

Is Causation Probabilistic?

Caren A. Frosch (c.frosch@reading.ac.uk)

Department of Psychology, University of Reading, RG6 6AL, UK

P. N. Johnson-Laird (phil@princeton.edu)

Department of Psychology, Princeton University, NJ 08540, USA

Abstract

One view of causation is deterministic: *A causes B* means that on any occasion in which *A* occurs, *B* occurs. An alternative view is that causation is probabilistic: it means that given *A*, the probability of *B* is greater than some criterion, such as the probability of *B* given not *A*. Evidence about the induction of causal relations cannot decide between these two accounts, and so we examined how people *refute* causal relations. Three experiments showed that they tend to be satisfied that a single counterexample of *A* and not-*B* refutes claims of the form, *A causes B* and *A enables B*. But, as a deterministic theory based on mental models predicted, when participants required more than one refutation they tended to do so for claims of the form, *A enables B*. Similarly, refutations of the form not-*A* and *B* were more frequent for enabling than causal claims. We interpret these results to imply that causation is a deterministic notion, and that causation and enabling conditions are distinct concepts.

Keywords: Causation; refutation; causes and enablers.

Introduction

The everyday concept of causation is puzzling. No-one seems sure about what it means, and some theorists even deny its coherence and seek to outlaw it from scientific discourse (Russell, 1912-13; Salsburg, 2001). The traditional view is that causation is deterministic (e.g., Hume, 1748/1988). But, a contrasting probabilistic conception developed in the twentieth century (Reichenbach, 1956; Suppes, 1970). Both views have current proponents in psychology. In what follows, we outline psychological theories of causation, consider some of the recent evidence for the role of covariation in inferring causation, and examine the role of refutation in elucidating the debate between deterministic and probabilistic theories.

Psychological theories of causation

The theory of mental models provides a deterministic account of the everyday meaning of causation (Goldvarg & Johnson-Laird, 2001; Johnson-Laird, 2006). *A causes B* refers to three possibilities:

A	B
not- A	B
not- A	not- B

in which *B* does not precede *A* in time. The theory acknowledges the role of probability in causal reasoning: the evidence supporting a causal relation may be probabilistic, but not the concept itself.

In contrast to deterministic theories, a probabilistic account of causation, such as the one proposed by Cheng and her colleagues, is based on the difference between the conditional probability of *B* given *A* and the conditional probability of *B* given not-*A* (Cheng & Novick, 1990). This difference enters into various computations in order to account for different causal tasks (e.g., Cheng, 1997; Novick & Cheng, 2004).

Other theories are less committed to either side of the debate and make provisions for both deterministic and probabilistic representations. For example, Sloman, Barbey, and Hotaling (2009) describe a theory, based on causal models, which allows causal assertions to be probabilistic or deterministic depending on whether an uncertainty parameter is included or ignored. Another recent theory postulates that causal relations are characterized by forces that either work together or oppose one another (Barbey & Wolff, 2007; cf. Wolff, 2007). And Sloman et al describe how Barbey and Wolff's (2007) transitive dynamics model can account for both deterministic and probabilistic causal relations.

Is covariation the key to causal inference?

Research on induction has shown that people infer causal relations from data about covariations among events (McArthur, 1972; Cummins, Lubart, Alksnis, & Rist, 1991; Shanks, 2004), and that they may do so when the probability of the effect given the cause is less than 1. Such results, however, do not establish that the concept of causation, as opposed to evidence for its applicability, is probabilistic. Indeed, the view that causal relations are inferred from covariation information alone is controversial and questioned by Lagnado, Waldmann, Hagmayer, and Sloman (2007). This idea receives further support from Luhmann and Ahn (2003), who demonstrated that when individuals consider physical causation they are willing to attribute causal roles to unobserved events. They explain cases of not-*A* and *B* by referring to alternative causes of *B*, and they explain cases of *A* and not-*B* by referring to inhibitory causes ('disabling' conditions). Hence, individuals make causal inferences without all of the relevant covariation data. Luhmann and Ahn (2005) went on to argue that probabilistic causality is psychologically implausible.

Schulz and Sommerville (2006) addressed the question of whether children have a deterministic view of causation. In a series of experiments, they showed that children infer unobserved causes when an observed cause does not always produce an effect, and that children can distinguish between

unobserved inhibitory causes and absent unobserved generative causes. These authors concluded that children have a deterministic view of causation. In sum, the tendency of children and adults to explain cases of *A* and *not-B* and *not-A* and *B* by invoking disablers and alternative causes casts doubt on a probabilistic interpretation of causation, which by its very nature should tolerate counterexamples without the need for explanation.

The role of refutation

The search for counterexamples is an integral part of the mental model theory of reasoning (Johnson-Laird & Byrne, 1991). For example, in reasoning from sentential connectives, individuals tend to refute putative conclusions that are consistent with the premises but that do not follow of necessity, by identifying a counterexample (Johnson-Laird & Hasson, 2003). Likewise, if causation is deterministic, then individuals should seek a single refutation to refute a causal claim. In contrast, if causation is probabilistic, then individuals should seek multiple refutations to refute a causal claim.

The model theory also draws a clear distinction between the meaning of causal claims and enabling claims. A causal assertion, such as, ‘emotions cause individuals to pay attention,’ refers to the following three possibilities in a temporal order:

emotion	attention
not-emotion	attention
not-emotion	not-attention

An enabling assertion, such as, ‘emotions enable individuals to pay attention’, refers to what emotions make possible, and so it refers to the following temporally-ordered possibilities:

emotion	attention
emotion	not-attention
not-emotion	not-attention

A weaker sense of enabling is consistent with all four contingencies, but there is often an implicature that only the antecedent, emotion in this case, makes the consequent possible (Goldvarg & Johnson-Laird, 2001).

To hold three distinct possibilities in mind is difficult (Bauer & Johnson-Laird, 1993; Bucciarelli & Johnson-Laird, 1999), and so the model theory postulates that individuals normally represent only the case in which both clauses are true (Johnson-Laird & Byrne, 1991). Hence, both causal and enabling claims have the same mental models:

A	B
...	

where the ellipsis denotes other implicit possibilities. One corollary is that individuals should have difficulty in

distinguishing between the meanings of causal and enabling claims, which may account for the common view that they do not differ in meaning (e.g., Mill, 1874).

According to the model theory, causal claims with the structure *A causes B* are refuted by a single occurrence of *A* without *B*:

A	not-B
---	-------

Enabling claims of the form *A enables B* granted the implicature that *only A* makes *B* possible are refuted by a single occurrence of *B* without *A*:

not-A	B
-------	---

However, in the absence of the implicature, the only way to refute the weak claim that *A* makes *B* possible is to observe that *B* never occurs in the presence of *A*. Hence, the theory predicts that individuals should seek multiple observations to refute enabling claims more often than to refute causal claims.

Assertions about prevention, *A prevents B*, are equivalent to *A causes not-B*, and so they refer to these three possibilities:

A	not-B
not-A	B
not-A	not-B

Hence, a claim about prevention should be refuted by a single occurrence of *A* with *B*:

A	B
---	---

Wolff (personal communication 06/12/2008) acknowledges that the force dynamics theory makes no clear predictions about how causes and enablers should be refuted, but suggests that a combination of the theory and lexical semantics might predict that a claim of the form, *A causes B*, is refuted by an observation of *A* and *not-B*, whereas an observation of *not-A* and *B* would be more damaging for an enabling relation. The theory accordingly makes much the same predictions as the model theory. But, other current theories take a different point of view. Sloman’s (2005) causal model theory stresses the importance of mechanisms, and so it implies that refutations establish either the absence of a mechanism relating cause to effect or that the mechanism is malfunctioning or broken. Hence, causes are refuted by the absence of an enabler or the presence of a disabler; and enablers are refuted by the absence of additional enablers (Sloman, personal communication, 09/29/2008). In some cases, a single refutation suffices; it depends on how many other causes, enablers, or disablers, an assertion brings to mind (Sloman, personal communication, 06/24/2008). However, the model theory predicts an asymmetry: individuals should be more likely to seek multiple refutations for enabling than for

causal assertions. We carried out three experiments to test these predictions.

The Experiments

In three experiments, the participants were presented with causal assertions made by different individuals, such as:

Peter says: Following this diet causes a person with this sort of metabolism to lose weight.

Their task was to state what would refute Peter's claim. Experiments 1 and 2 were conducted online, and, to strengthen their results, Experiment 3 was conducted in a face to face test of each participant in the laboratory. There were also differences in the materials and procedure over the three experiments. Participants in Experiment 1 were asked an open-ended question about the sort of evidence they required to refute causal, enabling, and prevention assertions. In general, the participants sought observations of *A* and *not-B* and *not-A* and *B*, but with no clear difference between causes and enablers. Hence, Experiments 2 and 3 explained the difference between causes and enablers in more detail, and the participants selected one of two options (*A* and *not-B* or *not-A* and *B*) and stated whether or not one observation was sufficient evidence for a refutation.

Method

Design The participants acted as their own controls in all three experiments. In Experiment 1, they carried out the task for five assertions about causes, five assertions about enabling conditions, and five assertions about preventions. We created the three sorts of assertion from each of fifteen contents, but each participant saw just one assertion with each of the contents, which were rotated over the participants in order to counterbalance them. In Experiments 2 and 3, the assertions described eight causes and eight enabling conditions. In all three experiments, the assertions occurred in a different random order for each participant.

Participants The participants were as follows: Experiment 1: 18 Princeton University students and staff (mean age = 22 years). Experiment 2: 20 Princeton University students and staff (mean age = 23 years). Experiment 3: 21 University of Reading undergraduates (mean age = 22 years).

Materials The materials for the experiments were drawn from five domains: physiological, e.g., 'regular exercise of this sort causes a person to build muscle', physical, e.g. 'contact between these two sorts of substance causes an explosion to occur', mechanical, e.g., 'tuning this kind of engine in this special way causes a reduction in its fuel consumption to occur', socio-economic, e.g., 'introducing these health care reforms causes more people to seek medical attention', and psychological, e.g., 'a person's positive attitude towards you causes you to like that person'. Each content occurred with three sorts of verb in Experiment 1: *causes*, *enables*, and *prevents*, but only with

the first two of these verbs in Experiments 2 and 3. People distinguish between causing and enabling relations whether causal relations are expressed using *causes*, *forces*, or *makes*, and whether enabling relations are expressed using *enables*, *allows*, or *helps* (Wolff & Song, 2003). But, there can be slight difference of meaning amongst these verbs, e.g., it would be odd to assert, *the weak brackets enabled the shelves to collapse*, and so we used the most general verbs. Individuals are more likely to test a hypothesis when they consider another person's claim rather than a self-generated one (Cowley & Byrne, 2005; under review; Sanbonmatsu, Posavac, Vanous, & Ho, 2005), and so each assertion was presented as made by a different person, and the participants had to say what would refute the assertion.

Procedure Experiments 1 and 2 were conducted online at Princeton University and Experiment 3 was conducted face to face at the University of Reading. In Experiment 1, participants were asked an open-ended question:

What sort of evidence would you require to refute this statement? Please describe one or more possibilities that would show that Peter's claim is false.

In Experiments 2 and 3, the participants were asked:

Which of these two possibilities provides more decisive evidence against this assertion?

And they were presented with two options of the form: *A* and *not-B*, *not-A* and *B*. They were then asked:

Would this observation suffice to show that the claim is false? And, if not, what other observation would be necessary?

In these two experiments, we explained that the causing event 'brought about the outcome' whereas the enabling event 'made the outcome possible'.

Results

Figure 1 presents the percentages of trials over the three experiments on which the participants required a single refutation, and the remaining responses were for multiple refutations. We used a stringent criterion in Experiment 1, e.g., assertions such as, 'I would have people perform the exercise and see whether they consistently built muscle or not,' were classified as requiring multiple refutations. There was an explicit question on the matter for Experiments 2 and 3.

As the figure shows, the participants required more single than multiple refutations for all three sorts of claim in Experiment 1 (Wilcoxon tests: for *causes*, $z = 2.43$, $p < .01$, one-tail probability here and throughout; for *enables*, $z = 2.52$, $p < .01$; for *prevents*, $z = 3.37$, $p < .005$). There was no reliable difference in the number of single refutations requested for *causes* (70%) and *enables* (77%; $z = .78$, $p >$

.22). In Experiment 2, we asked the participants whether the one observation they had chosen (either *A* and *not-B*, or *not-A* and *B*) was sufficient to refute the assertion. On 87% of the trials the participants stated that one observation was sufficient. But, as Figure 1 shows, they were more likely to request multiple observations for *enables* (21% of trials) than for *causes* (6% of trials; Wilcoxon test, $z = 2.35$, $p < .01$). In Experiment 3, as Figure 1 also shows, the participants were satisfied with a single observation for *causes* (91% of trials, Wilcoxon test, $z = 4.07$, $p < .00005$) and for *enables* (63% of trials, Wilcoxon test, $z = 2.1$, $p < .025$), and the increase in multiple refutations for *enables* compared to *causes* was reliable (Wilcoxon test, $z = 3.35$, $p = .0005$).




Figure 1. The percentages of single refutations for *causes*, *enables* in the three experiments; the balances of the percentages were for multiple refutations. The data for *prevents* are from Experiment 1.

Figure 2 presents the percentages of the two critical sorts of refutation for assertions based on *causes* and on *enables*: *A* and *not-B*, and *not-A* and *B*, respectively. In all three experiments, the predominant response was an observation of *A* and *not-B*. For Experiment 1, we derived these percentages by categorizing the responses, and the coding was verified by a second coder who was blind to the hypotheses, and the 9% of disagreements were resolved through discussion. In Experiments 2 and 3, the percentages in the Figure are based on the participants' explicit choices about which contingency would refute the assertions. In Experiment 1, the participants tended to require an observation of *A* and *not-B* to refute both *causes* and *enables*, whereas on 70% of trials they required an observation of *A* and *B* to refute *A prevents B*. They tended to require an observation of *not-A* and *B* more often for *enables* (11% of trials) than for *causes* (7% of trials), but the trend was not reliable (Wilcoxon test, $z = 1.2$, $p > .1$). Experiment 2 yielded the same pattern of results, but the participants selected an observation of *not-A* and *B* more often for *enables* (25% of trials) than for *causes* (10% of trials, Wilcoxon test, $z = 2.53$, $p < .01$). Those who chose the refutation of *not-A* and *B* for *enables* tended not to

request multiple observations: ten participants required them on fewer than half of such trials, only two participants required them on more than half of such trials, and there were two ties (Binomial test, $p < .02$). Experiment 3 replicated the results: the participants opted for an observation of *not-A* and *B* more often for *enables* (38% of trials) than for *causes* (12% of trials; Wilcoxon test, $z = 2.5$, $p < .01$).

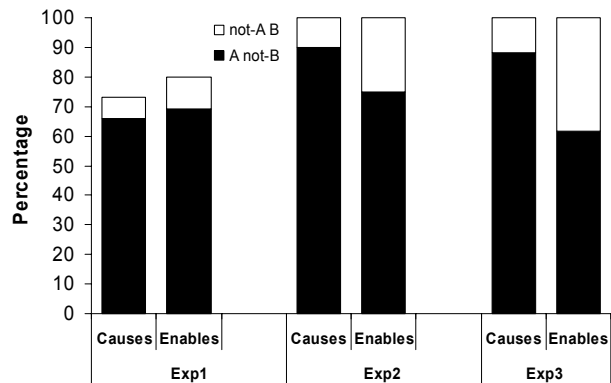


Figure 2. The percentages of *A* and *not-B* and *not-A* and *B* refutations for *causes* and *enables* across the three experiments.

Among the requests for multiple refutations in Experiment 1, 88% referred to the observation of conjunctions between the antecedent and the outcome. Some requested evidence of *A* and *not-B* as well as *not-A* and *B*, whereas others requested multiple observations of the same kind, e.g. 'Showing that, through scans of muscle or strength tests, that the particular sort of regular exercise does not aid in building muscle'. The remaining 12% of multiple refutations were of a different nature, requesting information about other potential factors that might contribute to the outcome, e.g. 'would need to see how many people seek medical attention regardless of reforms'.

General Discussion

The three experiments showed that single observations were likely to be judged sufficient to refute causal claims, such as: 'Regular exercise of this sort causes a person to build muscle'. In Experiment 1, participants tended to suggest single refutations, such as: 'The person did regular exercise of the particular sort and didn't gain muscle'. They also required single observations to refute assertions about *enables* and *prevents*. Experiments 2 and 3 clarified the difference between the relations by adding a rider that *causes* means *brings about*, and *enables* means *makes possible*. Participants continued to request single refutations of *A* and *not-B*, but there were now reliable differences between the two assertions. As the model theory predicts, the participants tended to require single refutations for *causes* but a greater proportion of multiple refutations for *enables*. Similarly, they were more likely to require

observations of *not-A* and *B* to refute *enables* than to refute *causes*. Both the preference for single refutations and the difference between *causes* and *enables* are difficult to explain from a probabilistic standpoint, which implies that refutations should be statistical: the weight of evidence should accumulate to depress the conditional probability of *B* given *A* below some given criterion (e.g., Cheng, 1997). Probabilistic accounts also have little to say about enabling relations. One such view is that an enabling condition is constant in the situation, whereas the cause is inconstant (Cheng & Novick, 1991). But, this constraint is not invariable, e.g.:

Mary threw a lighted cigarette into a bush. Just as the cigarette was going out, Laura deliberately threw petrol on it. The resulting fire burnt down her neighbor's house.

Naïve individuals identify Mary's action as the enabler of the fire and Laura's action as its cause, but the enabler is not a constant (Frosch, Johnson-Laird, & Cowley, 2007).

Why, then, should anyone suppose that causation is probabilistic? Luhmann and Ahn (2005) suggest that causal inferences sometimes appear to be probabilistic, because individuals make no explicit reference to hidden causes responsible for counterexamples. We propose three further factors that enhance the popularity of the probabilistic view of causation.

The first factor is philosophical. The success of quantum mechanics in the Twentieth century replaced Newtonian determinism with an irreducible probabilistic physics. This view, in turn, has inculcated a probabilistic metaphysics (Suppes, 1984).

The second factor is methodological. Systematic evidence pertinent to causation is often statistical, in part because noise and erroneous observations are bound to occur, and in part because hidden causes and disabling factors may be uncontrolled in samples of data. For instance, if you observe that 99 out of 100 smokers develop cancer, whereas only 9 out of 100 nonsmokers from the same population do, then you have *prima facie* evidence that smoking causes cancer, but it is not the whole story. Some hidden disabling component is at work sparing the single survivor, granted that the observation is not spurious.

The third factor is psycholinguistic, and perhaps the most relevant to psychological theories of causation. Many causal claims are couched in the form of *generic* assertions. A 'generic' assertion contains a noun phrase as its subject that lacks a specific quantifier, such as 'all' or 'some', e.g.: *Ducks lay eggs* (Leslie, 2008). And generic assertions are compatible with counterexamples, e.g., *drakes don't lay eggs* (Khemlani, Leslie, Glucksberg, & Fernandez, 2007). An assertion, such as:

Smoking causes cancer

is generic, and it too tolerates counterexamples. But, if we introduce an explicit universal quantifier, e.g.:

Smoking always causes cancer

then individuals are likely to judge that a single counterexample refutes the claim. Generic causal assertions tolerate exceptions, but that is because they are generic, not because they are causal. With a universal quantifier, the deterministic nature of everyday causality is so obvious that we deliberately chose to use generic assertions in our experiments. Yet, in all three of them, the participants tended to require just a single observation of the form *A* and *not-B* in order to refute claims of the form: *A causes B*. This result is contrary to a probabilistic concept of causation.

One defense of the probabilistic view is that it allows that the probability of an effect given a cause could be 1, and so at the limit it is a deterministic theory. But, granted a probabilistic concept, why should individuals make the limit interpretation so often? The probabilistic theory needs to explain the occurrence of this phenomenon. Consider, for instance, these two assertions: "The spark will probably cause a fire" and "The spark will cause a fire". If the spark doesn't cause a fire, the second assertion is false, but not the first one. On a probabilistic account, the two assertions should be synonymous. Likewise, as the model theory predicts, participants require a greater proportion of multiple refutations for enabling assertions than for causal assertions. Probabilistic theories offer no ready account for this phenomenon, either. Our findings add to the mounting evidence that causation is, not probabilistic, but deterministic. When one event causes another, the antecedent suffices for the consequent to occur.

Acknowledgments

This research is supported in part by a Postdoctoral Research Fellowship awarded to the first author by the UK Economic and Social Research Council (PTA-026-27-1688) and in part by a National Science Foundation grant, SES 0844851. Deductive and probabilistic reasoning, to the second author. We thank Steve Sloman and Philip Wolff for their helpful advice.

References

- Barbey, A., & Wolff, P. (2007). Learning causal structure from reasoning. In D. S. McNamara & J. G. Trafton (Eds.), *Proceedings of the 29th Annual Cognitive Science Society* (pp. 713-718). Austin, TX: Cognitive Science Society.
- Bauer, M. I. & Johnson-Laird, P. N. (1993). How diagrams can improve reasoning. *Psychological Science*, 4, 372-378, 1993.
- Bucciarelli, M. & Johnson-Laird, P.N. (1999). Strategies in syllogistic reasoning. *Cognitive Science*, 23, 247-303.
- Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, 104, 367-405.
- Cheng, P. W., & Novick, L. R. (1990). A probabilistic contrast model of causal induction. *Journal of Personality and Social Psychology*, 58, 545-567.

- Cheng, P.W., & Novick, L.R. (1991). Causes versus enabling conditions. *Cognition*, 40, 83-120.
- Cowley, M., & Byrne, R. M. J. Hypothesis falsification and opponents in the 2-4-6 task, under review.
- Cowley, M. & Byrne, R. M. J. (2005). When falsification is the only path to truth. In B. G. Bara, L. Barsalou, & M. Bucciarelli (eds.). *Proceedings of the 27th Annual Conference of the Cognitive Science Society* (pp. 512-517). Mahwah, NJ: Erlbaum.
- Cummins, D. D., Lubart, T., Alksnis, O., & Rist, R. (1991). Conditional reasoning and causation. *Memory & Cognition*, 19, 274-282.
- Frosch, C.A., Johnson-Laird, P.N., & Cowley, M. (2007). It's not my fault, Your Honor, I'm only the enabler. In D. S. McNamara & J. G. Trafton (Eds.), *Proceedings of the 29th Annual Cognitive Science Society* (p. 1755). Austin, TX: Cognitive Science Society.
- Goldvarg, E., & Johnson-Laird, P. N. (2001). Naive Causality: a mental model theory of causal meaning and reasoning. *Cognitive Science*, 25, 565-610.
- Hume, D. (1988). *An enquiry concerning human understanding*. Ed. A. Flew. La Salle, IL: Open Court. (Originally published 1748.)
- Johnson-Laird, P. N. (2006). *How we reason*. Oxford: Oxford University Press.
- Johnson-Laird, P.N., & Byrne, R.M.J. (1991) *Deduction*. Hillsdale, NJ: Erlbaum.
- Johnson-Laird, P.N., & Hasson, U. (2003) Counterexamples in sentential reasoning. *Memory & Cognition*, 31, 1105-1113.
- Khemlani, S., Leslie, S. J., Glucksberg, S., & Fernandez, P. (2007). Do ducks lay eggs? How people interpret generic assertions. In D. S. McNamara & J. G. Trafton (Eds.), *Proceedings of the 29th Annual Cognitive Science Society* (pp. 395-401). Austin, TX: Cognitive Science Society.
- Lagnado, D. A., Waldmann, M. R., Hagmayer Y., & Sloman, S. A. (2007). Beyond covariation: Cues to causal structure. In A. Gopnik & L. Schulz (Eds.), *Causal learning: Psychology, philosophy, and computation*. Oxford: Oxford University Press, pp. 154-172.
- Leslie, S. J. (2008). Generics: cognition and acquisition. *Philosophical Review*, 117, 1-47.
- Luhmann, C. C., & Ahn, W. K. (2003). Evaluation the causal role of unobserved variables. In R. Alterman & D. Kirsh (Eds.), *Proceedings of the 25th Annual Conference of the Cognitive Science Society* (pp. 734-739). Mahwah, NJ: Erlbaum.
- Luhmann, C. C., & Ahn, W. K. (2005). The meaning and computation of causal power: Comment on Cheng (1997) and Novick and Cheng (2004). *Psychological Review*, 112(3), 685-693.
- McArthur, L. (1972). The how and what of why: some determinants and consequences of causal attribution. *Journal of Personality and Social Psychology*, 22, 171-193.
- Mill, J.S. (1874) *A system of logic, ratiocinative and inductive: Being a connected view of the principles of evidence and the methods of scientific evidence*. Eighth Edition. New York: Harper. (First edition published 1843.)
- Novick, L.R. & Cheng, P.W. (2004). Assessing interactive causal influence. *Psychological Review*, 111, 455-485.
- Reichenbach, H. (1956). *The direction of time*. Berkeley: University of California Press.
- Russell, B.A.W. (1912-13). On the notion of cause. *Proceedings of the Aristotelian Society*, 13, 1-26.
- Salsburg, D. (2001). *The lady tasting tea: How statistics revolutionized science in the Twentieth century*. New York: W.H. Freeman.
- Sanbonmatsu, D. M., Posavac, S. S., Vanous, S., & Ho, E. A. (2005). Information search in the testing of quantified hypotheses: How "all," "most," "some," "few," and "none" hypotheses are tested. *Personality and Social Psychology Bulletin*, 31 (2), 254-266.
- Schulz, L. E., & Sommerville, J. (2006). God does not play dice: Causal determinism and preschoolers' causal inferences. *Child Development*, 77(2), 427-442.
- Shanks, D. R. (2004). Judging covariation and causation. In Koehler, D.J., and Harvey, N. (Eds.) *Handbook of judgment and decision making*. (pp. 220-239) Oxford: Blackwell.
- Sloman, S. A. (2005). *Causal models: How we think about the world and its alternatives*. New York: Oxford University Press.
- Sloman, S. A., Barbey, A. K., & Hotaling, J. (2009). A causal model theory of the meaning of cause, enable, and prevent. *Cognitive Science*, 33, 21-50.
- Suppes, P. (1970). *A probabilistic theory of causality*. Amsterdam: North-Holland.
- Suppes, P. (1984). *Probabilistic metaphysics*. Oxford: Basil Blackwell.
- Wolff, P. (2007). Representing causation. *Journal of Experimental Psychology: General*, 136, 82-111.
- Wolff, P. & Song, G. (2003). Models of causation and causal verbs. *Cognitive Psychology*, 47, 276-332.