

CHAPTER 9

PROBABILISTIC THEORIES

JON WILLIAMSON

1. INTRODUCTION

Causal relationships are typically accompanied by probabilistic dependencies—normally when *A* causes *B* the former raises or lowers the probability of the latter. Probabilistic theories of causality usually try to characterize or analyse causality in terms of these probabilistic dependencies: they try to provide probabilistic criteria for deciding whether *A* causes *B*, and often maintain that causality just *is* the corresponding pattern of probabilistic relationships. This chapter provides an introduction to and criticism of such accounts. While it will be argued that probabilistic theories are ultimately unsuccessful, work on probabilistic causality has shed a great deal of light on the relationship between causality and probability and hence these theories repay a thorough understanding.

The chapter is organized as follows. In sect. 2 some key distinctions are introduced: these are helpful for categorizing theories of causality. The idea of a probabilistic theory of causality is discussed in sect. 3; then some of the key early theories are presented: Reichenbach's in sect. 4, Good's in sect. 5, and Suppes's in sect. 6. More recently, major steps have been taken by researchers in Artificial Intelligence and these have led to the causal net formalism (sect. 7), perhaps the most important contemporary probabilistic theory. The influence of this formalism is sketched in sect. 8. It is argued in sect. 9 that probabilistic theories flounder because they admit counterexamples, and because they fail to accommodate the important connection between causality and physical mechanisms (sect. 10). In sect. 11 the epistemic theory of causality is introduced. This theory, I claim, allows one to take advantage of the insights of probabilistic causality while avoiding its pitfalls.

This work has been produced with the aid of grants from the Leverhulme Trust and the British Academy. I am grateful to Chris Hitchcock, Phyllis McKay Illari, and Federica Russo for very helpful comments.

2. CATEGORIZING PHILOSOPHICAL THEORIES OF CAUSALITY

There are a wealth of theories of causality and these theories can be categorized according to the way they answer a range of key questions.

Some questions concern the causes and effects that are related by causality. First, are the

causal relata single-case or generic? A philosophical theory of causality might hold that a cause or effect concerns only a single occasion and so either obtains or fails to obtain: for example, *Audrey's letter will cause Balthasar anguish when he reads it*. Or it might hold that causes and effects can obtain and fail to obtain on different occasions: *smoking causes cancer*. In the former case, causes and effects are called *single-case, particular, or token-level*; in the latter case they are *generic, repeatably instantiatable, or type-level*. Second, are the causal relata *population-level* or *individual-level*? In the former case a cause or effect concerns a group of individuals: *an increase in inequality of wealth in Britain in 2007 caused a reduction in happiness* (note that cause and effect are single-case here). On the other hand, an individual-level cause or effect concerns only one individual at a time: *viral infection causes an immune response* (cause and effect are generic in this example). Of course *all* these kinds of causal relata occur in our causal claims, apparently without any great problem, so any theory that considers one kind to the exclusion of the other kinds provides only a partial account of causality.

Other questions concern the causal relation itself. First, is causality some kind of *physical* connection between cause and effect? Or is it purely *mental* in the sense that it is a feature of some (possibly idealized) individual's epistemic state? Second, is the causal relation *objective*, in the sense that if two agents disagree as to causal relationships, at least one of them must be wrong, or is it *subjective*, admitting a degree of personal choice? Third, does the theory in question attempt to understand actual or potential causality? A golf ball bouncing off a tree is in general a preventative of it going into the hole, though it may, in fact, cause it to go into the hole. The general case is sometimes known as *potential causation* or *possible causation* while the factual case is called *actual causation*.

While these clearly do not exhaust the questions one might ask, they are often viewed as central questions in the debate about causality. It is certainly useful to be aware of the distinctions they draw when trying to understand an unfamiliar philosophical theory of causality.

3. PROBABILISTIC THEORIES OF CAUSALITY

Most probabilistic theories of causality are motivated by the following central intuitions: (1) changing a cause makes a difference to its effects, and (2) this difference-making shows up in probabilistic dependencies between cause and effect. (There are exceptions, e.g. the early Reichenbach as outlined in sect. 4.) Many proponents of probabilistic theories will go further by maintaining that probabilistic dependencies *characterize* the causal relation, i.e. provide necessary and sufficient conditions for causal connection, of the form: *C causes E if and only if* appropriate probabilistic dependencies obtain. They often go further still by maintaining that the probabilistic dependencies *analyse* the causal relation: '*C causes E*' *just means that* the corresponding probabilistic dependencies obtain.

When characterizing the causal relation the probabilistic dependencies may themselves be formulated using causal terms, but in an analysis of causality any reference to causal terms in the probabilistic conditions should in principle be eliminable. With a probabilistic characterization of causality, a characterization of probability must also be provided, and to

complete a probabilistic analysis of causality, one needs an analysis of probability. This gives rise to questions about probability: is probability physical or epistemic? Does probability arise from indeterminism? Indeed all the questions of sect. 2 have probabilistic analogues, and when assessing probabilistic theories of causality these probabilistic analogues need to be settled in order to answer the causal questions.

The key concern is to come up with a set of probabilistic conditions that yields a plausible probabilistic theory of causality. In the following sections we shall consider several important attempts to do just that. A number of specific problems have been raised in the literature for each of the attempts given; rather than wade through a slurry of technical concerns we shall focus, in sects. 9 and 10, on two very general problems that face probabilistic theories.

4. REICHENBACH

Russell (1913) offered a compelling critique of the notion of cause, arguing that the folk notion is incoherent in several respects and that the fundamental sciences do not appeal to causal terms but instead to functional equations. Following attacks on causal language by Mach (1905) and Pearson (1892, 3rd edn. 1911), this created a context of some wariness towards causality (cf. sect. 8).

Thus Reichenbach (1959, first published 1923) went against the grain by taking causal relationships seriously and interpreting them as the lawlike functional equations of physics. He argued against aprioristic and conventional views of causality, and in favour of his ‘probabilistic conception of causality’ (ibid. }6–8). It should be emphasized, though, that Reichenbach’s probabilistic conception of that book was not a metaphysical account of causality but instead an account of the epistemology of causality. Since most probabilistic accounts of causality are intended as metaphysical accounts, and since Reichenbach did at a later stage offer a metaphysical probabilistic account of causality, this is apt to cause some confusion.

The epistemological account of Reichenbach (1959) proceeded along the following lines. Causal discovery does *not* take the form: ‘we notice that certain laws hold in *particular* instances, and we infer that the laws will hold in *all* instances’ (ibid. 131). This is because we cannot suppose causal laws obtain in particular instances without other causal knowledge: ‘if the probability that causality holds in a specific case could not be independently established for at least one case, it could not be increased by an appeal to other cases. Such an appeal presupposes the probability that causality holds has been determined for these other cases’ (ibid. 132). *No causes in, no causes out*, to use a contemporary slogan. Instead, ‘the correct inference has the following form: since we have observed that the same function governs a finite number of observations, we conclude that it governs *all* observations’ (ibid.). Reichenbach’s account is probabilistic in the sense that it attaches probabilities to functional equations and thus to causal laws, not in the sense that the causal laws themselves are reducible to probability relationships.

The approach of the later Reichenbach (1971, first published 1956) is very different in that it attempts to analyse causality itself: ‘in the present book, I wish to study the cause–effect relation in itself; that is, to treat it no longer as primitive, but to reduce it to other relations’

(ibid. 25). Causality is no longer viewed just as a functional relationship, because functional equations are symmetric in cause and effect while causal relations are asymmetric (ibid. 28). Reichenbach (ibid. sect. 19) attempts to analyse the direction of time in terms of the direction of causality, and appeals the following Principle of the Common Cause to determine the direction of causality:

If coincidences of two events A and B occur more frequently than would correspond to their independent occurrence, that is, if the events satisfy relation $[P(AB) > P(A)P(B)]$, then there exists a common cause C for these events such that the fork ACB is conjunctive, that is, satisfies relations $[A \perp\!\!\!\perp B|C, P(A|C) > P(A|\bar{C}), P(B|C) > P(B|\bar{C})]$. (ibid. 163)

In this quotation, as elsewhere in this chapter, I have substituted contemporary mathematical notation for the original notation: $A \perp\!\!\!\perp B | C$ means that A and B are probabilistically independent, conditional on C . The latter relations imposed by Reichenbach imply that $P(AB) > P(A)P(B)$, so the common cause can be said to explain the dependence between A and B .

Reichenbach's key idea was that the Principle of the Common Cause should be used to determine the direction of causality and thus the direction of time: where there is a fork ACB (see [Fig. 9.1](#)) such that the relations in the Principle of the Common Cause hold, and where there is no other C' that satisfies these conditions with respect to A and B , then C is the common cause of A and B and is earlier than A and B .

Reichenbach (1971: §22) extends the analysis to deal with causal betweenness: C is *causally between* A and B (as in [Fig. 9.2](#)) if $1 > P(B|C) > P(B|A) > P(B) > 0$, $1 > P(A|C) > P(A|B) > P(A) > 0$ and $A \perp\!\!\!\perp B | C$.

The aim is that these probabilistic conditions can be used to determine a complete causal graph. Here 'determine' is used in a metaphysical rather than epistemological sense: the conditions are not a description of how one ought to discover causal relationships but rather they offer a characterization and analysis of the causal relation. According to this characterization, a cause raises the probabilities of its direct effects, and no other event renders the cause and a direct effect probabilistically independent.

Reichenbach had developed his own frequency theory of probability and these frequencies were supposedly non-causal, so causality is ultimately analysed in terms of frequency relations.¹ In terms of the distinctions of sect. 2, causality is thus a physical and objective relation; since frequencies handle generic attributes, Reichenbach's analysis deals only with generic relata and potential causation (leaving open the question as to whether the causal relata are population-level or individual-level).




Fig. 9.1 A fork or v-structure.




Fig. 9.2 C is causally between A and B.

Note that Reichenbach represented causal connections graphically, using diagrams like those depicted in [Figs. 9.1](#) and [9.2](#). He called these diagrams *causal nets*. [Figure 9.2](#), which represents a sequence of events with a direct causal connection from one event to the next, is called a *causal chain*. This graphical representation has become an important feature of recent work on probabilistic causality, as will be explained in sect. 7.

Reichenbach's probabilistic analysis of causality was criticized by Salmon (1980: sect. 2), although in later works Salmon (1984; 1998) went on to develop the mechanistic ideas to be found in Reichenbach's writings. See Martel (2000) for a defence of Reichenbach's approach.

5. Good

Good (1959: 307) attempted to provide an analysis of causality in terms of Popperian propensity, a kind of physical probability.² As with Reichenbach, Good avoids reference to time in order to leave open the possibility of analysing time in terms of causality. For an event F to be a cause of event E , Good (1959) requires that (a) E and F both occur; (b) $P(E|H) > 0$ and $P(F|H) > 0$ where H consists of all the laws of nature (H_1) and all the true background conditions that are taken for granted (H_2); (c) $P(E|FH) > P(E|\bar{F}H)$; (d) FH_2 does not logically imply E ; (e) there is an event G that did not occur, that could have prevented F , $P(\bar{F}|GH) \approx 1$, and whose absence would not ensure F , $P(F|\bar{G}H) > 0$; and (f) $P(E|GH) \approx P(E|\bar{F}GH) \approx P(E|\bar{G}\bar{F}H) \approx P(E|\bar{F}H)$. Good ultimately rejected this account on the grounds that conditions (e) and (f) in fact hold trivially in the presence of the other conditions (a–d). This led Good to seek a better probabilistic account.

Good (1961a; 1961b) provided an updated account. Here Good no longer avoids appealing to time in the analysis of causality, explicitly requiring that a cause be earlier than its effects. This account analyses causality in terms of physical probability, but it goes beyond his previous account in that it attempts to give quantitative measures of both potential and actual causation. The (potential causal) *tendency of F to cause E* is measured by

$$\log \frac{P(\bar{E}|\bar{F}H)}{P(\bar{E}|FH)}$$

where H consists of all laws of nature and the background conditions before F started. Thus for F to be a potential cause of E , the two must be probabilistically dependent conditional on H . The (actual causal) *degree to which F caused E* is the limit, as the sizes of the events tend uniformly to zero, of the strength of the network of causal connections between E and F . Here the strength of a link from F to E is measured by the tendency of F to cause E ; the strength of the network as a whole is a function of these link strengths which takes into account interactions amongst causes.

Good quite explicitly develops the notion of a *causal net*, which is represented by a directed acyclic graph whose nodes are events and whose arrows signify causal connections from cause to effect and chart the probabilistic dependencies that obtain amongst the events. Good's approach is devised to be able to cope with token events, and clearly construes causality as an objective, physical relation. Good preferred his own account over Reichenbach's on the grounds that Reichenbach's account turns out to be vacuous: for event F to be a cause of event E , Reichenbach requires that F raise the probability of E and that there is no event G that renders E and F probabilistically independent; Good (1961b: app. 1) argued that one can always gerrymander some event G that renders the two variables probabilistically independent, in which case there are no causal relations at all under Reichenbach's account.

Good's account of causality is criticized in Salmon (1980: }1 and 1988).

6. SUPPES

Unlike the later Reichenbach but like the later Good, Suppes (1970) appeals to time in his analysis of the causal relation. Suppes (1970: 12) has three central steps to his analysis. First he introduces the notation A_t to signify that event A occurs at time t , and gives a preliminary definition: 'the event $B_{t'}$ is a prima facie cause of the event A_t if and only if (i) $t' < t$, (ii) $P(B_{t'}) > 0$, (iii) $P(A_t|B_{t'}) > P(A_t)$ '. Suppes (ibid. 25) then defines a prima facie cause $B_{t'}$ of A_t to be a *spurious cause* of A_t if there is a prior partition $\pi_{t''}$ of events that screens off $B_{t'}$ from A_t : for all elements $C_{t''}$ of $\pi_{t''}$, (i) $P(B_{t'}|C_{t''}) > 0$, (ii) $P(A_t|B_{t'}|C_{t''}) = P(A_t|C_{t''})$. Finally, a *genuine cause* is a prima facie cause that is not spurious.

Suppes goes on to define *direct cause* and *negative cause*, and then extends his account to deal with causal relations between variables, not just events. Suppes is clear that he intends his account as an analysis of causality, rather than merely a characterization of the causal relation.

He leaves open the interpretation of probability (ibid. }2, [pp. 79–80](#)) and intends that the account apply to both generic and single-case relata (ibid. 79). Suppes does not think that causality is entirely objective, in the sense that the causal relation is not uniquely determined. First, causal relationships are relative to the conceptual framework (the set of events or variables) under consideration (ibid. 75). Second, ‘the analysis of causes is always relative to a particular conception of mechanism, and it does not seem satisfactory to hold that the analysis of mechanism is ever complete or absolute in character’ (ibid. 72). For example, the analysis depends on whether time is continuous or discrete. While Suppes discusses mechanisms in some detail, he is not explicit about the way mechanisms fit into the analysis. The connection with mechanisms would suggest that Suppes considers the causal relation to be a physical relationship; however, the fact that he admits the possibility of a mental interpretation of probability suggests that a mental notion of cause may also be adopted. In sum, while Suppes offers a single formal reduction of causality to probabilities, he is a pluralist about causality—causality varies according to the interpretation of probability, the conceptual framework, and the notion of mechanism under consideration.

See Salmon (1980: §3) for critical discussion of Suppes’s probabilistic account of causality.

Note that Salmon (1984: 190) explicitly advocates ‘probabilistic causality’, but under Salmon’s view causality is probabilistic only in the sense that it is not deterministic (*not* in the sense that causal relations can be characterized or analysed in terms of probabilistic relations):

I cannot think of any reason to suppose that ordinary causal talk would dissolve into nonsense if Laplacian determinism turned out to be false. I shall therefore proceed on the supposition that probabilistic causality is a coherent and important philosophical concept.

In advocating the notion of probabilistic causality, neither Suppes nor I intend to deny that there are sufficient causes; indeed, Suppes explicitly introduces that concept into his theory (1970, [p. 34](#), def. 9). On our view, sufficient causes constitute a limiting case of probabilistic causes. On the sufficiency/necessity view, which we reject, this limiting case includes all bona fide cause–effect relations. The latter approach to causality seems needlessly restrictive.

Thus Salmon’s ‘probabilistic causality’ is a reaction to Mackie’s account of causality in terms of *inus* conditions (which holds that a cause is an Insufficient but Necessary component of a condition that is Unnecessary but Sufficient for an effect—see Ch. 7 on regularity theories of causation in this volume).

Suppes’s account was developed by Cartwright (1979) who put forward the principle:

CC: C causes E iff $P(E|CK) > P(E|K)$ for all states K of the E ’s other causes that are not between C and E .

Cartwright did not take this to be an analysis of causality because cause appears on both sides of ‘iff’. The *context unanimity* of this condition was relaxed by Skyrms (1980: 108–9) who claimed that a cause need only raise the probability of the effect relative to *some* state K , rather than all such states (though there should be no state K for which C lowers the probability of E). Eells and Sober (1983) discussed Cartwright’s condition and argued that this

kind of condition only holds of generic causal claims, not of single-case claims. Eells (1991) provided a detailed defence of generic causality using this type of probabilistic account, and went on to provide a probabilistic theory of single-case causality, appealing to *probability trajectories*, which trace the evolution of single-case probabilities over time.

7. CAUSAL NETS

The 1980s saw a very influential mathematization of the notion of cause. This line of work stemmed from efforts among Artificial Intelligence researchers to automate reasoning in the face of uncertainty. In the 1970s so-called *expert systems* were developed in order to automate reasoning tasks that would normally require human expertise. For example, expert systems were developed for the diagnosis of intestinal problems and for mineral prospecting. Early expert systems tended to take the form of *rule-based systems*: the expert knowledge was encoded by logical rules and on inputting a set of facts (e.g. a patient's symptoms) the rules would be applied to generate inferences (e.g. diagnoses). It was soon realized that rules for diagnosis and other common expert tasks are rarely exceptionless, and that expert systems needed a way of handling this uncertainty. One way is to invoke the notion of probability: a particular pathology is more or less probable given a patient's symptoms, though rarely certain or certainly ruled out. In lieu of tractable computational procedures for handling probabilities, several non-probabilistic formalisms for handling uncertainty were developed. Then, in the 1980s, a formalism was developed which at one fell swoop offered the possibility of tractably representing and reasoning with probabilities on the one hand, and representing and reasoning with causal connections on the other. This is the formalism of Bayesian networks (Pearl 1988; Neapolitan 1990).

A Bayesian network consists of a directed acyclic graph whose nodes are variables in the domain of interest, together with the probability distribution of each variable conditional on its parents in the graph (or its unconditional distribution if the variable has no parents). The graph and the probabilities are tied together by a fundamental assumption known as the *Markov condition*: each variable is probabilistically independent of its non-descendants conditional on its parents in the graph, written $A_i \perp\!\!\!\perp ND_i \mid Par_i$ for each i . Then one can calculate any probability involving the variables in the domain via the identity $p(a_1 \dots a_n) = \prod_{i=1}^n P(a_i \mid par_i)$ where each a_i is an assignment of a value to variable A_i and par_i is the assignment of values to its parents Par_i consistent with a_1, \dots, a_n .

For example, suppose there are four two-valued variables, A, B, C, D . A Bayesian net can be formed by taking the directed acyclic graph of [Fig. 9.3](#) and specifying the probability distribution of each variable conditional on its parents, e.g.

$$\begin{aligned}
P(a^1) &= .3, P(a^0) = .7 \\
P(b^1) &= .9, P(b^0) = .1 \\
P(c^1|a^1b^1) &= .4, P(c^0|a^1b^1) = .6 \\
P(c^1|a^0b^1) &= .5, P(c^0|a^0b^1) = .5 \\
P(c^1|a^1b^0) &= .1, P(c^0|a^1b^0) = .9 \\
P(c^1|a^0b^0) &= .2, P(c^0|a^0b^0) = .8 \\
P(d^1|c^1) &= .8, P(d^0|c^1) = .2 \\
P(d^1|c^0) &= .7, P(d^0|c^0) = .3
\end{aligned}$$

Then for instance

$$\begin{aligned}
P(a^1b^0c^1d^1) &= P(a^1)P(b^0)P(c^1|a^1b^0)P(d^1|c^1) \\
&= .3 \times .1 \times .1 \times .8 = .0024.
\end{aligned}$$

In recent years a whole host of techniques has been developed for efficiently calculating probabilities from a Bayesian net and for constructing a Bayesian net to match the probability distribution of a dataset.




Fig. 9.3 An example of a directed acyclic graph.

Bayesian nets are often used to represent and reason with causal relationships. A *causally interpreted Bayesian net* or *causal net* is a Bayesian net in which the arrows of the graph are interpreted as denoting direct causal relationships. Thus for example under a causal interpretation [Fig. 9.3](#) says that *A* and *B* cause *C* and also cause *D* indirectly via *C*. Under a causal interpretation, the Markov Condition—now called the *Causal Markov Condition*—says that each variable is probabilistically independent of its non-effects conditional on its direct causes. It is normally assumed that if the graph in the net correctly portrays the causal relationships amongst the variables, and no causally relevant variable is omitted, then the Causal Markov Condition must hold, though in Williamson (2005) I argue that this condition can at best be justified as a *default* assumption.

Causal nets can be invoked as a characterization of the causal relation as follows. Suppose physical reality is conceptualized using some domain *V* of variables, and that P^* is the physical probability function (propensity or frequency function) over this domain. Let C^* be

the smallest directed acyclic graph on V that satisfies the Markov Condition with respect to P^* . Then, it may be claimed, C^* characterizes the causal relationships. This type of characterization is invoked by many advocates of causal modelling, e.g. Spirtes, Glymour, and Scheines (1993) and Pearl (2000), and is very close to those developments of Suppes's account which do not invoke context unanimity (sect. 6). Often restrictions are made on V (e.g. that V should pick out all and only the physical events) and further assumptions are made concerning P^* (e.g. that all the independencies of P^* are representable by some directed acyclic graph; that P^* is defined over interventions as well as observations) in order to ensure that C^* is independent of V and is uniquely determined. The details of this formalism can be found in Ch. 14 on causal modelling in this volume.

Some proponents of this kind of characterization go further by using it to analyse the causal relation itself (see e.g. Spohn 2002). Under this view, the causal relation *just is* the relation picked out by C^* : causal relationships are a chart of the independencies satisfied by physical probability.

8. CAUSAL TALK

Probabilistic causality—in particular the causal net formalism—provides a mathematical calculus for handling causality that has been enormously influential. Prior to the late twentieth century ‘cause’ was a dirty word to most research scientists, who, since the time of Pearson and Russell, had been warned that correlation need not imply causation and that causality is a muddled notion devoid of mathematical treatment. But the increasing mathematization of causality that accompanied the rise of probabilistic causality began to change all this. The causal net formalism reached a wide audience with the books of Pearl (1988) and Neapolitan (1990) and this dissemination helped bring causal talk back in vogue.

Evidence of this transition can be obtained by analysing titles of books and research papers. [Figure 9.4](#) shows a decline in causal talk in the nineteenth century and into the twentieth century: the bars show the proportion of published books in the English language with a word in the title beginning with ‘caus-’ (from the British Library database). The proportion is of course very small, so percentage figures are omitted, but the general downward trend can be observed. It can be seen, though, that this trend lasts only until the 1970s. [Figure 9.5](#) shows the period since the 1970s in detail: here the bars show the proportion of published research papers with a word in the title beginning with ‘caus-’ (from the Web of Science database). This graph shows a steady increase in causal talk.

While this rise might be attributable to the influence of probabilistic causality and causal nets, another possible explanation of this phenomenon is that the rise is due to an increase in the number of papers in medicine, a subject that is inherently causal.³ This is a plausible hypothesis because, as Bauer (1998) shows, from the 1930s to the 1990s there were more articles about physics than medicine in quality newspapers, but in the 1990s medicine overtook physics. Perhaps, then, more research is being done in medicine, which uses more causal talk. But [Fig. 9.6](#) shows that while the portrayal of science by the media may have changed around 1990, the balance of science did not itself change. Here the black bars show the numbers of papers in journals whose titles contain a word beginning with ‘med-’ while the

shaded bars show the number of papers in journals whose titles contain a word beginning with ‘phys-’. While the numbers of papers are growing year by year, it is clear that research in medicine is not overtaking research in the physical sciences. [Figure 9.7](#) shows the same phenomenon, but with ‘bio-’ journals compared to ‘phys-’ journals.




Fig. 9.4 Proportion of books whose title includes causal terms.

In sum, causal talk is coming back into vogue and this is no doubt at least partly attributable to the rise of probabilistic causality and causal net modelling.

9. COUNTEREXAMPLES TO PROBABILISTIC CAUSALITY

In this section and the next we shall take a look at two general kinds of worry about the probabilistic approaches to causality outlined above. In this section we shall see that the connections between probability and causality that are invoked by the above theories admit counterexamples. In sect. 10 it is argued that physical mechanisms are neglected by the above accounts but are crucial to our understanding of causality.




Fig. 9.5 Proportion of papers whose title includes causal terms.




Fig. 9.6 Numbers of medical (black) and physical (shaded) papers.




Fig. 9.7 Numbers of biological (black) and physical (shaded) papers.

Note that the Causal Markov Condition implies the following version of the Principle of the Common Cause:

PCC: If variables A and B are probabilistically dependent then one causes the other or there is a set U of common causes in V which screen off A and B , i.e. render them probabilistically independent, $A \perp B \mid U$.

This version of the Principle of the Common Cause is also a consequence of Reichenbach's own version (under a suitable mapping between events and variables) and of developments of Suppes's account that do not invoke context unanimity.⁴

Unfortunately, PCC is oversimplistic. It says that any probabilistic dependence can be fully accounted for by causal connections. In fact, though, probabilistic dependencies may be

attributable to other kinds of relationships between the variables. A and B may be dependent not because they are causally related but because they are related logically (e.g. where an assignment to A is logically complex and logically implies an assignment to B), mathematically (e.g. mean and variance variables for the same quantity are connected by a mathematical equation), or semantically (e.g. A and B are synonymous or overlap in meaning), or are related by non-causal physical laws or by domain constraints. In such cases there may be no common cause to accompany the dependence, or if there is, the common cause may fail fully to screen off A from B . To take a simple example, if a logically implies b then $P(b|a) \geq P(b)$ while $P(b)$ may well be less than 1. In such a case variables A and B (where A takes assignments a and $\emptyset a$ and B takes assignments b and $\emptyset b$) are probabilistically dependent; however it is rarely plausible to say that A causes B or vice versa, or that they have a common cause. Further examples can be found in e.g. Williamson (2005: §4.2). A typical response to this objection takes the form: as long as you take care to ensure that your variables are not related logically, mathematically, semantically etc. then PCC will hold. But this is to render PCC devoid of content, for it is tantamount to saying: as long as you ensure that there are no probabilistic dependencies that are not attributable to causal relationships, then all dependencies will be attributable to causal relationships. Moreover, this response renders PCC much less useful as an epistemological tool, for it is often very hard to tell whether a dependency is attributable to a non-causal relationship.

PCC also hinges on the particular physical interpretation of probability that is adopted. Under an *actual frequency* interpretation of probability, probabilistic dependencies may be entirely accidental, having no underlying explanation and in particular no causal explanation. For example, consider observations of vehicles at a particular road junction today; it may be the case that the proportion of green vehicles that turn right happens to be larger than the proportion of red vehicles that turn right. These proportions are probabilities and hence the direction of turn and the colour of vehicle are probabilistically dependent. This dependency is not attributable to a common cause—it is accidental and in no need of explanation at all. In order to save PCC one might switch interpretation of probability, perhaps maintaining that if one were to have examined vehicles indefinitely today then direction of turn and colour would be independent with respect to *limiting relative frequency*. But—aside from reservations one might have about the move to a counterfactual account of physical probability—it is a purely metaphysical assumption that accidental dependencies will disappear in the limit, and by no means a plausible assumption if one considers time series (Yule 1926; Sober 1988, 2001; Reiss 2007). As Sober observes, British bread prices and the Venetian sea level are correlated, not because they are causally connected but simply because they are both increasing over time for quite separate reasons.

10. MECHANISMS AND PROBABILISTIC CAUSALITY

In some cases causal relationships are not accompanied by the raising of probabilities. Consider the following example, adapted from examples of Salmon (1984: 196–202) and Dowe (2000: §II.6) and presented in more detail in Williamson (2005: §7.3). A potentially unstable atom can decay as follows: radioactive isotope a^1 will decay to either b^1 or b^0 each of which will decay to either c^1 or c^0 ; on the other hand isotope a^0 is stable and will not decay.

Moreover c^1 and c^0 can only be obtained by decay from b^1 or b^0 and b^1 and b^0 can only be obtained from a^1 . We have that $P(c^1|b^1) = \frac{3}{4}$ (so $P(c^0|b^1) = \frac{1}{4}$), $P(c^1|b^0) = \frac{1}{4}$, $P(b^1|a^1) = \frac{1}{4}$, $P(b^1|a^0) = \frac{1}{4}$, $P(a^1) = \frac{1}{2}$. Let variable A take assignments a^1 or a^0 , B take assignments b^1 or b^0 , and C take c^1 or c^0 . Then the causal picture is depicted in [Fig. 9.8](#). Clearly b^0 is a potential cause of c^1 (and may be the actual cause of c^1) even though it lowers the probability of c^1 .

Causal relationships need not even be accompanied by probabilistic dependencies. Consider the same example, but where $P(c^1|b^1) = \frac{1}{2}$ and $P(c^1|b^0) = \frac{1}{2}$. In this case B is still a cause of C —indeed the only cause of C —but C and B are probabilistically independent. This is a serious problem for the probabilistic accounts of causality considered above, all of which require that causal relationships be coextensive with certain probabilistic dependencies. If causal relationships are not characterizable in terms of probabilistic relationships, that clearly puts paid to a probabilistic analysis of causality.

In these examples it is not the probabilities that tell us about the causal relationships but rather physical knowledge—knowledge about the physical entities and physical mechanisms that link them. One might think, then, that probabilistic accounts of causality should be abandoned in favour of mechanistic accounts—indeed, this was the line of argument taken by Wesley Salmon (see also Ch. 10 below on mechanistic accounts (‘Causal Process Theories’)).

But this is too quick. Mechanistic accounts fare no better, for there are cases in which it is clearly not mechanisms that tell us about causal relationships, but rather probabilistic dependencies. One important problem for mechanistic theories is to do with *absences*. There can hardly be mechanisms linking non-existent entities, yet absences are prime examples of causes and effects: e.g. missing the ferry causes Celia to be absent from work. Thus causal relationships can occur without mechanisms (though with accompanying probabilistic dependencies). Another important problem for mechanistic theories occurs where there is a mechanism of the appropriate sort but no causal relationship. There are any number of mechanisms linking a beating heart with liver failure, in cases where such failures occur. Many of these mechanisms qualify as causal mechanisms according to contemporary mechanistic theories of causality, yet one would not want to say that a beating heart is a cause of liver failure, because there is no probabilistic dependence between the two. Instead one would cite those events that are mechanistically connected to liver failure *and* which make a difference to the probability of liver failure as the causes of liver failure. A third problem for mechanistic theories is that mechanisms are heterogeneous across the sciences—mechanisms in physics are quite different from mechanisms in economics, for example; a physical-mechanism account is unlikely to cope with causal relations in economics, while an economic-mechanism account is unlikely to cope with causes in physics, and an account which cashes out causes in physics in terms of physical mechanisms and causes in economics in terms of economic mechanisms is unlikely to be able to explain the apparent homogeneity of the causal relation.

In sum, an account that seeks to characterize or analyse causality just in terms of probabilistic dependencies, or just in terms of physical mechanisms, will be inadequate. This conclusion led Hall (2004) and others to argue in favour of a pluralistic account of causality—sometimes causal relationships are probabilistic, in other cases they are mechanistic. But causal pluralism, discussed later in this volume, has its own set of problems (Russo and

Williamson 2007; Williamson 2006a). Pluralism has trouble accounting for the uniformity of our causal talk—there is no apparent ambiguity when I say ‘smoking causes cancer’, and it would not make sense to require that I clarify my claim by saying ‘smoking mechanistically causes cancer’ or ‘smoking probabilistically causes cancer’ since both mechanisms and probabilities are important evidence for this causal claim. This latter point is worth spelling out. If ‘smoking causes cancer’ is to be understood in terms of a probabilistic relationship between smoking and cancer, then there is an epistemological problem: it is hard to explain why, given that there was excellent probabilistic evidence in favour of smoking being a cause of cancer, the causal claim was not generally accepted until a plausible physiological mechanism linking smoking and cancer was discovered (Fisher, for one, argued that a physiological mechanism had to be found before the causal claim could be substantiated). If the causal claim is probabilistic, why should evidence of mechanisms be required over and above evidence of probabilities? On the other hand if the causal claim is mechanistic, it is hard to explain why, were a plausible mechanism known, one would normally still require evidence that the cause made a difference to the effect before the claim could be said to be substantiated. This epistemological problem besets probabilistic, mechanistic, and pluralist theories of causality alike—under a pluralist account ‘smoking causes cancer’ must be given one or other interpretation, but then it is not clear why evidence both of mechanisms and of probabilistic dependence is required.




Fig. 9.8 Radioactive decay example.

One reason why mechanistic evidence is often required over and above evidence of probabilistic dependence is that causal claims need to be explanatory. Causal claims have two kinds of use: an *inferential* use, for making predictions, diagnoses, and strategic decisions, and an *explanatory* use, to give an account of why an effect occurred. In order to be put to the inferential use, it is crucial that a cause make a difference to its effect—that is, that cause and effect are probabilistically dependent—otherwise one could not predict the effect and could not intervene on the cause to produce the effect. But difference-making is not enough for the explanatory use of causal claims. It is not enough when asked ‘why did the effect occur?’ to answer ‘because an event that makes a difference to the effect occurred’, because that is no explanation at all—it still leaves the question ‘why did the event make a difference to the effect?’ In order to explain an event one needs to invoke some kind of theoretical knowledge—facts about the events and their linking mechanisms. Hence causal claims need to be associated, where possible, both with probabilistic dependencies and with mechanisms.⁵

Reichenbach thought that his probabilistic notion of causality is coextensive with a mechanistic account (Reichenbach 1971; Otte 1986) but the above examples show that this is not the case. It is hard to deny that both probabilities and mechanisms are important to our understanding of causality, yet it seems that probabilistic, mechanistic, and pluralist accounts all fail. There remains the thorny question of how to provide a viable account of causality that integrates probabilistic and mechanistic considerations.

11. THE EPISTEMIC THEORY OF CAUSALITY

I will argue that this question can be answered if we interpret causality as an *epistemic relation*. According to this view, causality is to be analysed neither in terms of physical probabilities nor in terms of physical mechanisms, but in terms of an agent's epistemic state. This type of view can be traced to Hume and Kant (Williamson 2005: §9.2) and also Mach:

There is no cause nor effect in nature; nature has but an individual existence; nature simply *is*. Recurrence of cases in which *A* is always connected with *B*, that is, like results under like circumstances, that is again, the essence of the connection of cause and effect, exist but in the abstraction which we perform for the purpose of mentally reproducing the facts. (Mach 1883: 483).

Much of the authority of the ideas of cause and effect is due to the fact that they are developed *instinctively* and involuntarily, and that we are distinctly sensible of having personally contributed nothing to their formation. We may, indeed, say, that our sense of causality is not acquired by the individual, but has been perfected in the development of the race. Cause and effect, therefore, are things of thought, having an economical office. (ibid. 485)

Thus to say that the causal relation is an epistemic relation is to say that causality is a feature of the way we represent the world rather than a feature of the agent-independent world itself.

An epistemic theory of causality can be developed using the following recipe.

First, take an ideal causal epistemology. In particular consider the way that one's evidence (including one's background knowledge) should constrain the causal beliefs that one has. (Here causal beliefs are not to be construed as beliefs *about* causality. Instead a causal belief is a certain *type* of belief, namely a directed relational belief, representable using directed acyclic graphs, and one that is put to the inferential and explanatory uses that typify causal reasoning.) Certain causal belief graphs are compatible with evidence, others are ruled out. The uses to which causal beliefs are put determines this mapping from evidence to a set of possible causal belief graphs. It is this ideal mapping that constitutes the required ideal causal epistemology.

I won't say much here about this mapping. In practice the epistemology of causality is much less controversial than the metaphysics of causality, and I shall just presume that this mapping is well defined. Current methods for causal discovery offer approximations to this ideal causal epistemology—one such can be found in Williamson (2006a: app. A) and there are many others. Presumably as science progresses the approximations to the ideal mapping will improve. What is clear from sects. 9 and 10, though, is that evidence both of probabilistic dependencies and of mechanisms will play a role in constraining the set of viable causal belief graphs. (The question of exactly how these are to be understood—which notion of mechanism

and which physical interpretation of probability is required—may be left to the ideal causal epistemology to decide.)

The second step of the recipe for epistemic causality is to take all empirical facts. By this I mean all facts about physical reality, but not facts about rational epistemology—for example, neither facts about the ideal causal epistemology, nor facts about causality itself (we shall turn to the question of what these facts are shortly).⁶

Third, apply the causal epistemology to this ideal set of evidence. This results in a set of ideal causal belief graphs. If a rational agent had as evidence the ideal evidence set, knew the ideal causal epistemology, and were able to apply the latter to the former then her causal beliefs would be representable by one of these ideal causal belief graphs. This set of ideal causal belief graphs characterizes the causal relation. If this set is a singleton, the causal relation is fully objective, otherwise there is some subjective choice as to what the causal relationships are—the extent of this subjective choice is proportional to the cardinality of the set of ideal causal belief graphs. The facts about causality are just the facts about *all* the graphs in the ideal set. Thus it is fact that *A* causes *B* just if *A* causes *B* in each ideal causal belief graph.

Fourth, extend this characterization to an analysis of causality. Causal relationships just *are* the result of applying the ideal causal epistemology to the ideal evidence set. They are the set of causal beliefs one should have were one to know all physical facts and the ideal causal epistemology and were one able to apply the latter to the former. (Here, as before, it must be emphasized that a causal belief is a kind of belief, a relational belief that is put to the inferential and explanatory uses that are associated with causal claims. It should not be confused with a belief about causality, which is a non-relational belief about the set of ideal causal graphs.) It is thus the uses to which causal claims are put that determines the nature of causality itself.

All this is obviously highly idealized, and while it may satisfy our need for a metaphysical account of causality, the application of an ideal causal epistemology to an ideal evidence set doesn't say much about how we can discover causal relationships in practice. The best we can do, of course, is to apply the causal epistemology of the moment to the evidence of the moment. Quite plausibly, some of our more entrenched causal claims of the moment will remain entrenched as our evidence and causal epistemology improves. Thus there is no reason why much of what we think of now as knowledge of causal relationships should not in fact be such knowledge (i.e. knowledge of the ideal causal belief graphs).

Note that steps 1–3 of this recipe are rather uncontroversial. Proponents of a probabilistic or mechanistic analysis of causality will surely agree with the claim that the application of the ideal causal epistemology to the ideal evidence set will characterize the causal relation—though of course they may differ as to what might constitute the ideal causal epistemology. Indeed proponents of any view of causality in which causal relationships are not radically unknowable will concur. Substantial disagreement only comes at step 4—the analysis. The proponent of a probabilistic/mechanistic analysis holds that '*A* causes *B*' says something about probabilities/mechanisms respectively, while the proponent of the epistemic theory holds that it says something about rational belief.

It is not hard to see how this epistemic view of causality gets round the problems that beset probabilistic causality, mechanistic causality, and a pluralist combination of the two. The

counterexamples to probabilistic causality are not counterexamples to epistemic causality because PCC is not an assumption of epistemic causality. Under the epistemic view, PCC may hold in certain circumstances (Williamson 2005), but it is not guaranteed to hold: if in a particular case one should not posit a common cause to account for a probabilistic dependence, then the ideal causal epistemology, when applied to the ideal evidence set, will not yield a common cause. Similarly, under the epistemic view, cause and effect need not be probabilistically dependent. If, as in the example of sect.10, one should conclude that *B* causes *C* even though the two are probabilistically independent, then by construction of the epistemic theory, it will. Thus problems with probabilistic causality do not carry over to epistemic causality. Turning to difficulties with the mechanistic approach, we see that absences are no problem for the epistemic theory. Since the link between cause and effect is not physical, causes and effects need not be physical entities either. Further, if one should not conclude from the fact that there are certain mechanisms linking a beating heart and a failing liver that the former is causing the latter, then by construction the epistemic theory won't. Heterogeneity of mechanisms across the sciences is no problem because the causal relation is not analysed in terms of those mechanisms but in terms of rational belief, an account that is not specific to particular sciences. So problems with the mechanistic theory are not problems for the epistemic theory. Moving on to pluralism, it is clear that the epistemic approach is not pluralist, so it can account for the homogeneity of causal talk. Finally, the epistemological problems with pluralism and the two monistic accounts cannot carry over to epistemic causality. If the claim that smoking causes cancer requires both probabilistic and mechanistic evidence then it requires both kinds of evidence in the ideal causal epistemology, and hence under the epistemic account.

Of course the epistemic theory of causality may be subject to problems of its own. Here we shall consider only a few possible objections; further discussion can be found in Williamson (2005; 2006a; 2006b; 2007); Choi (2006). First, if in the ideal causal epistemology causal beliefs are constrained by known mechanisms as well as probabilistic dependencies, and mechanisms are themselves causal, then won't the epistemic account offer an account of causality in terms of causality—a circular account?² Of course it is a matter of debate as to whether mechanisms—or more generally whatever does the explanatory work in science—are themselves causal. Russell (1913) argued that at base science is not causal; others disagree. But we do not need to decide this question here, because even if mechanisms are causal, that does not mean that the epistemic account is viciously circular. I don't dispute that in the ideal causal epistemology evidence of mechanisms as well as dependencies helps to constrain appropriate causal beliefs. But this constraining relation is epistemological rather than ontological: if you grant certain facts, then certain beliefs are appropriate. For there to be any vicious circularity, it would have to be the case that you could not grant those facts without having the beliefs in the first place. That is, there would have to be an epistemological circularity. But there is no such circularity: we know about mechanisms linking a beating heart to a failing liver without needing to know whether or not the former causes the latter. Perhaps, in order to know about this mechanism one needs to know about other causal relationships—lower-level relationships that concern the circulation of blood. But even if that were the case, there would be no vicious circularity because these other causal relationships are not in question here: it is the higher-level relationship that is in question. There is no

reason why the ideal causal epistemology should not involve feedback—that is, certain evidence warrants certain causal beliefs which in turn allow the prospect of further evidence which then warrants other causal beliefs. It is only a problem if one already needs to believe that *A* causes *B* in order to believe that *A* causes *B*; but that is not the case here.

The preceding response applies equally to a second objection. Causal relations can themselves be causes and effects (Williamson and Gabbay, 2005). For instance, smoking causing cancer causes governments to restrict tobacco advertising. Again, there is a whiff of circularity here: do we not need to know causal relations in order to determine causal relations? But again, there is no vicious circularity. While it seems reasonable to hold that one needs to know whether smoking causes cancer before one can tell whether or not that causal relationship is a cause of advertising restrictions, the former causal relation is at a lower level to the latter. As before, the causal epistemology is incremental—some causal beliefs are required before others can be acquired. It is not circular in the sense that a particular causal belief presupposes itself.

There is a third related objection. Under the above account there is a category difference between causal relationships on the one hand, which are epistemic, and probabilistic dependencies and mechanisms on the other, which are physical. But what is the basis for this sharp distinction? If there is no substance to it and dependencies and mechanisms turn out to be epistemic, then the concept of ideal evidence would be difficult to delineate. In response I would say this: I have argued that one needs to plump for epistemic causality because of the failure of analyses in terms of probabilities or mechanisms, or a pluralist combination of the two. Purely on grounds of simplicity it would be very nice to have some such analysis, but unfortunately it is not possible, and we are forced to turn to an epistemic account. I have not argued for epistemic causality on the basis of general considerations in favour of epistemic accounts or general considerations that tell against physical accounts. Thus unless there are compelling reasons why one can't construe evidence of dependencies and evidence of mechanisms in terms of physical probabilities and physical processes respectively, it is quite natural to maintain a sharp distinction between epistemic and physical entities. Now I grant that there are problems with physical accounts of probability—it is implausible that a physical interpretation of probability will be successful in underpinning all our uses of probability. For example, one may rightly say that there is probability 0.1 that the trillionth digit of π is 9, but a physical account is unlikely to make sense of this claim. However, there seems no problem construing probabilistic *evidence* in terms of physical probabilities—the dependence between smoking and cancer for instance can quite easily be interpreted in terms of relative frequencies. Similarly there does not seem to be any insurmountable problem with construing evidence of the mechanism from smoking to cancer in terms of the physical features of the human body and of smoke. Hence there do not seem to be any compelling reasons for abandoning physical accounts of probabilistic and mechanistic evidence, and the distinction between epistemic and physical can be maintained.

A fourth related objection proceeds as follows. One may grant that both probabilities and mechanisms are physical, but claim that there is no sharp ontological distinction between probabilistic dependencies and mechanisms. (Perhaps on the grounds that mechanisms are ultimately reducible to low-level chains of probabilistic dependencies; perhaps because probabilistic dependencies are somehow analysable in terms of mechanisms; or, more

plausibly, because both are analysable in terms of the basic make-up of the physical universe.) If probabilities and mechanisms are at base of the same stuff, and probabilistic dependencies and mechanisms support causal claims, then surely causality is of this stuff too—that is, causality is itself at base physical. In response, I would say that this objection is sound but misplaced. It is sound in the sense that indeed the epistemic theory makes no commitment to non-physical entities in order to analyse causality. According to the epistemic theory, the causal relation is determined by its uses—inference and explanation; causality is a map of optimal causal beliefs, and these beliefs are optimal in the sense that they chart the optimal inferences and explanations. Now it must be the physical world that makes these inferences and explanations optimal, so the physical world is at base the truthmaker for causal claims. The objection is misplaced, however, because the job of a philosophical theory of *X* is to do more than say what kind of thing the truthmakers of *X*-claims are *at base*. It is easy to say that fundamental physical entities and their spatio-temporal locations make *X*-claims true; it is of course much harder to say *how* they make them true. If *Zs* make *X*-claims true but only via *Ys*, then a theory of *X* should point this out.⁸ According to the epistemic account the physical world makes causal claims true, but only via probabilistic dependencies, mechanisms, and rational beliefs. Some maps directly map the world, others map inferences—as the epistemic theory makes clear, causality is a map of the latter kind.

This point can be put another way. If the epistemic account of causality is right, it must in principle be possible to come up with some characterization of causality that just appeals to the indicators of causality, along the lines of ‘*C causes E* iff there is a dependency and a mechanism and *C* is prior to *E* unless *C* or *E* are absences in which case ...’ However, the resulting characterization would be so complicated that it would be very hard to see why it is a correct characterization and why we should have a concept of cause at all. The answer to these questions, I think, must invoke the uses of the causal relation and the idea of inferential and explanatory success. So the epistemic account is to be preferred on the grounds (1) that it is clear that its characterization of causality must be correct and (2) that it tells us the full story, while the above kind of characterization is just a part of the picture.

Here’s an analogy. Consider a travel graph whose nodes are towns and where one node is linked to another by an edge if normally you can travel between them in two hours. This kind of graph allows one to make a whole host of useful inferences and explanations related to travelling. Let’s call the binary relation that is depicted by the edges of this map ‘travelity’. Now one could try to characterize this relation in terms of its evidential indicators as follows ‘*C travelizes E* iff there is some kind of mechanism for travelling between *C* and *E* for which the mean travel time is less than two hours unless it is a Sunday or bank holiday or ...’ Already it’s getting a bit pointless. It’s just a map. It’s a very useful kind of map precisely because it over-simplifies and because it overloads a simple binary relation with connotations of travel mechanisms (which are explanatory) and travel times (which enable inferences). While travelity may supervene on physical entities and their spatio-temporal locations, any viable account of travelity should talk a bit about the map and its uses.

I have argued in this chapter that probabilistic theories of causality are inadequate in two respects: they admit counterexamples and they fail to account for the relationship between causality and mechanisms. But mechanistic theories are no better, nor are pluralist accounts

that are part probabilistic, part mechanistic. The right way to integrate probabilistic and mechanistic considerations into an account of causality is to embed them in an epistemic account.

In terms of the distinctions of sect. 2, the epistemic account can cover both single-case and generic causal claims, since we can have both single-case and generic causal beliefs. Similarly causes and effects can be population-level or individual-level. Causality turns out not to be directly physical on the epistemic account, since it is analysed in terms of the way we should represent the world, rather than directly in terms of the world itself. The epistemic account can accommodate an objective notion of cause or indeed a subjective notion, though I have suggested in Williamson (2005: §9.7) that it is likely that ideal causal beliefs are so highly constrained that they are fully objective or close to fully objective. The epistemic account can handle both potential and actual causation; the former is predominantly associated with the inferential uses of causality while the latter is primarily used for explanation.

FURTHER READING

A good grounding in probabilistic causality can be had by studying the following texts. Reichenbach (1971) is a readable and historically important first port of call (see sect. 4). Salmon (1988) gives a recent exposition of his own view of causality: while Reichenbach gives primacy to probabilistic relationships over mechanisms in his account, Salmon does the opposite. Pearl (2000) offers a comprehensive picture of the causal net approach outlined in sect. 7. The motivation behind the epistemic theory of causality is presented in Williamson (2005). Hitchcock (1997) is also recommended as a survey of probabilistic causation that covers several issues not taken up here.

REFERENCES

- BAUER, M. (1998). 'The Medicalization of Science News—from the "Rocket-Scalpel" to the "Gene-Meteorite" Complex', *Social Science Information* 37/4: 731–51.
- CARTWRIGHT, N. (1979). 'Causal Laws and Effective Strategies', *Noûs* 13: 419–37.
- CHOI, S. (2006). Review of 'Bayesian Nets and Causality', *Mind* 115: 502–6.
- DOWE, P. (2000). *Physical Causation*. Cambridge: Cambridge University Press.
- EELLS, E. (1991). *Probabilistic Causality*. Cambridge: Cambridge University Press.
- and SOBER, E. (1983). 'Probabilistic Causality and the Question of Transitivity', *Philosophy of Science* 50: 35–57.
- GILLIES, D. (2000). *Philosophical Theories of Probability*. London: Routledge.
- GOOD, I. J. (1959). 'A Theory of Causality', *British Journal for the Philosophy of Science* 9: 307–10.
- (1961a). 'A Causal Calculus I', *British Journal for the Philosophy of Science* 11: 305–18. Errata *ibid.* 13: 88.
- (1961b). 'A Causal Calculus II', *British Journal for the Philosophy of Science* 12: 43–51. Errata *ibid.* 13: 88.
- HALL, N. (2004). 'Two Concepts of Causation'. in J. Collins, N. Hall, and L. Paul (eds.),

- Causation and Counterfactuals*. Cambridge, Mass.: MIT, 225–76.
- HITCHCOCK, C. (1997). ‘Probabilistic Causation’ *Stanford Encyclopedia of Philosophy*. Rev. edn. 2002.
- MACH, E. (1883). *The Science of Mechanics*. Stanford, Calif.: University of Stanford Press. 4th edn. 1919. La Salke, Ill.: Open Court.
- (1905). *Knowledge and Error*. Dordrecht: Reidel.
- MARTEL, I. (2000). ‘*Probabilistic Empiricism: In Defence of a Reichenbachian Theory of Causation and the Direction of Time*’. PhD thesis, University of Colorado.
- NEAPOLITAN, R. E. (1990). *Probabilistic Reasoning in Expert Systems: Theory and Algorithms*. New York: Wiley.
- OTTE, R. (1986). ‘Reichenbach, Causation, and Explanation’, in *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, i. 59–65.
- PEARL, J. (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Mateo, Calif: Morgan Kaufmann.
- (2000). *Causality: Models, Reasoning, and Inference*. Cambridge: Cambridge University Press.
- PEARSON, K. (1892). *The Grammar of Science*. 2nd edn. 1900. London: Black. 3rd edn. 1911. New York: Macmillan.
- REICHENBACH, H. (1959). ‘The Principle of Causality and the Possibility of its Empirical Confirmation’, in *Modern Philosophy of Science*. London: Routledge & Kegan Paul, 109–34.
- (1971). *The Direction of Time*. Berkeley and Los Angeles: University of California Press.
- REISS, J. (2007). ‘Time Series, Nonsense Correlations and the Principle of the Common Cause’, in *Causality and Probability in the Science*. London: College Publications, 179–96.
- RUSSELL, B. (1913). ‘On the Notion of Cause’, *Proceedings of the Aristotelian Society*, 13: 1–26.
- RUSSO, F., and WILLIAMSON, J. (2007). ‘Interpreting Causality in the Health Sciences’, *International Studies in the Philosophy of Science* 21/2: 157–170.
- SALMON, W. C. (1980). ‘Probabilistic Causality’, *Pacific Philosophical Quarterly* 61: 50–74; repr. in Salmon (1998), 208–32.
- (1984). *Scientific Explanation and the Causal Structure of the World*. Princeton, NJ: Princeton University Press.
- (1988). ‘Intuitions—Good and Not-so-good’, In B. Skyrms and W. L. Harper (eds.), *Causation, Chance, and Credence* Dordrecht: Kluwer, i. 51–71.
- (1998). *Causality and Explanation*. Oxford: Oxford University Press.
- SKYRMS, B. (1980). *Causal Necessity: A Pragmatic Investigation of the Necessity of Laws*. New Haven, Conn.: Yale University Press.
- SOBER, E. (1988). ‘The Principle of the Common Cause’, in J. H. Fetzer (ed.), *Probability and Causality: Essays in Honour of Wesley C. Salmon*. Dordrecht: Reidel, 211–28.
- (2001). ‘Venetian Sea Levels, British Bread Prices, and the Principle of the Common Cause’, *British Journal for the Philosophy of Science* 52: 331–46.
- SPIRITES, P., GLYMOUR, C., and SCHEINES, R. (1993). *Causation, Prediction, and Search*.

2nd edn. 2000. Cambridge Mass.: MIT.

- SPOHN, W. (2002). 'Bayesian Nets Are All There Is to Causal Dependence', in M. C. Galavotti, P. Suppes, and D. Costantini (eds.), *Stochastic Causality*. Chicago: University of Chicago Press.
- SUPPES, P. (1970). *A Probabilistic Theory of Causality*. Amsterdam: North-Holland.
- WILLIAMSON, J. (2005). *Bayesian Nets and Causality: Philosophical and Computational Foundations*. Oxford: Oxford University Press.
- (2006a). 'Causal Pluralism versus Epistemic Causality', *Philosophica* 77: 69–96.
- (2006b). 'Dispositional versus epistemic causality', *Minds and Machines*, 16: 259–76.
- (2007). 'Causality', in D. Gabbay and F. Guenther (eds.), *Handbook of Philosophical Logic*. Berlin: Springer, xiv. 95–126.
- and GABBAY, D. (2005). 'Recursive Causality in Bayesian Networks and Self-fibring Networks', in D. Gillies (ed.), *Laws and Models in the Sciences*. London: King's College Publications, 173–221. With comments, 223–45.
- YULE, G. U. (1926). 'Why Do We Sometimes Get Nonsense-Correlations Between Time Series?' A Study in Sampling and the Nature of Time Series, *Journal of the Royal Statistical Society*, 89/1: 1–63.