed, given that we are assuming some minimal variety of reliabilism, as long as the rule in question is reliable, then it doesn't really matter, epistemically speaking, whether one knows that she is following the rule. All that matters is that one is successfully following the rule.

The experiential account bypasses the issue of inexpressibility completely. Saying that the contents of vision are inexpressible is required, on Byrne's account, only because saying otherwise threatens to make See look insufficiently reliable. But—and this is the important point—See is *perfectly reliable* on the experiential account. That is, it follows from the experiential account that one successfully follows See and forms the belief that she sees *o only if she actually sees o*. The reason for this is simple. To follow See, on must have a visual experience with *o* as

16. An anonymous referee wonders what the relationship is between the experiential account and *constitutivist* accounts of self-knowledge (e.g. Shoemaker 1994). On such accounts, S's being in a mental state M in some sense constitutes (perhaps via something like a part-whole relation) her knowing that she is in M. This might seem to some extent incompatible with the account I offer since I rely on causal relations between mental states to underpin self-knowledge. However, the two views might actually be compatible. One could hold that S's seeing *x partially* constitutes S's knowing that she sees *x* in the following sense: *what it is* to know that one sees *x just is* (1) to have a belief that one sees *x* and (2) for this belief to be appropriately causally related to the seeing of *x*. This seems to be compatible with both the experiential account and constitutivism. Much more needs to be said about this, of course, but I lack the space to do so here.

17. Of course, one might point out that there is something like an inference at play since there is a transition from one state with propositional content to another. It's just that one of the states is non-doxastic. This is a perfectly reasonable observation. And if one wants to use 'inference' so as to include transitions between doxastic and non-doxastic states alike, I do not object. But that is not how Byrne and others party to this debate use the term. At any rate, what is important is not how we label the account I offer, but that the account is different from Byrne's and others—different enough that it bypasses the problems they face. See Siegel 2018 (particularly ch. 5) for discussion of inference in general.

18. Thanks to an anonymous referee for requesting clarification on this.

part of its v-content. But having a visual experience with *o* as part of its v-content *constitutes* seeing *o*—at least given that one is (as the account requires) sufficiently attentive.[19] Thus, to follow See, with respect to an object *o*, one must actually see *o*. And, for this reason, whenever one follows See and comes to believe that she sees *o*, she is guaranteed to have a true belief. See is thus perfectly reliable when followed in this non-inferential way.

To illustrate, imagine someone who *tries* to follow See while merely having (say) an auditory experience of $[\ldots o \ldots]_V$ (if such an experience is possible). This person would not *actually be following* See—following See requires having a *visual* experience. Likewise, someone with congenital blindness who tried to follow See by means of merely recognizing $[\ldots o \ldots]_V$ through testimony would not be following See either. And your friend in Leadbelly, although he sees *something*, does not have a visual experience with *Leadbelly* as part of its content. For this reason, he cannot follow See and thereby come to believe that he sees Leadbelly. To repeat, if the experiential account of rule-following is right, then one can successfully follow See and come to believe that she sees *o only if* she in fact sees *o*. See is therefore perfectly reliable.

In addition to providing a perfectly reliable rule, thus bypassing the problem of inexpressibility, the experiential account also solves a number of problems that Byrne raises for his own variant of the rule-following account. I'll address two.

First, Byrne notes that there is a problem due to visual imagination and memory (Byrne 2018: 140-141). Plausibly, when remembering a viewed scene, one can entertain $[\ldots o \ldots]_V$ in visual imagination. But when one does this, that might put her in a position to follow See even though she no longer sees *o*. This would be bad. It would, in an ordinary circumstance, yield the result that following See doesn't issue in a true belief. We once again have a threat to See's reliability. To solve this problem, Byrne appeals to the notion of a *degraded* and *transformed* v-content (ibid. 141). Degraded/transformed v-contents are (by stipulation) not quite v-contents. And since following See requires, on Byrne's view, recognizing a v-content, one cannot follow See by visually imagining $[\ldots o \ldots]_V$ simply because one *cannot* visually imagine this sort of content.

But there is a much simpler solution that doesn't appeal to the existence of degraded/transformed v-contents (on whose nature I shall remain silent). The problem posed by visual imagination is basically the same problem posed by testimony: it is a worry about v-contents being conveyed by something other than a visual experience. And just as Byrne's account requires that testimony can't convey v-contents, so too does it require that imagination can't convey v-contents. Thus, my response to the problem of imagination is similar to my response to the problem of testimony. Following See, on the experiential account, requires having a visual experience with a particular v-content. By 'visual experience' I mean *perceptual* visual experience and not an *imaginative* visual experience. When one sees a red apple, one has a perceptual experience. When one imagines or visualizes a red apple, one has an imaginative experience. And someone who tries to follow See when they are not having a perceptual experience and are merely having an imaginative experience *is not actually following* See. Granted, they may be *trying* to follow it. But trying to follow a rule is not the same thing as actually following it. Thus, visual imagination poses no serious challenge to See on the experiential account. And note that in giving this explanation, I did not need to appeal to the notion of a degraded/transformed v-content. Nor did I need to claim that visual perception and visual imagination never share contents. Instead, I simply appealed to an independently plausible account of what it takes to follow See and the difference between perceptual and imaginative experiences (whatever that difference ultimately amounts to).

Another problem can be solved. Byrne worries that See falters in

19. In other contexts, this might be a controversial assumption. For example, one might want to say that seeing is a *non-content-involving* or *non-representational* relation between a subject and object (perhaps, for example, Fish 2009). Alternatively, one might opt for a purely *causal* theory of seeing (Grice 1961). However, this assumption is ultimately inessential and, if necessary, both accounts of seeing are compatible with the experiential account I offer.

the case of known illusions (2018: 142). Suppose that I see a blueberry on my kitchen table but, for one reason or another, I know that my present visual experience is illusory and the blueberry is not as it appears. In short, I recognize that the v-content of my experience is false. Barring severe cognitive dissonance, I won't believe let alone recognize or know this content. Byrne finds this puzzling, writing that "if I do not believe the relevant v-propositions, I cannot even try to follow See. Hence, cases of known-illusion threaten to blow the transparency proposal entirely out of the water" (ibid.). Byrne's solution is to say that visual perception is *belief-dependent* (ibid. 143-146). If I have a visual experience involving a v-content, then there must be some sense in which I believe that v-content. That is what enables me to follow the rule.

There is no need, however, to claim that visual perception is belief-dependent on the experiential account. According to the experiential account, one needn't know or even believe a v-content in order to follow See. Instead, one need only *have* a visual experience with that v-content (and attend to that v-content, of course). Belief is not required. To illustrate, suppose again that I see blueberry *b* on my kitchen table. Due to a powerful drug that I have knowingly consumed, the blueberry looks like a red apple *i.e.* it looks red and apple-shaped. The content of my visual experience is thus something to the effect that [. . . *b* is red and apple-shaped . . .]$_V$. But because I know that I have consumed the powerful drug, I haven't the slightest temptation to believe this content. Despite this, I can still know that I see *b* by following See. I have an illusory experience with *b* as part of its v-content, attend to this content, and then this causes me to form the belief that I see *b*. Thus, I am still following See. Even if problematic for inferentialism, known illusions are utterly unproblematic for the experiential account.

In a way, I hope the preceding experiential account seems obvious. Maybe it even seems like cheating! But it isn't. What we are trying to do in specifying a rule like See is make clear how we *do* in fact form beliefs about our own seeing and *why* what we do generates knowledge. We're trying to understand why our belief-forming behavior is epistemically acceptable. Showing that what we in fact do is conform to a perfectly reliable rule vindicates our belief-forming behavior, even if we are not regularly aware that this is the rule we are following. If I appear to have cheated, it is only because knowledge of seeing is so easily attained, and the game—when we pay attention to the rules—isn't that hard to win.

References

Aydede, Murat (2003). Is introspection inferential? In Brie Gertler (Ed.), *Privileged Access: Philosophical Accounts of Self-Knowledge* (pp. 65-82). Routledge.

Bar-On, Dorit (2004). *Speaking My Mind: Expression and Self-Knowledge*. Oxford University Press.

Barnett, David James (2016). Inferential justification and the transparency of belief. *Noûs*, 50(1), 184-212. DOI:10.1111/nous.12077

Boghossian, Paul (2014). What is inference? *Philosophical Studies*, 169(1), 1-18. DOI:10.1007/s11098-012-9903-x

Boyle, Matthew (2011). Transparent self-knowledge. *Aristotelian Society Supplementary Volume*, 85(1), 223-241. DOI:10.1111/j.1467-8349.2011.00204.x

Buchanan, Ray (2010). A puzzle about meaning and communication. *Noûs*, 44(2), 340-371. DOI:10.1111/j.1468-0068.2010.00746.x

Byrne, Alex (2018). *Transparency and Self-Knowledge*. Oxford University Press.

Conee, Earl & Richard Feldman (1985). Evidentialism. *Philosophical Studies*, 48(1), 15-34. DOI:10.1007/BF00372404

Dretske, Fred (1994). Introspection. *Proceedings of the Aristotelian Society*, 94, 263-278. DOI:10.1093/aristotelian/94.1.263

Dretske, Fred (1995). *Naturalizing the Mind*. MIT Press.

Dretske, Fred (1996). Phenomenal externalism or if meanings ain't in the head, where are qualia? *Philosophical Issues*, 7, 143-158. DOI:10.2307/1522676

Dretske, Fred (1999). The mind's awareness of itself. *Philosophical Studies*, 95(1), 103-124. DOI:10.1023/A:1004417407415

Dretske, Fred (2003a). Experience as representation. *Philosophical Issues*, 13(1), 67-82. DOI:10.1111/1533-6077.00009

Dretske, Fred (2003b). How do you know you are not a zombie? In Brie Gertler (Ed.), *Privileged Access: Philosophical Accounts of Self-Knowledge* (pp. 1-14). Routledge.

Fish, William (2010). *Perception, Hallucination, and Illusion*. Oxford University Press.

Gow, Laura (2016). The limitations of perceptual transparency. *Philosophical Quarterly*, 66(264), 723-744. DOI:10.1093/pq/pqw014

Grice, H. Paul (1961). The causal theory of perception. *Proceedings of the Aristotelian Society*, Supplementary Volume 35, 121-152. DOI:10.1093/aristoteliansupp/35.1.121

Harman, Gilbert (1990). The intrinsic quality of experience. *Philosophical Perspectives*, 4, 31-52. DOI:10.2307/2214186

Hawthorne, John (2002). Deeply contingent a priori knowledge. *Philosophy and Phenomenological Research*, 65(2), 247-269. DOI:10.1111/j.1933-1592.2002.tb00201.x

Hofweber, Thomas (2017). Are there ineffable aspects of reality? In K. Bennett and D. Zimmerman (Eds.), *Oxford Studies in Metaphysics: Volume 10*. Oxford University Press.

Jackson, Frank (1977). *Perception: A Representative Theory*. Cambridge University Press.

James, S. P. (2014). Hallucinating real things. *Synthese*, 191(16), 3711-3732. DOI:10.1007/s11229-014-0495-6

Johnston, Mark (2004). The obscure object of hallucination. *Philosophical Studies*, 120(1-3), 113-183. DOI:10.1023/B:PHIL.0000033749.18014.12

Kind, Amy (2003). What's so transparent about transparency? *Philosophical Studies*, 115(3), 225-244. DOI:10.1023/A:1025131015434

Lackey, Jennifer (1999). Testimonial knowledge and transmission. *Philosophical Quarterly*, 49(197), 471-490. DOI:10.1111/1467-9213.00154

Lewis, David (1997). Naming the colours. *Australasian Journal of Philosophy*, 75(3), 325-342. DOI:10.1080/00048409712347931

Martin, M. G. F. (2002). The transparency of experience. *Mind and Language*, 17(4), 376-425. DOI:10.1111/1468-0017.00205

Martin, M. G. F. (2004). The limits of self-awareness. *Philosophical Studies*, 120(1-3), 37-89. DOI:10.1023/B:PHIL.0000033751.66949.97

Moore, G. E. (1903). The refutation of idealism. *Mind*, 12(48), 433-453. DOI:10.1093/mind/XII.48.433

Pautz, Adam (2007). Intentionalism and perceptual presence. *Philosophical Perspectives*, 21(1), 495-541. DOI:10.1111/j.1520-8583.2007.00138.x

Quilty-Dunn, Jake & Eric Mandelbaum (2018). Inferential transitions. *Australasian Journal of Philosophy*, 96(3), 532-547. DOI:10.1080/00048402.2017.1398723

Ryle, Gilbert (1949/2009). *The Concept of Mind*. Routledge.

Shoemaker, Sydney (1994). Self-knowledge and 'inner sense'. *Philosophy and Phenomenological Research*, 54(2), 249-314. DOI:10.2307/2108490

Siegel, Susanna (2018). *The Rationality of Perception*. Oxford University Press.

Siewert, Charles (2004). Is experience transparent? *Philosophical Studies*, 117(1-2), 15-41. DOI:10.1023/B:PHIL.0000014524.37758.8b

Speaks, Jeff (2005). Is there a problem about nonconceptual content? *Philosophical Review*, 114(3), 359-398. DOI:10.1215/00318108-114-3-359

Strawson, Galen (1989). Red and 'red'. *Synthese*, 78(2), 193-232. DOI:10.1007/BF00869282

Thau, Michael (2002). *Consciousness and Cognition*. Oxford University Press.

Tye, Michael (1995). *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*. MIT Press.

Tye, Michael (2002). Representationalism and the transparency of experience. *Noûs*, 36(1), 137-151. DOI:10.1111/1468-0068.00365

Williamson, Timothy (2002). The contingent a priori: Has it anything to do with indexicals? *Analysis*, 62(3), 113-117. DOI:10.1093/analys/62.3.216