8

Causation in Commonsense Realism

Johannes Roessler⋆

We think of perceptual experience as a source of propositional knowledge of the world around us. So we think of perceptual experience as a reliable source of true beliefs. Therefore, we must be thinking of perceptual experience as causally dependent on perceived objects. This is a rough sketch of a line of reasoning that plays a significant role in P. F. Strawson's defence of the causal theory of perception. In Strawson's words, it provides the 'rationale' for the causal theory of perception. The 'rationale' has been criticized by Paul Snowdon in a series of papers pioneering a 'disjunctivist' account of perceptual experience. Snowdon's central complaint is that the causal requirement has not been shown to be necessary, given that the causal theory is not without alternatives. For example, one might account for the reliability of perception by invoking a relational view of perceptual experience, on which mind-independent objects are constituents of the experience we enjoy in perceiving them.

Some philosophers have argued that the causal theory of perception and the relational view of experience are not in fact in conflict.[1] I think they may well be right, but the point does not resolve the dispute between Strawson and Snowdon. Even if we can consistently hold both views, we still face the question whether Strawson is right that a causal requirement on perception is somehow implicit in everyday explanations of what someone knows in terms of what they perceive. If he is, this would help to motivate, or even lend support to, the idea that such conditions are in some sense part and parcel of our very concepts of perception, or, as Strawson sometimes puts it, that they are a 'pre-theoretical commitment of commonsense realism'. It would present a challenge to sceptics, such as Snowdon, who think that an interest in the causal requirements on perception is, as it were, an acquired taste, not

⋆ I am grateful to an anonymous reader for OUP for valuable comments on a previous version of this chapter. I would also like to thank Bill Brewer, Steven Butterfill, Naomi Eilan, Christoph Hoerl, Hemdat Lerman, Guy Longworth, and Matthew Soteriou for helpful comments and suggestions.

[1] For one version of this move, see Child (1994) and Child (this volume). For another, see Steward (this volume). For discussion of some of the differences between them, see my Introduction (this volume).

something that is integral to our very grasp of concepts of perception.[2] Now it seems to me that while Snowdon has certainly identified a gap in Strawson's argument, it may well be possible to amend the argument. On the other hand, I think Snowdon's response—at least the part of his response that turns on the availability of alternative ways of finding the epistemic role of perception intelligible—may be too concessive. An unargued assumption of Strawson's 'rationale' is that insofar as we, even 'pre-theoretically', think of perception as a source of knowledge, we are committed to conceiving of it as a source of *beliefs*, indeed as a source whose nature makes it *intelligible* that beliefs deriving from it constitute knowledge. My aim in this chapter is to question this assumption. Everyday explanations of perceptual knowledge, I will argue, may be more simple-minded than Strawson allows. My alternative suggestion will be this: we think of seeing an object as an immediate causal enabling condition of *seeing that* the object has certain features and falls under certain types. Such explanations invoke a primitive causal link between experience and knowledge, which should not be assumed to be intelligible in terms of causal relations between experience and beliefs. In section 2, I clarify Strawson's questionable assumption and set out my alternative analysis. In section 3, I argue that under this analysis, reflection on everyday explanations of perceptual knowledge provides a 'rationale' for the *relational view* of perceptual experience (rather than for the causal theory of perception). In section 4, I consider the broader issue of whether the simple-minded scheme can be defended against philosophical criticism. I begin by reviewing the dispute between Strawson and Snowdon.

1 Strawson's rationale for the causal theory of perception

Strawson's line of reasoning (to repeat) may be summarized as follows:

(1) We think of perceptual experience as a source of knowledge of the world around us.

So

(2) We think of perceptual experience as a generally reliable source of true beliefs.

Therefore,

(3) We must be thinking of our perceptual experiences as causally dependent on perceived objects.

I begin by examining Strawson's view of the relation between (2) and (3), and Snowdon's objection. In the next section, I will look at the credentials of (2).

Consider the following parenthetical remark, addressed to someone who accepts (2) but is doubtful about (3): 'Otherwise [if it were not part of "our ordinary scheme of thought" that perceptual awareness in general causally depends on the way perceived things objectively are] the normal truth or correctness of perceptual judgments would

[2] See Snowdon (this volume) for the most recent defence of his scepticism. In that paper Snowdon also makes some further points about Strawson's 'rationale' which I will not be considering here.

seem to be something inexplicable, an extraordinary coincidence' (Strawson 1992: 60). The idea here is that if the causal condition were not integral to our naïve concepts of perception, we would be committed to thinking of perceptual experience as generally reliable, yet lack the resources to make sense of its reliability. For example, we would be thinking of someone's experience of the shape of physical objects around her as a reliable source of true beliefs about the shape of physical objects. We would assume that if someone's perceptual belief that a certain object is square is caused by her experience as of a square object, then (supposing the experience to be a perception) the belief will very likely be correct. But the fact that this is so would be bound to look mysterious. It is not clear that this would amount to any formal inconsistency in the commonsense scheme.[3] But certainly our confidence in the reliability of perception would not look like reasonable confidence. It would resemble something like blind faith. The causal theory of perception enables us to give a more charitable interpretation of commonsense realism. Suppose we wonder how it is that visual experience is a reliable source of true beliefs about the shapes of objects in our visual field. The matter has a ready explanation if it is a conceptual truth that the experience involved in a visual perception of a physical object causally originates with the object, and that its phenomenal character causally depends on certain features of the object. For presumably beliefs generated by visual experiences of objects reflect the character of the experience. Given the causal condition, the character of the experience in turn can be seen to reflect the way the object is—for example, its shape. So there is nothing (in principle) mysterious about the normal truth of visually based beliefs about the shape of things.

One problem with this argument, as Snowdon points out, is that the correctness of the causal analysis is not the only way to make the reliability of perception intelligible (Snowdon 1998), Occasionalism (with its appeal to a common cause) would be one alternative. A relational view of perception would be another. According to the latter, a perceived mind-independent object is a constituent of the experience one enjoys in perceiving it. Accordingly, certain features of the object make an immediate difference

[3] Some of Strawson's formulations do encourage this stronger reading. In 'Causation in Perception' (1974) he writes that 'dependability in this sense *entails* dependence, causal or non-logical dependence on appropriate M-facts'(p. 79, my emphasis). Now perhaps it would indeed be incoherent simultaneously to think of perception as a reliable source of true beliefs and to regard the normal truth of perceptual beliefs as an 'extraordinary coincidence'. But the fact that in a certain situation we would have to regard the reliability of perception as *mysterious* does not license the conclusion that we would be committed to regarding it as a mere *coincidence* that perceptual beliefs tend to be true. For we might assume that there is in fact some kind of explanation of the reliability of perception, though not one (currently) within our ken.

One might wonder whether the weaker reading of the argument is not too weak for Strawson's purposes; in particular, too weak to warrant the conclusion that it is a *conceptual* truth that perception causally depends on perceived objects and their features. But notice that Strawson's 'rationale' is not intended to work in isolation. Its purpose might be characterized as that of linking the causal theory, originally supported by the traditional methods of conceptual analysis, to commonsense explanatory practice. It aims to show that the causal condition plays an important *role* in the commonsense psychology and epistemology of perception. (See Snowdon (1998) for illuminating critical discussion of the point of the rationale.) Relative to that aim, it is not essential that a denial of the causal condition can be shown to be incoherent.

to the phenomenal character of one's experience, not in virtue of being causally responsible for the occurrence of an experience with that character but in virtue of being (partly) constitutive of its character. Snowdon's point is that the relational conception of perceptual experience would enable us to make the reliability of perception intelligible without invoking a causal requirement on perception.[4] For example, the normal truth of perceptual beliefs about the shape of physical objects would be unsurprising if such beliefs were caused by experiences whose character is partly constituted by the actual shape of the perceived objects. Such beliefs would reflect the character of the experiences: *i.e.* the shape of perceived mind-independent objects.

In his response to Snowdon, Strawson expresses incredulity at the idea that mind-independent objects figure as constituents of perceptual experience. As Strawson understands it, this would amount to envisaging *logical* relations between 'natural items' (rather than between their descriptions); a claim he proceeds to dismiss as a 'category howler' (1998, p. 314). I take it this response is inadequate. Certainly the notion of constitution in general, and the idea of objects being constituents of experience in particular, stand in need of clarification, but they are not obviously nonsensical. One immediate clarification is provided by highlighting the modal implications of the relational view: if the character of one's visual experience of a square object is partly constituted by the shape of the object, the subject could not have enjoyed a qualitatively identical experience in the absence of the object or in a situation in which the object was differently shaped (even if, for some reason, it would still have looked square in that situation).

A better response, on behalf of Strawson's rationale, it seems to me, is the following. The response consists of two observations and one suggestion. The first observation is that Snowdon offers no reason to think that the relational view correctly articulates the way commonsense epistemology *in fact* understands the reliability of perception. The second observation is that there does not appear to be any such reason: either of the two explanatory schemes would enable us to demystify the normal truth of perceptual beliefs. The suggestion is that bearing in mind the principle of charity, the causal interpretation of commonsense epistemology may be defended on the grounds of greater parsimony. Seeing a square object in front of one is not the only kind of visual experience that, other things being equal, tends to give rise to the belief that there is a square object in front of one; optical illusions, or visual hallucinations, as of square objects have that same tendency. The causal theory would seem to offer a parsimonious conception of the causal-explanatory link between these sorts of experience and beliefs: experiences falling under one common kind are causally responsible for beliefs concerning the presence of square objects. The relational view, in contrast, is committed to positing at least two different kinds of explanatory

[4] My formulations here draw on Campbell (2002) and Martin (2006). Snowdon speaks of 'visible facts' being 'constituents of the experience'.

relations, corresponding to two different kinds of experience: in the genuine case, the belief is caused by one kind of experience (partly constituted by the perceived environment); in the other cases, by an experience of a different kind. The multiplication of kinds seems inelegant and unwarranted. (The point is familiar from the literature on disjunctivism: see e.g. Sturgeon 1998.)

I conclude, tentatively, that although Snowdon identifies an important lacuna in Strawson's rationale, a causal theorist remains on strong ground so long as our concern is merely with demystifying the reliability of perception. Even if this is disputed, it is clear that Snowdon's response does not (and is not intended to) amount to a positive case for the relational view. But perhaps Snowdon's response concedes too much. I want to argue that there are grounds for scepticism concerning (2); and that these grounds help to make a case for thinking that it is the relational view to which commonsense explanatory practice is committed.

2 A simple view of perceptual knowledge

We can distinguish two readings of (2). I will call these the analytic and the explanatory reading. Let it be agreed, at least for the sake of the argument, that the following are necessary conditions for A's knowing that p: A believes that p; it is true that p; and A is not merely 'flukishly right' in believing that p (Strawson 1974: 79). Suppose we also agree that it follows from this last condition that if perception is a source of knowledge, it has to produce a high ratio of true beliefs. (This is not to say that propositional knowledge can be reductively analysed as reliably true belief.) Then it may be argued that anyone who thinks of perception as the source of a given piece of knowledge is committed to thinking of that knowledge as involving a belief whose truth is unsurprising, given its derivation from perception. Let us call this the analytic reading of (2), and let us suppose that there is much to be said for (2), thus read. (2) might also be associated with a stronger claim. The idea here would be that there is a certain structure in the way we explain knowledge in terms of perception. If we find someone's knowing that p intelligible in the light of their current perceptual experience, this is because we understand two things: that her experience is causally responsible for her believing that p, and that the explanatory relation between the experience and the belief meets a certain general condition required for propositional knowledge (e.g. the condition that the belief not be 'flukishly' correct). On this view, our conception of the role of experience in explaining what someone believes is more basic than our conception of its role in explaining what someone knows: we find the latter *intelligible* in terms of the former. Call this the explanatory reading of (2).

To see why the distinction matters, consider an imaginary commonsense realist who rejects the causal requirement on perception. Strawson reasons that this person would have to regard the normal truth of perceptual beliefs about mind-independent objects as a coincidence, or at least as something inexplicable. He seems to suppose that if we confronted the anti-causalist with the question 'but *how is it* that perception is a

dependable source of true beliefs about the mind-independent world?' she could only shrug her shoulders. Even setting aside the concerns raised by Snowdon, though, this is implausible. Here is something she could say: 'Of course perception is a source of reliably true beliefs—that is entirely unsurprising. After all, perception is a source of *knowledge* (and knowledge entails something like the absence of flukishness).' Note that this simple-minded response respects, indeed appeals to, the conceptual entailment affirmed by the analytic reading of (2). If the response strikes us as inadequate or question-begging this is because we assume the correctness of the explanatory reading of (2). We assume that understanding the epistemic role of perception must be underpinned by an *independent* explanation of the normal truth of perceptual beliefs, an explanation that does not simply help itself to the idea that perception is a source of knowledge.

The simple-minded response points to a further gap in Strawson's rationale for the causal theory. The analytic reading of (2) is plausible but too weak. If there is to be any prospect of reaching (3) it is the explanatory reading of (2) that is needed. The problem is that Strawson offers no support for (2) under that reading. Nor, I want to suggest, does the explanatory reading follow trivially or obviously from the analytic reading. It might be said that if we think of the acquisition of perceptual knowledge as a matter of acquiring reliably true beliefs, it can hardly fail to occur to us to ask what *explains* the normal truth of such beliefs, and once we ask *that* question, simply falling back on the knowledge-yielding role of perception would be an expression of blind dogmatism, hardly more rational than regarding the reliability of perception as 'inexplicable'. But I now want to argue that this line of reasoning begs the central question, of how we do find knowledge intelligible in terms of the subject's perceptual experience.

Consider Strawson's example of a 'non-philosophical observer gazing idly through a window' (1988). Suppose we ask him, not, this time, for a description of the character of his visual experience, but for an explanation of the source of his knowledge of certain facts. Suppose we ask him 'How do you know that the deer are grazing?' One thing he might say in reply is: 'I see that they are grazing.' This reply has its virtues. As Snowdon puts it, we treat it 'as totally unproblematic that someone's knowledge that p can be explained by saying they saw that p' (Snowdon 1998: 301). Part of what makes such explanations reassuring may be that they are, in Quassim Cassam's terms, 'knowledge-entailing'. (See Cassam 2008.) On the other hand, it has to be said that the reply is not particularly informative. The notion of epistemic seeing (or seeing-that) is used in many ways. Even discounting clearly metaphorical uses, appeal to epistemic seeing by itself has relatively little explanatory value. True, when you see that p, the source of your knowledge that p has something to do with vision. But it can draw on much else besides. One way of seeing that your neighbour is at home is to infer her presence from the observed presence of her bicycle. One way of seeing that the piano will not fit through that door is by engaging in a certain kind of imaginative exercise. Appeal to someone's seeing that p naturally prompts the question of how the subject is in a position to see that p—what *enables* her to see that p. Suppose we put this latter

question to our non-philosophical observer. The natural reply is the same reply he gave to Strawson's request for a description of his experience: 'I see the deer.' Or: 'I see the deer grazing.' This provides a more illuminating explanation of the source of his knowledge. The reply leaves some important details implicit, though. Someone may see the deer grazing without being able to recognize the deer or their activity. They may be too far away or lighting conditions may be poor or he may lack the relevant visual recognitional capacities. Making the non-philosophical observer's point slightly more explicit, we might say this: he is able to see that the deer are grazing because (a) he sees the deer, and the event of the deer grazing, (b) he sees certain features of the deer, as well as relations to other objects, and (c) he has the visual recognitional capacities required for recognizing grazing deer and, in virtue of the truth of (a) and (b), is able on this occasion to exercise those capacities in the normal way.

We can summarize this sketch of the commonsense epistemology of perceptual knowledge by saying that we think of object perception as an enabling condition of epistemic perception. (Perhaps the *basic* enabling condition of epistemic perception, in the sense that the condition is indispensable to *any* explanation of how someone is in a position to perceive that p.) Now on the face of it, what is conspicuous by its absence in this picture is the notion of perceptual belief. If the sketch is at all on the right lines, we explain someone's non-inferential perceptual knowledge that p without even touching on the question 'why does he believe that p?'. We exploit explanatory relations between object perception and epistemic perception, and in that way explain the source of his knowledge. The attitude we aim to explain is conceived as someone's possession of a piece of knowledge or, more specifically, as a case of epistemic perception. The explanatory project does not require or involve thinking of the attitude as a case of believing that p. Of course, we can be induced to acknowledge certain entailments, such as that if the observer possesses perceptual knowledge that the deer are grazing he believes that they are grazing, and he must be non-flukishly right in believing this. But it does not follow that we account for his perceptual knowledge *by* explaining how he came by his belief. It is not clear that the structure revealed by conceptual analysis corresponds to any structure in our explanatory practice. If it does not, there will be no reason to suppose that a charitable interpretation of commonsense realism has to equip it with the resources to make sense of the reliability of perception. The question of what explains the normal truth of perceptual beliefs does not normally arise. If it is raised, commonsense realists may answer it by appeal to the fact that perceptual experience yields epistemic perception.

But how does perceptual experience explain knowledge, on the present analysis of commonsense realism? The natural answer is that reference to object perception provides a causal explanation of epistemic seeing and propositional knowledge. At least, this would be the most straightforward reading of the word 'because' in explanations such as 'He knows the deer are grazing because he sees them'. Note that we can manipulate his capacity to acquire knowledge by intervening on his visual experience of the deer. For example, we might intervene on his experience by improving lighting

conditions, which in turn would enable the observer to acquire more detailed knowledge of the deer's activities. The obvious explanation of how this sort of manipulation is possible is that the experience is a causal enabling condition of the observer's coming to know various things. This gloss on the commonsense scheme would suggest that Strawson was right about the crucial importance of causal understanding to the commonsense realist conception of perception as an immediate awareness of mind-independent objects. Where the present analysis—call it the simple view of perceptual knowledge—differs from Strawson's is in envisaging just *one* sort of causal relation, between perceptual experience and knowledge, in place of the two causal relations featuring in Strawson's account, with perceptual beliefs produced by experience and the latter conceived as causally dependent on objects and their features.

Does the simple view of perceptual knowledge correctly articulate ordinary explanatory practice? It is of course not easy to settle this issue decisively. Here I can only sketch what seems to me a promising way to address it. I think an illuminating perspective on the issue is provided by Austin's well-known remarks about the difference between perception and evidence. Austin contrasted two cases: believing that there is a pig in the vicinity on the basis of clues, such as buckets of pig-food or the characteristic noises and smell of pigs, vs being confronted by the pig itself, standing there 'plainly in view'. In the latter case, Austin contended, our belief that the animal is a pig is not based on evidence. Rather, we 'just *see* that it is [a pig], the question is settled' (1962: 115). I think it is plausible that Austin has identified a salient feature of the 'manifest image' of perceptual knowledge. The question is, how should the putative contrast be articulated? According to one suggestive proposal, we credit perceptual experience with a distinctive causal power: we take it to cause perceptual beliefs in such a way as simultaneously to 'silence' any competing considerations. (See Campbell forthcoming.) It is not just that perceptual experience presents peculiarly powerful evidence of the pig's presence, outweighing other bits of evidence. Rather, other considerations cease to be even relevant once you see the pig. This analysis, though, raises the following question: how is it that perception has the power to 'silence' the rest of one's evidence? What is the ground of this remarkable disposition? It is natural to think that the answer has to do with the distinctive credentials of claims to direct perceptual knowledge. Thus it might be said that experience provides a special kind of justifying reason. As Martin puts it, 'for him [Austin], one might suggest, the reason for thinking that there is a pig there is simply the pig itself' (2002: 390). The idea that the source of perceptual knowledge can lie in the perceived object itself is one that seems to play an important role in Austin's discussion. But putting the point in terms of objects providing *reasons for* believing something has two unattractive consequences. One is that we would face the awkward task of explaining the sense in which mere *objects* can be thought of as justifying reasons for belief. The other is that we would have to make sense of the distinctive significance commonsense attaches to what would then appear to be merely a distinction between two kinds of justifying reasons (those provided by objects and those provided by what Austin describes as evidence). The

simple view of perceptual knowledge offers a more straightforward gloss on Austin's contrast. The difference between the two cases turns on our conception of what is *explained* by appeal to evidence or experience, respectively. Possession of evidence concerning the presence of a pig most immediately explains your *belief* that there is a pig nearby (by providing a justifying reason for which the belief may be held). If the evidence is good enough, it may also, in turn, help to explain why your belief qualifies as knowledge. In contrast, your encounter with the pig itself, standing before you in full view, explains your *seeing that* this is a pig. Appeal to the experience yields an immediate answer to the question of how you *know that* the animal is a pig. It is in this way that it 'settles the question'.

3 A rationale for the relational view

Is the simple view of perceptual knowledge committed to a relational conception of the character of perceptual experience? The case for an affirmative answer may be set out as follows.

Suppose the correct account of what constitutes the character of your visual experience of a deer is neutral with respect to whether your experience is a veridical perception or an illusion or a hallucination. Given this account, the fact that your experience has the character it does will not, on its own, provide much of an explanation of how the experience puts you in a position to see that a certain animal in your environment is a deer. It can yield such an explanation, if at all, only in combination with what would be a further fact, that the experience in question is a veridical perception. Call this a composite account of the explanans. There is, of course, some explanatory work that the character of your experience, thus conceived, can do on its own. Even illusory or hallucinatory experiences have the power to generate *beliefs*. Moreover, the fact that your experience is a veridical perception does not cancel this explanatory potential of its character. So a defender of the composite account now faces a question. On her view, the character of your experience figures in two kinds of explanations: in the explanation of your *belief* that there is a deer in front of you and in the explanation of your *seeing* that there is a deer in front of you. How are these explanations related to one another? The natural analysis is that even in the second case, what's causally explained by the character of your experience is your belief. The explanation of your belief forms a proper part of the explanation of your seeing-that. The latter adds to this an account of what makes your belief knowledge. This suggests that a defender of a composite account of the explanans will find it difficult to resist a composite account of the explanandum. The latter, though, is inconsistent with the idea of a primitive explanatory link between experience of objects and epistemic perception, as posited by the simple view.

On the other hand, suppose we reject the 'highest common factor' conception of the character of perceptual experience and adopt a relational conception. Then the idea of a primitive explanatory link becomes intelligible; or at least, reflection on the

explanatory potential of illusions and hallucinations no longer poses a threat to its intelligibility. If the character of your experience is determined or constituted by the real visible features of the deer, it provides for an immediate explanation of how you are in position to see that the deer has certain features. That your experience is veridical will not be a further fact, to be invoked alongside the character of your experience. So there will be no room for a composite account of the explanans, hence no pressure towards a composite account of the explanandum. Of course we still face a question over the relation between explanations of perceptual knowledge and belief. An illusory or hallucinatory experience that is subjectively indistinguishable from your visual experience of the deer can explain the subject's believing that there is a deer in front of her. Indeed one might insist that even in the veridical case, your belief can be made intelligible in the same way as in these other cases. But these points do not show that explanations of perceptual belief are *more basic* than explanations of perceptual knowledge. On the contrary, our understanding of the explanatory link between experience and *belief* may be said to be parasitic on our more basic grasp of the link between (veridical) experience and knowledge. In the case of illusion or hallucination, you have an experience that is subjectively indiscriminable from one in which you see a deer and which would put you in a position to see that there is a deer in front of you. No wonder, we might say, that you end up *believing* there is a deer in front of you—for you intelligibly *take yourself* to be able to *see* that this is so. The explanatory link between the illusory experience and your first-order belief turns on the fact that there is an intelligible link between enjoying an experience that is subjectively indiscriminable from a perceptual experience of a deer and believing that one is able to *see certain facts*, which in turn commits one to certain first-order beliefs, such as that there is a deer in front of one.

How successful is this argument? The argument is reminiscent of claims that have been made in other areas, to the effect that the explanatory role of some relational explanans is not to be understood in terms of a conjunction of facts. For example, it has been argued that the distinctive explanatory role of propositional knowledge cannot be matched by the role of belief conjoined with certain other conditions. (Williamson 1994) This latter argument, though, has no disjunctivist implications. The distinctive explanatory role of knowledge is consistent with the idea that knowledge and 'mere opinion' share a common element. Nor does it require that our understanding of the explanatory role of belief is parasitic on that of knowledge. That argument may block a composite account of the role of knowledge but it does not lead to a disjunctivist picture of belief. One might wonder, therefore, whether reflection on the explanatory role of experience (as conceived by the simple view) can yield a convincing rationale for the relational view of experience—a view notorious for its commitment to a disjunctivist conception of experience. The natural response is that there is a major disanalogy between the two cases. Propositional knowledge is not a type of experience. Questions about the nature of the subjective character of the explanans simply do not arise in the context of reflection on the explanatory role of knowledge. This response

leads to a more serious concern, though. It highlights the importance of an assumption underpinning the whole argument, that everyday explanatory practice assigns a causal role to the character of experience. This assumption certainly requires clarification and defence.

One might suspect that the assumption reads into commonsense explanatory practice a distinctively philosophical doctrine, deriving from the foundationalist tradition in epistemology. Robert Brandom, for example, voices just this suspicion when he characterizes the idea as a 'residually Cartesian intuition'. (Brandom 2002: 98) One might elaborate the suspicion by noting that the range of properties and types whose instantiation we take to be manifest to us in perception far outstrips the sorts of features that can plausibly be held to constitute the character of perceptual experience. It cannot be right that when you gain direct perceptual knowledge that b is F, b's being F must be part of what constitutes the sensory character of your experience. As John Campbell puts it, 'there are plenty of cases in which we have perceptual knowledge of properties that are not making a difference to the nature of our experience' (in press). Recognizing a deer would seem to be a case in point.

To understand the disagreement here it is important to consider what might be the rationale for insisting that we ordinarily regard the character of experience as a relevant factor. I suggest the rationale is this: we take the explanatory link between your visual experience of b and your ability to see that certain things are true of b to be one that is intelligible not just to philosophers or cognitive scientists but, for example, to you. In recognizing a deer it is evident to you not only that the animal is a deer but also that you can see that it is a deer because you see the animal (in good lighting conditions, not too far away etc.). And it is not as if you merely think of your experience as something that, in virtue of some unknown underlying mechanism, has the disposition to facilitate epistemic perception of certain truths. The grounds of that disposition are manifest to you: as you might put it, you can see that it is a deer because it *looks like one*. This does not mean that the character of your experience has to be constituted partly by the species to which the object belongs. Your recognitional capacity can be intelligible to you through its dependence on certain lower-level features of the object that do make a direct difference to your experience—the characteristic shape of the animal, say, its colour pattern, its languid movements etc. One might say that this makes your knowledge depend on inference. I think that is the opposite of the truth. As I will argue in a moment, it is precisely because your recognitional capacity is intelligible to you in terms of the character of your experience that you are able to see without inference that the object is a deer.

The obvious test case for the proposed rationale is the familiar legend of the chicken-sexers, expert perceivers who have been trained to discriminate male from female hatched chicks. As Brandom tells the story, the chicken sexers 'have no idea what features of the chicks they are presented with they are responding differentially to' (2002: 97). They find themselves thinking 'this is female' or 'this is male' but their judgements are not intelligibly related to any features that make a difference to the

character of their experience.[5] Of course they know that they have a reliable recognitional capacity. Otherwise they would attach no authority to their judgements (and would see no point in making them). I think the issue raised by this case is not whether the chicken sexers can have knowledge. Surely they can. The question is whether their knowledge has the same kind of source as ordinary perceptual knowledge. It might be said that it does: the difference is merely that in the ordinary case the source is evident to us insofar as we experience the features enabling us to perceive certain facts whereas the chicken sexers have no such immediate understanding of the source of their knowledge. I think resistance to this picture need not be underpinned by a dogmatic insistence that as a completely general matter, sources of knowledge have to be intelligible 'from within'. Suppose we allow that there might be wholly unreflective chicken sexers (for example, specially trained *chickens*). There is no reason to deny that their recognitional capacity might yield knowledge, despite their utter ignorance of its source. But the fact is that a rational chicken sexer is not like that. Mindful of his special skill, he will use his own spontaneous judgements as *evidence* concerning the sex of a chicken. As we have seen, without this background he could hardly be expected to make such guesses in the first place. Importantly, his grasp of this evidence, i.e. his ability to *infer* the chicken's sex from his spontaneous judgements, has to be a key factor when it comes to explaining the source of his knowledge. It would be quite irrational for him to reach beliefs in this area by relying on guesswork when he is actually in possession of compelling evidence. Such irrational beliefs would not amount to knowledge. Thus a reflective chicken sexer's knowledge will inevitably depend on inference. In this respect it differs from ordinary recognition-based knowledge that a certain animal is a deer. In the latter case, there is no need to rely on evidence. For in attending to the experienced object and its experienced features, you are aware of how you are in a position to acquire knowledge of it: you are aware that you see the object, that it looks like a deer, and that in this way you are able to see that it is a deer. There is no need to rely on evidence precisely because your capacity for epistemic perception is intelligible to you in this way. And of course there is no irrationality in your acquiring the non-inferential belief that it is a deer when it is evident to you how you can *see* it is a deer. That your recognition relies on certain experienced lower-level features does not necessarily make the resulting knowledge inferential. For it is not as if you first judged the animal to have a certain definite shape and colour and then concluded that therefore it is probably a deer. Your best effort at articulating the set of experienced features to which you are responding may well be in terms such as 'the characteristic shape and colour of a deer'. Your identification of the set of features may well draw on the very non-inferential recognitional capacity they help to make intelligible.

[5] Contrary to the philosophical legend, real chicken sexers do find their judgements intelligible in terms of their experience—specifically, in terms of the visible shape of the chicks' genital eminences. (Males tend to be described as 'round', females as 'pointy'.) See Biederman and Shiffrar (1987).

4 Knowledge and explanation

There are several reasons why it can seem hard to take seriously the idea of a direct causal-explanatory connection between perceptual experience and knowledge. I want to conclude by looking at two sources of resistance to the idea. We may call them the charge of psychological and epistemological irrelevance, respectively.

The first source of resistance stems from a particular conception of psychological explanation, encapsulated in Stephen Stich's remark that 'what knowledge adds to belief is psychologically irrelevant' (Stich 1978: 574). Discussion of this claim tends to focus on the role, real or apparent, of knowledge as an explanans in causal explanation; its role, for example, in explaining intentional actions. But Stich's slogan also bears on the role of knowledge as an explanandum of causal explanations. If the ideas behind the slogan are correct, apparently direct causal explanations of knowledge had better be something else, for example causal explanations of *belief*, in terms that simultaneously illuminate the epistemic status of the belief. I will not enter this debate here but it seems to me that the 'internalist' conception of psychological explanation has no more plausibility in the case of action explanation than in the case of explanations of perceptual knowledge. One way to challenge it, in the action case, is to advert to the distinctive sets of counterfactuals sustained by explanations in terms of knowledge, bringing out their distinctive explanatory value (Williamson 1994). A more radical move is to insist on the indispensable role of knowledge in reason-giving explanation: only explanations in terms of factual states, such as knowledge, accord an explanatory role to the facts that constitute justifying reasons for action (Hornsby 2008). It is not implausible to think that if we have learned to live with the psychological relevance of knowledge in the action case, we should not be too flustered by the idea of causal explanations *of* knowledge in terms of perceptual experience.

I think it is the charge of epistemological irrelevance that poses the deeper and more serious challenge to the simple view. In a discussion of the aspirations of a philosophical explanation of perceptual knowledge, Barry Stroud makes the following unflattering remarks about the explanatory scheme of what I have called the simple view:

> To say simply that we see, hear, and touch the things around us and in that way know what they are like, would leave nothing even initially problematic about that knowledge. Rather than explaining how, it would simply state that we know. There is nothing wrong with that; it is true, but it does not explain how we know even in those cases in which (as we would say) we are in fact seeing or hearing or touching the object. (Stroud 2000: 145)

It might be said that Stroud overstates the point. If you want to know how I know the deer are grazing, there *is* something that is initially problematic about my knowledge, viz. how I got it. If I say 'simply that I see the deer', or perhaps that I see (the event of) the deer grazing, this should put your mind to rest. Without some such account of the source of my knowledge, you cannot be sure that I was not merely guessing. My simple answer should do something to convince you that I am not; it can help to

establish the credentials of my claim to knowledge. Of course, you may still be dissatisfied. My answer falls short of a full account of the enabling condition of my seeing that the deer are grazing, nor does it address potential worries about the satisfaction of relevant *disabling* conditions. What these points illustrate is that the explanatory role of object perception depends on various sorts of background conditions. But they provide no support for the prima facie surprising conclusion that appeal to object perception has no explanatory value whatsoever—that it fails to 'explain how we know even in those cases in which (as we would say) we are in fact seeing or hearing or touching the object'.

As Stroud himself emphasizes, it is relative to a certain *philosophical* conception of what counts as a good explanation of perceptual knowledge that the stronger conclusion appears plausible and perhaps inescapable. There are a number of philosophical reasons for thinking that the simple view is completely unsatisfactory, and that any genuine explanation of perceptual knowledge has to consist of an account of the explanatory relation between perception and *belief*, along with an account of how beliefs that are thus related to perception earn the 'status of knowledge'. These considerations support the following conditional: *if* object perception does indeed play some role in giving us knowledge of mind-independent objects it has to be possible to make that role philosophically intelligible by tracing it to the way object perception helps to satisfy certain general conditions on propositional knowledge. Now this demand is not trivial. For reasons discussed earlier, it cannot be motivated simply by reference to the entailment between knowing that p and believing that p. The question is why an *explanation* of someone's knowing that p should be expected to proceed via an explanation of her believing that p. One suggestion at this point might be that the answer lies in the 'slightest philosophy' Hume declared to be sufficient to expose the vulgar error that we are directly aware of mind-independent objects. What the arguments from illusion or hallucination show, it might be said, is that perceptual experience, subsuming as it does both veridical and illusory cases, can be *directly* explanatory only of belief, not knowledge. Arguably, though, the philosophical demand has sources that are deeper and perhaps less easy to resist. A central theme of Stroud's work is the connection between the traditional epistemological project and what is sometimes taken to be an implication or requirement of philosophical realism, the possibility of a *completely general* account of the credentials of claims to knowledge of the mind-independent world.[6]

But I think we can see the force of the philosophical demand even without a full understanding of its philosophical sources. The simple view appears to be committed to a position Quassim Cassam calls minimalism. According to minimalism, 'the connection between knowledge and perception is primitive and cannot be explained any further.' Cassam argues that minimalism is implausible, given that 'surely we want to

[6] See Stroud (1984: esp. ch. 2) and Stroud (2000: esp. chs. 8 and 10). Martin (2006: part 3) offers an illuminating discussion of Hume's 'slightest philosophy' and its relation to Cartesian scepticism.

say that perceiving is a way of knowing because, and only because, there are more general conditions on knowing that p that one satisfies in virtue of perceiving that p' (2008: 20). Why do we want to say this? One good reason is dialectical. It is not just that minimalism would deprive us of the intellectual satisfaction of understanding what it is in virtue of which the connection between perception and knowledge holds. We should like to think of our confidence in the epistemic role of perception as reasonable confidence, different in kind, say, from the confidence some people place in clairvoyance. The obvious way to debunk some putative source of knowledge is to show that beliefs flowing from that source violate certain completely general necessary (and jointly sufficient) conditions on knowledge. Correlatively, the obvious way to establish the epistemic credentials of perceptual experience is to satisfy ourselves that it does provide for the satisfaction of the conditions in question.

Cassam argues for what he calls a middle way between minimalism and the idea that 'a prior reductive analysis of the concept of knowledge is needed to make it intelligible that perceiving is a way of knowing' (2008: 21). The middle way proposes a limited elucidation of the epistemic role of perceptual experience, by reference to (merely) necessary conditions for knowledge. This is certainly plausible in as much as a good account of the necessary conditions for knowledge (e.g. reliability) would suffice to debunk certain putative sources (e.g. clairvoyance). Nevertheless, the middle way would fail to make the epistemic role of perception fully transparent to us.[7] One way in which this lack of transparency manifests itself is this. As far as the middle way is concerned, there could, on the face of it, be two sources of knowledge that are indistinguishable as regards the way they provide for the satisfaction of the necessary conditions for knowledge, yet intuitively count as importantly different sources. Consider the intuitive difference between knowing that a certain object is a pig by visually experiencing the object and knowing this on the basis of something like blindsight. On the face of it, the difference between explanations of knowledge in terms of the two sources could hardly be more profound. In one case, you know because the pig is there right in front of you, 'plainly in view'. The explanatory link between your visual experience and knowledge is peculiarly direct and intelligible. It is part of the 'pre-theoretical scheme of commonsense realism'. In the other case, you know because your guesswork is unwittingly controlled by your visual system. The explanatory link between blindsight and knowledge, such as it is, is part of a theory that is fully intelligible only to professional cognitive scientists. What is not clear, though, is that the difference between the two sources can be explained in terms of how they provide for the satisfaction of general necessary conditions for knowledge. Suppose it cannot be so explained.[8] The middle way, being concerned only with necessary conditions for knowledge, would have to acknowledge that the two sources are

[7] Cassam acknowledges that the middle way is 'very close to minimalism', insofar as it 'helps itself to the intuitive distinction between good and bad answers to the question "How does S know?"' (2007: 22).

[8] See Roessler 2009 for more detailed discussion of these issues.

distinct, but would provide no help in elucidating their difference. We would be reduced once more to the minimalist mantra that 'the connection between knowledge and perception (specifically, perceptual *experience*) is primitive and cannot be explained any further'.

There is no time for a detailed discussion of the numerous issues raised by minimalism. But I want to end by sketching a suggestion. It may be possible to go beyond minimalism, not by meeting the philosophical demand (even halfway), but rather by presenting grounds for resisting it. A basic assumption informing the demand may be put by saying that the concept of belief enjoys a certain explanatory independence, in the following sense. Understanding how someone knows what they know about a certain matter requires understanding the epistemic status of their beliefs in this area; specifically, understanding what confers the status of knowledge on their beliefs. An essential first step, then, is to identify the beliefs whose epistemic status is to be examined. Importantly, this is only a first step. As far as an explanation of the source of the knowledge in question is concerned, nothing has been settled by identifying the relevant beliefs. We can think about these beliefs in purely 'doxastic' or 'psychological' terms. Whether and how they have secured the status of knowledge will be a substantive further question, to be addressed by investigating whether they satisfy the general conditions for knowledge. Now in any particular case this seems a very sensible assumption to make. The question that I think is worth pressing is whether the assumption is still plausible when we reach the dizzyingly high level of generality at which the philosophical demand operates. Here is one reason to think the matter may not be straightforward. Consider what is involved in attributing to oneself and others beliefs involving perceptual demonstratives, such as 'this lemon is yellow'. To think of you as holding such a belief it is not enough to say that you believe of a certain object that it is yellow. We need to acknowledge that you think of the thing in a distinctive way, made available to you by your perceptual experience of it. Thus in the case of beliefs involving perceptual demonstratives, the idea of an explanatory link between your experience and your belief is integral even to the preliminary project of identifying the beliefs to be examined. (See Campbell 2002.) Suppose this is right. Then two questions arise. Can we coherently think of experience as a source of understanding perceptual demonstratives without simultaneously thinking of it as providing a good explanation of propositional knowledge? And if we cannot, what would be the implications for the philosophical demand?

Evidently grasping a perceptual demonstrative is consistent with having a lot of false beliefs about the object in question. It may even be consistent with having no propositional knowledge whatsoever of the object. Nevertheless, grasp of a perceptual demonstrative may essentially involve being in a position to gain some propositional knowledge of the object through one's experience of it. One way to support this claim would be to argue that perceptual demonstrative identification requires not just experience but selective attention, putting one in a position (at least under favourable conditions) to keep track of the object, and hence to gain propositional knowledge of

its boundaries and identity over time. The upshot of such an argument would be that recognizing the explanatory link between experience and perceptual demonstratives commits one to acknowledging an explanatory link between experience and propositional knowledge. Then the second question is whether this result would be consistent with accepting the conditional demand, that if there is a genuine explanatory link between perceptual experience and knowledge, it has to be possible to make it intelligible in the light of the way experience provides for the satisfaction of the general conditions for knowledge. On the face of it, the result would call into question what I suggested is a basic assumption behind this demand, that identifying perceptual beliefs is merely a first step towards explaining how we know what we know. It would be mere pretence to proceed as if the idea that perceptual experience is a source of knowledge is to be justified (or debunked) by an investigation of whether and how experience helps to satisfy general conditions on knowledge. If we get as far as identifying the beliefs whose epistemic status is to be scrutinized we would already be committed to recognizing the explanatory link between perceptual experience and propositional knowledge. We would be so committed in virtue of the role experience plays in presenting us with the world around us and enabling us to have perceptual demonstrative thoughts about objects, independently of any underlying conception of how experience is linked to the general conditions for knowledge.[9] This would not prove minimalism correct. But it would disarm one kind of resistance to minimalism, and hence to the simple view of perceptual knowledge and the relational conception of experience.

References

Austin, J. L. 1962: *Sense and Sensibilia*. Oxford: Oxford University Press.

Biederman, Irving and Shiffrar, Margaret 1987: 'Sexing Day-Old Chicks: A Case Study and Expert Systems Analysis of a Difficult Perceptual-Learning Task', *Journal of Experimental Psychology: Learning, Memory and Cognition* 13, pp. 640–645.

Brandom, Robert 2002: 'Non-inferential Knowledge, Perceptual Experience, and Secondary Qualities: Placing McDowell's Empiricism', in N. Smith (ed.), *Reading McDowell*. London and New York: Routledge.

Campbell, John 2002: *Reference and Consciousness*. Oxford: Oxford University Press.

——in press: 'Visual Attention and the Epistemic Role of Consciousness', in C. Mole, D. Smithies, Wayne Wu (eds.), *Attention: Philosophical and Psychological Essays*. Oxford: OUP.

——forthcoming: 'Does Perception Do Any Work in an Understanding of the First Person?', in Annalisa Coliva (ed.), *The Self and Self-Knowledge*. Oxford: Oxford University Press.

[9] This account of the source of our commitment would not imply that the philosophical demand could not be met. (One elegant, if implausible, way to meet it would be to argue that something like acquaintance or intuition is a general condition for propositional knowledge.) But it would undermine an outlook that takes the rationality of the commitment to be dependent on the possible satisfaction of the demand. See Roessler 2009 for further discussion of this line of argument.

Cassam, Quassim 2007: 'Ways of Knowing', *Proceedings of the Aristotelian Society* 107, 339–358.
—— 2008: 'Knowledge, Perception and Analysis', *South African Journal of Philosophy* 27.
Child, William 1994: *Causality, Interpretation, and the Mind*. Oxford: Clarendon Press.
Hornsby, Jennifer 2008: 'A Disjunctive Conception of Acting for Reasons', in A. Haddock and F. Macpherson (eds.), *Disjunctivism*. Oxford: OUP.
Martin, Michael 2002: 'The Transparency of Experience', *Mind and Language* 17, 376–425.
—— 2006: 'On Being Alienated', in T. Szabó Gendler and J. Hawthorne (eds.), *Perceptual Experience*, pp. 354–410.
—— 2009: 'The Limits of Self-Awareness', in A. Byrne and H. Logue (eds.), *Disjunctivism*. Cambridge, Mass.: MIT Press.
Roessler, Johannes 2009: 'Perceptual Experience and Perceptual Knowledge', *Mind* 118, pp. 1013–1041.
Snowdon, Paul 1998: 'Strawson on the Concept of Perception', in L. E. Hahn (ed.), *The Philosophy of P. F. Strawson*. Chicago and Lasalle, Illinois: Open Court.
Stich, Stephen 1978: 'Autonomous Psychology and the Belief-Desire Thesis', *The Monist* 61, 573–591.
Strawson, Peter F. 1974: 'Causation in Perception', in *Freedom and Resentment*. London: Methuen.
—— 1988: 'Perception and its Objects', in G. McDonald (ed.), *Perception and Identity: Essays presented to A. J. Ayer*. London: Macmillan, pp. 41–60 (1979). Reprinted in J. Dancy (ed.), *Perceptual Knowledge*. Oxford: OUP.
—— 1992: *Analysis and Metaphysics*. Oxford: OUP.
—— 1998: 'Reply to Paul Snowdon', in L. E. Hahn (ed.), *The Philosophy of P. F. Strawson*. Chicago and Lasalle, Illinois: Open Court.
Stroud, Barry 1984: *The Philosophical Significance of Scepticism*.
—— 2000: *Understanding Human Knowledge*. Oxford: Oxford University Press.
Sturgeon, Scott 1998: 'Visual Experience', *Proceedings of the Aristotelian Society* 80, 185–216.
Williamson, Timothy 1994: 'Is Knowing a State of Mind?', *Mind* 104, 533–565.