

## *Causation and perception: the puzzle unravelled*

ALVA NOË

1. Attempts to give a philosophical analysis of perception haven't come very far since Grice's causal theory of perception (Grice 1961). According to the causal theory, *S* sees that *o* is *F* if and only if *S* has a visual experience as of *o*'s being *F*; *o* is *F*; and *o*'s experience depends, causally, on *o*'s being *F*. The point of the causal-dependence clause is to rule out what came, later on, to be known as veridical hallucination, i.e. cases in which you enjoy a veridical visual experience which nevertheless falls short of being genuinely perceptual (Lewis 1980). Grice's example went like this: suppose there is a clock on the shelf and that you have a veridical experience as of its being there. But suppose that the cause of your experience is not the clock's being on the shelf, but, say, post-hypnotic suggestion. In such a case, you do not succeed in seeing that the clock is on the shelf, even though your experience is veridical. What explains this failure, so Grice proposed, is the fact that in this case the requirement of causal dependence is not met.

Almost everyone immediately recognized that Grice was on to something. Perception *is*, after all, a causal concept. But it quickly became clear (to Grice as well) that the causal theory was too weak, and that reasonable efforts to strengthen it were likely to make it too strong. Suppose that what produces your experience as of the clock on the shelf is a manipulative neurosurgeon. Suppose further that this neurosurgeon makes it look to you as if there is a clock on the shelf *because* there is a clock on the shelf. In a case such as this, your experience does satisfy the causal-dependence clause, but it does so in the wrong sort of way. The standard move in response is to strengthen the theory by constraining what counts as the right sort of causal dependence (e.g. Strawson 1974). The idea, roughly, is that for an experience to count as an instance of perception, it must depend on what it is as of in the normal way that perceptual experiences depend on what they are as of. Whatever else is true, it seems clear that the dependence mediated by the neurosurgeon's intentions is not normal.

Sadly, the strengthened causal theory faces two problems. First, it makes the philosophical analysis of perception beholden to empirical findings about what the correct causal relation is. This isn't such a big deal in our post-Putnamian intellectual climate; it may even be a virtue that the theory makes room for developing empirical science. The second problem is more serious: the strengthened causal theory is way too strong; it rules out as non-perceptual any cases in which the dependence

of experience on how things are is abnormal. But we don't want to rule out all such cases. We can certainly make sense of the possibility of wildly abnormal forms of prosthetic or artificial perception.<sup>1</sup>

In the face of these difficulties, most philosophers despaired of giving a coherent analysis of 'the perceptual relation' and moved on to other problems. I am no fan of the project of philosophical analysis: it's doubtful that there has ever been an analysis (that is, a breakdown into necessary and sufficient conditions) of any philosophically interesting concept. But I do think that the causal theory is obviously right in certain ways, and it is obviously wrong in others, and it would be worthwhile to explain why this is so, even if we reject the project of analysis.

I now believe it is possible to do this. The problem with the causal theory is not that it fails to articulate with sufficient detail the right kind of causal relation, or that it relies on empirical rather than merely a priori considerations about what perception really is. The problem, rather, is that it relies on a much too simplistic account of the content of perceptual experience. Once we enrich our account of perceptual content, we'll get a better understanding of what's right in the causal theory, what's wrong, and how the causal theory needs to be amended to give a more illuminating account. Or so, at least, I argue.

2. The truth in the causal theory is this: it is a necessary condition on your seeing that such and such is the case not only that things are that way, and that you have a visual experience as of them being that way, but that you wouldn't have had that experience if they hadn't been that way. The causal theory captures the fact that perceptual experience – how things look, say – depends on how things are, and it does so in a way that supports the relevant counterfactual considerations.

A little reflection reveals that how things look does not *only* depend on how things are, and this is why the conditions mentioned in the previous paragraph are not jointly sufficient for genuine perception. There are aspects of perceptual content – aspects of how the experience presents things as being – that are not determined by how things are alone, but that depend, in addition, on the perceiver's relation to how things are.

Perception is a way of keeping track of how things are, but it is also a way of keeping track of one's relation to how things are (Hurley 1998; Noë 2002a). This point has two aspects. First, 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

<sup>1</sup> Recent work on 'sensory substitution' and prosthetic perception in the perceptual and neural sciences shows that these are not idle points. For discussion, see O'Regan and Noë 2001.

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. If we count perspectival properties (such as the elliptical look) as belonging to the representational content of experience (Harman 1990; Noë 2002b), then we are implicitly counting ourselves (or at least our vantage points or bodily locations) among those contents. Second, it is hard to understand how one could keep track of how things are if one were not also capable of keeping track of the ways in which one’s perceptual experience depends on what one does. The perspectival shape of the plate changes as one moves. Indeed, it seems likely that our practical grasp on the way it changes as we move is precisely the way we succeed in experiencing its roundness. More generally, how things look (what one sees) changes with every movement of the eye or turn of the head.

Here’s the crux: these *perspectival* aspects of perceptual content are only partly determined by how things are. They depend further on one’s relation to how things are. Any account of perception that ignores this dependence of how things look on one’s movements (that is, on changes in one’s relation to how things are), and that in effect ignores the distinctively perspectival aspects of perceptual content, will fail to provide an adequate account of what perception is. This is the problem with the causal theory.

Let’s back up and go over this more carefully. Perceptual content – what philosophers call representational content: *how* the experience presents the world as being – 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.<sup>2</sup>

Now there is no general requirement that for an experience to be veridical, it must be completely veridical. You can succeed in seeing a spoon, for example, even if the spoon non-veridically looks bent (because it is in water, say). And so, likewise, there is no hard requirement that for a perceptual experience to be veridical, it must be veridical along both dimensions of its content. We frequently enjoy visual experiences that are perspectivally non-veridical but factually veridical. This is the case, for example, when you spy on someone from under water using a periscope. The experiences you have with the periscope represent how things are, but they *mis*represent your relation to how things are, that is, they mis-

<sup>2</sup> This is true of hearing and touch as well, although in the case of touch, the term ‘perspectival’ seems somewhat less appropriate.

represent the perspectival content of your experience.<sup>3</sup> When you watch a live sporting event on television, you are able to track what's happening, but you do so in a perspectively non-veridical way. Perhaps you adopt the standpoint of one or more cameras. Crucially, you don't correctly or veridically experience the event's spatial relation to yourself. It would be dogmatic to deny that you *see* the sporting event, that you see it 'through' or by means of the television cameras. But it would be just as dogmatic to insist that there is no difference whatsoever between normal perception, in the here and now, and *televisual* perception. The difference is one that is now easy to explain: when you witness events in person, your experiences track not only how things are, but also how things are in relation to you.<sup>4</sup>

For an experience to be veridical *simpliciter*, then, is for it to be veridical along both factual and perspectival dimensions. But we have already drawn attention to the fact that veridicality is not sufficient for perception. Just as there are two dimensions along which an experience can be veridical, so there are two dimensions along which it can be *veridically hallucinatory*. Things can turn out to be the way they seem, even though their seeming that way is independent of the fact that they are that way, and they can turn out to stand in the relation to you in which they seem to stand, even though your relation to how things are has no effect on how they seem.

For an experience to be not merely veridical (*simpliciter*), but genuinely *perceptual* (*simpliciter*, and, as we might put it, normal), the right sort of counterfactual supporting dependence must be maintained along *both* dimensions of content. We know what that requires as far as factual content is concerned: that things would have looked different had they been different. But what does it require as far as perspectival content is concerned? The answer is ready to hand: it requires that had one's relation to how things are been different, things would have looked different (even if how things in themselves are were unchanged). It requires, therefore, that how things look depends in a fine-grained, closely-coupled sort of way, on movements. In addition to a causal-dependence requirement, we need

<sup>3</sup> Mirrors may provide a further example. You see the police in the rear-view mirror, but in seeing them, you mis-experience their location relative to you. We rarely notice the perspectival non-veridicality in cases such as this, because they are so familiar. In other cases, this is more apparent. For example, it is rather difficult (for some of us at least) to *see*, of the car in the mirror, which side is the driver's side. Thanks to Casey O'Callahan for this example.

<sup>4</sup> I take it that it matters that you are watching the game on TV 'live,' i.e. in real time. The security guard, for example, really watches (and sees) the crowd on closed-circuit monitors. Exactly what the role of time is in perceptual experience is tricky. Can you see the stars in the heavens, even though they may no longer exist?

a further movement-dependence requirement to the effect how things look depends (in a counterfactual supporting way) on movements (of the relevant type).

When you take your experience at face value, you encounter it as raising questions not only about how things are, but about how we stand in relation to how things are. To be a perceiver, then, you must understand, implicitly, that your perceptual content varies as things around you change, and that it varies in different ways as you move in relation to things around you.

3. This is the key to understanding what went wrong with the causal theory. Consider again problem cases such as that of the manipulative neurosurgeon. The way the case was initially described, it was reasonable to suppose that although factual dependence is preserved (it wouldn't have looked as if there were a clock on the shelf if there hadn't been one there), perspectival dependence fails. Perspectival dependence fails because it is natural to assume, given the way the case is described, that how things look wouldn't depend, in a nuanced and fine-grained way, on your movements and changing relation to how things are. How, after all, could the neurosurgeon preserve that kind of dependence? It's hard to imagine. Wouldn't the electrodes impede movement? So we naturally assume that, in a case such as this, head movements and eye movements make no difference to what is seen.<sup>5</sup> But this is just to say that dependence along the perspectival dimension fails and that the experience is, at least on this dimension, not genuinely perceptual. An experience which is unchanged as the eyes and head move, as the relation to the environment alters, is not *visual*, whatever else it is.

If this diagnosis of our inclination to view the manipulative neurosurgeon case as falling short of bona fide perception is along the right lines, then it ought to be possible to construct cases that do rise to genuine perception but that are equally deviant in their causal chains. Consider an example. An angel hovers near and makes it the case that your experiences depend, in a counterfactual supporting way, on how things are and on what you do. It looks to you as if there is a clock on the shelf, only because there is one there. And for every little twitch of your eye, there is a corresponding change in how things look. Moreover, the angel, let's

<sup>5</sup> Actually, the non-perceptuality of this case may have been overdetermined. Not only is there a failure of perspectival dependence, it is hard to imagine that factual dependence could be really maintained by the fiddlings of a surgeon. We naturally suppose that it won't be the case that experience changes with every little visible change in how things are around you. As is frequently the case in philosophy, our intuitions are driven by the way in which cases are *under*-described.

say, is committed to maintaining these regularities come hell or high water. This is a way of seeing, I would say. It's an unnatural way of seeing, to be sure, but it's seeing nonetheless. And this would be so even if, as we might imagine, your visual cortex and retina are damaged so that, but for the angel's interventions, you would be blind. This would be a kind of divine prosthetic vision.

A second example is due to David Sanford (unpublished lecture). Sanford asks us to imagine the case of 'Chris the amazing human hearing aid'. Chris has superhuman powers of mimicry. She listens to sounds and conversations and is able to repeat them for you in a way which is qualitatively indistinguishable from the sounds themselves. Imagine that she is able to perform this feat in real time. According to Sanford, when you use a normal hearing aid to hear the people across the table, it is truly those people you are thus enabled to hear. But when you employ Chris the amazing human hearing aid, it is actually Chris, and not the people across the table, that you hear. Chris does not, in other words, provide prosthetic hearing, even though, subjectively speaking, the experiences of using Chris and a top-notch hearing aid are indistinguishable. Counterfactually, Sanford is quick to reassure, Chris and the mechanical hearing aid are exactly on a par.

But this is not quite right. Sanford is mistaken when he claims that Chris and the authentic hearing aid are on a par as far as counterfactual dependence goes. They are only on a par as far as *factual* content is concerned. To appreciate this, consider that, as the case is described, it is natural to suppose that there is no dependence (of a counterfactual supporting sort) between how things sound (to you, the user of Chris-the-human-hearing-aid) and your spatial relation (say) to the people across the table. The relation of your body (and your ears) to the people across the table makes no difference to what you hear, so long as Chris is able to speak into your ears. With normal hearing, in contrast, or when you use a conventional hearing aid, this counterfactual supporting dependence is in evidence. Normally, that is, your auditory experiences of the people across the table will change if, for example, you turn away, or get up and walk across the room. I propose, then, that the basis of our judgement that Chris does not provide a means of genuine prosthetic perception – a judgement that I think is right – is the fact that in this case how things sound depends not on one's spatial relation to the object of hearing, but only on one's spatial relation to Chris. One hears the sounds as they are heard by Chris. But to hear how they sound from her vantage point is not to hear how they sound (or how they would sound) from one's own vantage point. It is, in short, the tacit presupposition of the violation of the principle of the dependence of experience on perspective and movement which explains the strength of the intuition that Chris doesn't en-

able genuine prosthetic hearing. Importantly, the fact of Chris's own agency, or the fact that it is she who is the object of hearing (and not the people across the table), make no difference to the relevant facts of the case.

Suppose that Chris learns to produce the sounds that you would hear exactly as you would hear them if your ears were not defective. She produces the sounds in a way that is modulated to account for your movements and shifting spatial relation to the sound source and also in such a way as to account for ambient conditions. With Chris's assistance, you now hear exactly what you would hear if your ears were normal and in such a way that both the factual and perspectival content of the resulting experiences are genuinely perceptual. I think it is clear that Chris now enables prosthetic hearing (even though she is an agent with a mind of her own, and even though you hear the people across the room *by way* of hearing her voice). This becomes much more intuitively acceptable if we imagine that Chris is very small, so small that she can fit snugly in the ear. Chris is now causally responsible for its being the case that you have auditory experiences which depend, counterfactually, on what is going on around you *and* on your relation to what is going on. Chris-dependent hearing is full-fledged prosthetic perception.

The problem with the causal theory is not that it can't specify or constrain the causal relation. As the examples of divine prosthetic perception and 'Chris the-human-hearing-aid' show, no causal relation is so strange or unnatural that it is incompatible with genuine perception.<sup>6</sup> The problem is that the standard hard cases are under-described. They are presented in such a way that they leave the full content of the experiences unspecified, and so leave the sense that perspectival content is unaccounted for. This is what explains our judgement that cases such as that of the manipulative surgeon are not bona fide cases of perception. Crucially, it never had to do with causation, but with a failure to acknowledge the perspectival content of perceptual experience, and with a failure to recognize that perception is answerable not only to how things are, but to what we do.

4. In conclusion, perception is a causal concept, but it is also a concept of a kind of action. Perception depends on how things are, and it depends

<sup>6</sup> In the natural world, however, it may be that the only way to preserve the right kinds of dependence relations between experience and the world is by the standard, biologically realized ways with which we are familiar. From a certain standpoint, artificial perceptual systems are very much like biological ones. Importantly, the inadequacy of such systems is in direct proportion to the degree to which they are different from ordinary, biological systems.

on what we do. When we perceive, we keep track not only of how things are, but of our changing relation to how things are. The upshot of this is that the problem with the causal theory is not that it can't account for the right kind of causal relation, but that it neglects the perspectival content of perceptual experience. In perception, the world acts on us, and we act right back.<sup>7</sup>

University of California  
Santa Cruz CA 95064, USA  
anoe@cats.ucsc.edu

### References

- Grice, H. P. 1961. The causal theory of perception. *Proceedings of the Aristotelian Society*, Suppl. Vol. 35: 121–152.
- Harman, G. 1990. The intrinsic quality of experience. In *Philosophical Perspectives 4: Action Theory and Philosophy of Mind* ed. J. E. Tomberlin, 31–52. Atascadero, CA: Ridgeview Press.
- Hurley, S. 1998. *Consciousness in Action*. Cambridge, MA: Harvard University Press.
- Lewis, D. 1980. Veridical hallucination and prosthetic vision. *Australasian Journal of Philosophy* 58: 239–49. Repr. in *Vision and Mind: Selected Readings in the Philosophy of Perception*, ed. A. Noë and E. Thompson. Cambridge, MA: The MIT Press, 2002.
- Noë, A. 2002a. Is perspectival self-consciousness nonconceptual? *Philosophical Quarterly* 52: 185–94.
- Noë, A. 2002b. 'On what we see.' *Pacific Philosophical Quarterly*. 83: 57–80.
- O'Regan, J. K. and A. Noë. 2001. A sensorimotor approach to vision and visual consciousness. *Behavioral and Brain Sciences* 24: 939–1031.
- Sanford, D. Unpublished lecture. Some puzzles about prosthetic perception. Presented at the 1997 meetings of the Society for Philosophy and Psychology in New York.
- Strawson, P. F. 1974. Causation in perception. In *Freedom and Resentment and Other Essays*, 66–84. London: Methuen.

<sup>7</sup> For helpful discussion about the ideas in this paper, I would like to thank David Chalmers, John Doris, Matthew Henken, Jim Higginbotham, Jerome Neu, Casey O'Callahan, Ric Otte, David Sanford and Evan Thompson.