

CAUSATION

*Christopher Hitchcock***Introduction**

In a paper read before the Aristotelian Society, Bertrand Russell (1913: 1) claimed:

All philosophers, of every school, imagine that causation is one of the fundamental axioms or postulates of science, yet, oddly enough, in advanced sciences such as gravitational astronomy, the word “cause” never appears... To me, it seems that ... the reason why physics has ceased to look for causes is that, in fact, there are no such things. The law of causality, I believe, like much that passes muster among philosophers, is a relic of a bygone age, surviving, like the monarchy, only because it is erroneously supposed to do no harm.

Russell was hardly alone in that opinion. Other writers of the period, such as Ernst Mach, Karl Pearson, and Pierre Duhem, also rejected as unscientific the notion of causation. Their view was shared also by most of the logical positivists. Indeed, the concept of causation was regarded with suspicion by philosophers, as well as by many statisticians and social scientists, throughout much of the twentieth century.

Contrary to Russell’s claim, however, the most casual perusal of the leading scientific journals reveals that causal locutions are commonplace in science. The 2006 volume of *Physical Review Letters* contains articles with titles like “Inverse Anderson Transition Caused by Flatbands” (by Masaki Goda, Shinya Nishino, and Hiroki Matsuda) and “Softening Caused by Profuse Shear Banding in a Bulk Metallic Glass” (by H. Bei, S. Xie, and E. P. George). Indeed, physicists refer to a variety of phenomena as “effects”: the “Hall effect,” the “Kondo effect,” the “Lamb-shift effect,” the “Zeeman effect,” and so on. Presumably where there are effects, there are causes as well. Causal claims are even more common in the medical sciences: for example, a 2005 editorial by E. K. Mulholland and R. A. Adegbola in the *New England Journal of Medicine* bore the title “Bacterial Infections – a Major Cause of Death among Children in Africa.” Given the ubiquity of causal claims in the sciences, *causation* deserves to be a concept of great interest to philosophers of science.

Analyses of causation

Diverse attempts have been made to analyze causation, and many of the debates that surround the concept of causation stem from fundamental disagreements about the best way to go about the project. Proposed analyses of causation can be divided into two broad categories: *reductive* and *non-reductive*. Reductive analyses of causation aim to provide truth-conditions for causal claims in non-causal terms. Non-reductive analyses of causation aim to establish systematic relationships between causation and other concepts of interest to philosophers; those relationships can then be used to derive interesting non-causal consequences from causal claims, even when the causal claims cannot themselves be paraphrased without causal remainder.

Pressure to provide a reductive analysis of causation comes from at least two sources: epistemology and metaphysics. Epistemological pressure stems from the unobservability of causal relations: we may observe the hot sun and the soft wax, but we do not observe the sun's *causing* the wax to soften. Thus, it seems that in order to assess the truth-value of a causal claim, it must be possible to translate that claim into one that does admit of direct epistemic access. Metaphysical pressure stems from Ockham's razor: in metaphysical system-building, it is preferable to analyze causal relations away rather than posit them as additional ingredients of the world.

Both of these pressures are capable of being resisted. Epistemologically, causal claims may be treated as akin to claims about theoretical entities such as electrons. We do not expect to be able to translate a claim such as that "every hydrogen atom contains one electron" into purely observational terms. All that a reasonable epistemology can demand of us is that such claims be susceptible to empirical confirmation or disconfirmation, for example, by entailing various observational consequences or by rendering some observations more probable than others. Causal claims are regularly subjected to empirical test in the sciences. In the medical sciences, for example, causal claims are often tested using controlled clinical trials. Such tests are capable of providing strong evidence in support of causal claims without the need to reduce those claims to non-causal claims. Metaphysically, systems that include causation as a basic feature of our world need not be unnecessarily complex: causal relations may well be the sorts of basic constituents of our world into which other relations are analyzed.

Challenges

There are a number of challenges that an adequate account of causation must meet. First, an account of causation must be able to distinguish between genuinely causal relationships and merely accidental relationships. Suppose, for example, that only a small handful of human beings eat a particular kind of fruit before the species of plant that bears it becomes extinct. By sheer coincidence, all of these people die shortly after eating the fruit. A theory of causation should not then rule that consumption of this particular fruit causes the death: the relationship between eating the fruit and death is merely accidental. In other words, an adequate theory of causation should entail that *post hoc ergo propter hoc* is, at least sometimes, a fallacy.

A second challenge is to distinguish causes from effects. Typically, perhaps even universally, when one event *C* causes another event *E*, it is not also the case that *E* causes *C*. In such typical cases, an adequate theory of causation must correctly rule that *C* causes *E*, but not vice versa. Some philosophers have attempted to address this problem by stipulating that, *by definition*, causes occur earlier in time than their effects. Thus if we have two events *C* and *E* that are related as cause and effect, we can identify the cause as the one that occurs earlier, and the effect as the one that occurs later. This solution to the problem has the disadvantage that it renders claims of backward-in-time causation false by definition. For example, there are solutions to the general field equations of general relativity that permit closed causal curves: time-like trajectories along which an object could travel from spatio-temporal region *A* to the distant spatio-temporal region *B*, and then back to *A*. Along such a trajectory, it may happen that the state of the object at *A* causes the state of the object at *B*, and the state of the object at *B* causes the state of the object at *A*. While such models may not describe the actual universe, that would seem to be an empirical matter, and not one to be settled *a priori* by our definitions of “cause” and “effect.” Thus it would be desirable for a theory of causation to provide an independent account of the directionality of causation.

A third challenge is to distinguish causes and effects from effects of a common cause. It may be, for example, that smoking causes both stained teeth and lung cancer, with the former occurring before the latter. If so, then it may be common for individuals with stained teeth to develop lung cancer later in life. But stained teeth do not cause lung cancer; rather, stained teeth and lung cancer are effects of a common cause. An adequate theory of causation had better be able to mark the distinction.

Finally, an account of causation ought to be able to distinguish between genuine causes and pre-empted backups. Suppose, for example, that a building receives its electricity from the city’s main power grid. In addition, the building has a backup generator that will kick in if there is a power failure. When the city’s power grid is functioning properly, it is that power source, and not the backup generator, that causes the lights in the building to be on. A successful theory of causation must be able to mark the difference.

Regularity theories of causation

Perhaps the best-known attempt to analyze causal relations is that of David Hume: “we may define a cause to be *an object, followed by another, and where all the objects similar to the first, are followed by objects similar to the second*” (Hume 1977 [1748]: 76; italics in original). Hume, then, analyzes causation in terms of constant conjunction: a cause is always conjoined with its effect. According to Hume, our experience of such a constant conjunction produces in us a customary transition in the mind. Thus “[w]e may ... form another definition of cause; and call it, *an object followed by another, and whose appearance always conveys the thought to that other*” (*ibid.*: 77; italics in original). It is our impression of that mental operation from which our idea of causation is derived.

In the nineteenth century, John Stuart Mill pointed out that simple causes will not invariably be followed by their effects. Thus, for example, smoking will not always be

accompanied by lung cancer: some smokers may not be susceptible, or may die of other causes before cancer develops. In order to account for this sort of case, John Mackie (1974) developed his theory of INUS conditions. An INUS condition is an *insufficient* but *non-redundant* part of an *unnecessary* but *sufficient* condition. Thus *C* will be an INUS condition for *E* if there is a conjunction of factors *ABCD* ... such that whenever these factors occur together, they are followed by *E*, but where the factors *ABD*... without *C* are not invariably followed by *E*. This account allows that *C* may sometimes occur without *E* and vice versa.

One problem with this account is that it may be an accident that all conjunctions of *ABCD* ... are followed by *E*. One strategy for dealing with this problem is to require that the regularity be a consequence of *laws of nature*; that is, it must be possible to derive *E* from *ABCD* ... together with statements describing laws of nature. This strategy is essentially that adopted by Carl Hempel in his Deductive–Nomological model of scientific explanation. There is a sense, however, in which this approach simply relocates the problem, for now we must have an account of laws that distinguishes genuine laws of nature from mere accidental generalizations.

As Hume defined them, causes precede their effects in time. It is hard to see how a regularity theory of causation can capture the asymmetry between causes and effects without this stipulation. For example, critics of Hempel's *deductive–nomological model of explanation* have pointed out that the same laws that can be used to deduce the length of a shadow from the height of a flagpole and the angle of the sun can also be used to derive the height of the flagpole from the length of its shadow; but only the former derivation captures the right causal direction. Similarly, regularity theories of causation have difficulties with effects of a common cause. If there are conditions that when conjoined with smoking are invariably followed by lung cancer, then there may well be further conditions that, when conjoined with stained teeth, are always followed by lung cancer (these further conditions would include, for example, the absence of factors other than smoking that might account for stained teeth).

Finally, regularity theories have trouble distinguishing genuine causes from pre-empted backups. For example, it may well be that whenever a backup generator is in good working order, the lights in a certain building will be on – either because the generator itself is powering them or because the city's power grid is working effectively. But only in the former case would we consider the backup generator to be a cause of the lights being on. These difficulties with regularity theories of causation have led some philosophers to search for alternative accounts of causation.

Probabilistic theories of causation

The success of quantum mechanics in the twentieth century raises the possibility that our world may be indeterministic at the most fundamental level. If so, then causes need not be constantly conjoined with their effects, even if we specify all of the other relevant conditions. It may be that a complete specification of relevant factors *ABCD* ... suffices only to fix a certain probability for *E* to occur. Probabilistic theories of causation embrace this possibility. The central idea is that causes need

not be sufficient for their effects, but need only raise the probabilities of their effects. The most natural way to make this precise is through conditional probability: C raises the probability of E just in case $Pr(E|C) > Pr(E)$, where $Pr(E|C)$ is defined to be $Pr(E \& C)/Pr(C)$.

One worry with this approach is that E may chance to happen more often in the presence of C than in its absence, even though there is no causal relationship between C and E . This is the analog of the problem of accidental generalizations that plagues regularity theories of causation. In order to guard against this possibility, the function Pr must refer to the true underlying probabilities, and not merely to statistical frequencies. This gives rise to the question of how to interpret the relevant probability claims. In particular, since causal relations are objective features of the world, the probabilities should correspond to objective features of the world, and not just to our state of uncertainty about the world.

The basic idea that causes raise the probabilities of their effects does not, by itself, do anything to solve the problems associated with the direction of causation. Indeed, it is easy to show that if $Pr(E|C) > Pr(E)$, then $Pr(C|E) > Pr(C)$. Moreover, if A and B are effects of a common cause, then typically we will have $Pr(A|B) > Pr(A)$ and $Pr(B|A) > Pr(B)$. For example, if A represents lung cancer, and B stained teeth, we would expect to find a greater prevalence of lung cancer among people with stained teeth than in the population at large, for the former group will have a higher proportion of smokers. If we look only at the probability relations among pairs of events, those problems are insoluble; matters change, however, once we consider the probability relationships between three or more events. If C is a common cause of A and B , then it will typically be the case that C screens-off A from B , that is, $Pr(A|BC) = Pr(A|C)$. (Screening-off will fail, however, if A and B share a further common cause in addition to C .) Thus while B might raise the probability of A overall, it does not raise the probability of A conditional on the common cause C . Thus, in judging whether C is a cause of E , we need to consider not the simple probabilities $Pr(E|C)$ and $Pr(E)$ but more complicated conditional probabilities of the form $Pr(E|C \& K)$ and $Pr(E|K)$, where K represents various other causal factors that need to be held fixed. Screening-off relations can also help us to distinguish causes from effects. If C is a common cause of A and B , then, as we have noted, C will typically screen-off A from B . On the other hand, if E is an effect of both A and B , then typically E will not screen-off A from B . We can thus appeal to these distinctive probabilistic signatures to determine whether the causal arrows are pointing into or out of A and B .

Most recent probabilistic approaches to causation are non-reductive. The reason for this is that in order to assess whether C is a cause of E , we must look at the conditional probabilities $Pr(E|C \& K)$ and $Pr(E|K)$, where K includes common causes of C and E . If we cannot specify which factors must be included in K in non-causal terms, then we will not be able to analyze the claim that C causes E into probabilities without causal remainder.

Probabilistic approaches to causation have problems discriminating genuine causes from pre-empted backups. Suppose, for example, that the connection between the city's power grid and a particular building is faulty, so that the building might fail to

receive electricity even when the power grid is otherwise running properly. Then the presence of the backup generator might raise the probability that the lights will be on in the building, even when we hold fixed the functioning of the power grid. Yet on a given occasion it might still be the power grid, rather than the backup generator, that is powering the lights. In such a case, probabilistic approaches to causation would incorrectly rule that the backup generator is also causing the lights to be on.

Counterfactual theories of causation

Counterfactual approaches to causation take from jurisprudence the central idea that causes are conditions *sine qua non* for their effects. In other words, when *C* causes *E*, then the counterfactual conditional “If *C* had not occurred, *E* would not have occurred” is true. This counterfactual then becomes the test for causation. According to the standard possible-world semantics for counterfactuals, this counterfactual will be true just in case there is at least one possible world in which *C* does not occur and *E* does not occur that is closer to the actual world than any possible world in which *C* does not occur but *E* does occur. In other words, the counterfactual will be true just in case *E* does not occur in the closest possible worlds in which *C* does not occur. Thus, to specify the truth-values of counterfactual claims, it is necessary to specify the metric that determines the relative closeness of possible worlds.

Suppose that as a matter of accident, conjunctions of events of type *ABCD* ... are always followed by events of type *E*, while conjunctions of events *ABD* ... without *C* are not. Now consider one particular incident in which a conjunction of events of type *ABCD* ... occurs, and is followed by an event of type *E*. In this case, *C* is not a genuine cause of *E*. Consider the counterfactual “If *C* had not occurred, then *E* would not have occurred.” In order for this counterfactual to be true, the closest not-*C* worlds where *E* does not occur would have to be closer to actuality than any not-*C* worlds where *E* does occur. The (not-*C*, not-*E*) worlds might seem to be further from actuality than the (not-*C*, *E*) worlds, because the (not-*C*, not-*E*) worlds differ from the actual world with respect to the occurrence of *E*, while the (not-*C*, *E*) worlds do not. But there is another sense in which the (not-*C*, not-*E*) worlds might seem to be closer to actuality: in these worlds, the conjunction *ABD* ... is not followed by *E*. In order to avoid the conclusion that *C* is a cause of *E*, the relevant metric of similarity must put more weight on similarity with respect to the occurrence of *E* than on similarity with respect to accidental generalizations. On the other hand, if the connection between *C* and *E* is lawful, then the closest worlds in which *C* fails to occur and *E* occurs anyway would involve a violation of the laws of the actual world, and this sort of difference would be accorded a much greater significance. Indeed, the ability to support counterfactuals is often taken to be a feature that distinguishes genuine laws from accidental generalizations.

In order to capture the directionality of causation, the relevant counterfactuals must themselves be directional in the appropriate way. Suppose, for example, that Julian smokes, and as a result his teeth become stained, and he develops lung cancer. Then it seems plausible to say that if he had not smoked, he would not have stained

teeth and he would not have lung cancer. These counterfactuals correctly entail that Julian's smoking caused his stained teeth and his lung cancer. But we must not say that if Julian did not have stained teeth, it would have to be because he did not smoke, and hence he would not have had lung cancer either. If counterfactuals are allowed to *back-track* in this way, then our counterfactual criterion will rule that *C* is a cause of *E* when in fact *C* is an effect of *E* or *C* and *E* are effects of a common cause. One challenge, then, is to provide an account of the metric of similarity over possible worlds that preserves this directionality. If this cannot be done in non-causal terms, then it will not be possible to provide a reductive analysis of causation in terms of counterfactuals.

Counterfactual theories of causation face problems with pre-emption. Unlike regularity and probabilistic theories, the problem is not that counterfactual theories judge pre-empted backups to be causes, but rather that they fail to recognize pre-empting causes. Suppose, for example, that the city's power grid is functioning properly, causing the lights in the building to be on. Now it is false that if the power grid were not functioning properly, the lights would not be on; for if the power grid were not functioning, the backup generator would come on. There are a number of attempts to rescue the counterfactual approach to causation from the problem of pre-emption: this is currently a lively area of research.

Manipulability theories of causation

Manipulability approaches to causation take as their point of departure the idea that causes are means for producing their effects. This means that agents can exploit the link between *C* and *E* as a handle for bringing about *E*. Agents are not merely passive observers, but intervene in the normal course of nature to bring about events that would not otherwise have occurred. The relationship between *C* and *E* can be used as a means for producing *E* only if it remains stable under this sort of intervention. Suppose, for example, that *E* is in fact a cause of *C*, rather than vice versa. It may well be that events of type *C* are typically accompanied by events of type *E*. Nonetheless, if an agent were to intervene in order to produce an event of type *C*, we would no longer expect it to be accompanied by its usual cause *E*. This is because the intervention is by itself sufficient to produce *C*; it breaks the customary link between *C* and *E*. Similarly, if *A* and *B* are both effects of a common cause *C*, we would not expect that an intervention to produce *A* would result in the occurrence of *B*. Once again, the intervention breaks the link between *A* and its usual cause *C*. Similarly, if the relationship between *C* and *E* is accidental, there would be no reason to expect that a novel event of type *C* produced by an intervention would be accompanied by an event of type *E*.

One worry is that this account makes reference to the interventions of an agent. This might seem to make the account of causation too anthropocentric: what of causal relationships where intervention is not practicable or even possible; for instance, causal relationships in astrophysics or in the early universe? While reference to the actions of an agent is a useful heuristic, it is possible to characterize the relevant notion of intervention

without making reference to human beings or other agents. The important feature of an intervention is not its origin in the intentions of an agent, but rather its status as an independent cause that overrides the customary causal mechanisms for the production of *C*. The notion of an intervention is itself a causal notion, hence an account of causation in terms of interventions will be non-reductive.

Manipulability approaches to causation face problems with pre-emption in much the same way that counterfactual theories do. It may be that the city's main power grid is causing the lights to be on in a certain building, even though, due to the presence of the backup generator, the lights cannot be controlled by intervening on the city's power grid. Many of the strategies that have been proposed for counterfactual theories to deal with this problem may be adapted for manipulability theories as well.

Difference-making

All four approaches to causation discussed above share a common idea: causes are difference-makers for their effects, in the sense that the cause makes a difference to whether or not the effect occurs. The various approaches differ over precisely how the notion of *making a difference* is to be understood. According to regularity theories, the presence or absence of the cause *C* makes a difference for whether the effect *E* regularly follows from the conjunction of additional factors *ABD*... According to probabilistic theories of causation, the presence or absence of the cause *C* makes a difference to the probability of the effect *E*. In the counterfactual framework, the presence or absence of the cause *C* in nearby possible worlds makes a difference to whether the effect *E* occurs in those worlds. And in manipulability theories, interventions that make *C* occur or fail to occur make a difference to whether or not *E* occurs.

Process theories of causation

Process theories of causation are quite different from the difference-making approaches to causation already described. Instead of focusing on causal relationships between discrete events, process theories focus on continuous causal process. Causal processes include ordinary physical objects like baseballs and automobiles, more esoteric objects like photons and neutrinos, as well as various kinds of waves, such as sound waves and water waves. These processes need to be distinguished from pseudo-processes, such as shadows and spots of light. One important difference between them is that causal processes are restricted by the *first-signal principle* of the special theory of relativity, whereas pseudo-processes are not. For example, if one were to shine a very bright light on the wall of a large circular stadium, it would be possible in principle to rotate the light source so that the spot of light traveled along the wall with a velocity greater than the speed of light. By contrast, no causal process can be accelerated across the speed of light.

A central challenge for process theories of causation is to distinguish between causal processes and pseudo-processes. According to one leading approach, causal processes differ from pseudo-processes in their ability to transmit conserved quantities, such as

energy, linear momentum, and charge. Baseballs, automobiles, photons, neutrinos, and sound waves are all capable of carrying energy from one place to another. Shadows and spots of light are not capable of transmitting conserved quantities. Here the process theorist must take care to distinguish between the *transmission* of a conserved quantity and the mere *presence* of a conserved quantity at various locations. For example, as a spot of light moves along a wall, energy will be present at each point along the wall as it is illuminated. Nonetheless, energy is not *transmitted* from one point on the wall to another; rather the energy is supplied to the various points along the wall from the central source. The spots of light on the wall are related not as cause and effect, but as effects of a common cause. The challenge for the *conserved-quantity theory* is to characterize the relevant notion of transmission in order to make this distinction.

Process theories of causation can easily solve the problem of pre-emption. We know that it is the city's power grid rather than the backup generator that is causing the lights in a building to be on because there are causal processes – electrons, which transmit the conserved quantity charge – that connect the city's power grid to the light sources in the building. There are no analogous processes connecting the backup generator to the lights. On the other hand, process theories offer little that is new to the problem of the direction of causation. If there is a causal process connecting *C* to *E*, then there will be a causal process connecting *E* to *C*. The process theorist can, of course, define the cause to be the earlier of the two events, a strategy that is available to all of the approaches to causation that we have canvassed.

One approach to causation, which is closely related to the process theories, analyzes causal relationships in terms of the *mechanisms* that connect causes with their effects.

Conclusion

It is fair to say that there is no one account of causation that has won the allegiance of the majority of philosophers who have thought about these issues. Nonetheless, sufficient progress has been made that few philosophers today continue to regard the concept of causation with the same suspicion voiced by Russell and his contemporaries.

See also Determinism; Explanation; Laws of nature; Mechanisms; Physics; Probability.

References

-
- Dowe, P. (2000) *Physical Causation*, Cambridge: Cambridge University Press.
- (2004) "Causation: Causal Processes," in Edward N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy* (winter 2004 edition); available: <http://plato.stanford.edu/archives/win2004/entries/causation-process>.
- Eells, E. (1991) *Probabilistic Causality*, Cambridge: Cambridge University Press.
- Gasking, D. (1955) "Causation and Recipes," *Mind* 64: 474–87.
- Hitchcock, C. (2001) "The Intransitivity of Causation Revealed in Equations and Graphs," *Journal of Philosophy* 98: 273–99.
- (2002) "Causation: Probabilistic," in Edward N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy* (fall 2002 edition); available: <http://plato.stanford.edu/archives/fall2002/entries/causation-probabilistic>.

- Hume, D. (1977 [1748]) *Enquiries Concerning Human Understanding and Concerning the Principles of Morals*, ed. L. A. Selby-Bigge, 3rd edn, rev. P. H. Nidditch, Oxford: Clarendon Press.
- (1978 [1739–40]) *A Treatise of Human Nature*, ed. L. A. Selby-Bigge, 2nd edn, rev. P. H. Nidditch, Oxford: Clarendon Press.
- Lewis, D. K. (1973) "Causation," *Journal of Philosophy* 70: 556–67; reprinted with postscripts in Lewis, *Philosophical Papers*, Oxford: Oxford University Press, 1986, Volume 2, pp. 159–213.
- Mackie, J. (1974) *The Cement of the Universe*, Oxford: Oxford University Press.
- Menzies, P. (1989) "Probabilistic Causation and Causal Processes: A Critique of Lewis," *Philosophy of Science* 56: 642–63.
- (2001) "Causation: Counterfactual Theories," in Edward N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy* (spring 2001 edition); available: <http://plato.stanford.edu/archives/spr2001/entries/causation-counterfactual>.
- Mill, J. S. (1843) *A System of Logic: Ratiocinative and Inductive*, London: J. W. Parker.
- Reichenbach, H. (1956) *The Direction of Time*, Berkeley and Los Angeles: University of California Press.
- Russell, B. (1913) "On the Notion of Cause," *Proceedings of the Aristotelian Society* 13: 1–26.
- Spirtes, P., Glymour, C. and Scheines, R. (2000) *Causation, Prediction, and Search*, 2nd edn, Cambridge, MA: MIT University Press.
- Woodward, J. (2001) "Causation and Manipulability," in Edward N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy* (fall 2001 edition); available: <http://plato.stanford.edu/archives/fall2001/entries/causation-mani>.
- (2003) *Making Things Happen: A Theory of Causal Explanation*, Oxford: Oxford University Press.

Further reading

Russell's critique of the concept of causation is presented in Russell (1913). Hume presents his account of causation in Book I, Part III of *A Treatise of Human Nature* (1739–40) and in section VII of *An Enquiry Concerning Human Understanding* (1748). Mill presents his regularity theory of causation in Volume I, Chapter V of his *System of Logic* (1843). Mackie (1974: Ch. 3) contains a detailed discussion of Mill's theory, and presents the theory of INUS conditions. Reichenbach (1956, Part IV) is an influential early presentation of a probabilistic theory of causation, including the central idea of screening-off relations. Eells (1991) is a book-length treatment. Spirtes, Glymour, and Scheines (2000) is a very technical but important work on the connection between causation and probability. The problem of pre-emption for probabilistic theories of causation is presented in Menzies (1989). Hitchcock (2002) is a survey article covering probabilistic approaches to causation. Lewis (1973) is the classic presentation of the counterfactual approach to causation. Hitchcock (2001) is one recent attempt to address the problem of pre-emption. Menzies (2001) surveys the counterfactual framework. Gasking (1955) is an early defense of the manipulability theory of causation; Woodward (2003) is a recent book-length treatment while Woodward (2001) is an article-length survey. Dowe (2000) presents the conserved quantity theory of causal processes, and Dowe (2004) is a survey of process theories of causation.

THE
ROUTLEDGE COMPANION
TO PHILOSOPHY OF
SCIENCE

Edited by
Stathis Psillos and
Martin Curd

First published 2008
by Routledge
2 Park Square, Milton Park, Abingdon, OX14 4RN

Simultaneously published in the USA and Canada
by Routledge
270 Madison Ave, New York, NY 10016

This edition published in the Taylor & Francis e-Library, 2008.

“To purchase your own copy of this or any of Taylor & Francis or Routledge’s collection of thousands of eBooks please go to www.eBookstore.tandf.co.uk.”

Routledge is an imprint of the Taylor & Francis Group, an informa business

© 2008 Stathis Psillos and Martin Curd for selection and editorial matter;
individual contributors for their contributions

All rights reserved. No part of this book may be reprinted or reproduced or utilised in any form or by any electronic, mechanical, or other means, now known or hereafter invented, including photocopying and recording, or in any information storage or retrieval system, without permission in writing from the publishers.

British Library Cataloguing in Publication Data

A catalogue record for this book is available from the British Library

Library of Congress Cataloging in Publication Data

The Routledge companion to philosophy of science / edited by Stathis Psillos and Martin Curd.

p. cm. -- (Routledge philosophy companions)

Includes bibliographical references and index.

1. Science--Philosophy. I. Psillos, Stathis, 1965- II. Curd, Martin.

Q175.R68 2008

501--dc22

2007020000

ISBN 0-203-00050-1 Master e-book ISBN

ISBN10: 0-415-35403-X (hbk)

ISBN10: 0-203-00050-1 (ebk)

ISBN13: 978-0-415-35403-5 (hbk)

ISBN13: 978-0-203-00050-2 (ebk)