

Seeing Reasons

JENNIFER CHURCH

Vassar College

“I should like to say ‘I experience the because’” (Wittgenstein, P.I. §177)

We frequently speak of seeing the reason for an event or a state of affairs: “I see why the chair tipped over”, “You see why I’m angry”, “Now she sees why she should stay.” Does such talk of seeing simply register our success at acquiring knowledge of reasons, or does it refer to a distinctively perceptual, or perceptual-like way of knowing reasons? This paper defends an account of perception that focuses on the experience of objectivity, and it shows how, on such an account, our knowledge of reasons can indeed be perceptual.

There are many different modalities of perception – visual, auditory, tactile, etc. -- and many interesting differences between these modalities.¹ Since the claims of this paper are not specific to any particular perceptual modality, however, I will (for the most part) use “seeing” and “perceiving” interchangeably.

Throughout, I assume that we can only perceive what is real or actual (although we can certainly *seem* to perceive much that is not real); I cannot perceive a flower in the fountain there unless there really is a flower in the fountain, and I cannot perceive that you are angry because you were pushed unless you actually were pushed. In this fairly minimal sense, perceiving something requires us to get things right. I remain neutral, however, on the question of whether the relevant sense of ‘getting things right’ depends on attributing properties or making a judgment. (It all depends on what one means by attributing or judging.) Thus, I also remain neutral on the question of whether perception is a state of knowledge that precedes judgment, a state of knowledge that already includes a judgment (or many judgments), or a state that only gives rise to knowledge with the addition of an appropriate judgment. For our purposes, it is enough to acknowledge that there is a difference (however fuzzy it may be at times) between perceptual knowledge and non-perceptual knowledge. This paper is an attempt to clarify that difference and to show how it applies to our knowledge of various sorts of reasons.

The immediacy of perception

Most would agree that the knowledge we gain through perception is more immediate than the knowledge we gain in other ways, and many have maintained that

¹ See, for example, discussions by O’Shaughnessy (1990), O’Callaghan (2007), and Keeley (2002).

such immediacy is definitive of perception. There is considerable dispute, however, over what constitutes the relevant sort of immediacy, and when it is present.

If it is the *experiential* immediacy of our knowledge that makes it perceptual – the fact that we are not *aware* of any other states of mind playing a mediating role, then any knowledge that *seems* to come “out of the blue” or “in a flash” will have the requisite sort of immediacy. If, for example, tomorrow I suddenly realize who was disguised as the clown at today’s party, without being aware of any mediating thoughts or memories, then my delayed recognition of the clown’s identity will have just as much experiential immediacy as my current perception of his red nose. If experiential immediacy suffices for perception, then, my delayed recognition of the clown’s identity qualifies as perceptual knowledge. Likewise, if the reason for the clown’s duck-like walk comes to me in a flash (“Aha! He walked like that because he wanted us to think he was shorter than he is”), and if experiential immediacy suffices for perception, then my knowledge of this reason is also perceptual. ²

Why should *experiential* immediacy be so important, though? Isn’t it possible that my knowledge of a particular state of affairs (Jack playing the clown) to be mediated by other knowledge (knowledge of the clown’s crooked smile, or knowledge of the clown’s high-pitched laugh) without my ever *realizing* its mediating role? Equally, isn’t it possible to recognize the mediating role of certain other knowledge without the resulting knowledge thereby being non-perceptual – as when other people’s observations prompt me to look more closely at the clown’s smile, leading me to see its distinctive shape? Questions such as these recommend an alternative understanding of the immediacy of perception – one that focuses on *justificatory* immediacy rather than experiential immediacy. For even though I do not consciously remember the clown’s smile or the clown’s voice, I may need to bring them into my awareness in order to justify my claim about his identity. And even though it was your remarks about his duck-like walk that made me notice it, I do not rely on your remark to justify my claim. Let us suppose, then, that the sort of immediacy that is relevant to perception is *justificatory* immediacy:

² Equating perception with experiential immediacy makes perception dependent on habit in some interesting ways. Not only will it be the case that some people will be able to perceive a bird where others see only leaves on account of what they are used to attending to; it will also be the case that some people will see a cardinal where others can only infer from its color and its crest that the bird they see is a cardinal – on account of their lack of practice in identifying birds. Or, to take the example of Brownian motion discussed by Brandom (1994), not only will some experts see movement where others see only a cloud of dust, but they will see mu mesons where the rest of us can only infer the presence of mu mesons from the darting dust particles that we see. The very same belief, resulting from the very same assumptions, will or will not constitute *perceptual* knowledge, then, depending on whether or not the relevant assumptions are entertained consciously or not (whether I am conscious of contemplating the cardinal’s crest and its color, for example, and whether crucial assumptions about Brownian motion are entertained consciously or not).

perceptual knowledge is immediate knowledge in the sense that it does not depend on any other knowledge for its justification.³

There are two, somewhat different, ways to understand justificatory immediacy. On the first reading, perceptual knowledge *cannot* be justified by reference to any other knowledge. On the second reading, perceptual knowledge *need not* be justified by reference to any other knowledge. Both of these readings are problematic. On the one hand, it seems that knowledge that is justified directly always *can* be justified indirectly as well. If I see a bird in a tree (and thus know, without further justification, that there is a bird in that tree), I *could* also provide indirect justification for my conviction by appealing to other observations (of color and movement, of falling debris, etc.) or by appealing to other people (with better vantage points, more expertise, etc.). On the other hand, the *need* to turn to other knowledge for justification seems to depend on contextual matters and not on the perceptual versus non-perceptual character of one's knowledge. In certain contexts (where the local birds have been poisoned, or where the light is dim), I may need to invoke other knowledge in order to justify my claim that there is a bird in the tree (knowledge that I am an experienced birdwatcher, knowledge that I have excellent night vision). And in many contexts, there is no need to appeal to any other knowledge for justification despite the fact that my knowledge is not perceptual (my knowledge of where I was born, for example, or my knowledge of where Greenland is located).⁴ So the attempt to understand the special immediacy of perception as justificatory immediacy appears to flounder on either reading.⁵

³ Note that if one holds that the relevant sort of mediation must be inferential (and thus justificatory) and that there is no inference without consciousness (a thesis defended by Searle (1992), for example), then phenomenal immediacy will be entailed by justificatory immediacy.

⁴ There is a rich and recent literature that distinguishes between conditions that entitle one to claim knowledge and considerations that justify that claim. See, especially, Burge (2003). It is not clear to me, however, why entitlement is not simply a subset of justification – supporting an argument of the form: S's judgment that p constitutes knowledge under conditions C. Conditions C obtain. Therefore, S knows p. But even if this is wrong, my other objection stands: non-perceptual knowledge does not always need further justification either, so not needing further justification cannot be what makes knowledge perceptual.

Note how these points cut across disputes between internalists and externalists regarding justification – i.e. disputes about who can or must offer the justification. According to internalists, perceptual knowledge requires knowledge of justification or entitlement on the part of the perceiver; according to externalists, it is enough that there is justification or entitlement, whether recognized or not.

⁵ A third reading of justificatory immediacy might distinguish between claims for which one *need not be able* to offer further justification in order to have knowledge and claims for which one *must be able* to offer further justification. I am always wary of analyses that rely on capacities, given the difficulty of determining just when one does or does not have the requisite capacity. But, in any case, the stated divide seems to mark a divide

One could, of course, accept that the need for justification depends on context and maintain that the perceptual versus non-perceptual character of our knowledge varies accordingly. When no further justification is needed, our knowledge is perceptual; when further justification is needed, it is not perceptual. This would result in what I would call a “deflationary” understanding of perception. A deflationary understanding of a category like perception explicates what the category is by detailing how and why we *use* it as we do, resisting the urge to explain our usage by appeal to its underlying *referent*. (The deflationist will not deny that there is a referent, but will ‘deflate’ its role, insisting that it cannot *explain* our use but can only be reflective of that use.⁶) So if the purpose of categorizing certain experiences as perceptions is simply to discharge the need for any further justification, then it is not surprising that the same phenomenal states could count as a perception in one context but fail to count as a perception in another; the call for further justification will be appropriate in some contexts but not others. In the absence of a more explanatory account of perceptual immediacy, the deflationary response can look very attractive.

The best (and perhaps only) counter to the deflationary position is the development of a more substantive theory of perception – one that describes the intrinsic character of perceptual experiences in such a way as to explain our use of the category “perception” in the justificatory contexts, one that not only describes our use of the term “perception” but also seeks to ground our use of the term in something more fundamental. In what follows, I attempt to offer such a theory – one that *explains* the importance of both experiential immediacy (of a certain sort) and justificatory immediacy (of a certain sort) in terms of something more fundamental -- namely, the experience of objectivity. Then, drawing on this account of perception, I indicate how it is indeed possible to see various sorts of reasons.

Self-evidence and the experience of objectivity

It has sometimes been suggested that the truth of our perceptions is self-evident. If this were so, then the justificatory immediacy of perception would be a part of the phenomenology of perception and not merely a contextual matter of when and whether a demand for further justification is legitimate. Even self-justifying experiences may be doubted and further justification may sometimes be needed; but the self-evident character of perceptual knowledge would distinguish it from non-perceptual knowledge that is not self-evidently true (even if it needs no further justification).

To better understand what the self-evident character of perceptions might amount to, consider a case where knowledge has both experiential immediacy and justificatory immediacy but lacks something that perception has – a case of ‘blind’ conviction:

between internalists and externalists about knowledge in general rather than between perceptual and non-perceptual types of knowledge.

⁶ Deflationists describe the rules we follow but reject any attempt to establish foundations or justifications for following those rules. See, for example McDowell (1984).

Five-year old Sam regularly becomes convinced that a particular ant is in the room; and Sam is almost always right. Asked to give a reason for his conviction, Sam says “I don’t know; I just thought it was here again”. Asked whether there was any special sound or sensation that led to his conviction, he answers “no”. A series of experiments is conducted, and they establish Sam’s reliability without, however, suggesting any explanation for it: Sam’s location relative to the ant seems irrelevant (he can be sitting on the chair under which the ant is crawling, or he can be facing in the opposite direction); covering Sam’s eyes and ears makes no difference, and he reliably reports that this particular ant is in the room even when it is enclosed in a lead box. Eventually, everyone (including Sam) is convinced that he does know, in some mysterious way, when that ant is present, and they trust his reports.

This seems to be a case of knowledge that has experiential immediacy and justificatory immediacy without, however, being a case of perception. Sam’s knowledge comes to him directly, without relying (consciously, anyway) on any other knowledge, and once Sam is established as a reliable reporter of the ant, there is no need to justify his claims by appeal to any other facts. Yet Sam’s knowledge comes to him in the form of a ‘blind’ conviction rather than in the form of a perception. Just what, though, is lacking in this case? Just what does it mean to call his knowledge ‘blind’?

It is common to suppose that perceptions must be acquired through our senses – that states of affairs that do not impact on our sense organs, or that do not prompt sensations of some sort, cannot be perceived.⁷ Perhaps it is because Sam’s information is not received through his eyes or ears or nose or skin, or because he denies that there is any particular sensation that accompanies his knowledge, that his knowledge is not perceptual knowledge. This suggestion gets things the wrong way around, however. For we identify the organs of perception after we ascertain the fact of perception, not before. If we were to discover that the center of Sam’s brain is sensitive to a wavelength that is activated by the nervous system of a particular ant, that would not tell us that he does indeed perceive the ant; it would only serve to demystify his non-perceptual knowledge. And if my current experience of the act were not dependent on my eyes, that would not establish that my knowledge is non-perceptual; it would only show that some other sense organ is operative. Likewise, the notion of a sensation is derivative from the notion of a perception, not *vice versa*. The sensation of loudness, for example, does not make sense apart from the perception of loudness any more than the notion of seeming to be loud (or seeming to be red, etc.) makes sense apart from knowing to be loud (or knowing to be red, etc.). Even if Sam’s conviction that the ant is present were reliably accompanied by a distinctive sensation, it would only count as a sensation of the ant (or of a brain

⁷ Goldie (forthcoming), for example, stipulates two conditions on perception – the first of which requires the involvement of sense modalities: “I will define a non-inferential perceptual belief (and, *mutatis mutandis*, judgment) as (i) a belief that something has a certain property, from the way it appears, relative to one or more sense modality; and (ii) this belief arises in a way that is experientially immediate; in other words it must not be the product of a conscious process of inference”. That sense modalities require sense organs seems clear from his statement that mathematical knowledge cannot meet this first requirement.

disturbance) if we were independently convinced that Sam can perceive the ant (or the brain disturbance).⁸

It is also common to suppose that perceptual knowledge must be knowledge that *that* ant is in the room as opposed to knowledge that *some* ant is in the room. Perhaps it is the fact that Sam's conviction is sensitive to any ant that emits the relevant wave length that prevents his knowledge from being perceptual – the fact that Sam is insensitive to the possibility of duplicate ants. This points in the direction of something right but, again, it gets things the wrong way around. For it is perceptions that secure the particularity of the objects of our thoughts, not the other way around. Even if Sam were known to be sensitive to the presence of just one particular ant – his internal sensor somehow glued to that ant alone, disregarding duplicate ants in other locations, and even if, knowing this, Sam was able to have demonstrative thoughts about *that* ant, his knowledge would not be perceptual. Rather, it is because of a prior perception of the ant (by Sam, or by others to whom he defers) that Sam is able to have a demonstrative thought about that particular ant.⁹

Michael Ayers, in his book on Locke, considers some imaginary cases that are similar to our case of Sam (as well as the case of innate knowledge, which is of particular interest to Locke). Ayers maintains that, in contrast to such cases, “when I believe my eyes, there is no mystery about the source and basis of my belief.”¹⁰ But even if we were to remove the mystery, discovering that Sam's conviction is a result of an unusual neural sensor and belief generator in his brain, I don't think we would consider it perception – because the truth of his conviction is not evident from *within* the experiences that comprise, and give rise to, that conviction. (That, I take it, is what Locke meant by the “perspicuity and evidence” of perception.)¹¹

What is missing from Sam's experience, I suggest -- what (literal) perception has that Sam's knowledge lacks -- is the experience of objectivity. An experience is a perceptual experience precisely when the independent reality of its object is evident from *within* that experience; perceptual experience is self-justifying because it includes the experience of its own objectivity. This is also (at least part of) what it means for perception to be “transparent”: it presents its object in such a way that one is unaware of any distinction between how an object appears and how it is in fact. Sam knows that a particular ant is in the room, and he may also know that he knows, but the objectivity of

⁸ There are many and varied disputes about the nature of sensations (and related disputes about qualia) – about their intentionality, about their physicality, about their epistemic role. What I say here follows the lead of Sellars (1956), among others. The onus is on those with a different view of sensations to show how it helps distinguish perceptual from non-perceptual knowledge.

⁹ See Brewer (1998). Perception is needed to determine spatial location (relative to the subject making the reference), and only spatial location picks out particulars.

¹⁰ Ayers (1993) pp. 171-2.

¹¹ See, also, work Snowden (1979-80,1990) on the self-verifying character of perception.

the ant is not evident from within his experience. (Again: the fact that an experience has a self-evident, or transparent, character does not guarantee its accuracy; not every perception-like experience is, in fact a case of perceptual knowledge. I am not trying to defend the epistemic standing of perceptions or perception-like experiences, only to clarify how they differ from non-perceptual forms of knowledge.) It is not enough, then, for Sam's knowledge to come 'in a flash', without any conscious reliance on inference; for it to be a case of literal perception, he must experience the fact itself, and not just his knowledge of it, as arriving without phenomenal intermediaries. Likewise, it is not enough for Sam to be entitled to his conviction without the need for further justification; for it to be a case of literal perception, the justification for his conviction must be integral to his experience of knowing. In the end, it is the experience of objectivity – an experience that includes a particular sort of experiential immediacy (transparency) and a particular sort of justificatory immediacy (self-justification) – that lies at the heart of perception.

Little progress will have been made, however, unless we can say something more about what is involved in the experience of objectivity. It is fine to claim that the sorts of experiential and justificatory immediacy that are relevant to perception are the sorts that come with the experience of objectivity; but this only puts us closer to a substantive account of perception insofar as a further explication of the experience of objectivity is possible. For a state of affairs to *be* objective, it must exist or obtain independently of how (and whether) it is experienced. To *experience* a state of affairs *as* objective, then, must be to experience it as existing independently of how and whether it is experienced. But how is this possible?

It is easy to how we might experience one *state of affairs* as independent of various other states of affairs if we observe that it remains constant despite changes in those other states of affairs. And it is easy to see how we might experience one *experience* as independent of various other experiences if we notice that it remains constant despite changes in those other experiences. But how can we separate the object of our experience and the experience itself sufficiently to recognize the constancy of the first despite changes in the second? We can, of course, stand back from our experiences enough to notice that the accuracy of some of them is inconsistent with the accuracy of others; and assuming that the world is not inconsistent, we may rightly conclude that the world is independent of our experiences. But this will, at best, tell us that there is *some* independent state of affairs; it won't give us an *experience* of any particular state of affairs as objective. Furthermore, knowledge that is reflective in this way seems already to presuppose a less reflective sort of knowledge – perceptual knowledge, which is precisely what we are seeking to understand.

Alva Noë's account of perception offers a revealing (though finally unsatisfactory) response to this challenge. He distinguishes two aspects or components of perception and argues that experiencing the dependence of one aspect on the other amounts to an experience of objectivity. According to Noë, the contents of perception include a factual dimension and a perspectival dimension, an aspect that indicates how things are and an aspect that indicates how things appear.

“We experience not only how things are, but also how they look from here. We experience that the plate is round and that it looks elliptical from here. Its elliptical look from here is a genuine property of the plate – we see the shape and we see the perspectival shape from here – but it is also a relational property, one that depends on where ‘here’ is.”

“Perceptual content ... is two-dimensional. It can vary along a *factual* dimension, in regard to how things *are*. And it can vary along a *perspectival* dimension, in regard to how things *look* from the vantage point of the perceiver. Visual experience always has both these dimensions of content.”¹²

Noë treats instances of immediate but non-perceptual knowledge (as in our case of Sam) as cases in which our experience has factual content (the ant is in the room) but not perspectival content (how it looks from here), contents which are dependent on the facts they pick out but not on the perspective from which those facts are known. Perception, he argues, requires both sorts of content. And whereas factual content’s dependence on the world (its truth) must be discovered from an external point of view, we can discover the dependence of perspectival content on factual content by moving through space and noticing how the way things look depends both on how they are and on where we are.¹³ Indeed, we only become capable of distinguishing looks from facts insofar as we become capable of tracking such dependencies.

While I agree with the suggestion that we experience objectivity only when we discover constancy across perspectival change, Noë’s insistence that we experience appearances (e.g. the elliptical look of a plate) as well as facts (e.g. its roundness) creates more contents (and more puzzles) than are needed.¹⁴ When we look at a plate from the side, we do not see an elliptical

¹² (2003) pp. 2-3. Elsewhere (2002), p. 74, Noë makes a similar distinction between the representational and sensory factors that make up the qualitative character of experience: “The qualitative character of experience, as we have seen, depends on two factors. First, it depends on the qualities that we experience (e.g. looks, sounds, etc). This is a representational feature. Second, it depends on the character of the activity in which the temporally extended activity may consist. So, for example the fact that we do not make eye movements when we explore the environment haptically makes a difference to what it is like to touch. These differences in the sensorimotor contingencies governing the different sensory modalities are differences in the qualitative character of experience that do not correspond, directly at least, to differences in what is perceived.”

It is not clear (to me, anyway) whether Noë intends the second, non-representational aspect of experience to count as part of its content – akin to the perspectival contents discussed above. The term “content” doesn’t matter, perhaps, but the ability to recognize the dependence of perception on fact, through the dependence of one factor on another, does.

¹³ The question of whether these movements must be actively initiated or whether mere movement with respect to an object (which may be initiated by the object rather than the viewer) is left open.

¹⁴ His talk of “dimensions” of content is less problematic than his talk of experiencing “looks” as well as facts. And he is clearly aware of the dangers of falling into a sense data type account of perception. Nonetheless, he is determined to preserve what seems right

appearance of the place; we either see the round plate or we mistake the round plate for something elliptical. And when we experience the shapes around us as if they were flattened onto a screen before our eyes, we do not attend to an already present aspect of experience (as Noë suggests); we alter the content of our experience. We do not see what is variable (the plate's "appearances"), and recognize its dependence on the invariant (the plate's actual shape); rather, we see what is invariant (its shape) from various different perspectives as we move about it. The variation is in perspective or in modes of access rather than in the contents of our experiences; the experience of objectivity depends on being able to negotiate different ways of accessing the same content, not on the dependence of some contents (or some aspects of content) on others. I agree with Noë's ultimate reliance on the coordination of abilities to explain perception, but I disagree with his reliance on a second type of content, or a second aspect of content, that is perspectival. The convergence of abilities (which do not themselves need to be the objects of some further experience) on the same content is enough for the experience of objectivity. Indeed, the very specification of content depends on such a convergence – i.e. it depends on our ability to move between different ways of accessing the same fact.

Consider the perception of a table tilting relative to the floor. I may acquire knowledge of this particular state of affairs both by looking at it and by touching it, by looking at it from one side and from another, from close up and from a distance. Insofar as I use these different perspectives and modalities to check one another, and expect them to agree, then I take them to be accessing the same state of affairs – the table tilting. And, equally, insofar as I take them to access the same state of affairs, I use these different perspectives and modalities to check one another and I expect them to agree. The relationship between using one judgment to check another, on the one hand, and taking both judgments to be about something objective, on the other, is one of mutual dependence; neither one has causal or explanatory priority over the other. That a table's tilting can be known in any number of ways ensures that it is independent of any particular 'pathway' between our knowledge and its objects; and the independence of the table's tilting ensures that it can be known in several different ways. ¹⁵

Imagining Alternatives

To what extent is the plurality of ways in which we can know a state of affairs, and thus the independence of that state of affairs, something that is evident from *within* a given experience? If the experiences that are used to check one another occur separately, how can the

about the idea that we see appearances as well as facts, and that is something I want to deny. (One way to put the complaint: Noë's account abandons the transparency of perception.) Following Lewis and others, I maintain that how things appear is not part of the *content* of our perceptions; rather, it is a *way* in which we may come to know things. On this view, adding perspectives can add to the epistemic standing of one's judgment without adding content.

¹⁵ The importance of mastering multiple perspectives in order to fix the content of an experience is familiar from the work of Dretske (1981), Millikan (2000), Davidson (2001), and others. What is distinctive about Noë's analysis is his attempt to analyze the phenomenology of perception as the experience of objective reference.

commonality of their object be a part of any one perceptual experience? ¹⁶ While it is true that we can experience only a limited number of ways of knowing at a time, it is certainly possible simultaneously to touch and to look, to touch from two different angles, or to look from the slightly different angles of two separate eyes. It is also possible to imagine many different perspectives and modalities – perspectives that we do not now occupy, and modalities that we do not now engage – and the juxtaposition of those different perspectives and modalities in imagination can also create an experience of objectivity. As I look at the tilted table from here, I can also imagine looking at it from over there; and I can imagine what it would feel like to rest my hands on the table, or to lie on its surface; and so on. And each of these combinations can serve to secure our experience of the table as something objective. ¹⁷

It is important to realize that merely having the *capacity* to imagine alternative perspectives is not enough for an experience of objectivity. Various authors insist that perceivers be *able* to imagine other perspectives on an object, but it is not clear how this capacity could generate an experience of objectivity as opposed to the idea of objectivity. ¹⁸ I am claiming something much stronger – namely, that we *actively* imagine alternative perspectives whenever we experience something as an objective state of affairs. We are not usually attentive to our imagining of alternative perspectives when we are perceiving the objects before us, but there is good reason to suppose – on the basis of the phenomenology of perception as well as the arguments presented above – that such imagining is always present. For it cannot be denied that there is a difference between seeing an object as two-dimensional and seeing it as three-dimensional, and this difference is neither purely cognitive – a matter of merely endorsing the

¹⁶ Noë (2006), in contrast, claims “experience is a temporally extended phenomenon; it is an activity of skillful probing” (p.430) and “experiential presence is virtual all the way in ... The rear side is present virtually, but the present side is present *simpliciter*. Notice, however, that you do not, as a matter of fact, have the whole of the facing side of the tomatoes in consciousness all at once.” (p.427)

¹⁷ Note that the alternative perspectives I imagine must be fairly accurate in order to produce perceptual *knowledge*. Their accuracy is not secured by my current causal connection to the object, however, as I do not currently occupy any of the alternative perspectives (and perhaps never have). A recent article by Alan Miller (2007) offers an alternative to theories that merely require causal contact with objects of perception, and theories that disregard causal interaction as irrelevant. Miller argues that what matters for perception is the “exercise of the relevant capacities” – those that “prime one to form suitable judgments in suitable environments -- where relevant capacities do not count as exercised unless knowledge is acquired. I would suggest that the capacity to imagine alternative perspectives might be such a capacity.

¹⁸ Bill Brewer (1998), for example, claims that a perceiver “has materials to construct, in imagination, the systematically varying appearances of those things being just the way they are from various alternative, possible but non-actual, perspectives.”(pp. 23-4; see, also, his footnote 21)) What results, then, is said to be a “recognition” of objects as such. But “recognition” is an interestingly neutral term; it does not distinguish between the thought of, or belief in, objectivity and the experience of objectivity. I claim that it is only by actively imagining alternatives that an experience of objectivity is possible.

proposition that the object before us is a house with inner spaces and not simply a façade – nor purely dispositional – a mere readiness to imagine other perspectives if prompted. Rather, as Strawson suggestively remarks, “the actual occurrent perception of an enduring object ... is, as it were, soaked with or animated by, or infused with ... other past or possible perceptions of the same object.”¹⁹

To experience something as knowable in a plurality of ways, and thus to experience it as objective, is also to experience it as existing in space. This seems obvious in the case of multiple *perspectives* in so far as one thinks of visual perspectives and visual perspectives differ precisely when the spatial position of the perceiver differs. This is true for other sensory modalities as well, however. We hear sounds as objective, at least in part, because we know how to track the same sounds from different spatial positions; we tilt our heads from side to side in order to get a different perspective on the same sound – and, indeed, in order to determine the objectivity of a sound. We also imagine hearing the sound from a closer location, or from a location farther from other, distracting noises; and it is these imaginings that give us a clearer sense of just what the sound is, objectively. Likewise with touching and tasting and smelling; by moving through different spatial positions (both actually and imaginatively) while tracking the same texture or flavor or odor, we come to experience the objectivity of those qualities.

Less obvious is the way that distinguishing between modalities of experience – sight versus touch versus hearing, and so on – depends on spatial distinctions. For there is no way to distinguish between different *types* of sensory qualities (smells, tastes, sounds, etc.) apart from their relation to different *types* of sensory organs (noses, mouths, ears), and there is no way of distinguishing between these different types of sensory organs except by reference to the different *types* of spatiotemporal routes by which they give us knowledge. There is no one-to-one correlation between sensory qualities and sensory organs; for each of our sense organs gives us knowledge of many different sensory qualities (we see color and brightness and sheen, we taste sweetness and bitterness and sourness, we hear pitch and timbre and volume), and many sensory qualities are perceived through more than one sense organ (we can both see and touch the smoothness of a tomato, we can both feel and hear the beat of a drum, we can both taste and smell the sourness of a lemon). Indeed, we wouldn’t speak of different sensory “modalities” at all unless there were more than one way of perceiving the same quality in an object. Nor would we speak of different sensory modalities merely on the basis of the different locations of our sensory organs (receptors for touch are located throughout our body, after all, and seeing would still be distinct from hearing even if our eyes were located inside our ears). What makes one sensory modality different from another is not the uniqueness of what is sensed or the location of the sensory organ, then; what does seem to distinguish one sensory modality from another is, rather, differences in the enabling conditions for one versus the other (due to differences in the causal pathways that support each). Sight but not hearing requires intervening light; touch but not smell requires an absence of intervening surfaces; taste but not hearing requires close proximity; and so on. Our bodily recognition of these differences is evident in the ways we negotiate our environment to improve our sight, our hearing, our touch, our smell. When encountering a new environment or trying to learn more about one that we are already familiar

¹⁹ Strawson (1974), p. 53. For further discussion of these and related issues, see Church () and (2008).

with, we may move to the side in order to hear better, move forward to smell better, move away in order to get a better view; different sensory modalities demand different spatial maneuvering.

²⁰ We distinguish between sensory modalities, then, not (in the first instance, anyway) on the basis of differences in what is sensed or on the basis of differences in the bodily locations of the sensing organs but, rather, on the basis of the different causal pathways – through space – that connect the two.

Returning, then, to the case of Sam: Sam does not perceive the ant in the room, nor does he perceive the presence of special emissions given off by the ant (even if we determine that his reliable predictions are caused by his special sensitivity to such emissions), because Sam does not experience the objectivity of the ant (or the emissions). He may come to know that certain convictions (and, perhaps, certain sensations) are reliable indicators of an objective state of affairs, but their objectivity is not evident from within the experience itself. In order for Sam to experience the objectivity of the ant, he must experience its constancy across multiple perspectives and multiple modalities; and in order to experience the fact of such constancy at a given moment, he must experience the ant as located at the intersection of various imagined perspectives or modalities. Suppose, for example, that Sam came to realize that the strength of his convictions changed depending on how close he was to the walls of a room and depending on how deeply he filled his lungs; and suppose that he became adept at knowing just what changes in position and in breath resulted in just what changes in conviction. He could then check the predictive accuracy of a particular conviction by varying his position relative to the walls and by varying his breath to see whether his convictions changed in the appropriate ways; and he could experience the ant's presence as objective in so far as the memory and imagination of these other perspectives was a part of his current experience. In which case, Sam's experience *would* be a case of perception.

I am not claiming that merely imagining multiple perspectives on an object (even if those imaginings are accurate, and even if they are integrated with each other) suffices for perception. Whether we experience something as an objective presence will depend on just how (and how well) we relate our imaginings to what we take to be our own position in space and time. (Whether it *is* in fact an objective presence, and in fact a perception, will also depend on whether we are right about our own position.) I am, however, claiming that it is the limits of our imaginations and not the limits of our sense organs (except insofar as accurate imagining depends on the possibility of accurate sensing) that determine what we can or cannot perceive. As long as we have a reasonably accurate understanding of the spatiotemporal position from which we are experiencing an object, and a reasonably accurate understanding of how it is related to other possible perspectives on that same thing, there is no *logical* limit on the size or

²⁰ Sean Kelly (2005) has nicely documented the way in which perception implicitly relies on the recognition of ideal viewing positions, and movements that approach that ideal. Brian O'Shaughnessy (1990) makes interesting distinctions between visual properties that 'attach' to objects versus sounds and smells, e.g., that traverse or pervade a space. It is interesting that Dominic McIver Lopes (2000) insists that different sensory modalities have different phenomenologies, but his descriptions of these differences are largely dependent on their different spatial functioning.

the distance of a perceived object. While I have not said much about how we identify one particular perspective as our actual perspective, I doubt that it has much to do with special sensations, extra vividness, or greater (or lesser) control. Rather, I suspect that the perspective we take to be our actual perspective is the perspective from which most of our experiences can be integrated. (If my imagined position in a dream life were better integrated with the totality of my experiences than my actual position in my waking life, I would identify my actual perspectives with what is, in fact, a merely imagined perspective.) If this is right, I do not determine our present position and then fit my experiences around it so much as I determine my present position by determining how best to fit my experiences together.

There is not a sharp line between what is perceived and what is not perceived, between what is experienced as objective and what is not. This is true for a number of reasons. First, the *number* of different perspectives that we are able to remember and project -- the different perspectives that give us a sense of a thing's objectivity -- varies widely across circumstances. I easily imagine multiple perspectives on the desk before me but I struggle to imagine even one other perspective on the dark fleck that has just crossed my visual field; and, for that reason, the former is more clearly an instance of perception. Second, our ability to *coordinate* the different perspectives we remember or imagine -- with each other, and with the perspective we currently occupy -- varies quite widely. Even if you show me photos of the dust that crossed my visual field -- photos that enable me to imagine what it looks like from many different perspectives, this will only enhance my perception of that dust insofar as I am able to relate those other perspectives to my present perspective and to each other; I must have a sense of how the different imagined perspectives fit together in this world, and just how I would now have to position myself differently in order to occupy the other, imagined perspectives. The more coordinated the imagined alternatives, the more perceptual the experience. Third, there are differences in the extent to which we can imagine different perspectives *simultaneously*. Even if I come to understand just how my current perspective on the dust is related to the perspectives of the many photos you show me, it may be harder for me to conjure up those other images while observing the dust than it is for me actively to imagine other perspectives while observing my desk. Imagining alternatives simultaneously is the temporal counterpart to imagining the coordination of perspectives in space; and both are needed for the experience of objectivity. Given the close connection, then, between our perception of a state of affairs and our imagination of alternative perspectives on that state of affairs, some experiences will count more clearly than others as cases of perception. As long as our intuitions about clear versus unclear cases accord with the predictions of this analysis, the resultant fuzziness counts for rather than against the analysis offered here.

Perceiving why

So far, we have been considering the perception of particular objects or states of affairs -- the ant in a room, the desk's tilt relative to the floor, a bit of dust drifting before my eyes, and so on. To perceive the *reason* for a state of affairs -- the reason the ant is in the room, the reason the desk is tilted as it is -- requires something more, however; it requires us to see what *explains* or *justifies* that particular state of affairs, and it requires us to see it *as* explaining or justifying. In this section, I discuss the possibility of

perceiving the *explanatory* reasons for various states of affairs. The possibility of seeing *justificatory* reasons is addressed in the next section.

One of the most common forms of explanation is, of course, a causal explanation.²¹ Assuming that causal relations are real rather than illusory, and external rather than internal²², we can still wonder whether causal relations are the kind of relations that can be perceived. Clearly, we can perceive at least some spatial relations, as when we perceive that the cat is on the mat; and we can perceive at least some temporal relations, as when we perceive that the cat stepped onto the mat before it lay down. But, as Hume pointed out, causation involves something more than spatiotemporal continuity. Can we perceive the fact that the cat *moved* the mat, or that its movements *influenced* the mat's position? As in the case of Sam's judgments about ants, it is not enough that our judgments about causality are immediate and reliable; in order to *perceive* a causal relation, we must experience the objectivity of the causal relation itself.²³ But how is that possible? The objectivity of an ant may be experienced through the juxtaposition of multiple spatial perspectives, but multiple perspectives on the moving cat and the moving mat can only show us their objectivity, not the objectivity of the causal relation between them. More needs to be said, then, about what a causal relation *is* before we can decide whether causal relations are perceivable.

Consider three different analyses of what the causal relation is: (1) An event C causes an event E just in case C-type events and E-type events are regularly conjoined in space and time (with C-type events preceding E-type events).²⁴ (2) C causes E just in

²¹ Aristotle describes four different sorts of causes or reasons – material cause, formal cause, efficient cause, and final cause. In keeping with modern usage, I use the term “cause” for efficient causes only. But note that what I consider what I call “constitutive reasons” and “justificatory reasons”, below, are roughly aligned with Aristotle's formal causes and final causes, respectively.

²² If causal relations are illusory, then they cannot be perceived. In Michotte's (1963) famous experiments with moving shapes, one projected shape appears to cause the other to move, even though it is not in fact the cause; and viewers seem to see a causal relation between the two, even though they do not. If causal relations are internal – relations between subjective impressions, as Hume is often read as claiming – then we do not perceive causal relations between objects, but only (at best) within ourselves. See Costa (1989) for a nice review of the current state of debate on Hume's view.

²³ Dretske (1969) introduces a notion of “secondary seeing” to cover cases where seeing p (in the primary sense of seeing, where things would not look the way they do unless they were so) plus a belief that p entails q leads one to believe q without, however, inferring q from p. This notion is supposed to cover such cases seeing that a screwdriver is magnetized, or seeing that a cake is not yet done given the toothpick result. These are not cases of seeing a causal relation, though – of seeing that the magnetization causes the screw to jump, or seeing that the underdone cake causes the toothpick to stick.

²⁴ This is, basically, the Humean view – on a realist reading of Hume, which is controversial.

case there is an appropriate transfer of energy from C to E.²⁵ (3) C causes E just in case E would not have happened if C had not happened.²⁶ Not all causal relations must be of the same type, of course, and not all causal relations must be perceivable; but the very possibility of perceiving causality, on any of the above analyses, can seem doubtful. For on analysis (1), the possibility of perceiving causal relations depends on perceiving regularities, and regularities extend across long periods of time so cannot be perceived in a single moment. On analysis (2), perceiving a causal relation depends on perceiving a transfer of energy, but the transfer of energy occurs at the level of subatomic particles and forces, and they can't be perceived through the senses. And on analysis (3), perceiving a causal relation seems to depend on perceiving a non-actual possibility, whereas we can only perceive what is actual.

Helen Beebee has argued that regularity analyses of causation (type (1) above) do not preclude the possibility of perceiving causes because our current perceptions are often shaped by regularities in our past experiences (or by the reports of others). We now see the cat causing the rug to move – we see the rug's movement as following with a kind of inevitability – because we have seen so many similar sequences in the past. Furthermore, the fact that we now experience the sequence as inevitable gives us good reason to believe that there are similar sequences of events both in the past and in the future.²⁷ Others have argued that we perceive our own power to put our intentions into action as a kind of transfer of energy (type (2) above)²⁸, and/or that we can feel (through our fingers, for example) the transfer of energy from one object to another. If this is right, and if a background of such experiences shapes our observation of more distant movements as well, then we may succeed in perceiving causation in these cases as well. I do not think that either of these accounts accurately describes most experiences of

I assume that events are the relata; but nothing turns on this. The relata could also be states or facts, for example.

²⁵ This is another traditional account, discussed in Collins, *et al* (2004) pp. 12-15. The details of what energy transfers are “appropriate” do not matter here.

²⁶ This sort of analysis currently dominates the field, largely due to the work David Lewis (1986). The relevant counterfactual, as I state it, concerns individual events. But it could also be stated as a counterfactual about types of events (which would make it into a law-based notion of cause).

²⁷ Beebee (2003) Note that she does not attempt to defeat the skeptic about induction. Beebee considers Peacocke's (1986) claim that we can't perceive a cause because a single perception cannot meet the rationality requirement for knowledge – i.e. acquisition via a method that is appropriate to “canonical acceptance conditions” and “the canonical acceptance conditions for the holding of a causal relation ... have to do, ultimately, with laws.” As I understand her response, she views the background conditions that enable us to experience one event as causing another as fulfilling the canonical acceptance conditions. I am less clear about her claim that a single experience – e.g. of pain upon touching a fire – can provide a rational basis for a causal judgment, on a regularity analysis of cause.

²⁸ See Reid (1792/2001).

causation, however. Furthermore, given the prevalence of counterfactual analyses of causation (type (3) above), it is important to determine whether causation so understood could be perceived or not.

Recall that the problem with perceiving causes, analyzed counterfactually, was the non-existence of the relevant counterfactual possibilities. If C causes E just in case E would not have occurred if C had not occurred, seeing C cause E seems to require that we see the non-actual case in which C does not occur; and that seems absurd. This argument overlooks the crucial role of imagination in perception, however. Earlier I argued that perceiving anything at all requires us to imagine multiple ways of knowing that thing – different angles from which we might view the tilt of a table in space and/or multiple modalities through which we might discern its position. Otherwise, we don't experience it as an object in space at all. But most of these perspectives or modalities are non-actual for the perceiver at the time of perception: if I am looking at a table off to the side and at a distance, I can only imagine seeing it straight on and close up; and I can only imagine feeling its surface. And yet, I have argued, my experience must be informed and “infused” by these merely imagining alternatives in order for me to experience the table and its properties as objective – i.e. in order for me to perceive them at all.²⁹ Likewise, I suggest, the non-actual alternatives that constitute certain relations as causal relations are alternatives that, when actively imagined, can inform and infuse our current experience in such a way as to make causality perceivable.

Consider, for example, what happens as we watch a paddle move through the water. The paddle moves back and the water curls to the side; the paddle moves deeper and a small funnel is formed in the water; the paddle lifts up and drops of water follow it into the air; as it reenters the water, there is a sharp slapping sound; and as it is pulled back, there is a lower sucking sound. We can observe these things as a disconnected array of events but, more typically, we observe them as causally connected – the movement of the paddle causing the curling of the water, the reentry of the paddle causing the slapping sound, and so on. Typically, what enables us to see the causal connections, I suggest, is our simultaneous imagining of alternatives in which the paddle remains still, alternatives in which the paddle moves forward rather than back, alternatives in which the paddle remains in the water, and so on -- alternatives that show us what depends on what by showing us what varies with what. What I am claiming is a certain parallel between the imagining that is required in order to perceive objects as objects and the imagining that is required in order to perceive causal relations as causal relations. In the case of seeing an object, I imagine alternative perspectives on it; in the case of seeing causal relations, I imagine alternative arrangements of objects and events. Put another way, in the first case I imagine different relations between myself and a particular state of affairs, and in the latter I imagine different relations between one state of affairs and various others.

²⁹ The term “infusion” is Strawson’s (1974). Similarly, Wittgenstein (1980) writes: “It is as if one had brought a concept to what one sees, and one now uses the concept along with the thing. It is itself hardly visible and yet it spreads an ordering veil over the objects.” (§483)

What are the non-actual alternatives that must be imagined when we perceive one thing as a cause of the other? To support the judgment that the water would not have swirled without the movement of the paddle, we must consider alternatives in which the paddle does not move. There are many alternatives in which the water swirls even though the paddle stays fixed – because there is an undercurrent, because there is a propeller nearby, and so on -- but these are not the alternatives that enable us to see the spark as the cause. The alternatives that must be imagined when perceiving one event as a cause of another are alternatives in which various other conditions remain constant; if there is not an undercurrent or a nearby propeller in the actual situation, then the imagined alternatives should not include these. (*Ceteris paribus* clauses added to causal laws indicate the need to fix conditions in this way, but it is never possible to provide an exhaustive list of what the relevant conditions are.³⁰) Insofar as many different factors contribute to a given event, we can perceive the causal power of each only by imagining alternative in which most other factors remain constant. What counts as *the* cause may depend, at least in part, on the particular interests around which our imagination turns. (If we are interested in predicting future fires, we will focus on the leaves; if we are interested in holding someone responsible, we will focus on the spark.) On the other hand, insofar as other conditions are already sufficient for the event in question (i.e. insofar as there is causal over-determination), the factor that we imaginatively subtract will not be perceived as a cause.

When we perceive one event as the cause of another, the distinction between background conditions and cause is not as neat as this suggests, however. When we perceive the swing of the hammer as the cause of the moving nail, not only do we imagine what would happen if the hammer were not swung; we also imagine what would happen if the hammer were swung in some other way – with less force, at a different angle, and so on. Similarly, we do not imagine all other conditions as fixed; for as we imagine the hammer swinging with greater force, we also imagine swinging it further back, and holding it a bit differently, and so on. The imagining that supports the perception of causes requires us to entertain many different possibilities, more or less simultaneously; for only then do the relevant dependencies become apparent to us. (This is, again, parallel to the imagining of multiple perspectives that makes state of affairs seem objective.) Furthermore, there will be cases in which we fix the background conditions and discover that although the presence of one event is not necessary for another, it makes it much more probable – in which case, if it is perceived at all, the causal relation will appear as a somewhat weaker force or tendency, the effectiveness of which may depend on background conditions of which we remain unaware.³¹ Ideally,

³⁰ In the words of Yablo (2004), C is the cause of E if and only if had C not occurred “and had suitably chosen other factors been held fixed”, then E could not have happened.

³¹ This may be the phenomenological counterpart of seeing causes on a probabilistic account of causal relations. Suppes (1970) is a classic defense of probabilistic theories of causation; Glymour (1998) and Reed (1998) are informative on the complex relations between causation and probability.

we fix background factors fully enough to reveal the cause as both necessary and sufficient for the effect; but, of course, this ideal is never wholly realized.

Given the amount of work that is being attributed to the imagination in this account, and given the fact that we are often unaware of imagining the alternatives in question, it is worth reviewing some of the considerations in favor of this hypothesis. The role of imagination in the perception of both objects and causal relations can be demonstrated in a couple of ways – an empirical way, and a transcendental way. Introspectively, we can transform our ordinary experiences of objects into experiences of patches of light and color (as though spread across a canvas, for example) by ceasing to imagine alternative perspectives (the back side of the house, the close up view of a tree, and so on); and we can transform our experiences of causal relations by blocking the infusion of non-actual alternatives into our experience (alternatives in which the paddle moves in some other way, or alternatives in which the paddle is entirely absent). Such exercises help to demonstrate the essential role of the imagination in ordinary perception.³² There is also third-person empirical evidence that we are pretty much always imagining alternative perspectives and alternative arrangements of what we see insofar as we are clearly primed for some developments and surprised by others. We are surprised to discover that what we took to be a building is merely the backdrop for a movie set because we were already imagining alternative perspectives; and we are primed to intervene in some ways rather than others because we have already imagined alternative possibilities that show us what depends on what.³³ The central argument that I have been offering, however, is, first, a transcendental argument to the effect that it is *only* by imagining alternative perspectives and imagining non-actual possibilities that we *could* perceive states of affairs, because only then could the objectivity of what is seen be evident from *within* perception. Assuming that this is right, and assuming that general skepticism about the objectivity of our experiences is not a tenable position, it follows that we *must* be imagining such alternatives. The second part of the argument, then, is an argument to the effect that a very similar sort of imagining can explain the possibility of perceiving causal relations. So while it may not be incoherent to deny that we perceive causes (in the way that it is incoherent to deny that we see states of affairs), the similarity

³² Some have used this sort of evidence, along with evidence of extensive “filling in” during visual saccades, to argue that we actually see very little of the world, that most of what we see is virtual rather than actual. This is a pervasive theme in Akins (1996). This conclusion leads to some very counterintuitive results, however. Rather than claiming that most of what we see is merely imagined, I claim that most of our seeing is heavily dependent on imagining.

³³ Much empirical work in psychology assumes that (a) quicker responses indicate that certain representations are already formed and unconsciously 'present', and (b) the experience of surprise also indicates that certain representations are already formed and unconsciously 'present'. Other third-person empirical evidence for the role of imagination in perception comes from brain scans that indicate that activation of the visual cortex must added to retinal activation to produce perception. Church (2008) for further discussion of relevant empirical evidence.

in the imagining required for each gives us no reason to doubt the accuracy of ordinary claims to have seen what caused what.³⁴

Not all explanatory reasons are causes, of course. Some explanations invoke the constituents rather than the antecedents of a state as the reasons for that state – as when we explain the shape and hardness of a crystal by describing its atomic structure, or when we explain the volume of a cube by appeal to the dimensions of its sides. Many such explanations can be known *a priori* in a way that causal explanations usually cannot; they develop out of mathematical calculations rather than observed correlations.³⁵ Like causal explanations, though, constitutive explanations explain one fact in terms of another by demonstrating the dependency of one on the other: by showing how the shape of a diamond crystal would not be what it is unless each carbon atom of which it is composed were symmetrically surrounded by four other carbon atoms, or by showing how a change in the dimensions of a cube would alter its volume. And like causal dependencies, constitutive dependencies can be seen to the extent that we can simultaneously imagine the relevant alternatives. To perceive its atomic structure as the reason for the diamond's shape and hardness, we must imagine other atomic structures creating other shapes and other bonds; and to see the dimensions of the cube as the reason for its volume, we must imagine the different volumes of cubes with different dimensions. Not all constitutive dependencies are easy to perceive, of course – even when the dependencies in question are geometrical dependencies. The shape of a diamond crystal is constitutively dependent on particular alignments of carbon atoms, but this is not something that is easily seen. We must look at diagrams and read explanations and become familiar with various aspects of chemistry more generally before we can, using our imagination, come to see the lineup of carbon atoms as the reason for the diamond's shape and hardness. The atomic structure of a diamond is too small to be apparent to the naked eye, of course, and yet we may get to the point where we automatically see that structure 'behind' or 'within' the diamond crystal before us. (In order to see the reason for the diamond's shape and hardness, it is not enough merely to believe that certain atomic alignments create certain shapes and bonds, nor is it enough to imagine such geometrical structures when contemplating a diamond; our view of the diamond must be informed and infused with our imagining of various changes in atomic structure and their various effects.)

Psychological explanations often invoke reasons that are both causal and constitutive : causal insofar as wanting to catch a someone's attention is the *cause* of one's gesture (the wanting occurred first, and one would not have gestured that way otherwise); and constitutive insofar as being a gesture at all (as opposed to an stretch or a tic, say) depends on being caused by an appropriate desire.³⁶ Whether all mental explanations are both causal and constitutive, as some functionalists claim, is not important for our purposes. What is important is the

³⁴ If Kant is right in thinking that the perception of causal connections is just as necessary for experience as the perception of objects in space, then the second part of this argument could be part of a transcendental argument as well.

³⁵ I do not mean to rule out the possibility of *a posteriori* necessities.

³⁶ Note how constitutive relations work in the other direction as well: wanting to catch someone's attention entails making an appropriate gesture (other things being equal).

recognition that the account of perceiving reasons defended above can be extended to mental explanations as well: it is possible to perceive the reason for someone's gesture because it is possible simultaneously to imagine relevant counterfactual alternatives. In particular, in the context of a party, I might imagine how your behavior would have been different if you were uninterested in the people around you, I might imagine how you would have behaved if the other person were either closer or farther away, or if your vision had been blocked, or if you wanted to warn them away; it is because of the way that your hand movement varies with your interests, as I imagine them across a range of counterfactual possibilities, that I come to see your wanting to catch someone's attention as the reason for the gesture you use. Or, to take another case: we may come to see why someone avoids certain kinds of conversation by imagining what she would have done if the conversation had taken some other turn, what she would have said if she were less worried about being rejected, how she would respond if we continued to press a point, or if we pressed it in a different way, and so on. For it is only by imagining such alternatives that we come to see just what depends on what.

The accuracy of our imagining, and thus the possibility of actually seeing (as opposed to merely seeming to see) the reason for someone's behavior may arise out of past observations of others or they may arise out of reliable projections of one's own behavior in a variety of circumstances, or both. For one's knowledge to be perceptual rather than inferential, though, the relevant alternatives must be actively imagined in such a way as to infuse the very look of the gesture (as jaunty and inviting rather than slow and absentminded, as it would be if it were merely a stretch) and the very sound of the conversation (as uneasy and evasive rather than impatient and restless, as it would be if the speaker were merely bored). Indeed, in order to see an action as an action of any sort (versus merely a bodily movement), we must see it as both caused and constituted by a particular desire or fear or hope or plan. We must sense the gesturer's want and the avoider's insecurity as palpable powers in a field of forces that is both causal and constitutive.

Perceiving oughts

So far we have been discussing reasons that explain the way things are. Let's turn next to reasons for how things *ought* to be (whether or not they are actually that way), reasons that *justify* rather than explain.³⁷ Is it possible to perceive relations of this

³⁷ Note, though, how Aristotle's final causes (like evolutionary explanations) combine explanations and justifications. Aristotle thought that all natural states of affairs have a final cause, but it is now more common to claim that only conscious beings have purposes and aims, so explanations that appeal to reasons of this narrower sort are only appropriate when the state to be explained belongs to a conscious being. The underlying thought, here, is something like the following: a future aim can only be efficacious insofar as it is contemplated and acted on in the present, and only conscious beings are capable of such contemplation and action. Or, put another way: it is not the outcome as such, but rather our *representation* or *anticipation* of an outcome that is causally effective and therefore explanatory.

sort?³⁸ (Note the analog to Sam’s blind conviction here too: There are cases in which we automatically, and reliably, do the right thing – without, however, experiencing the action’s rightness as something objective, without experiencing ourselves as having a reason for what we do.)

Two features of justificatory relations seem to stand in the way of our perceiving them. First, the correspondence between an act and its justification is much less reliable than the correspondence between an event and its cause; acts of kindness or revenge, for example, often occur regardless of their justification or lack thereof. So it cannot be on the basis of observed regularities that we come to see an act and its justification as related (indeed, if people always did what they should do, there would be little need for justifications); and counterfactuals that claim that a given act would not occur in the absence of a particular justification seem doubtful. Second, whereas our interest in explanations typically concerns either past or present events -- events that have already been observed, our interest in justifications typically concerns future possibilities – events that are not yet observable. But if the state that is justified is not observed, how is it possible to perceive a relation between it and its justification? I address these two challenges in turn.

Causal dependencies, it will be recalled, always assume certain background conditions against which one state of affairs will not occur unless preceded by the precipitating state of affairs; and the alternatives that we imagine in coming to perceive the effect as dependent on the cause are alternatives in which the background conditions remain fixed. The nail would not move without being hit by the hammer on the assumption that it is not surrounded by an unusual magnetic field, and being hit by the hammer necessitates the nail’s movement on the assumption that it is not welded to the surrounding surfaces, and so on. Likewise one act justifies another only against the background of certain assumed conditions. A threat to someone’s life justifies forceful intervention on the assumption that such intervention will help rather than hurt, for example. For justifications, like explanations, must support certain counterfactuals (if x is the reason you should do y, then you should not do y in the absence of x); and there are ethical necessities just as surely as there are physical or natural necessities.³⁹ In each case, the relevant counterfactuals only hold *ceteris paribus*. Even so, the objector may insist, the fact that I should intervene to save another’s life, all things being equal, does not ensure that I *do* intervene to save another’s life, all things being equal – whereas the case that the hammer swings does ensure that the nail moves, all things being equal. But now add the reasonable assumption that what we *should* do must be something we *can* do, and that what we can do must conform to physical necessities (causal laws, for example). Then facts about what we should do – ethical necessities

³⁸ There is an important, if minority, view that insists that all moral knowledge is perceptual – that no other sort of knowledge works for the sort of thing that morality is. See, for example, Murdoch (1970, 1993) and McDowell (1985).

³⁹ See Fine (2002) on the formal parallels between moral necessities and other sorts of necessities. Whether the relevant counterfactual, for moral oughts, is “If state S did not obtain, then act A should not be undertaken (*ceteris paribus*)” or “If state S did not obtain, then act A is not required (*ceteris paribus*)” is an interesting and complex question. Either option is compatible with what is said here about the perceivability of justificatory relations.

– must also be facts about what we would do under appropriate conditions. Just what those conditions might be is a matter of considerable dispute, of course: some will say we would do what we should if we were knowledgeable about our own desires and about relevant worldly facts, others will say we would do what we should if our ‘ears’ were open to the word of God, if we were responsive to the claims of reason, if our will were strong, or if we could successfully resolve our Oedipal complex. As in the case of causal explanations, the relevant background conditions needn’t be normal in the sense of usually obtaining; rather, they must specify the “norm” in the sense of identifying an ideal or standard within which causal laws do hold and against which deviations can be registered as such. My point, then, is just this: ethical reasons that justify, like causal or constitutive reasons that explain, only support counterfactuals against the backdrop of certain fixed conditions, and as long as the imagined alternatives respect those fixed conditions, the relevant dependencies can be perceived. As long as I restrict my imagining to cases where my will is strong and my knowledge is complete, for example, I will perceive the threat to another’s life as necessitating my intervention.⁴⁰

The other worry about the possibility of perceiving justificatory reasons as opposed to explanatory reasons concerned the future status of the justified event. It should be noted that while justifications are more likely to concern not-yet-observed events, and explanations are more likely to concern already-observed events, this is not always the case. We frequently justify our past actions – to our friends, or before the law, for example; and we frequently offer an explanation of something that has not yet happened – an upcoming eclipse, for example. If seeing reasons depended on the past versus future status of the event in question, this would not rule out the possibility of seeing justificatory reasons for events that have already occurred and have been observed. It would not rule out the possibility of my seeing why you are helping your mother up the stairs, for example – because she is in pain, because you care for her, because she needs you (these are all things that can be seen). Furthermore, it is not entirely obvious that we cannot perceive future events. We often *say* that we can see what is going to happen – that the child is going to fall, that the sun is about to appear from behind the clouds, that the pineapple will be ready to eat tomorrow, that the house will fall without additional support – and these reports are often based on experiences that have both phenomenological and justificatory immediacy. On the account I have been defending, if such foresight also includes present imagining of the changes leading to that future event, then our knowledge may indeed be perceptual. Whether this means that our perception of a given state of affairs does not need to be caused by that state of affairs⁴¹, that future events can be the cause of present perceptions⁴², or

⁴⁰ Brandt (1996) was a prominent defender of basing ethics on the imagining of ideal conditions. There is an interesting recent literature, especially Gendler (2000) exploring what a moral person *cannot* or *will* not imagine. See also Goldie (forthcoming), who introduces the analogy of chess, maintaining that, with enough practice at imagining and eliminating failing strategies, we may get to the point that only one option remains as viable. But seeing only one option as available is different than seeing a justificatory relation – i.e. feeling the force of a particular fact on one’s actions/ seeing the necessity of a certain act given that fact, which is what I seek to explain.

⁴¹ See Dretske (1969), pp. 108-9, for example.

that future states of affairs can be a part of the present in some way⁴³ is not something to be decided here; for present purposes, it is enough to note that the perception of future events is compatible with a variety of positions regarding causation and perception.

Some advantages of perceiving reasons

What may not be evident at this point is why the ability to see reasons matters. There is, I take it, no puzzle about why having accurate beliefs about reasons is important: having reasons enables us to discover connections, formulate unities, make predictions, and increase our power to change things. But none of these advantages seems to depend on *seeing* reasons – knowing them directly and imaginatively rather than knowing about them in a more indirect or reflective way. Indeed, the attempt to imagine counterfactual possibilities in an attempt to see why things happened as they did or why things should happen in a particular way may seem more misleading than helpful; relying on our imaginations can cause us to overlook important distinctions, to be distracted by extraneous details, and to confuse what is spatial with what is not. These are all legitimate worries – worries that should temper any preference we might have for perceptual rather than non-perceptual ways of knowing reasons. By clarifying the process involved, though, the account offered above has clarified some of the advantages of perceiving reasons, and these advantages are worth noting.

One advantage of perceiving reasons – one that was recognized by Descartes, for example -- is that the immediacy of perception makes our knowledge less reliant on the vagaries of memory.⁴⁴ Knowledge that is inferential rather than perceptual will be less secure insofar as it depends on a series of steps that may be forgotten by the time a conclusion is reached. With forgetting, we lose access to the reasons we had, which makes us more susceptible to doubt and less flexible in the face of challenges. I may rightly believe, on the basis of a series of inferences about molecular structure, that ice has a more stable structure than water; but if I can't remember the relevant inferential steps, my claim to knowledge will be more vulnerable to doubt – by me and by others -- than it would be if I could 'see' the reasons for its greater stability. Likewise, I may come to know, based on past inferences from your behavior, that you do not like to talk about your work; but if I can't remember what my inferences were and what they were based on, my knowledge will be less secure than if your past behavior appears as alternate manifestations of an underlying trait that I imagine as existing 'within' or 'behind' this particular dislike.

⁴² Dummett (1954) defends the possibility of what he calls "quasi-cause" (quasi because not tied up with agency in the way that cause is). He argues for the intelligibility of precognition. (Flew's response to Dummett argues that the possibility of precognition is incompatible with the presence of memory.)

⁴³ It might be argued, for example, that a baby's future as a human being is part of its present moral status. Harman (2003) offers an interesting version of such an argument.

⁴⁴ Descartes (1637/2001) recommends going over the steps of proofs faster and faster until they can be seen all at once as a solution to the problem of memory. He also (1681/1952) expounds on the importance of visualizations in math for new discoveries, which is relevant to the second advantage, discussed below.

What I am calling the greater security of perceptual knowledge should not be equated with its greater reliability. A sequence of calculations concerning the size of a room or the path of the Big Dipper, for example, may give us more reliable knowledge than an all-at once perception of the same.⁴⁵ Indeed, the sort of perception at issue here (what Kant called “reproductive imagination”) depends on imagination that itself relies on memory, so it may seem that seeing reasons is equally vulnerable to the vagaries of memory. Even so, I would argue, there is an advantage to having the contributions of memory combined in a present perception, for by occurring simultaneously rather than in sequence, it is easier to ensure the overall consistency of one’s beliefs, thereby counteracting the incoherence if not the inaccuracies of memory. A report on an arrangement of objects in a room allows inconsistencies to escape notice in a way that imagining or drawing the arrangement does not. And while it is relatively easy to equivocate on a term across several premises considered in sequence, it is quite difficult to miss a shift in attention between simultaneously imagined objects. Seeing reasons is a more secure form of knowledge, therefore, insofar as it ensures consistency across a broader range of beliefs.

A second advantage of seeing reasons is that perception enhances our ability to generate new knowledge from old (using what Kant would call “productive” imagination). If we can see the molecular structure of ice, we can also see how an increase in molecular movement – i.e. an increase in temperature – would break down that structure and result in the ice melting. Prior to visualizing the molecular structure of ice, we might infer (given our belief about the molecular structure of ice and our belief that a rise in temperature causes ice to melt) that a temperature rise that causes ice to melt must cause a change in its molecular structure, but we wouldn’t yet know why; thus, we wouldn’t yet know that the transition from ice to water is abrupt, or that changing ice to 32 degree water requires more heat than changing 32 degree water to 33 degree water.⁴⁶ These are facts that are immediately evident once we picture the internal process involved, and otherwise not. Although these facts may be derived (rather laboriously) from various propositions about temperature and molecular structure, they may not be discovered unless the relevant reasons are perceived. Of course, it may also be true that we are more likely to discover other things through non-perceptual calculations. I am not claiming that all new discoveries

⁴⁵ It could be argued that a non-perceptual judgment is more certain only if it is derived from a greater number of perceptual judgments. I’m not sure how to count here. But, in any case, this means that a non-perceptual judgment that p could be more trustworthy than a perceptual judgment that p. Furthermore, the probability of p can be greater than the probability of any individual observation (q, r, s, t) from which it is derived if p is the only possible explanation of each of q, r, s, and t and if the probability that at least one of these observations is accurate (the probability of q or r or s or t) is greater than the probability of any one of them (the probability of q, the probability of r, the probability of s, or the probability of t).

⁴⁶ Some other examples: When we see the molecular structure of crystals, we can also see that electricity will be released when pressure is applied along the axis of symmetry. If I see the fear in a friend’s behavior (or in a painting), I am usually better able to interpret certain other aspects of her behavior (or other features of a painting) that were previously either unnoticed or unintelligible to me.

depend on perceptual ways of knowing – only that many do, and that this constitutes an advantage of being able to see reasons.

A third advantage to seeing reasons is that perceptions usually motivate us more effectively than beliefs that depend on inference.⁴⁷ This is particularly important in moral contexts where seeing another's suffering can move us to act when merely believing that another is suffering does not. It can also be valuable in non-ethical contexts where immediate action is called for; visualizing the internal structure of melting ice can strengthen one's motivation to retreat from a thawing lake quickly – a retreat that may not feel so urgent otherwise. Here the important thing is not the more secure epistemic standing of perception but its visceral urgency. Many of our movements arise quite forcefully and automatically in response to things that we see; we dodge fast-approaching objects without even having the time to think about it, we reach out to protect a falling vase before remembering that we hoped it would break, and we make our way through a crowded room thinking only of the person at the other end. Without perception of the relevant objects, these actions are often possible but usually less likely and less effective. Emotionally also, we are more responsive to things that we see than we are to things that we merely believe: seeing that a painting depicts war affects us with a force that merely believing that it depicts war does not, and seeing a child's delight moves us in a way that merely believing that it is delighted does not. Insofar as increased responsiveness to the world and to others is desirable, then seeing reasons will be preferable to understanding that remains non-perceptual.

Conclusion

I have defended a particular account of what perception involves – an account that elaborates on the experience of objectivity as a way to make sense of the distinctive kind of immediacy or transparency that perception has. The experience of objectivity requires us to actively imagine multiple perspectives on a given fact, for it is only through such a multiplicity of perspectives that the independence of a fact can be recognized, and it is only by actively imagining alternative perspectives that the objectivity of what we experience is evident within that experience itself.

This account of perception is then extended to make sense of the claim that we can perceive reasons of various sorts – causal reasons, constitutive reasons, and justificatory reasons. Again the emphasis is on discerning invariance across multiple alternatives; but whereas, in the case of perceiving a tipping table, it is alternative perspectives on that state of affairs that must be imagined, in the case of perceiving a broken leg as the cause of tilting table, it is alternative arrangements of the table's parts and its surroundings that must be imagined. The basic idea here is that we can come to see the necessity of one state or event, given the presence of a certain other state or event (in a given context) by imagining relevant counterfactual possibilities. Despite the differences between explanatory reasons and justificatory reasons, it was possible to give similar accounts of what makes each perceivable.

⁴⁷ See Doggett and Egan (2007) for an interesting discussion of the motivational force of imagining.

Finally, I described three respects in which perceiving reasons is preferable to knowing reasons in some other way: it is easier to check the consistency of our beliefs when they are presented perceptually, it is easier to generate new knowledge when the relations between existing knowledge can be considered simultaneously, and our actions are more likely to accord with our reasons when those reasons are perceived.⁴⁸

⁴⁸ I have learned from comments received in response to earlier versions of this paper presented at CUNY Graduate Center, at the CSMN Center of the University of Oslo, and at Vassar College. This paper has also benefited from the encouragement and suggestions of Quassim Cassam, Peter Goldie, Ásta Sveinsdóttir, David Velleman, and an anonymous reviewer for this journal.

REFERENCES

- Aikins, Kathleen, ed. 1996. Perception. Oxford: Oxford University Press.
- Ayers, Michael. 1993. Locke: epistemology and ontology. London: Routledge and Kegan Paul.
- Beebe, Helen. 2003. "Seeing Causing". *Proceedings of the Aristotelian Society* 103: 257-280.
- Brandt, Richard. 1996. Facts, Values, and Morality. Cambridge: Cambridge University Press.
- Brandom, Robert. 1994. Making it Explicit. Cambridge, Mass: Harvard U Press.
- Brewer, Bill. 1998. "Experience and Reason in Perception". In A. O'Hear, ed., Current Issues in Philosophy of Mind. Cambridge: Cambridge University Press.
- Burge, Tyler. 2003. "Perceptual Entitlement". *Philosophy and Phenomenological Research* 67(3): 503-548.
- Callaghan, Casey. 2007. *Sounds*. Oxford University Press.
- Church, Jennifer. 2000. "'Seeing As' and the Double Bind of Consciousness". *Journal of Consciousness Studies* 7 (8/9): 99-111.
- Church, Jennifer. 2008. "The Hidden Image: A Defense of Unconscious Imagining and Its Importance". *American Imago* 65 (3): 379-404.
- Collins, John and Ned Hall and L.A. Paul, eds.. 2004. Counterfactuals and Causation. Cambridge, Mass: MIT Press.
- Costa, Michael. 1989. "Hume and Causal Realism". *Australasian Journal of Philosophy* 67(2): 172-190.
- Davidson, Donald. 2001. Subjective, Intersubjective, Objective. Oxford: Clarendon Press.
- Descartes, Rene. 1637/2001. Discourse on Method, translated, with an introduction, by Paul J. Olscamp. Indianapolis: Hackett Publishing.
- Descartes, Rene. 1681/1952. Rules for the Direction of the Mind, translated by Elizabeth Haldane and G.R.T. Ross. Chicago: Encyclopaedia Britannica.

Doggett, Tyler and Egan, Andy. 2007. "Wanting Things You Don't Want: The Case for an Imaginative Analogue of Desire". *Philosophers Imprint* 7 (9): 1-17.

Dretske, Fred. 1969. Seeing and Knowing. Chicago: University of Chicago Press.

Dretske, Fred. 1981. Knowledge and the Flow of Information Cambridge, Mass: MIT Press.

Dummett, M. 1954 "Can an Effect Precede its Cause". *Proceedings of the Aristotelian Society Supplement* 28: 27-44.

Dummett, M. 1964."Bringing about the Past". *Philosophical Review* 73: 338-359.

Fine, Kit. 2002. "The Varieties of Necessity". In T.S. Gendler and J. Hawthorne, eds., Conceivability and Possibility. Oxford: Oxford University Press.

Gendler, Tamar Szabo. 2000."The Puzzle of Imaginative Resistance". *Journal of Philosophy* 97(2): 5-81.

Gendler, T.S. and Hawthorne, J., eds. 2006. Perceptual Experience. Oxford: Clarendon Press.

Glymour, Clark. 1998. "Psychological and Normative Theories of Causal Power and Probabilities of Causes". In G. F. Cooper and S. Moral, eds., Uncertainty in Artificial Intelligence. San Francisco: Morgan-Kaufmann.

Goldie, Peter. (forthcoming) "Seeing what is the kind thing to do: perception and emotion in morality".

Harman, Elizabeth. 2003. "The Potentiality Problem". *Philosophical Studies* 114(1-2): 173-198.

Hume, David. 1777/1978. Enquiries Concerning Human Understanding and Concerning the Principles of Morals, L.A. Selby-Bigge and P.H. Nidditch, eds. Oxford: Clarendon Press.

Keeley, Brian L. 2002. "Making sense of the senses: Individuating modalities in humans and other animals". *The Journal of Philosophy* 99: 5-28.

Kelly, Sean. 2005. "Seeing Things in Merleau-Ponty". In The Cambridge Companion to Merleau-Ponty, ed. Carman Taylor. Cambridge: Cambridge U Press.

Lewis, David. 1986. Philosophical Papers. Oxford: Oxford University Press.

- Lopes, Dominic McIver. 2000. "What Is It Like to See with Your Ears? The Representational Theory of Mind". *Philosophy and Phenomenological Research* 60(2): 439-53.
- Martin, M. G. F. 2002. "The Transparency of Experience". *Mind & Language* 17(4): 376-425.
- McDowell, John. 1984. "Wittgenstein on Following a Rule". *Synthese* 58:325-363.
- McDowell J. 1985. "Values and Secondary Qualities". In Morality and Objectivity: a Tribute to J.L. Mackie. London: Routledge and Kegan Paul. Reprinted in his Mind, Value, and Reality. Cambridge, Mass: Harvard University Press (2002).
- Michotte, A. 1963. The Perception of Causality. London: Methuen.
- Miller, Alan. 2007. "What the Disjunctivist is Right About". *Philosophy and Phenomenological Research* 74(1):176-198.
- Millikan, Ruth. 2000. On Clear and Confused Ideas. Cambridge: Cambridge University Press.
- Murdoch, Iris. 1970. The Sovereignty of Good. New York: Schocken Books.
- Murdoch, Iris. 1993. Metaphysics as a Guide to Morals. London: Penguin.
- Noë, Alva. 2002. "On What We See". *Pacific Philosophical Quarterly* 83:57-80.
- Noë, Alva. 2003. "Causation and perception: the puzzle unraveled". *Analysis* 63(2):93-100.
- Noë, Alva. 2006. "Experience without the Head". In T.S.Gendler and J. Hawthorne, eds., Perceptual Experience. Oxford: Clarendon Press.
- O'Shaughnessy, Brian. 1990. "The Appearance of a Material Object". *Philosophical Perspectives* 4 :131-151.
- Peacocke, Christopher. 1983. Sense and Content: Experience, Thought, and their Relations. Oxford: Oxford University Press.
- Peacocke, Christopher. 1986. Thoughts: An Essay on Content. Oxford: Blackwell.
- Reed, Lisa. 1999. "Necessary versus Probably Cause". *Journal of Philosophical Logic* 28:289-326.
- Reid, Thomas. 1792/2001. "Of Power". *The Philosophical Quarterly* 51(202):3-12.

- Searle, John. 1992. The Rediscovery of the Mind Cambridge, Mass: MIT Press.
- Sellars, Wilfred. 1956. "Empiricism and the Philosophy of Mind". In H. Feigl and M. Scriven, eds., Minnesota Studies in the Philosophy of Science. Reissued (1997), with commentary by Robert Brandom, as Empiricism and the Philosophy of Mind. Cambridge, Mass: Harvard University Press.
- Snowdon, Paul. 1979-80. "Perception, Vision and Causation". *Proceedings of the Aristotelean Society* 81:17-192.
- Snowdon, Paul. 1990. "The Objects of Perceptual Experience". *Proceedings of the Aristotelean Society Supplement* 64:121-150.
- Stanley, Jason and Williamson, Timothy. "Knowing How". *The Journal of Philosophy* 98(8):411-44.
- Strawson, P. F. 1974. "Imagination and Perception". In his Freedom and Resentment, and other Essays. London: Methuen.
- Suppes, Patrick. 1970. A Probabilistic Theory of Causation. Amsterdam:North-Holland Publishing Co..
- Yablo, Stephen. 2004. "Advertisement of a Sketch of an Outline of a Prototheory of Causation". In J. Collins and N. Hall and L.A. Paul, eds., Counterfactuals and Causation. Cambridge, Mass: MIT Press.
- Wittgenstein, Ludwig. 1980. Remarks on the Philosophy of Psychology, edited by G.E.M. Anscombe and G.H. von Wright, translated by G.E.M. Anscombe. Oxford: Blackwell.