Causal Reasoning

by Christoph Hoerl

This is an electronic version of an article forthcoming in *Philosophical Studies*. The definitive version will be available at www.springerlink.com. DOI: 10.1007/s11098-009-9474-7

Abstract:

The main focus of this paper is the question as to what it is for an individual to think of her environment in terms of a concept of causation, or causal concepts, in contrast to some more primitive ways in which an individual might pick out or register what are in fact causal phenomena. I show how versions of this question arise in the context of two strands of work on causation, represented by Elizabeth Anscombe and Christopher Hitchcock, respectively. I then describe a central type of reasoning that, I suggest, a subject has to be able to engage in, if we are to credit her with causal concepts. I also point out that this type of reasoning turns on the idea of a physical connection between cause and effect, as articulated in recent singularist approaches of causation.

Keywords: Causal reasoning, singular causation, concepts

Causal Reasoning

What is it for an individual to think of her environment in terms of a concept of causation (or, more generally, in terms of causal concepts[1])? My aim in this paper is to draw attention to this question, which is not often discussed in contemporary debates about causation, and show how it might in fact be of some relevance to such debates.

In the first two sections, my main focus will be on two, rather different, strands of work on causation, represented by Elizabeth Anscombe and Christopher Hitchcock, respectively. I will show how each, in its own way, raises the need to distinguish between the ability to use causal concepts and other, more primitive ways in which a subject might be able to pick out or register what is in fact a type of causal phenomenon without actually using causal concepts. I will also draw on some empirical work in psychology, especially developmental psychology, to illustrate the distinctions at issue.

In the final section, I will sketch one way of connecting the issue as to what it is for an individual to have causal concepts with current debates about causation. I will suggest that a central type of reasoning that a subject has to be able to engage in, if we are to credit her with causal concepts, turns on the idea of a physical connection between cause and effect as articulated in recent singularist approaches to causation.

1. Thick Causal Concepts

Elizabeth Anscombe's inaugural lecture 'Causality and Determination' is often seen as a paradigmatic statement of what has become known as singularism about

[1] See the next section for the relevance of this distinction.

causation, i.e. of the view that the truth of singular causal claims (such as 'David's smoking caused his heart disease') does not depend on the existence of universal regularities or laws, and that general causal claims (such as 'Smoking causes heart disease') are generalizations over singular ones.[2] In what follows, I will be interested primarily in two claims, or stages in Anscombe's argument, and their relationship to one another. Consider the following, often-cited passage:

> The word 'cause' itself is highly general. How does someone show that he has the concept *cause*? We may wish to say: only by having such a word in his vocabulary. If so, then the manifest possession of the concept presupposes the mastery of much else in language. I mean: the word 'cause' can be *added* to a language in which are already represented many causal concepts. A small selection: *scrape, push, wet, carry, eat, burn, knock over, keep off, squash, make* (e.g. noises, paper boats), *hurt*. But if we care to imagine languages in which no special causal concepts are represented, then no description of the use of a word in such languages will be able to present it as meaning *cause* (Anscombe, 1971, p. 248).

The issue at the forefront in this passage is not singularism as such. Anscombe's main interest here is in a distinction that actually cuts across the distinction between singular and general causal claims. Following Cartwright (2004), I will characterise the distinction Anscombe has in mind as that between the *thin* concept 'cause', on the

[2] There are, in fact, several quite different strands to Anscombe's paper, of which this is only one. One of the others, which I will set aside, is a denial of determinism, as it is not clear to me that the falsity of determinism would need to have the singularist consequences Anscombe thinks it does (see also Hitchcock 1995).

one hand, and *thick* concepts such as 'scrape', 'push', 'wet', etc., on the other.[3] That this distinction does indeed cut across the one between singular and general causal claims can be seen from the fact that the thin concept 'cause', despite being 'highly general' in Anscombe's sense, can still figure in singular causal statements (e.g. 'The explosion of the space shuttle Challenger was caused by the failure of the o-rings on the right hand booster'); conversely, thick causal concepts can be used to make general causal statements (see, e.g., Cartwright's, 2004, description of the workings of a carburettor).

Nevertheless, the argument in the passage quoted above plays a key role in a larger argument for singularism. For Anscombe thinks that thick concepts such as 'scrape', 'push', 'wet', etc. can be applied on the basis of observing particular instances of causal interactions without presupposing a relevant general causal claim. Thus, we seem to have here the elements required for fleshing out a version of the view that singular causal claims have a priority over general ones. Basically, the idea would be that thick concepts, applied in the context of observing particular instances of causation, are basic in our understanding of causal relationships in general, and that general causal claims should be seen as generalizations from such particular instances.

Yet, it is possible to agree with the idea of a priority of thick concepts in our understanding of causal relations, and also with the idea that observation of circumstances in which such concepts apply plays a crucial role in such understanding, whilst nevertheless thinking that there is an important issue left unaddressed by Anscombe. Loosely speaking, the worry is that Anscombe has failed to show us what it is about a subject's ability to recognize instances of scraping,

[3] Cartwright, in turn, adopts the "thick/thin" distinction from Williams (1985, p. 129), who introduces it in a discussion of ethical concepts, distinguishing "'thicker' or more specific ethical notions . . . such as *treachery* and *promise* and *brutality* and *courage*" from thin ethical notions such as *good* or *ought*.

pushing, wetting, etc., in virtue of which the subject can count as thinking of these instances in genuinely *causal* terms. For, intuitively, what her story leaves out is precisely what it is in virtue of which the subject understands that these are all forms of *causing* something.[4]

To make the issue I have in mind more concrete, it might help to consider two studies on children's understanding of causation carried out by Gelman, Bullock, and Meck (1980), and by Das Gupta and Bryant (1989), respectively. In Gelman et al.'s study, children were shown two pictures of an object, depicting the object before and after it had undergone a causal transformation (e.g. a picture of an apple and a picture of a wet apple), and they had to choose from an array of pictures of objects (e.g. a towel, a knife, a water jug) the right one to bring about the relevant transformation. Gelman et al. found that 3-year-olds were able to match objects to outcomes, and, conversely, were also able to match outcomes to objects.

Das Gupta and Bryant's (1989) study seems, on the face of it, only slightly more complex, but achieved rather different results. In this study, children again had to choose from an array of pictures of objects the right one to complete a sequence of three pictures, but this time the first picture in the sequence showed an object that had already undergone a causal transformation. Thus, the first picture might, for instance, be of a cut apple, and the last of a wet and cut apple. Das Gupta and Bryant tested both 3-year-olds and 4-year-olds and found that only the latter reliably succeeded in the task, whereas the former failed.

---

[4] Godfrey-Smith (forthcoming) raises a related worry in his critical discussion of what he calls 'causal minimalism'. As he explains, "causal minimalism holds that to assert that some connection is causal is not to attribute some special empirical or modal feature to it. To call a connection causal is to say that this connection is one that can also be described using some unspecified member of [a set of causal verbs representing 'special causal concepts' in Anscombe's sense]". Like Godfrey-Smith, I don't believe that minimalism is what Anscombe was actually after, so the issue I am raising is how she might manage to avoid it. Cartwright (2004) is probably the philosopher with the most explicit minimalist leanings, and I suspect she would accuse my way of framing the worry question-begging.

Now, Anscombe's claim was that there are types of causation that are observable, and that there is a class of concepts, thick concepts, which can be applied on the basis of such observation without knowledge of a relevant general causal claim. The issue that Das Gupta and Bryant's findings seem to make pressing, though, is how these two claims are related to each other. For it seems plausible to explain 3-year-olds' ability to pass Gelman et al.'s task by saying that they have seen instances of the relevant causal transaction before, and, on that basis, possess a type of knowledge that can guide selection of an object suitable for bringing about a certain result. Yet, Das Gupta and Bryant's study suggests that this type of knowledge is, as it were, quite piecemeal, and does not equip the children with the ability to engage in certain, rather basic, forms of inferential reasoning about causal relations. And, at least on the face of it, this should hold us back from ascribing to the children concepts of the relevant causal relations, because it is concepts that are required for inferential reasoning.

It might be useful to note a parallel between the issue I have just sketched and a similar issue that arises in connection with the relationship between shape perception and shape concepts. In 'Things without the mind', Gareth Evans writes:

> [I]t does not seem to be possible to regard the conception of the shape of a material thing – with all the propositions about its characteristic behaviour and interaction with other bodies which that implies – as the same as whatever shape concepts might be grounded in the colour mosaic thought to be given in immediate visual experience (Evans, 1985, p. 270).

Evans does not mean to deny, here, that shape properties can be perceived. Rather, his point is that our grasp of what it is for an object to be round, for instance, cannot be explained simply in terms of an ability to recognize perceptually instances in which that property is instantiated. Instead, such a grasp requires that one "master a set of interconnected principles which make up an elementary theory" (ibid., p. 269) into which the relevant property fits.

A somewhat similar line of thought can be found in John Campbell (1995), who contrasts two different types of representation of shape, which differ according to the types of operations that can be performed on them.

> At the level of perception we have representations which are involved in imagistic reasoning: the rotations and zooms with which Shepard and others have made us familiar. And these representations are used in guiding action. Even an animal incapable of conceptual thought can engage in perceptually guided action. At the conceptual level we have representations which are involved in deductive reasoning: reasoning involving the application of the laws of logic. […] To have a shape concept is to be able to form representations of shape which can figure in deductive reasoning (Campbell, 1995, p. 362).

Intuitively, the findings on children's developing understanding of causation mentioned above suggest that we need to make a similar kind of distinction between two types of reasoning, and correspondingly two types of representation, in the case of causal relations. Perceptually-based imagistic reasoning abilities may be sufficient

to pass Gelman et al.'s task, whereas Das Gupta and Bryant's task may require deductive reasoning involving conceptual thought about causal relations.

To make good such a suggestion, we need to say more about what might be the distinguishing features of reasoning that involves the use of causal concepts, and show how these features are present in Das Gupta and Bryant's task. I will return to this issue in the final section of this paper. Before doing so, however, I want to turn to quite a different type of philosophical work on causation and show that it actually raises the same type of issue that I have just discussed in connection with Anscombe's account, though in a slightly different guise.

2. Causal Models

In several recent papers, Christopher Hitchcock has extolled the virtues of causal models as tools for representing causal relationships (see, e.g., Hitchcock, 1995, 2001, 2007). Causal models, which were originally developed in areas such as epidemiology and artificial intelligence (Pearl, 2000; Spirtes, Glymour & Scheines, 2000), typically consist in a directed graph, identifying a number of variables that are connected by directed edges ('arrows'), and an accompanying set of structural equations, which specify the exact nature of the connection between variables. As such, they can be used to represent what is often called the *causal structure* of a given physical system by making explicit (often quite intricate) patterns of counterfactual dependence holding between different variables in the system.

One important reason why Hitchcock is interested in causal models is that he thinks that, depending on context, common-sense judgements about causation typically focus only on one or another subset of the counterfactual relationships that hold in a given causal scenario, and using causal models can thus help us

disambiguate such judgements. There has also been increasing interest, however, in the connected idea that a basic form of mental representation governing causal learning embodies a causal model. In other words, the thought is that we can explain certain features of causal learning by postulating cognitive structures that represent causal structure in the manner of a causal model.[5]

For the purposes of this paper, this latter idea is of particular interest because, as a device for representing causal relations, causal models might be seen to possess one feature that Anscombe's account seemed to have difficulties providing for. Note that there appears to be a straightforward, and general, answer to the question as to why we should say that the relationships represented by causal models are represented as causal relationships. To quote Hitchcock: "The equations in the model represent *causal* structure […] because the equations provide information about the effects of hypothetical interventions, and about counterfactuals" (Hitchcock, 2007, p. 509, emphasis in the original). Obviously, there is considerable room for debate here as to what exactly this claim presupposes (and what it leaves open) about the relationship between causation, interventions and counterfactuals. For the sake of the argument, though, I will assume that it can be made good.

Rather, I will focus on another feature of causal models that Hitchcock (2007, p. 510) himself notes: Just like Anscombe's thick causal concepts, they can be used to capture both singular and general causal relationships. In this respect, causal models bear some resemblance to another sort of representation, namely maps.[6] Perhaps we typically think of maps as representations of particular places, because we take as our paradigm geographical maps, which do represent particular places, but there are also

[5] Cf., e.g., Glymour (2001), Gopnik et al. (2004).
[6] Perhaps unsurprisingly, some theorists therefore also talk of 'causal maps' (see, e.g., Gopnik et al., 2004). Another type of representation which shares this feature are scripts as described originally by Shank & Abelson (1972), and as explored in detail in developmental psychology by Nelson (1996); see also Hoerl & McCormack (forthcoming).

maps that represent general types of spatial arrangement. For instance, a budget hotel chain may have a map of the standard room layout they use, and there are anatomical atlases that show such things as maps of the human circulatory system. In other words, as a type of representation, maps are intrinsically neither particular nor general; they only become so if put to specific uses by us. The same, it seems, goes for causal models. I will call representations that share this feature 'generic', and I want to argue that being generic, in the sense just described, is a feature of representations that a subject can possess, and that allows that subject to be sensitive to the obtaining of certain relations in the absence of her possessing a concept of the relation in question.

For an argument to that effect, applied to the case of maps, we can again look to work by John Campbell. Campbell (1994) picks up on work in comparative psychology, where animals are sometimes described as having a 'cognitive map' of a certain terrain. An animal might be said to have a cognitive map of a maze, for instance, because it is able to find the shortest way back to a food source at a particular location in the maze independently of where in the maze it is put. Campbell's question is in what sense an animal in possession of a cognitive map, thus described, can itself be said to appreciate the connectedness of space, i.e. the fact that all the places represented on the cognitive map are spatially related to each other, and whether its doing so involves explicit reasoning involving spatial concepts. Taking one particular type of account of a cognitive map (involving the technical notion of a centroid) as his example, Campbell argues that

> the animal is not using reflective thought about its targets to assign causal significance to spatial relations: it does this rather through the fact of its own

> engagement in the space. […] It is true that the vectors from the centroid to various cues do not depend on the location of the animal. But when we subtract the animal, we also subtract any physical meaning for those vectors. The vectors get their meaning only when the animal is plugged in. They have causal significance only through their relations to the animal's perception and action (Campbell, 1994, p. 32).

It seems to me that if it is legitimate to draw the type of distinction Campbell has in mind here in the case of cognitive maps, it should also be legitimate to draw a parallel distinction in the case of causal models. Roughly speaking, the idea would be that the grounds we might have for ascribing a causal model to a subject might not be sufficient for ascribing causal concepts to that subject, if the only way in which the model is put to use by the subject is in her own practical engagements with the world, without the subject herself being able to give significance to the idea that the same causal relations exploited in those practical engagements can also be present in circumstances in which there is no agent.

To make the issues at stake here more precise, it may help to note a difference between the kind of distinction I have in mind, between a purely practical grasp of causal relations and a capacity to employ causal concepts, and a perhaps similar-sounding distinction drawn by James Woodward, whose own work makes use of causal models. Woodward (2007) distinguishes between three levels of causal understanding, or 'causal viewpoints', as follows:

> (1) [A] purely *egocentic* [causal viewpoint: 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.
>
> (2) [An] *agent causal* viewpoint: the agent grasps that the very same relationship that it exploits in intervening also can be present when other agents act.
>
> (3) [A] *fully causal* viewpoint: the agent grasps that the very same relationship that he exploits in intervening also can be present both when other agents intervene and in nature even when no other agents are involved (Woodward, 2007, p. 32).

Woodward's distinctions are meant to capture one way in which a subject's causal understanding may be more or less sophisticated, depending on how closely tied up it is with the subject's own capacities to act. What matters for the distinction between an egocentric and a non-egocentric causal point of view, in Woodward's sense, is the type of information a subject can *draw on* in learning about causal relationships – whether it can only learn about them from its own interventions and their outcomes, or whether it can also use information from situations in which it did not intervene itself.

There is, however, also another way in which a subject's causal understanding may be more or less sophisticated, depending on how closely tied up it is with the subject's own capacities to act. And this is the one at issue in what I have characterized as the distinction between being able to engage in reasoning involving a use of causal concepts, on the one hand, and more primitive ways of being sensitive to

causal relations, on the other.[7] Consider Woodward's description of a subject who possesses a fully causal viewpoint, in his sense, as "grasp[ing] that the very same relationship that he exploits in intervening also can be present both when other agents intervene and in nature even when no other agents are involved". What I want to argue is that when it comes to spelling out the sense in which a subject might be said to have a grasp of the sameness of the relevant causal relationships, here, we need to draw a distinction that is analogous to the one at issue in the paragraph from Campbell cited above, where he contrasts two quite different ways in which a subject in possession of a cognitive map might be said to have a grasp of the connectedness of space. In both cases, we need to make a distinction between what could be called a purely practical grasp, and a grasp that involves the subject having a concept of the relationships in question (spatial and causal, respectively).

A brief consideration of Woodward's 'agent causal viewpoint' might help to bring out the point. A hallmark of an 'agent causal viewpoint', as Woodward explains, is the ability to learn about causal relations from the observation of other's actions. Developmentally, some such ability already appears to be in place relatively early on. In experiments carried out by Andrew Meltzoff (1990), for instance, 14-month-old children saw an adult touching the top of a box with his forehead, whereupon a light came on inside the box. They saw this action only once, and the box was removed as soon as the adult had performed the action, without giving the child the chance to touch the box herself. What Meltzoff found was that the 14-month-olds, when re-introduced to the box after one week's delay, reproduced the action they had seen the adult perform.

[7] Initially I thought that what Woodward was aiming at with his description of a 'fully causal viewpoint' was just the viewpoint of an agent who possesses causal concepts, as I understand it. However, an anonymous reviewer has persuaded me that this is not the case, and that the two issues are in fact orthogonal.

Meltzoff's particular interest lies in the connection between children's imitative abilities, as demonstrated in this study, and the development of what is typically called a 'theory of mind', i.e., a capacity to reason about oneself and others using mental and action-concepts. One possible position one might take here is that imitative abilities require a theory of mind, in the sense of an ability, on the part of the child, to work out through reflective reasoning that the same action that she sees the adult perform is also within her ability. However, on the basis of studies such as the one described above, as well as others that demonstrate basic imitative abilities already in newborns, Meltzoff argues against this view. Rather, he takes such studies to provide evidence for the innate presence of a common mental code, i.e. an underlying set of neural representations coding for certain aspects of goal-directed actions that are activated both when a child carries out an action herself, and when she sees the same action being performed by another person.

Meltzoff likes to frame his own view in terms of the idea that children, more or less from birth, have a grasp of others as being 'like me', which is manifested in their imitative abilities. If the above is correct, however, it shows a way of understanding how a child might have a grasp of others being 'like me' that is not a matter of any ability, on the part of the child, to engage in reflective thought about the relationship between her own actions and those of others. Rather, in the case of imitation drawing on representations involving a common code for self-performed as well as observed actions, the child's grasp of the other as being 'like me' may be a purely practical one, in the sense that the only way in which the child can put such representations to use may be in her own practical engagements with the world. In other words, the child's sense of the other as being 'like me' may simply be exhausted by her ability to acquire certain pieces of practical knowledge through observation,

e.g. how to activate the light-box in Meltzoff's (1990) experiment with 14-month olds.[8]

In a similar way, I would argue that there could possibly be a number of ways in which a subject might acquire practical knowledge from observation of naturally occurring events in the environment without necessarily engaging in reflective thought about the relationship between those events and her own actions. Thus, when it comes to Woodward's idea of a 'fully causal viewpoint', and the idea that a subject with such a viewpoint can grasp that the same relationship governing her own actions can also be present when no agent is involved, we can again distinguish between two forms such a grasp might take. The subject might be able to give substance to that sameness in reflective thought, or she might have a more basic, purely practical grasp of that sameness, which is simply exhausted by the ability to acquire practical knowledge through observation of naturally occurring events.

3. Causing and Preventing

In the preceding two sections, I have tried to bring into focus the question of what it is for a subject to possess causal concepts by highlighting two ways in which a subject might fall short of having such concepts despite being able to recognize or being sensitive to what are in fact causal phenomena. In the remainder of the paper, I will sketch a suggestion for an answer to the question as to what it is to possess causal concepts.

I shall take my lead from the final sentence of Anscombe's 'Causality and Determination', which runs: "The most neglected of the key topics in this subject are: interference and prevention" (Anscombe, 1971, p. 258). The main focus in what

[8] I discuss the issues at stake here in more detail in Hoerl, 2002.

follows will be on the notion of prevention, though I should make clear at the outset that I am treating this notion as an example of a larger class. My suggestion, put very briefly, is that ascribing a grasp of causal concepts to an individual is, at least in part, a matter of thinking of that individual as being able to engage in a distinctive type of causal reasoning, of which reasoning about how a certain outcome might be prevented provides an example. In this type of reasoning, causal relations are themselves being conceived of as subject to conditions, such that it is possible to ask how A might be prevented from causing B, how one might enable A to cause B, or how else the relationship between A and B might be interfered with, such that A might bring about one effect rather than another.[9]

To spell out this suggestion in more detail, it is important to distinguish between two quite different ways in which the notion of prevention might be understood. One of them is expressed in the following remark by Hitchcock:

> In our interventions in the world, we typically seek to promote those outcomes that we deem desirable and to prevent or inhibit those outcomes that we wish to avoid. Thus, we typically aim to prevent or inhibit thrombosis, while we may seek either to prevent or promote pregnancy, depending on our desires at a particular stage in our lives. […] Unfortunately, our language can be treacherous here, for we often use the word *cause* specifically to mean *promote*. In this usage, cause and prevent (or inhibit) are antonyms.

[9] The resulting account of causal concepts is likely to share certain features with manipulability theories of causation such as Woodward's, in that it connects our grasp of causal relations to our ability to intervene in the world ourselves. But the type of intervention it focuses on is not the one captured by the notion of an intervention that is commonly used in manipulability theories. Put very crudely, the basic idea behind manipulability theories is that A causes B only if there is an association between A and B that is invariant under interventions on A (see Woodward, 2003, for an account that offers a considerable amount of refinement of this basic idea). By contrast, the notion of an intervention that is crucial in the present context is not that of an intervention on A, but rather that of an intervention that determines, on a given occasion, what effect A has, even if A cannot be (or is not) intervened on.

> Nonetheless, prevention is a kind of causal relationship (Hitchcock, 2007, p. 104).

Hitchcock's point is that causes can come in varieties that raise the chances of a certain outcome or varieties that lower the chances of a certain outcome, and that, for instance, the relationship between birth control pills and pregnancy is clearly a causal one, even though we ordinarily would not say that the former cause the latter. Within the framework used by Hitchcock, this ambiguity is resolved by distinguishing between preventing causes, which are conceived of as ones that lower the probability of a certain outcome, and promoting causes, which are conceived of as ones that increase the probability of the outcome.

The understanding of the notion of prevention at issue in Anscombe's 'Causality and Determination' is arguably very different from that just sketched. A key part of Anscombe's argument is that singular causal truths cannot be understood as instances of universal generalizations because, for any putative causal relationship between an event A and an event B, we can imagine a situation in which A happens, but B does not because of some further cause, and there is no way of limiting the number of such other causes that might prevent B from happening.[10]

Setting aside the issue as to what exactly this argument establishes about the relationship between causation and universal generalizations, it brings out that there is a further dimension to common-sense uses of the notion of prevention that Hitchcock glosses over. An understanding of the notion of prevention that is more complex than that captured by the idea of a preventing cause, understood along Hitchcock's lines, is expressed in sentences with a particular sort of grammatical structure, following the

[10] Russell (1913) makes what amounts to more or less the same point, but uses it to a different effect, *viz.* in support of an eliminativism about causation. See Field (2003) for discussion.

schema "A prevents B from x-ing C". (Indeed, *pace* Hitchcock, it is likely that we actually have a version of this schema in mind when we think of the relation between taking birth control pills and getting pregnant.) The grammatical structure here describes not a relationship between two events (or event types), such as that of one event's lowering the probability of another, but a relationship that obtains between an event (or type of event) and an occurrence of an instance (of a type) of causation. In what follows, I will mark that difference by speaking of the schema "A prevents B from x-ing C" as embodying the idea that causal relations are preventable.[11]

Now, one interesting feature of the schema "A prevents B from x-ing C" is that it can be filled using both the thin concept 'cause' or one of Anscombe's thick concepts to describe the prevented relationship between B and C. This might help us see a possible answer to a question we raised in connection with Anscombe's account, i.e. the question of what it is about a subject's ability to recognize instances of scraping, pushing, wetting, etc. in virtue of which the subject can be said to think of those instances in genuinely causal terms. One idea that Anscombe might avail herself of, in response to this question, is the idea that instances of wetting, pushing, scraping, etc. are thought of as causal by the subject in so far as the subject understands, for instance, that they are preventable (at least in principle), or that they can otherwise be interfered with.

Again, we can perhaps make this suggestion more concrete by looking at the two studies on children that I mentioned in the section on Anscombe's account. In Gelman et al.'s study, children were able to match an outcome to an object with the right kind of causal power to produce that outcome. Having some kind of grasp of the causal powers of objects in this sense, though, might be possible without grasping that

[11] Dowe offers a similar analysis of prevention in chapter 6 of Dowe (2000).

something might prevent the object from exerting its causal power, or how else the presence of other causes might affect the outcome. Das Gupta & Bryant's study, by contrast, whilst not specifically dealing with children's grasp of preventability, can nevertheless be seen as measuring their ability to appreciate the particular conditions that instances of causal relations may be subject to, such that, for example, whether an instance of pouring water over an apple produces a whole or a cut wet apple depends on what conditions the apple starts out in. Arguably, the reasoning here, involving a grasp of one sense in which a causal relationship can be interfered with, is structurally of the same type as that involved in grasping relations of preventability.

Consider next the issues I raised in the section on causal models. The key issue was that a subject may possess a causal model and yet fall short of possessing causal concepts, if the only way in which the causal model is put to use by the subject is in her own practical engagements with the world. How, if at all, might the question as to whether the subject can grasp the idea of the preventability of instances of causation be relevant here? Consider the relationship between two events (or event types) as represented within a causal model. There may actually be nothing about that representation (specifying certain counterfactual relationships between the two events) that the idea of something preventing one of them from causing the other could get any purchase on.[12] This is not to say that causal models are unable to represent a situation to which the schema "A prevents B from causing C" applies – doing so within a causal model usually involves introducing a further variable in

[12] The causal model might, of course, imply that we can, say, bring it about that the second event doesn't happen by bringing it about that the first event doesn't happen. But this is different from the notion of prevention as considered here, which is concerned with the question of how the second event might be prevented, even if the first does happen.

between B and C for A to influence.[13] Rather, the point here is that the thought that any causal relation is in principle preventable involves a substantial theoretical commitment that goes beyond what is implied by the information about counterfactual relationships captured by a causal map.

It is at this point, I believe, that the issues I have been raising connect in an interesting way with recent work on singular causation. I believe it is arguable that thought of causal relations as in principle preventable involves a commitment to a view of causation that is at least broadly of the same type as that advanced in recent philosophical accounts of singular causation, according to which singular causation involves, for instance, the interaction of continuous processes and the exchange of physical quantities such as energy or momentum (Salmon, 1984; Dowe, 2000). Following Schaffer (2004), I will say that a common feature of such accounts is that they think of causes and effects as physically connected. It is this idea of a physical connection, however, that seems to be required to give substance to the thought that causal relations are in principle preventable.[14] Above, I said that thinking of causal relationships as preventable is an instance of a distinctive type of reasoning about causation in which causal relations are themselves being conceived of as being subject to certain conditions. The idea of a physical connection, of an empirical feature by which causes need to be connected to their effects, is precisely what is needed to ground the idea of such conditions, and thus for this type of reasoning to get a grip.

To summarize, then, I have suggested that there is a particular type of reasoning that a subject needs to be able to engage in if she is to be credited with

[13] Interestingly, Woodward (2007) also notes the significance of a grasp of intermediate variables for what he calls a fully causal viewpoint, but the rest of his argument is somewhat different from the one presented here.

[14] and also for making intelligible, in the kind of case at issue in the Das Gupta & Bryant study, how the outcome of a causal intervention can differ depending on the presence of other causes.

causal concepts. In this type of reasoning, causal relations are themselves conceived of as being subject to certain conditions, such that it makes sense to ask how A might be prevented from causing B, how one might enable A to cause B, or how else the relationship between A and B might be interfered with. I have argued that there is a primitive ability to grasp the causal power of objects and there are abilities involving the use of a causal model that a subject can possess without being able to consider such questions and thus to think of causal relations as subject to conditions in this way. Conversely, I have tried to show that the latter abilities involve thought of causes and effects as being physically connected. Thus, if what I have argued is at least roughly along the right lines, having causal concepts turns on being able to engage in a type of reasoning that involves a commitment to the type of view of causal relations articulated by current singularist theories of causation.

Acknowledgements

Many of the ideas in this paper go back to discussions I have had with Teresa McCormack, who introduced me to the psychological literature on children's causal reasoning. I have also benefited a great deal from comments and suggestions by Stephen Butterfill, Monika Dullstein, Naomi Eilan, Hemdat Lerman, Guy Longworth, Johannes Roessler, Matthew Soteriou and an anonymous referee for the journal. Work on this paper was carried out within the context of the Arts and Humanities Research Council (AHRC) Project on Causal Understanding, and some of it while I held a Research Leave Grant from the AHRC.

References

Anscombe, G.E.M. (1971). Causality and determination. Reprinted in: P. Van Inwagen & D.W. Zimmermann (eds.). Metaphysics: The big questions (pp. 244-258). Oxford: Blackwell 1998.

Campbell, J. (1994). Past, space and self. Cambridge, MA: MIT Press.

Campbell, J. (1995). Shape properties, experience of shape and shape concepts. Philosophical Issues 7, 351-363.

Cartwright, N. (2004). Causation: One word, many things. Philosophy of Science 71, 805–819.

Das Gupta, P. & Bryant, P.E. (1989). Young children's causal inferences. Child Development 60, 1138-1146.

Dowe, P. (2000). Physical causation New York: Cambridge University Press.

Evans, G. (1985). Things without the mind. In: Collected papers (pp. 249-290). Oxford: Clarendon Press.

Field, H. (2003). Causation in a physical porld. In: M. J. Loux & D. Zimmerman (eds.): The Oxford handbook of metaphysics (pp. 435–60). Oxford: Oxford University Press.

Gelman, R., Bullock, M., & Meck, E. (1980). Preschoolers' understanding of simple object transformations. Child Development 51, 691-699.

Glymour, C. (2001). The mind's arrows: Bayes nets and graphical causal models in psychology. Cambridge, MA: MIT Press.

Godfrey-Smith, P. (forthcoming). Causal pluralism. In: H. Beebee, C. Hitchcock & P. Menzies (eds.). The Oxford handbook of causation. Oxford: Oxford University Press.

Gopnik, A., Glymour, C., Sobel, D., Schulz, L., Kushnir, T., & Danks, D. (2004): A theory of causal learning in children: Causal maps and Bayes nets. Psychological Review 111, 1-31.

Hitchcock, C. (1995). The mishap at Reichenbach Fall: Singular vs. general causation. Philosophical Studies 78, 257-291.

Hitchcock, C. (2001). The intransitivity of causation revealed in equations and graphs. The Journal of Philosophy XCVIII, 273-299.

Hitchcock, C. (2007). On the importance of causal taxonomy. In: A. Gopnik & L. Schulz (eds.). Causal learning: Psychology, philosophy and computation (pp. 101-114). New York: Oxford University Press.

Hoerl, C. (2002). Reply to Jean Decety: Perceiving actions and understanding agency. In: J. Dokic and J. Proust (eds.): Simulation and Knowledge of Action (pp. 73-85), Amsterdam: John Benjamins.

Hoerl, C. & McCormack, T. (forthcoming). Time in cognitive development. In: C. Callender (ed.): The Oxford handbook on time. Oxford: Oxford University Press.

Meltzoff, A.N. (1990). Towards a developmental cognitive science: The implications of cross-modal matching and imitation for the development of representation and memory in infancy. In: A. Diamond (ed.): Annals of the New York Academy of Sciences: Vol. 608: The development and neural bases of higher cognitive functions. New York: New York Academy of Sciences.

Nelson, K. (1996). Language in cognitive development: Emergence of the mediated mind. Cambridge: Cambridge University Press.

Pearl, J. (2000). Causation. Cambridge: Cambridge University Press.

Russell, B. (1913). On the notion of cause. Proceedings of the Aristotelian Society

13, 1–26.

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. Oxford: Blackwell.

Schank, R. & Abelson, R (1977). Scripts, plans, goals, and understanding: An inquiry into human knowledge structures. Hillsdale, NJ: Lawrence Erlbaum.

Spirtes, P., Glymour, C., & Scheines, R. (2000). Causation, prediction and search (2nd ed). Cambridge, MA: MIT Press.

Williams, B. (1985). Ethics and the limits of philosophy. Cambridge, MA: Harvard University Press.

Woodward, J. (2003). Making things happen: A theory of causal explanation. Oxford: Oxford University Press.

Woodward, J. (2007). Interventionist theories of causation in psychological perspective. In: A. Gopnik & L. Schulz (eds.). Causal learning: Psychology, philosophy and computation (pp. 19-36). New York: Oxford University Press.