

Causation and the Time-Asymmetry of Knowledge

Thomas Blanchard

University of Cologne

This paper argues that the knowledge asymmetry (the fact that we know more about the past than the future) can be explained as a consequence of the causal Markov condition. The causal Markov condition implies that causes of a common effect are generally statistically independent, whereas effects of a common cause are generally correlated. I show that together with certain facts about the physics of our world, the statistical independence of causes severely limits our ability to predict the future, whereas correlations between joint effects make it so that no such limitation holds in the reverse temporal direction. Insofar as the fact that our world conforms to the causal Markov condition can itself be explained in terms of the initial conditions of the universe, my view is compatible with Albert's well-known account of the origins of temporal asymmetries, but also provides a more illuminating way to derive the knowledge asymmetry from those initial conditions.

Keywords: Temporal Asymmetry, Causation, Knowledge, Causal Markov Condition

1. Introduction

We know more about the past than the future. I can tell a far more precise and reliable account of the last ten years of my life than of the next ten ones; figuring out who won the last literature Nobel Prize is much easier than predicting the next winner; we know when the Etna last erupted, but not when the following eruption will be. Of course, knowledge of future events is not impossible—we know quite a lot about the near

future, for example. Indeed, our survival depends on our ability to predict imminent happenings in our surroundings and adjust to them accordingly. But as we try to peak into the more distant future, prediction becomes harder and rapidly impossible, except when it comes to events pertaining to closed systems (for example, future eclipses). No such severe limitations hold toward the past: witness how archaeology, palaeontology and other historical sciences offer detailed and comprehensive reconstructions of the deep past, whereas no comparably successful sciences of the remote future exist.

What explains this epistemic time-asymmetry? An initially tempting answer is that the future is 'open' (hence unpredictable) while the past is 'closed' (therefore knowable). But the open/closed distinction is scientifically and metaphysically dubious, and at any rate this answer makes it mysterious how we can know anything about the future at all. A more plausible hypothesis is that the knowledge asymmetry emerges from another, more basic asymmetry in the physics of our world. In this paper I argue that this more basic asymmetry is the fact that effects of a common cause are correlated, but causes of a common effect are not. This well-known asymmetry, first identified by Reichenbach (1956), is connected to the causal Markov condition (Spirtes, Glymour, and Scheines 2000; Pearl 2009), arguably the most general principle connecting causation and probabilities. In a nutshell, I will argue that the statistical independence of causes embodied in the causal Markov condition severely limits our ability to predict the future consequences of present macroscopic events, whereas the statistical dependence of effects means that no such limitations hold in the reverse temporal direction.

The idea that the knowledge asymmetry stems from this statistical asymmetry between causes and effects is not new. It is already there in Reichenbach's (1956) seminal work, and has been defended more recently by Frisch (2014) and Stradis (2021).¹ (I discuss these accounts in section 8.) But it remains a minority view. Albert's (2000) account—by far the most popular theory of temporal asymmetries nowadays—makes no mention of a statistical asymmetry of causation in its derivation of the knowledge asymmetry. In

¹ Horwich (1987) also claims to derive the knowledge asymmetry from the statistical asymmetry embodied in the causal Markov condition. But the way his derivation works is very opaque (see Healey 1991).

claiming that a causal asymmetry is essential to understanding why we know more about the past than the future, my account diverges from Albert's. But as we will see, that causal asymmetry is itself derivable from certain facts about initial conditions of the universe which, on Albert's view, are the source of all temporal asymmetries. My view is therefore compatible with Albert's stance on the ultimate origins of the knowledge asymmetry, though I believe it offers a more illuminating derivation of that asymmetry from initial conditions of the universe.

2. **Albert's Account of the Knowledge Asymmetry**

As many authors (including Reichenbach (1956), Earman (1974) and Horwich (1987)) have noted, the knowledge asymmetry is closely linked to the fact that the world contains many more local macroscopic traces of the past than of the future. True, some events or entities do provide reliable information about future facts. (The smoke rising from my cooking pan tells me that the fire alarm will start beeping shortly; by checking a store's website I can learn that they close at 6pm tonight; etc.). But they are far less common than signatures of the past – memories, footprints, photographs, etc. That asymmetry explains why agents like us, who can only observe local macroscopic events, know more about the past. I know more about the last decade of my life than the next because I have many records (memories, documents, etc.) of the former but not of the latter; we know past but not future Nobel prizes because lists of past winners are easily accessible, whereas no lists of future winners currently exist; traces of past but not future eruptions abound; and the profusion of fossils, ruins and other leftovers of the remote past compared to the paucity of signs of the faraway future explains why there are successful sciences of the former and not of the latter (Cleland 2011). To understand why we know more about the past, we thus need to explain why the world contains many more local macroscopic signs of the past than of the future.

On Albert's view, this 'asymmetry of records', like other major temporal asymmetries, stems from two posits about the initial state of the universe: a 'Past Hypothesis' (PH) to the effect that the universe started in a low-entropy macrostate, and an initial randomness postulate that consists in a uniform probability distribution over possible

initial microstates of the universe compatible with PH (Albert 2000; see also Loewer 2012; Ismael 2016; Fernandes forthcoming). As I understand it, his explanation of the knowledge asymmetry proceeds in two steps. First, Albert argues that the reliability of records of the past depends crucially on PH. To see why, take the standard statistical-mechanical distribution over current microstates compatible with the present macrostate of the world. Strikingly, on that distribution all apparent records of the past contained in the present macrostate are likely spurious. For example, footprints in the sand are much more likely to have evolved spontaneously from an anti-thermodynamic fluctuation than from somebody walking on the beach earlier. (This is one upshot of Loschmidt's famous 'reversibility objection' against Boltzmann's account of the entropic arrow.) PH blocks this sceptical scenario by excluding trajectories in which the past has higher entropy than the present and thereby ensuring the accuracy of records of the past. In a second step, Albert argues that to be reliable, inferences from records must proceed on the assumption that the record was in a certain 'ready state' prior to the recorded interaction. For instance, inferences from footprints to earlier walking presuppose that the sand was smooth and damp before somebody walked on it. Albert contends that because PH underlies the reliability of all records of the past, it plays the role of ultimate ready state. Because no such ultimate ready state exists in the future, we can only have records of the past.

This account certainly has many virtues. In particular, the argument that PH plays a crucial role in making records of the past reliable is compelling. (Indeed my own account, while differing from Albert's in crucial ways, will take this point on board.) And because initial conditions of the universe also explain the thermodynamic asymmetry, Albert's story captures the widespread sense that the thermodynamic and knowledge asymmetries are linked, without making the connection as implausibly tight as earlier 'entropic accounts' of the time-asymmetry of knowledge.² But to my mind, his account

² The earliest such account was offered by Reichenbach (1956) in §19 of *The Direction of Time*. On that account, the key feature of records is their low entropy. Because the second law implies that low-entropy records cannot have evolved in isolation from a higher-entropy state, such records can provide information about interactions in the past but not in the future. Earman (1974) convincingly argues that this proposal draws much too tight a connection between

does not truly explain the key asymmetry described at the start of this section: the fact that the present contains many more local macroscopic signs of the past than the future.³ After all, showing that PH and initial randomness are needed to explain the reliability of the many present traces of the past is not the same thing as explaining why there are so many such traces in the first place. Nor does Albert's account of inferences from records shed light on that issue. What that account shows, if successful, is that indicators of the future are not *records* of the future, as inferences from them do not rely on any assumption of ready state. (To infer imminent alarm beeping from the smoke rising, we do not need to assume that the alarm or the smoke will be in a certain ready state *after* they interact.) But even so, this does not explain what we would like to know, namely why there are fewer such indicators of the future than of the past. To explain that fact, I will argue, we need to look at the causal structure of the world and the statistical asymmetries associated with it.

3. The Knowledge Asymmetry and the Causal Markov Condition

That the knowledge asymmetry may be linked to a causal asymmetry is suggested by the fact that our inferences from the present toward the past or the future are virtually always *causal* inferences. That is, inferences of past events generally involve deriving a past cause from one of its present effects (as when I infer earlier walking by noticing the footprint), whereas inferences of future events generally involve deriving an effect from a cause. (The inferred event may be an effect of the observed event, as when I predict imminent alarm beeping from the smoke in my kitchen. Or the inference may proceed by first deriving a past cause of the observed event and then deriving another future effect of that cause, as when I infer that the store closes at 6 from reading it on their

entropy and knowledge. His objections have been widely regarded as decisive, although see Fernandes 2022 for a recent reassessment.

³ Stradis (2021: 9525) makes the same diagnosis.

website, the common cause being the storeowner's hours policy.) That inductive inferences toward the past or the future are based on causation is unsurprising. An event a reliably indicates another event b only if they are correlated, and this requires a causal connection between them⁴: that is, a causes b , b causes a , or they have a common cause. The fact that we know little about the future compared to the past thus suggests as a hypothesis that inferring an effect from a cause is generally more difficult than inferring a cause from an effect.

Now, that hypothesis needs to be qualified. As it stands, it is in tension with the familiar idea that causes can be manipulated to control their effects (Woodward 2003), which could not be correct unless causes sometimes made their effects highly likely. But note that stock examples of causal relations that can be reliably exploited for control (scratching a match to light it up, throwing a rock to break a window, etc.) generally involve effects that lie in the short-term future of their causes. More generally, cases in which a single cause strongly predicts its effect tend to be limited to temporally nearby causal relata. When we consider more temporally distant causal relata, the idea that causes are easier to infer from their effects than vice versa becomes much more plausible, and the contrast becomes starker as we consider longer stretches of time. (This fits nicely with the fact noted above that while we do know quite a lot about the close future, predicting consequences of present events quickly becomes harder and often impossible altogether as we look further away into the future.) If we could identify an asymmetry between causes and effects that explains this contrast, this would shed light on the knowledge asymmetry. (You might worry that to do this we first need to settle what causation is. But the story I will present does not hinge on controversial issues of causal metaphysics, as will become clear later.)

To identify this asymmetry it is useful to consider the following example. The left panel of the Pilgrim Monument in Plymouth (Massachusetts) states that the *Mayflower* was chartered in 1620 by Robert Cushman, a member of the Leiden congregation of

⁴ With perhaps some rare exceptions that we can ignore here, such as entangled components of quantum systems and causally independent quantities obeying similar laws of evolution (see Arntzenius 2010).

separatists and also likely one of the ship's passengers. Cushman's chartering of the ship (event s) was a cause of the present existence of the inscription on the panel (event p). This is a clear case where no one could have predicted the effect upon observing the cause, whereas the cause is easily inferable from observing the effect. Why is that so?

It is easy to see why no witness of s in 1620 could have predicted the occurrence of p 403 years later. Obviously, the chartering of a ship at some time t is causally insufficient to ensure the existence of a memorial of that event centuries later: for the latter to occur, an enormous number of additional causal factors besides the chartering also have to obtain at t . However, the probability that all of the relevant factors line up in the right way is minuscule, meaning that ship charterings are statistically very unlikely to be publicly commemorated centuries later. In the *Mayflower* case, the relevant additional factors that led from s to p included all of the weather conditions that would later ensure safe travel of the ship and its landing in Plymouth, various facts about Native American people who would later provide vital assistance to the Pilgrims, the existence in 1620 of the ancestors of the people who commissioned and sculpted the monument at the end of the 19th century, etc. That particular line-up of favourable circumstances was highly unlikely, and so was any other possible scenario by which s might have led to p . (For example, a scenario in which the *Mayflower* sunk at sea and an eccentric North American descendant of Cushman later decided to build a memorial for their ancestor in Plymouth would have required a similarly unlikely arrangement of circumstances back in 1620.)

These remarks highlight some key reasons why consequences of present events in the faraway future are often unpredictable. Suppose one observes an event c that will have an effect e in the distant future. Generally, there is no true claim of the form ' C s are regularly followed by E s' that one might use to predict e upon observing c . For consider any kinds of events C and E through which a human observer might conceivably pick out c and e (that is, kinds that can be described in everyday macroscopic terminology).⁵

⁵ If the laws are deterministic, c will fall under a description from which the later occurrence of e follows, namely a description of the form 'event such that the complete microstate of the world

As a rule, *Cs* alone are only partial causes of *Es*, and whether an *E* occurs always depends on further causal factors besides the occurrence of a *C*. This pervasive underdetermination of effects by individual causes is built into the physics of our world. The laws of physics are global: they entail that a local event's occurrence always depends on a multitude of events at any prior point in time—indeed, on what happens at every point in the relevant cross-section of the backward light-cone (Field 2003). In particular, the laws always make it possible for a local causal chain to be disturbed by external circumstances, and thus entail that a local event can produce another only if no such disruption occurs. This means that whether a local event will be followed by another always depends on further causal factors – and the more temporally distant the two events, the more additional factors are involved.⁶ Thus an instance of *C* can cause a later instance of *E* only if one of several possible *enabling conditions* $X_1, X_2\dots$ is also instantiated at t , each one consisting of a conjunction of other factors causally relevant to *E*, such that $C \& X_i$ guarantees *E* or at least makes it highly likely.⁷ Consequently, *Cs* can be frequently followed by *Es* only if it is highly probable that one of those enabling conditions obtains when an event of type *C* occurs. Now, if the effect lies in the immediate future of the cause this may easily be so, for in that case the effect depends

at that time is such and such'. But this description is not one that agents like us could ever use to pick out *c*.

⁶ Perhaps not *every* factor in the backward light-cone of an event counts as a cause of it. The exact microscopic state of Mars lies in the present current cross-section of the backward light-cone of the dinner party I hold tonight, but one might deny that the former is a cause of the latter, especially if one takes causation to be an essentially macroscopic phenomenon (see e.g. Papineau 2022). But note that virtually every account of causation would count a more coarse-grained state such as 'no extremely powerful pulse of energy traveling from Mars to my dining room' as a cause of my dinner party (although certainly not a salient cause).

⁷ Note that even if the laws are deterministic, $C \& X_i$ may fail to fully guarantee *E*. By assumption *C* is a macroscopic, and if causation is essentially macroscopic the causal factors that figure in X_i must be as well. But there are good reasons to think no macroscopic description of the state of the world can nomically guarantee a future event – there are always at least some 'maverick' microstates compatible with that description which do not nomically entail *E* (see Albert 2000).

only on a few other factors whose joint occurrence may be highly probable. For example, smoke rising from my pan is regularly followed by the alarm going off because the relevant additional required circumstances (alarm battery is charged, no incoming nuclear explosion, etc.) are likely to obtain. But for effects in the faraway future, each enabling condition typically contains a very large number of different factors spread out over an enormous spatial area. And those factors are highly unlikely to all occur simultaneously, so that the occurrence of an enabling condition is highly improbable. *Cs* can then only be infrequently (if at all) followed by *Es*. Consequently, the mere observation of a *C* does not suffice to predict the later occurrence of an *E*. To predict *e*, an observer would thus have to not only witness *c*, but also somehow ascertain that an enabling condition obtains, and figure out the future consequences of *c* and that enabling condition. But typically this is far beyond our observational and computational powers. For example, no witness of *s* in 1620 could have observed all of the additional causes of *p* in place at that time, which were distributed over a gigantic spatial region. Neither could they have had enough background knowledge of the world to figure out the existence of all of these causes. And even if they somehow were aware of all those factors, they would still have been utterly unable to trace out the dizzyingly complex chain of interactions by which those factors together with *s* would lead to *p* four centuries later.⁸

Why, by contrast, is it so easy for anybody visiting the Pilgrim monument today to infer *s* by observing *p* alone? The answer, I take it, is that in this case there *is* a true generalization that one can use to infer *p* from *s*—a generalization stating that inscriptions on public monuments are reliable. This, note, is a causal claim stating that a certain type of effect (an inscription on a monument) is regularly preceded by a certain

⁸ Of course, not all effects in the *close* future can be predicted either. Riding a bike is rarely followed by getting into an accident 10 seconds later only very infrequently, because the additional factors needed for an accident are fortunately rare. Predicting an imminent accident is therefore impossible unless one can observe the relevant factors (e.g. one sees the incoming car in the cross street).

type of cause (viz. an event of the sort represented by the inscription). What makes that generalization true?

Albert's Past Hypothesis provides part of the answer: without PH, inscriptions on public monuments would more likely be the results of anti-thermodynamic fluctuations than of the events they claim to commemorate. But that is not all. For even holding PH fixed, the laws of physics allow for many mundane scenarios in which a public monument records an event that never occurred. (The engraver might have gotten their facts wrong, or a prankster may have replaced the original panel with a fake one.) These scenarios, note, are ruled out or at least made highly unlikely if we consider other present *effects* of the recorded event. Consider all present effects of s contained in the current macrostate of the world: not only p , but also other monuments to the pilgrims, statements about Cushman in history books and online sources, memories of s in the brains of history scholars, etc. Conditional on all those effects taken together, the laws make s highly likely. The flipside, however, is that according to these laws, whether s happened depends not on p alone, but also on the existence of those other effects. If none of those other records s existed, the laws would entail that p is most likely the result of a mistake or a prank.

More generally, just like local causes underdetermine their effects, local effects underdetermine their causes. This is because the physical laws are global in both temporal directions. Just like the only prior states that nomically determine an event (or fix its probability) are global states of the world, the only posterior nomic determinants of an event are global states at later times, for example complete cross-sections of future light-cones in relativity (Field 2003: 436–7).⁹ These laws thus entail not only that a cause can be followed by an effect only if many other factors obtain simultaneously with the cause, but also that an effect can have been preceded by a cause only if some enabling condition obtains simultaneously with the effect, each such condition

⁹ As shown by Elga (2001) and Frisch (2005), this fact implies that Lewis's 'asymmetry of overdetermination'—the alleged fact that an event rarely has more than one nomically sufficient cause, but typically has many localized nomically sufficient effects—does not exist.

consisting of a conjunction of possible other effects of the cause.¹⁰ (Typically there are many such possible enabling conditions. In our example, some conditions include Wikipedia stating that Cushman chartered the *Mayflower*, while in others Wikipedia currently says nothing about Cushman but somebody also remembers having deleted the *Mayflower* article a few minutes ago.) Moreover, because each such enabling condition is a very long conjunction of highly specific events, it is unconditionally highly unlikely to obtain. (Among all physically possible ways the current macrostate of the world might be, very few contain records of Cushman chartering the *Mayflower*: in most of those worlds, Cushman never existed, or never chartered the ship, or the ship sank, etc.). There is thus a deep symmetry between effect-to-cause and cause-to-effect inference: in both cases, the observed event severely underdetermines the inferred event, and the additional circumstances on which the inferred event depends are unconditionally very improbable. However, while in the cause-to-effect direction this underdetermination entails that as a rule causes do not raise the probability of their temporally distant effects by very much, no such consequence follows in the other direction: a single effect can perfectly well make one of its remote causes highly likely, as the *Mayflower* example illustrates. What explains the difference?

The answer is that causes of a common effect are generally statistically independent of one another, whereas effects of a common cause are generally correlated. When explaining why predicting the remote future is hard, I claimed that when the enabling conditions needed for *C* to cause *E* are statistically rare, *C*s can be only rarely followed by *E*s. But that claim presupposes that those enabling conditions are statistically independent of *C*, and hence remain as improbable when a *C* happens as when it doesn't. And that presupposition is correct because enabling conditions are other *causes*

¹⁰ For reasons similar to those outlined in fn. 7, the conjunction of *E* and other effects of *C* likely fails to determine *C* even under determinism, as *E* is macroscopic and other effects of *C* may have to be as well if causation is an essentially macroscopic phenomenon. (Certainly the other effects of *p* mentioned in the main text are all macroscopic states.) Note that this is so even if one holds PH fixed: even then, every macroscopic description of the current state of the world remains compatible with certain maverick microstates that make the cause absent (Elga 2001; Loewer 2007: 315).

of E and thereby uncorrelated with C . Consider for example the myriad of events happening in North America in 1620 that would later lead to the meeting between Native Americans and the Pilgrims, or the myriad of decisions taken by certain people in 1620 that later led to the birth of the people who commissioned the Pilgrim monument. Cushman's chartering of the ship did not in any way raise the probability of those events. (Their joint co-occurrence was, as we say, a *coincidence*.) More generally, whether a ship is chartered at a certain time t does not in any way raise the probability that background circumstances at t are of the right sort to ensure that the chartering will be memorable four centuries later. Because the occurrence of such favourable circumstances is highly unlikely to begin with, it remains so on occasions when a ship is chartered, and this is why ship charterings are only rarely followed by commemorative monuments. Suppose for a minute that this wasn't the case, and that on most occasions when a ship is chartered circumstances happen to line up in the right way to make that event memorable four centuries later. Ship charterings would then be regularly associated with memorials of them later on, and one could confidently predict the latter just by observing the former. Thus, a crucial reason why predicting the future is so hard is that causes of a joint effect are uncorrelated. (Or at least, they are uncorrelated if they are causally unconnected—that is, they neither cause each other nor have a common cause. That condition seems satisfied in my *Mayflower* example, but may not be in other cases. I will table that issue for now and return to it in section 5.)

In the reverse direction, however, no similar independence condition holds: effects of a common cause can be and generally are correlated. This is why an effect E can make its cause C very likely despite the fact that C also nomically depends on many other present effects. This only requires that E be very strongly correlated with those other effects. For example, the fact that inscriptions on public monuments make the events they report highly probable implies the existence of very strong correlations between these inscriptions and the contents of encyclopedias, the brain states of historians, and so on. But those correlations are not surprising or improbable: the causal structure of our world forbids the existence of such correlations between causes, but allows them between effects, including effects of causes in the remote past. It is that fact that explains why observing a single event is sometimes enough to infer a long-gone cause of it. (Suppose for a minute that effects of a common cause had to be uncorrelated. It

would then take a giant coincidence for monuments, encyclopedias and other records of human history to have the same content: most likely their respective contents would be incompatible with one another, and almost all of them would be false. It would then be impossible to infer anything about the past from any single record by itself. Instead, to figure out whether a past event occurred, one would have to somehow inspect all potential traces of that event and figure out what kind of causal history might have led to that distribution of traces. For events far away in the past, this task would be far beyond our observational and computational abilities. Thus, if effects of common causes were independent, we would be facing the exact same predicament toward the past that we face when we try to predict the future.)

It may be helpful to recast the reasoning more formally. Consider two kinds of events A and B such that A is either a cause or an effect of B . (Assume both are local macroscopic kinds describable in terminology that agents like us can understand.) And consider what conditions must obtain for A s to make B s highly likely, so that

$$(1) P(B/A) \approx 1$$

(The relevant probabilities are meant to be objective, physical probabilities attached to event-types.) As we saw, the physics of our world implies that B always depends on other factors concomitant with A , so that B is likely to follow or precede A only if some enabling condition is co-instantiated with A . If A causes B , each enabling condition is a conjunction of other factors causally relevant to B , and if B causes A each one is a conjunction of other effects of B . Letting $\{X_1, \dots, X_n\}$ be the set of such enabling conditions that together with A make B very likely (that is, close to 1), and \mathbf{X} their disjunction, we have:

$$(2) P(B/A\&\mathbf{X}) \approx 1$$

(Note that (2) presupposes that P is conditioned on the Past Hypothesis: otherwise, when B is the cause and A the effect, B would still remain unlikely given A and other putative effects of B – the most likely hypothesis would be that all apparent records of B

are products of anti-thermodynamic fluctuations. (In that respect my account takes on board one of Albert's key points.) (1) can then hold only if¹¹

$$(3) P(\mathbf{X}/A) \approx 1$$

But when A is the cause of B , it is statistically independent of B 's other causes figuring in those enabling conditions, so that

$$(4) P(\mathbf{X}/A) = P(\mathbf{X})$$

If so, (1) can hold only if

$$(5) P(\mathbf{X}) \approx 1$$

that is, if the presence of an enabling condition is unconditionally highly likely. But unless B is in the near future of A the probability of an enabling condition occurring is typically low, so the probability of B given A has to be low as well. On the other hand, if A is the effect of B , (4) does not hold: instead, B makes it likely that other effects of A are also present and are thus with \mathbf{X} . So even if each enabling condition is unconditionally very unlikely, (1) can still hold. In those cases, we will find that the A and B 's other effects are very strongly correlated, so that $P(\mathbf{X}/A)$ is close to 1. But the causal structure of our world allows for such correlations.

The fact that effects of a common cause are correlated but causes of a given effect are not was first identified by Reichenbach (1956), who called it the 'fork asymmetry'. It is connected to a number of other conditions like the common cause principle (also discovered by Reichenbach) and the physical principle that 'incoming' but not 'outgoing' influences are uncorrelated (Frisch 2014). That asymmetry also follows from the causal Markov condition, the basis of the strikingly successful causal modelling framework (Spirtes, Glymour, and Scheines 2000; Pearl 2009) and arguably the most general principle linking causation to probabilities. The causal Markov condition says that

¹¹ To see this note that the law of total probability entails that

$$P(B/A) = P(B/A \& \mathbf{X}) P(\mathbf{X}/A) + P(B/A \& \neg \mathbf{X}) P(\neg \mathbf{X}/A)$$

By assumption $P(B/A \& \neg \mathbf{X}) \ll 1$, so (1) can be true only if (3) is.

conditional on its direct causes, a variable is statistically independent of all other variables except its effects. This directly entails that causally unconnected causes are statistically independent of one another. Together with other causal modelling axioms, the causal Markov condition also entails that effects of common causes are correlated.¹² On the view I propose, then, the causal Markov condition is the key causal principle that underlies the knowledge asymmetry.

To summarize, I have argued that together with certain facts about the physics of our world, the statistical independence of causes implied by the causal Markov condition places severe limits on predicting the future. Whether an effect E follows a cause C always depends on whether further causes of E also occur. Because these further causes are statistically independent of C , they are not made more likely to occur just because C is instantiated. So unless those other causes are highly probable to begin with, C s can only be rarely followed by E s, and one cannot predict a future instance of E just by observing a present instance of C . The only way to make a prediction would be to find out the values of other factors causally relevant to E , and work out the joint consequences of C and those factors in the future. But limited agents like us are typically not in a position to do this, unless perhaps E lies in C 's near future (though even then prediction might be impossible).¹³ But because the causal Markov condition allows for

¹² Together with another modelling axiom, the 'minimality condition', the causal Markov condition entails that if two effects have a single common cause and they do not cause one another they must be correlated. Those two conditions do allow for more complex causal structures in which effects of a common cause are uncorrelated, but this requires highly unlikely correlation cancelations ruled out by the 'faithfulness condition'.

¹³ As noted above, causal inferences to the future can proceed either by deriving an effect from a cause, or by inferring a joint effect (i.e. one first infers a past cause c of a present event e_1 and then infer another effect of c e_2 that lies in e_1 's future). Provided that e_2 lies in the distant future of e_1 , whether e_2 will occur typically depends not only on c but also on a myriad of other causal factors at e_1 's time of occurrence whose joint occurrence is very improbable, and independent of c 's (and e_1 's occurrence). Thus it will be hard or impossible to predict e_2 from e_1 alone. As a reviewer noted, there might seem to be a tension here with my claim that in the *Mayflower* example the other present effects of s are very strongly correlated with p . After all, each such

correlations between effects, these limits do not hold in the reverse direction. An effect E can be tightly correlated with a cause C in the remote past, even if the laws of physics entail that C depends on vastly more events in the present macrostate than E alone. As long as those other parts of the macrostate are tightly correlated with E itself, E by itself can make C highly probable. A human observer can then infer a past instance of C by observing a present instance of E alone, provided they know that E s are regularly preceded by C s.

4. Dispelling some Worries

That last remark shows that my account can address a challenge that every account of the knowledge asymmetry faces. Any such account must explain why agents like us can more easily infer the past than the future from the present. But what inferences we can draw depends on the background knowledge we have. And because of the knowledge asymmetry, agents like us have much more background knowledge about the past than the future. In explaining why we can more easily read the past than the future from the present, we must take care not to implicitly rely on this aspect of background knowledge, on pain of presupposing the very asymmetry that we seek to explain. The explanation I proposed satisfies this constraint. In the *Mayflower* case, to learn s from p one only needs to know a certain regularity – that inscriptions on public monuments are

effect (e.g. Wikipedia’s statement that Cushman chartered the *Mayflower*) is also the product of many factors besides s that are independent of p (e.g. the existence of the Internet, Jimmy Page’s decision to create Wikipedia, etc.) So we should expect p to be only very weakly correlated with these other effects. My reply is that this is correct, and that p is indeed only very weakly correlated with each of the possible enabling conditions that together with it make s highly likely. But it is still very strongly correlated with the *disjunction* of those conditions (the X in equation (2) above). That is, p by itself does not predict any of the other particular effects of s in the current macrostate of the world, but strongly predicts that this macrostate must be arranged in some way or other that is also indicative of s . This is all that is needed for p by itself to make s highly likely.

reliable. Of course, a typical visitor to the Pilgrim monument presumably knows a lot about the past, including perhaps a number of facts about the *Mayflower*, North American history, etc. But none of that knowledge is needed for the inference to go through. A creature who knew nothing about the past could still easily make the inference, as long as they knew that public monuments are reliable. (That piece of knowledge, note, is not specifically about the past, for it applies equally well to future monuments.) This is not to deny that an agent who already knows a lot about the past will also be able to extract more information about the past from the present, because their prior knowledge already puts constraints on what the past must have been like.¹⁴ But my account does not rely on that fact.

Now, one may worry that my account still presupposes an unexplained knowledge asymmetry. To see why, consider how agents like us manage to learn the generalizations we use to predict the past (for example, inscriptions on monuments are reliable) or the future (for instance, kitchen fires cause smoke alarms to go off). One method is induction over past observed instances of those generalizations.¹⁵ This, however, requires an ability to form *memories* of those past instances. Or we may come to know those generalizations via testimony, but again this arguably requires a capacity to remember those testimonies. Or one might learn them by deriving them from broader regularities (for example, inferring that monuments are reliable from more basic facts about human psychology and sociology). But presumably those broader regularities would themselves have to be learned through induction or testimony, which once again requires memories. And memories are of course time-directed: we can remember the past, but not the future. But because the account I proposed only concerns inferences, it seems unable to explain this asymmetry, as memory-based knowledge is not inferential. When I remember an event, I do not first observe the

¹⁴ For example, knowing what I know about the *Mayflower* and its history, I may infer from *p* not only that Cushman chartered the *Mayflower*, but that he likely was a Puritan, may have emigrated to Holland, etc.

¹⁵ Induction is actually not a plausible method to learn that monuments are reliable since we rarely witness both a monument and the event it commemorates. But it is certainly one way in which one might learn that kitchen fires make smoke alarms beep.

memory and infer the event from it. Rather, I simply come to believe that the event happened by having the memory.

A sensible response to this worry is that while our knowledge of causal generalizations rests on memories, nothing depends on the fact that those memories are about the *past*. A creature who could somehow remember the future but not the past could still learn that public monuments are reliable or that kitchen fires make smoke alarms go off based on induction over *future* instances of those generalizations – or, perhaps, memories of future testimonies. (So even such a time-reversed creature would end up learning more regularities relating events to their causes than to their effects.) But this response does not put the worry fully to rest. For there is another way in which my account presupposes a time-asymmetry of memory. If we could form memories of the future, those memories would put further constraints on possible future histories, and we could infer much more about the future from the present than we actually can. (If I ‘remembered’ going to the hospital in 20 minutes, I could also infer that going on a bike ride right now will lead to an accident.) So the inferential asymmetry that my account identifies and seeks to explain would disappear, or at least be diminished.

Fortunately, the line of reasoning presented in the previous section can be extended to explain the time-asymmetry of memory itself. For note that memories can provide knowledge only if they are reliable indicators of remembered events. That is, conditional on the memory the occurrence of the remembered episode must be highly likely. Having memories of the future would thus require that certain present brain states be strongly correlated with localized events in their future. And the upshot of my argument is that the independence of causes makes such correlations highly unlikely – whereas no such limits exist on correlations between present brain states and past events. Like the time-asymmetry of inferences, the time-asymmetry of memory finds its roots in the causal Markov condition. (One loophole here is that my account does allow for strong correlations with effects in the *immediate* future. But that loophole can be closed once we take into account another feature of memories, namely their richness: we remember past events, especially those in the recent past, in very specific and fine-grained detail. And a consequence of my account is that the more precise the description of a future effect, the harder it is to predict from any of the present events correlated with it: for a more specific effect is also requires more specific and hence also

more unlikely enabling conditions. For example, during a conversation I may be able to anticipate what my interlocutor is about to say, but finer details – their exact words, tone of voice, exact facial expressions as they speak, etc. – are outside of my reach as they depend on a myriad of present unknowable factors. (Contrast how vividly I may recall my interlocutor’s last spoken words in all of their fine-grained texture.) This suggests that the very rich and fine-grained correlations between brain states and past recollected events characteristic of memories simply could not hold with future events, even imminent ones.)

The view that emerges from this discussion can be helpfully recast in the form of the following ‘genealogy’ of the knowledge asymmetry. Imagine an agent who can only observe small portions of the macrostate of the world at a time, and who has little or no background knowledge about the world. The causal Markov condition and the statistical asymmetry to which it gives rise will ensure that this agent’s brain states can become reliably correlated with events in her past, but not in her future, enabling the agent to remember one but not the other. Through those memories, our agent will over time come to learn a number of regular associations between causes and effects. But the causal Markov condition will also ensure that those regularities will more often relate present events to past causes than to future effects, unless these effects are very close in time. Our agent will then be able to use those generalizations to draw many inferences toward the past, but comparatively fewer about the future. Moreover, the knowledge of past events our agent will gain through those inferences will itself facilitate further inferences toward the past, by putting strong constraints on epistemically possible past histories. Over time, then, that agent’s state of knowledge will come to display the sharp temporal asymmetry characteristic of agents like us.

5. The Possibility and Limits of Prediction: A Closer Look

The line of reasoning presented in the last two sections calls for several remarks and clarifications. First, as noted above, the causal Markov condition says that causes are

statistically independent when they are causally unconnected. Although my lead example seems to satisfy that condition, it is easy to imagine cases that don't, and in which a cause strongly predicts an effect because it is correlated with that effect's other causes. For example, suppose that today you hear a member of the majority group in Parliament declare their intention to vote for a measure presented tomorrow. This is evidence that other group members will vote likewise. (Votes of parliamentary group members have common causes – shared beliefs, party leadership instructions, etc. – and are therefore correlated.) You may therefore predict that the measure will pass upon hearing the declaration alone. Such cases, however, are (I believe) relatively rare. More precisely: it is not uncommon for a cause c of an effect e to be correlated with some of its other causes. But typically there are also further causes that are causally unconnected and hence uncorrelated with c : and the further away in the future e is, the greater the number of such causes. So unless e lies in the close future, c cannot raise the probability that some enabling condition obtains by very much, because each one still includes many factors independent of c . In my example, the measure passing depends not only on other members' votes, but also on many further facts causally independent of the group member's voting intentions, for example the absence of a political scandal or natural disaster before the vote. Prediction can go through because the probability of such a disruption between now and tomorrow is low, but if the vote were scheduled for (say) next year, that probability would sharply go up, and the inference would become unreasonable. Successful prediction involving causally connected causes thus appears mostly limited to the short-term future.

Second, although I have emphasized how difficult it is to know the faraway future, my account also explains why and how certain kinds of long-term prediction are possible. For one thing, my view implies that there is a trade-off between confidence and precision in prediction—a more specific effect requires more specific and hence more unlikely causes, so that more detailed predictions are also necessarily less reliable. Accordingly it is sometimes possible to confidently predict long-term consequences of present events as long as the prediction remains highly unspecific. (We may not know how the Russian invasion of Ukraine will affect European geopolitics in the long term, but we can tell that it will, in one way or another.) My account also entails that precise prediction of the distant future is possible in special circumstances, for example if we

are dealing with isolated or quasi-isolated systems that are effectively shielded from the influence of outside factors. A paradigmatic example is the solar system, whose long-term future can be reliably and precisely predicted by observing current positions and velocities of a small number of celestial bodies. (Possible outside factors such as strong incoming gravitational disturbances might disrupt the system, but are sufficiently unlikely that they can be safely ignored.) Note, however, that isolated systems are far and few between, at least in the 'sublunar' portion of the world where we live our lives. A pervasive feature of our condition is that the events that have the most significance for us and that we would therefore dearly like to know about—election results, wars, life-changing encounters, and so on—always lie at the intersection of a myriad of independent causal lines, and are therefore bound to remain unpredictable, except perhaps in barest outline. This underscores how impoverished our ability to read the future from the present is, especially in comparison to the wealth of knowledge we can amass about the past.

6. Further Conditions on Knowledge of the Past

Of course, although the past is more easily accessible than the future, we do not know everything about the past. One reason is that many past events have left no macroscopic traces in the present. In particular, as Turner (2005) has argued, many facts about the deep past are lost to us because information-destroying processes have erased all visible effects of those facts. But that is not all. On the view I have proposed, past-directed inferences rely on generalizations of the form 'Es are regularly preceded by Cs'. Many of the events involved in the causal history of a present fact cannot be inferred from it because no such regular association exists. Imagine that in 1620 a certain storm sunk a pirate ship that would have otherwise pillaged the *Mayflower*. Although the storm played a role in the Pilgrim monument coming to existence, it cannot be inferred from the latter as there is no relating causal regularity relating storms to monuments. The causal Markov condition allows for the existence of regularities relating effects to past causes, but does not explain why some regularities hold but not others. One may therefore ask what further conditions must be in place for such

regularities to exist. This is a difficult question; my best guess is that there may be no general answer to it. Perhaps all we can do is offer particular explanations for why certain causes but not others regularly precede their effects. For example, we can explain why inscriptions on public monuments are generally caused by the events they represent in terms of certain sociological facts about monuments (in particular, social structures and incentives that explain why mistakes on public monuments are rare). Likewise, we can appeal to facts about thermodynamics and anatomy to explain why footprints of a certain shape must have been caused by human beings (rather than by spontaneous fluctuations or by other animals). And so on. But there may be nothing more general we can say about why certain causes but not others are quasi-necessary for their effects. If that is correct, then global constraints on the causal and statistical structure of the world can explain why it is easier for effects to be reliable indicators of their causes than vice versa; but any explanation of why a particular effect is a reliable indicator of a given cause but not others will have to proceed locally, on a case-by-case basis.

7. Causation and Initial Randomness

While the account I have developed traces the knowledge asymmetry back to facts about causation, it should now be clear that it does not presume any particular view of what causation is. Instead, my approach only relies on two minimal and uncontroversial assumptions about causal structure. The first is that local causes nomically underdetermine their effects and vice versa. This fact is a consequence of the physics of our world, and does not presuppose any particular understanding of causation. The second is that the causal structure of our world conforms to the causal modelling axioms, chief among them the causal Markov condition. That assumption is compatible with all major accounts of causation (be they in terms of regularities, probabilities, counterfactuals, interventions, processes or dispositions). Indeed, given the striking successes of the causal inference techniques based on those axioms (Spirtes, Glymour, and Scheines 2000; Pearl 2009), any serious account of causation has better be

compatible with it.¹⁶ Nor does my view rely on any particular assumptions about the status of the causal Markov condition. Like many, I am inclined to regard that condition as something like a conceptual or metaphysical constraint on causation that captures the core feature distinguishing causes from effects. But my explanation of the knowledge asymmetry is compatible with more modest stances on the causal Markov condition, such as a view on which the condition is a default principle about causation that is generally correct but fails in some special (such as quantum) contexts.

Assuming that causes and effects display the asymmetry embodied in the causal Markov condition, a further question is why that asymmetry is aligned with the arrow of time. That is, why is it that in our world the factors that count as causes according to the causal Markov condition come earlier in time than those that count as effects? (If those two arrows pointed in different directions, my account implies that we would know more about the future than the past.) The leading answer to that question appeals to initial randomness or ‘initial micro-chaos’ (Horwich 1987; Arntzenius 2010; Frisch 2014; Stradis 2021). This is the fact (or assumption) that the microscopic parameters of the initial macrostate of the universe were randomly distributed, resulting in a uniform initial probability distribution over the microstates compatible with that macrostate. Roughly, initial randomness implies that spatially separated variables that represent initial conditions of a system must be statistically independent (whereas no such constraint holds for final conditions), and hence count as causes by the light of the causal Markov condition. (See Arntzenius (2010: §3.3) for a more formal argument). If this is so, on my account the knowledge asymmetry ultimately traces back to this initial randomness assumption.¹⁷ And as I noted, my account must also appeal to PH to block a

¹⁶ One might even demand that a good account of causation should be able to explain why those axioms are true (Papineau 2022).

¹⁷ Note that on this story, the fact that we can derive the time-asymmetry of causation from initial randomness is taken to mean that the former explains the latter. As Frisch (2014) observes, however, the derivation also works the other way: if causes precede their effects, and the causal Markov condition holds, we should find that initial conditions are randomly distributed. So another option is to regard the time-asymmetry of causation as a primitive fact about our world that explains initial randomness. For broadly Humean reasons, I find that

certain sceptical threat against our knowledge of the past, in the manner identified by Albert. And so it turns out that on the analysis I have developed, the knowledge asymmetry finds its ultimate source in the same facts which, according to Albert, underlie all temporal asymmetries: the low-entropy initial state of the universe, and an initial randomness postulate.¹⁸

While my view thus lends support to Albert's stance of the ultimate origins of time-asymmetries, our two accounts are not equivalent. True, they both give the PH a similar role, that of ruling out 'anti-thermodynamic' trajectories that undermine the reliability of records of the past. But whereas Albert's account gives PH the central role in explaining the knowledge asymmetry, on my analysis it is initial randomness that does the bulk of the explanatory work. It does so by explaining the phenomenon which (I argued earlier) Albert's account leaves mysterious: the fact that localized macroscopic signs of the past are more prevalent than signs of the future. It is the statistical independence of initial conditions that explains why causes provide reliable evidence of their effects in restricted contexts only, whereas correlations between final conditions explain why a single record can provide decisive evidence for its cause.

I suspect that one reason why Albert's account misses this connection between initial randomness and the knowledge asymmetry is that it proceeds on the idealization that we have epistemic access to the entire present macrostate of the world (Albert 2000: 114). But as I have argued, the epistemic asymmetry is tied crucially to the fact that we can only ever observe a very *limited* portion of the macrostate. Indeed, our inability to predict the future is due precisely to the fact that to confidently predict the future we need to know what happens in an enormous region of the current macrostate—much

option unattractive, but it is worth noting that those better inclined toward it can still take on board the explanation of the knowledge asymmetry I have proposed. This is another respect in which my account presupposes little by way of causal metaphysics.

¹⁸ The idea that a causal Markov account of the knowledge asymmetry is compatible with Albert's view of temporal asymmetries a central theme in Stradis 2021. See his paper for an in-depth discussion of how the state of initial randomness that underlies causal asymmetries in time relates to Albert's postulates about initial conditions of the universe.

larger than what we can directly observe. By contrast, because effects of common causes are correlated, acquiring knowledge about the past does not require access to the whole macrostate of the world: in that case, highly partial observations are enough. Once we recognize this crucial difference, it becomes clear that the fact that causes but not effects are independent is the key fact that underlies the knowledge asymmetry. But this can come into view only by starting from the fact that we have access only to a very limited portion of the full macrostate of the world.

8. Other Causal Accounts of the Knowledge Asymmetry

The account of the knowledge asymmetry presented in this paper was anticipated by Reichenbach (1956). *The Direction of Time* contains two explanations of the knowledge asymmetry. §19 contains the famous entropic account mentioned in fn. 2. But §21 also contains another discussion of the knowledge asymmetry that connects it to the fork asymmetry and contains several of the ideas discussed above. Thus, Reichenbach notes that if c_1 and c_2 have a joint effect e (thereby forming a fork 'open towards the past'), the statistical independence of c_1 and c_2 means that one can generally not predict e by observing c_1 alone (1956, 181). He also imagines a time-reversed device that reliably records future barometric pressure and notes that this would require many odd correlations between the device's state and other present features of the world, and a temporal reversal of the fork asymmetry (183-6). However, Reichenbach's discussion omits some critical points. In particular, missing is the key idea that virtually *every* forward causal inference involves a fork open towards the past (with c_1 the observed cause and c_2 the other, unobserved causes of e). And his discussion of the time-reversed barometer does not make it explicit that in that situation we can infer future barometric pressure from the device alone precisely *because* there are fine correlations between the device's behaviour and other parts of the macrostate. (Consequently, the fact that in our world the correlation of effects is what enables reliable inference toward the past from localized events does not appear clearly either.) For these reasons, the general connection between the epistemic and the causal arrows does not emerge entirely

clearly from Reichenbach's discussion. This is perhaps why, to my knowledge, subsequent literature on the knowledge asymmetry contains virtually no reference to the ideas presented in §21 of Reichenbach's work.

Recently, Frisch (2014: 224–8) has also offered an account of the knowledge asymmetry along lines similar to mine. His account posits that a record r^* of a past state s^* must satisfy a 'shielding condition' stating that for most possible values of external causes of r^* , the probability of a false positive (r^* without s^*) or a false negative (s^* without r^*) is low. The 'no false positive' aspect of this condition allows for special circumstances in which external causes conspire to produce r^* without s^* . But such causes would then have to all take very specific values at the exact same time. The causal Markov condition makes this scenario improbable and thus explains how r^* can reliably indicate s^* . To show that the reasoning does not work toward the future, Frisch takes the example of a bomb about to be dropped on a city. Here the bomb drop can easily be a 'fake record' of the city's destruction. For instance, a missile may destroy the bomb mid-air, so that the city remains unscathed. That scenario also requires highly specific correlated states in the future, for example scattered traces of the bomb-missile collision. But the causal Markov condition allows for such correlations. This explains why the bomb drop is not a record of the city's destruction, and more generally why no or few records of the future exist.

While Frisch's and my account derive the knowledge asymmetry from the same causal asymmetry, they also differ in important ways. And while Frisch's account is illuminating in various respects, I believe these differences count in favour of mine. Contrary to his, my account does not assume that r^* can be a record of s^* only if the probability of a false negative is low, and shows clearly that this assumption is incorrect. In the *Mayflower* example, external factors could easily have prevented s to cause p , and yet p is a record of s all of the same. It is actually unclear whether Frisch's account really needs this assumption: his actual derivation relies only on the 'no false positive part' of the shielding condition. But by making it clear that this assumption is false, my account better identifies the conditions that records must satisfy, and provides a more satisfactory picture of the knowledge asymmetry. For as we have seen, we can often infer a cause from one of its later effects that would itself have been utterly unpredictable by observing the cause alone. But Frisch's 'no false negative condition'

implies that r^* is a record of s^* only if s^* itself makes r^* highly probable, and hence cannot make sense of this aspect of the knowledge asymmetry. A further advantage of my account is that it more clearly identifies the conditions under which we can infer the future from a local present event. While Frisch recognizes that such inferences are possible (stating that ‘there are no (*or at least many fewer*) records of the future’ (2014: 227, my emphasis)), his account does not explain under which conditions an event can act as a ‘record’ of a future one.

Finally, further reflection on Frisch’s bomb example brings into focus a third advantage of my account. In this example, the effect (city destruction) lies in the immediate future of the cause (the bomb dropping). But as noted above, inferring the near future is often fairly easy. And imminent city destruction can indeed be reliably inferred from the bomb dropping provided one knows (via observation or background knowledge) that no missile is incoming. Frisch’s example thus fails to highlight the most severe obstacle to knowledge of the future: the fact that as a rule, local causes massively underdetermine their effects, except for those in the immediate future. And consequently, it does not fully bring to light what I have argued to be the main mystery behind the knowledge asymmetry: the fact that this underdetermination holds in the reverse direction as well, and yet is no obstacle to backward inference. This is also a lacuna in Stradis’s (2021) recent fork asymmetry account of the knowledge asymmetry. To explain the reliability of a record Stradis proposes to think of a record of c as composed of many correlated sub-components that together (although perhaps not individually) make c highly likely. (Think for instance, of the many ink blobs that make up a photograph.) But this assumes that a single localized record such as a photograph makes its cause highly likely, leaving it mysterious how this can be so despite the fact that the cause depends on many other events in the present macrostate.¹⁹ By showing how the causal Markov condition makes underdetermination an obstacle to knowledge

¹⁹ I hasten to add that Stradis’s paper also makes several important points that usefully complement my account. For instance, Stradis notes that our access to the distant past is facilitated in part by the fact that past events often leave many records around, which increases the chance that some of those records persist for a long time.

toward the future only, my account more precisely identifies how that condition explains the knowledge asymmetry.

Acknowledgements

I would like to thank Tobias Henschen, Andreas Hüttemann, Kian Salimkhani, Bram Vaassen, Martin Voggenauer, the Editor of this journal and two anonymous referees for very helpful comments.

ORCID

0000-0001-5002-7517

References

- Cleland, Carol (2011) 'Prediction and Explanation in Historical Natural Science', *British Journal for the Philosophy of Science* **62**: 551–82. doi:10.1093/bjps/axq024.
- Earman, John (1974) 'An Attempt to Add a Little Direction to "The Problem of the Direction of Time"', *Philosophy of Science* **41**: 15–47. doi:10.1086/288568.
- Elga, Adam (2001) 'Statistical Mechanics and the Asymmetry of Counterfactual Dependence', *Philosophy of Science* **68**: 313–24. doi:10.1086/392918.
- Fernandes, Alison (2022) 'How to Explain the Direction of Time', *Synthese* **200**: 389. doi:10.1007/s11229-022-03818-4.
- Fernandes, Alison (Forthcoming) 'Time, Flies, and Why We Can't Control the Past', in Barry Loewer, Eric Winsberg and Brad Weslake, eds., *Time's Arrows and the Probability Structure of the World*. Harvard University Press.
- Frisch, Mathias (2005) *Inconsistency, Asymmetry and Non-Locality: A Philosophical Investigation of Classical Electrodynamics*. Oxford University Press.
- Frisch, Mathias (2014) *Causal Reasoning in Physics*. Cambridge University Press.
- Healey, Richard (1991) 'Review of Asymmetries in Time: Problems in the Philosophy of Science', *The Philosophical Review* **100**: 125–30. doi:10.2307/2185522.
- Ismael, Jenann (2016) *How Physics Makes Us Free*. Oxford University Press.
- Loewer, Barry (2012) 'Two Accounts of Laws and Time', *Philosophical Studies* **160**: 115–37. doi:10.1007/s11098-012-9911-x.
- Papineau, David (2022) 'The Statistical Nature of Causation', *The Monist* **105**: 247–75. doi:10.1093/monist/onab034.
- Pearl, Judea (2009) *Causality: Models, Reasoning, and Inference*, 2nd edition. Cambridge University Press.
- Reichenbach, Hans (1956) *The Direction of Time*. University of California Press.
- Spirtes, Peter, Clark Glymour, and Richard Scheines (2000) *Causation, Prediction, and Search*, 2nd edition. MIT Press.
- Stradis, Athamos (2021) 'Memory, the Fork Asymmetry, and the Initial State', *Synthese* **199**: 9523–47. doi:10.1007/s11229-021-03214-4.
- Woodward, James (2003) *Making Things Happen*. Oxford University Press.

