

CAUSATION IN THE SPECIAL SCIENCES: THE CASE FOR PRAGMATISM

Huw Price
Department of Philosophy
University of Sydney

1. Introduction

One of the jobs of philosophers of the special sciences is to connect the local concerns of particular disciplines with those of philosophy in general. The two-way complexities of this task are well-illustrated by the case of causation. On the one hand—from the outside, as it were—philosophers interested in general issues about causation are prone to turn to the special sciences for real-life examples of the use of causal notions. On the other hand, from the inside, the special disciplines themselves throw up philosophical puzzles in which the notion of causation plays a role. When does correlation indicate causation, for example? Physics and economics both generate hard cases of this kind.

In principle, then, a philosopher of a discipline such as physics or economics occupies a rather exposed position, liable to be called as expert witness about causation from inside and out—by people who know more about the special discipline, and by people who know more about relevant parts of philosophy at large. This dual role may seem a trifle self-contradictory. After all, to address a causal puzzle within physics (say) a philosopher of physics needs to hold fixed a general philosophical account of causation; and therefore hasn't the luxury of remaining uncommitted about the philosophy of causation in general, the kind of cautious player who awaits the verdict of physics on such matters. But progress is surely possible by a kind of reflective equilibrium, an informed interplay between the two kinds of constraints, philosophical and scientific. The difficult task, the task which falls especially to the philosopher of the special disciplines, is to open and keep open the dialogue.

This paper is an attempt at that task. In particular, I want to recommend a particular way of thinking about causation—a kind of pragmatism—to philosophers of sciences such as physics and economics. The term pragmatism means several things in contemporary philosophy, and so I'll say more in a moment about the sense I have in mind. In the sense in question, pragmatism about causation is an uncommon view—and particularly so, I think, among philosophers of the special sciences. But it deserves to be taken much more seriously, in my view, especially within the philosophy of physics and economics.

For pragmatism about causation throws important light on puzzles within these disciplines, and considerations stemming from these disciplines provide some of the main reasons for endorsing the pragmatic approach.

I suspect that some of the resistance to this pragmatic approach within the philosophy of the special sciences stems from the feeling that it is too 'subjective', not sufficiently 'realist', for the purposes of the sciences concerned. Philosophers of the special sciences often show a commendable concern to defend the 'seriousness' of the sciences on which their own discipline is parasitic. However, the usual maps of the relevant philosophical territory are sometimes rather crude, I think. There are several varieties of realism, and several varieties of subjectivism, not often well distinguished. One of my aims here is philosophical cartography: I want to draw some maps to clarify on the one hand, what it means to be a realist about causation, in the philosophy of physics or economics; and on the other, the sense in which the kind of pragmatism I have in mind is and is not a subjectivist view. Roughly, I hope to show both that the view is not subjectivist in the sense most inimical to realism, and that to the extent that realism is an attractive position about causation, it is compatible with the kind of pragmatism I have in mind.

I also want to call attention to a fallacy, as I see it, in a popular defence of realism about notions such as causation. The argument in question relies on a supposed analogy between causation and the theoretical postulates of the special sciences. The suggestion is that postulating a causal relation is 'just like' postulating electrons or any other theoretical entity, so that (among other things) realism about causation is no more problematic than—or at least no different in kind from—realism about electrons. In my view, this supposed analogy is seriously flawed, for reasons which turn out to support a pragmatic approach to causation.

2. Three kinds of pragmatism

What is a pragmatic account of causation, or of any other philosophical topic? The term pragmatism is associated with at least three distinct philosophical views. One view, stemming from Peirce, takes pragmatism to be primarily a doctrine about truth: the view that truth is to be identified with some epistemological notion such as warranted assertibility. Another view, more closely associated with James and Dewey, takes the key to pragmatism to be an appeal to a notion of success. Like Peirce's pragmatism, this second view can issue in an account of truth. For example, it might be said that true beliefs are those which lead to successful behaviour. But the approach itself is not tied to such an account of truth: a pragmatist of this second sort might equally give a success-based account of other notions, such as mental content.

A third form of pragmatism, and the one which interests me here, is the view that a philosophical account of a problematic notion—that of causation itself, for example—needs to begin by playing close attention to the role of the concept concerned in the practice of the creatures who

use it. Indeed, the need to explain the use of a notion in the lives of ordinary speakers is often the original motivation for an account of this kind. Causal notions and their kin are ubiquitous in the everyday talk of ordinary people. Pragmatists argue that we cannot hope to explain this anthropological fact if we begin where metaphysics traditionally begins, at the level of the objects themselves—if we ask what causation is, if we begin by looking for something for causation to be, which will explain all these uses. Instead, pragmatists think, we need to start with the practice of using such notions, and to ask what role such notions play in the lives of the creatures concerned—why creatures like us should have come to describe the world in these causal terms.

The pragmatic approach to causality and related notions is certainly not new. Ramsey's views on probability, and his late views on causation and physical necessity, are of this kind, for example (1978a, 1978b). Take probability, for instance. Ramsey's view is that an account of probability needs to begin with the role of probability judgments in decision making. This connection needs to be there from the very beginning. We could not give a satisfactory account of probability which left it out initially, and then proceeded to note that the notion thus characterised had some relevance to decision contexts.

However, I think the distinctive character of the pragmatic approach is not well understood in contemporary philosophy. This is partly a result of some confusing terminology left to us by tradition. Pragmatism of this kind is thought to be subjectivist, or non-objectivist. But these terms too apply to several distinct approaches, often poorly distinguished, and pragmatism is easily mischaracterised. This kind of confusion contributes to misunderstanding of the relationship between pragmatism and realism, in my view. Roughly, realism requires objectivism in one sense, while pragmatism of this kind denies objectivism in another sense. The two views are thus less in conflict than commonly assumed.

In fact, it turns out to be useful to distinguish three kinds of objectivity—three somewhat independent respects in which a view may vary between objectivism and subjectivism. In one of these respects, subjectivism is indeed in conflict with realism, as intuitively understood. As we shall see, however, a philosophical view may be objectivist in this realism-related respect and yet subjectivist in either of the other two respects. In this way, pragmatism about causation remains compatible with realism: the respect in which pragmatism is subjectivist is not the respect in which realism requires objectivism, and the two views are not in conflict.

3. Three kinds of objectivity

What are the three varieties of subjectivity and objectivity? For present purposes, applied to the case of causation, they may be characterised as follows:

1. Causation is ontologically subjective if the existence of causal relations depends on the presence of minds, speakers, observers or

the like; if causal states of affairs are in this ontological sense mind- or observer-dependent.

2. Causation is topic-subjective if talk of causation is in part talk about speakers, agents, or humans—for example, if all causal claims are in part about our own psychological states.
3. Causation is practice-subjective if an adequate philosophical account of causation needs to make central reference to the role of the concept in the lives and practice of creatures who use it.

Some cautionary notes. Obviously, these characterisations are somewhat rough. In particular, I have not been careful about marking the fact that some of these characterisations apply more naturally to causation itself—the object or ‘thing’ in the world—and some to the term or concept. (More on this later.) And the list may not be exhaustive, in the sense that there may be other interesting varieties of subjectivity and objectivity which could usefully be added. Still, this three-way classification provides some useful distinctions. In particular, it will enable me to maintain that concerns about realism relate mainly to the ontological axis, whereas the pragmatism I advocate about causation involves practice subjectivity.

More on the last point in a moment. First, let me emphasise that pragmatism about causation is not the view that when we talk of causation we are talking about ourselves, in whole or in part. The latter view would involve topic-subjectivity, not practice-subjectivity. The corresponding distinction is often missed, I think, in the analogous case of subjectivism about probability. As a result, the view I attributed to Ramsey—the interesting and in my view plausible claim that probability cannot be characterised without reference to its role in decision-making—is confused with the view that speakers use probabilistic terms to talk about their own mental states.

If the difference between these views isn’t clear, think of what the two views say about a claim such as, “It will probably rain tomorrow.” According to the topic-subjective account, someone who says this is talking about her own state of mind—she is saying, in effect, that she is confident that it will rain tomorrow. All probability claims are thus psychological in content, at least in part. Not so according to the practice-subjective account: probability is not a psychological concept, on this view. True, an adequate philosophical account of the concept will need to mention psychological notions, in virtue of mentioning the role of the concept in decision making. But to say that an explication of the function of a concept needs to mention psychological notions is not to say that that concept has a psychological content: of course not, for otherwise all human language would have a psychological content, simply in virtue of the fact that an account of the function of language needs to be grounded in psychology. (To put it more simply, we might explain the function of the utterance “It will rain” by saying that it is used to give voice to the belief that it will rain; but this does not imply that the utterance is about that belief.)

I want to say a corresponding thing about causation. In my view, an adequate philosophical account of causation needs to begin with its role in the lives of agents, creatures who have the primitive experience of intervening in the world in pursuit of their ends. I have defended such a view in a number of places (see Price 1991, 1992a, 1992b, Menzies & Price 1993, and especially Price 1996, ch. 6), and noted its affinity to a late view of Ramsey (1978b). Here, I simply want to emphasise that the view is not the topic-subjective claim that talk of causation is talk about agents or agency, but rather the practice-subjective doctrine that we don't understand the notion of causation—as philosophers, as it were—until we understand its origins in the lives and practice of agents such as ourselves.

So much for the distinction between topic-subjectivism and practice-subjectivism. It remains to point out that neither commits us to ontological subjectivism, in any interesting sense. Because topic-subjectivism treats causation (or whatever) as in part a matter of psychology, it is ontologically subjective in the boring sense in which psychology itself is: if there were no thinking creatures, there would be no psychological states. As for practice-subjectivism, it is not even ontologically subjective in this uninteresting way. If the concept of causation is essentially tied to our experience as agents, as my kind of practice-subjective pragmatism maintains, then of course the concept would not arise in a world without agents. But this does not make it appropriate to say that if there had been no agents there would have been no causation. Pragmatism does not conflict with realism in that sense. In other words, as I noted above, realism seems to require ontological objectivism, and my kind of pragmatism about causation requires practice-subjectivism. These two requirements are not in conflict.

4. Realism responds?

At this point, some realists will feel that the compromise has been achieved at the cost—from my point of view—of making pragmatism a rather uninteresting view. These realists have no objection to the view that an account of the concept of causation is practice-subjective, but see the compromise as resting on the fact that this is quite compatible with a demand for realist account of causation itself—that is, of the 'thing in the world', rather than simply the concept by means of which we refer to it.

I agree that by the lights of such a realist account, if there could be such a thing, the pragmatist view would be at best a supporting act, an interesting prelude to the main game. However, I want to cast doubt on the claim that there could be a main attraction of the kind the realist promises. I think there are good reasons for doubting such claims in philosophy in general, but here my goal is more specific: I want to show that there are particular reasons for being suspicious of this promise in the case of causation—for thinking that pragmatism retains centre stage in this case, whatever its fate elsewhere.

My argument turns on two main ingredients: first, a distinction between two species or grades of realism, one stronger than the other, and second, a rejection of an often-claimed analogy between causation

and the theoretical entities of first-order science. Roughly, I want to argue that the realist move just canvassed requires the stronger of the two grades of realism, but that this is even more problematic in the case of causation than it is in the supposedly analogous case of the theoretical entities of first-order science. In the sense in which realism about causation is defensible, then, it is also compatible with the view that pragmatism is the central ingredient in a philosophical account of causation.

In view of the importance of the supposed analogy, it will be useful to characterise the two grades of realism initially as two different responses to antirealism about theoretical entities. And it will be instructive, I think, to do this against a brief sketch of the most relevant historical background. I include this sketch for two reasons. First, I commented earlier on a feeling in contemporary philosophy of science that antirealism shows a lack of due respect for science, a lack of seriousness about the special disciplines. Sometimes antirealism seems to be thought of as the product of a recent, aberrant, anti-scientific turn. Yet a little history makes clear that the very opposite is true: for three hundred years, antirealism has often seemed the proper consequence of the kind of deep respect for science embodied in the empiricist tradition—of the desire, especially, to rid science of the vestiges of a degenerate and unempirical metaphysical tradition. This impression may have been wrong, of course, but contemporary friends of science—philosophers of physics and economics, for example—do well to be careful. Prescientific metaphysics may lurk in the guise of scientific realism.

The second virtue of the following sketch lies in the fact that another feature of this historical tradition, a feature present from its beginnings in the seventeenth century, is the desire to distinguish what is 'objective' in the world from what in some sense 'comes from us'. I want to take the opportunity to link some of the milestones of this long project to the three-way classification of kinds of subjectivity set out above, and hence to the kind of pragmatism I advocate about causation.

5. The origins of the debate

The issue of objectivity is a major concern of many of the great natural philosophers of the seventeenth century, writers such as Descartes, Galileo, Boyle and Locke. One issue which concerns these thinkers is the distinction between what Boyle and Locke call the primary and secondary qualities. Roughly, the primary qualities are those which exist and have their natures independently of sentient observers, and are therefore apt to be studied by a science of the natural world. The secondary qualities are those of which this is not true—those which are 'observer-dependent', in some sense. Here, for example, is a beautiful early account of the distinction, from Galileo:

“I feel myself impelled by the necessity, as soon as I conceive a piece of matter or corporeal substance, of conceiving that in its own nature it is bounded and figured in such and such a

figure, that in relation to others it is large or small, that it is in this or that place, in this or that time, that it is in motion or remains at rest, that it touches or does not touch another body, that it is single, few, or many; in short by no imagination can a body be separated from such conditions; but that it must be white or red, bitter or sweet, sounding or mute, of a pleasant or unpleasant odour, I do not perceive my mind forced to acknowledge it necessarily accompanied by such conditions; so if the senses were not the escorts, perhaps the reason or the imagination by itself would never have arrived at them. Hence I think that these tastes, odours, colours, etc., on the side of the object in which they seem to exist, are nothing else than mere names, but hold their residence solely in the sensitive body; so that if the animal were removed, every such quality would be abolished and annihilated.”

Il Saggiatore (quoted by Burt (1932), p. 75).

In this passage, the distinction is drawn in terms of the ‘location’ of the two kinds of qualities. Primary qualities lie in the world, secondary qualities (despite appearances and common beliefs to the contrary) lie in us. Thus secondary qualities are seen as topic-subjective, in the terms distinguished above. As a result, they are also ontologically subjective: their existence depends on the presence of the ‘sensitive body’.

In the tradition which follows, however, this ontological issue soon becomes blurred. Perhaps secondary qualities are in the world, after all, but are of a subject-involving nature—in ascribing such a property to the world we say that the world is such as to affect sentient creatures in certain ways. This is closer to Locke’s view, and while it remains topic-subjective, but is less ontologically subjective, for presumably the world might have such dispositional properties in the absence of any actual sentient creatures.

A Lockean account of this kind is easily transformed into a view which is practice-subjective, rather than topic-subjective. For example, suppose that the Lockean view says that to be red is to be disposed to produce a certain response R in humans. Then there is a corresponding practice-subjective account of the concept red, a central feature of which will relate the use of this concept to a speaker’s having response R. (On the most crude account, the concept is simply used to give voice to this response.) Thus as in the case of subjectivism about probability, the topic-subjective view interprets “That’s red” as in part about human psychology, whereas the practice-subjective view does not. (I have argued elsewhere (Price 1998) that this practice-subjective approach to pragmatism has general advantages compared to its Lockean twin. Here, the important point is just that there are two distinct possibilities.)

However the subjectivism of the secondary qualities is characterised, the issue arises as to whether any properties are purely primary, in the original sense. Perhaps the properties in terms of which we describe the world are never the essential properties of the things in themselves;

perhaps we only ever know and describe things in terms of their appearances from our own point of view. This thought crops up in later philosophy of science in 'structuralist' or 'dispositionalist' views of theories, which claim that science always describes only the general patterns or structures in which reality manifests itself to us, not reality itself. One source of this view is Kant, whom Putnam (1981, p. 60) suggests may be read as arguing that all properties are secondary. (Putnam calls the result a kind of pragmatism.)

These structuralist views illustrate a very deep tension in empiricism since the seventeenth century, between the metaphysical desire to know the world as it really is, via empirical science, on the one hand; and the epistemological desire that our knowledge claims be well-grounded in observation, on the other. Metaphysics pulls us outwards, epistemology pulls us inwards. In structuralist views, the view that science searches for inner natures has been dropped, in favour of the view that it searches for something more accessible, by empiricist lights. This is far from an isolated case. In empiricism in general, the advantage of relatively subjective approaches is thought to be that they avoid unjustified inferences to things 'beyond' the observable.

From Hume onwards, moreover, the inward pull of epistemology has been reinforced by a concern about meaning. If, as many empiricists have held, the acquisition of meaning needs ultimately to be grounded in a novice speaker's observational experiences, then meaningful talk about unobservable matters seems impossible. Metaphysical claims seem in danger of being empty, rather than merely unjustified.

These considerations are the main source of the problem of the status of theoretical terms in science, and the motivation for antirealism about such terms. As I noted earlier, these empiricist antirealists are not enemies of science, as a casual reader of contemporary debates might assume. On the contrary, the antirealists' goal is to save the truly scientific core of what passes for science, by ridding it of the remnants of the old metaphysics—to free their beloved empirical science from the grip of a degenerate and anti-empiricist style of philosophy.

One of the more interesting forms of antirealism is instrumentalism, which would save science from metaphysics by denying that the function of theoretical discourse is to describe a reality 'beyond' observation (and so denies that theoretical terms have a referential function). On the contrary, an instrumentalist says, the significance of theoretical terms lies in their role in a particular human practice—a practice whose function lies in its ability to predict future observations. So this approach is pragmatic, in the third sense I distinguished at the beginning.

Instrumentalism thus illustrates one way in which pragmatism may become a global position, or at least a position with very wide application. I mentioned another view of this kind a moment ago, in Putnam's interpretation of Kant, as a philosopher who believes that all properties are secondary properties. It may seem that these views pose a problem for the project of defending a pragmatic approach to causation, in the sense that they promise an embarrassment of riches. If

pragmatism promises to become a global view, what need is there for argument which picks out the case of causation?

However, there is no substantial tension here with my present project. My present argument turns on the fact that realism and objectivism about causation are importantly different from, and more problematic than, the corresponding views about the theoretical postulates of science. Put another way, I am arguing that there are reasons for pragmatism about talk of causation which are not reasons for pragmatism about (say) talk of electrons. The differences on which this argument depends are compatible with the possibility that there might be other reasons for endorsing pragmatism of some global variety.

6. Two grades of realism

There are two popular ways of trying to avoid antirealism about the theoretical entities of science. One, which I shall call strong realism, relies on so-called ‘inference to best explanation’. On this view, our reason for believing in the reality of (at least some of) the theoretical entities and states of affairs postulated by science is that if these things exist or obtain, they provided the best causal explanation of the available observational data, broadly construed. This approach thus rejects Russell’s (1917) view that science should avoid inferred entities wherever possible, and defends the legitimacy of such inferences, at least in causal–explanatory contexts. The causal–explanatory project provides the framework which legitimates what many empiricists saw as idle and ill-founded metaphysics.

The second popular brand of realism in contemporary philosophy of science is what may be called weak or minimal realism. This view simply takes the existence claims of science at face value, and rejects any ‘additional’ metaphysical or philosophical viewpoint from which it would make sense to ask ‘Do these things (electrons, for example) really exist?’ The key to weak realism is the rejection of a standpoint for ontology beyond that of science. The most famous proponent of this view is therefore perhaps Quine, for whom the rejection of such an ‘external’ ontological standpoint is the key doctrine of papers such as ‘On What There Is’ (Quine 1953, ch. 1). More recently, for example, it is prominently embodied in Arthur Fine’s (1986, ch. 7) advocacy of what he terms the ‘natural ontological attitude’.

I am following convention in calling this view a species of realism. However, it is also instructive to see the view as rejecting the traditional realist–antirealist debate altogether, at least as that debate arises within the empiricist tradition. Roughly, a minimalist wants to say that within this debate, both sides presuppose that there is a legitimate standpoint for philosophy, apart from that of science, from which questions such as the following make sense: Do the theoretical entities postulated by our best scientific theories really exist? Do the theoretical terms of these theories really refer, and if so to what? Do theoretical claims really have truth values? Realists and antirealists give different answers to these questions, but agree that the questions

themselves (or variants of them) are well-posed, in the intended ‘supra-scientific’ sense. But minimalism itself challenges this assumption, and thus rejects the debate. As Fine (1986, ch. 8) makes clear, the view is neither realism nor antirealism, in the old sense.

Among the challenges that weak realism offers to its stronger cousin is a dilemma concerning inference to best explanation itself. Is the employment of this inference itself part of science? If so, then it doesn’t take us beyond what the weak realist already accepts—it doesn’t yield an ontological standpoint beyond science. (An instrumentalist of the above sort should have no objection to it, for example.) If not, then how is it to be justified? How could it have a justification which wasn’t already available to science?

My sympathies lie with the weak realist on this point. However, my present purpose is not to adjudicate on the issue of the proper form of realism about theoretical postulates, but to use that issue as background to the issue concerning realism about causation. In particular, I want to argue that the case against strong realism is even stronger in the case of causation than in the case of theoretical entities such as electrons. Someone who wants to argue for strong realism about causation by analogy with theoretical entities thus faces two kinds of challenge: those intrinsic to the debate in the philosophy of science in general, and those which depend on the fact that the analogy itself is problematic, in the way I want to explain.

7. Is causation a theoretical postulate?

As I have noted, some philosophers suggest that causation and other modal notions are on a par with the theoretical entities postulated by science. (If there is a difference, it is simply that unlike electrons, causation is already a part of prescientific ‘folk theory’.) On this view, our reasons for treating causation realistically are essentially the same as those for treating electrons realistically.

If what is at issue is strong realism about causation, however, then we run into problems. As we have seen, the main argument for strong realism about theoretical entities goes in terms of inference to explanatory causes. But this reason simply takes the notion of causation for granted, and therefore can’t be applied in support of realism about causation. In this context, the supposed role of inference to best explanation is epistemological—it is supposed to justify a belief in the reality of entities of a certain kind. My point is that such an attempt at justification would be viciously circular in the case of causation itself, in virtue of the fact by the realist’s own lights, the inference presupposes realism about explanatory causes.

A similar point can also be made in a metaphysical key. Suppose we are interested in investigating the nature of causation, in the manner of the realist we imagined in Section 4, who saw pragmatism as a conceptual prelude to metaphysics. Again, the analogy with the theoretical entities of science breaks down. In science we learn more about an entity by learning more about its causes and effects; we postulate theoretical entities in terms of causes and effects. These

techniques reduce to nonsense if we try to apply them to investigate causation itself. Causation is the medium in which the investigation is conducted, the thread which leads to unseen objects. Treat the thread itself as unknown, and we are simply blind.

To make the same point in philosophical vocabulary which will be familiar to many, consider David Lewis's account of theoretical reduction (1970, 1972). According to Lewis, electrons are whatever entities actually play the 'causal roles' defined by electron theory. To accept electron theory is to accept that there is something in reality which plays these roles. What we know about what electrons actually are is what we find out by empirical means about the entities which do these causal jobs. Once again, this approach would be circular if we were to try to apply it to causation itself. The term 'causation' is employed in the reductive machinery itself, and therefore cannot be processed by that machinery.

It is true that the immediate problem can be avoided if causal roles are replaced by semantic roles, so that, for example, electrons are said to be the referents of the appropriate terms in electron theory. But in this case the circularity re-emerges if the approach is applied to the semantic notions themselves—or if reference is defined in causal terms!

The immediate problem can also be avoided if causation is analysed in terms of other modal notions—if causation is reduced to counterfactuals and possible worlds, for example. However, it then re-emerges if we attempt to use the same argument in favour of realism about those notions themselves. To avoid this complexity, I shall assume for the moment that our topic is not causation but modality in general.

The general problem for modal realists may then be characterised as follows, in two keys. In the epistemological key, it is that if the justification of theoretical beliefs in science relies on modal presuppositions, such as those that underlie inference to best explanation, then the justification of modal beliefs cannot be of the same kind, on pain of a threat of vicious circularity—the threat that the proposed justification presupposes some of the modal claims in question. In the metaphysical key, it is that the investigation of modal reality cannot rely on the standard tools of empirical science, to the extent that these tools are themselves constructed of modal materials.

Thus the supposed analogy between causation (and modal notions generally), on the one hand, and the theoretical entities of first-order science, on the other, breaks down in important ways. Even if strong realism were unproblematic with respect to the theoretical postulates of science, that would not count in favour of strong realism about causation. Indeed, the conclusion seems even worse than that, from the strong realist's point of view. The argument suggests that strong realism about theoretical entities is bound to be in trouble, by the realist's own lights, unless strong realism about causation is already secure. And the point applies whether the realist's concerns are epistemological or metaphysical: whether she wants to justify a belief in the reality of causal relations, or whether she wants to investigate their nature.

So much for the bad news for strong realism about causation. Now to the good news for pragmatism. In the epistemological key, the point just made is that the inference to modal conclusions from non-modal premises does not parallel that to theoretical beliefs in science. Inferences of the latter kind presuppose the relevant modal background in a way in which inferences of the former kind could not. Let us call this the problem of upward inference (inference to modal conclusions).

Modal realists also face a problem about inferences in the opposite direction. We may call it the problem of downward inference, or the application problem. It turns on the issue as to why objective modal facts should 'matter' to us in the way that they do. Why should we take them into account in decision-making, for example? Consider, for example, an account of possibility in terms of the existence of real 'possible worlds', distinct from our own world. If there were such worlds, why should they matter to us when we make decisions in this world? Wouldn't everything be just the same for us even if those other worlds didn't exist? Similar problems arise for other objectivist views about modal notions. If probability is objective, why should it have its assumed relevance to human decision making?

As in the case of the problem of upward inference, this downward or application problem doesn't arise with respect to theoretical entities in science. In that case, roughly speaking, application or relevance is simply defined in terms of causality. The relevance of theoretical commitments lies in their causal consequences. But again, this would be viciously circular if we tried to apply it to causation itself.

In my view, the pro-pragmatist lesson to be drawn from these considerations is that in considering causation, probability, and other modal notions, philosophy needs to begin with the practice of which the upward and downward inferences are already a part. If we think of the issue of the status of this practice as an adjunct to a strongly realist account, an account which begins by postulating the existence of a realm of entities of a certain kind, then it will always seem a mystery why talk of such entities should play the role that it does in our lives (why probability judgements should have the role they do in decision, for example).

In a sense, the supposed analogy with theoretical entities simply highlights this difficulty. It encourages us to think of causation and the like as part of the furniture of the world, on a par with the aspects of the world investigated by the sciences. And this very parity then becomes the problem. Why should particular pieces of furniture—the modal pieces, not the others—have the special significance that they do? (Don't say that we can ask the same question about electrons. In that case the answer is in causal terms.)

To sum up. The supposed analogy with the theoretical entities of science fails as a defence of 'strong' realism about causality and other modal notions. Moreover, because the failure turns on the role that the modal notions themselves play in arguments for strong scientific realism, their failure exposes a deep problem for modal realism, and a deep advantage of the pragmatist approach.

What is the realist about causation to do at this point? One option is to be less ambitious, by espousing only weak or minimal realism. Won't the analogy with theoretical entities survive, so long as we are weak realists in both cases?

In my view, its survival is doubtful, even in this context. If inference to best causal explanation remains an important tool within science, it remains a tool which cannot be used by a scientist to lay bare the causal relation itself. So the analogy still breaks down. However, for present purposes I want to emphasise a different point concerning weak realism about causation and modality, namely that this form of realism is perfectly compatible with pragmatism, in the third sense outlined at the beginning.

8. Pragmatism and minimal realism

The defining doctrine of weak or minimal realism is that the first-order claims of science or ordinary usage don't need to be 'validated' from some external philosophical perspective. This view is sometimes characterised as a kind of quietism. It denies the need for—and, usually, the coherency of—a certain kind of philosophical enquiry. Clearly, this kind of quietism is as intelligible with respect to our commitments about causation as it is with respect to those about electrons. In both cases it amounts to the same advice: Stop doing what tradition has done. Walk away from the realist–antirealist debate, at least in this case.

Sometimes quietism of this kind goes as far as recommending that we walk away from philosophy, or at least from a major part of its concerns. However, this stronger recommendation is inessential, and in my view misguided. To reject an externalist ontological viewpoint need not be to reject all kinds of philosophical reflection on the topics in question. In particular, it is quite compatible with taking seriously questions like these: Why do we humans go in for talk of causation and other modal notions? What role does this 'talk' play in our lives, and to what extent does this role depend on particular features of our situation (on the fact that we are agents, for example)? How, if at all, does causal and modal talk differ in these respects from talk of electrons, or tables, or numbers, or whatever?

These are precisely the questions that a pragmatist seeks to answer—a pragmatist in my third sense, the sense I take to exemplified by Ramsey's account of probability, and especially by his late account of causality and laws of nature. In my view, this kind of pragmatism embodies what turn out to be the philosophically interesting issues about causation and the modal notions generally, once we set aside strong realism, and the ill-grounded idea that the status of modal notions is 'just like' that of the theoretical postulates of the special sciences.

One aspect of the realism thus rejected is the view that in order to validate talk of causation in science and ordinary life, we need to find something in the physical world with which to 'identify' causation. I have argued that this reductive project cannot simply ride on the back of scientific reductionism, for the latter takes causal notions for granted in a way in which the former cannot. However, the project has many admirers in contemporary philosophy, and in rejecting it in the case of

causation we are swimming against a strong tide. So I want to close by mentioning another kind of objection to such reductive accounts of causation, which I have developed in much more detail elsewhere (see especially Price 1996, ch. 6). This objection turns on the claim that at least one central feature of causation—roughly, its ‘directedness’—simply isn’t there to be ‘found’ in the appropriate place in physical reality. Again, pragmatism does well where realism fails. The feature in question falls naturally into place in an account of the human origins of the concept of causation.

9. Realism, pragmatism and the arrow of causation

Causation seems ‘directed’ in two senses. First, the causal relation is asymmetric, in the sense that if A is a cause of B then B is not a cause of A. Second, this ‘causal arrow’ is strongly aligned past-to-future. Effects don’t occur before their causes. As a result, causation comprises one of the main ways in which the world seems asymmetric in time. So there is a puzzle here: Where does this asymmetry come from, especially in the light of the temporal symmetry of the underlying physical laws. But this puzzle is secondary to that of the causal arrow itself: What distinguishes cause and effect? What determines which way the causal arrow ‘points’?

Let us think about these puzzles from the point of view of a strong realist, who thinks that there is a well-posed issue about the real nature of causation, of the ‘thing in the world’. From such a point of view, an account of the causal arrow seems to require two things: an understanding of the asymmetry, or ‘directedness’ of the real relation itself; and an answer to the question as to why this intrinsic directionality should show the temporal preference that it does.

Most contemporary realists would be reluctant to treat these aspects of causation as brute facts about the world. (One reason for reluctance is the difficulty of connecting brute facts with practice. How could we tell if the brute facts were otherwise, for example?) Many, therefore, would prefer to seek some reductive account of causation, and some appropriate connection with other temporal aspects of the world.

The approach then becomes what I have called the third arrow strategy (Price 1996, ch. 6). This strategy tries to account for the directedness of causation by reducing causation to some feature of the physical world which is itself asymmetric or directed in the appropriate way—something which both has its own internal ‘arrow’ to give us the distinction between cause and effect, and something such that this internal arrow turns out to have the temporal alignment we find in the case of causation. This ‘third arrow’ then links those of causation and time.

Unfortunately for strong realism, the third arrow strategy is hard to reconcile with contemporary physics. Modern physics doesn’t seem to provide a third arrow, of the appropriate kind. The favoured candidate is the time-asymmetry associated with the second law of thermodynamics, but this turns out to be inadequate for the task at hand (see Price 1992a, 1996). For example, it doesn’t work at a sub-statistical level, to explain our causal intuitions about microphysics.

Realists have a fall-back position. They can concede that the causal relation is not really 'directed', in any non-conventional sense. The relation remains asymmetric in the logical sense—if A is a cause of B then B is not a cause of A—but the distinction between cause and effect is now regarded as merely conventional. In other words, it is analogous to the conventional distinction between up trains and down trains (on a railway line with no significant overall gradient). So long as the labelling convention imports a temporal asymmetry, by fiat or in fact, it saves the phenomena. Cause and effect are distinguished, and aligned in time.

This conventionalist strategy plays into the pragmatists' hands, however. The strategy comes in three main versions:

1. Follow Hume in stipulating that causes are the earlier and effects the later of pairs of events related in some otherwise time-symmetric way (e.g., as in Hume, by a relation of constant conjunction).
2. Follow Reichenbach (1956), so that the stipulation refers not to earlier and later, but to orientation with respect to some physical temporal 'signpost', such as the thermodynamic asymmetry.
3. Take the relevant 'signpost' to be provided by our perspective as agents: say that the cause is the element of the pair which we could 'manipulate', in principle, if the other element were our desired end. In this case the temporal orientation follows from our temporal orientation as agents.

The three options come to much the same thing in the end, I think, but because (3) promises to explain why such a labelling convention should be useful to us, it gets the genealogy right on the first pass. (1) and (2) need an additional step, to connect the convention with our own practice.

One way to see this is to ask why the causal arrow matters, on this view. Why should we bother to treat causation as temporally asymmetric? In other words, why should temporal information be embedded in the labelling convention for causation in this way? Option (3) contains the beginnings of an answer to this question, because it connects the asymmetric labelling convention to something which matters to us, and something which has connections with time. As agents, we act for the future, on the basis of information about the past.

However, the attractiveness of (3) is just the attractiveness of pragmatism. In backing-off as far as this, in other words, a strong realist approach has conceded that strong or reductive realism has nothing distinctive to tell us about one of the most distinctive features of causation, its the directedness and temporal orientation. Pragmatism thus has an important advantage over strong forms of realism in this case, an advantage which is quite independent of general issues between realism and its critics.

10. Conclusion: The case for pragmatism

I remarked earlier that within the philosophy of the special sciences, seriousness about science is sometimes taken to require realism about causation, of a kind which conflicts with pragmatism. I have tried to counter this thought, in a number of ways. In particular, I have argued that pragmatism is compatible with the kind of realism which perhaps takes science most seriously, in not taking it to be subject to ‘validation’ from some distinctively philosophical vantage point. I have also claimed that the breakdown of supposed analogy between causation and theoretical entities not only makes any stronger realism more problematic for causation than it is in science itself, but also highlights the central role of pragmatic considerations in any adequate account of causation.

If there is a single conclusion to be emphasised, it is this one. Being realist about causation is not like being realist about electrons, or any of the other postulates of the special sciences. Causation is different, and it is pragmatism, not realism, which offers an account of the difference.

REFERENCES

- Burtt, E. A. (1932). *The Metaphysical Foundations of Modern Physical Science*. London: Routledge & Kegan Paul.
- Fine, A. (1986). *The Shaky Game*. Chicago: University of Chicago Press.
- Lewis, D. (1970). How to Define Theoretical Terms. *Journal of Philosophy*, 67, 427–446.
- Lewis, D. (1972). Psychophysical and Theoretical Identifications. *Australasian Journal of Philosophy*, 50, 249–258.
- P. Menzies & Price, H. (1993). Causation as a Secondary Quality. *British Journal for the Philosophy of Science*, 44, 187–203.
- Price, H. (1991). Agency and Probabilistic Causality. *British Journal for the Philosophy of Science*, 42, 157–176.
- Price, H. (1992a). Agency and Causal Asymmetry. *Mind*, 101, 501–520.
- Price, H. (1992b). The Direction of Causation: Ramsey’s Ultimate Contingency. In D. Hull, M. Forbes and K. Okruhlik (Eds.), *PSA 1992*, Vol. 2. Philosophy of Science Association, pp. 253–267.

- Price, H. (1996). *Time's Arrow and Archimedes' Point*. New York: Oxford University Press.
- Price, H. (1998). Two Paths to Pragmatism II, in Casati, R. and Tappolet, C. (Eds.), *European Review of Philosophy*, Vol. 3, pp. 109–147.
- Putnam, H. (1981). *Reason, Truth and History*. Cambridge: Cambridge University Press.
- Quine, W. V. O. (1953). *From a Logical Point of View*. Cambridge, Mass.: Harvard University Press.
- Ramsey, F.P. (1978a). Truth and Probability. In D.H. Mellor (Ed.) *Foundations: Essays in Philosophy, Logic, Mathematics and Economics*. London: Routledge & Kegan Paul, pp. 58–100.
- Ramsey, F.P. (1978b). General Propositions and Causality. In D.H. Mellor (Ed.) *Foundations: Essays in Philosophy, Logic, Mathematics and Economics*. London: Routledge & Kegan Paul, pp. 133–151.
- Reichenbach, H. (1956). *The Direction of Time*. Berkeley and Los Angeles: Univ. of California Press.
- Russell, B. (1917). The Relation of Sense-data to Physics. In *Mysticism and Logic and Other Essays*. London: George Allen & Unwin, pp. 145–179.