Causes and Laws

[Link to publication record in Manchester Research Explorer](#)

Citation for published version (APA):

Beebee, H. (2015). Causes and Laws: Philosophical Aspects. In J. D. Wright (Ed.), *International Encyclopedia of the Social and Behavioral Sciences* (2nd ed., Vol. 3, pp. 266-273). Elsevier BV.

Published in:

International Encyclopedia of the Social and Behavioral Sciences

Citing this paper

Please note that where the full-text provided on Manchester Research Explorer is the Author Accepted Manuscript or Proof version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version.

General rights

Copyright and moral rights for the publications made accessible in the Research Explorer are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Takedown policy

If you believe that this document breaches copyright please refer to the University of Manchester's Takedown Procedures [<http://man.ac.uk/04Y6Bo>] or contact uml.scholarlycommunications@manchester.ac.uk providing relevant details, so we can investigate your claim.

Causes and Laws: Philosophical Aspects

Helen Beebee

Pre-published version: do not cite. Published version: ‘Causes and Laws: Philosophical Aspects’, in J. D. Wright (editor-in-chief), *International Encyclopedia of the Social and Behavioral Sciences* (2nd edn., vol. 3), 266–73 (Oxford: Elsevier)

Abstract

The questions ‘what makes it the case that one event causes another?’ and ‘what makes it the case that something is a law of nature?’ are highly controversial ones for which, amongst contemporary philosophers, there is no answer that can claim orthodoxy. This article sketches some of the most influential theories of causation and lawhood.

1 Preliminaries

‘Cause’ and ‘law’ are perhaps two of the most important and fundamental concepts that human beings deploy in their attempts to understand and intervene in their environment, both in everyday life and in their scientific endeavors. When we want to explain why an event occurred, we seek out its causes. When we act, we do so because we believe the action will have certain effects. When we want to know why our environment and our fellow human beings behave in regular, predictable ways, we look for the laws that govern that behavior.

It is part of the scientist’s job to discover what causes what, and to investigate what the laws of nature are. The philosopher’s job, however, is a more abstract one: the philosopher wants to know what makes it the case that something—anything, be it the decay of a subatomic particle or the assassination of Archduke Ferdinand—causes something else, or that something is a law of nature—whether it is a physical, biological, psychological, economic or any other kind of law.

We know that when one strikes a match and it lights, the first event—the striking—caused the second. But this causal fact depends for its truth on more than the mere fact that the two events occurred one after the other: there must be some extra feature of the world that binds these two events together as cause to effect. Similarly, it seems that laws of nature must be more than mere

regularities: it is doubtless true that all lumps of gold are smaller than a mile in diameter, but it is not a law that this is so. So genuine laws must have some extra feature, apart from regularity, which marks them off from mere ‘accidental’ regularities. The philosopher’s job, or at least part of it, is to say what, if anything, this ‘extra feature’ might be.

1.1 Determinism vs. Indeterminism

A question that is relevant to any discussion of the nature of causes and laws is that of whether the universe is fundamentally deterministic or indeterministic. One way of putting the thesis of determinism is this: the complete state of the universe at any given time together with the laws of nature determines precisely what the complete state of the universe will be at all future times. Indeterminism is simply the denial of determinism.

The view that the universe is fundamentally indeterministic has been widely (though by no means universally) held within science since the advent of quantum mechanics; however, most philosophers only began to take the possibility of indeterminism seriously in the 1960s and 1970s. One consequence of this reluctance to take indeterminism on board is that some current theories of causes and laws were originally formulated under the assumption of determinism, and were then modified to take into account the possibility that at least some laws and causal processes are fundamentally indeterministic. The theories of laws and causation discussed below are mostly presented in their simpler, deterministic form. While the indeterministic versions of the theories present additional complications and problems—for example, because they employ the notion of probability—the fundamental philosophical issues addressed, and problems faced, by the indeterministic and deterministic versions are mostly the same or similar (see *Probability and Chance: Philosophical Aspects*).

1.2 The particular and the general: events, facts and properties

Many of our causal claims are claims about individual occurrences: the assassination of Archduke Ferdinand, the Great Recession, the birth of a particular child, a specific meteor shower, or whatever. Such occurrences are thus the ‘relata’ of particular or ‘token’ causal relations. There is an ongoing debate about whether such individual occurrences should be conceived as events (the assassination of Archduke Ferdinand) or facts (the fact that Archduke

Ferdinand was assassinated) – see for example Mellor 1995 and Lewis 1986 – but for simplicity we shall assume that they are events.

Sometimes, however, we make general causal claims about whole populations rather than their individual members. Thus it is doubtless true that smoking causes cancer, but it does not follow that every individual who smokes will be caused to get cancer as a result; and it can be true that a particular person dies as the result of a bee sting without its being true that bee stings in general cause death. General or ‘population-level’ causation is generally taken to apply to properties or repeatable features that different individuals can share, such as smoking, getting cancer, being red, or weighing 3kg.

Laws of nature are likewise concerned with general relationships between properties – whether it is the relationship between force, mass and acceleration enshrined in $f=ma$ or the economic ‘law’ of diminishing marginal utility. In what follows, general features or properties will be denoted ‘ F ’, ‘ G ’, and so on, while individual events will be denoted ‘ c ’, ‘ e ’, etc.

2 Laws of Nature

2.1 Laws vs. accidents

What is the difference between laws of nature and merely ‘accidental’ regularities? (Suppose that it is a law that all metals expand when heated and that it is merely accidentally true that all (past, present, and future) Queens of England are less than two meters tall.) A crucial difference is that laws ‘support counterfactuals,’ whereas accidental regularities do not. It is true of any unheated piece of metal that if it were to be heated it would expand. But it need not be true of any non-Queen of England that if she were to be Queen she would be less than 2 meters tall. If someone is currently, say, 2.05 meters tall, becoming Queen of England would not result in her getting any shorter.

Another difference is that violations of law are said to be physically impossible, whereas violations of accidental regularities are physically possible. It is physically impossible for a heated metal to fail to expand when heated; but it is perfectly possible for a Queen of England to be taller than 2 meters. Finally, it is sometimes claimed as a third difference that laws, unlike accidental regularities, ‘govern’ what goes on in the universe. We tend to suppose that statements of laws of nature are rather like pieces of divine legislation: decrees that the universe must obey rather than mere general descriptions of what in fact happens.

2.2 Realist vs. regularity accounts of laws

Can an analysis of lawhood be given that does justice to these intuitions about the nature of laws? A popular view (Armstrong 1983) is that laws are relations of necessitation between universals. To say that it is a law that all F s are G s is to say that there is a necessary relation, N , that holds between the universals (properties) F and G . (Following Armstrong, we can write this ' $N(F,G)$.)' When an object instantiates F , the instantiation of F guarantees, via N , that G will also be instantiated. Thus, $N(F,G)$ —its being a law that all F s are G s—guarantees that all F s will in fact be G s.

The basic point of this 'realist' theory of laws is that it provides an 'ontological ground' for the difference between laws and accidents: accidental regularities just happen, whereas lawful regularities are grounded in and explained by a necessary relation holding between the co-occurring properties. And, on this view, sense can be made of the idea that laws govern what happens: if $N(F,G)$ holds, then G must be instantiated whenever F is instantiated.

The main rival to this conception of laws is the 'Humean' or 'regularity' conception, according to which laws are really no more than regularities. The classic (though not the only) motivation for the regularity view is an 'empiricist' epistemology according to which we ought not to believe in entities that we cannot perceive. Since our evidence that there are any laws of nature comes solely from the observation of regularities—when we see a law (as opposed to an accidental regularity) being instantiated we do not see any extra ontological feature—we ought not to believe that there is any extra ontological feature which laws have but accidental regularities lack.

A 'naïve' regularity theory, however—according to which statements of law are merely true universal generalizations—is untenable, for it simply collapses the distinction between laws and accidental regularities. If statements of law are merely true universal generalizations, then accidental regularities are themselves laws: it turns out to be a law of nature that no Queen of England is taller than 2 meters after all.

A more sophisticated version of the regularity theory is provided by the ‘Ramsey–Lewis view’, after F. P. Ramsey and David Lewis (see Lewis 1973b, pp. 73–5; J. S. Mill also held a similar view). The basic idea of the Ramsey–Lewis view is as follows. When we seek, as scientists do, to find out how the universe behaves, we are not satisfied with isolated facts about what happened in the laboratory on Tuesday afternoon, what happened when a particular patient took a particular drug, and so on. Rather, what we seek are wide-ranging generalizations that abstract as far as possible from the particularities of the laboratory or the patient. Generalizations about the heights of Queens of England or the diameters of lumps of gold apply to an extremely restricted number and range of phenomena. On the other hand, generalizations about subatomic particles, chemical reactions, the relationship between force, mass and acceleration, and so on, apply to an enormous number of diverse phenomena, and are hence much more useful and explanatorily powerful. According to the Ramsey–Lewis view, the difference between laws of nature and merely ‘accidental’ regularities amounts to roughly this difference between generalizations that are wide-ranging and powerful and those that are not. It is not that there is any ontological distinction between laws and accidents—laws do not have an extra feature (*N*, say) that accidental regularities lack; rather, it just so happens that some generalizations (the laws) are interesting and useful ways of codifying what happens in the universe, and others (the accidents) are not.

The Ramsey–Lewis view (and the Humean view of lawhood in general) diverges from a conception of laws of nature according to which laws ‘govern’ what happens in the Universe, since on the Humean view the laws just are regularities, rather than features of the world that explain why the world is regular (Beebe 2000). Some philosophers argue that Humeanism about laws makes it impossible for us to know any laws of nature (Armstrong 1983, 52-9; Bird 2008; Beebe 2011). On the other hand, some philosophers claim that Humeanism about laws takes the sting out of the alleged conflict between determinism and free will (see Swartz 1985, chaps. 10 and 11, Beebe & Mele 2001, Berofsky 2012) (see also *Free Will and Action*).

2.3 *Ceteris paribus* laws

The accounts of laws canvassed above are perhaps best suited to the laws we find in, say, physics and chemistry – the places where we are most likely to find truly exceptionless generalizations that hold across the whole of space and time (though it is often argued that even in the ‘hard’ sciences, truly exceptionless generalizations are extremely rare). There are plenty of generalisations in, say, economics, sociology and psychology that, while not exceptionless, are clearly lawlike in some sense. The law of diminishing marginal utility, for example, does not hold in all cases, so if lawhood requires an exceptionless regularity, there will turn out to be no laws in economics. On the other hand, it is surely no accident that, by and large, the utility of something eventually decreases as one obtains a greater quantity of it. Similarly, there are plenty of generalisations in biology and psychology that are not exceptionless but nonetheless provide a good foundation for prediction and control.

Such cases are generally thought of as ‘*ceteris paribus*’ laws (‘cp-laws’ for short): generalisations that hold not in all cases *simpliciter*, but only ‘other things being equal’, that is, in cases where there is no additional, interfering factor. We might try to subsume cp-laws under our existing accounts of laws simply by building in the absence of such interfering factors, so that our candidate law has the form ‘all *F*s are *G*s, except in circumstances *X*, *Y* and *Z*’. This approach is problematic, however, since in many, if not most, cases, we either have no idea what the relevant *X*, *Y* and *Z* are, and so – even if there is a law in the offing – we will not be in a position to formulate it, or else there will be indefinitely many ways in which the generalization in question could in principle fail, so there is no way in principle of specifying what the relevant law is.

We might try to avoid these problems by claiming that cp-laws are only supposed to hold in ‘ideal’ or ‘normal’ conditions, thus removing the need to specify exactly what those conditions are within the content of the law itself. However, unless we can independently specify which conditions count as normal or ideal, the threat of circularity looms: we cannot allow ‘normal’ or ‘ideal’ conditions to simply be those in which the regularity obtains, or cp-laws will turn out to be trivially true (all *F*s are *G*s – except when they’re not). (See Hausman 1992, Ch.8; Pietroski & Rey 1995; Earman & Roberts 1999; Schrenk 2007a.)

2.4 Stability and invariance

Some contemporary accounts of lawhood take a different starting point to the idea that a law is an exceptionless regularity of a certain kind (the Ramsey-Lewis view) or else is the obtaining of some metaphysical relation that entails the existence of an exceptionless regularity (Armstrong's view). Instead, they start from the idea that the distinctive feature of laws is that they are stable or invariant under a range of different background conditions. Suppose that the generalization 'all the pens currently on Obama's desk are blue' is true. Despite being true, it is highly unstable: there are very many possible situations in which it would be false. Obama might, for example, have left the red pen he was using earlier today on his desk rather than putting it in his pocket, or one of his aides might have left their own black pen on the desk. By contrast, 'nothing travels at the speed of light' is highly stable: there are (so we currently believe) no physically possible situations at all that would render it false. The notion of stability is thus a 'modal' notion: the stability or otherwise of a generalization is a matter of whether it would still be true in various different, non-actual situations. And, since stability is not all-or-nothing but rather comes in degrees – one generalization can be more or less stable than another – we can think of lawhood itself as coming in degrees, so that some generalizations are more lawlike than others, rather than thinking that a generalization either expresses a law or it does not.

The notion of stability (or invariance) may be a good place to start if we want to make sense of cp-laws, because it seems to capture what is 'lawlike' about, say, the law of diminishing marginal utility. While the generalization in question only holds *ceteris paribus*, it arguably still exhibits a good deal more stability than, say, the generalization concerning the pens on Obama's desk: it would remain true (*ceteris paribus*) given a wide range of changes in prices, consumer preferences, and so on. Of course, we still need to find a way to spell out the restrictions imposed by the *ceteris paribus* clause; it has been suggested that we think of cp-laws as generalisations that are stable across circumstances that are of particular interest to a given science, such as psychology or economics (see Lange 2002; Woodward & Hitchcock 2003a, b).

2.5 Laws and dispositions

A final kind of account of the nature of laws – again, one that is perhaps well suited to dealing with cp-laws – is a 'dispositionalist' account. According to a naive version of the so-called 'counterfactual' account of dispositions (such as fragility, which, when applied to panes of glass and china cups at least, we can think of as the disposition to break when dropped), an object *o* (a

china cup, say) has disposition to manifest feature *M* (being broken) in ‘stimulus conditions’ *S* (being dropped) just if, if *o* were to be in situation *S*, it would manifest *M*. As it stands, however, this account cannot be right. For example, if I were to drop a china cup that is wrapped in a lot of cotton wool, it would not break – but the cup is nonetheless fragile. The presence of the cotton wool ‘masks’ the disposition – it prevents the disposition from being manifested. Antidotes provide another counterexample to the naïve counterfactual account: the redback spider is deadly (its bite has the disposition to kill people) despite the fact that, thanks to the existence of an antivenom, nobody has died from a redback bite for many years.

Whether the counterfactual account can be amended to deal with problem cases such as these is a controversial question (see e.g. Lewis 1997, Bird 1998). What is relevant to the discussion of laws, however, is that we can think of laws of nature as describing not how things actually behave, but how they are disposed to behave. So a candidate law of nature might have the form ‘all *F*s are *D*s’, where property *D* is the disposition or capacity to be *G*, rather than ‘all *F*s are *G*s’ (Cartwright 1989, Ellis 2002, Bird 2005).

A dispositionalist account of laws may be well suited to deal with cp-laws because – as we have just seen – something’s having a disposition is compatible with its failing to manifest that disposition even though the stimulus condition is present. If we think of laws as generalisations about dispositions, then we can think of cp-laws as those laws that describe dispositions that are susceptible to masks, antidotes and the like. Thus, for example, ‘all china cups break when dropped’ is false: I can easily falsify it by dropping a china cup wrapped in cotton wool. But this does not falsify ‘all china cups are disposed to break when dropped’, since, as we have seen, a china cup is still disposed to break when dropped – it is still fragile – even if, thanks to the cotton wool, it is not in a position to manifest that disposition when I drop it. (See Schrenk 2007b.)

3 Causation

Theories of causation have proliferated in the last fifty years or so, with different theories speaking to different underlying metaphysical commitments and to different views about how the ontology of causation (that is, how the world has to be arranged in order for causal claims to be true) ought to related to its epistemology (that is, how we can come to know, or justifiably

believe, that a causal claim is true). This section surveys some of the leading theoretical approaches.

3.1 The naïve regularity theory

The agenda for much of the contemporary discussion about the nature of causation was set by Hume (1748/1975, Sects. IV–VII). Hume was an empiricist, believing that there could be no ‘ideas’ in the mind that do not somehow come from our senses. How, then, do we come to have an idea of causation: what features of the world could furnish us with that idea? Famously, Hume maintained that we do not perceive any intrinsic connection or relation between two events we judge to be causally related: we see the match being struck, and we see it light, but we do not see any causation between those two events. Causation, then, cannot be a mysterious intrinsic relation between events—for if it were, since we cannot perceive any such relation, we could have no idea of causation. (See Menzies 1998 and Beebe 2009 for discussion of the question of the observability of causation.)

The exact nature of Hume’s positive theory of causation is a matter for interpretative dispute (see Beebe 2006a). However, the standard interpretation – and therefore the one that has had the most impact on contemporary discussion – characterizes Hume as a ‘naïve regularity theorist’. On this view, two events are related causally in virtue of contiguity (causes and their immediate effects are right next to each other in space and time), temporal priority (causes precede their effects), and constant conjunction (events similar to the cause are always followed by events similar to the effect). So what makes it true that the striking caused the lighting is that the two events are contiguous (or, if there is some spatio-temporal gap between them, the events are mediated by further events so that there is a chain of contiguous events starting with the striking and ending with the lighting); the striking occurs before the lighting; and similar strikings are always followed by similar lightings. The constant conjunction requirement ensures that for Hume (and for most modern day Humeans) causation is an extrinsic relation: a causal relation obtains between two events not purely in virtue of how those events themselves are, but in virtue of features of other events: the striking only gets to count as a cause of the lighting because other strikings are also followed by lightings.

The other notable feature of the naïve regularity theory is that it is a reductionist view: causal claims, like ‘the striking of the match caused the lighting,’ are made true, at bottom, not by any primitive causal feature of the world, but by non-causal features (since the fact that two events are contiguous, the fact that one occurs before the other, and the fact that they are similar to other events are not themselves causal facts).

The naïve regularity theory is now universally regarded as untenable, although, as we shall see, elements of it are clearly discernable in many contemporary theories. It is untenable for a variety of reasons. For example, the contiguity requirement rules action at a distance (one event’s causing another, spatially or temporally distant event without there being a chain of events ‘hooking up’ the first to the last) to be impossible, while many philosophers think that is at least conceptually possible. And the constant conjunction requirement is obviously flawed: there can be constant conjunction without causation—for example, joint effects of a common cause are constantly conjoined but neither causes the other—and, if determinism is false, it seems there can be causation without constant conjunction. My betting on number 13 on the roulette wheel is a cause of my winning some money, even if in exactly the same circumstances the following day I bet on number 13 and lose.

Nevertheless, the general issues that Hume raised are still among the main foci of dispute in contemporary philosophy of causation. In particular, the question of whether causal facts reduce, somehow or other, to facts about regularities is the driving force behind many of the theories of causation currently on offer.

3.2 Counterfactual Analyses

Arguably the most influential analysis of causation in recent years has been Lewis’s counterfactual analysis (Lewis 1973a). According to Lewis’s analysis, an event e causally depends on event c if e counterfactually depends on c —which is to say, if, had c not occurred, e would not have occurred either (see *Counterfactual Reasoning, Quantitative: Philosophical Aspects*; *Counterfactual Reasoning, Qualitative: Philosophical Aspects*). A ‘chain of causal dependence’ is a series of events $\langle a, b, c, \dots, n \rangle$ such that b causally depends on a , c causally

depends on *b*, and so on. Finally *c* causes *e* if there is a chain of causal dependence (perhaps involving only *c* and *e* themselves, but perhaps involving many hundreds of intermediate events) from *c* to *e*.

The need for chains of dependence arises when there is no counterfactual dependence between cause and effect because of a backup mechanism or ‘preempted alternative.’ Suppose that the bus stop is next to the taxi rank. Susan takes the bus to work (*c*); but if the bus had been late she would have taken a taxi. She attends a 9.30 meeting (*e*). *e* does not counterfactually depend on *c*, since if Susan had missed the bus she would have taken a taxi and attended the meeting just the same. But there is some intermediate event *d*—Susan’s getting off the bus at the stop outside her office, say—which counterfactually depends on *c* (if she had not got on the bus in the first place, she would not have got off it either) and upon which *e* counterfactually depends (if she had not got off at that stop—given that she was already on the bus, and it was already nearly 9.30—she would have had to walk back to her office from the next stop, and she would have missed the meeting).

Lewis’s counterfactual analysis of causation has been found to be subject to a number of counter-examples (e.g. Schaffer 2000a). Much recent work on causation has been devoted to devising alternative analyses which, while remaining true to Lewis’s basic claim that causation is to be analysed in terms of counterfactuals, add further complexities in an attempt both to circumvent the problem cases and to deal with indeterministic causation (see, e.g., Menzies 1989, Noordhof 1999, Lewis 2000).

Both Lewis’s original analysis and its more complex successors fall squarely within the Humean tradition, since the counterfactual analysis is a reductionist project: it seeks to show how causal facts are made true by noncausal features of the world. Indeed, since Lewis’s account of counterfactuals makes the truth conditions of counterfactuals depend primarily on the laws of nature, and the laws of nature themselves, on Lewis’s view, are merely a subset of the regularities, one can think of counterfactual analyses of causation generally as more sophisticated descendants of the naïve regularity theory of causation.

Not everyone, however, is a Humean. Many philosophers deny Hume’s central motivating claim that causation cannot be perceived (see Anscombe 1971, Armstrong 1997,

Chap. 14, Menzies 1998, Beebe 2009). Some argue that the features of the world to which reductionist analyses seek to reduce causal facts are themselves really causal features; hence reductionist analyses are circular and therefore unsuccessful (see Carroll 1994). Others argue that if a broadly Humean analysis of causation were true it would render the pervasive regularities that the universe exhibits a sort of cosmic fluke (see Strawson 1989, Chap. 5; Beebe 2006b). And many simply think that the persistent failure of reductionist accounts to deal with counterexamples, despite ingenious amendments, is good reason to think that the whole reductionist project is doomed to failure.

3.3 Process and mechanistic theories

We generally take it for granted that causal influence is propagated through mechanisms or processes of various kinds. When the impact of the white snooker ball causes the black ball to roll into the pocket, we can grant that the latter event counterfactually depends on the former, but may yet be inclined to say that this dependence is symptomatic rather than constitutive of the causal relation between the two. This is because there is an identifiable process linking the two, involving the transfer of energy-momentum from the white to the black ball and subsequent movement of the black ball towards the pocket. Some philosophers have attempted to analyse causation by focusing on the distinction between a genuine causal process and a mere ‘pseudo-process’. Salmon (1984), for example, takes the distinguishing feature of causal processes to be their ability to ‘transmit a mark’. A mouse walking across the floor is a process that can transmit a mark; for example, if I paint a blue patch on the mouse’s fur, the blue patch will ‘transmit’ from earlier to later stages of the process (since it will go wherever the mouse goes). By contrast, if I make a spot of light move across the floor by means of a torch, the movement of the light across the floor is a pseudo-process: if I paint the first illuminated patch blue, that mark will not transmit to later stages of the ‘process’; it will stay just where I put it.

Salmon’s ‘mark transmission’ theory is problematic for various reasons (see Dowe 2008, §6). Dowe’s suggested replacement, the ‘Conserved Quantity Theory’ (2000), identifies the causal relation with the transfer of whatever physical quantities are conserved according to the conservation laws of physics, e.g. energy-momentum. A problem with this account is that it only tells us what causation actually consists in. It seems that we can imagine possible situations where there is causation in the absence of any transfer of a conserved quantity (e.g. by the

casting of spells). Other analyses of causation – such as the counterfactual analysis – tell us why both the movement of snooker balls does, and the effective casting of spells would, count as cases of causation, since they both involve counterfactual dependence. Dowe's account fails to do this.

One solution to the problem is to attempt to spell out the 'folk theory' of causation: those features that, according to our ordinary concept of causation, are essential for a relationship to be causal, and hence those features that will apply to snooker balls and spells alike. And we can then claim that the transfer of conserved quantities counts as causation precisely because, as things actually are, it is such transfer that satisfies our folk theory or definition of causation – while at other possible worlds, such as one in which magic spells are effective, other kinds of relation count as causal (see Dowe 2000, Chap.1; Menzies 1996; Lewis 2004).

Like process theories, 'mechanistic' accounts of causation also seek to identify causation with the process by which the cause produces the effect (Glennan 1996; Machamer, Darden and Craver 2000; see Williamson 2011 for a survey). However, mechanistic accounts are more pluralist in spirit, and do not try to locate a single feature of the world (such as the transfer of a conserved quantity) that can be identified with the causal relation. Machamer et al. note that 'terms like "cause" and "interact" are abstract terms that need to be specified with a type of activity and are often so specified in typical scientific discourse. Anscombe [1971] ... noted that the word "cause" itself is highly general and only becomes meaningful when filled out by other, more specific, causal verbs, e.g., scrape, push, dry, eat, burn, knock over. An entity acts as a cause when it engages in a productive activity' (Machamer et al. 2000, 6; see also Cartwright 2004).

Dowe's account is explicitly reductionist in spirit, in two senses. First, it attempts to identify causation with a feature of the world that can itself be understood in non-causal terms: the 'transfer' of conserved quantities is not supposed to be conceived as, for example, the presence of a energy-momentum in the white ball causing its subsequent presence in the black ball. It is also reductionist in the sense that all (actual) causation is, at bottom, causation at the level of fundamental physics: the presence of a causal process is determined by features of the world that are visible, as it were, only from the perspective of physics. Machamer, Darden and Craver's account is anti-reductionist in both senses. 'Productive' activities are not themselves to be analysed in non-causal terms (the instantiation of regularities or the transfer of conserved

quantities); nor are the mechanisms that underpin causal relations to be found only at the level of fundamental physics. Indeed, part of the purpose of their view is to make sense of causal explanation in biology.

3.4 Probabilistic theories

Frequently it is claimed on cigarette packets that smoking causes heart disease; this is a claim about ‘general’ or ‘type-level’ or ‘population-level’ causation. What makes (or would make) such a claim true? We might try to regard it as some kind of generalization over individual instances of particular smokers whose heart disease has been caused by their smoking. But it is clearly not a universal generalization: the fact that Jane – a lifelong smoker – had no heart disease when she died from other causes (falling off a cliff, say) does not serve to refute the claim that smoking causes heart disease. Nor is our general causal claim merely an existential claim to the effect that at least one smoker has been caused by their smoking to develop heart disease. If that were so, then it would take just one unlucky person to die because of choking on a Mars Bar to make it true that eating Mars Bars causes imminent death.

This being so, we should perhaps regard general and ‘token’ or ‘individual’ causation as two different relations, the former relating types or properties, and the latter relating individual events or instantiations of properties – something that is suggested in any case by the grammatical form of general causal claims (‘smoking causes heart disease’, and not ‘instances of smoking cause instances of heart disease’). Indeed, some authors have argued that we have not only two different relations, but two relations that are independent of one another (e.g. Eells 1991, Ch.1). Others argue that the two ‘levels’ of causation are in fact closely related (Hitchcock 1995), so that, once we have the right account of causation on the table, we will see that we do not really have two distinct species of causation at all.

What is generally accepted, however, is that general causation should be understood in terms of probabilistic or statistical relationships, so that, in essence, a claim of the form ‘*C* causes *E*’ is made true by *C*’s raising the probability of *E*. (Similarly, ‘*C* prevents *E*’ is made true by *C*’s lowering the probability of *E*’.) According to one account (Eells 1991; see also Suppes 1970), we should understand *C*’s raising the probability of *E* as the obtaining of $\Pr(E \mid C \ \& \ B_i) > \Pr(E \mid \sim C \ \& \ B_i)$, where $B_1, B_2, \dots$ are all the different combinations of the presence and absence of other relevant factors within a given population. So, for example, suppose we divide our

population (UK citizens, for example) into ‘homogeneous reference classes’, so that within each reference class everyone is exactly alike with respect to all other risk factors for heart disease (alcohol consumption, cholesterol levels, relevant genetic factors, etc.). Then ‘smoking causes heart disease’ is true if and only if in each such reference class, $\Pr(\text{heart disease} \mid \text{smoker}) > \Pr(\text{heart disease} \mid \text{non-smoker})$.

This condition is extremely hard to satisfy, however. For example, imagine that there is a very rare genetic trait that interacts with smoking in such a way as to decrease the risk of heart disease. Then on the above account it will be false that smoking causes heart disease for any population that includes possessors of this genetic trait as a sub-population. A rival proposal (Dupré 1984, 1990) is that C causes E if and only if C raises the probability of E in a fair sample of the population (i.e. in a sample of the population in which the relative frequency of other relevant factors is the same as it is for the population as a whole). This analysis yields the desired result that smoking causes heart disease in the above example, since in any fair sample those with the genetic trait will be vastly outnumbered by those who lack it; so C will still raise the probability of E in a fair sample. An alleged benefit of Dupré’s analysis is that it makes the truth conditions for general causal claims match up in an obvious way with a standard way in which general causal claims are in fact established, namely randomized control trials, since the point of randomization is precisely to make it likely that other relevant factors are present with equal frequency in both the trial group and the control group; and an increase in frequency of the effect under investigation in the trial group relative to the control group is supposed to be good evidence that the drug, or whatever, is a cause of that effect.

The accounts just discussed are reductionist – indeed, Humean – in spirit, in that they seek to analyse (general) causation in terms of regularities (although the regularities in question are statistical rather than exceptionless). A more recent and similar approach – the ‘manipulability’ or ‘interventionist’ approach – shuns reductionism, but nonetheless aims to shed light on how causal structure can be inferred from statistical evidence (Woodward 2003, 2008).

The need for the appeal to either homogenous reference classes or to fair samples stems in part from the need to rule out ‘spurious causation’: cases of probability-increase in the absence of causation. One factor F can easily increase the probability of another, G , in the population as a whole without causing G , as when, for example, F and G are effects of a common cause H . For example, the barometer needle pointing to ‘rain’ (F) is positively correlated with rain (G), but

both are effects of a common cause, namely low atmospheric pressure (H). (We can think of the population in this case as comprising ‘weather episodes’.) There will be no such spurious correlation in a homogeneous reference class, however, since by hypothesis all members of any such reference class will be alike with respect to whether or not they possess H . ‘Holding fixed’ factor H screens off the statistical relationship between F and G . Similarly, if we compare a fair sample of F s with a fair sample of $\sim F$ s, by definition H will have the same relative frequency within each of the samples, and so, again, there will be no correlation between F and G .

Interventionist accounts provide a different way of ‘breaking’ spurious probabilistic relationships. The basic idea is very simple. Take the barometer case again. F and G are positively correlated within the population as a whole. But if I intervene on F , say by manually moving the pointer so it either points to ‘rain’ (f_1) or to ‘sunny’ (f_2), the correlation disappears: there is no positive correlation between F and G if we consider just those cases where I have directly brought about either f_1 or f_2 . (so long as my manually moving the pointer is not itself correlated with the atmospheric pressure). In effect, my intervening breaks the causal (and hence probabilistic) connection between the putative cause we’re interested in – here, the position of the barometer needle – and its usual causes (atmospheric pressure).

A great advantage of the interventionist account is that, with the notion of an intervention in place, we can define causation in terms of increase in probability without the need to make any stipulations about other (relevant) background factors – unlike the probabilistic accounts considered previously. Interventionist accounts are not, however, reductionist accounts, since the notion of an intervention is itself defined in causal terms: an intervention is, by definition, the ‘setting’ of a variable X (e.g. needle position) at a particular value x (e.g. ‘rain’) in a way that breaks the connection between (or ‘screens X off’ from) its usual causes.

3.5 One concept of causation or two?

We can think of the various accounts of causation surveyed above as falling into two distinct categories. On the one hand, we have process and mechanistic accounts (§3.3), which take causation to be a matter of what we might call ‘production’; on the other, we have counterfactual and probabilistic accounts, which take causation to be a matter of ‘difference-making’. Other analyses of causation not considered here are generally pretty easy to categorize as either ‘production’ or ‘difference-making’ accounts.

To see how these two notions can in principle come apart, consider a ‘possible world’ where much of what goes on is a result of spells cast by witches and wizards. Let’s stipulate that the spells can operate at large spatial and temporal distances, and that they do so directly: there is no process or mechanism that connects the casting of a spell to its implementation. It seems right to say that when Merlin, at midday, casts a spell to turn the prince into a frog at midnight – which he duly does – the former caused the latter. Merlin’s spell made a difference to what happened, but there was no process connecting the two: the former did not ‘produce’ the two, in just this sense of ‘produce’ (Schaffer 2000a). Or consider a less far-fetched scenario: Linda the lifeguard sees a swimmer in trouble and rushes to help, intent on preventing him from drowning. Tanya deliberately trips Linda up as she runs towards the sea, thus preventing her from preventing the swimmer from drowning. The swimmer duly drowns. (Such cases are known as cases of ‘double prevention’.) It seems plausible to say that Tanya’s action is a cause of the swimmer’s death; after all, had Tanya not tripped Linda up, the swimmer would have been saved. But again there is no process connecting the two: we have the situation on the beach involving Linda and Tanya, and the swimmer way out at sea, with nothing linking them together. So again we seem to have difference-making without production (Schaffer 2000b).

It looks as though we can also have production without difference-making. There are two routes of the same length to the shop – route *A* and route *B*. Which route I take makes no difference to the time and manner in which I arrive at the shop, but I am intent on going to the shop and will definitely take one or other of these routes; in fact, today I take route *A*. It looks as though my taking route *A* is not a difference-maker: if I hadn’t taken route *A*, I would have taken the route *B*, and I would have got to the shops just the same. However, there is clearly a causal process linking my setting out on route *A* and my subsequent arrival: I put one foot in front of the other, repeated a lot of times, and as a result reached the shop. So we have production without difference-making.

Ned Hall (2004) argues that we really have two distinct concepts of causation: the production concepts and the difference-making concept. He diagnoses the repeated failure of various analyses of causation to cover the full range of standard problem cases by noting that analyses fall into either the production category or the difference-making category; so analyses of the former kind will inevitably fail to deal with cases of difference-making without production, and analyses of the latter kind will inevitably fail to deal with cases of production

without difference-making. And, he claims, it is easy to fail to distinguish the two kinds of causation because cases of one without the other are pretty rare: in our ordinary dealings with the world, causes are both difference-makers and producers of their effects. (When I sink the black ball in snooker, my taking the shot causes the black ball to drop into the pocket in both senses of ‘cause’.)

Many philosophers, however, have been reluctant to accept that there is no univocal concept of causation (Lewis 2004). For example, Dowe (2001), whose Conserved Quantity Theory is clearly a ‘production’ account, argues that cases of difference-making without production are not really cases of causation at all. Meanwhile, defenders of difference-making accounts argue that alleged cases of production without difference-making can in fact be dealt with by their accounts. For example, we might respond to the shop case just described by appealing to the idea that causation is context-sensitive, in the sense that causal claims are true or false relative to contrast class (which is normally only implicit). Thus my taking route *A* rather than route *B* really is not a cause of my subsequent arrival, since my taking route *A* rather than route *B* made no difference to the time and manner of my arrival. My taking route *A* to the shops rather than taking no route at all, however, is a cause of my arrival. So we can save the intuitive judgement that my taking route *A* is a cause of my arrival by finding a context within which the former does make a difference to the latter; hence we don’t have a clear counter-example to a broadly difference-making account after all (Menzies 2007).

5 Causes and Laws

The above discussion of causation and laws of nature has treated them as separate phenomena, to be given distinct analyses; yet intuitively there is some relationship between, for instance, the law that all metals expand when heated and the fact that a particular piece of metal’s expansion is caused by its being heated. How, then, are causation and laws related?

Different theories of causation connect with laws of nature in different ways. Hume’s constant conjunction requirement on causation, for example, can be seen as a requirement that causes and effects are lawfully correlated (or ‘fall under’ a covering law). Laws of nature enter into Lewis’s counterfactual analysis of causation indirectly, via his analysis of counterfactuals (see Lewis 1973b): for Lewis, the truth of ‘if *A* had been the case, *B* would have been the case’

requires that B is true at the ‘closest possible world(s)’ in which A is true; and closeness of possible worlds depends in part on the extent to which they share the same laws of nature. Dowe’s Conserved Quantity Theory of causation (§3.3 above) identifies the causal relation with the transfer of any quantity that is subject to conservation laws, so that, again, facts about causation depend on facts about the laws of nature – though in a rather different way. Heathcote and Armstrong (1991) argue that the causal relation is identical with N , the relation of nomic necessity in virtue of which laws of nature obtain (see §2.2 above).

On the other hand, it has been argued that lawful correlation is not necessary for causation. Anscombe (1971), for example, holds that c may cause e even though c and e do not fall under any regularity. (Causation of this kind is sometimes referred to as ‘singular causation’; but note that ‘singular causation’ can also be used to refer to particular causation, in contrast to general causation.) Heathcote and Armstrong (1991) claim that even though the causal relation is in fact identical with nomic necessity, this identity does not itself hold as a matter of necessity. Heathcote, Armstrong and Anscombe all agree, then, that ‘higgledy-piggledy’ worlds are conceivable – that is, possible worlds where there is plenty of causation but where causally related events are not subsumable under any laws (and so under any regularities) at all.

See also:

Causal Inference and Statistical Fallacies; Causation (Theories and Models): Conceptions in the Social Sciences; Causation: Physical, Mental, and Social; Counterfactual Reasoning, Qualitative: Philosophical Aspects; Counterfactual Reasoning, Quantitative: Philosophical Aspects; Determinism: Social and Economic; Empiricism, History of; Logical Positivism and Logical Empiricism; Natural Law; Probability and Chance: Philosophical Aspects

Bibliography

- Anscombe, G. E. M., 1971. *Causality and Determination*. Cambridge University Press, London.
- Armstrong, D. M., 1983. *What Is A Law of Nature?*. Cambridge University Press, Cambridge.
- 1997. *A World of States of Affairs*. Cambridge University Press, Cambridge.
- Beebe, H., 2000. The non-governing conception of laws of nature. *Philosophy and Phenomenological Research* 56, 571-94.
- 2006a. *Hume on Causation*. Routledge, Abingdon.

- 2006b. Does anything hold the universe together?. *Synthese* 149, 509-33.
- 2009. Causation and observation, in: Beebe, H., Hitchcock, C. and Menzies, P. (Eds.), *The Oxford Handbook of Causation*. Oxford University Press, Oxford, pp. 471-97.
- 2011. Necessary connections and the problem of induction. *Nous* 45, 504-27.
- , Mele, A. R. 2002. Humean compatibilism. *Mind* 111, 201-23.
- Berofsky, B., 2012. *Nature's Challenge to Free Will*. Oxford University Press, New York.
- Bird, A., 1998. Dispositions and antidotes. *The Philosophical Quarterly* 48, 227-34.
- 2005. Laws and essences, *Ratio* 18, 437-61.
- 2008. The epistemological argument against Lewis's regularity view of laws. *Philosophical Studies* 138, 73–89.
- Carroll, J. W., 1994. *Laws of Nature*. Cambridge University Press, Cambridge.
- Cartwright, N., 1989. *Nature's Capacities and their Measurement*. Cambridge University Press, Cambridge.
- 2004. Causation: one word, many things. *Philosophy of Science* 71, 805-19.
- Dowe, P., 2000. *Physical Causation*. Cambridge University Press, Cambridge.
- 2001. A counterfactual theory of prevention and 'causation' by omission. *Australasian Journal of Philosophy* 79, 216-26.
- 2008. Causal Processes, in: Zalta, E. N. (Ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2008 Edition). <http://plato.stanford.edu/archives/fall2008/entries/causation-process> (accessed 31.07.13).
- Dupré, J. 1984. Probabilistic causality emancipated. *Midwest Studies in Philosophy* 9, 169–175.
- Dupré, J. 1990. Probabilistic causality: A rejoinder to Ellery Eells. *Philosophy of Science* 57, 690–98.
- Earman, J, Roberts, J. 1999. Ceteris paribus, there is no problem of provisos. *Synthese* 118, 439–78.
- Eells, E. 1991. *Probabilistic Causality*. Cambridge University Press, Cambridge.
- Ellis, B. 2002. *The Philosophy of Nature: A Guide to the New Essentialism*. Acumen, Chesham.
- Glennan, S. 1996. Mechanisms and the nature of causation. *Erkenntnis* 44, 49-71.
- Hall, E. J. 2004. Two concepts of causation, in: Collins, J., Hall, E. J., Paul, L. A. (Eds.), *Causation and Counterfactuals*. MIT Press, Cambridge, MA.

- Hausman, D. 1992. *The Separate and Inexact Science of Economics*. Cambridge University Press, Cambridge.
- Heathcote, A., Armstrong, D. M. 1991. Causes and laws. *Noûs* 25, 63–73.
- Hitchcock, C. 1995. The mishap at Reichenbach Fall: singular vs. general causation. *Philosophical Studies* 78, 257–91.
- Hume, D. 1748/1975. *Enquiry Concerning Human Understanding*. In Selby-Bigge, L. A., Nidditch, P. (Eds.), *Enquiries Concerning Human Understanding and Concerning the Principles of Morals*, third ed. Clarendon Press, Oxford.
- Lange, M. 2002. Who's afraid of ceteris paribus laws? Or: How I learned to stop worrying and love them. *Erkenntnis* 52, 407–23.
- Lewis, D. K. 1973a. Causation. *Journal of Philosophy* 70, 556–67.
- 1973b. *Counterfactuals*. Blackwell, Oxford.
- 1986. Events, in his *Philosophical Papers*, Vol. II. Blackwell, Oxford, pp. 241–69.
- 1997. Finkish dispositions. *The Philosophical Quarterly* 47, 143–58.
- 2000. Causation as influence. *The Journal of Philosophy* 97, 182–97.
- 2004. Void and object, in: Collins, J., Hall, E. J., Paul, L. A. (Eds.), *Causation and Counterfactuals*. MIT Press, Cambridge, MA.
- Machamer, P., Darden, L., Craver, C. 2000. Thinking about mechanisms. *Philosophy of Science* 67, 1–25.
- Mellor, D. H. 1995. *The Facts of Causation*. Routledge, London.
- Menzies, P. 1989. Probabilistic causation and causal processes: a critique of Lewis. *Philosophy of Science* 56, 642–63.
- 1996. Probabilistic causation and the preemption problem. *Mind* 105, 85–117.
- 1998. How justified are Humean doubts about intrinsic causal links?. *Communication and Cognition* 31, 339–64.
- 2007. Causation in context, in: Corry, R., Price, H. (Eds.), *Causation, Physics and the Constitution of Reality: Russell's Republic Revisited*. Oxford University Press, Oxford.
- Noordhof, P. 1999. Probabilistic causation, preemption, and counterfactuals. *Mind* 108, 95–125.
- Pietroski, P., Rey, R. 1995. When other things aren't equal: saving ceteris paribus laws from vacuity. *British Journal for the Philosophy of Science* 46, 81–110.

- Salmon, W. 1984. *Scientific Explanation and the Causal Structure of the World*. Princeton University Press, Princeton, NJ.
- Schaffer, J. 2000a. Trumping preemption. *The Journal of Philosophy* 97, 165-81.
- 2000b. Causation by disconnection. *Philosophy of Science* 67, 285-300.
- Schrenk, M. 2007a. *The Metaphysics of Ceteris Paribus Laws*. Ontos, Frankfurt.
- 2007b. Can capacities rescue us from ceteris paribus laws?, in Kistler, M., Gnassounou, B. (Eds.), *Dispositions and Causal Powers* Ashgate, Aldershot, pp. 221–47.
- Strawson, G. 1989. *The Secret Connexion*. Clarendon Press, Oxford.
- Suppes, P. 1970. *A Probabilistic Theory of Causality*. North-Holland, Amsterdam.
- Swartz, N. 1985. *The Concept of Physical Law*. Cambridge University Press, New York.
- Williamson, J. 2011. Mechanistic theories of causality. *Philosophy Compass* 6, 421-47.
- Woodward, J. 2003. *Making Things Happen*. Oxford University Press, Oxford.
- 2008. Causation and Manipulability, in: Zalta, E. N. (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2008 Edition).
<http://plato.stanford.edu/archives/win2008/entries/causation-mani> (accessed 31.7.13).
- , Hitchcock, C. 2003a. Explanatory generalizations, part I: a counterfactual account. *Noûs* 37, 1–24.
- 2003b. Explanatory generalizations, part II: plumbing explanatory depth. *Noûs* 37, 181–99.