

# CAUSAL PLURALISM VERSUS EPISTEMIC CAUSALITY

Jon Williamson

To appear in *Philosophica*  
Draft of January 9, 2008

## ABSTRACT

It is tempting to analyse causality in terms of just one of the indicators of causal relationships, e.g., mechanisms, probabilistic dependencies or independencies, counterfactual conditionals or agency considerations. While such an analysis will surely shed light on some aspect of our concept of cause, it will fail to capture the whole, rather multifarious, notion. So one might instead plump for pluralism: a different analysis for a different occasion. But we do not seem to have lots of different kinds of cause—just one eclectic notion. The resolution of this conundrum, I think, requires us to accept that our causal beliefs are generated by a wide variety of indicators, but to deny that this variety of indicators yields a variety of concepts of cause. This focus on the relation between evidence and causal beliefs leads to what I call *epistemic causality*. Under this view, certain causal beliefs are appropriate or rational on the basis of observed evidence; our notion of cause can be understood purely in terms of these rational beliefs. Causality, then, is a feature of our epistemic representation of the world, rather than of the world itself. This yields one, multifaceted notion of cause.

## CONTENTS

§1	The Indicators of Causality	2
§2	A Plurality of Causality?	3
§3	The Epistemic View of Causality	5
§4	Epistemic Metaphysics	9
§A	A Formal Epistemology	11

## THE INDICATORS OF CAUSALITY

The indicators of causality are several and disparate. We base our causal claims on observed associations, observed independencies, temporal cues, known mechanisms, theoretical connections, experiments, controlled trials, other causal knowledge, intuitions about subjunctive conditionals, and more. In trying to understand the nature of causality it is reasonable to attempt to analyse causal connections in terms of one or other of these indicators. Thus we have a range of contemporary theories including mechanistic, probabilistic, counterfactual, and agency-based (the last three of which are often classed as *difference-making* accounts of causality, since they are based on the intuition that a cause should make a difference to its effects).

Unfortunately, these monistic theories of causality have great difficulty in accounting for the epistemology of causality. While such a theory does well at explaining how the particular indicator used in the analysis can be taken as evidence for causal claims, it has trouble explaining how other indicators can also count as evidence. Thus mechanistic theories (Salmon, 1998; Dowe, 2000) have little problem explaining how knowledge of mechanisms and physical theory can ground causal claims, but they have their work cut out explaining how, even when we know about the salient mechanisms, we seek further evidence, e.g., evidence of probabilistic dependencies or independencies. It is not enough to know of some chain of connections linking exchange rate and inflation, we want to know that exchange rate makes a difference to inflation before we claim that it is a cause. In contrast, probabilistic theories (see, e.g., Suppes, 1970) do well at accounting for probabilistic indicators, but poorly when it comes to mechanistic indicators. It was not enough to know that in samples the prevalence of smokers made a difference to the prevalence of lung cancer, we needed to know that the two are linked by a physical mechanism before the claim that smoking causes cancer could be accepted. Similarly, counterfactual (Lewis, 1973) and agency (Price, 1992) approaches struggle with respect to mechanisms. This epistemological problem is presented in more detail in Russo and Williamson (2007a).

Perhaps the main reason why we seek varied evidence for a causal claim is this. Causal claims have two uses: they are used for inference on the one hand and explanation on the other. The inferential use—making predictions, diagnoses and strategic decisions on the basis of causal claims—requires that a cause should typically make a difference to its effects, for otherwise information about the presence of a cause would tell us nothing about the presence of its effects and vice versa, and instigating a cause would not be a good strategy for achieving its effects. The explanatory use requires something more, namely some physical account of why the event in question happened. When asked for an explanation of an event, it is not enough to say that some other event occurred and that there exists a difference-making relationship between the two, for that is no explanation at all—it leaves the question, why did the explaining event make a difference to the explained event? Inasmuch as we can answer ‘why’ questions at all, we do so at root by invoking physical theory, physical events and physical processes. If a causal story is to offer an explanation, it had better fit with physical theory and tell us a bit about the ultimate physical explanation.

Hence causal claims need to say something about physical mechanisms as well as about difference-making.

The above epistemological problem for monistic accounts motivates the move to a less simplistic account of causality—an account that takes the full variety of causal indicators seriously. Pluralism is a step in this direction. However, I shall argue that it is the wrong step (§2). Instead, I shall argue in §3 that an *epistemic* theory of causality offers the right way to handle the full range of indicators of causality. In §4 I shall suggest that, in general, an epistemic theory of a complex concept can have more to offer than a simple-minded analysis of the concept in terms of a single indicator, or even a more sophisticated pluralist analysis. The appendix, §A, outlines a formal causal epistemology that forms a component of the epistemic theory of causality.

## §2

### A PLURALITY OF CAUSALITY?

The move to causal pluralism is often motivated by the inadequacies of contemporary monistic accounts of causality and incompatibilities between mechanistic and difference-making accounts. In §1 I suggested that contemporary accounts lack a viable epistemology: mechanistic accounts make it a mystery as to why we back up our causal claims with evidence of difference-making *over and above* evidence of mechanisms, while difference-making accounts cannot explain why we seek evidence of mechanisms *as well as* evidence of difference-making. But there are other paths to pluralism that pick up on other inadequacies of monistic accounts. Hall (2004) argues that his favoured difference-making approach, the counterfactual theory, cannot account for basic features of causality (namely its transitivity, the spatio-temporal continuity of causal processes, and the causal character of a process being determined by its intrinsic non-causal features), and that while a mechanistic approach can account for the latter properties it cannot account for counterfactual dependence being sufficient for causation, nor can it account for absences being causes and effects. Consequently—Hall (2004, §6) claims—there are two concepts of cause, one which corresponds to counterfactual dependence and the other which corresponds to mechanistic production. Cartwright (2004, §2) argues that contemporary accounts are incompatible with one another and that no individual account has universal applicability. She concludes that each account specifies a different kind of causal law.

Not only are there several paths to pluralism but there are also several varieties of pluralism. Pluralists agree that there is no single thing that is picked out by causality,<sup>1</sup> but that leaves plenty of scope for disagreement. Some (e.g., Psillos, 2006; Godfrey-Smith, 2008, §6) do not think much more can be said about what causality is, while others (e.g., Hall, 2004) argue that there are

---

<sup>1</sup>Note that under this conception, pluralism does not encompass the view that there are different *forms* of causal relationship—e.g., ‘type cause’ versus ‘token cause’, or ‘component effect’ versus ‘net effect’ (Hitchcock, 2001)—except where it is maintained that these different forms appeal to entirely different concepts of cause (Sober, 1985; Hitchcock, 2003); for a single concept of cause might be used in different ways to shed light on the various forms of relationship. Nor does this conception of pluralism encompass the view that there are different forms of causal *explanation* (Weber et al., 2005)—except where it is maintained that different forms of explanation require different notions of cause.

distinct and coherent senses of cause and would like to understand each of these senses.

I take it that the former, nebulous variety of pluralism is a last resort. If one can't say much about the number and kinds of notions of cause then one can't say much about causality at all; this stance should only be adopted if there is no viable alternative. I do not think the latter, determinate variety of pluralism includes any viable alternative, for the reasons set out below. But I do think that there is a viable monistic account, as developed in §3, so there is no need to resort to nebulous pluralism.

Of course those in the latter, determinate-pluralism camp differ substantially as to the number and nature of the senses of cause. They also differ as to the task at hand. One might think that it is enough to shed light on the alleged different notions of cause, e.g., by saying that there are two notions of cause, one mechanistic and the other difference-making. But one might want to go further by rendering the various notions precise, in order to explicate the notion of cause in the sense of Carnap (1950, §2), or to provide a reductive analysis of cause. Further, one might want to delimit the proper zone of application of each concept of cause, e.g., by saying that mechanistic causality is appropriate in the natural sciences while difference-making causality is appropriate in the social sciences (a move analogous to the pluralism about probability of Gillies (2000, Chapter 9)).

There are a number of problems with determinate pluralism. First, pluralism is not parsimonious—if, as I suggest in §3, there is an adequate monistic account of cause, then arguably that account should be preferred purely on the grounds of parsimony. Second, while we have many words that are suggestive of causation, e.g., 'push' and 'pull' (Anscombe, 1971; Cartwright, 2004, §3), we have only one word-stem 'cause' for the fully general notion.<sup>2</sup> If causality were a plural concept then one would think that we would have several word-stems, including one for each general notion. Perhaps, the pluralist might reply, this is just a case where our language has not adequately evolved to match our world. But if so, one would still expect some qualifiers (e.g., 'mechanistic', 'difference-making') to the word 'cause' to be routinely used to distinguish the types of cause under consideration. At the very least, one would expect clarificatory questions to be used to disambiguate uses of the word 'cause' (Godfrey-Smith, 2008, §3). But all this is absent.

Third, it is clear that determinate pluralism will not do justice to the worry with which we began, namely the problem of accounting for causal epistemology. How can one explain the fact that there was excellent evidence that smoking and lung cancer sat in the right sort of difference-making relation, yet some suitable physical mechanism linking smoking and lung cancer was required before the causal claim could be established? How can one explain the fact that there is excellent evidence of a mechanism linking smoking and breast cancer (involving the presence of breast cancer carcinogens in tobacco smoke), yet that causal claim has hitherto not been established because the evidence concerning difference-making is equivocal? The pluralist is stuck. If she says that some particular claim invokes a difference-making account of causality then she can-

---

<sup>2</sup>Arguably 'prevent' is also fully general. However there is general consensus that 'cause' and 'prevent' are not different concepts. Either they are two sides of the same coin, causation being a positive relationship and prevention negative, or causation is taken to subsume prevention.

not explain the requirement of a mechanism. If she says that the claim involves a mechanistic account of causality then she can not explain the requirement of difference-making. If she says that the claim simultaneously involves two notions of cause—mechanistic and difference-making—then she is in danger of not being a pluralist at all, but of espousing a single concept of cause that has two necessary conditions, one mechanistic and one difference-making. The different aspects of causality—mechanistic and difference-making—are clearly connected, since a causal claim requires both as evidence. But pluralism sheds no light on this connection; if anything, it pushes these two aspects apart, viewing each as evidence for a different claim.

If determinate pluralism doesn't cut muster and contemporary monistic accounts also fail, then we need to go back to the drawing board. Our options are a more elaborate form of monism, or, as a last resort, nebulous pluralism. As a first attempt, as suggested above one might try to develop a monistic conception which takes a causal connection to require *both* an underlying mechanism *and* that the cause make a difference to the effect. Unfortunately, this will not do either, for the simple reason that not all of our causal claims have an underlying mechanism and not all of our claims reflect difference making—see, e.g., Hall (2004), who dismisses monism on these grounds. While we *seek* evidence of a mechanism as well as evidence of difference-making, such evidence is sometimes unattainable—this fact puts paid to a monistic analysis of causality in terms of one or other or both of these notions.

Perhaps, then, we must look to some less determinate account which appeals to a vague cluster of different notions that underlie a single concept of cause. This nebulous monism may be marginally more appealing than nebulous pluralism, but again not one to which we need resort, since, as we shall see, the epistemic theory offers a more determinate kind of monism.

### §3

## THE EPISTEMIC VIEW OF CAUSALITY

In a sense causality is a very simple concept—it is just an asymmetric binary relation.<sup>3</sup> Therefore, it can only carry so much information. But we demand a lot of this relation. Causality is used throughout the sciences and in daily life for inference and for explanation: we represent the world causally so that we can make predictions, diagnose faults, make strategic decisions, explain events and apportion blame and praise. Thus we overload a simple relation with connotations both of difference-making and of mechanisms. As pluralists have observed, there are some tensions between the inferential and the explanatory uses of the causal relation. This explains the multi-faceted epistemology of causality and the apparent complexity of the notion of cause.

The *epistemic theory of causality* (Williamson, 2005a, 2007a, 2006a) takes causal epistemology as primary and builds up causal metaphysics from this

---

<sup>3</sup>Some have argued that causality is slightly more complicated than that—e.g., Schaffer (2005) argues that it is not binary, Mellor (1995) that it is not a relation, Spirtes (1995) that it is not acyclic—but these alternative conceptions are still simple enough to make the subsequent point, and the analysis of this section can be modified to take these complications into account if need be. Pluralists argue that causality is not a *single* relation, but I have argued in §2 that pluralism is unacceptable.

epistemology. Arguably only by this process of reverse engineering can one address the epistemological problem of §1.

The epistemic theory takes an epistemology of *rational belief* as its starting point. The idea is that an agent’s evidence constrains the range of *causal beliefs* that it would be rational for her to adopt. Some possible causal beliefs are incompatible with the evidence, others are suggested by the evidence; the agent should choose from the latter. These beliefs are just that—they are highly defeasible in the light of new evidence and nothing like as stable as causal knowledge. Nevertheless, this relation between evidence and rational causal belief is enough both to develop a full causal epistemology and to isolate the concept of cause itself.

Just what is this relation between evidence and rational causal belief? What causal beliefs should an agent adopt on the basis of her evidence? The answer to a question about what an agent should do hinges, of course, on what the action is intended to achieve—in this case, on the uses to which she will put her causal beliefs. The explanatory use of causality requires that the causal relation should typically fit with known physical theory and evidence of mechanisms: typically, some cause of an event will be invoked by its physical explanation. The inferential use of causality requires that the causal relation should typically fit with evidence of difference-making: typically, cause and effect should be probabilistically dependent, when intervening to fix the cause and controlling for the effect’s other causes. There is also the rather general use of beliefs to systematise one’s evidence: an agent’s beliefs should typically be able to offer some kind of explanation of her experience and evidence. For example, if the agent discovers that two events are probabilistically dependent, and she knows of no non-causal explanation of this dependence (the events are not known to be overlapping, for instance) then she should (tentatively) believe that some causal connection between the events gives rise to the dependence, because dependencies between physical events are typically explained causally. This sketch involves a lot of ‘typically’s, because none of these features of causality hold invariably; if they did, a more straightforward analysis of causality in terms of one or more of these features might be possible; yet ‘typically’ is quite enough for causal beliefs to be useful from an inferential and explanatory point of view.

One way of making this sketch more precise proceeds as follows—a more detailed exposition is given in the appendix §A and the motivation behind some of the assumptions can be found in Williamson (2005a). An agent’s causal beliefs can be represented by a directed acyclic graph whose nodes are the variables of interest in her domain and whose arrows correspond to direct causal connections. As explained in §4, her *evidence* or *epistemic background*,  $\beta$ , which contains everything that the agent takes for granted in the context in hand, can be used to determine a probability function,  $p_\beta$ , over the variables in her domain—namely the probability function that satisfies constraints imposed by evidence but that is otherwise as non-committal as possible, i.e., that has maximum entropy.  $p_\beta$  represents the degrees of belief that the agent should adopt on the basis of  $\beta$ . The causal belief graph,  $\mathcal{C}_\beta$ , that the agent should adopt on the basis of  $\beta$  is determined as follows. First,  $\mathcal{C}_\beta$  should be compatible with the constraints  $\kappa$  that are imposed by the agent’s mechanistic and theoretical knowledge: e.g., causes should not occur after their effects; if physical theory treats two variables symmetrically then neither can be a cause of the other (for otherwise each would be a cause of the other, breaking the asymmetry of

causality); if mechanisms indicate a common cause of two variables rather than a direct causal relation from one to the other then this should be reflected in the causal belief graph. Second, as long as it is not prohibited by the mechanistic-theoretical constraints  $\kappa$ , there should be an arrow  $A \longrightarrow B$  in the causal graph to account for each *strategic dependence* from  $A$  to  $B$ , i.e., whenever  $p_\beta$  renders  $A$  and  $B$  probabilistically dependent when intervening to fix  $A$  and controlling for  $B$ 's other direct causes (i.e., whenever  $A$  and  $B$  are probabilistically dependent conditional on  $B$ 's other direct causes and  $A$ 's direct causes). Third, the agent's causal beliefs should otherwise be as non-committal as possible: there should be no arrows in  $\mathcal{C}_\beta$  that are not warranted by evidence  $\beta$ .

Interestingly, standard methods can be used to determine  $\mathcal{C}_\beta$ :

**THEOREM 3.1**  *$\mathcal{C}_\beta$  is a minimal graph that satisfies  $\kappa$  and the Causal Markov Condition (c.f., Definition A.7), if there is such a graph at all.*

PROOF: See §A. □

There are a whole host of algorithms for finding minimal causal graphs that satisfy the Causal Markov Condition and some set of causal constraints (Korb and Nicholson, 2003, Appendix B). The system Hugin, for instance, offers a commercial implementation of a range of techniques (see Andersen et al. (1989), www.hugin.com). Thus these methods fit well with the above epistemology.

Once we have an epistemology that elucidates the relationship between evidence and rational causal belief, one can use this epistemology to determine the concept of cause itself via the following identity: the causal relation is just the causal belief graph of an omniscient rational agent (an agent whose evidence is exhaustive).

This identity can be understood in two ways. It could be thought of as a fact about a concept of cause on which we have an independent handle. For example, the proponent of a mechanistic analysis of cause might want to claim that, if we had full empirical evidence, our rational causal beliefs would coincide with this mechanistic relation. But this claim would be very hard to maintain, thanks to the epistemological problem of §1: it is implausible that, if we had full evidence of mechanisms, further probabilistic evidence should not alter our causal beliefs. Alternatively, one might think of this identity as constitutive of causality—there is no independent handle, causality just is a set of rational beliefs. It is this second understanding that forms the crux of the epistemic theory of causality. According to this epistemic view, the epistemology of causality is determined by the uses we put this relation to—inference and explanation—and causality itself is determined by this epistemology, and so is ultimately reducible only to its uses.

According to the epistemic theory, then, the causal relation is characterised by the causal beliefs that an omniscient rational agent should adopt. It should be clear in principle how this characterisation can overcome the epistemological problem that besets other accounts. The epistemological problem is that of developing an account of causality that fits with the following epistemological fact: in certain cases, one should not infer a causal connection solely on the basis of evidence of difference-making, or solely on the basis of evidence of mechanisms.<sup>4</sup> I.e., in such a case one would not be rational to hold the corre-

<sup>4</sup>As alluded to at the end of §2, in other cases one kind of evidence is sufficient for the causal claim.

sponding causal belief. If so, and if in such a case there were difference-making but no mechanism, or vice versa, then an omniscient rational agent would not hold the causal belief. Hence according to the epistemic theory there would be no such causal connection. On the other hand, if in such a case there were both difference-making and a mechanism then there would be a causal connection. So in such a case there is a causal connection if and only if there is both a mechanism and difference-making. Thus there is a tight fit between the epistemic theory of causality and the epistemological fact. The epistemological fact is accounted for by the uses of causal beliefs: the explanatory and inferential uses of causal claims require that, where possible, causal claims should coincide both with mechanisms and with difference-making.<sup>5</sup>

The epistemic view leaves us with an elaborate epistemology but a simple metaphysics: many indicators of causality but one concept of cause. That causal epistemology is pluralistic and somewhat elaborate is no news to anyone. That causality itself is monistic and rather straightforward, yet fits with this epistemology, is perhaps more surprising. Moreover, this is not nebulous monism. One of the advantages of this view is that it is somewhat easier to agree on an appropriate causal epistemology, elaborate though it is, than to agree on an appropriate understanding of cause, simple though it may turn out to be. Since the latter task can be reduced to the former, metaphysical progress becomes possible and a determinate monism is within reach. Indeed, as I hope appendix §A shows, a causal epistemology may be made very precise, in which case the monistic concept of cause is precisely defined too.

There is an interesting question concerning the objectivity of the causal relation under the epistemic view. How much choice does an agent have when deciding which causal beliefs to adopt? Clearly an agent can not choose just any directed acyclic graph as her causal belief graph  $\mathcal{C}_\beta$ . In fact, if there is a graph that satisfies  $\kappa$  and the Causal Markov Condition, and if  $p_\beta$  is faithful, then  $\mathcal{C}_\beta$  must be chosen from a Markov equivalence class—a set of directed acyclic graphs that have the same independencies via the Causal Markov Condition (see Proposition A.17). Gillispie and Perlman (2002) carried out studies that suggest that on average a Markov equivalence class has four members, i.e., the agent will be able to choose the directions of two arrows in the graph, on average, and all other arrows will be determined by evidence. Thus causality is very highly determined on the epistemic account. This seems to be just what we want—by and large the causal relation is objectively determined, but there are cases (see, e.g., Hitchcock, 2003) that suggest that causality is not fully objective.

---

<sup>5</sup>A referee astutely pointed out that it is unclear how an omniscient agent could have any inferential needs for causal beliefs to satisfy. It is important to note that the uses of causal beliefs outlined here—explanation and inference—are not the uses to which every hypothetical agent puts such beliefs. Indeed they are not the uses to which every real agent puts such beliefs, since there are agents that make no inferences or proffer no explanations, and many agents have ulterior (non-explanatory, non-inferential) motives for adopting certain causal beliefs. The uses emphasized here are the general epistemological uses to which the bulk of us (non-omniscient) humans put our causal beliefs.

## §4

### EPISTEMIC METAPHYSICS

We have seen, then, that traditional monistic accounts of causality, which seek to analyse causality in terms of just one of its indicators, face a crucial epistemological problem, namely that of accounting for the variety of evidence required for causal claims. Determinate pluralism fares no better: it also succumbs to this epistemological problem. The epistemic account, on the other hand, provides a monistic theory that has causal epistemology at its base, and is not beset by this problem. If one can provide a determinate epistemology, such as that outlined in appendix §A, then this leads to a determinate monism about causality; there is no need to resort to indeterminate monism or indeterminate pluralism.

The epistemic theory of causality is an example of a general strategy for developing a determinate, monistic metaphysics that is true to the epistemology of a concept. Sometimes, attempts to explain a concept by positing a single mind-independent entity that corresponds to the concept meet fundamental difficulties, including counterexamples and epistemological problems. For example, probability faces an analogous epistemological problem: our probability judgements are based on knowledge of frequencies, knowledge of symmetries, indifference in the face of lack of knowledge and so on; if we try to analyse probability in terms of one of these indicators it is hard to explain the relevance of the others. More generally, mathematics faces an epistemological problem: our mathematical claims are based on a vast panoply of evidence, including proofs, patterns and pictures; monistic views such as platonism face well known difficulties in accounting for this epistemology (Benacerraf, 1973). In these cases, simple-minded realism offers a poor account of the concept in question and some other kind of account is needed.

An *epistemic theory* of concept  $X$  is a good strategy in such cases, one that can be used to provide an account of  $X$  that is truer to its epistemology and less prone to counterexamples. According to such a theory,  $X$  is to be interpreted in terms of a rational agent's epistemic state: rational  $X$ -beliefs are determined by an agent's epistemic background (and the uses to which  $X$ -beliefs are put);  $X$ -facts are characterised by those  $X$ -beliefs that an omniscient agent ought to adopt.

Objective Bayesianism provides an example of an epistemic theory of probability (see, e.g., Williamson, 2005a, Chapter 5). Here  $X$ -beliefs are an agent's degrees of belief, which, as Bayesian argue, satisfy the axioms of probability. These  $X$ -beliefs are determined by an agent's epistemic background  $\beta$  as follows. Degrees of belief are used for inference, e.g., prediction and decision. Given the predictive use, one's degrees of belief should be calibrated with one's evidence. Thus knowledge of frequencies directly constrains degrees of belief (if you know just that 80% of days like today are followed by rain you should believe today will be followed by rain to degree 0.8). So does knowledge of symmetries (if you know that accepted physical theory treats the different possible values of a particle's spin symmetrically, you should believe a particle has spin up to the same degree that you believe it has spin down). Given the decision-making use of degrees of belief, they should not be susceptible to a Dutch book and should not be bolder than is warranted by evidence (for otherwise one opens oneself

up to unnecessary risk—see Williamson (2007b)). Thus on a finite domain, an agent’s degrees of belief are represented by the probability function  $p_\beta$ , from all those that satisfy constraints imposed by  $\beta$ , that is most non-committal (i.e., has maximum entropy). Probability facts are then determined by these probability beliefs. The probability facts at time  $t$  are characterised by the probability function that an agent with knowledge of everything up to time  $t$  should adopt as her belief function. Note that there are differences between this epistemic theory of probability and the epistemic theory of causality. In particular, since probability has been axiomatised and shown to have several models, a certain pluralism is inevitable. Thus objective Bayesianism may be used to provide an account of the probability of a single case while the frequency theory of von Mises (1928) may be used to explicate the probability attaching to an indefinitely repeated sequence of outcomes (see also Russo and Williamson (2007b) on this point).<sup>6</sup>

An epistemic theory of mathematics proceeds similarly (Williamson, 2006b). Like causal beliefs, mathematical beliefs are used for explanation as well as inference. The explanatory use requires that mathematical beliefs be justifiable by means of proofs and interpretations, and that mathematical beliefs account for the evidence and be non-committal in other respects. The inferential use also imposes constraints: e.g., if a proof of a proposition is available as evidence then the proposition should be believed. Under the epistemic view, mathematical facts are characterised by the mathematical beliefs that an agent with full evidence should adopt. This gives a grounding to mathematics that is radically different to contemporary accounts such as platonism, neologicism, structuralism and nominalism.

The epistemic formula may also be applied to other problematic  $X$ s, e.g., logic, ethics. It is the failure of standard accounts of  $X$  which motivates the move to an epistemic account, not some global pragmatism or some modified criteria for adopting philosophical theories. In particular inference to the best explanation, a favoured mode of inference of the monistic realist, can be used to motivate an epistemic view of  $X$ . A realist conception of  $X$  may simply be untenable—prone to counterexamples or unable to account for the epistemology of  $X$ , for instance. In which case it does not offer the best explanation for our having the concept. There is thus room for an epistemic theory to provide the best explanation of our having  $X$ : we have concept  $X$  because of its utility (e.g., for inference and explanation), not because  $X$  corresponds to some single non-epistemic thing, just as we have hands because of their utility, not because of some kind of correspondence. Of course, the world must be such that  $X$  is a useful concept—just as the world must be such that hands are useful—so our having  $X$  says something about the world. It just does not say that there is something  $X$ -like in the world that is picked out by our concept. According to this stance, epistemic theories are to be judged by the same criteria as realist theories (see Williamson (2006a) for potential criteria). If a realist theory of  $X$  is viable, it may then be preferred over an epistemic theory of  $X$  on the grounds of simplicity. Thus we have the concept of table because there are tables that

---

<sup>6</sup>Under the epistemic view,  $X$ -beliefs are a *type* of belief (a directed relational belief in the causal case, a degree of belief in the probabilistic case) and are to be distinguished from beliefs *about*  $X$ , which are beliefs about which  $X$ -beliefs one should adopt if one had total evidence. An epistemic account understands  $X$  in terms of  $X$ -beliefs, not in terms of beliefs about  $X$  (the latter approach would lead to problems of circularity—see Williamson (2007a, §7)).

the concept picks out, but we have the concept of cause because of its inferential and explanatory utility.

## §A

### A FORMAL EPISTEMOLOGY

This appendix provides the details of a formal epistemology that can be integrated with the epistemic theory of causality. See Williamson (2005a, Chapter 9) for further discussion of the motivation behind this framework.

As in §3, let  $\beta$  be the agent's epistemic background. We shall suppose that  $\beta$  can be represented by two components,  $\pi$  which is a set of probabilistic evidence, and  $\kappa$  which is set of causal constraints, determined by the agent's other knowledge, including mechanistic and theoretical knowledge. For example,  $\kappa = \{A \longrightarrow B, A \not\rightarrow C, A \rightsquigarrow C, C \not\rightarrow D\}$ , where  $A \longrightarrow B$  signifies that  $A$  is a direct cause of  $B$ ,  $A \not\rightarrow C$  that  $A$  is not a direct cause of  $C$ ,  $A \rightsquigarrow C$  that  $A$  is a cause of  $C$ , and  $C \not\rightarrow D$  that  $C$  is not a cause of  $D$ . Let  $p_{\kappa, \pi}$  be the probability function, from all those that satisfy constraints imposed by  $\kappa, \pi$ , that has maximum entropy (see Williamson (2005a, §5.8) for an account of how causal knowledge constrains a probability function). All probability assertions will be made with respect to this probability function. For sets of variables  $X, Y, Z$ ,  $X \perp\!\!\!\perp Y|Z$  signifies that  $X$  and  $Y$  are probabilistically independent conditional on  $Z$ , while  $X \rightleftharpoons Y|Z$  signifies the opposite, that  $X$  and  $Y$  are probabilistically dependent conditional on  $Z$ . We are interested in determining  $\mathcal{C}_{\kappa, \pi}$ , a directed acyclic graph on the domain of  $p_{\kappa, \pi}$  that represents the causal beliefs that the agent should adopt on the basis of  $\kappa$  and  $\pi$  (arrows in the graph correspond to direct causal connections). Any causal graph  $\mathcal{C}$  will be assumed to be a directed acyclic graph (dag). With respect to such a graph,  $D_A$  is the set of direct causes of variable  $A$  and  $NE_A$  is the set of  $A$ 's non-effects. The question arises first as to how the agent's probabilistic knowledge  $\pi$  constrains choice of causal graph: what set  $\kappa'$  of causal constraints is imposed by probabilistic knowledge  $\pi$ ?

**DEFINITION A.1 (STRATEGIC DEPENDENCE)** *There is a strategic dependence from variable  $A$  to variable  $B$  with respect to probability function  $p$  and causal graph  $\mathcal{C}$ , written  $A \Rightarrow B$ , iff  $A$  and  $B$  are probabilistically dependent conditional on  $B$ 's other direct causes and  $A$ 's direct causes,  $A \Rightarrow B \mid D_B \setminus A, D_A$ .<sup>7</sup>*

The following definition and principle allow one to translate probabilistic constraints  $\pi$  into causal constraints  $\kappa'$ :

**DEFINITION A.2 (CAUSAL TRANSFER)** *Let  $\kappa^* \stackrel{\text{df}}{=} \{A \longrightarrow B : A \Rightarrow B\}$ . Given a causal graph  $\mathcal{C}$  that satisfies  $\kappa$ , a causal transfer of  $\pi$  with respect to  $\kappa$  and  $\mathcal{C}$  is a maximal subset  $\kappa'$  of  $\kappa^*$  such that  $\mathcal{C}$  satisfies  $\kappa$  and  $\kappa'$  (i.e.,  $\mathcal{C}$  satisfies  $\kappa$  and  $\kappa' \subseteq \kappa^*$ , and there is no  $\kappa''$  such that  $\kappa' \subset \kappa'' \subseteq \kappa^*$  and  $\mathcal{C}$  satisfies  $\kappa$  and  $\kappa''$ ).*

**PRINCIPLE A.3 (PROBABILISTIC TO CAUSAL TRANSFER)**  *$\mathcal{C}$  satisfies  $\kappa$  and  $\pi$  if and only if  $\mathcal{C}$  satisfies  $\kappa$  and  $\kappa'$ , where  $\kappa'$  is some causal transfer of  $\pi$  with respect to  $\kappa$  and  $\mathcal{C}$ .*

<sup>7</sup>This definition is a bit simpler than that given in Williamson (2005a, §9.5), but nothing very much hangs on the difference between the two definitions, other than the simplicity of some proofs.

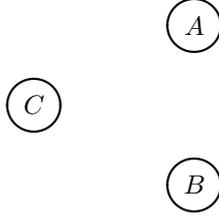


Figure 1: An empty graph.

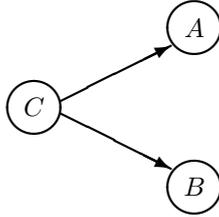


Figure 2: A common cause.

DEFINITION A.4 (EXPLANATORY RESIDUE)  $\bar{\kappa} \stackrel{df}{=} \kappa^* \setminus \kappa'$  is the (explanatory) residue of  $\mathcal{C}$ , with respect to  $\kappa$  and  $\pi$ .

The smaller the residue the fewer the strategic dependencies that have no causal explanation in  $\mathcal{C}$ . Let  $\mathbb{C}[\kappa, \pi]$  be the set of all causal graphs that satisfy constraints imposed by  $\kappa$  and  $\pi$  and that have smallest residue,  $\mathbb{C}[\kappa, \pi] = \{\mathcal{C} : \mathcal{C} \text{ satisfies } \kappa \text{ and } \pi, \mathcal{C} \text{ minimises } |\bar{\kappa}|\}$ .

PRINCIPLE A.5 (RATIONAL CAUSAL BELIEF) *An agent's rational causal belief graph  $\mathcal{C}_{\kappa, \pi}$  should be chosen from the set  $\mathbb{C}_{\kappa, \pi}$  of all minimal graphs in the set  $\mathbb{C}[\kappa, \pi]$  of all minimum-residue causal graphs that satisfy constraints imposed by  $\kappa$  and  $\pi$ ,  $\mathbb{C}_{\kappa, \pi} = \{\mathcal{C} \in \mathbb{C}[\kappa, \pi] : \mathcal{C} \text{ has fewest arrows}\}$ .*

Thus  $\mathbb{C}_{\kappa, \pi}$  is determined by first isolating the graphs that satisfy the constraints, then eliminating those that do not have minimum residue, then eliminating those that do not have the minimum number of arrows, then choosing one of the remaining graphs.

Here we have a minor point of departure from the approach of Williamson (2005a, §9.5). There it was suggested that  $\mathbb{C}_{\kappa, \pi}$  be determined simply by choosing a minimal graph from all those that satisfy the constraints. Here, we have an extra condition, namely that the residue be minimised. This condition is motivated by the following example.

EXAMPLE A.6 *Suppose the domain consists of three binary variables,  $V = \{A, B, C\}$ , with possible assignments  $a^0, a^1$  for  $A$ ,  $b^0, b^1$  for  $B$ , and  $c^0, c^1$  for  $C$ . Suppose that  $\pi = \{p(b^1|a^1) \geq p(b^1) + 0.3\}$ , and that  $\kappa = \{A \not\rightarrow B, B \not\rightarrow A\}$ . Then  $A \rightleftharpoons B$  is the only dependence in  $\mathbb{C}_{\kappa, \pi}$ , and the empty graph, Fig. 1, satisfies the constraints with residue  $\bar{\kappa} = \{A \rightarrow B, B \rightarrow A\}$ . On the other hand, the graph Fig. 2 also satisfies the constraints with no residue. Intuitively the latter graph is to be preferred, even though it has more arrows, because it*

includes an explanation of the dependence between  $A$  and  $B$ —it attributes the dependence to a common cause.<sup>8</sup> Thus residues should be taken into account.

Having isolated a rational causal belief graph, we turn to the question of how to find such a graph in practice. Clearly an exhaustive search through the space of all directed acyclic graphs will not be practical. Practical methods will make use of the following condition:

**DEFINITION A.7 (CAUSAL MARKOV CONDITION)** *The Causal Markov Condition (CMC) is said to hold if each variable  $A$  in the domain is probabilistically independent of its non-effects, conditional on its direct causes,  $A \perp\!\!\!\perp NE_A \mid D_A$ .*

**LEMMA A.8** *For any causal graph  $\mathcal{C}$ ,  $\mathcal{C}$  satisfies  $\kappa^* \Leftrightarrow \mathcal{C}$  satisfies CMC.*

**PROOF:** [ $\Rightarrow$ ] Suppose  $\mathcal{C}$  satisfies  $\kappa^*$ . Suppose for contradiction that  $\mathcal{C}$  does not satisfy CMC. Then there is some variable  $A$  and non-effect  $B$  such that  $A \rightleftharpoons B \mid D_A$ . This implies  $A \rightleftharpoons B, D_B \mid D_A$  by the contrapositive of the Decomposition property of conditional independence (see, e.g., Williamson, 2005a, §3.2), which in turn implies  $A \rightleftharpoons B \mid D_A, D_B$  by the contrapositive of the Contraction property.

Since  $D_B = D_B \setminus A, A \rightleftharpoons B$ . But this contradicts the assumption that  $\mathcal{C}$  satisfies  $\kappa^*$ , since  $A \not\rightarrow B$  in  $\mathcal{C}$ . Thus  $\mathcal{C}$  does satisfy CMC after all.

[ $\Leftarrow$ ] Suppose  $\mathcal{C}$  satisfies CMC. Suppose for contradiction that  $A \rightleftharpoons B$  but that  $A \not\rightarrow B$  in  $\mathcal{C}$ . There are four cases:

(i) If  $B$  is an (indirect) effect of  $A$  then  $\text{CMC} \Rightarrow B \perp\!\!\!\perp A, D_A \mid D_B \Rightarrow B \perp\!\!\!\perp A \mid D_A, D_B$  (by the Weak Union property) which contradicts  $A \rightleftharpoons B$ .

(ii) If  $A$  is an indirect effect of  $B$  then  $\text{CMC} \Rightarrow B \perp\!\!\!\perp A, D_B \mid D_A \Rightarrow B \perp\!\!\!\perp A \mid D_A, D_B$  contradicting  $A \rightleftharpoons B$ .

(iii) If  $A$  is a direct effect of  $B$  then  $A \rightleftharpoons B$  implies  $A \rightleftharpoons B \mid D_B, D_A$  which is impossible since  $B \in D_A$ .

(iv) If neither is a cause of the other then  $\text{CMC} \Rightarrow B \perp\!\!\!\perp A, D_B \mid D_A \Rightarrow B \perp\!\!\!\perp A \mid D_A, D_B$  contradicting  $A \rightleftharpoons B$ .

Thus in each case we have the required contradiction.  $\square$

We come now to a restatement of Theorem 3.1:

**THEOREM A.9** *Suppose there is some graph  $\mathcal{C}$  that satisfies  $\kappa$  and CMC. Then  $\mathcal{C}_{\kappa, \pi}$  is a minimal such graph.*

<sup>8</sup>Two remarks are in order. First, it might be objected that the epistemic background may include a non-causal explanation of the dependence, in which case one would not want any form of causal explanation of the dependence. But if that were the case then  $\kappa$  would rule out a common-causal explanation too, e.g.,  $\kappa = \{A \not\rightarrow B, B \not\rightarrow A, \forall X \neg(X \rightarrow A \wedge X \rightarrow B)\}$ , and Fig. 2 would then not be adopted. Second, one might object that the common causal explanation doesn't really account for the dependence between  $A$  and  $B$  because the cause  $C$  is independent of these variables. But it must be remembered that this independence is with respect to rational degree of belief, i.e., with respect to current evidence. This leaves open the question of whether there are dependencies between  $C$  and  $A$  and between  $C$  and  $B$  with respect to their frequencies. If so, as evidence improves, rational degree of belief may be expected to reflect those dependencies; the explanation of the dependence then improves. Thus a causal picture is only by itself part of an explanation of a dependence—a full explanation would need to appeal to probabilities based on good evidence.

PROOF: If  $\mathcal{C}$  satisfies  $\kappa$  and CMC then by Lemma A.8 it satisfies  $\kappa$  and  $\kappa^*$ . Thus  $\mathcal{C}$  has null residue. Hence  $\mathbb{C}[\kappa, \pi] = \{\mathcal{C} : \mathcal{C} \text{ satisfies } \kappa \text{ and CMC}\}$ . The result follows by Rational Causal Belief, Principle A.5.  $\square$

DEFINITION A.10 (STRATEGIC CONSISTENCY) *If there is a causal graph that has no residue with respect to  $\kappa$  and  $\pi$  (equivalently, if there is a graph that satisfies  $\kappa$  and CMC) then  $\kappa$  and  $\pi$  are said to be strategically consistent.*<sup>9</sup>

If  $\kappa$  is not strategically consistent with  $\pi$  then all is not lost. Suppose  $\kappa$  is consistent (i.e., there is some directed acyclic graph that satisfies  $\kappa$ ), contains only direct-causal constraints (i.e., constraints such as  $X \rightarrow Y$  or  $X \not\rightarrow Y$  that involve only direct causal connections), contains only atomic constraints (i.e., no logically complex constraints such as  $(X \rightarrow Y) \vee (Y \not\rightarrow Z)$ ), and contains no repetitions. Then we can write  $\kappa = \kappa^+ \cup \kappa^-$  where  $\kappa^+$  is the set of positive constraints in  $\kappa$  and  $\kappa^-$  is the set of negative constraints in  $\kappa$ . Consider the following algorithm:

ALGORITHM A.11

INPUT:  $\kappa$  (consistent; atomic direct-causal constraints; no repetitions),  $\pi, p_{\kappa, \pi}$

1. Choose a maximal set of constraints  $\lambda$  such that  $\kappa^+ \subseteq \lambda \subseteq \kappa$  and there is some  $\mathcal{C}$  that satisfies  $\lambda$  and CMC.
2. Take a minimal such  $\mathcal{C}$ .
3. Remove arrows from  $\mathcal{C}$  to satisfy the constraints in  $\kappa \setminus \lambda$  and yield a graph  $\mathcal{C}'$ .

OUTPUT:  $\mathcal{C}'$

THEOREM A.12 *Suppose  $\kappa$  is consistent and contains atomic direct-causal constraints with no repetitions. Then  $\mathcal{C}_{\kappa, \pi}$  can be taken to be  $\mathcal{C}'$  produced by the above algorithm.*

PROOF: First we need to show that such a  $\lambda$  exists. By Lemma A.8,  $\mathcal{C}$  satisfies  $\lambda$  and CMC if and only if  $\mathcal{C}$  satisfies  $\lambda$  and  $\kappa^*$ . There is such a  $\lambda$  because the complete directed acyclic graph is bound to satisfy  $\kappa^+$  and  $\kappa^*$ .

By consistency of  $\kappa$  and construction of  $\mathcal{C}'$ , the graph  $\mathcal{C}'$  satisfies the constraints in  $\kappa$  and  $\pi$ .

By maximality of  $\lambda$  and atomicity of  $\kappa$ , each constraint in  $\kappa \setminus \lambda$  violates a single constraint in  $\lambda \cup \kappa^*$ . Since  $\kappa$  contains no repetitions, each constraint in  $\kappa \setminus \lambda$  violates a *different* constraint in  $\lambda \cup \kappa^*$ . Since  $\kappa$  is consistent, each constraint in  $\kappa \setminus \lambda$  must violate a different constraint in  $\kappa^*$ . Hence the size of the residue of  $\mathcal{C}'$  (with respect to  $\kappa, \pi$ ) is  $|\kappa \setminus \lambda|$ .

$\mathcal{C}'$  must be a minimum-residue graph because  $\lambda$  is maximal.

Finally,  $\mathcal{C}'$  must be a minimal minimum-residue graph. This is because the choice of maximal  $\lambda$  makes no difference to the size of  $\mathcal{C}'$ .

Note that if there is some graph  $\mathcal{C}$  that satisfies  $\kappa$  and CMC, then the algorithm reduces to that of Theorem A.9.  $\square$

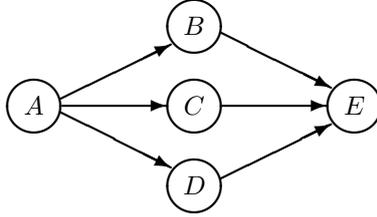


Figure 3: One graph satisfying CMC.

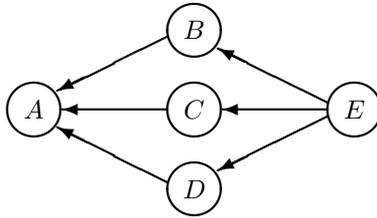


Figure 4: Another graph satisfying CMC.

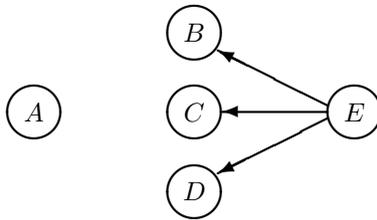


Figure 5: The result of the algorithm.

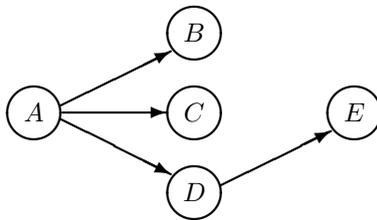


Figure 6: The rational causal graph.

The assumption that  $\kappa$  contain only atomic direct-causal statements is quite restrictive: while we often know claims of the form ‘ $X$  is a cause of  $Y$ ’ or ‘ $X$  is not a cause of  $Y$ ’, it is rarer that causal knowledge takes the form ‘ $X$  is a direct cause of  $Y$ ’ or ‘ $X$  is not a direct cause of  $Y$ ’. Thus it would be much more useful to be able to include atomic causal statements, so that  $\kappa$  contains atomic statements of the form  $X \rightsquigarrow Y, X \not\rightsquigarrow Y, X \longrightarrow Y, X \not\longrightarrow Y$ . Unfortunately the above algorithm is not guaranteed to succeed if we extend  $\kappa$  in this way. Suppose  $\kappa = \{E \not\rightsquigarrow A, B \not\rightsquigarrow E, C \not\rightsquigarrow E\}$  and the only minimal graphs that satisfy CMC are Fig. 3 and Fig. 4. Then  $\lambda = \{B \not\rightsquigarrow E, C \not\rightsquigarrow E\}$  and  $\mathcal{C}'$  is determined by removing arrows from Fig. 4 to give, e.g., Fig. 5, which has a residue of size 3. However, the rational causal graph is obtained by removing arrows from Fig. 3 to give Fig. 6 which has residue of size 2. The more general form of  $\kappa$  therefore requires a modified algorithm:

ALGORITHM A.13

INPUT:  $\kappa$  (atomic),  $\pi, p_{\kappa, \pi}$

1. Let  $\mathbb{C} = \{\mathcal{C} : \mathcal{C} \text{ satisfies CMC and } \kappa^+\}$ .
2. Set  $\mathbb{D} = \emptyset, n_{min} = \infty$ .
3. For each  $\mathcal{C} \in \mathbb{C}$ ,
  - (a) remove as few arrows as possible from  $\mathcal{C}$  to satisfy the constraints in  $\kappa$ , yielding graph  $\mathcal{D}$ ,
  - (b) let  $n =$  number of arrows removed,
  - (c) if  $n < n_{min}$ , set  $n_{min} = n, \mathbb{D} = \emptyset$ ,
  - (d) if  $n = n_{min}$  and  $\mathcal{D}$  satisfies  $\kappa$ , add  $\mathcal{D}$  to  $\mathbb{D}$ .
4. Let  $\mathbb{E} = \{\mathcal{D} \in \mathbb{D} : \mathcal{D} \text{ is minimal}\}$ .

OUTPUT:  $\mathbb{E}$

THEOREM A.14 *Suppose  $\kappa$  is atomic. Then  $\mathbb{E} \subseteq \mathbb{C}_{\kappa, \pi}$ , where  $\mathbb{E}$  is produced by the above algorithm. Moreover, if  $\kappa$  is consistent then  $\mathbb{E} \neq \emptyset$ .*

PROOF: Each  $\mathcal{E} \in \mathbb{E}$  satisfies  $\kappa$ . Further,  $\mathcal{E}$  is produced in step 3a by deleting as few arrows as possible from a graph that satisfies CMC, i.e., satisfies  $\kappa^*$ , so  $\mathcal{E}$  satisfies  $\pi$  as well as  $\kappa$ . Step 3c ensures that  $\mathcal{E}$  has minimum residue, while step 4 ensures that  $\mathcal{E}$  is minimal. Thus  $\mathcal{E} \in \mathbb{C}_{\kappa, \pi}$ .

If  $\kappa$  is consistent then there is some graph  $\mathcal{C}$  that satisfies  $\kappa$ . Any complete supergraph of  $\mathcal{C}$  satisfies  $\kappa^+$  and CMC, and hence is in  $\mathbb{C}$ . Thus  $\mathcal{C}$  will result from step 3a and is a candidate for admission in  $\mathbb{E}$ ; if  $\mathcal{C}$  is not in  $\mathbb{E}$  then that is because a graph that is smaller or has smaller residue is in  $\mathbb{E}$  instead. Hence  $\mathbb{E}$  is non-empty.  $\square$

Note that it may not be the case that  $\mathbb{E} = \mathbb{C}_{\kappa, \pi}$ , since different ways of carrying out step 3a may lead to different graphs in  $\mathbb{C}_{\kappa, \pi}$ , some of which are omitted from  $\mathbb{E}$ . The algorithm can be modified to output  $\mathbb{C}_{\kappa, \pi}$  by altering step 3a so that the algorithm runs through the *set* of maximal  $\mathcal{D} \subseteq \mathcal{C}$  that satisfy  $\kappa$ .

<sup>9</sup>This differs from the definition of strategic consistency given in Williamson (2005a, §9.6) but the sentiment is the same— $\kappa$  is strategically consistent with  $\pi$  if  $\kappa$  does not block the transfer of strategic dependencies to arrows in Principle A.3.

If  $\kappa$  is not atomic, then it may not be possible to divide  $\kappa$  into  $\kappa^+$  and  $\kappa^-$ . The above algorithm can be modified to cope with this more general type of causal constraint just by letting  $\mathbb{C} = \{\mathcal{C} : \mathcal{C} \text{ satisfies CMC}\}$  in step 1.

We see then that practical methods for finding a minimal graph satisfying CMC and a set of further constraints can be applied to the problem in hand, namely determining  $\mathcal{C}_{\kappa,\pi}$ . Note that such methods invariably involve querying the probability function  $p_{\kappa,\pi}$ . However, it may be possible to construct  $\mathcal{C}_{\kappa,\pi}$  directly from the epistemic background  $\kappa$  and  $\pi$  itself, without having to determine  $p_{\kappa,\pi}$  as an intermediary, via the following algorithm:

ALGORITHM A.15

INPUT:  $\kappa$  (atomic),  $\pi$  (strategically consistent)

1. Construct an undirected graph  $\mathcal{G}$  on the variables in  $\kappa$  and  $\pi$  by linking each pair of variables with an edge if they occur together in the same constraint in  $\pi$ .
2. Find a minimal (in terms of fewest edges) triangulation  $\mathcal{G}^T$  of  $\mathcal{G}$ .
3. Form the set  $\Omega = \{\omega : \omega \text{ is a maximum cardinality ordering of the variables, } \omega \text{ is a causal ordering consistent with } \kappa\}$ . (N.b.  $\omega$  is a maximum cardinality ordering if each variable  $A_i$  is a variable from  $\{A_j : j \geq i\}$  that is adjacent in  $\mathcal{G}^T$  to the largest number of variables in  $\{A_1, \dots, A_{i-1}\}$ .)
4. For each ordering  $\omega$  form a directed acyclic graph  $\mathcal{H}_\omega$  as follows:
  - (a) Let  $D_1, \dots, D_l$  be the cliques of  $\mathcal{G}^T$ , ordered according to highest labelled vertex.
  - (b) Let  $E_j = D_j \cap (\bigcup_{i=1}^{j-1} D_i)$  and  $F_j = D_j \setminus E_j$ , for  $j = 1, \dots, l$ .
  - (c) Take the variables as vertices of  $\mathcal{H}_\omega$ .
  - (d) Add an arrow from each vertex in  $E_j$  to each vertex in  $F_j$ , for  $j = 1, \dots, l$ .
  - (e) Add further arrows, from lower numbered variables to higher numbered variables, to ensure that there is an arrow between each pair of vertices in  $D_j$ ,  $j = 1, \dots, l$ .
  - (f) Add arrows corresponding to the positive constraints  $\kappa^+$  in  $\kappa$ .
5. Let  $\mathbb{H} = \{\mathcal{H}_\omega : \mathcal{H}_\omega \text{ satisfies } \kappa, \mathcal{H}_\omega \text{ is minimal}\}$ .

OUTPUT:  $\mathbb{H}$

THEOREM A.16 *Suppose  $\kappa$  is atomic and strategically consistent with  $\pi$ , that  $\mathcal{G} = \mathcal{G}^T$ , and that  $\pi$  does not on its own imply any probabilistic independencies. Then  $\mathbb{H} \subseteq \mathbb{C}_{\kappa,\pi}$ .*

PROOF: Each graph produced by the end of step 4e is acyclic and satisfies CMC with respect to  $p_{\kappa,\pi}$  (Williamson, 2005a, Theorem 5.6), hence so does each graph produced by the end of step 4f. Since  $\pi$  does not imply any independencies and  $\mathcal{G}$  is already triangulated, the graphs resulting from step 4e include

the minimal graphs satisfying CMC (Williamson, 2005a, Theorem 5.4).<sup>10</sup> Thus if  $\mathbb{H}$  contains a graph at all, it is a minimal graph satisfying  $\kappa$  and CMC, hence, by Lemma A.8, it is a minimal graph satisfying  $\kappa$  and  $\kappa^*$ . Since there is no residue (by assumption  $\kappa$  is strategically consistent with  $\pi$ ), it is a minimal graph satisfying  $\kappa$  and  $\pi$  and is in  $\mathbb{C}_{\kappa,\pi}$ .  $\square$

Note that  $\mathbb{H}$  may be empty: in Example A.6, the smallest graph that satisfies  $\kappa$  and CMC, Fig. 2, can not be obtained by removing arrows from the smallest graph that satisfies CMC (which has only one arrow, between  $A$  and  $B$ ).

Computational considerations may motivate simplifications of this algorithm. In particular, the second step, finding a optimal triangulation, is NP-hard (Yan-nakakis, 1981). Thus rather than demanding that there be no smaller triangulation in step 2, one might demand instead that the triangulation be minimal in the sense that no subgraph is a triangulation—as Berry et al. (2004) show, this is much more feasible (see also Neapolitan (1990, §3.2.3) for a fast triangulation algorithm). Similarly, one might want to stop step 4 when one directed acyclic graph has been found that satisfies  $\kappa$ . If such modifications are made, or if it is not known whether  $\pi$  implies any probabilistic independencies, then a resulting graph in  $\mathbb{H}$  can be viewed as an approximation to  $\mathbb{C}_{\kappa,\pi}$ .

If  $\kappa$  is not strategically consistent with  $\pi$ , this algorithm can be combined with Algorithm A.11 or Algorithm A.13 to try to identify a rational causal belief graph.

Finally, the set of rational causal belief graphs satisfies some interesting properties:

PROPOSITION A.17 *Suppose  $\kappa$  is strategically consistent with  $\pi$ .*

1. *If  $\kappa$  provides a causal ordering of the variables then  $\mathbb{C}_{\kappa,\pi}$  is uniquely determined if and only if  $p_{\kappa,\pi}$  is strictly positive.*
2.  *$[\mathbb{C}_{\kappa,\pi}] \subseteq \mathbb{C}_{\kappa,\pi}$ , where  $[\mathbb{C}_{\kappa,\pi}]$  is the Markov equivalence class of  $\mathbb{C}_{\kappa,\pi}$ .*
3. *If  $p_{\kappa,\pi}$  is faithful then  $\mathbb{C}_{\kappa,\pi} = [\mathbb{C}_{\kappa,\pi}]$ .*

See Williamson (2005a, §9.7) for the relevant definitions and proofs.

## ACKNOWLEDGEMENTS

This work was carried out as a part of the ‘Mechanisms and causality’ project supported by the Leverhulme Trust. I am very grateful to Federica Russo and an anonymous referee for helpful comments.

## REFERENCES

- Andersen, S., Olesen, K., Jensen, F., and Jensen, F. (1989). Hugin—a shell for building Bayesian belief universes for expert systems. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, pages 1080–1085.

<sup>10</sup>Such a graph can be used as the graph in a Bayesian net representation of  $p_{\kappa,\pi}$ . This is called an *objective Bayesian net*—see Williamson (2005b).

- Anscombe, G. (1971). *Causality and determinism*. Cambridge University Press.
- Benacerraf, P. (1973). Mathematical truth. In Benacerraf, P. and Putnam, H., editors, *Philosophy of mathematics: selected readings*, pages 403–420. Cambridge University Press, Cambridge, second (1983) edition.
- Berry, A., Blair, J., Heggernes, P., and Peyton, B. (2004). Maximum cardinality search for computing minimal triangulations of graphs. *Algorithmica*, 39(4):287–298.
- Carnap, R. (1950). *Logical foundations of probability*. Routledge and Kegan Paul, London.
- Cartwright, N. (2004). Causation: one word, many things. *Philosophy of Science*, 71:805–819.
- Dowe, P. (2000). *Physical causation*. Cambridge University Press, Cambridge.
- Gillies, D. (2000). *Philosophical theories of probability*. Routledge, London and New York.
- Gillispie, S. B. and Perlman, M. D. (2002). The size distribution for Markov equivalence classes of acyclic digraph models. *Artificial Intelligence*, 141:137–155.
- Godfrey-Smith, P. (2008). Causal pluralism. In Beebe, H., Hitchcock, C., and Menzies, P., editors, *The Oxford handbook of causation*. Oxford University Press, Oxford.
- 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.
- Hitchcock, C. (2001). A tale of two effects. *The Philosophical Review*, 110(3):361–396.
- Hitchcock, C. (2003). Of Humean bondage. *British Journal for the Philosophy of Science*, 54:1–25.
- Korb, K. B. and Nicholson, A. E. (2003). *Bayesian artificial intelligence*. Chapman and Hall / CRC Press, London.
- Lewis, D. K. (1973). Causation. In *Philosophical papers*, volume 2, pages 159–213. Oxford University Press (1986), Oxford.
- Mellor, D. H. (1995). *The facts of causation*. Routledge, London and New York.
- Neapolitan, R. E. (1990). *Probabilistic reasoning in expert systems: theory and algorithms*. Wiley, New York.
- Price, H. (1992). Agency and causal asymmetry. *Mind*, 101:501–520.
- Psillos, S. (2006). Causal pluralism. [www.phs.uoa.gr/~psillos/Publications\\_files/TheNaturalCausal\(Final\).doc](http://www.phs.uoa.gr/~psillos/Publications_files/TheNaturalCausal(Final).doc).
- Russo, F. and Williamson, J. (2007a). Interpreting causality in the health sciences. *International Studies in the Philosophy of Science*, 21(2):157–170.
- Russo, F. and Williamson, J. (2007b). Interpreting probability in causal models for cancer. In Russo, F. and Williamson, J., editors, *Causality and probability in the sciences*, Texts in Philosophy, pages 217–241. College Publications, London.
- Salmon, W. C. (1998). *Causality and explanation*. Oxford University Press, Oxford.
- Schaffer, J. (2005). Contrastive causation. *The Philosophical Review*, 114(3).
- Sober, E. (1985). Two concepts of cause. In Asquith, P. and Kitcher, P., editors, *PSA 1984*, volume 2, pages 405–424, East Lansing. Philosophy of Science Association.
- Spirtes, P. (1995). Directed cyclic graphical representation of feedback models.

- In Besnard, P. and Hanks, S., editors, *Proceedings of the 11th Annual Conference on Uncertainty in Artificial Intelligence (Montreal)*, pages 491–498. Morgan Kaufmann, San Mateo CA.
- Suppes, P. (1970). *A probabilistic theory of causality*. North-Holland, Amsterdam.
- von Mises, R. (1928). *Probability, statistics and truth*. Allen and Unwin, London, second (1957) edition.
- Weber, E., van Bouwel, J., and Vanderbeeken, R. (2005). Forms of causal explanation. *Foundations of Science*, 10:437–454.
- Williamson, J. (2005a). *Bayesian nets and causality: philosophical and computational foundations*. Oxford University Press, Oxford.
- Williamson, J. (2005b). Objective Bayesian nets. In Artemov, S., Barringer, H., d’Avila Garcez, A. S., Lamb, L. C., and Woods, J., editors, *We Will Show Them! Essays in Honour of Dov Gabbay*, volume 2, pages 713–730. College Publications, London.
- Williamson, J. (2006a). Dispositional versus epistemic causality. *Minds and Machines*, 16:259–276.
- Williamson, J. (2006b). From Bayesianism to the epistemic view of mathematics. *Philosophia Mathematica (III)*, 14(3):365–369.
- Williamson, J. (2007a). Causality. In Gabbay, D. and Guenther, F., editors, *Handbook of Philosophical Logic*, volume 14, pages 89–120. Springer.
- Williamson, J. (2007b). Motivating objective Bayesianism: from empirical constraints to objective probabilities. In Harper, W. L. and Wheeler, G. R., editors, *Probability and Inference: Essays in Honour of Henry E. Kyburg Jr.*, pages 151–179. College Publications, London.
- Yannakakis, M. (1981). Computing the minimum fill-in is NP-complete. *SIAM Journal of Algebraic and Discrete Mathematics*, 2:77–79.