

13

Perception, Causal Understanding, and Locality

*Christoph Hoert**

1 Introduction

Locality, as I will understand the term in this paper, is the view that causal relations are always subject to spatial constraints in the following sense: For any case in which there is a causal relation between two spatiotemporal particulars,¹ that causal relation would not have obtained if certain spatial parameters of the situation had been different (e.g. had one of the two particulars been located elsewhere).

The issue as to whether locality is true is arguably central to recent philosophical debates about causation. Theories that seek to analyse causal relationships in terms of the notion of a mechanism, or so-called process theories of causation appear committed to locality. Put briefly, such theories maintain that causation involves the presence of a spatiotemporally continuous empirical feature connecting cause and effect, i.e. a feature which is also present at points intermediate between cause and effect—be this a concrete object travelling between cause and effect, or radiation, or the existence of a magnetic or gravitational field (Salmon, 1984, pp. 208ff.).² This rules out what is typically referred to as action-at-a-distance, that is, alleged cases of causation that violate locality. Thus, such theories rule out, for instance, the idea that a ‘powder of

* An early version of this paper was presented at Sydney University, a later one to the Wednesday philosophy group at Warwick University. By and large, the comments I received on each of the two occasions came from two very different sides, but I found them equally helpful. I am also grateful to the AHRC for its support—both for the AHRC Project on Causal Understanding, in the context of which I was prompted to write this paper, and for a term’s research leave, during which much of the writing took place.

¹ I am using the phrase ‘a causal relation’ in a loose sense here, which is not meant to imply a commitment to a specific philosophical view of the ontology of causation. The kinds of cases I have in mind might, for instance, be described by saying that x’s ϕ -ing caused y’s ψ -ing, or by saying that the fact that x ϕ -ed caused the fact that y ψ -ed, where x and y are the relevant spatiotemporal particulars. Thus, localism, as I understand it, is a thesis about token or singular causation, but it is otherwise neutral as to what entities such causation ultimately relates (e.g. events vs facts).

² This characterisation will do for the moment. In section 3, I will suggest that the key intuitions behind causal process views of causation can also be accommodated by a slightly more liberal formulation.

sympathy', applied to a weapon that had caused a wound, could heal the wound purely in virtue of the past contact between the knife and the wound, and irrespective of where the wounded person was located now.³

By contrast, if regularity theories or counterfactual theories of causation are correct, the truth of locality, and the non-existence of cases of action-at-a-distance, seem at best interesting empirical facts about our world. The truth of locality has nothing to do with what, according to these theories, causation is. Indeed, the alleged conceivability of action-at-a-distance, as illustrated by the example above, has sometimes been claimed to provide an argument in favour of adopting one of these latter types of theory.

In this paper, I want to suggest that we can extract from the existing philosophical literature on perception, and perceptual judgement, materials for an argument to the effect that a commitment to locality *is* part of our common-sense concept of causation. I will also suggest that this literature can help to bring out a specific kind of role this commitment plays in our reasoning about causal relationships. However, in order to make the connection between the two different areas of philosophy—the philosophy of perception and the philosophy of causation—we also need to have in place a particular way of thinking of what is at stake, in general, in philosophical debates on causation. Thus, a second aim of the paper is metaphilosophical: to articulate and defend a specific philosophical approach to causation.

Along with other philosophers working on causation, I will take my task to be that of providing (part of) an account of the meaning of 'our' concept of causation. In the existing literature, authors often explicitly state their aim in those terms.⁴ However, it also seems implicit in the predominance of a particular type of methodological tool: this is an area of philosophy where the use of examples and thought-experiments proliferates, i.e. where many of the arguments proceed by trying to show that the theory put forward is one that best accords with our intuitive causal judgements about a range of actual or imagined scenarios.⁵ Use of this method would seem to be to the

³ Another example would be the idea that a person could be brought to harm by damaging an object representing them, purely in virtue of the representational relationship between the object and the person. It is important to note, though, that not all historical accounts of such 'sympathetic' action do in fact conceive of it as involving action-at-a-distance in the sense at issue here. (Thanks to John Sutton for pointing this out to me.) For instance, Kenelm Digby's (1658) account of the effectiveness of the 'powder of sympathy' is cast entirely in terms of corpuscular-mechanical principles. According to him, as the powder mixes with the blood on the weapon, particles of medicine mixed with blood are released in the air, which then find their way to the wound (see Mercer, 1998). For more discussion of the notion of action-at-a-distance, in the philosophically relevant sense, see Hall (2002) and Frisch (2010).

⁴ One exception here is Dowe (2000), who describes himself as offering an 'empirical analysis of causation' (ibid., p. 2), aimed to uncover 'what causation is in the objective world' (ibid., p. 1). He specifically contrasts this with the task of 'elucidat[ing] our normal concept of causation' (ibid., p. 1). The latter, he claims, is exemplified by Strawson's (1959) project of 'descriptive metaphysics'. Some of what follows is inspired by Strawson's project, and I will make use of Dowe's theory in a way that is against his own stated metaphilosophical intentions. However, I think this would be objectionable only if I ended up criticising Dowe's theory for failing to do a job that it is not intended for, which I will not.

⁵ Collins, Hall, and Paul (2004) include a separate index for the examples used in that volume, which runs to several pages.

point only if the target of the theory is the concept of causation mobilised in such intuitive judgements.

Yet, whilst the widespread use of examples and thought-experiments in the philosophical literature on causation may indicate that the target explanandum is indeed our common-sense concept of causation, we can still ask how good they ultimately are as a tool for probing that concept. When the question as to what it is to have a certain concept is raised in philosophy, it is usually assumed that a subject can possess a concept and yet misapply it on occasion. If this is true, perhaps some of the scenarios appealed to in the existing literature mislead us about the nature of our concept of causation, if there are features of such scenarios that make us more prone to misapplying that concept. Suppose, on the other hand, we take the particular causal judgements people do in fact make as the ultimate arbiter in determining the nature of their concept of causation. In that case, it is unclear whether we can still be said to be carrying out a distinctively philosophical project of conceptual analysis. Instead, we seem to have strayed into the territory of anthropology.

Put very broadly, what we need, it seems, is some principled way of singling out specific types of reasoning about causation as central to our possessing a concept of causation, or at least a way of distinguishing between more or less central types of such reasoning. This, I believe, is where some thoughts familiar from the literature on perception might be brought to bear on the debate, in the philosophy of causation, about the status of locality.

Beyond these programmatic remarks, to see how exactly the connection might be made, some stage setting is required. In the next section, I present one way of looking at a key debate in contemporary work about causation. In section 3, I argue that the (or, at any rate, a) central issue on which that debate turns is best described in terms of the question as to what place (if any) locality has in our ordinary thinking about causation. In section 4, I discuss a suggestion recently made by James Woodward as to how questions such as this are best approached, and offer an alternative. It is not until section 5 that I get to the link with the philosophy of perception. In that section, I draw on an influential type of argument to the effect that perceptual judgement recruits a certain basic form of causal understanding. In the context of the present paper, the aspect of the argument that will be of particular interest is that the causal understanding that, according to the argument, is recruited in perceptual judgement does seem to involve a commitment to locality. In the final two sections, I show how analogous arguments in other areas might be used to draw a more general conclusion.

2 The debate between difference-making and causal process theories of causation

Contemporary philosophical debates on causation are dominated by the two broad types of theoretical approach to causation that I mentioned at the beginning of this

paper, which are sometimes called difference-making and causal process approaches, respectively. Both types of approach have been developed to a considerable level of sophistication, but the basic idea behind each is relatively simple to explain. What is proving to be a rather recalcitrant problem, by contrast, is how exactly the dialectic between the two types of approach is to be understood. There is an obvious sense in which difference-making and process theories make quite different claims about causation. What exactly the difference between them comes to, however, depends on what we take to be their respective explanatory aims. Indeed, it has been suggested that there may be no conflict between them, because they are trying to account for two different types of thing (see the opening methodological remarks in Dowe, 2000, and Woodward, 2003, for discussion).

In what follows, I will suggest that perhaps a more fruitful way of construing the dialectical situation between difference-making and causal process theories of causation is by seeing them as articulating two conflicting sets of ideas about the meaning of ordinary causal thought and talk.

According to one prominent way of thinking, a key part of the task of a philosophical theory of causation is to provide an account of what distinguishes causation from mere correlation. Typically, when the need to distinguish causation from correlation is raised in the context of scientific investigation, the thought is as follows. A statistical correlation between A and B implies only that there is *some* sort of causal connection between the two; but it is compatible with a number of quite different causal scenarios—e.g. A causing B, B causing A, or both A and B being effects of a common cause.⁶ This is the key problem that has led, for instance, to the development of the Bradford Hill criteria in epidemiology (Hill, 1965). These are concerned with the evidence on which we ought to decide which particular causal scenario is present in a given case. However, we can also ask what the difference between these different causal scenarios consists in in the first place.

As I will interpret them, process theories of causation start off from a line of thought that runs roughly as follows (though see the next section for an important qualification). If the choice is, e.g. between a situation in which A causes B and a situation in which A and B are effects of a common cause, A causes B only if there is a process, a spatiotemporally continuous empirical feature, that directly links A with B. In the relevant sense of ‘directly’, there is no such process linking A directly with B if A and B are effects of a common cause; instead there are separate processes linking the common cause with A and with B. Further details will need to be added, e.g. to distinguish between a situation in which A and B are effects of a common cause, and a situation in which A causes an event which, in turn, causes B. But the basic thought is that spelling out these details is just a matter of characterising in more detail the process in operation (e.g. introducing the idea of a direction intrinsic to a process). It is thus that a person’s

⁶ I leave to one side here the possibility of statistical correlations in the absence of any sort of causal connection.

grasp of the notion of a process is implicated in their ability to distinguish mere correlation from causation; or so the thought goes.

As soon as the basic thought behind process theories of causation is spelled out in these terms, however, it may look as though the proponent of such a theory starts off on the back foot. For one thing, it may seem that the general notion of a process, as used in the preceding paragraph, is little more than a placeholder. For another, there are in fact a variety of philosophical accounts that seem to be able to distinguish, say, between a cause–effect relation and a common cause relation, without making any appeal to the notion of a process, thus understood.⁷ Instead, the distinction is couched in terms of the idea that a particular kind of counterfactual relationship holds between A and B only if A causes B, and is absent if A and B are effects of a common cause. The most developed account of this kind is typically referred to as interventionism or manipulationism. The basic intuition behind interventionism is that for A to cause B is for there to be an association between A and B that remains invariant under (a range of) interventions on A—where this is to be contrasted specifically with three other cases in which any association between A and B will be broken by interventions on A: (a) there is no causal connection; (b) B actually causes A, rather than vice versa; (c) A and B are joint effects of a common cause. Of course, a lot more has to be said about the key notion of an intervention here, but I will assume that we can find the necessary refinements in the work of theorists such as Pearl (2000) and Woodward (2003), who have offered detailed accounts of that notion. In other words, I will assume that interventionism does indeed provide a way of distinguishing, say, between the case in which A causes B and the case in which they are both effects of a common cause, without using the notion of a causal process.

Where does this leave the dialectic between difference-making and causal process accounts of causation? According to the process theorist, our common-sense view of causation involves a metaphysical commitment to a particular kind of empirical feature the presence of which makes causal statements true, viz. a causal process. This, by itself, does not force the process theorist to deny that interventionism can provide a way of distinguishing a case in which A causes B from a case in which both are the effects of a common cause without making reference to that empirical feature. What the process-theorist will maintain, though, is that interventionism should be seen as giving an account of certain (idealised) ways of deciding whether certain causal statements are true, rather than of what makes those statements true. In other words, the process theorist can invoke a distinction between an epistemic account of the methods by which to arrive at certain causal judgements, on the one hand, and our grasp of what it

⁷ At least that is what I will assume for the sake of the argument. One stubborn problem for counterfactual accounts such as Woodward's is how to account for certain cases of preemption. Woodward's own analysis of such cases involves the idea of an intervening variable in between the cause and effect. The idea that there *must* be such intervening variables, however, seems to involve theoretical commitments that go beyond the basic idea of analysing causation in terms of certain counterfactuals and is in tension with the view that, e.g. action-at-a-distance is possible.

is for those judgements to be true, on the other. In this respect, from the causal process theorist's point of view, the debate between causal process theories and difference-making theories of causation resembles the debate between, say, a realist concerning statements about the past (of the type represented by, e.g. Peacocke, 2005), and a justificationist approach to such statements (of the type represented by Dummett, 2005).

From the point of view of an interventionist, by contrast, the existence of relationships that are invariant under interventions is not just evidence for the existence of causal relationships, but is the very feature of reality that causal talk aims to get at. The interventionist's diagnosis of where the process theorist goes wrong runs roughly as follows. There are, the interventionist can admit, legitimate ways of spelling out the notion of causal process, such that many instances of causation will in fact involve the operation of such processes. In the interventionist literature, instances of causation that do involve causal processes in the relevant sense are often also referred to as ones where there is a *mechanism* present. The interventionist can also admit that there are various reasons why identifying such mechanisms can be of scientific interest. Yet, he will insist that the existence of a mechanism mediating the causal connection between A and B is simply a matter of there being separate parts of the relationship between A and B which can themselves be described in terms of interventionist counterfactuals (cf. Woodward, 2007). The mistake made by the process theorist, thus, is to confuse our legitimate interest in mechanisms with an insight into what makes the relevant relationship causal. If A causes B, the existence of a mechanism between A and B is simply a (contingent) further fact about the causal relation between them. Indeed, on this way of construing the notion of a mechanism, it is actually ruled out that causal relationships are always mediated by mechanisms. As Campbell (2007, p. 65) puts the view, '[e]xplanation by means of mechanisms must bottom out somewhere, and then we are left with the bare facts about what would happen under interventions'.

3 Locality and negative causation

I have sketched one way of understanding of what is at stake in the debate between difference-making and causal process approaches to causation. Put schematically, according to the first type of approach, our common-sense concept of what it is for A to cause B involves a commitment only to the holding of certain counterfactual relationships involving A and B. According to the second, the relevant counterfactuals point to some of the reasons why we might endorse the claim that A causes B; to capture what it is for A to cause B, though, we have to appeal to the idea of a causal process.

In the previous section, I spelled out the latter thought in a way that implies a commitment, on the part of the causal process theorist, to the idea that causes are physically connected to their effects, as it is sometimes put. In other words, as I have characterised the causal process theorist so far, she holds that appealing to the idea of a

causal process mediating between cause and effect must take the form of claiming that there is an actual spatiotemporally continuous empirical feature connecting cause and effect. In fact, however, I believe that the substance of the dialectic between difference-making and causal process approaches that I have sketched is not well described by saying that it turns on the question as to whether causes are physically connected to their effects in this sense. The dialectic still arises for a weaker form of causal process theory that is not committed to this claim, which I will outline below. Correspondingly, I will suggest that rather than concentrating on the question as to whether causes are physically connected to their effects, a better question to focus on in evaluating the two sides of the dialectic is that of locality.

Against specific theories fleshing out a causal process approach, such as the conserved quantity theory (Salmon, 1984; Dowe, 2000), it is sometimes held that there is a class of cases which common-sense seems happy to classify as cases of causation, but which do not involve any physical connection along which a physical quantity such as energy or momentum could be transmitted between cause and effect. The cases at issue here are ones that exemplify what is sometimes called ‘negative causation’, that is, causation in which absences act as causes or effects. A paradigmatic example are cases of double prevention, such as the following. I remove from the side of a road a traffic sign that warns of a narrow bend further down the road; minutes later a car takes the bend too fast and crashes, because the driver was unaware that there was a bend coming up. There is an intuition here that I caused the accident, or was at least in some sense causally involved—an intuition that might manifest itself in the feeling that I should be punished for what I did.⁸ Yet, were we to ask how I caused the accident, we could point to nothing concrete in the space between where the traffic sign used to be and where the car is travelling that could provide the answer; there is no physical connection between my actions and what happened to the car.⁹

One option for a defender of a causal process theory, in response to examples of this kind, is to deny that they are *bona fide* examples of causation. Dowe (2000), for instance, takes this line, arguing that intuitions in fact vary across different cases of alleged causation involving absences. However, if our main interest is in the dialectic between difference-making and causal process approaches as described in the previous section, it is not obvious that siding with a causal process approach requires taking this (somewhat heroic) line.

As Dowe himself notes, in as far as there is indeed an intuition, in response to examples of double prevention like that given above, that we are dealing with a case of causation, it is because the initial action disrupts a physical connection that would otherwise have been in place. To make intelligible the sense in which I am causally

⁸ Examples of structurally identical cases that invite even stronger intuitions (if they are not strong enough in this case) can be found in Woodward (2003) and Schaffer (2004).

⁹ Or so I will assume. Frisch (2010) explores one interpretation of the notion of a field that might allow us to speak of the obtaining of a physical connection even in cases such as this.

involved in the occurrence of the car crash, it seems we have to bring in the idea that, had I left the traffic sign alone, there *would have been* a physical connection between the sign and the subsequent behaviour of the driver (involving, e.g. the physical processes that make it possible for the driver to *see* the sign), such that the car would have made it safely around the bend.

What this observation suggests, in my view, is that a simple-minded version of a causal process approach, which demands that cases in which A causes B must involve an actual physical connection between A and B, is untenable, but that we can nevertheless hold on to the spirit of a causal process approach if we formulate the relationship between the idea of a causal process and that of a physical connection in slightly more complex terms.

Part of what causal process accounts of causation are trying to get at, on the view I have in mind, is the idea that, with respect to any case in which A causes B, we assume that it makes sense to raise the question as to *how* A causes B. On this view, causal process accounts can accommodate cases of negative causation, provided that we allow that the answer to the how-question sometimes needs to be framed in counterfactual terms. Note that the relevant counterfactuals here are still quite different from the counterfactuals appealed to by proponents of difference-making approaches such as interventionism. The counterfactuals needed to explain our causal intuitions regarding cases of negative causation seem to be able to do so only because they mention physical connections that would have obtained if a certain event had not taken place. In other words, we can hold on to the view that causal relations always involve causal processes, and also that detailing a causal process by which A causes B is always a matter of appealing to some sort of physical connection, as long as we allow that the latter can sometimes take the form of appealing to a physical connection that would have obtained had a certain event not happened. Thus, even though the resulting theory has to introduce a counterfactual element into the analysis of negative causation, it is still a theory that takes the notion of a physical connection as primitive to our understanding of causation.

On this way of dealing with cases of negative causation, one particular issue emerges as central to the debate between causal process approaches and difference-making approaches to causation, viz. the issue of locality. Note that there is one crucial feature that cases of negative causation have in common with cases of positive causation, in which there is an actual physical connection between cause and effect: due to the role that actual or counterfactual physical connections play in each case, it will be true for both of them that A would not have caused B, had certain spatial parameters of the situation been different. As I will also express it, on the kind of view I have sketched above, both positive and negative causation are subject to spatial constraints. By contrast, at least on the face of it, there is no room for mentioning spatial constraints in counterfactual theories such as Woodward's interventionism. In its most simple form, interventionism simply says that what there is to causation is the holding of certain interventionist counterfactuals. It seems to be a consequence of such a view that it

should be possible for there to be causal relationships that are not subject to any spatial constraints—the theory itself makes no mention of space at all.¹⁰

Thus, one way to pursue the basic dialectic between difference-making and causal process approaches to causation, understood along the lines sketched in the previous section, is by asking whether we need to see it as an integral part of our concept of causation that causation is always subject to spatial constraints, or whether we should treat this feature, which is present in the cases of both positive and negative causation that we are familiar with, as a dispensable feature. Specifically, my interest in the following will be in the question as to the grounds on which one might adopt one rather than another answer to this question.

4 ‘Practical’ and ‘impractical’ accounts of causation: A misleading dichotomy

I have suggested that we should think of one, if not *the*, key issue at stake in the debate between causal process accounts and difference-making accounts of causation in terms of the question as to whether a commitment to locality is part of our common-sense concept of causation, i.e. whether that concept implies that causal relations are always subject to spatial constraints. It might of course be agreed upon on all sides that, as a matter of fact, the only kinds of causal relations we are familiar with are ones that involve such spatial constraints, and that action-at-a-distance, violating locality, does not in fact exist. The issue, though, is whether this is merely a contingent, empirical matter regarding the world we live in, or whether action-at-a-distance is somehow ruled out by our very concept of causation.

In outlining a possible defence of the claim that locality is indeed an ingredient in our common-sense concept of causation, I want to make use of a type of methodological consideration used by Woodward, even though it will lead me to a position that disagrees with his. Woodward, perhaps slightly tongue-in-cheek, writes the following:

Broadly speaking, philosophical theories of causation fall into two main camps [...]. Some theories, call them ‘impractical,’ provide accounts that fail to make it understandable how knowledge of causal relationships has any practical utility at all. Other theories, call them ‘practical,’ do provide (or are naturally interpretable as providing) accounts that connect causal knowledge with some goal that has practical utility [...]. My suggestion is that [...] we have strong prima-facie reasons to prefer practical to impractical theories. (Woodward, 2003, p. 30)

He then goes on to argue that causal process theories should be counted as impractical theories. The example he uses in this context is the conserved quantity theory of Salmon (1984) and Dowe (2000). He writes:

¹⁰ Indeed, Woodward explicitly takes it to be a virtue of his theory that it can allow for ‘action-at-a-distance’. See, for instance, Woodward, 2003, p. 148 and p. 353.

[T]he conserved quantity theory, as it stands, fails to explain the practical point of (the benefit that would follow from) the ability to distinguish between interactions that involve the exchange of energy and momentum and those that do not. Needless to say, animals, small children, and many adults do not possess the concepts of energy, momentum, and conservation laws and are not motivated by a disinterested intellectual concern to classify interactions into those that do and those that do not involve the transfer of energy and momentum. (Woodward, 2003, p. 30)

I take it that Woodward's remarks here are meant to extend also to other theories that imply locality, as opposed to theories such as his own version of interventionism, which he thinks provides a straightforward and plausible answer to the question as to what the point of our having the notion of causation is. Put crudely, the thought is that the distinction between causal and noncausal relationships matters to us because it is only causal relationships that we can exploit practically, in the sense of, say, bringing about the effect by bringing about the cause.

I want to take from Woodward the idea that an account of the concept of causation has to earn its keep, at least in part, by making plausible why we have that concept in the first place. I want to argue, though, that there is something potentially quite misleading about the way Woodward sets up the alternatives here, in particular if the passages above are read as a general critique of causal process accounts of causation.

First of all, it is not obvious that a defender of the type of causal process view I have outlined in the previous section cannot tell a story as to how distinguishing between causal and noncausal relationships might be of practical relevance for us. For instance, if causal relationships are subject to spatial enabling conditions, grasp of such enabling conditions might, on occasion, allow us to prevent A from causing B, or to enable it to do so, by changing some spatial parameters of the situation. This is an issue I will return to.

However, I think there is also a deeper methodological problem with the way in which Woodward appeals to a dichotomy between disinterested intellectual concern and practical interests. Woodward describes himself as being motivated by questions such as the following: 'What is the point of our having a notion of causation (as opposed to, say, a notion of correlation) at all? What role or function does this concept play in our lives?' (Woodward, 2003, p. 28). I want to suggest that these questions in fact call for an answer that comes in two parts, only one of which Woodward seems to focus on, to the exclusion of the other. The point that Woodward makes, which I am happy to agree on, is that any plausible account of the concept of causation needs to show what sort of practical payoff there might be for us in carving up the world into causal vs noncausal relations. Yet, an issue Woodward seems to ignore is that being able to carve up the world into causal vs noncausal relations may not always require the use of a *concept* of causation. Thus, when we ask about the role or function that concept has in our lives, part of what needs to be explained is also what benefit there is for us in drawing the contrast between causal and noncausal relations in conceptual terms. In other words, what I am urging is that we need to draw a further distinction between what might be called a '*purely practical*' and a '*theoretical*' grasp of causal relationships,

where this distinction turns not on the idea that the latter is impractical, but on the idea that, in contrast to the former, it involves an exercise of the *concept* of causation.¹¹

If what I have just said is along the right lines, it suggests a certain strategy in tackling the question as to whether locality is indeed part of our common-sense concept of causation. In deciding this question, we need to look at contexts in which it is plausible to think that the subject does not just have to be sensitive to certain *de facto* causal relations, but in which the subject's employing a *concept* of causation seems to play a crucial cognitive role. The question we should then examine is whether, in these contexts, the role that the concept of causation plays is tied to a commitment to the idea that the causal relations at issue are subject to spatial constraints.

5 Causal understanding and perceptual judgement

I think it is at least arguable that one context in which it is crucial for us not just to be sensitive to certain causal relations, but also to be able to think of them in causal terms is that of perceptual judgement. I should make clear that what I will say on this issue does not aspire to any originality, and is heavily based on the work of others.¹² My aim is not to provide an argument of my own for the claim that perceptual judgement recruits causal reasoning abilities, but to find out what the motivation might be for making such a claim, and then see whether this motivation might give us a lead on the question as to what role causal concepts, more generally, play in our cognitive lives.

The precise way in which one might think of perception itself as involving a causal relation is of course the subject of considerable debate in philosophy, as some of the other chapters in this volume attest. But not all aspects of this debate are relevant for our present concerns. I will focus primarily on two types of considerations, both of which can be extracted from the following quotation from Peter Strawson:

[W]e think of perception as a way, indeed the basic way, of informing ourselves about the world of independently existing things: we assume, that is to say, the general reliability of our perceptual experiences; and that assumption is the same as the assumption of a general causal dependence of our perceptual experiences on the independently existing things we take them to be of. . . . It really should be obvious that with the distinction between independently existing objects and perceptual awareness of objects we already have the general notion of causal dependence of the latter on the former, even if this is not a matter to which we give much reflective attention. (Strawson, 1979, p. 54)

One question that Strawson is concerned with here is the question as to what it is for perception to *be* a reliable way of informing ourselves about the world around us, and there is clearly a suggestion in play that, in order to answer that question, we need to

¹¹ On related issues, see Hoerl (2011) and Peacocke (forthcoming).

¹² Some key influences here are: Campbell, 1995; Eilan, 1998; Roessler, 1999. See also Johannes Roessler's Introduction to this volume, and the chapters by Naomi Eilan and William Child.

appeal to something like the idea that perceptual experience causally tracks the state of the world. The first thing to note about this suggestion is that it ties in quite naturally with an interventionist account of causation such as the one offered by Woodward. Arguably, what matters for reliability is the obtaining of what Woodward calls invariant relationships, i.e. a relationship between experience and the world such that differences in the way the world is make a systematic difference to what we experience. However, the question that concerns us at the moment is not just whether a certain *de facto* feature of perceptual experience, such as its reliability, calls for an explanation in causal terms (perhaps along interventionist lines), but whether there are any reasons for thinking that the subject herself must be able to think of certain relations in causal terms in order to, say, arrive at perceptual judgements. And here it does not look likely that considerations about the reliability of perception can supply us with such a reason. Specifically, it cannot be that the subject must first establish, say, the obtaining of certain invariant relationships between her experiences and the world before she can rely on her perceptual experiences in informing herself about the world around her. Otherwise, it could not be true that perception is the most basic way we have of informing ourselves of the world, and on this Strawson is surely right.

I think a more promising context in which to look for a reason for thinking that the ability to make perceptual judgements involves a grasp of causal concepts is provided by Strawson's remark that 'with the distinction between independently existing objects and perceptual awareness of objects we already have the general notion of causal dependence'. In other words, one key motivation behind the idea that perceptual judgement involves a grasp of causal concepts is that we need to use such concepts to make intelligible to ourselves, in making perceptual judgements, how perception can present us with objects whose existence is independent of being perceived. An idea along those lines stands behind Gareth Evans' claim that perceptual judgement recruits a *simple theory of perception*. Evans puts the claim as follows:

Any thinker who has an idea [...] of a world of objects and phenomena which can be perceived but which are not dependent on being perceived for their existence [...] must be able to think of his perception of the world as being simultaneously due to his position in the world, and to the condition of the world at that position. The very idea of a perceivable, objective, spatial world brings with it the idea of the subject as being in the world, with the course of his perceptions due to his changing position in the world and the more or less stable way the world is. (Evans, 1982, p. 222)

There is considerable scope for debate as to what exactly we should take to be the upshot of the line of thought Evans is describing here. What seems clear is that the kind of reasoning described by Evans, in terms of which the subject makes intelligible to himself how it can be that an object he currently perceives does not depend for its existence on being perceived, employs some basic causal concepts and involves a capacity to give causal explanations of a certain kind. Amongst philosophers who invoke the notion of a simple theory of perception, however, we can find a range of

different answers as to the exact role that such causal reasoning abilities are meant to play in the context of perception. At one end of the spectrum is the idea that there is a sense in which the very ability to perceive mind-independent objects as such requires possession of (at least a rudimentary) simple theory of perception. Or the thought is that making perceptual judgements in which concepts of particular mind-independent objects figure requires a grasp of such a theory. In its weakest form, at the other end of the spectrum, the claim is that possession of the general concept of an objective reality, at any rate, requires a simple theory of perception.

My own interest, for present purposes, is not in where along that continuum the right answer lies (assuming it does lie on that continuum).¹³ That question is one about the particular role that accounts invoking the idea of a simple theory of perception assign to the capacity to engage in causal reasoning in the context of perceptual judgement. My point is rather that, irrespective of what the answer to that question is, these accounts can also be seen to provide us with a model of what the capacity to engage in such causal reasoning consists in. That is to say, they give us one way of spelling out what is involved in going beyond merely relying on the *de facto* obtaining of certain causal relations, and instead engaging in reasoning about such relations.

What kind of theory is the simple theory of perception? Note that, as described by Evans, the theory is not so much concerned with the fact that, generally, our perceptual experiences causally track the way the world is, and are, in this sense, reliable. In a way, that sort of causal dependence is already taken for granted. Instead, what the simple theory of perception is concerned with is why some particular instances of that general causal dependence obtain and others not, i.e. why I can perceive certain items but not others.¹⁴ And here it seems that the idea that causation, at least in the case of perception, is subject to spatial enabling conditions plays a crucial role: On each given occasion, what I can perceive is determined, in part, by my own position in space, my orientation, and such things as the absence of barriers between me and the thing perceived. In other words, the way we give substance to the general idea that there are further enabling conditions to be fulfilled, apart from the actual existence of an object, before we can perceive that object, is in spatial terms.

We can perhaps clarify further the nature of the reasoning capacity at issue here, and the role that a grasp of spatial enabling conditions plays in it, by considering a passage from William Child's contribution to this volume. In the context of defending the view that our ordinary thinking about seeing is a form of causal thinking, Child discusses the idea that

what makes a concept a causal concept is just that it has to do with explaining why something happened; why an event or state of affairs occurred, or came about, or persisted; what produced some event or state of affairs; why a particular thing behaved as it did, or why that kind of thing

¹³ See Naomi Eilan's contribution to this volume for a discussion of these issues.

¹⁴ Compare Campbell, 1995, p. 208: 'The actual explanations [provided by the simple theory of perception] are all at the level of the concrete particular.'

generally behaves as it does; and so on. [...] And by that standard, what makes our ordinary thinking about vision a form of causal thinking is that the ‘because’ in our reasoning about seeing (‘She couldn’t see it because it was too far away’, and so on), has to do with the explanation of why something happened (or did not happen). In our ordinary thought about vision, we are concerned with the occurrence or non-occurrence of natural phenomena: someone’s seeing this, or failing to see that. In the same way, when we reason about the enabling and defeating conditions of vision, we are reasoning about why something happened or persisted, or why something of a certain sort failed to happen. That is enough for this reasoning to be a form of causal reasoning. (Child, this volume, p. 174)

The issue Child is concerned with in this passage is somewhat different from the one at issue in the present section. However, I believe the connection he draws between causal reasoning and the explanation of particular occurrences is also relevant to our current concerns. Indeed, the difference might be put as a mere difference in emphasis. What Child, in effect, is suggesting is that a subject is engaged in *causal* reasoning to the extent that she is concerned with the explanation of *particular occurrences*. An idea that is implicit in the line of thought that I have sketched above might similarly be put by saying that a subject is engaged in causal *reasoning* to the extent that she is concerned with the *explanation* of particular occurrences. Thus, the strength of the thought that perceptual judgement, for instance, involves causal reasoning abilities turns on the idea that it requires that the subject be able to grasp why the same object can sometimes be seen and sometimes not.

The idea behind the claim that perceptual judgement recruits a simple theory of perception, then, is that a simple theory of perception can furnish the subject with the required kind of explanation. However, it seems to be able to do so only because within the simple theory of perception, perception is conceived of as involving causal relations that are subject to spatial constraints. What does explain why I can sometimes see an object and sometimes not, according to the simple theory of perception, is that, apart from the object being present, I must also stand in the right sort of spatial relation to it. Thus, in as far as accounts of perceptual judgement that invoke the notion of a simple theory of perception make a plausible case for the idea that (an element of) perceptual judgement recruits causal reasoning abilities, rather than just a sensitivity to the *de facto* obtaining of causal relations, those reasoning abilities turn on a commitment to a version of locality, at least in connection with the causal relations involved in perception.

6 Causal understanding and agency

It is easy to imagine someone agreeing, at least in broad outline, with the line of thought just sketched, but arguing that it tells us more about the concept of perception than about that of causation. In other words, it may well be that, in the context of talking of the way in which we think of perception as a causal phenomenon, we also need to bring in the idea that it is subject to spatial enabling conditions. However, it

clearly does not follow from this that this idea always has to be brought in when we think of a phenomenon in causal terms.

Furthermore, it might be held that a more natural (or at least equally natural) starting point in considering questions about subjects' grasp of causal concepts is by looking at their ability to recognise their own actions as being effective. For instance, Menzies and Price (1993, p. 187) claim that 'the ordinary notions of cause and effect have a direct and essential connection with our ability to intervene in the world as agents'. More specifically, on their view, we can explain how subjects come to have causal concepts by pointing to features of their own experiences as agents. They explain:

[F]rom an early age, we all have direct experience of acting as agents. That is, [...] we all have direct personal experience of doing one thing and thence achieving another. We might say that the notion of causation thus arises not, as Hume has it, from our experience of mere *succession*; but rather from our experience of *success*: success in the ordinary business of achieving our ends by acting in one way rather than another. It is this common and commonplace experience that licenses what amounts to an ostensive definition of the notion of 'bringing about'. In other words, these cases provide direct non-linguistic acquaintance with the concept of bringing about an event. (ibid. p. 194f.)

At least part of the proposal Menzies and Price put forward here fits in quite well with an interventionist approach to causation. We can make sense of the idea of a distinctive experience of success, as they call it, if we think of that experience as reflecting the detection of patterns of invariant relationships of the type that interventionist accounts of causation focus on.¹⁵ Arguably, some such ability to detect patterns of invariant relationships is key to the very possibility of agency. Using the technical notion of 'wiggling' (Lewis, 2004, p. 91), Price elsewhere explains that '[t]he task of an agent is to find out what else wiggles, in response to her possible present actions' (Price, 2007, p. 13). Without being able to register, at least on occasion, whether or not her actions make a systematic difference to aspects of her environment, a subject would not have any grip at all on how to bring about events.

In what follows, I thus want to grant the idea that agency involves something like direct acquaintance with the phenomenon of bringing something about, understood along interventionist lines. However, I want to argue that in the context of agency, too, grasp of causal *concepts* involves theoretical commitments that go beyond those recognised by interventionism. More specifically, my claim will be that an agent's ability to use causal concepts in thinking about her own actions turns on her possessing a simple theory of action akin to the simple theory of perception discussed in the previous section. Furthermore, as with the simple theory of perception, the specific role that the simple theory of action plays in an agent's reasoning seems tied to a commitment to a version of locality, i.e. the thought that the causal relations exploited in one's actions are subject to spatial constraints.

¹⁵ See also Woodward, 2007, p. 29.

Perhaps the most vivid way of bringing out the key issue that is at stake here is by considering the possibility of what Woodward (2007) calls *purely egocentric* instrumental behaviour and learning. In such purely egocentric instrumental behaviour and learning the agent grasps (or behaves as if it grasps) that there are regular, stable relationships between its manipulations and various downstream effects but stops at this point, not recognizing (or behaving as though it recognizes) that the same relationship can be present even when it does not act, but other agents act similarly or when a similar relationship occurs in nature without the involvement of any agents at all. (Woodward, 2007, p. 32)

As Woodward points out, it has been suggested that primates, for instance, do not go beyond this basic level of learning. Consider the following passage from Tomasello and Call (1997, p. 389):

[S]uppose that an individual ape, who has never before observed such an event, for the first time observes the wind blowing a tree such that the fruit falls to the ground. If it understands the causal relations involved, that the movement of the limb is what caused the fruit to fall, it should be able to devise other ways to make the limb move and so make the fruit fall. [...] We believe that most primatologists would be astounded to see the ape, *just on the basis of having observed the wind make fruit fall*, proceed to shake a limb, or pull an attached vine, to create the same movement of the limb.

The type of limitation Tomasello and Call talk about here is meant to stand in stark contrast with primates' other learning abilities, e.g. when it comes to trial and error learning. To use Price's phrase, primates are in fact very capable of finding out what else wiggles, in response to their own actions.

Woodward contrasts purely egocentric instrumental behaviour and learning with what he calls a fully causal viewpoint, which involves, for instance, the ability to move between classical conditioning and instrumental conditioning, or, in other words, 'the ability to use information learned about causal relationships through passive observation to guide interventions and vice versa' (Woodward, 2007, p. 35). This is one way of characterising what the limitation here comes to, which turns on the question as to which sources of information a subject can exploit. However, if Tomasello and Call's description is correct,¹⁶ there is clearly also a sense in which we are dealing with a limitation on the conceptual level. That is to say, the ape described by Tomasello and Call lacks the conceptual resources to give substance to the sameness between the causal relations that are involved in its own interactions with the world and those involved in natural occurrences in which it is not itself the agent.

This poses a difficulty for the idea that the general concept of bringing about an event, or similar causal concepts, could simply be ostensively defined with reference to a characteristic experience of success in the way suggested by Menzies and Price. Arguably, a creature in possession of a such a concept should have some way of giving

¹⁶ See Povinelli, 2000, p. 314, for some scepticism as to whether it is in fact empirically true.

substance to the idea that the same causal relationship that governs its own actions can also be present in events in which it is not itself involved. But this is precisely what the ape described by Tomasello and Call seems to lack. And the trouble is that it is hard to see how having a characteristic experience of success could play any role in overcoming that lack. If there is such a thing as an experience of success, as envisaged by Menzies and Price, there seems to be no reason to deny that the ape has such experiences, and that they play a key role in the learning abilities it does possess. Yet, they play a role only in enabling it to learn from its own actions.

I think it is worth noting, at this point, a structural similarity between the issue that has come to the fore in the preceding paragraphs, and the issue that prompted the idea, in the last section, that perceptual judgement recruits a simple theory of perception. There, the question was how a subject, in making perceptual judgements, can make sense of the idea that the very objects that she can perceive are not dependent for their existence on being perceived. The issue that has emerged from the above discussion of Menzies and Price's proposal, in effect, is what it takes for a subject to make sense of the idea that the very kinds of causal relations that she can exploit in acting on her environment can obtain without her acting.

To put the matter slightly differently: The point Evans and others have been making in the case of perceptual judgement is that we need a substantive account of what it is in virtue of which the subject can appreciate that the same object can exist perceived or unperceived. In a similar way, I am now arguing that, in the case of causal judgement, we need a substantive account of what it is in virtue of which the subject can appreciate that the very same causal relationships governing her own actions can also be present in natural occurrences that do not involve agents. Only then can we say that she is not merely sensitive, in some way, to the *de facto* obtaining of certain causal relations, but also that she is capable of thinking about them in terms of causal concepts.

If the line of thought discussed in the previous section is at least roughly along the right lines, a natural suggestion, at this point, is this: it is a grasp of a simple theory of action, akin to the type of simple theory of perception appealed to in the accounts of perceptual judgement I have discussed, that is required to give substance to the idea that the same kinds of casual relations that are exploited in one's own acting can also govern events in which no agent is involved. Again, the idea would be that such a simple theory of action furnishes the subject with an explanation of particular occurrences (and non-occurrences), reflecting the fact that, e.g. one's own actions have certain outcomes only under certain conditions. It is this idea that then also allows for the thought that the same type of outcome could depend in the same way on causes that are not actions, if its occurrence is subject to the same conditions. In other words, what makes available to the subject the thought that the same type of event can be brought about in the same way by her own actions as well as by naturally occurring events is the idea that the same explanation can be given for the occurrence or non-occurrence of this type of event in each case.

Just as in the case of the simple theory of perception, though, it seems that the idea that causal relations are subject to spatial enabling conditions has a crucial role to play in how a simple theory of action can provide for the type of explanation called for here. Consider again the example of Tomasello and Call's ape. What seems to be missing in this example is any kind of indication that the ape grasps *why* the wind, in the particular situation it witnesses, has the power to make the fruit fall. This is why it cannot connect the familiar fact of the wind blowing with the falling of the apple in a way that might also make it realise implications for its own actions. At its most basic, though, such a transfer between observation and action would seem to require a grasp of the fact that the wind can only make the fruit fall because it is acting on the limb. Thus, the reasoning required for the ape to give substance to the idea that the same causal relations at work in the events it witnesses could also be exploited in its own actions turns crucially on the idea that those causal relations are subject to certain spatial enabling conditions.

If this line of thought is along the right lines, it suggests that, just as in the case of perception, giving causal reasoning capacities and a grasp of causal concepts a role to play in agency is tied to the idea that the agents in question can engage in a form of reasoning that trades on an assumption of locality. My argument has been that we require a substantive account of how a subject with causal concepts makes intelligible to herself that the same such concepts can be applied to its own interactions with the world as well as to events in which it is not an agent. What I have suggested, in short, is that we can see how the subject can do so if we credit it with the idea that the same spatial constraints govern the occurrence of particular types of outcome, no matter whether they are being brought about by the subject herself or in some other way. Thus, in the context of agency too, it seems, making plausible the idea that a subject engages in genuine causal reasoning, rather than merely being sensitive in some more primitive way to the *de facto* obtaining of certain causal relations, seems to be bound up with the idea that the reasoning in question involves a commitment to locality.

7 Causation, space, and practical interests

In section 4, above, I outlined a strategy for resolving the question as to whether a commitment to locality should be seen as being part of our common-sense concept of causation. The strategy I recommended was to look at types of reasoning that are arguably central to having causal concepts in the first place, and to see to what extent, if any, these types of reasoning turn on an assumption of locality. I subsequently discussed two reasoning contexts that might, with some justification, be held to be fundamental to our grasp of causal concepts. I suggested that making perceptual judgements and insightful causal learning both mobilise a simple theory (of perception and action, respectively), in which the idea that the causal relations at issue are subject to spatial constraints plays a crucial role. Thus, I have argued that at least two types of causal

understanding that have some claim of being central to our understanding of causation generally turn on the assumption of locality.

I do not mean to suggest that what I have said about the case of perception and the case of agency is already sufficient to establish that an assumption of locality is part of our common-sense concept of causation, quite generally. However, in this final section, I wish to draw out three key issues that emerge from what I have argued, which at least offer starting points for working towards a more general conclusion.

First, when discussing Woodward's distinction between what he calls 'practical' and 'impractical' theories of causation, I said that Woodward's way of setting up this distinction ignores a further distinction that we might want to draw between what I called a 'purely practical' and a 'theoretical' grasp of causal relations, where only the latter involves the use of causal concepts. I think that both the case of perception and that of action provide plausible examples of contexts in which we do indeed need to invoke a distinction of this kind.

I take it that there are creatures who do not possess what I have called a simple theory of perception, but who can nevertheless use their sense organs to find their way around the terrain, locate food, etc. Similarly, I take it that there are creatures who do not possess a simple theory of action, but can nevertheless engage in purposive behaviour (if Tomasello and Call, 1997, are right, apes are an actual example). Arguably, there is a sense in which these creatures exploit just the same causal connections in perception and action that we do in making perceptual judgements and engaging in insightful causal learning. But they do so in a way that does not need to involve the use of causal concepts at all. My argument has been that it is not perception or agency as such that demand the possession of causal concepts and the ability to engage in the kinds of causal reasoning that a simple theory of perception or action provides for, but more specific abilities, such as the ability to form judgements about a mind-independently existing world on the basis of perception, or the ability to think about one's own actions in the same terms as about happenings that do not involve agents.

Second, just how does grasp of a simple theory of perception and action go beyond a mere practical sensitivity to causal relations and instead involve the subject herself thinking of those relations in causal terms? To repeat, grasping a simple theory of perception is a matter of grasping that what we can perceive is the joint upshot of what there is in the world and our own position in the world, thus making intelligible how the very objects we can perceive are, at the same time, not dependent for their existence on being perceived. Similarly, grasp of a simple theory of action is, in part, a matter of grasping that what particular outcomes we can bring about is a matter of what kind of thing can cause what kind of other thing and where we are located, thus making possible the idea that the very same causal relationships we can exploit in our actions can obtain independently of us acting. I think one way of describing the general kind of causal reasoning at issue here is by saying that it is concerned with the relationship between general and singular causation in the case of perception and action. It involves, for instance, both the idea of a general causal dependence of

perceptual experience on the world, and the idea of particular instances of a feature of the world causing a particular perceptual experience, and provides an explanation of what else it takes, apart from the former, for the latter to occur.

I think the general moral in the offing here is that causal concepts generally only come into play once a subject faces an explanatory task of this particular type. That is to say, there may be various ways in which a subject is able to exploit the fact that As cause Bs without the subject herself using causal concepts. But such concepts do seem to be required if the subject has to grasp what else has to be the case, given the fact that As generally cause Bs, for a particular A to cause a B. Yet, note that this explanatory task already presupposes that there must, in each case in which an A causes a B, be an answer to the question how it does so that also makes intelligible how, sometimes, an A can occur but without causing a B. Thus, the subject's use of causal concepts, in this context, is tied to the idea that the question of how A causes B always has an answer—and, as I have suggested in section 3, above, this is just one way of understanding the idea that causation involves causal process.

Finally, once we put the point of having a concept of causation in the way just sketched, we can perhaps also see how what I have called a theoretical grasp of causal relations need not just be a matter of mere disinterested intellectual concerns, as Woodward calls it, but can be of practical relevance. Understanding why, even though As generally cause Bs, they can only do so if certain conditions are met, also involves understanding what can prevent A causing B on a given occasion, or how a change in conditions might enable it to do so, and the notions of preventing and enabling, in this sense, are clearly of practical relevance.¹⁷ Note also, though, that the notions of preventing and enabling, understood in this sense, do not seem to apply to putative cases of action at a distance, such as the workings of the 'powder of sympathy' described at the beginning of this paper. Or, at any rate, we have no idea as to how we could apply them. That is to say, there does not seem to be any intelligible account we could give of why such actions might work on a given occasion but not another. And this might indicate a sense in which we do not just think of the possibility of action-at-a-distance as far-fetched, but in which the idea of action-at-a-distance violates a basic feature of our common-sense concept of causation.

References

- Anscombe, G. E. M. (1971). Causality and determination. Reprinted in: P. Van Inwagen and D. W. Zimmermann (eds.) *Metaphysics: The Big Questions* (pp. 244–258). Oxford: Blackwell 1998.
- Campbell, J. (1995). *Past, Space and Self*. Cambridge, MA: MIT Press.

¹⁷ I believe that, understood in the way indicated here, concepts such as 'prevent' and 'enable' form an important, and generally neglected, middle ground between what Anscombe (1971) calls special causal concepts, such as 'scrape, push, wet, carry, eat, burn, knock over', and the general concept 'cause'. This issue is discussed in further detail in Hoerl (2011).

- (2007). An interventionist approach to causation in psychology. In: A. Gopnik and L. E. Schulz (eds.), *Causal Learning: Psychology, Philosophy, and Computation* (pp. 258–266). Oxford: Oxford University Press.
- Collins, J., Hall, N., and Paul, L. A. (eds.) (2004). *Causation and Counterfactuals*. Cambridge, MA: MIT Press.
- Digby, K. (1658). *A late discourse made in solemn assembly of nobles and learned men at Montpellier in France, touching the cure of wounds by the powder of sympathy*, 2nd edn, trans. R. White. London: R. Lowdes.
- Dowe, P. (2000). *Physical Causation*. New York: Cambridge University Press.
- Dummett, M. (2005). The justificationist's response to a realist. *Mind* 114, pp. 671–688.
- Eilan, N. (1998). Perceptual intentionality, attention and consciousness. In: A. O'Hear (ed.) *Contemporary Issues in the Philosophy of Mind*. (pp. 181–202) Cambridge: Cambridge University Press.
- Evans, G. (1982). *The Varieties of Reference*. Oxford: Clarendon Press.
- Frisch, M. (2010). Causes, counterfactuals and non-locality. *Australasian Journal of Philosophy* 88, pp. 655–672.
- Hall, N. (2002). Non-locality on the cheap? A new problem for counterfactual analyses of causation. *Nous* 36, pp. 276–294.
- Hill, A. B. (1965). The environment and disease: association or causation? *Proceedings of the Royal Society of Medicine* 58, pp. 295–300.
- Hoerl, C. (2011). Causal reasoning. *Philosophical Studies* 152, pp. 167–179.
- Lewis, D. (2004). Causation as influence. In: J. Collins, N. Hall, and L. A. Paul (eds.) *Causation and Counterfactuals*. Cambridge, MA: MIT Press.
- Menzies, P. and Price, H. (1993). Causation as a secondary quality. *British Journal for the Philosophy of Science* 44, pp. 187–203.
- Mercer, C. (1998). Digby, Kenelm. In E. Craig (ed.), *Routledge Encyclopedia of Philosophy*. London: Routledge.
- Peacocke, C. (2005). Justification, realism and the past. *Mind* 114, pp. 639–70.
- (forthcoming). Representing causality. In: T. McCormack, C. Hoerl, and S. Butterfill (eds.), *Tool Use and Causal Cognition: Issues in Philosophy and Psychology*. Oxford: Oxford University Press.
- Pearl, J. (2000). *Causation*. Cambridge: Cambridge University Press.
- Povinelli, D. J. (2000). *Folk Physics for Apes: The Chimpanzee's Theory of How the World Works*. Oxford: Oxford University Press.
- Price, H. (2007). The effective indexical [Draft at November 23, 2007]. Retrieved from <http://philsci-archive.pitt.edu/archive/00004487/01/EffectiveIndexical.pdf>
- Roessler, J. (1999). Perception, introspection and attention. *European Journal of Philosophy* 7, pp. 47–64.
- Salmon, W. (1984). *Scientific Explanation and the Causal Structure of the World*. Princeton: Princeton University Press.
- Schaffer, J. (2004). Causes need not be physically connected to their effects. In: C. Hitchcock (ed.), *Contemporary Debates in Philosophy of Science* (pp. 197–216). Oxford: Blackwell.
- Strawson, P. F. (1959). *Individuals: An Essay in Descriptive Metaphysics*. London: Routledge.
- (1979). Perception and its objects. In: G. McDonald (ed.), *Perception and Identity: Essays presented to A. J. Ayer* (pp. 41–60). London: Macmillan & Co.

Tomasello, M. and Call, J. (1997). *Primate Cognition*. Oxford: Oxford University Press.

Woodward, J. (2003). *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press.

—— (2007). Interventionist theories of causation in psychological perspective. In: A. Gopnik and L. Schulz (eds.). *Causal Learning: Psychology, Philosophy and Computation* (pp. 19–36). New York: Oxford University Press.