RECENT WORK

CAUSATION

HELEN BEEBEE
The University of Manchester

THE CONCEPT OF CAUSATION

There was a time when analytic philosophers of an empiricist persuasion wanted to have as little to do with causation as possible. Prompted by Hume's scepticism about causation,[1] and perhaps also by Russell's remark that 'the law of causality' was "a relic of a bygone era, surviving, like the monarchy, only because it is erroneously supposed to do no harm",[2] they tried to do without causation, lest they be accused of metaphysical hocus pocus.

Causation, however, is very difficult to live without. It is well nigh impossible to think about a whole raft of other issues—for instance responsibility, knowledge, perception, and freedom—without invoking the concept of causation; and in recent years appeals to causation in theorising about other areas have gone hand in hand with attempts to analyse causation, thereby rendering it distinctly less disreputable than it once was.

Nonetheless, the concept of causation is a murky one at best. When one considers allegedly paradigm cases—billiard balls colliding with one another, shootings, poisonings, agents' satisfying their desires and so on—it might seem that there is reasonable agreement about the extension of the concept. Unfortunately, however, there is very much less agreement than one might suppose. For example there is no consensus in the literature about whether alleged cases of chance-decreasing causation are really causation, whether alleged causation by absence is really causation, and whether hasteners and delayers are causes. Such cases are not at all obscure, borderline or even rare; they crop up routinely in the actual world. There is even less agreement when it comes to considering the extension of the concept in other possible worlds. There is no consensus, for example, about whether there could be backwards causation, or causation at a temporal distance, or causation without laws of nature.

---

1. Or, at least, what was widely held to be Hume's scepticism. Hume has recently been reinterpreted as a realist about causation; see for instance Galen Strawson, *The Secret Connexion: Causation, Realism, and David Hume* (Oxford University Press, 1989), and Simon Blackburn's riposte in 'Hume and Thick Connexions', *Philosophy and Phenomenological Research*, supp. 50 (1990), pp. 237–250, reprinted in his *Essays in Quasi-Realism* (Oxford University Press, 1993).
2. Bertrand Russell, 'On the Notion of Cause', *Proceedings of the Aristotelian Society* 13, (1912–13), pp. 1–26; reprinted in his *Mysticism and Logic* (Allen & Unwin, 1917); the quoted passage is on p. 180 of the latter.

One major difficulty for the project of theorising about causation is therefore the elusiveness of the subject matter: aside from some basic and uncontroversial cases, it is not even clear where to look for the phenomenon under investigation. Or—to put the point a different way—it is difficult to judge the adequacy of a given analysis when there is so little agreement about what would count as an extensionally adequate theory.

However, some recent work on causation has explicitly tried to describe the fundamental nature and purpose of the concept of causation, and to see what follows. Hugh Mellor, for instance, in his book *The Facts of Causation*, discusses at some length what he describes as the 'connotations' of causation (causes precede their effects; causes are contiguous to their immediate effects; causes and effects are evidence for each other; causes explain their effects; and causes are means of bringing about their effects) and argues that these connotations entail that causes raise the chances of effects.[3] Peter Menzies, in his 'Probabilistic Causation and the Pre-emption Problem',[4] similarly attempts to identify the 'central tenets' of the 'folk theory of causation' and uses that 'folk theory' to fix the reference of 'causes' in much the same way as functionalist theories of the mind attempt to use 'folk' theories of pain, belief and so on to fix the reference of mental concepts.

More recently still, Daniel Hausman, in his *Causal Asymmetries*, argues that the roots of the concept of causation lie in our interest, as agents, in being able to intervene in and manipulate the world around us; and he offers a theory of causation that is based on the notion of an intervention.[5] Grounding the concept of causation in the concept of agency is a feature not just of Hausman's and Mellor's theories, but also in so-called 'agency theories' of causation.[6]

DETERMINISM VS. INDETERMINISM

One notable feature of most recent analyses of causation has been a commitment to indeterminism—or rather, a commitment to the view that an adequate analysis of causation must apply equally to deterministic and indeterministic worlds.[7] Mellor argues that indeterministic causation is consistent with the connotations of causation;[8] Hausman, on the other hand, defends the view that in indeterministic settings there is, strictly speaking, no indeterministic causation, but rather deterministic causation of probabilities.[9]

---

3. See Mellor, *The Facts of Causation* (Routledge, 1995), especially Chs. 5–7.
4. *Mind*, 105 (1996), pp. 85–116.
5. See D.M. Hausman, *Causal Asymmetries* (Cambridge University Press, 1998), especially Chs. 7 and 8.
6. See for instance Huw Price, 'Agency and Probabilistic Causality', *British Journal for the Philosophy of Science*, 42 (1991), pp. 157–175, and Peter Menzies and Price, 'Causation as a Secondary Property', *British Journal for the Philosophy of Science*, 44 (1993), pp. 187–203. Hausman discusses agency theories in detail in Chs. 5 and 7 of *Causal Asymmetries*.
7. One exception is David Lewis's most recent analysis in 'Causation as Influence', *Journal of Philosophy*, 97 (2000), pp. 182–197, which is deterministic.
8. See *The Facts of Causation*, Ch. 5.
9. See *Causal Asymmetries*, Ch. 9.

Following Patrick Suppes and David Lewis,[10] the standard approach has been to replace the thought that causes are sufficient for, or determine, their effects with the thought that a cause need only raise the probability of its effect. This shift of attention has raised the thorny issue of which kind of analysis of probability, if any, is up to the job of underpinning an analysis of indeterministic causation. The most common notion of probability employed in analyses of singular causation is that of single-case chance;[11] in analyses of general causation, conditional probabilities are standardly used.[12]

REALISM, REDUCTIONISM AND SUPERVENIENCE

The most fundamental question one can ask about causation is: what kind of thing *is* it? The central dispute here, as in other areas of metaphysics, is between realism and reductionism. The realist claim that causal facts fail to supervene on non-causal facts is upheld by, for instance, David Armstrong and Michael Tooley.[13] Tooley argues that 'causation' is a theoretical term and therefore susceptible to analysis despite the failure of supervenience, while Adrian Heathcote and Armstrong argue that the causal relation is (as a matter of *a posteriori* necessity) identical with the instantiation of N, the contingent relation of necessitation that, on Armstrong's view, is constitutive of laws of nature.[14] Armstrong and Tooley therefore differ on the issue of whether causation is in any sense primitive or irreducible, while agreeing about the failure of supervenience.

Reductionists generally hold (though usually only implicitly) that some sort of supervenience relation holds between causal facts and more basic facts, although they differ with respect both to whether they take the supervenience relation to hold contingently or necessarily, and to what they take the relevant supervenience base to be.

10. See P. Suppes, *A Probabilistic Theory of Causality* (North-Holland, 1970) and D.K. Lewis, 'Postscripts to "Causation"', in his *Philosophical Papers Volume II* (Oxford University Press, 1986), pp. 172–213.
11. For one account of single-case chance, see Mellor, *The Facts of Causation*, Chs. 2–4. See Dorothy Edgington's 'Mellor on Chance and Causation', *British Journal for the Philosophy of Science*, 98 (1997), pp. 411–433, for some criticisms of Mellor's view. For a different account, see Lewis, 'A Subjectivist's Guide to Objective Chance', in R.C. Jeffrey (ed.), *Studies in Inductive Logic and Probability*, Volume II (University of California Press, 1980), reprinted with postscripts in Lewis's *Philosophical Papers Volume II*, pp. 83–132; and Lewis, 'Humean Supervenience Debugged', *Mind*, 103 (1994), pp. 473–490 (1994), reprinted in his *Papers in Metaphysics and Epistemology* (Cambridge University Press, 1999).
12. See for instance Eells, *Probabilistic Causality* (Cambridge University Press, 1991), Ch. 1.
13. See Adrian Heathcote and D.M. Armstrong, 'Causes and Laws', *Nous*, 25 (1991), pp. 63–73, Armstrong, *A World of States of Affairs* (Cambridge University Press, 1997), Ch. 14, and Tooley, *Causation: A Realist Approach* (Oxford University Press, 1987) for realist views. For some arguments against reductionism, see Tooley, 'Causation: Reductionism versus Realism', *Philosophy and Phenomenological Research*, 50 (1990), pp. 215–235, reprinted in E. Sosa and M. Tooley (eds.), *Causation* (Oxford University Press, 1993).
14. For detailed exposition of Armstrong's theory of laws, see his *What is a Law of Nature?* (Cambridge University Press, 1983).

Classic reductionist analyses—for instance older-style regularity theories and more recent counterfactual analyses—are Humean in spirit, in the sense that they deny the existence of intrinsic causal relations: the truth-makers for causal facts are to be found, at least in part, in facts that are extrinsic to the cause and effect.[15] If they are to avoid circularity, such analyses require a suitably 'wholesome' supervenience base. One objection to theories claiming to reduce the causal to the non-causal is that there is no such base, since the analyses appeal to concepts, like 'law of nature' and 'chance', that are tainted by causal or related notions.[16]

The issue of whether the alleged supervenience of the causal on the non-causal is supposed to be contingent or necessary has not been much discussed in the literature, although Peter Menzies's 'Probabilistic Causation and the Pre-emption Problem'[17] presents a theory that explicitly claims that causal facts contingently supervene on the non-causal. Transference theories are also most naturally understood as embodying a contingent supervenience claim: such theories generally aim only to describe what makes causal claims true in worlds with the same laws of physics as our own, and not to describe what would make them true in other, more distant worlds.

COUNTERFACTUAL ANALYSES OF CAUSATION

Most philosophers agree that, by and large, causation goes hand in hand with counterfactual dependence. It's no coincidence that 'the short circuit caused the fire' and 'had the short circuit not occurred, the fire would not have occurred' are both true, and that 'Jack's scratching his nose caused the earthquake' and 'had Jack not scratched his nose, the earthquake would not have occurred' are both false.

There is far less agreement about the reason for this close (though imperfect) correlation between causation and counterfactual dependence. Some philosophers hold that causal facts are what make the counterfactuals true: 'had the short circuit not occurred, the fire would not have occurred' is true just *because* the short circuit caused the fire.[18] For such philosophers, the project of providing a counterfactual analysis of causation is doomed from the outset. Others believe that some form of counterfactual analysis provides the best prospect for a fully reductionist, 'Humean' analysis of causation. According to counterfactual analyses, causal facts depend on counterfactual facts, which in turn depend in large part on regularities (via a suitably Humean analysis of laws of nature). Such analyses have sought to refine the basic idea that truths

15. For a discussion of, and attack on, the motivation for the Humean claim that causal relations are extrinsic, see Menzies, 'How Justified are Humean Doubts about Intrinsic Causal Relations?', *Communication and Cognition*, 31 (1998), pp. 339–364.
16. This issue is discussed in detail in John Carroll's *Laws of Nature* (Cambridge University Press, 1994); see especially Chs. 3 and 5. The expression 'wholesome base' is Carroll's.
17. *Mind*, 105 (1996), pp. 87–117.
18. Versions of this position are defended by Frank Jackson, 'A Causal Theory of Counterfactuals', *Australasian Journal of Philosophy*, 55 (1977), pp. 3–21, Igal Kvart, *A Theory of Counterfactuals* (Hackett, 1985), and Stephen Barker, 'Counterfactuals, Probabilistic Counterfactuals and Causation', *Mind*, 108 (1999), pp. 427–469.

36

about causation reduce to truths about counterfactual dependence in various, increasingly complex, ways in order to avoid the myriad counter-examples and problem cases that have beset earlier attempts.[19] I suspect that still others believe that the counterfactual-analysis research program has run its natural course, and that the mind-bending complexity of recent counterfactual analyses shows, like Ptolemaic epicycles, that the research program is in terminal decline.

Counterfactual analyses began with David Lewis's 'Causation'[20] and the Postscript added some thirteen years later.[21] Lewis's original analysis ran as follows (where c, d, e are events, '$O(c)$' says 'c occurred', and $\Rightarrow$ is the counterfactual conditional):

> An actual event e *causally depends* on actual event c if and only if e *counterfactually depends* on c—which is to say, if and only if $\neg O(c) \Rightarrow \neg O(e)$. A 'chain of causal dependence' is a series of events $<a, b, c, \ldots, n>$ 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 'pre-empted alternative'. Suppose that had c_1 not occurred, e would have occurred anyway, caused instead by c_2. The e does not counterfactually depend on c_1. However, so long as the actual process leading from c_1 to e cut off the alternative, pre-empted process at some point before e occurred, then there will be some event d, occurring after the alternative process has been stopped, such that $<c, d, e>$ is a chain of causal dependence. (Cases that follow this pattern are generally known as cases of 'early pre-emption'.)

In the Postscript, Lewis extends the analysis to cover indeterministic causation by substituting a weaker notion of counterfactual dependence: e counterfactually (and hence causally) depends on c iff had c not occurred, e's chance of occurring would have been lower than it actually was. Deterministic causal dependence thus becomes the extreme case of causal dependence so defined: the case where e's chance of occurring is raised from 0 (what it would have been without c) to 1. (Hugh Mellor adopts a similar approach, although for Mellor the most basic causes and effects are facts rather than events.)[22]

SOME STANDARD PROBLEMS FOR COUNTERFACTUAL ANALYSES

(i) Chance increase without causation

Both Mellor and Lewis face the problem that there can be chance increase without causation, whereas on their accounts chance increase is sufficient for

19. Recent analyses include Murali Ramachandran, 'A Counterfactual Analysis of Causation', *Mind*, 106 (1997), pp. 263–277; Paul Noordhof, 'Probabilistic Causation, Pre-emption and Counterfactuals', *Mind*, 108 (1999), pp. 95–123; and Lewis, 'Causation as Influence'.
20. In *Journal of Philosophy*, 70 (1973); reprinted in his *Philosophical Papers Volume II*, pp. 159–172.
21. 'Postscripts to "Causation" '.
22. Mellor, *The Facts of Causation*.

causation.[23] For example, Fred and Ted both want Jack dead. Fred poisons Jack's soup and Ted poisons his coffee, and each act increases Jack's chance of dying. Jack eats the soup but (feeling rather unwell) leaves the coffee, and dies later. Ted's act raised the chance of Jack's death but was not a cause of it.

One response to this problem is to require that there be a continuous causal process—defined in terms of a continuous chain of counterfactual dependence, say—between c and e in order for c to be a cause of e. There is no such chain between Ted's act and Jack's death; for example, there are no events occurring after Jack has eaten the soup which both counterfactually depend on what Ted did, and upon which Jack's death in turn counterfactually depends. Unfortunately, however, this response rules out causation at a temporal distance.[24]

(ii) Chance decreasing causes?

When c lowers the chance of e, sometimes it seems right to say that c caused e, and sometimes it doesn't. To use some stock examples: suppose a golfer makes a good shot, so that her ball has a very high chance of landing in the cup. Improbably, a squirrel kicks the ball away from its path; and, again improbably, the ball then ricochets off a branch and lands in the cup. Here it seems right (according to many philosophers' intuitions) to say that the kick caused the ball to land in the cup, even though the kick lowered the ball's chance of so doing. On the other hand, suppose I spray a plant with defoliant, thereby lowering its chance of being alive in a month's time. Improbably, the plant survives. Here it seems wrong to say that the spraying was a cause of the survival; rather it seems right to say that the plant survived *despite* being sprayed.[25]

On Lewis's view, and indeed in general on analyses that take causation to be transitive, chance decreasers like the squirrel's kick and the defoliant come out as causes, since in both cases there is a chain of (chance-increasing) counterfactual dependence leading from the first event to the second. On accounts (like Mellor's) that require *all* causes to raise the chances of effects, on the other hand, no chance decreasers are causes. Such accounts are forced to deny that causation is transitive.[26]

23. This problem was originally raised in Peter Menzies, 'Probabilistic Causation and Causal Processes: A Critique of Lewis', *Philosophy of Science*, 76 (1989), pp. 642–663.
24. See Menzies, 'Probabilistic Causation and Causal Processes', for an analysis that defines causation in terms of continuous chains of counterfactual dependence.
25. The squirrel example is originally due to Deborah Rosen ('In Defense of a Probabilistic Theory of Causality', *Philosophy of Science*, 45 (1978), pp. 604–613; the defoliant example is due to Nancy Cartwright, 'Causal Laws and Effective Strategies', *Nous*, 13 (1979), pp. 419–437, reprinted in her *How the Laws of Physics Lie* (Clarendon Press, 1983). For further discussion see for instance Eells, *Probabilistic Causality*, Ch. 6; David Papineau, 'Causal Factors, Causal Inference, Causal Explanation II', *Aristotelian Society Supplementary Volume LX* (1986), pp. 115–136; and Dowe, 'The Conserved Quantity Theory of Causation and Chance-Raising', *Philosophy of Science*, 66 (Proceedings) (1999), pp. S486–S501.
26. Such accounts include Eells's theory of singular causation and Menzies and Price's 'agency' theory; see Eells, *Probabilistic Causality*, Ch. 6, and Menzies and Price, 'Causation as a Secondary Quality'. For further discussion of transitivity see Beebee, 'Taking Hindrance Seriously', *Philosophical Studies*, 88 (1997), pp. 59–79; and Ned Hall, 'Causation and the Price of Transitivity', *Journal of Philosophy*, 97 (2000), pp. 198–222.

(iii) Redundant causation[27]

In general, counterfactual analyses get into trouble in cases where there is causation without counterfactual dependence: c causes e, but had c not occurred, e would have occurred anyway. Such cases are usually known as cases of 'redundant' causation.

Counterfactual analyses generally accommodate 'early' pre-emption by assuming that pre-emption works by what Lewis has called 'cutting':[28] at some time t before the effect e actually occurs, the process leading from the pre-empted non-cause c_2 is 'cut' by the process leading from the pre-empting cause c_1. In such cases, counterfactual analyses can appeal to some event d that occurs as part of the c_1-process between t and the time of e. Such an event will generally counterfactually depend on c_1, and e will in turn counterfactually depend on d, since by the time d occurs there is no longer any alternative process ready to cause e should the c_1-process not run to completion.

A harder problem is 'late pre-emption', the most commonly discussed form of which occurs when the c_1-process runs faster than the pre-empted c_2-process. Here, the pre-empted process is prevented from running to completion only by the occurrence of the effect e itself; hence there is no room between the termination of the c_2-process and the time of e in which to find a suitable event that counterfactually hooks c_1 and e together. In ordinary cases of late pre-emption, 'cutting' still occurs: since e occurs too soon, there are events missing from the c_2-process that would be needed for c_2 to be a cause of e. Some recent analyses have exploited this feature in order to solve the problem.[29] However there is a range of more problematic cases where cutting does not appear to occur, and hence where the above move will not save the analysis. For example, one can add action at a temporal distance into the mix to get cases of pre-emption where there is no cutting.[30]

Another problem case is deterministic overdetermination, where both the process leading from c_1 and the process leading from c_2 run to completion, and hence both c_1 and c_2 seem to be genuine causes of e. However there is no chain of counterfactual dependence running to e either from c_1 or from c_2. So standard counterfactual analyses yield the wrong result that neither c_1 nor c_2 is a cause of e.

27. For some general discussions of, and proposed solutions to these problems, see Michael McDermott, 'Redundant Causation', *British Journal for the Philosophy of Science*, 46 (1995), pp. 323–344; Ehring, *Causation and Persistence*, Chs. 1 and 2; and Lewis, 'Causation as Influence'.

28. See Lewis, 'Causation as Influence'.

29. See for example Ramachandran, 'A Counterfactual Analysis of Causation'; Jonardon Ganeri, Noordhof and Ramachandran, 'Counterfactuals and Pre-emptive Causation', *Analysis*, 56 (1996), pp. 219–225, and 'For a (Revised) PCA Analysis', *Analysis*, 58 (1998), pp. 45–47; and Noordhof, 'Probabilistic Causation, Pre-emption and Counterfactuals'.

30. See for example L.A. Paul, 'Problems with Late Pre-emption', *Analysis*, 58 (1998), pp. 48–53; Jonathan Schaffer, 'Trumping Pre-emption', *Journal of Philosophy*, 97 (2000), pp. 165–181; and McDermott, 'Redundant Causation'.

(iv) Problems with counterfactuals

One problem for counterfactual analyses of causation is that they take for granted what many authors believe to be an erroneous theory of counterfactuals, namely Lewis's.[31] There have been many objections to Lewis's theory of counterfactuals, both in general and in the context of analysing causation. Some authors object, for instance, to the idea that the closest world where some actual event *c* fails to occur is 'miraculous' relative to the actual world, thus making it true that had *c* not occurred the laws of nature would have been different.[32] Some object that Lewis's analysis faces special problems in an indeterministic setting.[33]

WHAT ARE CAUSES AND EFFECTS?

The current market offers a range of alternative candidates for the role of cause and effect, the most popular choices being facts, events, and tropes. Davidson's argument that causation must be a relation between events[34] has been rebutted by Mellor, but his own argument that the most basic causes and effects must be facts (and that the most basic kind of causation is not relational) has also been criticised.[35] Lewis-style events seem to be the overwhelming favourite amongst those offering counterfactual analyses of causation,[36] although the precise nature of those events is much disputed. For example, taking the time of an event to be an essential property of it provides an easy solution to the problem of late pre-emption, but unfortunately entails that events cannot be hastened or delayed—merely replaced by different,

---

31. For Lewis's theory of counterfactuals, see his *Counterfactuals* (Blackwell, 1973), and his 'Counterfactual Dependence and Time's Arrow', *Nous*, 13 (1979), pp. 455–476, reprinted in his *Philosophical Papers Volume II*. Most counterfactual analyses of causation at least tacitly presuppose Lewis's analysis of counterfactuals. Mellor's theory of causation also employs counterfactuals, although it is not clear whether or to what extent he endorses Lewis's analysis. For a critical discussion of Lewis's analysis of counterfactuals in the context of analysing causation, see Hausman, *Causal Asymmetries*, Ch. 6.
32. See for example Carroll, *Laws of Nature*, Appendix B.
33. See for example Jackson, 'A Causal Theory of Counterfactuals', and Barker, 'Counterfactuals, Probabilistic Counterfactuals and Causation'.
34. Donald Davidson, 'Causal Relations', *Journal of Philosophy*, 64 (1967), reprinted in his *Essays on Actions and Events* (Oxford University Press, 1980), pp. 163–180.
35. Mellor provides a long discussion of the facts vs. events dispute—arguing in favour of facts—in *The Facts of Causation*, Chs. 9–14. For some criticisms of Mellor's arguments, see Noordhof, 'Critical Notice: Causation, Probability, and Chance', *Mind*, 107 (1998), pp. 855–875.
36. See Lewis, 'Events', in his *Philosophical Papers Volume II*, pp. 241–269, for the standard theory of events employed; the theory of events is tailor-made to fit Lewis's theory of causation. Unfortunately some critics of the view that the primary causal relata are events have aimed their objections not at Lewis-style events but at events as characterised by Davidson (in 'Causal Relations' and in 'The Individuation of Events' in N. Rescher (ed.), *Essays in Honor of Carl G. Hempel* (1969), reprinted in Davidson's *Essays on Actions and Events*) or Jaegwon Kim ('Causation, Nomic Subsumption, and the Concept of Event', *Journal of Philosophy*, 70 (1973), pp. 217–236); see for example Mellor, *The Facts of Causation*, Ch. 11, and Ehring, *Causation and Persistence*, Ch. 3.

40

earlier or later events.[37] Tropes are relative newcomers to the scene, but have enjoyed some recent popularity.[38]

One issue that provides a clear distinction between event- and trope-based accounts on the one hand and fact-based accounts on the other concerns the causal status of absences. If (to use an example of Mellor's) Kim's having no children was caused by her use of contraceptives, then it seems that causation cannot be a relation between events, since, while there is a perfectly good candidate for a factual effect (namely the fact that Kim has no children), there is no *event* (or trope) that is Kim's having no children. Supporters of a relational, event- or trope-based conception of causation cannot allow absences to be causes or effects, while supporters of a non-relational, fact-based account of causation can (or rather, they can allow *facts* about absences to be causes and effects).

TRANSFERENCE THEORIES

'Transference' theories of causation hold that causation consists in—or at least has as its basic ingredient—the transfer or transmission or continued possession of something—a particular kind of property, for instance—between cause and effect. Such theories are directly opposed to the central Humean claim that there is literally nothing in the world that connects causes and effects; on the other hand, transfer theorists are generally Humean in the sense that they do not take causation to be a fundamental, irreducible constituent of reality.[39]

Douglas Ehring's theory is a transference theory in the above sense: for Ehring, the most basic kind of causation is the literal transfer of a trope between cause and effect. However, other transference theories identify causation with the sorts of thing physicists (as opposed to ontologists) talk about.[40] Early transfer theorists included Jerrold Aronson, David Fair, and Wesley Salmon.[41] Aronson claimed that causation amounts to the transfer of some physical quantity (like momentum or heat); Fair identified causation with 'energy-momentum flow', and Salmon characterised a causal process as a process with the ability to transmit a mark.

37. See Penelope Mackie, 'Causing, Delaying and Hastening: Do Rains Cause Fires?', *Mind*, 101 (1992), pp. 483–500; and L.A. Paul, 'Keeping Track of the Time: Emending the Counterfactual Analysis of Causation', *Analysis*, 58 (1998), pp. 191–198.
38. See Ehring, *Causation and Persistence*, Ch. 3; Hausman, *Causal Asymmetries*, Ch. 2; and L.A. Paul, 'Aspect Causation', *Journal of Philosophy*, 97 (2000), pp. 235–256. (For Paul, the relata of causation are 'property instances' rather than tropes.)
39. See Wesley Salmon, 'A New Look at Causality' in his *Causality and Explanation* (Oxford University Press, 1998).
40. They are also concerned primarily with causal processes, rather than with causation per se.
41. See Aronson, 'On the Grammar of Cause', *Synthese*, 22 (1971), pp. 414–430, and 'The Legacy of Hume's Analysis of Causation', *Studies in the History and Philosophy of Science*, 2 (1971), pp. 135–157; Fair, 'Causation and the Flow of Energy', *Erkenntnis*, 14 (1979), pp. 219–250; and Salmon, *Scientific Explanation and the Causal Structure of the World* (Princeton University Press, 1984).

More recently, Salmon and Phil Dowe have developed versions of the 'conserved quantity theory', according to which causal processes are characterised by the possession by an object, or the exchange between different objects, of a 'conserved quantity' (such as linear momentum, charge or mass-energy).[42]

The central distinction for both Dowe and Salmon is the distinction between a causal process and a 'pseudo-process'. Pseudo-processes include, for example, the movement of a spot of light on a wall, or a shadow. Salmon's earlier 'mark theory' ruled such processes as non-causal by appealing to the fact that they are incapable of transmitting a mark: modifications to earlier stages of the light spot or shadow are not preserved in later stages. However, Salmon's theory involved appeal to counterfactuals, since the test for a causal process was not whether the process *does* in fact transmit a mark, but whether it *can* do so: in other words, whether possible modifications to earlier stages of the process *would* be transmitted to later stages.[43] Dowe's theory, on the other hand, requires no such appeal: genuine causal processes *do*, while pseudo-processes do not, possess conserved quantities. Pseudo-processes like light spots and shadows often retain *some* properties over time—shape and size, for example—but not, according to Dowe, the right sorts of properties: that is, not conserved quantities.

THEORIES OF GENERAL CAUSATION

Some recent analyses of causation have sought to analyse not particular or singular causal facts (like 'John's smoking caused his heart attack') but general or population or type-level causal facts (like 'smoking causes heart attacks').[44]

The standard method has been to analyse population causation in terms of relative frequencies in 'homogeneous reference classes' or 'background contexts',[45] and to take 'C causes E' to be a matter of C's raising the probability of E within some or all of these background contexts. A background context is a subset of the population, whose members are all the same with respect to which factors relevant to E (apart from C itself) they possess. For example, if

42. See Salmon, 'Causality without Counterfactuals', *Philosophy of Science*, 61 (1994), pp. 297–312, reprinted in his *Causality and Explanation*, and 'Causality and Explanation: a Reply to Two Critics', *Philosophy of Science*, 44 (1997), pp. 461–477; Dowe, 'Process Causality and Asymmetry', *Erkenntnis*, 37 (1995), pp. 179–196, 'Wesley Salmon's Process Theory of Causality and the Conserved Quantity Theory', *Philosophy of Science*, 59 (1992), pp. 195–216, and 'Causality and Conserved Quantities: A Reply to Salmon', *Philosophy of Science*, 62 (1995), pp. 321–333.
43. For this objection see Philip Kitcher, 'Explanatory Unification and the Causal Structure of the World', in Kitcher and Salmon (eds.), *Scientific Explanation*, vol. 13, *Minnesota Studies in the Philosophy of Science* (University of Minnesota Press, 1989), pp. 410–505, section 6.
44. See for example Suppes, *A Probabilistic Theory of Causality*; John Dupré, 'Probabilistic Causality Emancipated', *Midwest Studies in Philosophy IX* (1984), pp. 169–175; Papineau, 'Causal Factors, Causal Inference, Causal Explanation II'; and Eells, *Probabilistic Causality*.
    There is no generally accepted word for either kind of causation. For example, Eells (*Probabilistic Causality*) contrasts 'type-level' with 'token-level' causation, and Mellor (*The Facts of Causation*) contrasts 'singular' with 'general' causation.
45. The expression 'homogeneous reference class' is Salmon's (see his *Statistical Explanation and Statistical Relevance* (Pittsburgh University Press, 1971)). Eells uses the term 'background context'.

there are two factors apart from C that are relevant to E (call them X and Y), there will be four background contexts: $X\&Y$, $X\&\sim Y$, $\sim X\&Y$, and $\sim X\&\sim Y$. Each member of the population under investigation will be a member of one and only one background context.

Looking at C's probabilistic impact on E within each background context rather than simply within the population as a whole is the standard way of ensuring that the probabilistic correlation between C and E is not spurious. Falling barometer readings (C) raise the probability of rain (E)—that is, $\Pr(E/C) > \Pr(E/\sim C)$—even though C does not cause E. Rather, C and E are both effects of a common cause: low atmospheric pressure (F). When we construct background contexts—that is, when we 'hold fixed' other relevant factors (in this case, F) when assessing C's probabilistic impact on E—we consider the correlation between C and E in the presence of F and, separately, in the absence of F. What we find in the barometer case is that $\Pr(E/C \& F) = (\Pr(E/\sim C \& F)$ and $\Pr(E/C \& \sim F) = \Pr(E/\sim C \& \sim F)$: the probabilistic correlation between C and E disappears when we hold fixed the relevant factor, and this reflects the fact that falling barometer readings do not cause rain.

According to Eells's analysis, C causes (or is a *positive causal factor for*) E in a particular population if and only if C raises the probability of E in every background context of that population; that is, if and only if $\Pr(E/C \& B_i) > \Pr(E/\sim C \& B_i)$ for each background context B_i. Similarly C is a *negative causal factor* for E if C lowers the probability of E in every background context, and C is *causally neutral* for E if C makes no difference to the probability of E in each background context.[46]

John Dupré has argued that analyses like Eells's are too strong: they make it too hard for general causal claims to be true.[47] Suppose that some tiny minority of the US population has some peculiar physiological condition P that makes smoking (C) a prophylactic against heart disease (E)—even though for everyone else in the population, smoking increases the risk of heart disease. It follows from Eells's analysis that smoking does not cause heart disease in the US population, since $\Pr(E/C \& P) < \Pr(E/\sim C \& P)$. This is a result that Dupré regards as highly counter-intuitive.

Dupré's rival analysis takes as its starting point a method that is actually used to test causal hypotheses in the social and medical sciences: that of the controlled experiment. If one wants to know whether C causes E in a given population, one way of trying to find out is to take a random sample of the population, induce C in a random subset of that sample, and compare the results. The point of having a *random* sample is to try to ensure that other factors that are relevant to E occur with the same relative frequency as they do in the population as a whole, and thus that the probabilistic correlation between C and E is not spurious. Dupré calls a sample that achieves this match of frequencies a *fair* sample. His claim is that C causes E if and only if

46. See Eells, *Probabilistic Causality*, Chs. 2–4.
47. See his 'Probabilistic Causality Emancipated' and his 'Probabilistic Causality: A Rejoinder to Ellery Eells', *Philosophy of Science*, 58 (1990), pp. 690–698. See also Eells, 'Probabilistic Causality: Reply to John Dupré', *Philosophy of Science*, 54 (1987), pp. 105–114.

C raises the probability of *E* in a fair sample of the population. This analysis yields the desired result that smoking causes heart disease in the above example, since in any fair sample those with the physiological condition *P* will be vastly outnumbered by those who lack *P*; so *C* will still raise the probability of *E* in a fair sample.

A problem with probabilistic theories of general causation, however, is that they appear to fail given the assumption that at least some features of the world are determined by prior circumstances. For example, suppose that *C* causes *E*, but that *C* itself is determined by the combination of factors *XYZ*. Then it isn't true that $\Pr(E/C \ \& \ XYZ) > \Pr(E/{\sim}C \ \& \ XYZ)$, since $\Pr(E/{\sim}C \ \& \ XYZ)$ is undefined; hence, according to Eells's analysis—and contrary to hypothesis—it isn't true that *C* causes *E*.[48] Dupré's analysis also falls prey to this objection. Suppose that 10% of the population have *XYZ*. Then in a fair sample of *C*s, 10% must have *XYZ*; similarly for a fair sample of ~*C*s. But there can be no such fair sample of ~*C*s, since by hypothesis there will be nobody at all who has *XYZ* but lacks *C*. So according to Dupré's analysis it isn't true that *C* causes *E*.[49]

Typical theories of general causation seek to analyse general causation solely in terms of probabilities, so that general causal truths are logically independent of singular causal truths.[50] But it is natural to think that general and particular causal facts are related in some way: it is natural to think that the fact that smoking causes heart disease has *something* to do with the fact that lots of individual smokers are caused by their smoking to get heart disease. This line of thought leads naturally to the view that general causal claims are not claims about a distinctive *kind* of causation—general causation—at all, but are rather merely generalisations about *particular* causal claims. This latter approach is adopted by Hausman and by Carroll, both of whom seek to show, in different ways, how singular causal facts determine general causal facts.[51]

CONCLUDING REMARKS

Perhaps the most striking feature of work on causation in the last decade has been its diversity. Different kinds of theory have been developed in relative isolation from each other, with the overwhelming majority of published debates and disagreements having a distinctly 'in house' character. Part of the explanation for this plurality of approaches doubtless lies in a corresponding plurality of aims that different kinds of theory are trying to achieve. Such aims include: to explicate the 'common sense' concept of causation; to explicate the 'scientific' concept of causation; to vindicate statistical methods in the

48. See John Carroll, 'The Unanimity Theory and Probabilistic Sufficiency', *Philosophy of Science*, 59 (1992), pp. 471–479.
49. Further objections are to be found in Carroll, 'Property-Level Causation?', *Philosophical Studies*, 63 (1991), pp. 245–270; and Hausman, *Causal Asymmetries*, Ch. 5.2*.
50. Eells argues for this position in the Introduction to *Probabilistic Causality*.
51. See Hausman, *Causal Asymmetries*, Ch. 5.2*; and Carroll, 'Property-Level Causation?'.

sciences; to identify causation in the actual world; to identify causation in all possible worlds; to show that causation is or is not a basic element of the world's ontology. In short, there is no consensus on the issue of what a theory of causation should accomplish.

Over the last decade or so, analyses of causation have reached a very high level of technical sophistication. But, given the apparently irresolvable differences of opinion (noted earlier) over what does and does not count as a case of causation, the adequacy of the machinery cannot be judged solely on whether, when one cranks the handle, one gets the 'right' result. Attention to the issue of what the appropriate adequacy criteria are for theories of causation is a line of enquiry that would therefore be well worth pursuing in the coming decade.[1]