

CAUSAL PLURALISM

STATHIS PSILLOS

*Department of Philosophy and History of Science
University of Athens, Greece
E-mail: psillos@phs.uoa.gr*

Most of the philosophical discussion about the metaphysics of causation has been dominated by what I call the ‘straightjacket’: the view that there is a single, unified and all-encompassing metaphysical story to be told as to what causation is. It has been presumed that the aim of philosophical inquiry is to tell this story. More specifically, it has been assumed that the aim of a philosophical theory of causation is to engage in conceptual analysis of the relation *c causes e*, where this analysis a) covers all and only cases in which intuitions determine that we correctly assert that *c causes e*; and b) is cast (preferably) in non-causal terms. This paper questions the plausibility and fruitfulness of the ‘straightjacket’ as a whole. It lays out a number of ways to deny the straightjacket, ranging from some mild ones to some genuinely pluralistic. It outlines and defends a version of causal pluralism according to which causation is very much like the *common cold*: a rather loose condition with no single underlying nature. What philosophers have taken to be the (competing) identifying characteristics of causation are, it is claimed, symptoms of causation. And though there is no unique nature of causation that these symptoms track, it can be traced reliably by its symptoms. Part of the argument for this causal pluralism will be what may be called Wittgensteinian pluralism, a view that can be traced back to G. E. M. Anscombe. The thrust of the argument is that explicit causal talk is dispensable, or almost dispensable, being useful for forming certain generalisations.

1. Introduction

Most of the philosophical discussion about the metaphysics of causation has been dominated by what I shall call the ‘straightjacket’: the view that there is a single, unified and all-encompassing metaphysical story to be told as to what causation is.^a It has been presumed that the aim of philosophical inquiry is to tell this story. More specifically, it has been assumed that the aim of a philosophical theory of causation is to engage in conceptual

^aIf this expression is taken to be too provocative, it may be replaced by the milder term ‘monism’.

analysis of the relation c causes e , where this analysis a) covers all and only cases in which intuitions determine that we correctly assert that c causes e ; and b) is cast (preferably) in non-causal terms.

There has been no shortage of such conceptual analyses and no shortage of counterexamples to all of them. The counterexamples exploit, at least partly, situations in which we are presumed to have clear intuitions about what causes what, but which intuitions are not being respected by the suggested philosophical analysis. The counterexamples typically lead to a battery of sophisticated attempts to revise or amend the philosophical analysis so that it is saved from refutation. These attempts, typically, either deny the intuitions on which the counterexamples are based or accommodate the problematic cases within the theory by adding further clauses to the original philosophical analysis. The result of all this is that where the original philosophical theory rested on a simple, forceful and intuitively plausible idea (e.g., that causation consists in a relation of counterfactual dependence between discrete events), the modified philosophical theory becomes very convoluted, somewhat ad hoc and implausible.

In this paper, my aim is not to review these theories. Anyone who has worked on the philosophy of causation is familiar with the problems they face.^b My aim is to question the plausibility and fruitfulness of the ‘straightjacket’ as a whole. I will lay out a number of ways to deny the straightjacket, ranging from some mild ones to some genuinely pluralistic. I will outline and defend a version of causal pluralism according to which causation is very much like the common cold: a rather loose condition with no single underlying nature. What philosophers have taken to be the (competing) identifying characteristics of causation, viz., regularity, counterfactual dependence, probability raising, presence of a process, presence of a mechanism, are, I claim, symptoms of causation. And though there is no unique nature of causation that these symptoms track, though that is, causation (like the common cold) can be many things, it can be traced reliably by its symptoms. Part of the argument for this causal pluralism will be what may be called Wittgensteinian pluralism, a view that can be traced back to G. E. M. Anscombe.² The thrust of the argument is that explicit causal talk is dispensable, or almost dispensable, being useful for forming certain generalisations. Causation comprises whatever conditions in fact exist in the general region corresponding to causal talk. But this talk is diverse and variegated and there is no conceptual or linguistic pressure to

^bCf. my *Causation and Explanation*.¹

have a theory of causation that ascribes to it a single and deep underlying nature.

2. The Straightjacket View

The ‘straightjacket view’ of causation is the view that there are facts of the matter as to the whole bunch of issues that relate to the nature of causation (that is, that causation has a determinate nature), that a philosophical theory of causation should aim to reveal these facts and that a good philosophical theory of causation should tell a unified and complete story that covers each and every aspect of the nature of causation. In particular, the straightjacket view assumes that there are facts of the matter as to:

- what kind of relation causation is.

Here, the debate is about whether causation is

- singular or general
- extrinsic to its relata or intrinsic to them
- irreducible to anything else or reducible to (or supervenient upon) non-causal facts.

If causation is taken to be irreducible, there is the further issue of whether it is a primitive relation or further analysable (though irreducibly so). If causation is taken to be further reducible, there is the issue of what the basis of reduction is:

- regularities
- necessary and/or sufficient conditions
- relations of counterfactual dependence
- probabilistic relations (e.g., contextual unanimity)
- the transference of some property, or trope, or ...
- the possession of a conserved quantity
- the fact that some event was the last change in the environment of another event just before the latter happened
- relations of invariance under interventions
- ... (fill in the dots with your preferred view, if it is none of the above).

Supposing that these issues have been settled, the straightjacket assumes that there are further facts as to

- what formal properties the causal relation has

- what the causal relata are.

Here, the debate is about whether the causal relata are events or facts or properties or tropes or... or a combination thereof.

Supposing that these issues have been settled, the straightjacket assumes that there are further facts as to

- the number (*adicity*) of the relata
- the proper form of causal statements

Finally, supposing that these issues have been settled, the straightjacket assumes that there are further facts as to

- what the direction of causation is.

There is no point in putting flesh on the above philosophical skeleton by detailing well-known philosophical theories of causation. But there is a point in reminding the reader that there is no theory of causation that is counterexample-free. Nor is there any theory of causation that tallies best with all our intuitions about what causes what. Nor are these intuitions always clear-cut and forceful. Hence, the current-state of play in the philosophy of causation is something like this: ingenious additions of epicycles to intuitively plausible theories and inconclusive (though suggestive) intuition-based arguments. This point could be brought home by examining any theory, but to illustrate it one could just look at the counterfactual theory of causation and the massive literature concerning the problems of pre-emption, early pre-emption, late pre-emption, trumped pre-emption, double prevention, etc.

The persistent failure to find an adequate philosophical theory of causation may well make us sceptical about the prospects of a theory that complies to the straightjacket: there is no single, unified and all-encompassing theory of causation. Perhaps, there is no metaphysical fact of the matter as to what causation is; no deep, simple and unified nature of causation. Obviously, this type of argument is not conclusive. But it is not meant to be so. Rather, it is meant to cast doubt on the plausibility and fruitfulness of the search for an adequate metaphysical account of causation. Why is there any hope that persistent failure will be, at some point, replaced by success?

This kind of scepticism is not of the sort associated (perhaps, mistakenly) with Hume. It is not based on qualms about the unobservability of the supposed necessary connections in nature, nor on claims about the

undefinability of causation. Nor is the kind of scepticism associated with Russell. It is not based on the claim that science does not search for causes, nor on the claim that the principle of causation is devoid of content. As will be explained later on, we can and do know a lot about what causes what. One need not be sceptical about causal knowledge, or indeed about causation, if one is sceptical about the possibility of a metaphysics of causation of the sort envisaged by the straightjacket. In the sequel, I will use this kind of scepticism as motivation for the development of causal pluralism. Then, in articulating and defending causal pluralism I will present an argument that aims to remove some of the rationale for trying to offer a metaphysical account of causation that reveals its deep and unique nature. Before we move on to this, let us in sections 3 and 4 examine some *prima facie* plausible ways to resist the straightjacket.

3. The Functional View

A popular view^c has been that causation should best be understood in a functional way. State a number of propositions that our folk theory of causation consists in (or a number of platitudes that the concept of causation should satisfy) and use them to fix the reference of CAUSATION: causation is whatever satisfies this folk theory (or the set of platitudes). This move can be worked out more formally in terms of Ramsey-sentences: causation is whatever satisfies the Ramsey-sentence of our folk theory of causation (or the set of platitudes).

Though compatible with the straightjacket, the functional view is a step forward since it is also compatible with a more *neutral* account of the metaphysics of causation. An advocate of the functional view does not have to accept that there is a *deep* metaphysical story to be told about the nature of causation. Nor does she have to accept that the folk theory of causation has something or other to say about *all* issues that the straightjacket says there is a fact of the matter. It is typical, however, of the advocates of the functional view to claim that there is a deep and unified metaphysical nature of causation that the functional approach identifies indirectly through the platitudes. Menzies, for instance, includes among the platitudes the claim that causation is a singular relation among events, thereby making it inevitable that only a certain metaphysical account of causation will be compatible with the functional approach.

^cDefended by Peter Menzies in Refs. 3– 4.

The problem with functionalism is that the folk theory of causation is not given in a forceful and intuitively compelling way. A lot of questions will be begged if the platitudes are chosen one way instead of another. If, for instance, the claim that causation is a singular relation among events is taken as part of the folk theory, no regularity (or generalist) account of causation can satisfy the folk theory. But who's to decide and how what the platitudes of causation are and in particular that the platitudes are such that some (but not others) metaphysical theories are excluded from being adequate for theories of causation?

This issue is far from trivial. Drawing a distinction between intuitions and platitudes of causation,^d we can assume that the folk theory of causation attributes some platitudinous features to causation. Here are four of them:

- ◇ The difference platitude: causes make a difference, viz., things would be different if the causes of some effects were absent.
- ◇ The recipe platitude: causes are recipes for producing or preventing their effects, viz., causes are the means to produce (or prevent) certain ends (effects).
- ◇ The explanation platitude: causes explain their effects, but not vice versa.
- ◇ The evidence platitude: causes are evidence for their effect, viz., knowing that c causes e , and knowing that c occurred, gives us (some) reason to expect that e will occur.

Arguably, each and every philosophical theory of causation should accommodate these platitudes, that is, show how each of them is brought out by whatever constitutes, according to the theory, the relation of cause and effect. But we can also assume that there are two firm pre-philosophical views about what causation is — what I have called ‘intuitions’ about causation.

- ◇ The regularity intuition: whether or not a sequence of two distinct events c and e is causal depends on whether or not events like c are regularly followed by events like e .
- ◇ The intrinsic-relation intuition: whether or not a sequence of two distinct events c and e is causal depends wholly on the events c and e and their own properties and relations, that is, it depends wholly

^dCf. Ref. 1, p. 6–7.

on the intrinsic and local features of the actual sequence of events.

The regularity intuition is mostly driven by folk epistemological considerations: how can causal relations be known or reliably manipulated unless they embody or instantiate regularities? This intuition is underpinned by the fact that we are unwilling to pronounce a sequence of events c and e causal unless there has been a regular association between events like c and events like e . If a causal relation were an one-off thing (this causing that here and now), causation would be of little usefulness and causal knowledge would require some kind of special non-inductive method. The intrinsic-relation intuition, on the other hand, is mostly driven by folk metaphysical considerations: causal relatedness is a matter of something in the cause bringing about the effect; it is a tie between cause and effect which is independent of things that happen at other places and other times. These intuitions are equally firm, I presume, but each of them is too controversial to be taken as a platitude of causation.

Although a functional account can (and should) accommodate the platitudes of causation, no functional account can accommodate both intuitions. To be sure, there can be compatibilist accounts of causation. That is, there can be accounts based on the intrinsic-relation intuition that can accommodate the thought that, as a matter of fact, causal relations give rise to stable regularities (that is, that the world is essentially nomological). But such accounts depend on putting a premium on the folk metaphysical intuition of intrinsic relation; a move that is certainly question-begging, since it presupposes that one of the two intuitions is really more central to our folk theory, while the other is derivative. An egalitarian view that gives, as it were, equal footing to both intuitions is excluded by a functional account.

4. The Two-Concept View

One way to deny the straightjacket view, as well as the functionalist presupposition that CAUSATION is a single and unitary concept, is to claim that there are more than one concepts of CAUSATION. The case for there being two concepts of CAUSATION has been made by Ned Hall.⁵ He distinguishes between causation as dependence and causation as production. Hall takes dependence to be simple counterfactual dependence, while he takes the concept of production (c produces e) as primitive.

This view is plausible. In fact, I have argued that there have been (historically and conceptually) two broad approaches to the metaphysical issue of causation.⁶ On the dependence approach, to say that c causes e

is to say that e suitably depends on c . On the production approach, to say that c causes e is to say that something in the cause produces (brings about) the effect or that there is something (e.g., a mechanism) that links the cause and the effect. There have been different ways to cash out the relation of dependence: nomological dependence (cause and effect fall under a law); counterfactual dependence (if the cause hadn't happened, the effect wouldn't have happened); probabilistic dependence (the cause raises the probability of the effect). Similarly, there have been different ways to cash out the concept of production, but the most prominent among them are cast in terms of something being transferred from the cause to the effect (e.g., a property, or some physical quantity — force, energy etc.). A key thought in the production approach is that cause and effect are connected by means of a local mechanism.

Why take seriously the two-concept view? One reason is that the two concepts we are discussing align quite naturally with distinct intuitions about causation: the production view aligns with the intrinsic-relation intuition, while the dependence approach aligns with the intuition that causation is an extrinsic relation between events (a species of which is a regularity). Another reason is that the two views set conceptually distinct constraints on causal relatedness. On the production view causal connectedness amounts to the presence of some tie between cause and effect, while no such tie is required by the dependence approach — just a robust dependence. Finally, dependence theories and production theories are extensionally distinct. There can be cases of causation licensed by dependence theories without being licensed by production theories and conversely. Most typical cases, however, concern situations based on no clear-cut intuitions, such as causal overdetermination and causation by disconnection.

If we take the two-concept view seriously, we have reason to be sceptical about the straightjacket view. If causation has a double nature, there is no single, unified and all-encompassing theory of causation to be had. Yet, ironically, the two-concept view ends up with two straightjackets. For each of the two concepts, there is supposed to be a metaphysical matter of fact as to what causation is and what features it has. For each of them, there is a single and unified story to be told. But now, each story is not comprehensive: it leaves out some facets of causation that the other account covers. Obviously, this kind of thought requires that we are clear on what the *facets* of causation are and that they are all facets of *causation*, even though they cannot be accommodated under a single and unitary concept. Our intuitions are still at play, but since these are not always forceful and

clear-cut, we may be left wondering whether some of the features that one or the other concept of causation accommodates really are worth accommodating.

In virtue of what is it the case that these two concepts are both concepts of CAUSATION? There may well be some interesting answer to this question. It may be argued that the two-concept view is genuinely egalitarian: since there is no way to privilege one set of intuitions concerning CAUSATION over the other, and since there is no other way to tie causation with either production or dependence, egalitarianism dictates that both approaches are equally acceptable accounts of causation. The difficulty with this answer is that though the two-concept view entails that there is a metaphysical fact of the matter as to what causation is (it is either production or dependence), there is no further way to tell when it is this rather than that. For we are not told when to apply the one concept and when the other. Here again, we have to rely on our intuitions.

If there were no other way to deny the straightjacket I would go for the two-concept view. But the difficulties of this view highlight, to me at least, the claim that once the straightjacket view is denied, there is little gain by replacing it with watered down versions of it, even if the replacements are in the right direction. This thought might lead to a more radical denial of the straightjacket view: causal pluralism.

5. Varieties of Pluralism

One way to be pluralist about causation has been explored recently by Christopher Hitchcock.⁷ He means to deny that there is a single thing that is the referent of the expression '*the* causal relation'. He motivates his pluralism by highlighting two distinct stages in causal analysis. In the first stage, some privileged class of entity is identified which pertains to causal relations, e.g., laws, relations of counterfactual dependence, probabilistic dependence, manipulability, causal processes. The thought here is that when it is the case that c causes e , some such entity is present. After this stage is completed (and in particular, after the privileged class of entity is discriminated from impostors — e.g., after genuine laws are distinguished from accidentally true generalisations, or genuine causal correlations are distinguished from spurious correlations), the second stage kicks in. This is an analysis of causation in terms of the privileged class of entity identified in stage one. For instance, causation consists in the ancestral of counterfactual dependence among events. Or causation consists in the exchange of a

conserved quantity etc. Here, the point is that among the privileged classes of entity identified in the first stage, one is selected as being constitutive of *the* causal relation.

Hitchcock's suggestion is that we should embrace causal pluralism at stage two, viz., there is no unique way to analyse the causal relation; there is no single thing such that *the* causal relation consists in. His point, then, is that there are assorted causal relations and none of them should be identified as the causal relation: causal analysis should consist in identifying some causal relation that is present in a particular case of causation.

Though suggestive, this idea needs further development. As it stands, it seems consistent with two opposing views: disjunctivism and the many-concept view.

The disjunctivist view: the causal relation is disjunctive. It is nomological dependence, or counterfactual dependence, or probabilistic dependence, or the presence of a causal process, or invariance-under-intervention, or . . .

This line admits a unique but multiply realised causal relation. A problem with the disjunctive view is that it is not clear how causation is identified in the first place. There must be some independent grasp of CAUSATION, which is then identified with a certain disjunction. But it is not clear what this independent grasp consists in. Perhaps, causation is identified in the functional way noted above. Then CAUSATION is a second-order concept whose realisers are given by the disjunction. If this line is taken, disjunctive causal pluralism becomes a species of the functional view and inherits its problems concerning what exactly should be included in the folk theory of causation. Another problem with this view is that it is not clear how the context can specify which disjunct is realised in particular cases. To say the least, it is typically the case that more than one disjuncts are realised in cases of causings. How, on this view, would one of the disjuncts be picked out as the one in which a particular causal relation consists in (it is realised by)?

The many-concept view: there are many concepts of CAUSATION each corresponding to a way of identifying the causal relation; none of them should be privileged in being *the* concept of CAUSATION.

This line is a development of the two-concept view. As with the case of the two-concept view, the many-concept view allows for a number of straightjackets. For in each and every case, there should be a fact of the matter as to what concept of CAUSATION applies. Yet, it is not clear any more why and in virtue of what all these concepts are concepts of CAUSATION.

Hitchcock's pluralism seems to go beyond both of these two views in denying, or at least aiming to deny, that there is anything deeper that unites the many and varied causal relations. Still, he wants to argue that all these are causal relations; hence, though there is something in virtue of which they are all *causal* relations, it is not clear what this is. It seems that it's left to our intuitions to play the role of classifying all these relations as causal. Different intuitions favour different relations as being causal and since there is no way in which we can privilege one set of intuitions over another we should be egalitarian, and hence pluralist, about causal relations.

5.1. *The symptoms of causation*

Here is another way to develop the pluralist line. In most typical cases, where causal talk has a bite, there are many ways to identify the presence of a causal connection — that is to ascertain that c causes e . This is because, generally, when c causes e , it will be the case that

- ◇ there is a law (deterministic or statistical) that links c and e
- ◇ if c hadn't happened, e wouldn't have happened
- ◇ $\text{prob}(e/c) > \text{prob}(e)$ in (all) relevant background contexts
- ◇ some causal process (mechanism) connects c and e
- ◇ something gets transferred from c to e .

This is obvious in many ordinary cases, e.g., when we say that the ball broke the window or that aspirin causes headache relief or that smoking causes lung cancer. But it is no less obvious when we turn to more 'scientific' cases, e.g., when we say that the tides are caused by the moon's attraction or that increases of unemployment rates cause a rise of the crime rates. In most typical cases, these entities (regularity, counterfactual dependence, probability raising, presence of a process, presence of a mechanism etc.) are correlated. This correlation explains why there are many ways to identify a causal fact and why there is, typically, agreement about what causes what, even if there is (philosophical) disagreement about what causation consists in.

Let us call the entities above 'symptoms' of causation. There are many and different symptoms of causation. It's not necessary that all of them are present in order to claim that a certain relation is causal. Nor is any of them privileged in identifying the presence of a causal relation. The case here is similar with diseases. Think of either measles or common

cold. In both cases, we have many symptoms of them. It is not necessary that all symptoms are present in order to claim that someone suffers from measles (or common cold). Sometimes, some typical symptom might be absent. On many occasions, it might be necessary to combine more than one symptom to assert that someone has measles (or common cold). Similarly, I'd claim, with causation and its symptoms. Yet, measles is a disease with a single underlying nature. It is a respiratory infection caused by the measles virus. With the common cold, however, things are more complicated (and more interesting). What we call 'common cold' is a rather loose condition with no single underlying nature. Several hundred cold-causing viruses have been found to cause the symptoms of the common cold (sneezing, sniffing, running/blocked nose, scratchy, sore, or phlegmy throat, coughing, headache, and a general feeling of unwellness). In light of this, there are two ways to develop causal pluralism, one along the lines of the measles analogy and the other along the lines of the common cold analogy.

Agnostic causal pluralism: there might be a deep and unique nature to causation — and hence a metaphysical fact of the matter as to what causation is — though there are many symptoms of it and many ways (none of which is privileged) to identify its presence.

Atheist causal pluralism: there is nothing single and deep that unites all the symptoms of causation and makes them track the unique nature of causation. Causation is a rather loose condition with no single underlying nature.

Agnostic causal pluralism is consistent with the claim that causation has a metaphysical nature, but differs from the straightjacket in that it does not search for this nature — if indeed there is. Rather, it counsels the use of the symptoms of causation in identifying causal facts — without bothering about whether there is anything deeper that these causal facts share in common. Atheist causal pluralism is more radical. On this view, we can meaningfully talk of causation, as we can meaningfully talk of common colds. We can identify cases of causation and discriminate them from cases of non-causation, as we can do the same with cases of common cold. But atheist causal pluralism denies that all cases of causation share something deep in common — a single and determinate common nature.

In what follows I will leave agnostic causal pluralism to one side and try to defend the more radical atheist position. The less radical agnostic view can always act as the pluralist's fallback position.

Part of the positive argument for atheist pluralism is the fact that the symptoms of causation, like the symptoms of common cold, are reliably

co-related in most ordinary cases. This allows us to group them together and claim that they track the same condition (causation; common cold) though, as it turns out, there is no single thing they track. The failures of the straightjacket are part of the negative argument for atheist pluralism. These failures suggest that no symptom of causation should be taken to be constitutive of it.

Against the two- or many-concept view, atheist pluralism claims that CAUSATION is a single concept. But one may wonder: what makes all these symptoms symptoms of *causation*? The atheist pluralist need not give a deep answer to this question. After all, what makes all the symptoms of common cold symptoms of *common cold*? As we have already noted, the atheist view is motivated by the fact that these symptoms are correlated in most typical cases in which we can ascertain causal facts. These correlations plus reasons of conceptual economy suggest that they are all symptoms of something, viz., causation. It turns out, however, that there is no single nature they track — this, if anything, is motivated by the failures we have had so far to identify it. Like the concept of common cold, the concept of causation has a history that goes back to antiquity. Tracing this history, thereby tracing the origins of the thought that causation (like common cold) is a single thing, would be a tall order.

It might seem that atheist pluralism is defeated by the fact that there are cases of genuine causation that do not involve some characteristic symptom of it, e.g., counterfactual dependence or regularity or the presence of a process or what have you. Cases such as these have produced well-known counterexamples to philosophical theories of causation. They show, presumably, that causation should not be identified with X, where X is some suitable entity. However, the attraction of atheist pluralism is that it does not identify causation with anything. The absence of a symptom is no reason to think that causation is not present, provided other symptoms are present. Besides, since there is some vagueness to CAUSATION, extreme and atypical cases do not turn the balance in favour of one symptom being more privileged than others. In most typical cases, the advantage of atheist pluralism is that there is no reason to choose one among the many symptoms of causation as being privileged (or constitutive of causation). Indeed, a problem faced by Hitchcock's pluralism, viz., that it cannot easily answer the question of how we choose among causal relations when more than one of them are present in a certain case, evaporates under the symptoms view.

Atheist pluralism, then, is not pluralism in the sense of 'two concepts or more'. It is not the view that there are two or more conditions called

‘causation’ and hence two or more deep metaphysical natures of causation. It is pluralism at the level of symptoms. At the level of the metaphysics of causation it is simply atheistic: there is no deep metaphysical nature of causation; there is no single thing that these symptoms track.

Atheist pluralism, no less than any other view on causation, presupposes that we know some causal truths. But where typical theories of causation rely on this knowledge to uncover the deep metaphysical nature of causation, atheist pluralism can take a more deflationary stance. Taking a cue from Arthur Fine’s *Natural Ontological Attitude*, we can claim that we can and do know a host of causal truths.⁸ That is, we do know what causes what. I will call the *Natural Causal Attitude* the stance that there is causal knowledge even if we do not know what causation is. According to it, in order to claim (correctly) that c causes e we don’t have to answer first some deep metaphysical questions about causation. In light of the above discussion, in order to know that c causes e we can appeal to some of the symptoms of causation and in particular to their correlation in most typical cases. The underlying idea is that causal truths are robust: they can be traced by means of regularities, relations of counterfactual dependence, relations of invariance under intervention, transference of energy-momentum etc.

The Natural Causal Attitude shifts the issue from the metaphysics to the epistemology of causation. It deflates the debate over the metaphysical nature of causation and stresses that we can get along with finding causal truths without being committed to any particular metaphysical view. But as it stands, it does not particularly favour atheistic pluralism, viz., the claim that there is no deep metaphysical nature of causation. To this end, the natural causal attitude needs to be supplemented by a relevant argument. The argument I will put forward suggests that there is no need to assume a deep and unique metaphysical nature of causation. It will lead us to what may be called Wittgensteinian pluralism.

6. Wittgensteinian Pluralism

It might be thought that the need for a deep metaphysical theory of causation, and in particular of a theory that reveals the single and determinate nature of causation, stems from the fact that there is a concept of CAUSATION and certain words in our languages (the verb ‘to cause’, the nouns ‘cause’ and ‘causation’) which aim to capture this concept. That is, it might be claimed that the very fact that there is explicit causal talk in our

language, that explicitly causal expressions are used, that certain generalisations are formed using explicitly causal language, is best explained by the admission of a condition in the world — causation — that answers to this talk. This line of thought, however, can be defeated. There is a sense in which explicit causal talk is dispensable, or almost dispensable. It is nonetheless useful for forming certain generalisations.

The relevant argument can be found in Anscombe's pregnant quotation:

“The truthful — though unhelpful — answer to the question: ‘How did we come by our primary knowledge of causality?’ is that in learning to speak we learned the linguistic representation and application of a host of causal concepts. Very many of them were represented by transitive and other verbs of action used in reporting what is observed. (...) 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*”.

(Ref. 2, p. 93)

Anscombe's focus in this particular quotation was on the issue of the observability of causation. Yet it seems clear that what she has in mind is the thought that there is a sense in which explicit causal language has only (or mostly) an *expressive* role to play vis-à-vis ordinary language: it expresses in a more abstract way facts that are already captured by cause-free vocabulary. Of course, for Anscombe this vocabulary is only explicitly cause-free. It is equipped to capture causal concepts such as those expressed by the verbs cited in the quotation. But her point, I take it, is that there is no need to invoke explicit causal expressions to capture and come to know causal truths.

Anscombe's claim rests on an idealisation. Ordinary language does already contain explicit causal expressions. But let us think in terms of a metaphor. Let us envisage a fragment of natural language NL_f which does not contain the word ‘cause’ (and relevant explicitly causal expressions).

Then we can think of an extension NL_{f+c} of this fragment which (explicitly) contains the word ‘cause’. What would the role of the word ‘cause’ be in NL_{f+c} vis-à-vis whatever can be expressed in NL_f ? Would substantial new truths not already captured in NL_f be expressed in NL_{f+c} ? The answer I want to explore is that explicit causal expressions play a useful role in forming certain generalisations, but do not add new content to whatever can already be expressed within NL_f . Suppose I say (truly) in NL_f that the hitting with the hammer *broke* the vase. If I were to make, in NL_{f+c} , the statement ‘The hitting with the hammer *caused* the vase to break’, I would not capture or prove any facts that were not already captured or proved by the cause-free formulation in NL_f . Similarly, for all causal verbs of NL_f .

Are there contexts in which the content of expressions that use explicit causal language cannot be captured by cause-free expressions in the way indicated above? Since I have no general answer to offer to this question, let us proceed by considering three important cases.

Case A: ‘ x causes ϕ ’, e.g., ‘Unemployment causes poverty’.

Suggestion: It can be captured by claims of the form ‘ x ϕ ’s’. For instance, ‘Unemployment impoverishes (people)’. This kind of move is impeded by the fact that we lack stock of relevant verbs. But though true, this is a stylistic point. There is no neat way to paraphrase the statement ‘Smoking causes lung-cancer’ as suggested above. But if we leave ugliness aside, we can always introduce new verbs. So, we can claim ‘Smoking lung-cancerises’. This would be on a par with perfectly acceptable causal statements such as ‘Smoking kills’ or ‘Aspirin relieves headaches’, which are true statements without explicitly causal expressions. The general idea, then, is that the expression ‘causes ϕ ’ can always be paraphrased by means of a concrete (new or already existing) concrete verb. Case A suggests that we should think of the verb ‘to cause’ as a placeholder for more specific verbs — sometimes they are available, in others they are not, and in yet other cases it is a twist of language to introduce them.

Case B: ‘the cause of e was ϕ ’, e.g., ‘The cause of Peter’ s death was hypothermia’.

Suggestion: Think of it intuitively. The content of the statement ‘The cause of Peter’ s death was hypothermia’ is fully captured by the statement ‘Peter died of hypothermia’. Statements such as ‘The cause of e was ϕ ’ can be analysed as definite descriptions: there is an x (caused- e x) and for all y (if caused- e y) then $x = y$ and ($\phi = x$). For instance, There is an x (caused Peter’ s death x) and for all y (if caused Peter’ s death y) then

$x = y$ and (hypothermia = x). The paraphrased sentence does not contain any constituent ‘the cause of’ for which we could substitute ‘hypothermia’. Yet, something true is asserted, viz., that Peter died of hypothermia. In the end, Case B suggests that we should think of the noun ‘cause’ as a placeholder for more specific events.

Case C: Consider the following claim:

‘John knocked the cup on the floor’. (i)

This, according to the above, implies

‘John caused the cup to be on the floor’. (ii)

The converse, of course, does not hold right away. John might have caused the cup to be on the floor in a different way. But statements such as (ii) do imply some kind of concrete statement like (i), e.g.,

‘John dropped the cup on the floor’. (iii)

So explicitly causal statements will always be made true by some concrete (implicitly) causal statement. The usefulness of (ii) consists in that it does not specify the way in which an effect was brought about and talks about it in an indefinite way. Case C highlights the fact that causal language is useful if we do not know how exactly an event was brought about, or if we do not want to be committed to any specific way.

If what said above is broadly right, then explicit causal talk is useful not because it enables us to talk about facts that we cannot, in principle, capture in another way. Its usefulness consists in the contingent fact that our languages are not rich enough to capture all causal truths by means of more specific verbs or expressions. This, however, is more of a practical difficulty. Besides, explicit causal talk is useful for forming generalisations (e.g., there are unknown causes of some phenomena or all events have causes), and for talking in an indefinite manner about the results of an action or an event-type.

The truth is that the verb ‘to cause’ is part of the language. Indeed, language has three general periphrastic causatives such as *cause*, *enable* and *prevent*. But note an interesting difference. Though we talk about enabling and prevention, there is no general theory of either enabling or prevention. Nor is there any presumption that enabling or prevention has a single and deep metaphysical nature; that all cases of enabling or prevention share something deep in common — a single and determinate common nature. Both ‘enable’ and ‘prevent’ play useful roles in language (and are

used to report truths) without requiring or implying a deep theory of what enabling or preventing is. Given the similarity of these three general verbs, there is no reason to treat ‘to cause’ differently. We may well treat it as an ordinary periphrastic causative, with no implication that something (some deep metaphysical nature or a deep and unique relation or activity) underwrites all correct applications of it. We do exactly this with other periphrastic causatives, such as ‘to enable’, or with concrete causatives, such as ‘to break’. Why should we not do it for ‘to cause’?

According, then, to this sort of Wittgensteinian pluralism, causation comprises whatever conditions in fact exist in the general region corresponding to causal talk — this talk being diverse and variegated. The linguistic expressions of causation are multiple and varied and do not require explicitly causal expressions, save for the need to form certain generalizations.^e There is no conceptual or linguistic pressure to have a theory of causation that ascribes to it a single and deep underlying nature. There is no pressure to assume that causal talk (which is mostly carried by implicitly causal verbs and other expressions) captures one and the same deep thing. Breaking, pushing, creating, capsizing, dissolving, decompressing, bonding and attracting (to name by a few) are all causings. But there is no pressure to think that they also share something deep in common, viz., that they are all causings in virtue of the (supposed) fact that they are all instances of a single and deep condition. To be sure, there can be (and there are) deep theories of each of these causings — e.g., in terms of molecular structure. If this is what we search when we look for a deep theory of causation, then well and good. But, apparently, it is not.

Activities such as the above are all causings without the need to assume that they share some hidden causal essence. And yet, it is also true that each of them gives rise to the symptoms of causation noted in the previous section. For instance, where there is breaking or pushing or ... there is transfer of energy, relations of counterfactual dependence, the presence of a law etc. This is reason enough to group them together as instances of cau-

^eLanguage already possesses a host of concrete (implicitly) causal expressions: causal conjunctions (e.g., because), prepositions (e.g., because of) and lexical causatives, that is verbs that encode causings (e.g., Peter *broke* the vase or Demetra *tore* the book). There are also the so-called periphrastic causative verbs (or auxiliary verbs such as *make*, *force*, *get* etc.). For a useful discussion of the linguistic representation of causation see Ref. 9. There are syntactic and semantic criteria to circumscribe the causal verbs. For instance, *to beg* is not a periphrastic causative, since I can meaningfully say: Mary begged John to leave, but he didn’t. But I cannot meaningfully say: Mary forced John to leave but he didn’t.

sation, even if, as stressed above, they do not necessarily track one and the same condition. That is how atheist causal pluralism and Wittgensteinian pluralism support each other.^f

7. Three Objections

Let us consider three objections to what said above.

First: Doesn't talking of *causal* truths imply that a) all these truths share something in common and b) that this something is a substantive property?

Reply: I agree with the first implication. As will be explained in the reply to the second objection, causal truths share some robustness which distinguishes them from another set of truths, viz., correlational truths. In this sense, they are substantive enough. What does not follow from this is that c) that there is a metaphysical fact of the matter as to what exactly causation is and that this fact underlies (and makes true) *all* causal truths. In other words, it does not follow that there is one single, unique, fully definite etc. truth-maker for all causal truths. But we need to be careful here. Atheist causal pluralism of the sort defended here is not causal anti-realism, *tout court*. It is anti-realist in so far as it denies that there is a single and deep truth-maker for all causal truths. But, in line with the natural causal attitude endorsed above, atheist causal pluralism admits that there are causal truths and that they can be (and are) known. Hence, it does not deny that causal truths have truth-makers. On the contrary, the thrust of atheist pluralism is that causal truths have a plurality of truth-makers that do not share anything deep in common.

Second: What do scientists search when they search for causes?

Reply: If answering this question presupposes an answer to the metaphysical question of what causation is (or a commitment to the view that causation has a deep and unique nature), clearly no progress has been made in finding out causes (since we don't know what causation is). But since there has been such progress, the argument is a *reductio* of the basic presupposition. In any case, the simple answer to the foregoing question is that scientists search for causal truths. What are these? They are a) truths and b) more robust than mere co-relations. (Recall that the statement of causal truths does not presuppose explicitly causal language. Consider: drug X

^fDue to lack of space I will not discuss here Nancy Cartwright's relevant views. See her *The Dappled World*.¹⁰

cures disease Y; drug Z *relieves* from pain; policy Q *reduces* urban crime etc; teaching method K *improves* the exam results of children; the earth *attracts* the moon.)

But why do we need (b)? I think this is a broadly empirical issue. We are interested in a special kind of truth (causal truth) because we are not interested solely in predictions. Co-relations (that is, truths that are not causal, but may appear to be causal) serve well in prediction. But they don't serve well in explanation; they break down upon manipulation; and they do not lead to effective strategies. So, we are interested in robust causal truths. No matter what the metaphysics of causation is (and no matter whether there is a fact-of-the-matter about this metaphysics), causal truths (should) satisfy a set of platitudes noted in section 3. Mere correlations do not satisfy these platitudes.

Third: Famously, Anscombe took it that her argument from causal verbs implies that causings (and hence causation) are observable.

Reply: Though I used to think differently, now Anscombe's point seems to me compelling. Suppose one says that a particular tree branch *bent* after having had pressure exerted on it. Then, by the very use of the verb 'to bend', one makes a *causal* claim. If this claim is true, since one had directly perceived the bending of the tree branch, one has thereby directly perceived the tree being caused to bend. So one has directly observed the causing. Anscombe's mistake, if I am allowed to talk in such terms, was the conclusion she wanted to draw from the observability of causings, viz., that some kind of non-Humean, that is singularist, (metaphysical) account of causation is true. This does not follow. The observability of causings is consistent with a number of metaphysical views about causation.⁸ It is also consistent with the kind of pluralism defended above.

Acknowledgments

Earlier versions of this paper were presented in seminars at: the *Centre for Logic and Philosophy of Science*, University of Ghent; the University of Lund; Bogacisi University; the University of Cyprus; at: the Symposium *Causal Inference and Probability in the Physical Sciences* in the meeting of the *Spanish Society for Logic, Methodology and Philosophy of Science*, Valladolid, Spain, November 2004; the fifth workshop of the *Metaphysics in Science Group*, University of Ghent, May 2005; and as an invited talk at

⁸This point has been ably defended by Helen Beebe.¹¹

the *Fifth European Congress for Analytic Philosophy*, University of Lisbon, August 2005. I have benefited from comments by a number of people. They are too many to be thanked all personally, but here is a partial list: Helen Beebe, Alexander Bird, Phil Dowe, Pascal Engel, Bengt Hansson, Carl Hoefer, Gurel Irzik, Antonis Kakas, Igal Kwart, Joke Meheus, Anna-Sofia Maurin, Demetris Portides, Rebecca Schweder, Mauricio Suarez, Erik Weber and Stephen Voss. Research for this paper was funded by the framework EPEAEK II in the programme Pythagoras II.

References

1. S. Psillos, *Causation and Explanation*. Acumen & McGill-Queens University Press (2002).
2. G. E. M. Anscombe, *Causality and Determinism*. Cambridge University Press (1971). Extracts reprinted as ‘Causality and Determination’. In: E. Sosa and M. Tooley (Eds.), *Causation*. Oxford: Oxford University Press, 88–104 (1993).
3. P. Menzies, Probabilistic Causation and the Pre-emption Problem. *Mind* **105**, 85–117 (1996).
4. P. Menzies, Intrinsic Versus Extrinsic Conceptions of Causation. In: H. Sankey (Ed.), *Causation and Laws of Nature*. Dordrecht: Kluwer, 313–29 (1999).
5. N. Hall, Two Concepts of Causation. In J. Collins, L. Paul and N. Hall (Eds.), *Counterfactuals and Causation*. MIT Press, 225–76 (2004).
6. S. Psillos, A Glimpse of the Secret Connexion: Harmonising Mechanisms with Counterfactuals. *Perspectives on Science* **12**, 288–319 (2004).
7. C. Hitchcock, Of Humean Bondage. *The British Journal for the Philosophy of Science* **54**, 1–25 (2003).
8. A. Fine, *The Shaky Game: Einstein, Realism and the Quantum Theory*. Chicago: University of Chicago Press (1986).
9. P. Wolff and G. Song, Models of Causation and the Semantics of Causal Verbs. *Cognitive Psychology* **47**, 276–332 (2003).
10. N. Cartwright, *The Dappled World*. Cambridge: Cambridge University Press (1999).
11. H. Beebe, Seeing Causing. *Proceedings of the Aristotelian Society* **103**, 257–80 (2003).