

★

Calibration for Epistemic Causality

★

Jon Williamson

Draft of June 7, 2019

Abstract

The epistemic theory of causality is analogous to epistemic theories of probability. Most proponents of epistemic probability would argue that one's degrees of belief should be calibrated to chances, insofar as one has evidence of chances. The question arises as to whether causal beliefs should satisfy an analogous calibration norm. In this paper, I formulate a particular version of a norm requiring calibration to chances and argue that this norm is the most fundamental evidential norm for epistemic probability. I then develop an analogous calibration norm for epistemic causality, argue that it is the *only* evidential norm required for epistemic causality, and show how an epistemic account of causality that incorporates this norm can be used to analyse objective causal relationships.

§1

Introduction

The epistemic theory of causality interprets causality in terms of rational belief (Williamson, 2005, Chapter 9). This view is analogous to that of various epistemic theories of probability (including subjective Bayesianism, objective Bayesianism and certain theories of imprecise probability), which interpret probability in terms of rational strength of belief.

Epistemic theories of probability need to provide some account of how these probabilistic beliefs ought to fit the available evidence. For example, epistemic theories of probability often hold that strengths of belief ought to be calibrated to chances, insofar as there is evidence of chances. Similarly, an epistemic theory of causality needs to provide an account of how causal beliefs ought to fit the available evidence. The main aim of this paper is to put forward such an account.

§2 introduces the analogy between epistemic causality and epistemic probability. §3 argues that one particular norm for calibrating probabilistic beliefs to chances is fundamental to epistemic probability. §4 develops an analogous calibration norm for epistemic causality, argues that this norm is the *only* evidential norm that is needed for epistemic causality, and develops an epistemic analysis of objective causation. Along the way, we encounter some disanalogies between epistemic causality and epistemic probability, which are summarised in §5.

Although this paper is not the place to mount a full defence of epistemic causality, an appendix outlines some motivation for the theory and discusses its relation

to other theories of causality, for the benefit of readers who are unfamiliar with epistemic causality.

§2

From epistemic probability to epistemic causality

Epistemic theories of probability interpret certain probabilistic claims as saying something about strength of belief: roughly speaking, a claim such as *the probability that the patient will recover is 0.7* is correct insofar as, given the available evidence, it is reasonable to believe to degree 0.7 that the patient will recover.

Epistemic theories of probability deem strength of belief to be reasonable or rational just when it satisfies certain norms. These norms are often chosen on the grounds that their adherence leads to beliefs that are epistemically or pragmatically optimal—beliefs that minimise epistemic inaccuracy (see, e.g., Pettigrew, 2016), or that avoid avoidable losses (e.g., Williamson, 2017, Chapter 9), for instance. Different epistemic theories of probability differ as to what counts as rational strength of belief, because they impose different sets of norms. Norms on strength of belief can be classified into four categories, as follows.

Structural norms impose constraints on the structure of one's strengths of belief. For example, subjective and objective Bayesianism hold that strengths of belief should be representable by a probability function, while imprecise probability holds that they should be representable by a non-empty set of probability functions.

Evidential norms explicate the way in which the available evidence imposes constraints on strength of belief. For example: if A is in one's body of evidence then one ought to fully believe A . Another such norm says that, if one's body of evidence E determines that the chance of A is x (which I shall write as $P^*(A) = x$), then one should believe A to degree x . This latter norm is sometimes called a *calibration norm*, because it requires that strengths of belief be calibrated to chances, insofar as one has evidence of them. Calibration norms will be explored in more detail in subsequent sections of this paper.

Equivocation norms are constraints on strength of belief imposed by lack of evidence: these tend to rule out extreme degrees of belief in favour of more equivocal degrees of belief, in situations where the evidential norms do not force extreme degrees of belief. For example, Regularity says that one should reserve probability 1 for those propositions that are logical consequences of one's evidence base. Subjectivists tend to disavow stronger equivocation norms, but some objective Bayesians advocate the Principle of Indifference, which says that, given a partition of basic expressible propositions that are treated symmetrically by the evidence, one ought to believe each such proposition to the same degree. Some objective Bayesians endorse a generalisation, the Maximum Entropy Principle, which says that strength of belief should be captured by probability functions with maximum entropy, from all those that satisfy structural and evidential norms. Some objectivist imprecise probabilists, on the other hand, impose an equivocation norm that says that strength of belief should be captured by the set of *all* probability functions which satisfy constraints imposed by evidential norms.

Finally, *diachronic norms* govern changes in strength of belief as evidence changes. Subjective Bayesians and imprecise probabilists tend to endorse Bayesian conditionalisation or some generalisation of Bayesian conditionalisation here, while objective Bayesians who advocate the maximum entropy principle do not need a further diachronic norm: they can say that one should simply maximise entropy afresh with

respect to the new evidence in order to determine new degrees of belief (Williamson, 2010, Chapter 4).

These norms ensure that rational strength of belief depends partly on the extent and limits of the evidence E , and, according to certain epistemic theories, partly on subjective inclinations of the agent.

Epistemic causality can be introduced by analogy to epistemic probability. Causal claims are interpreted as another kind of belief. According to an epistemic theory of probability, a probabilistic claim is interpreted as a relational belief of the form $P_E(A) = x$, where $x \in [0,1]$, or, in the case of imprecise probability, $P_E(A) = y$, where $\emptyset \neq y \subseteq [0,1]$. Similarly, according to an epistemic theory of causality, a causal claim F is a cause of G is interpreted as a relational belief of the form $C_E(F,G)$. These probabilistic and causal beliefs are not beliefs *about* probability or causality—not the belief *that* the probability of A is x , or the belief *that* F is a cause of G . Rather, these relational beliefs are *kinds* of belief: $P_E(A) = x$ expresses a quantitative belief that motivates certain predictions, bets and decisions; similarly, $C_E(F,G)$ expresses a qualitative belief that motivates a particular range of predictions, explanations and control inferences. These two kinds of belief are thus characterised by their distinctive links to inference and action.

One view has it that the right combination of norms on probabilistic beliefs is the combination that yields, on balance, optimal predictions, bets and decisions. Analogously, one might suggest that the right combination of norms on causal beliefs is the combination that yields, on balance, optimal predictions, explanations and control inferences. In either case, one can only theorise as to the right combination of norms.

A range of epistemic theories of causality can be developed by selecting norms analogous to those invoked by epistemic theories of probability, as follows. What we might call *precise* epistemic causality would include a structural norm that says that one's causal beliefs on evidence E should be representable by a single relation C_E induced by a directed acyclic graph (DAG) on the set of cause and effect variables. *Imprecise* epistemic causality, in contrast, might hold that causal beliefs should be representable by a set of such DAG relations. *Subjective* epistemic causality would not advocate an equivocation norm, while *objective* epistemic causality would: such a norm would rule a committal causal claim unreasonable in the absence of evidence that forces such a commitment.

While Williamson (2005, Chapter 9) develops an objective, precise epistemic theory of causality, in this paper I will defend an imprecise account. As in the case of the objective Bayesian theory of probability, the advocate of an objective epistemic theory of causality may well be able to do without a distinct diachronic norm, by endorsing the repeated use of the synchronic norms that track the changing evidence. On the other hand, no epistemic theory can avoid endorsing an evidential norm—any epistemic theory of causality must say something about the constraints that the available evidence imposes on rational causal beliefs. The main aim of the present paper is to make headway with this question of how to fit causal beliefs to available evidence.

We will also further the goal of understanding causality *simpliciter*. While epistemic accounts of probability and causality focus primarily on the probabilistic and causal claims of a particular agent with particular evidence, it might also be possible to use such accounts to analyse unrelativised probabilistic and causal claims. According to such an analysis, the present chance of A is x just in case, were one to have ideal evidence that includes the complete history so far (all particular mat-

ters of fact up to the present) and the ability to correctly infer chances from this complete history (i.e., were one able to correctly invoke all facts about how chances depend on this history), and were one to satisfy norms on rational strength of belief, one would believe A to degree x . Such an analysis would view chance facts as facts about rational degree of belief (Lewis, 1980, pp. 109–113). Analogously, one might hold that F is a cause of G just when, were one to have ideal evidence consisting of the complete history (all particular matters of fact) and the ability to correctly infer causal claims from this complete history (i.e., were one able to correctly invoke all facts about how causal relationships depend on this history), and were one to satisfy norms on rational causal beliefs, one would have the causal belief F is a cause of G . As I shall argue in §4.3, this characterisation has the potential provide an analysis of agent-independent causation. According to such an analysis, the causal facts are facts about rational causal beliefs.

§3

Calibration norms for epistemic probability

Having introduced epistemic causality and its relation to epistemic probability, let us now turn to the question of how to fit probabilistic beliefs to evidence. The aim will be to learn from the case of probabilistic beliefs, which will be explored in this section, in order to be in a position to develop an account of evidential norms for causal beliefs in the next section.

First, some general remarks about evidence. There is little agreement as to what evidence is. For example, some hold that an agent’s evidence is what she knows (Williamson, 2000), others that her evidence is constituted by her true beliefs (Mitova, 2017), or her credences that are set by observation (Jeffrey, 2004), or her information (Rowbottom, 2014), or what she rationally grants (Williamson, 2015). Here I shall remain neutral between these positions. Moreover, there is disagreement with regard to whether only true propositions qualify as evidence. Again, I shall take no stance on this question here. There is also little agreement as to whether, where a proposition is established on the basis of an agent’s evidence, that proposition is also evidence. Again, I shall remain neutral about this, calling the notion of evidence that is closed under establishing further propositions the agent’s *body of evidence*, E , and the notion of evidence that provides the primary grounds for whatever else is established the agent’s *evidence base*, B . I shall take the body of evidence E to be a consistent set of propositions about possible states of affairs up to the present time, such that a proposition is in E if (a) it is in the agent’s evidence base B , or (b) the agent is rationally required to infer it from B , if pressed, or (c) she is rationally permitted to infer it and she does in fact infer it or she would infer it if pressed. According to this picture, there is an inferential mapping $B \rightarrow E$ from the agent’s evidence base to her body of evidence that depends partly on what she is rationally required to infer and partly on what she is disposed to infer. These inferences may depend on the agent’s utilities: e.g., the greater the disutility of erroneously inferring a false proposition, the more cautious the agent might be in her inferences.

§3.1. Chance Calibration

As discussed above, one central evidential norm for epistemic probability requires calibration of strengths of belief to chances, insofar as there is evidence of these

chances. I shall explicate this norm as follows (see, e.g., [Williamson, 2010](#), §2.3):

Chance Calibration. If, according to E , the current chance function P^* lies within some set \mathbb{P}^* of probability functions, $P^* \in \mathbb{P}^*$, then one's belief function P_E should lie within the convex hull $\langle \mathbb{P}^* \rangle$ of that set, $P_E \in \langle \mathbb{P}^* \rangle$.

Chance Calibration is closely related to David Lewis' Principal Principle ([Lewis, 1980](#)):

Principal Principle. $P(A|XF) = x$, where X says that the chance at time t of proposition A is x , $P_t^*(A) = x$, and F is any proposition that is compatible with X and admissible at time t .

The Principal Principle can be thought of as a kind calibration norm if one invokes the commonly-held assumption that conditional beliefs are conditional probabilities:

CBCP. $P_E(A) = P(A|E)$, where P is a probability function that is rational in the total absence of evidence.

To see the connection between Chance Calibration and the Principal Principle, note first that one cannot currently possess evidence that is inadmissible or incompatible with the current chances: current evidence cannot tell one more about the truth of a proposition than can the current chances. Indeed, [Lewis \(1980, pp. 92–96\)](#) argues that matters of particular fact up to the present are always admissible with respect to the present chances. Thus if the present body of evidence $E = XF$ and t is the present time then F is admissible. Furthermore, F is compatible with X , since body of evidence E is a consistent set of propositions. Hence CBCP and the Principal Principle imply that $P_E(A) = P(A|XF) = x$. This is just what is required by Chance Calibration. More generally, if Y says that the chance of A at t lies in some set y of probabilities and $E = YF$ then CBCP and the Principal Principle imply that

$$\begin{aligned}
 P_E(A) &= P(A | YF) \\
 &= \int_{x \in [0,1]} P(A | XYF)P(X | YF) dx \\
 &= \int_{x \in y} P(A | XYF)P(X | YF) dx \\
 &= \int_{x \in y} P(A | XF)P(X | YF) dx \\
 &= \int_{x \in y} xP(X | YF) dx.
 \end{aligned}$$

This is in the convex hull $\langle y \rangle$ of y because $P(X | YF) \in [0,1]$ and $\int_{x \in y} P(X | YF) dx = 1$. Thus Chance Calibration can be viewed as a consequence of the Principal Principle and CBCP.

Conversely, Chance Calibration implies the evidential content of the Principal Principle, in the context of CBCP. Suppose body of evidence $E = XF$. As explained above, F is admissible and compatible with X , so E determines that $\mathbb{P}^* \subseteq \{P : P(A) = x\}$. Chance Calibration requires that $P_E \in \langle \mathbb{P}^* \rangle$, which implies that $P_E(A) = x$. By CBCP, $P(A|XF) = P_E(A) = x$. This is just what is required by the Principal Principle. Hence Chance Calibration and the Principal Principle can be viewed as different but equivalent explications of the same evidential norm.

Chance Calibration might be thought of as preferable to the Principal Principle as an explication of a norm requiring calibration to chances because it does not require CBCP. Arguably, CBCP is poorly suited to objective Bayesianism because it can conflict with the Maximum Entropy Principle (Friedman and Shimony, 1971; Williamson, 2010, Chapter 4). Moreover, Hawthorne et al. (2017) have argued that the Principal Principle in combination with CBCP is not suited to subjective Bayesianism because, under certain auxiliary assumptions about admissibility, it implies the Principle of Indifference, which is a kind of equivocation norm. Hence the Principal Principle, which requires CBCP if it is to play the role of an evidential norm, is on shakier ground than Chance Calibration, which doesn't depend on CBCP.

Chance Calibration has an additional advantage over the Principal Principle. The Principal Principle can only apply where chance statements are elements of the domain of the agent's belief function, while Chance Calibration applies independently of the agent's beliefs about chances. To insist that the agent has determinate beliefs about all chance propositions is neither realistic nor a normative ideal—quite the opposite, in fact. In the case of both artificial and human agents there are very good grounds for keeping the domain of expressible propositions under control: probabilistic inference is computationally complex and this complexity increases exponentially in the size of the domain. Therefore, the normative ideal is that the size of the domain of the probability function be no larger than strictly necessary. This favours Chance Calibration.

An advocate of the Principal Principle might put forward the following two considerations in its defence. Firstly, it is quite compatible with the Principal Principle that some chances are not in the domain of the belief function. In such a situation, the Principal Principle will be silent about which degrees of belief are appropriate. So there is a sense in which the Principal Principle is less demanding than Chance Calibration: it does not force any specific degrees of belief in this case.¹ In response, note that this behaviour is *not* what we want from an evidential norm. If A is in the domain of the belief function and there is evidence that $P^*(A) = x$ but $P(A) = x$ is not expressible in the domain of the belief function, it is nevertheless clearly appropriate to believe A to degree x . We would want an evidential norm to reflect this. Chance Calibration does provide this constraint but the Principal Principle does not. So the Principal Principle fares worse in such a situation after all.

Second, one might think that the Principal Principle has the following advantage over Chance Calibration: it does not presume that the domain of the chance function P^* is the same as the domain of the belief function P_E . This is important if certain propositions about which we can form beliefs do not have chances. However, this advantage is illusory: Chance Calibration does not require this presumption. Suppose \mathbb{P} is the set of belief functions. In the formulation of Chance Calibration, simply take \mathbb{P}^* to be that subset of \mathbb{P} that satisfies constraints imposed by evidence of chances. For example, if evidence determines just that the current chance of A is x , $P^*(A) = x$, then $\mathbb{P}^* = \{P \in \mathbb{P} : P(A) = x\}$. There is no need to assume that chances are defined on all elements of the domain of the belief function.

In sum, Chance Calibration requires neither CBCP nor that chance propositions be the objects of beliefs; these considerations make it a better general evidential norm for epistemic probability than the Principal Principle.

¹I am very grateful to an anonymous referee for this suggestion.

§3.2. Other evidential norms

Thus far we have considered Chance Calibration and the Principal Principle as two explications of a norm requiring calibration to chances and suggested that Chance Calibration is preferable. We shall now turn to other forms of evidential norm. I shall argue in this section that other evidential norms are less fundamental than Chance Calibration. An analogue of Chance Calibration will thus be a natural starting point when devising an evidential norm for epistemic causality in §4.

First let us consider a very simple evidential norm: *believe your evidence*, i.e., if $A \in E$ then one ought to fully believe A , $P_E(A) = 1$. This norm is routinely advocated by Bayesians of all stripes and it follows from CBCP. This norm is a consequence of Chance Calibration: if $A \in E$ then one would be reasonably expected to infer, if pressed, that the current chance of A is 1, so Chance Calibration requires that $P_E(A) = 1$. Thus the believe-your-evidence norm can be viewed as superfluous in the presence of Chance Calibration.

Chance Calibration can also be considered to be more fundamental than other deference norms. Chance Calibration says that one should defer to chances—insofar as one has evidence of them—when assigning degrees of belief. But in the absence of chances one should arguably defer to other quantities, such as long-run frequencies in appropriate reference classes. Similarly, one might also defer to the degrees of belief of experts on propositions within their domain of expertise.

Deference norms have the following general structure:

Deference. If body of evidence $E = YF$ where Y is $\varphi(A, y)$, then $P_E(A) = y$, as long as F is admissible with respect to Y .

Chance Calibration and the Principal Principle can be thought of as deference norms, where $\varphi(A, y)$ says that the current chance of A is y .² In the other cases considered above, $\varphi(A, y)$ says that the long-run frequency of A -like outcomes is y , or that an appropriate expert believes A to degree y .

There is a sense in which Chance Calibration is the most fundamental deference norm. In the absence of the precise chances themselves, it is rational to defer to quantities other than chance—e.g., long-run frequencies and expert credences—only because these quantities tell us something about chances. That long-run frequencies are good indicators of chances is witnessed by convergence theorems such as the Central Limit Theorem. Similarly, we defer to experts when their degrees of belief are likely to be better estimates of the chances than one's own. In general, we defer to quantities of which we have evidence when these quantities are indicators of chances.

There is another sense in which Chance Calibration is fundamental. Chance Calibration trumps these other deference norms where conflicts arise. Suppose there is evidence Y of a long-run frequency or expert degree of belief and $P_{YF}(A) =$

²One might worry that Chance Calibration does not qualify as a deference norm, because (1) it admits the possibility of chance propositions that are not in the domain of one's belief function, (2) one cannot entertain a chance proposition that is not in the domain of one's belief function, and (3) one cannot defer to a chance proposition if one cannot entertain it. (I am grateful to a referee for this suggestion.) Against (2), however, one does not have to be able to settle the question of how strongly one believes a proposition in order to entertain it. Arguably, it is enough that the chance proposition is in one's body of evidence for it to be entertainable: if it is in one's evidence base then—at least under most conceptions of evidence—it is entertainable; if one infers it (or would be disposed to infer it, if pressed) from one's evidence base, it is entertainable; and if one is rationally required to infer it, one should be able to infer it, so it is entertainable.

y . If, in addition, there is evidence X that determines that the chance of A is x , where $x \neq y$, then this chance information screens off the frequency or expert evidence from A , i.e., $P_{XYF}(A) = x \neq y$. To put it another way, YF is admissible with respect to X , but XF is inadmissible with respect to Y . This is because we only defer to frequencies or expert degrees of belief insofar as they are indicators of chances.

In response to this claim, one might object that there are other things to which one ought to defer and which cannot be viewed as screened off by present chances. For example, suppose $\varphi(A, y)$ says that one's degree of belief *tomorrow* in A is y . It is far from clear that this information would be trumped by present chances: one might have more or better evidence tomorrow, in which case one should arguably pay attention now to the consequences of that evidence, even if one knows the present chances. Tomorrow's credences may tell us about tomorrow's chances, and information about tomorrow's chances would trump evidence of the today's chances, where a conflict arises.

All this may be true, but there are limits as to the evidence one can have today. As noted above, current evidence is always admissible with respect to a claim X about current chances. One clearly cannot possess as evidence today information about tomorrow's chances that is not also information about today's chances. Thus Chance Calibration still reigns as the primary *evidential* norm: it governs the constraints imposed by current evidence. Perhaps one should defer to one's future credences—if so, this norm (sometimes called the *reflection principle*) is a diachronic norm, not an evidential norm, according to the classification of §2.

In sum, Chance Calibration is the most fundamental evidential norm. Given Chance Calibration, there is no need for either the Principal Principle or the believe-your-evidence norm, and Chance Calibration both motivates and trumps other evidential deference norms. Given this, our strategy for developing an evidential norm for epistemic causality in §4 will be to develop a causal analogue of Chance Calibration.

Before moving on, two points are worth noting. First, although I have argued that Chance Calibration is the most fundamental evidential norm for epistemic probability, this does not imply that other deference norms are redundant or eliminable. Consider the case where the truth of a proposition A of interest is already determined, because it is a proposition about the past. Then the present chance of A is 0 or 1 and, in the absence of evidence that decides between these two possibilities, Chance Calibration merely implies that $P_E(A) \in \{P : P(A) = 0 \text{ or } 1\} = [0, 1]$ and so provides no substantive constraint. In such a situation, it remains appropriate to defer to evidence of the long run frequency of A -like outcomes, or to the credence in A of someone with expertise relating to A . Such quantities cannot be viewed as estimates of the present chance: e.g., if the frequency or credence is 0.53, say, then it is clear that this value is nowhere near the present chance value, which is 0 or 1. Hence Chance Calibration does not say anything about whether one should calibrate one's degree of belief in A to such a value. Stand-alone deference norms are required in order to ensure deference to such quantities.

Second, although Chance Calibration is straightforward to state, this does not make the epistemology of chance a simple matter. Indicators of chances include confirmed theory (e.g., confirmed physical theory can inform estimates of the chances of radioactive decay) and physical symmetries (e.g., the symmetries and mass distribution of a die can inform estimates of the chances of rolling a 5) as well as long-run frequencies, actual frequencies and expert beliefs. The picture then is that

we have a complex epistemology of chances which appeals to wide variety of indicators of chances in order to infer the set \mathbb{P}^* from evidence base B . This picture is made more complex still because an agent's set \mathbb{P}^* of evidentially-compatible chance functions will typically depend on her utilities as well as her evidence base (Williamson, 2017, §7.2). For example, the greater the disutility of erroneously excluding the true chance function, the more inclusive \mathbb{P}^* will be. On the other hand, the greater the disutility of continuing to entertain many false chance hypotheses, the less inclusive \mathbb{P}^* will be. This agent-relative chance epistemology can be thought of as a mapping $B \mapsto \mathbb{P}^*$ from an evidence base to a set of possible chance functions that are compatible with that evidence base. This mapping is induced by the mapping $B \mapsto E$ from evidence base to body of evidence, discussed above. Chance Calibration itself says nothing about what the mapping $B \mapsto \mathbb{P}^*$ ought to look like; it is left to statistical theory to provide normative constraints on this mapping.

§3.3. The epistemic analysis of chance

At the end of §2 we saw that one view of chance has it that the present chances are rational degrees of belief on ideal evidence E^* , which contains ideal inferences to chances from the complete history so far. These ideal inferences can be thought of as the product of an ideal mapping $B^* \mapsto \mathbb{P}^*$, where B^* specifies all matters of particular fact up to the present. Then the chance $P^*(A) = x$ just when $P_{E^*}(A) = x$ for every belief function P_{E^*} that is rational on the basis of evidence E^* .

A circularity worry emerges on such an interpretation of chance: Chance Calibration appeals to chances to constrain rational degree of belief, yet chances themselves are interpreted as rational degrees of belief. However, this circularity need not be pernicious. The facts about chances that are presupposed by this account are encapsulated in the mapping \mapsto^* . Classical statistical theory aims to provide an independent characterisation of this mapping: it allows us to draw inferences about chances from an evidence base without appeal to rational degree of belief. There is thus an independent handle on what is presupposed about chance and no problematic circularity arises.

There is a second concern with this account of chance. We saw above that a mapping $B \mapsto \mathbb{P}^*$ may be agent-relative to some extent: it might depend on the disutility of erroneously excluding the true chance hypothesis, or of continuing to entertain many false chance hypotheses, for example. If so, then the possibility arises that the ideal mapping $B^* \mapsto \mathbb{P}^*$ also varies according to the epistemic context. This would induce relativity in the chance function $P^* = P_{E^*}$. The worry is that this analysis then fails to provide an adequate interpretation of chance, which is normally taken to be agent-independent.

In response, one can consider two different possibilities. The first possibility is that agent-dependence does not arise after all. Perhaps agent-dependent considerations get washed out as evidence increases. Or perhaps they fail to arise in the first place on certain accounts of evidence. For instance, on an account in which evidence is factive, the evidence cannot exclude the true chance hypothesis, so the disutility of erroneously excluding the true chance hypothesis is not a consideration. Under such scenarios, the ideal evidence E^* is indeed agent-independent and no difficulty arises for the analysis of chance here. The second possibility is that agent-dependent factors are not entirely eradicated. In such a case one can either bite the bullet and say that chances are agent-dependent to some extent, or one can take

the chances to be those credences on ideal evidence that are agent-independent, in which case the chance function is only partially defined. Neither option poses an insuperable problem for the analysis. (The characterisation of chance given above took the latter course.)

There is another key problem for this analysis of chance—the problem of undermining futures (see, e.g., Finn, 2014; Belot, 2016). The worry is that the ideal mapping \mapsto^* depends on the entire history of the universe, rather than merely on matters of particular fact up to the present, so there may be a non-zero present chance (as determined by this mapping) that the universe will turn out in such a way that the mapping, and hence the present chances, are different.

I shall set this problem aside. For the purposes of this paper, we only need to consider those aspects of the epistemic analysis of chance that carry over to the analogous analysis of causality, and, as we shall see, the problem of undermining futures does not carry over. Hence we shall not consider this problem in any detail and I shall leave open the question of whether this analysis of chance succeeds, focussing instead on the analogous analysis of causation.

Note that an epistemic analysis of chance—if viable—helps to explain why one should defer to expert credences. An expert's credences might be the product of better evidence (an evidence base B' that is more comprehensive with respect to the domain of expertise), or may be generated by better inferences (a more reliable mapping \mapsto than one's own mapping \mapsto), or both. For example, a professional weather forecaster has both more data about the weather and good statistical models for inferring chances from data. On the epistemic analysis of chance, it is very plausible that more evidence (an evidence base B' that is closer than one's own to B^*) and better statistical inference methods (a mapping \mapsto that is closer than one's own to the ideal mapping \mapsto^*) are likely to provide better estimates of the chances. This explains the rationality of deference to expert credences.

§4

Calibration norms for epistemic causality

We are now in a position to consider an evidential norm for epistemic causality. In the case of probability, I appealed to a distinction between evidence-relative probability, which I referred to as probabilistic beliefs or degrees of belief, denoted by P_E , and chance, denoted by P^* , which is apparently evidence-independent. I shall draw a similar distinction here between evidence-relative causal beliefs or claims, denoted by C_E , and an evidence-independent causal relation, denoted by C^* . I shall call C^* 'the' causal relation and call consequences of C^* *causal facts*.

§4.1 will put forward the core evidential norm. §4.2 will explain why this norm is core. The relation C^* can itself be analysed in terms of causal beliefs, as we shall see in §4.3.

§4.1. Causal Calibration

Pursuing the analogy between epistemic probability and epistemic causality, I shall develop an analogue of the fundamental evidential norm of epistemic probability, Chance Calibration. Here is a preliminary attempt at a causal analogue of this norm:

Causal Calibration (precise version). If, according to E , the causal relation C^* lies within some set \mathbb{C}^* of DAG relations, $C^* \in \mathbb{C}^*$, then the causal belief relation C_E should lie within the convex hull $\langle \mathbb{C}^* \rangle$ of that set, $C_E \in \langle \mathbb{C}^* \rangle$.

Let us clarify some of the terms that occur within this statement. Recall that a DAG relation is a binary relation that can be represented by a directed acyclic graph. Such a relation implies the particular causal relationship *F is a cause of G* when there is an arrow from F to G in the DAG that represents the relation.³ For a set \mathbb{C} of DAG relations, let $\bigwedge \mathbb{C}$ be the set of causal relationships that are implied by every DAG relation in that set, and let $\bigvee \mathbb{C}$ be the set of causal relationships that are implied by some DAG relation in \mathbb{C} . A DAG relation is in the convex hull $\langle \mathbb{C} \rangle$ of \mathbb{C} if it yields every causal relationship in $\bigwedge \mathbb{C}$ and no causal relationship that is not in $\bigvee \mathbb{C}$.⁴ I will call the causal relationships in $\bigwedge \mathbb{C}^*$ *established* and the causal relationships that are not in $\bigvee \mathbb{C}^*$ *ruled out*. These two sets constitute the what the agent takes to be causal facts.

In §2 we saw that one possible structural norm for epistemic causality says that one's causal beliefs should be representable by a DAG relation C_E . This is what I called *precise* epistemic causality. The Causal Calibration norm above is constructed with this structural norm in mind. However, it is perhaps implausible to hold that an agent's causal beliefs should decide every causal question, as would be the case if her beliefs were representable by a single DAG C_E . This consideration warrants turning to another representation of an agent's causal beliefs, and so another structural norm. A variant of Causal Calibration can be devised to correspond to the alternative structural norm.

A natural approach here is to represent the causal beliefs of an agent on evidence E by a set of causal relations rather than a single causal relation. This yields *imprecise* epistemic causality. As noted in §2, the corresponding structural norm requires that causal beliefs be representable by a non-empty set \mathbb{C}_E of DAG relations. The corresponding evidential norm would be:

Causal Calibration (imprecise version). If, according to E , the causal relation C^* lies within some set \mathbb{C}^* of DAG relations, $C^* \in \mathbb{C}^*$, then the causal belief set \mathbb{C}_E should lie within the convex hull $\langle \mathbb{C}^* \rangle$ of \mathbb{C}^* , $\mathbb{C}_E \subseteq \langle \mathbb{C}^* \rangle$.

A corresponding equivocation norm would then go on to require that $\mathbb{C}_E = \langle \mathbb{C}^* \rangle$, i.e., to only fully commit to causal relationships that are established by evidence and to only fully commit to the absence of those causal relationships that are ruled out by evidence.

§4.2. Other evidential norms

Drawing on the analogy between epistemic causality and epistemic probability, one might be tempted to consider further norms that use evidence to constrain causal beliefs. For example: if, according to the evidence, F is a cause of G then one should have the causal belief *F is a cause of G*. However, this believe-your-evidence

³According to some views of causality, a DAG relation can imply causal relationships that go beyond the arrows in the DAG. For example, if causation is taken to be transitive, F is deemed to be a cause of G just when G is a descendant of F in the DAG. I will not assume transitivity here, leaving this question open.

⁴Equivalently, $C \in \langle \mathbb{C} \rangle$ just when the reachability matrix of C is in the convex hull of the reachability matrices of the members of \mathbb{C} .

norm already follows from Causal Calibration. Suppose evidence E determines that C^* lies in a subset \mathbb{C}^* of causal relations that deem F to be a cause of G . Then $F \rightarrow G$ is in $\bigwedge \mathbb{C}^*$ and so \mathbb{C}_E deems F to be a cause of G . Hence this norm is superfluous in the context of Causal Calibration.

One might also think that one ought to defer to the causal beliefs of well-informed, competent experts in situations where one has little direct evidence of one's own to go on but one does have evidence of these experts' causal beliefs. Considerations about the primacy of Chance Calibration carry over to Causal Calibration. Where an expert's causal beliefs disagree with the causal facts, insofar as one has evidence of them, the latter trump the former in guiding our causal beliefs. Of course, in the absence of evidence of the causal facts, it may yet be appropriate to defer to the causal beliefs of experts. However, the motivation behind deference to the causal claims of appropriate experts is that their causal claims are more likely to accord with the causal facts than one's own. Thus one defers to these claims because they provide evidence of causal facts. Moreover, as in the case of probability, deference to one's future causal beliefs would be classified as a diachronic norm and not an evidential norm. So there is a sense in which Causal Calibration is the most fundamental evidential norm for epistemic causality.

We can go further: Causal Calibration is the *only* evidential norm required for epistemic causality. Interestingly, the question of whether there is a need for additional deference norms exposes a disanalogy between epistemic probability and epistemic causality. In the case of epistemic probability, we saw that a proposition whose truth is already determined has present chance 0 or 1 but that it is reasonable to defer to the credence of an expert, even though such a credence cannot normally be viewed as an estimate of the chance. A stand-alone deference norm was required in order to ensure that epistemic probability deems deference to expert credence to be reasonable. In the case of epistemic causality, in contrast, an expert's causal belief *can* be viewed as an indicator of the relevant causal fact: there is no mismatch corresponding to the mismatch between non-trivial expert credences and trivial chances. Given an expert's causal beliefs, if it is reasonable to infer something about the causal facts, namely that $C^* \in \mathbb{C}^*$ for some subset \mathbb{C}^* of DAG relations, then Causal Calibration itself forces $\mathbb{C}_E \subseteq \langle \mathbb{C}^* \rangle$, which is a kind of deference to the expert's beliefs. There is no need for an autonomous principle governing deference to an expert's causal beliefs. Other indicators of causal facts include probabilistic dependence and independence relationships, temporal cues, results of manipulations, and mechanistic connections. For exactly the same reasons, there is no need for stand-alone deference norms that deal with these other indicators. Causal Calibration is all we need.

Although Causal Calibration is simple to state, we have a complicated epistemology of causation which appeals to a wide variety of indicators in order to infer the set \mathbb{C}^* from evidence E . As in the case of epistemic probability (§3.2), the picture is made more complex still because the set \mathbb{C}^* of evidentially-compatible causal relations may depend on the agent's utilities as well as her evidence base. For example, the greater the disutility of erroneously excluding the true causal relation, the more inclusive \mathbb{C}^* will be. This agent-relative causal epistemology can be thought of as a mapping $B \rightarrow \mathbb{C}^*$ induced by the mapping $B \rightarrow E$ introduced at the beginning of §3. Causal Calibration itself says nothing about the mapping $B \rightarrow \mathbb{C}^*$; it is left to scientific methodology to tell us about this mapping. Parkkinen et al. (2018), for example, provide some guidance on this mapping that is applicable to the health sciences and that builds on the techniques of evidence-based medicine.

§4.3. The epistemic analysis of the causal relation

In §3.3 we considered an analysis of chance in terms of rational degrees of belief on ideal evidence E^* consisting of inferences obtained by applying an ideal chance epistemology to the complete history so far. Although I didn't endorse this analysis of chance (which may well fall to the problem of undermining futures), an analogous analysis of causation is more tenable.

The corresponding analysis of causation takes the causal facts to be determined by rational causal beliefs on ideal evidence E^* consisting of ideal inferences to causal claims from all matters of particular fact. These ideal inferences can be thought of as a product of an ideal mapping $B^* \mapsto C^*$, where B^* specifies all matters of particular fact (throughout time, rather than merely up to the present time as was the case with the analysis of chance). Then F is a cause of G just when it is deemed a cause of G by every relation in \mathbb{C}_{E^*} , i.e., by every causal belief relation that is rational on the basis of body of evidence E^* .

As in the case of chance, a circularity worry emerges on this analysis of the causal relation. Causal Calibration appeals to the causal relation C^* to constrain rational causal beliefs, yet the causal relation C^* is itself analysed in terms of rational causal beliefs. However, as in the case of chance, this circularity need not be pernicious. The facts about the causal relation that are presupposed by this account are encapsulated by the mapping \mapsto . Methodological theory aims to provide an independent characterisation of this mapping. Hence there is an independent handle on what is presupposed about causation, and no problematic circularity.

There is a second concern with this account of the causal relation. The crucial mapping $B^* \mapsto C^*$ may be agent-relative to some extent. In which case, so is E^* and so are the causal facts, which are defined in terms of \mathbb{C}_{E^*} . The worry is that this analysis would then fail to provide an adequate interpretation of causation, which is normally taken to be agent-independent. In response, one can consider two different possibilities. Perhaps agent-dependent considerations get washed out as evidence increases, or don't arise in the first place on certain accounts of evidence, and the ideal evidence E^* is agent-independent after all. If so, there is clearly no problem here for the analysis of causation. The other possibility is that agent-dependent factors are not entirely eliminable. In such a case one can either bite the bullet and say that the causal relation is agent-dependent to some extent, or one can take the causal facts to be those causal beliefs on ideal evidence that are agent-independent, in which case the causal relation is only partially defined. Neither option poses an insuperable problem for the analysis.

There is a further concern that one might have about any epistemic analysis of the causal relation. We often invoke causal relationships to explain the occurrence of effects. One might explain the occurrence of G by observing that F is a cause of G and that F occurred. The worry is that an epistemic notion of cause can't explain anything: there needs to be some causal 'biff' or 'oomph' out there in the world in order for causation to be explanatory (Williamson, 2013).

Again, the analogy with epistemic probability can help us to address this concern. There is an analogous worry about an epistemic analysis of chance, that can be allayed by noting that it is not the chance of A that explains the occurrence of A ; rather, it is the activation of a physical mechanism (aka 'chance setup') that has some chance of producing A that explains A . If the epistemic analysis of chance is right, the chance itself merely encapsulates some facts about which predictions,

bets and decisions are reasonable in such a situation. To the extent that we explain by appealing to chances, we do so in an elliptical way: there being a positive chance of A implies that there is some physical mechanism that can produce A . Similarly, it is not the fact that F is a cause of G that explains an occurrence of G , it is the mechanism responsible for G that explains G . The causal relationship merely encapsulates a range of reasonable predictions, explanations and control inferences. To the extent that we explain by appealing to causal relationships, we do so in an elliptical way: F being a cause of G implies that there is some complex of facts about underlying mechanisms and their activation and/or disactivation that explains G .

In sum, then, the epistemic analysis of the causal relation developed here can deal with charges of circularity, agent-relativity and explanatory deficiency in precisely the same way in which an epistemic analysis of chance can deal with them. On the other hand, while the problem of undermining futures presents a serious challenge for the epistemic analysis of chance, there is no analogous problem here, because the causal facts are not time-relative in the way that chances are. This puts the epistemic analysis of the causal relation on a stronger footing than the epistemic analysis of chance.

§5

Conclusion

I have argued for Chance Calibration as the fundamental evidential norm for epistemic probability and developed Causal Calibration, an evidential norm for epistemic causality that is analogous to Chance Calibration. The close analogy between epistemic probability and epistemic causality has been the guiding light throughout this paper. However, several differences have emerged between the two accounts.

Firstly, *imprecise* epistemic causality has been advocated here, because it is not possible to use a single DAG to represent non-committal causal beliefs. In contrast, it is possible to represent non-committal degrees of belief in the framework of precise epistemic probability: one can equivocate between A and $\neg A$ by adopting a middling degree of belief in A .

Second, we have seen that while epistemic probability requires deference norms other than Chance Calibration in order to ensure deference to quantities such as expert credences and long-run frequencies in situations where chances trivialise, such norms are superfluous in the case of epistemic causality. Causal Calibration is the only evidential norm that is needed for epistemic causality.

Third, I have suggested that it is possible to analyse the objective causal relation in terms of rational causal beliefs. This paper sat on the fence with respect to the analogous question about probability: can one analyse chance in terms of rational degrees of belief? This was because an epistemic analysis of chance faces a problem that is not faced by an epistemic analysis of the causal facts: the problem of undermining futures.

Fourth, while epistemic probability is invariably taken to be an interpretation of *single-case* probability, with probabilities attaching to propositions rather than repeatable events, I have not said anything about whether epistemic causal beliefs are single-case or generic. This is because that question is orthogonal to the issues considered here.

This paper has also remained neutral about other important questions. In particular, I have not endorsed any particular position on the nature or factivity

of evidence, appealing only to the distinction between an evidence base and a body of evidence. Moreover, as explained in the appendix, an epistemic account of causality is compatible with other accounts of causality, just as an epistemic account of probability is compatible with a physical account of chance, for example. Finally, I have not taken a stance as to whether to endorse an equivocation norm in addition to the structural and evidential norms considered here. It is hoped that, by remaining neutral about these questions, the epistemic theory of causality might achieve wider interest.

Acknowledgements

This paper is a product of the research project *Evaluating evidence in medicine*, funded by the UK Arts and Humanities Research Council. I am very grateful to Michael Wilde for helpful comments on an earlier version of this paper, and to two anonymous referees for many useful suggestions.

Appendix: Motivation and relation to other theories of causality

This appendix sketches some motivation for epistemic causality and compares epistemic causality to some other theories of causality. For further discussion of the epistemic theory of causality, see [Williamson \(2005, 2006a,b, 2013\)](#); [Russo and Williamson \(2007\)](#); [Wilde and Williamson \(2016\)](#).

There are two main ways to motivate epistemic causality: a positive way, which points to why such an account is natural on its own terms, and a negative way, which points to deficiencies of other theories of causality that are remedied by the epistemic theory. Let us sketch these lines of argument in turn.

From a positive point of view, it is clear that we do indeed have causal beliefs; that these beliefs are useful to us because they help us to predict, explain and control our world; and that some of these beliefs are rational insofar as they are structurally coherent, compatible with evidence, and don't go too far beyond the evidence. These grounds alone render the study of epistemic causality a worthwhile activity. None of this precludes the viability of other theories of causality. Just as proponents of epistemic theories of probability often admit objective chances, limiting relative frequencies, and/or actual frequencies, so too a proponent of epistemic causality can admit other kinds of causality to which epistemic causality should somehow relate. On the other hand, should epistemic causality provide a successful interpretation of our causal claims—perhaps via the analysis of §4.3—, other accounts of causality might be considered otiose.

This brings us to the more negative point of view. Other theories of causality are notoriously problematic. The principal accounts of causality can be divided into *association* (aka 'difference-making') theories and *mechanistic* (aka 'production') theories. Association theories include regularity, counterfactual, probabilistic and agency theories, all of which require as a minimum that cause and effect are probabilistically dependent ([Williamson, 2009](#)). These theories face certain particularly tricky cases of overdetermination where a cause cannot increase the probability of the effect because the effect will happen anyway (see, e.g., [Hall, 2004](#); [Williamson, 2009, §9](#)). Mechanistic theories analyse causality in terms of physical processes or complex-systems mechanisms. Unfortunately for mechanistic theories, there are cases where cause and effect cannot be linked by a physical mechanism because one or other is an absence ([Hall, 2004](#); [Williamson, 2011](#)). A disjunctive account, which

analyses causality in terms of either association or mechanism, will not work either, because there are cases of causality involving both overdetermination and absences, which are thus accompanied by neither association nor mechanism (Longworth, 2006). Moreover, causality is not obviously pluralistic in the way that probability is. When asked for the probability of surviving a certain disease, a doctor might say something like, ‘70% of people your age survive this disease, but you are particularly fit and healthy otherwise, so I’m more confident that you will survive.’ There is little debate that there are frequencies and credences, that these are both probabilities, and that they can differ. On the other hand, in response to the claim that the disease affects survival, it seems bizarre to ask whether this is a claim about association-causality or a claim about mechanistic-causality—this sort of pluralism about causality is at odds with our use of causal claims.

Association and mechanistic theories of causality also face an epistemological challenge. Russo and Williamson (2007) argue that, at least in the health sciences, in order to establish a causal claim one normally needs to establish both that the putative cause and effect are appropriately associated and that there is a suitable mechanism linking the two. If correct, this thesis poses a problem for both association and mechanistic theories. If an association theory were true, then association would be sufficient for causation: if the putative cause and effect are established to be associated in the appropriate way, that should be enough to establish causation, even if it is extremely doubtful that there is any mechanism linking the two. In practice, however, associations would not be enough to establish the efficacy of homeopathy or extra-sensory perception because current science casts strong doubt on the existence of an underlying mechanism of action. Thus association theories get judgements of causation wrong. Alternatively, if a mechanistic theory were true, then a mechanism would be sufficient for causation: if the putative cause and effect are found to be linked by an appropriate mechanism, that should be enough to establish causation, even if it is doubtful that there is any association between the two. In practice, however, the existence of a mechanism by which a drug interacts with a disease would not be enough to establish efficacy of the drug where studies show that the drug makes no significant difference to the disease. Thus mechanistic theories also get judgements of causation wrong. A simple-minded conjunctive account of causality, which requires both association and mechanism, would avoid this epistemological problem. However, the conjunctive account would fall to the counterexamples to association and mechanistic theories, outlined above: some causal relationships are not accompanied by an association; some have no mechanism.

The epistemic theory of causality is not susceptible to these problems that face association and mechanistic theories. This is because the epistemic theory does not analyse causality in terms of association or in terms of mechanism or in terms of a simple disjunction or conjunction of the two. According to epistemic causality, causal claims are relative to evidence; this evidence will normally include evidence of associations as well as evidence of mechanisms; certain causal beliefs will be rational on the basis of this complex combination of evidence, while others will not. Associations are indicators, not truth-makers, of causal claims; similarly for mechanisms.

Note that the positive and negative pathways to epistemic causality are independent: a reader wedded to an association or a mechanistic theory of causality may yet be moved by the positive motivation to take an interest in epistemic causality.

So far, I have considered the relationship between epistemic causality and the dominant theories of causality, namely association theories, mechanistic theories,

and pluralistic combinations of the two. Next I shall very briefly sketch the commonalities and differences between epistemic causality and two other theories of causality: Anscombian pluralism and inferentialism.

Anscombian pluralism takes causal relationships that are characterised by different activities—e.g., scraping and pushing—to be instances of different causal relations. Thus there are many kinds of causal relation—e.g., the scraping relation, the pushing relation—that are all *causal* only insofar as they stand in some family resemblance to one another (Anscombe, 1971, p. 93; Cartwright, 1999, §5.2.5; Machamer et al., 2000, §3.1; Cartwright, 2004; Psillos, 2009).

The proponent of an epistemic theory of causality can agree with the Anscombian pluralist that there are many kinds of causing and that there is no simple reductive analysis of causality to one or other of the indicators of causality. However, there is no need to endorse causal pluralism, which is implausible for the reasons alluded to above. Indeed, epistemic causality can be thought of as providing an account of why all the various kinds of causing are kinds of *causing*: causal claims motivate characteristic predictions, explanations and control inferences.

The fact that causal claims motivate certain inferences suggests a connection between epistemic causality and inferentialism. Reiss (2011, 2012) advocates an inferentialist theory of causality, according to which the meaning of a causal claim is characterised by the propositions from which one infers the claim (the *inferential base*) and the propositions that one infers from the claim (the *inferential target*).

However, there are some difficulties that face this view (Williamson, 2013, §6). For example, it is a radically pluralist view: ‘causes’ means different things in ‘inhaling tobacco smoke is a cause of cancer in mice’ and ‘inhaling tobacco smoke is a cause of cancer in humans’ because the inferential base differs in each case: randomised trials can be performed in mice but not in humans for ethical reasons, and so form a part of the inferential base of the former claim but not the latter claim. There are two main differences between epistemic causality and inferentialist causality, both of which avoid this awkward consequence. First, epistemic causality characterises causality in terms of the inferences typically warranted by causal claims, rather than in terms of the inferential base and target of a *particular* causal claim. A second difference between the inferentialist view and epistemic causality is that epistemic causality takes only the inferential target to characterise causal claims—not the inferential base. Thus the discrepancy between the evidential bases in the case of cancer in mice and cancer in humans does not lead to multiple concepts of cause. This second difference ensures that epistemic causality does not face a further problem for inferentialist causality: on the inferentialist view one cannot improve the epistemology of causality. According to inferentialism, when methods for establishing causal claims change, the inferential base changes, and so does the meaning of ‘cause’. Thus there is no sense in which one can have better methods for establishing the same causal propositions. For epistemic causality, on the other hand, the inferential base can change without changing the meaning of a causal claim.

Bibliography

- Anscombe, G. (1971). Causality and determination. In Sosa, E. and Tooley, M., editors, *Causation*, pages 88–104. Oxford University Press, 1993 edition.
- Belot, G. (2016). Undermined. *Australasian Journal of Philosophy*, 94(4):781–791.
- Cartwright, N. (1999). *The dappled world: a study of the boundaries of science*. Cambridge University Press, Cambridge.
- Cartwright, N. (2004). Causation: one word, many things. *Philosophy of Science*, 71:805–819.
- Finn, S. (2014). Humean supervenience rebugged. *Philosophia*, 42(4):959–970.
- Friedman, K. and Shimony, A. (1971). Jaynes’s maximum entropy prescription and probability theory. *Journal of Statistical Physics*, 3(4):381–384.
- Hall, N. (2004). Two concepts of causation. In Collins, J., Hall, N., and Paul, L., editors, *Causation and counterfactuals*, pages 225–276. MIT Press, Cambridge MA and London.
- Hawthorne, J., Landes, J., Wallmann, C., and Williamson, J. (2017). The Principal Principle implies the Principle of Indifference. *British Journal for the Philosophy of Science*, 68:123–131.
- Jeffrey, R. (2004). *Subjective probability: the real thing*. Cambridge University Press, Cambridge.
- Lewis, D. K. (1980). A subjectivist’s guide to objective chance. In *Philosophical papers*, volume 2, pages 83–132. Oxford University Press (1986), Oxford. With postscripts.
- Longworth, F. (2006). Causation, pluralism and responsibility. *Philosophica*, 77:45–68.
- Machamer, P., Darden, L., and Craver, C. (2000). Thinking about mechanisms. *Philosophy of Science*, 67:1–25.
- Mitova, V. (2017). *Believable evidence*. Cambridge University Press, Cambridge.
- Parkkinen, V.-P., Wallmann, C., Wilde, M., Clarke, B., Illari, P., Kelly, M. P., Norell, C., Russo, F., Shaw, B., and Williamson, J. (2018). *Evaluating evidence of mechanisms in medicine: principles and procedures*. Springer Briefs. Springer.
- Pettigrew, R. (2016). *Accuracy and the laws of credence*. Oxford University Press, Oxford.
- Psillos, S. (2009). Causal pluralism. In Vanderbeeken, R. and D’Hooghe, B., editors, *Worldviews, Science and Us: Studies of Analytical Metaphysics. A Selection of Topics From a Methodological Perspective*, pages 131–151, Singapore. World Scientific.
- Reiss, J. (2011). Third time’s a charm: causation, science and Wittgensteinian pluralism. In Illari, P. M., Russo, F., and Williamson, J., editors, *Causality in the Sciences*, pages 907–927. Oxford University Press, Oxford.
- Reiss, J. (2012). Causation in the sciences: An inferentialist account. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 43(4):769–777. Causality in the Biomedical and Social Sciences.
- Rowbottom, D. P. (2014). Information versus knowledge in confirmation theory. *Logique et Analyse*, 57(226):137–149.
- Russo, F. and Williamson, J. (2007). Interpreting causality in the health sciences. *International Studies in the Philosophy of Science*, 21(2):157–170.

- Wilde, M. and Williamson, J. (2016). Evidence and epistemic causality. In Wiedermann, W. and von Eye, A., editors, *Statistics and Causality: Methods for Applied Empirical Research*, pages 31–41. Wiley, Hoboken, New Jersey.
- Williamson, J. (2005). *Bayesian nets and causality: philosophical and computational foundations*. Oxford University Press, Oxford.
- Williamson, J. (2006a). Causal pluralism versus epistemic causality. *Philosophica*, 77:69–96.
- Williamson, J. (2006b). Dispositional versus epistemic causality. *Minds and Machines*, 16:259–276.
- Williamson, J. (2009). Probabilistic theories. In Beebe, H., Hitchcock, C., and Menzies, P., editors, *The Oxford Handbook of Causation*, pages 185–212. Oxford University Press, Oxford.
- Williamson, J. (2010). *In defence of objective Bayesianism*. Oxford University Press, Oxford.
- Williamson, J. (2011). Mechanistic theories of causality. *Philosophy Compass*, 6(6):421–447.
- 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*, 31(1):27–65.
- Williamson, J. (2017). *Lectures on inductive logic*. Oxford University Press, Oxford.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford University Press, Oxford.