

★

Evidence and Epistemic Causality

★

Michael Wilde & Jon Williamson

Draft of January 13, 2015

The epistemic theory of causality maintains that causality is an epistemic relation, so that causality is taken to be a feature of the way a subject reasons about the world rather than a non-epistemological feature of the world. In this paper, we take the opportunity to briefly rehearse some arguments in favour of the epistemic theory of causality, and then present a version of the theory developed in [Williamson \(2005, 2006, 2009, 2011\)](#). Lastly, we provide some possible responses to an objection based upon recent work in epistemology.

This paper provides a broad overview of the issues, and in a number of places readers are directed towards work which provides the relevant details.

§1

Causality and evidence

Standardly, there are *mechanistic* and *difference-making* theories of causality. On the one hand, mechanistic theories maintain that variables in a domain are causally related if and only if they are connected by an appropriate sort of mechanism. On the other hand, difference-making theories maintain that variables are causally related if and only if one variable makes an appropriate sort of difference to the other. Stated in these terms, there might be a worry that both types of theory will be uninformative or circular, since the appropriate sort of mechanism or the difference-making relationship might be less well understood than the causal relation, or only understandable in terms of causality. Therefore, a good mechanistic theory should attempt to provide an account of the appropriate sort of mechanism in better-understood and non-causal terms, e.g., the mechanism might be understood in terms of a process that possesses a conserved quantity ([Dowe, 2000](#)). Likewise, a good difference-making theory should attempt to provide an account of the appropriate sort of difference-making relationship in better-understood and non-causal terms, e.g., in terms of probabilistic dependence conditional upon other causes (cf. [Williamson, 2009](#)).

The two different types of theory—mechanistic and difference-making—have conflicting implications regarding the epistemology of causality. On a mechanistic theory, one's body of evidence is sufficient to establish a causal claim if and only if one's evidence is sufficient to establish that there exists an appropriate mechanistic connection. On a difference-making theory, one's evidence is sufficient to establish a causal claim if and only if that evidence is sufficient to establish the existence of an appropriate sort of difference-making relationship.

However, there are well-known proposed counterexamples to mechanistic and difference-making theories of causality. There are cases involving absences, which seem to be cases of causality but without any appropriate sort of mechanism (Williamson, 2011). For example, it seems that missing my flight in London is a cause of my talk being cancelled in Australia, even though they are not connected by any appropriate mechanism. The cases involving absences seem to demonstrate that the existence of a suitable mechanism is not a necessary condition for causality. But these cases also seem instructive regarding the epistemology of causality. Given these cases, it seems that establishing a causal claim does not require establishing the existence of an appropriate sort of mechanism, since here the causal claim is established and there exists no such mechanism. Rather, having established that there exists an appropriate difference-making relationship between missing my flight and my talk being cancelled here seems sufficient to establish the causal claim.

There are also cases of over-determination, which seem to be cases of causality but without any appropriate sort of difference-making relationship (Hall, 2004, pp. 232–241). For example, it seems that dropping the first bomb caused the end of the war, even though dropping this first bomb did not make the appropriate difference to the ending of the war, since a second bomb was dropped an instant later and would have ended the war regardless. The cases of over-determination seem to demonstrate that the existence of an appropriate difference-making relationship is not a necessary condition for causality. Once again these cases are instructive about the epistemology of causality. In over-determination cases, it seems that establishing a causal claim cannot require establishing the existence of an appropriate sort of difference-making relationship, since here the causal claim is established and there exists no such relationship. Instead, having established that there exists an appropriate sort of mechanism between the dropping of the first bomb and the ending of the war by itself seems sufficient to establish the causal claim.

How should one respond to these proposed counterexamples? There are two standard lines of response.

The first line of response is simply to dismiss the relevant suggested counterexamples. The proponent of a mechanistic theory could deny that cases involving absences are genuine cases of causality; see, e.g., Dowe (2000, pp. 123–145). Likewise, the proponent of a difference-making theory could deny that cases of over-determination are genuine cases of causality; see, e.g., Coady (2004). This line of response, however, looks implausible, since cases involving absences look like straightforward cases of causality. Indeed, Schaffer (2004) argues that cases involving absences are treated as genuine cases of causality in both ordinary and theoretical contexts, and rightly so since such cases have all the hallmarks of genuine cases of causality. Similarly, over-determination cases look like paradigmatic cases of causality, and accordingly proponents of a difference-making theory of causality typically accept such cases and attempt to accommodate them; see, e.g., Paul and Hall (2013, pp. 70–172). This suggests another way of dismissing the suggested counterexamples. The proponent of a difference-making theory could suggest that over-determination cases in fact do involve an appropriate sort of difference-making relationship, by suitably refining their account of the difference-making relationship. Likewise, the proponent of a mechanistic theory could suggest that cases involving absences do involve some appropriate sort of mechanism; see, e.g., Thomson (2003, pp. 84–86). However, it is generally agreed that there is currently neither a difference-making nor a mechanistic theory of causality that can accommodate all the proposed counterexamples in this manner (Paul and Hall, 2013, p. 1).

The second line of response is to advocate pluralism, e.g., by maintaining that there is both a mechanistic type and a difference-making type of causality; see, e.g., [Hall \(2004\)](#). The idea here is that cases involving absences are cases of the difference-making type of causality without the mechanistic type of causality, and vice versa for cases of over-determination. Of course, this line of response has its own implications regarding the epistemology of causality. Presumably, establishing that there exists an appropriate sort of difference-making relationship is sufficient to establish a causal claim about the difference-making type of causality, and establishing that there exists an appropriate sort of mechanism suffices to establish a causal claim about the mechanistic type of causality. Reasons to doubt this line of response are presented in [Williamson \(2006\)](#). For instance, it is argued there that non-pluralist theories of causality should be preferred on the grounds of simplicity.

The main problem is that both these lines of response to the counterexamples—attempting to rebut them or moving to pluralism—have difficulty accounting for the practice of scientists when establishing causal claims ([Williamson, 2006](#), pp. 73–74). In particular, when establishing a causal claim, health scientists typically require evidence both that there exists an appropriate difference-making relationship and that there exists an appropriate mechanism. (In this sense, cases involving absences and over-determination are atypical.) Firstly, establishing only that there exists an appropriate sort of difference-making relationship is typically not sufficient for a health scientist to consider the corresponding causal claim established. For example, smoking was not established as a cause of heart disease, despite strong evidence that smoking makes an appropriate difference to the prevalence of heart disease, until there was also strong evidence that there is an appropriate mechanism linking smoking and disease ([Gillies, 2011](#)). Secondly, establishing only that there exists an appropriate mechanism is also typically not sufficient for a health scientist to consider the corresponding causal claim established. It was not established that the micro-organism anthrax bacillus was the cause of anthrax, despite strong evidence of an appropriate mechanism, until there was also strong evidence that there exists an appropriate difference-making relationship between anthrax bacillus and anthrax ([Clarke et al., 2014a](#), p. 345).

The two lines of response to the counterexamples struggle to explain this need for both types of evidence. On the one hand, if a standard difference-making theory of causality is correct, it is difficult to explain the scientist's apparent need for evidence that there exists an appropriate mechanism, when there is already good evidence of an appropriate difference-making relationship available. On the other hand, if a standard mechanistic theory is correct, it is difficult to explain the scientist's apparent need for evidence that there exists an appropriate difference-making relationship when available evidence already establishes the existence of a suitable mechanism. The pluralist response is no better. The pluralist analyses some causal claims as mechanistic, others as referring to a difference-making type of cause. In the former case, the pluralist cannot explain the need for evidence of an appropriate difference-making relationship; in the latter case, the pluralist cannot explain the need for evidence of the existence of a suitable mechanism.

Now, it might be objected that it is not a desideratum of a theory of causality that it accounts for the practice of scientists when establishing causal claims. For instance, it might be objected that scientists are getting the epistemology of causality wrong, and thus that a theory of causality need not make sense of their practice. However, given the success of the sciences in establishing causal claims, it is likely that their practice is indicative of the correct epistemology of causality.

But there is another reason to believe that scientists are doing things right. In particular, causal claims are used for prediction, explanation and control. Russo and Williamson (2007) argue that the explanatory use of causal claims typically requires that there is an appropriate sort of mechanism linking the cause and the effect. This is because explanations are best given by appealing to mechanisms (cf. Williamson, 2013). Moreover, Russo and Williamson argue that the use of causal claims for prediction and control requires that a cause should typically make an appropriate sort of difference to its effects, for otherwise information about the presence of the cause would tell us nothing about the presence of its effects and vice versa, and also instigating a cause would not be a good strategy for achieving its effects (Russo and Williamson, 2007, p. 159). Given this, it is plausible that establishing a causal claim typically requires also establishing that there exists an appropriate difference-making relationship. Thus, it looks as if scientists are getting the epistemology right, and so a theory of causality should account for their practice.

Howick (2011) objects that sometimes establishing that there exists an appropriate difference-making relationship is sufficient to establish the corresponding causal claim. He says that '[i]n many cases, tightly controlled comparative clinical studies suffice to establish causation' (2011, p. 933). This is intended to constitute an objection to the thesis that establishing a causal claim requires establishing an appropriate difference-making relationship *and* an appropriate sort of mechanism. However, he assumes that tightly controlled comparative clinical studies provide evidence only that there exists an appropriate difference-making relationship. In fact, instances in which the results of tightly controlled comparative clinical studies suffice to establish a causal claim are plausibly instances in which those results establish also that some mechanism exists to explain the difference-making relationship. In other instances, it might be unreasonable to consider a causal claim established on the basis of the results of clinical studies, since the established difference-making relationship might be due to confounding, if the relata are correlated effects of a common cause. In these instances, the additional evidence that there exists an appropriate mechanism would help to rule out confounding as an explanation of the difference-making relationship. Plausibly, instances in which a causal claim can be established on the basis of the results of clinical studies are instances in which the studies are of sufficient quality that they establish that confounding is not the likely explanation of the difference-making relationship, and thus also provide evidence that the existence of a suitable mechanism is the likely explanation.

One might also object that establishing the existence of an appropriate mechanism is sometimes sufficient for establishing the corresponding causal claim. In these instances, however, it is plausible that the evidence which establishes the existence of the mechanism is also sufficient to establish an appropriate difference-making relationship. The problem with establishing the causal claim only on the basis of the existence of an appropriate sort of mechanism is that there might be undiscovered mechanisms that counteract the action of the known mechanism, so that overall there is no appropriate difference-making relationship. This is called the problem of *masking*. In some instances, one can know enough about a mechanism that one can establish that counteracting mechanisms do not exist, and thus that overall there exists an appropriate difference-making relationship. In other instances, additional evidence that there exists an appropriate difference-making relationship helps to overcome the problem of masking. Plausibly, instances in which a causal claim can be established on the basis of evidence of the existence of an appropriate mechanism are instances in which this evidence also suffices to overcome the problem of




Figure 1: Trihoral relationships involving Canterbury (C), London (L), Gatwick (G), Dunkirk (D), Paris (P) and Orléans (O).

masking, and thus also provide evidence that there exists an appropriate difference-making relationship. These considerations provide more reason that both types of evidence are required in order to establish a causal claim, since each type of evidence compensates for the limitations of the other (Illari, 2011, pp. 144–148).

To conclude this section, it seems that standard theories of causality are susceptible to counterexamples and also struggle to explain good evidential practice in establishing causal claims. Similarly, a pluralist theory struggles to explain the epistemology of causality. How should one respond to this state of affairs?

§2

The epistemic theory of causality

One response is to plump for an epistemic theory of causality. In this section, we introduce the epistemic theory of causality developed in Williamson (2005, 2006, 2009, 2011, 2013). According to this theory, causality is purely epistemic in the sense that our causal claims enable us to reason and interact with the world in certain ways; they are not claims about some causal relation that exists independently of us and our epistemic practices.

By way of analogy, consider the following relation, which we shall call the *trihoral* relation: two places stand in this relation if it is reasonable to expect to be able to travel between them within three hours. One can chart this relation, as in Fig. 1. It is not hard to see how such a chart might be useful—for parents of young children to plan breaks, for instance. Moreover, it seems clear that the edges between the nodes of such a graph do not correspond straightforwardly to a single sort of physical, or even non-epistemological, link between the places that correspond to the nodes. If the graph is correct, it is in virtue of a complex array of facts about the presence and absence of train, air, ferry and road connections, as well as normal conditions relating to travel. In that sense, the trihoral relation is purely epistemic. This is not to say that there is no fact of the matter as to whether two places stand in the trihoral relation—there is, at least in non-borderline cases. It is to say that such propositions are not made true by the existence of some single, unified, worldly (non-epistemological) connection between places that we can call ‘trihorality’.

Similarly with the causal relation. Our causal claims are extremely useful—particularly for prediction, explanation and control. It is this utility which accounts for our having the concept of cause: not the existence of some simple kind of worldly connection to which our causal claims refer that we can call causality. If a causal graph is correct, it is in virtue of a complex array of facts about the presence and absence of mechanisms, as well as the presence and absence of difference-making

relationships and their magnitude.

Consider another analogy, to Bayesian probability. Bayesian probabilities are epistemic—rational degrees of belief, not directly physical entities—and they underwrite certain predictions and bets. Moreover, at least on the objective Bayesian view, there is typically a fact of the matter as to what the correct Bayesian probabilities are, given the extent and limitations of the evidence available. In the version of objective Bayesianism developed in [Williamson \(2010\)](#), for example, three norms constrain the strengths of one’s beliefs:

Probability. One’s degrees of belief should be representable by a probability function P_E .

Calibration. One’s degrees of belief should fit evidence: $P_E \in \mathbb{E}$, the subset of probability functions that fit evidence. In particular, they should be calibrated to the corresponding empirical probabilities, insofar as one has evidence of them: if evidence determines just that the chance function $P^* \in \mathbb{P}^*$, then $P_E \in \langle \mathbb{P}^* \rangle$, the convex hull of the set of potential chance functions.

Equivocation. One’s degrees of belief should otherwise equivocate as far as possible between the elementary outcomes. In particular, if there are finitely many elementary outcomes then $P_E \in \text{maxent}\mathbb{E}$, the subset of those functions that fit evidence which have maximal entropy, as long as $\text{maxent}\mathbb{E} \neq \emptyset$.

These norms tend to be motivated by appealing to betting considerations, along the following lines. If one’s degrees of belief do not meet the norms and one places bets in accordance with these beliefs, then one opens oneself up to potential losses: the possibility of sure loss in the case of the Probability norm, long-run loss in the case of the Calibration norm, and worst-case expected loss in the case of the Equivocation norm ([Williamson, 2010](#), Chapter 3). On the other hand, one does not expose oneself to these losses if the norms are followed. Hence one’s degrees of belief must conform with the norms if one is to avoid avoidable losses. Arguably it would be irrational not to avoid avoidable losses. So the norms must hold for the strengths of one’s beliefs to be apportioned in a rational way.

This view of probability is different to the epistemic view of causality in that it is pluralist, positing empirical, non-epistemic probabilities (chances), in addition to epistemic, Bayesian probabilities. Nevertheless, it is instructive in that it suggests a particular connection between evidence and epistemic probabilities. An epistemic theory of causality can posit similar norms that constrain one’s causal claims:

Acyclicity. One’s causal claims should be representable by an acyclic graph \mathcal{C} .

Calibration. One’s causal claims should fit evidence: $\mathcal{C} \in \mathfrak{C}$, the subset of acyclic graphs that fit evidence.

Equivocation. \mathcal{C} should otherwise be as non-committal as possible about what causes what.

How might these norms be fleshed out? One simple recipe proceeds as follows. We can take \mathcal{C} to be a graph whose nodes correspond to variables, which contains an arrow from variable A to variable B if it is claimed that A is a cause of B , a gap (i.e., no connection) between A and B if it is claimed that neither causes the other, and an undirected edge between A and B if neither of the above two claims is made involving A and B . Such a graph is acyclic if there is some way of orienting the

undirected edges in the graph such that there is no chain of arrows in the graph that forms a cycle. It is plausible that causal claims can always be representable by an acyclic graph: in cases in which there are apparent causal cycles, one can eliminate these cycles by time-indexing the variables (see, e.g., Clarke et al., 2014b). We may suppose that evidence imposes certain constraints on \mathcal{C} . For example, if evidence establishes that A is a cause of B , represented by $A \rightsquigarrow B$, then there should be some chain of arrows from A to B in \mathcal{C} ; if evidence establishes that A is not a cause of B , $A \not\rightsquigarrow B$, then there should be no chain of edges and arrows from A to B . A causal graph \mathcal{C} is maximally non-committal, from all those in \mathfrak{C} , if there is no other causal graph \mathcal{D} in \mathfrak{C} which makes fewer causal claims (including both arrows and gaps) than \mathcal{C} .

How does evidence impose a constraint of the form $A \rightsquigarrow B$? As discussed in §1, in order to establish that A is a cause of B there would normally have to be evidence both that (i) there is an appropriate sort of difference-making relationship (or chain of difference-making relationships) between A and B —e.g., that A and B are probabilistically dependent, conditional on B 's other causes—, and that (ii) there is an appropriate mechanistic connection (or chain of mechanisms) between A and B —so that instances of B can be explained by a mechanism which involves A . (We saw above that there are some exceptions to this rule, which correspond to the counterexamples to standard theories of causality, discussed in §1.)

How does evidence impose a constraint of the form $A \not\rightsquigarrow B$? Typically, a causal relationship can be ruled out by either (i) evidence that there is no appropriate difference-making relationship (or chain of such relationships) between A and B , or (ii) evidence that there is no mechanism (or chain of mechanisms) that can account for B in terms of A . Thus a drug trial of sufficiently high quality that finds no association between treatment and cure would impose a $A \not\rightsquigarrow B$ constraint. On the other hand, we can rule out certain causal claims involving treatments (such as certain homeopathic treatments) where it is known that there is no possible mechanism by which the treatment can explain a cure.

In sum, then, the analogy with the trihoral relation suggests that causal claims are representational rather than real in the sense that they guide inference, explanation and action but do not refer to a non-epistemic connection that we can call causality. Moreover, the analogy with Bayesian probability can shed some light on the link between evidence and epistemic causality, by suggesting three norms by which the extent and limitations of available evidence constrains one's causal claims.¹

§3

The nature of evidence

It appears that the above recipe for arriving at one's causal claims requires that an ideally rational subject has perfect access to her evidence. The recipe suggests that one's causal claims are rational if and only if they are appropriately constrained by one's body of evidence in accordance with the Acyclicity, Calibration and Equivocation norms. Informally, if one's causal claims are thus constrained by what one takes to be one's body of evidence rather than what is in fact one's body evidence,

¹Note that it is important to distinguish this task of determining, on the basis of current evidence, which set of causal claims is established by that evidence, from the task of formulating, on the basis of current evidence, a set of more tentative causal hypotheses that can be tested by collecting further evidence. A formal approach to developing a set of causal hypotheses is developed in Williamson (2006, Appendix).

then one's causal claims are not appropriately constrained. Therefore, if an ideally rational subject's causal claims are appropriately constrained, it looks like an ideally rational subject must have perfect access to her body of evidence.

The problem is that certain recent work in epistemology purports to show that evidence is not as accessible as following the above recipe seems to require; see, e.g., (Williamson, 2000, pp. 93–113). In particular, it can be claimed that even an ideally rational subject may not have perfect access to her evidence. Thus, the critic might object that the epistemic theory of causality is committed to a false theory of evidence. How should the proponent of an epistemic theory of causality respond to this objection?

One possibility is for the proponent of an epistemic theory to simply deny that evidence is not perfectly accessible in the relevant sense: evidence is such that an ideally rational subject has perfect access to her evidence. Indeed, it looks like there is room for this line of response. This is because many theorists assume something like the *ought implies can principle*, which says that one must be in a position to accomplish anything that one ought to accomplish. The principle might seem a reasonable assumption, since intuitively one is not failing to fulfil one's obligations if one could not have possibly fulfilled those putative obligations. Crucially, this intuitively plausible principle implies that evidence is perfectly accessible in the relevant sense. In particular, if an ideally rational subject's causal claims are constrained by her evidence in accordance with the above recipe, then it must be possible that her causal claims be thus constrained, given the ought implies can principle. In turn, in order for it to be possible that an ideally rational subject's causal claims be appropriately constrained, her evidence must be perfectly accessible. Thus, the proponent of the epistemic theory of causality might maintain that, *contra* recent work in epistemology, evidence is in fact accessible in the relevant sense. Of course, then there is the pressing matter of pointing out where the recent work in epistemology has gone wrong; see Williamson (2015) on this point.

Alternatively, the proponent of an epistemic theory might want to endorse the view that evidence is not perfectly accessible. Then, the question is whether a viable epistemic theory of causality can be proposed that dispenses with the requirement that evidence is perfectly accessible. Arguably, such an epistemic theory of causality can be proposed. Once again, the analogy with Bayesian probability is instructive.

Objective Bayesian probabilities are degrees of belief appropriately constrained by the evidence in accordance with the Probability, Calibration, and Equivocation norms. The problem is that the objective Bayesian theory also seems to require that evidence is such that an ideally rational subject has perfect access to her evidence. Thus, if evidence is not perfectly accessible in this sense, it looks like this objective Bayesian theory of probability cannot be correct.

In response to this state of affairs, Timothy Williamson proposes an alternative *evidential* theory of probability, a theory that dispenses with the requirement that evidence is perfectly accessible (Williamson, 2000, pp. 209–237). On this theory, there exists an objective degree to which a claim is entailed by a given body of evidence, and it is evidential probabilities that measure this partial entailment relation between one's body of evidence and specific claims. But the objective degree to which a claim is entailed by the evidence is not reducible to an ideally rational subject's degree of belief in that claim, where an ideally rational subject follows the Probability, Calibration, and Equivocation norms. This is because such ideally rational subjects might disagree with regards to their degrees of belief on the same body of evidence, if this body of evidence is not perfectly accessible to each of them. Rather, degrees

of belief are rational given a body of evidence only insofar as they match the relevant evidential probabilities. Crucially, the theory of evidential probability remains an epistemic theory of probability, since evidential probabilities depend upon one's body of evidence rather than on purely non-epistemological features of the world.

In a similar manner, the proponent of the epistemic theory might propose an analogous evidential theory of causality, an epistemic theory that dispenses with the requirement that evidence is perfectly accessible in the relevant sense. This evidential theory of causality differs from the epistemic theory of causality presented in §2 in much the same way that the theory of evidential probability differs from the theory of objective Bayesian probability. In particular, this theory can hypothesize that one's body of evidence entails certain relationships between specific claims, relationships that license particular inferences concerning explanation, prediction, and control. These relationships can be charted by a unique causal graph given a body of evidence. This theory of causality remains epistemic, since the causal graph depends upon one's body of evidence rather than some non-epistemological feature of the world. But unlike the original epistemic theory, causality is not reducible to the causal claims arrived at by following a recipe which requires that evidence is perfectly accessible. Instead, one's causal claims are rational insofar as they match the unique causal graph. Arguably, then, a version of the epistemic theory of causality survives the objection based upon recent work in epistemology, viz., the evidential theory of causality.

The subject matter of this section is covered in more detail in forthcoming research of Michael Wilde.

§4

Conclusion

In this paper, we have argued that standard theories of causality have a hard time making sense of the epistemology of causality or are susceptible to counterexamples (§1). One response to this state of affairs is to adopt an epistemic theory of causality, such as that developed in Williamson (2005, 2006, 2009, 2011), as outlined in §2. While some might object that such an epistemic theory of causality conflicts with some recent work in epistemology, in §3 we have suggested two lines of response to this objection: either to rebut worries about the inaccessibility of evidence or to adapt the epistemic theory to fit the view that evidence may be inaccessible.

References

- Clarke, B., Gillies, D., Illari, P., Russo, F., and Williamson, J. (2014a). Mechanisms and the evidence hierarchy. *Topoi*, 33:339–360.
- Clarke, B., Leuridan, B., and Williamson, J. (2014b). Modelling mechanisms with causal cycles. *Synthese*, 191(8):1651–1681.
- Coady, D. (2004). Preempting preemption. In Collins, J., Hall, N., and Paul, L., editors, *Causation and Counterfactuals*, pages 325–339. MIT Press.
- Dowe, P. (2000). *Physical Causation*. Cambridge University Press.
- Gillies, D. (2011). The Russo-Williamson thesis and the question of whether smoking causes heart disease. In Illari, P., Russo, F., and Williamson, J., editors, *Causality in the Sciences*, pages 110–125. Oxford University Press.
- Hall, N. (2004). Two concepts of causation. In Collins, J., Hall, N., and Paul, L., editors, *Causation and Counterfactuals*, pages 225–276. MIT Press.

- Howick, J. (2011). Exposing the vanities—and a qualified defense—of mechanistic reasoning in health care decision making. *Philosophy of Science*, 78:926–940.
- Illari, P. (2011). Mechanistic evidence: Disambiguating the Russo-Williamson thesis. *International Studies in the Philosophy of Science*, 25:139–157.
- Paul, L. and Hall, N. (2013). *Causation: A User's Guide*. Oxford University Press.
- Russo, F. and Williamson, J. (2007). Interpreting causality in the health sciences. *International Studies in the Philosophy of Science*, 21:157–170.
- Schaffer, J. (2004). Causes need not be physically connected to their effects: the case for negative causation. In Hitchcock, C., editor, *Contemporary Debates in Philosophy of Science*, pages 197–216. Blackwell.
- Thomson, J. (2003). Causation: Omissions. *Philosophy and Phenomenological Research*, 66:83–103.
- Williamson, J. (2005). *Bayesian Nets and Causality*. Oxford University Press.
- Williamson, J. (2006). Causal pluralism versus epistemic causality. *Philosophica*, 77:69–96.
- Williamson, J. (2009). Probabilistic theories of causality. In Beebe, H., Hitchcock, C., and Menzies, P., editors, *The Oxford Handbook of Causation*, pages 185–212. Oxford University Press.
- Williamson, J. (2010). *In Defence of Objective Bayesianism*. Oxford University Press.
- Williamson, J. (2011). Mechanistic theories of causality part II. *Philosophy Compass*, 6:433–444.
- Williamson, J. (2013). How can causal explanations explain? *Erkenntnis*, 78:257–275.
- Williamson, J. (2015). Deliberation, judgement and the nature of evidence. *Economics and Philosophy*, in press.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford University Press.