

How to Explain the Direction of Time

Alison Fernandes

[16 July 2022]

Abstract: Reichenbach explains temporally asymmetric phenomena by appeal to entropy and ‘branch structure’. He explains why the entropic gradients of isolated subsystems are oriented towards the future and not the past, and why we have records of the past and not the future, by appeal to the fact that the universe is currently on a long entropic upgrade with subsystems that branch off and become quasi-isolated. Reichenbach’s approach has been criticised for relying too closely on entropy. The more popular approach nowadays is to appeal instead to a *particular* low-entropy initial state—Albert’s ‘Past Hypothesis’. I’ll argue that this neglect of Reichenbach’s approach is unwarranted. A Reichenbachian account has important advantages over Albert’s: it correctly identifies the minimal temporally asymmetric posit needed to derive key temporal asymmetries and it offers a more adequate account of the record asymmetry. While a Reichenbachian account needs to be supplemented, it provides the right foundations for explaining temporally asymmetric phenomena and what we might ultimately mean by ‘the direction of time’.

1. Introduction

Many phenomena at our world are temporally asymmetric: directed differently towards the past and future. Smoke disperses in a room but doesn’t recombine. We have records of the past but not the future. The laws of fundamental physics are, however, by and large temporally *symmetric*—they have no temporal orientation.¹ So, there is a challenge to explain why certain phenomena are temporally asymmetric.

One program, following the work of Ludwig Boltzmann, explains temporally asymmetric phenomena by relating them to a temporal asymmetry of entropy.² The Boltzmannian entropy of a system is proportional to the log of the measure of microstates in phase

¹ The laws of Newtonian mechanics and General Relativity as well as the Schrödinger equation are all ‘time reversal invariant’—they take the same form in both temporal directions. See Sklar (1993, pp. 246–8), Albert (2000, pp. 11–14) and North (2008) for precise formulations. While there are candidate theories of quantum mechanics whose dynamical laws aren’t time reversal invariant, there aren’t yet convincing accounts of how these violations explain the temporal asymmetries above.

² For general concerns with this Boltzmannian program see Earman (1974), Sklar (1993), Leeds (2003) and Maudlin (2007a).

space compatible with its macrostate, using the standard Lebesgue measure. Phase space is a continuous space containing six dimensions for every particle in the system, one for every dimension of position and momentum. The entropy of our universe, at least around here and now, increases towards the future and not the past.³ If one takes a Boltzmannian approach to explaining a temporally asymmetric phenomenon, one explains that phenomenon either using a) the *fact* that the entropy of the universe (at least with high probability and/or at least around here and now) increases towards the future and not the past, or using b) the entropic facts that *explain* the (perhaps probabilistic and/or local) asymmetric entropy increase of the universe. The temporally asymmetric phenomenon is due either to a universal or widespread entropic asymmetry or due to the entropic facts that explain that entropic asymmetry. For example, one might explain the fact that the entropy gradients of isolated systems are oriented towards the future and not the past, or the fact that we have records of the past and not the future, using *the asymmetric entropy increase of the universe* (option a) or using features that are taken to *explain* the asymmetric entropy increase of the universe—typically the universe’s low-entropy initial state (option b).⁴ Under either alternative, the temporal asymmetry to be explained is always *related to* a universal or widespread temporal asymmetry in the entropy of the universe, using time-symmetric laws and probabilities. Boltzmannian approaches may attempt to *derive* probabilistic (and even law-like) generalisations about how systems behave, but there is still only one ultimate source of the temporal asymmetry to be explained—an asymmetry in the entropy values of the universe (rather than a temporally asymmetric lawlike generalisation about how entropy behaves).⁵

³ This is not to suppose there is a direction of time independent of entropic asymmetry. The direction towards the future may simply be the direction in which entropy increases.

⁴ To explain the Second Law of Thermodynamics, one might also need to derive its status as a law or a ‘lawlike generalisation’ (Price 2002a, 2002b). For this reason, the explanandum I’ll focus on is the *fact* the vast majority of isolated subsystems have entropy gradients oriented towards the future, rather having entropy gradients oriented towards the past or having no gradient at all—a generalisation that covers most instances (and perhaps only one that holds in a limited part of the universe). One might also attempt to derive generalisations about the probable behaviour of isolated systems. I won’t assume one needs to derive probabilistic generalisations about behaviour in order to explain the entropy rise of isolated systems, so as not to rule out certain Boltzmannian approaches by fiat (Price 2002a, 2002b). My thanks to an anonymous reviewer for pressing me on these points.

⁵ In this sense, Boltzmannian approaches count as ‘one-asymmetry’ approaches in Price’s terminology (2002b), even though they may also attempt to *derive* law-like generalisations concerning how entropy behaves in isolated systems—*pace*, Price. The relevant contrast is ‘two-asymmetry’ approaches, which use both an asymmetric entropic condition and a temporally asymmetric posit about behaviour.

In *The Direction of Time* (1956),⁶ published posthumously, Hans Reichenbach adopts a Boltzmannian approach. He explains a range of temporally asymmetric phenomena using postulates about the entropy rise of the universe and ‘branch structure’—‘branch structure’ referring to particular ways in which subsystems of the universe become isolated.⁷ These posits are used to explain why isolated systems have entropy gradients oriented towards the future and not the past (in most cases), why there are records of the past and not the future, and other temporally asymmetric phenomena.

Part of Reichenbach’s purpose in explaining these phenomena is to give what he calls an ‘explication’ of the direction of time (p. 23). Some have thought time has a direction in the sense of there being an intrinsic asymmetry in time itself (Earman 1974; Maudlin 2007a, Ch. 4). Reichenbach doesn’t believe time has a direction in that sense. Instead, his explication is a suggestion for how we might *replace* our vague concept of the direction of time with a more precise concept that has *empirical* content and then explain that empirical content in scientific terms (pp. 23–4). He aims to explain a precisified sense of the direction of time. Explicating the direction of time will require explaining relatively basic temporally asymmetric phenomena, such as the entropy rise of isolated subsystems, as well as more complex temporally asymmetric phenomena that are more closely related to our ordinary concept of the direction of time—including the record asymmetry.⁸

Despite its acknowledged influence, Reichenbach’s approach to explaining temporal asymmetries is often given short shrift. Firstly, his explanation of the asymmetry of records is criticised for relying too closely on entropy. The complaint is that Reichenbach is unable to account for records that are not low-entropy or entropy-increasing (Earman 1974, p. 45; Horwich 1987, p. 83; Albert 2000, p. 124) (see Section 4). Secondly, his explanations of the entropy rise of subsystems are dismissed for their obscurity, even by sympathetic readers (Glymour and Eberhardt 2016). More generally, while Reichenbach’s explanations are acknowledged as important forerunners to contemporary accounts (Horwich 1987, pp. 68–70), they are rarely nowadays given serious treatment. Recent

⁶ All references to Reichenbach are to this work, unless otherwise noted.

⁷ For related approaches, see Schrödinger (1950), Smart (1967), Grünbaum (1973) and Davies (1974). For reasons of space, I won’t discuss the differences between these approaches, except to note that some objections to these accounts do not apply to Reichenbach’s—see footnote 11.

⁸ I won’t attempt to settle exactly what temporally asymmetric phenomena are required to explicate the direction of time. Such a project would presumably require empirical investigation into our temporal concepts. For recent work combining psychology and physics to explicate apparent features of time, see Callender (2017). For recent empirical work on our concept of the direction of time, see Lee et al (Forth.).

debates concerning Boltzmannian approaches (Earman 2006; Winsberg 2004a; Maudlin 2007a Ch. 4; Frisch 2007) typically focus on David Albert's account (2000, 2015), which aims to avoid Reichenbach's direct appeals to entropy. Even when branch systems accounts are discussed (Albert 2000, p. 89; Winsberg 2004b; Frigg 2008, pp. 128–30) the focus is typically not on Reichenbach's approach.

I will argue that there is a 'Reichenbachian' account worth recovering and that it has important advantages over Albert's. This Reichenbachian account correctly identifies the relevant temporally asymmetric posit needed to explain the entropy rise of isolated subsystems, as well as the record asymmetry. While the account needs to be supplemented, it provides the right foundation for making sense of two key temporally asymmetric phenomena that often form the basis for explaining other temporally asymmetric phenomena—particularly the temporal asymmetry of causation (Horwich 1987; Albert 2000, 2015; Loewer 2007; Fernandes 2017).⁹ Because the temporal asymmetry of causation is itself central to explaining many other temporally asymmetric phenomena, if one can explain the entropy rise of isolated system and the record asymmetry in Reichenbachian terms, one has grounds for thinking that a Reichenbachian approach can work for explaining temporally asymmetric phenomena more generally.¹⁰ If one further follows Reichenbach in thinking that all or some subset of these temporally asymmetric phenomena explicate direction of time, one will have the resources to explain (a precisified sense of) the direction of time.

The paper proceeds as follows. In Section 2, I present Reichenbach's explanation for why isolated systems have entropy gradients oriented towards the future and not the past. In Section 3, I evaluate Reichenbach's explanation by comparing it to Albert's. I argue that, when supplemented with a probabilistic posit, the Reichenbachian account is superior. In Section 4, I consider Reichenbach's explanation of the temporal asymmetry

⁹ Others are agnostic on the relation between the temporal asymmetries of records and causation (Callender 2017, Ch. 11). There are other relatively basic temporal asymmetries that I don't discuss, including those of radiation and electrodynamics. These asymmetries are typically less central to debates about the direction of time since they don't seem to play a role in explaining further temporal asymmetries and aren't taken to be part of our intuitive concept of time's direction. For discussion of these asymmetries, see Price (1996, Chs. 3–4), Frisch (2014, Ch. 7) and Callender (2021).

¹⁰ For recent work using causation to explain psychological temporal asymmetries, see Fernandes (2022) and references therein. See Fernandes (2017) for my take on how a range of temporally asymmetric phenomena are related. One could, of course, explain some temporal asymmetries using a Boltzmannian approach, but deny that approach works for other asymmetries—see, for example, Winsberg (2004b, p. 711). My thanks to an anonymous reviewer for pressing me on this point.

of records. After clearing up some confusion regarding its scope, I evaluate the Reichenbachian approach in comparison with Albert's (Section 5). I argue that only the Reichenbachian account explains the record asymmetry.

2. Reichenbach's Hypothesis of Branch Structure

Reichenbach attempts to explain why the vast majority of isolated subsystems have entropy gradients oriented towards the future, rather than having entropy gradients oriented towards the past or having no gradient at all. In this sense, he intends to explain why their entropies 'rise' towards the future. When I refer to the 'entropy rise of isolated systems', this is the explanandum I have in mind. Reichenbach bases his explanation of the entropy rise of isolated systems on the 'hypothesis of branch structure'. The rough idea is that the entropy rise of the whole universe will imply entropy rises in individual isolated subsystems. While later branch system approaches use elements of Reichenbach's account (Davies 1974), they don't all share Reichenbach's commitment to tracing temporal asymmetries back to an entropy gradient of the universe.¹¹ Reichenbach goes on to argue that these individual entropy rises imply a range of further temporally asymmetric phenomena. The hypothesis of branch structure consists of four posits—empirical claims that would be particularly relevant to the logical derivation of the relevant phenomena (p. 136). I'll present each of the four posits, followed by a brief commentary, before considering the derivation.

Posit 1. 'The entropy of the universe is at present low and is situated on a slope of the entropy curve'
(p. 136).

The universe is currently at non-maximal entropy and its entropy rises in one temporal direction (towards what we call 'the future') and decreases in the other temporal direction (towards what we call 'the past') over a significant period of time. Elsewhere,

¹¹ Later branch system approaches often apply a low-entropy constraint and/or a statistical postulate only at the time the branch system first comes into being in order to avoid reversibility objections concerning how the system began in its low-entropy state. As Sklar notes (1993, p. 325–32) they introduce a second temporal asymmetry as an explanans. Because Reichenbach avoids this move, he avoids objections to Davies's proposal made by Price (1996, pp. 44–6) and Albert (2000, p. 89), responded to by Frigg (2008, pp. 129–30) and Winsberg (2004b, pp. 712, 715–7), such as concerns about when exactly branch systems come into being. There are some problems shared by these accounts, concerning how to derive they entropy gradients of isolated subsystems from that of the universe—I discuss these in Section 3.1. The Reichenbachian account I defend makes use of both an entropy gradient and Albert's account of how a probability postulate applied to the whole universe can allow one to derive probabilities concerning isolated subsystems and so avoids Price's and Albert's objections.

Reichenbach talks of the universe being on an entropy ‘upgrade’ (pp. 118, 130, *passim*). While talk of ‘upgrade’ suggests a prior temporal direction, it is the direction in which entropy is increasing that determines what will count as an upgrade. The entropic upgrade must take place over a ‘long’ enough period of time (pp. 118, 131) to encompass all the temporally asymmetric phenomena to be explained.

Posit 2. ‘There are many branch systems, which are isolated from the main system for a certain period, but which are connected with the main system at their two ends’ (p. 136).

‘Branch systems’ are subsystems of the universe which are energetically isolated from the rest of the universe for a period of time and in energetic contact on either side of this period. Given the laws of our universe, this isolation will only ever be partial. What is needed is that ‘the process within the subsystem represents energy exchanges which are large compared with the interaction with the environment’ (pp. 117–8). Subsystems that are ‘quasi-isolated’ in this sense might include a box of gas, a thermos of coffee or a stone tablet—all insulated to some extent from interactions with the larger environment. Plausibly, subsystems may be quasi-isolated with respect to some variables but not others (Elga 2007).

Posit 3. ‘The lattice of branch systems is a lattice of mixture’ (p. 136).

This posit is the most difficult. To give the rough gist, Reichenbach intends to derive probabilities for the behaviour of *single* systems over time by considering the behaviour of systems of their *type*. If branch systems are ‘lattices of mixtures’, the probabilities for the behaviour of single systems *across time* are derivable from relative frequencies for systems of that same type at a *single time*. ‘Across-time’ probabilities for systems mirror their ‘across-system’ probabilities. So, even if one is skeptical of their being fundamental probabilities for the behaviour of single systems across time, as Reichenbach is, the probabilities for the behaviour of single systems can be derived.

To see how this might work, consider a sequence of events in a single system over time, such as the approximate location of gas particles in a box. Say we want to know the conditional probability of an event of type *B*, given an earlier event of type *B*. For example, we might want to know the conditional probability of finding most of the

particles on the left-hand side, given that they were mostly on the left-hand side a moment ago. If we don't have a long series of observations of this box, we might attempt to derive this conditional probability by looking to systems of the same *type*. Imagine creating a 'lattice' that records what event types occur, where the columns represent different moments in time and the rows represent different single systems falling under the same general type. We might look to determine the conditional probability of an event type *B*, given an earlier event type *B* in a single system by looking at the systems of the same type at a single time and considering *how many* of them have an event of type *B*, given they have an earlier event of type *B*. For example, we might consider, for boxes of gas of the same dimension with the same number of gas particles, what *proportion* of these have their particles mostly on the left-hand side given they were mostly on the left-hand side in the previous moment. If a system type is a 'lattice of mixture' for a given event type, the conditional probabilities for a single system's behaviour across time are given by the *proportion* (or relative frequency) of the same type of behaviour across systems of that same type at a single time.

In a lattice of mixture, furthermore, the absolute (non-conditional) probability for a single system being in a given state after some time will be given by the relevant frequencies across systems of that same type at that time. For example, if the lattice of gas boxes is a lattice of mixtures with respect to certain macroscopic parameters, then, after some time, the absolute probability of finding a box of gas with most of its particles in the left-hand side is given by the proportion of boxes of that same general type that are in that state type, after some time.

Why think that branch systems will be lattices of mixture? Reichenbach reasons that, for isolated systems, the laws of chance and other shuffling mechanisms operate in a similar manner on single systems and systems of that same type. So, conditional probabilities will be the same in either case and will tend to override any peculiarities of the single system, making for an equivalence in absolute probabilities after some time (p. 122). Note, however, that this kind of frequency matching only works *in one* temporal direction (in the case of conditional probabilities) or *after* some time in that temporal direction (in the case of absolute probabilities). This is because systems may have 'started out' in particular states that make their evolution towards their past states *unlikely*. Frequency matching only works towards what we call the future. Note, however, that the posit can

be stated, as I have done, in ‘neutral time’—without reference to the past or future. The requirement is that there is some (single) temporal direction in which the frequency matching works. I’ll consider whether this posit smuggles in a temporal asymmetry below (Section 3.1).

Posit 4. ‘In the vast majority of branch systems one end is a low-point [of entropy], the other a high point [of entropy]’ (p. 136).

Branch systems almost always start and end at different entropies, implying that they almost always have entropy slopes.¹² They do not typically start and end at maximal entropy.

From these four posits, Reichenbach derives a fifth (pp. 131–5; 143):¹³

Posit 5. ‘In the vast majority of branch systems, the directions toward higher entropy are parallel to one another and to that of the main system’ (p. 136).

In other words, branch systems tend to rise in entropy in the same temporal direction, a direction parallel to that in which the entropy of the universe rises. If Reichenbach can derive this posit, he can plausibly explain a generalisation sometimes taken to be captured by the Second Law of Thermodynamics—that the vast majority of isolated subsystems have entropy gradients oriented towards the future, rather than having entropy gradients oriented towards the past or having no gradient at all. Posit 5 also forms the basis of Reichenbach’s explanations of other temporally asymmetric phenomena, including the asymmetry of records (Section 4).

3. How to Explain the Entropic Rise of Isolated Systems

3.1 Problems of Parallelism and Derivation

To evaluate Reichenbach’s account, we need to first consider whether Posit 5 can be derived from the other posits. Surprisingly, Reichenbach claims Posit 5 can be derived from just Posits 3 and 4 (pp. 138–9). Insofar as branch systems are lattices of mixtures (Posit 4), and have entropy gradients (Posit 3), their entropy gradients will be parallel to

¹² It may be sufficient if merely a large *proportion* of branch systems have entropy slopes.

¹³ Reichenbach calls these ‘assumptions’ (p. 136). Even though some of these posits may be derivable from others, the idea is not to presuppose that issue by giving them different labels.

that of the main system (Posit 5). But there is an immediate problem. While Posit 5 states that the entropy rise of branch systems is parallel to that ‘*of the main system*’, nothing about the main system is mentioned in Posits 3 or 4—they refer only to branch systems.

In response, perhaps Posit 5 should be amended to omit reference to the main system, so that the entropy rises of branch systems are simply parallel to one another. But this doesn’t seem to be what Reichenbach has in mind. Consider the following quote, from just before Reichenbach states these posits (p. 135, my emphasis):

The direction of time is supplied by the direction of entropy, because the latter direction is *made manifest* in the statistical behavior of a large number of separate systems, generated individually in the general drive to more and more probable states.

A central feature of Reichenbach’s approach is that the entropic behaviour of the universe is supposed to be *reflected* in that of branch systems. It won’t do to simply omit the entropy rise of the universe from the account. Moreover, Reichenbach claims (p. 137) that Posit 5 would have to be altered if Posit 1 were to exclude reference to the entropy of the universe. So, Reichenbach seems to think Posit 5 *should* refer to the main system. While one could, nevertheless, revise Posit 5 to omit reference to the main system, doing so will not solve a second problem, considered below.

Sklar (1993, p. 325) suggests that Posit 3 could be revised to claim that the lattice of the main system is also a lattice of mixture. There are two problems with this suggestion. Firstly, as I’ll argue below (Section 3.2), it goes counter to Reichenbach’s views on probability. It can’t be what Reichenbach himself had in mind. More seriously, neither of these suggestions solves a second problem with the derivation. The problem, also noted by Sklar (*ibid.*, p. 324), is that Posit 3, revised or not, presupposes an unexplained ‘parallelism’ in the behaviour of branch systems. Posit 3 states that ‘*The lattice of branch systems is a lattice of mixture?*’. This posit presupposes that the temporal direction in which shuffling operates, and in which conditional probabilities for individual systems are derivable from the behaviour of systems of their same type, is *the same* for all branch systems. Posit 3 does almost all the work in explaining why the entropy of isolated systems rises towards the future. But Reichenbach’s derivation does not explain *why* these

systems behave in parallel. While assuming parallelism is logically distinct from assuming a temporal asymmetry in the behaviour of branch systems, if the problem regarding derivation is also to be solved, it seems Reichenbach must assume branch systems evolve in the same temporal direction as the entropy gradient of the universe (to what we call ‘the future’). Adopting this assumption would turn Reichenbach’s attempt to derive a thermodynamic asymmetry from a single temporal asymmetry into an example of a ‘two-asymmetry’ approach—with all the problems that entails: see Price (2002a, 2002b).

Neither of the suggested revisions answers this parallelism problem. If we leave Posit 3 unrevised and revise Posit 5 (first option), there is still the unexplained parallelism in Posit 3. If we revise Posit 3 to refer to the main system (Sklar’s suggestion), doing so only makes the problem worse, as Sklar notes (1993, p. 324). If we assume the universe is a lattice of mixture, we simply *assume* that the probabilistic behaviour of branch systems parallels that of the main system—without explaining why this is so.

While Sklar gives up on Reichenbach’s account at this point, I suggest there is a solution. To find it, we need to reconsider probabilities. I will argue that introducing a probabilistic postulate, of the kind Albert employs, allows one to derive the entropic behaviour of individual subsystems from that of the main system—solving the derivation and parallelism problems. While introducing such a postulate goes against Reichenbach’s views on probability, it provides a promising account of temporal asymmetries that is Reichenbachian in spirit.

3.2 Probability

Reichenbach adopts a hypothetical frequentist interpretation of probability (pp. 96, 123, 132). He thinks that, while we may loosely talk of probabilities that apply in the single case, these are really claims about the limit relative frequencies in infinite sequences of outcomes. Moreover, in order to justify a probability statement, he thinks there has to be a reasonably long observable actual finite sequence over which the relative frequencies are defined.¹⁴ In the case of branch systems away from equilibrium, Reichenbach thinks that such a sequence can be found by considering a long sequence of systems of that

¹⁴ I don’t endorse Reichenbach’s reasoning here. There needn’t be *actual* and *observable* sequences in order to justify statements about limits in hypothetical infinite sequences. For more on Reichenbach’s views on probability, see Reichenbach (1949) and Eberhardt and Glymour (2011). For general concerns with hypothetical frequentist approaches to probability, see Hájek (2009).

same *type* (at a single time) that are also away from equilibrium.¹⁵ Considering the long term behaviour of a single system won't work, as its long term behaviour (near equilibrium) isn't representative of its behaviour when away from equilibrium. In the case of the whole universe, however, Reichenbach thinks there is no long finite observable sequence that will work. There is only one universe and it spends too long at equilibrium. So, there can't be justified statements about the probable behaviour of the universe across time and Reichenbach can't assume the universe is a lattice of mixture.

The general form of Reichenbach's entropic explanation of temporal asymmetries, however, does not rest on his views about probabilities. One could adopt a different interpretation of probability, replace Posit 3, and then go on to derive Posit 5—solving the derivation and parallelism problems. Such an account would still be Reichenbachian in that it appeals to posits about entropy and branch structure (Posits 1, 2 and 4) to explain temporal asymmetries.

I suggest Posit 3 ('The lattice of branch systems is a lattice of mixture') should be replaced with:

Posit 3: 'The probabilistic behaviour of branch systems is derivable from a suitable probability postulate applied to the universe.'*

If the entropic behaviour of individual systems is to be derived from that of the universe, some probability postulate is required. There are different probabilistic postulates one could use. Here are two straightforward options.¹⁶ One could use a flat probability measure (the Lebesgue measure) applied to phase space over the microstates of the universe compatible with its initial macrostate—what Albert calls the 'Statistical Postulate'. Alternatively, one could use the same flat probability measure but applied to the *entire* phase space of the universe at any time—the 'Lebesgue Postulate'—without conditionalising on the universe's macrostate. Because the Lebesgue Postulate does not

¹⁵ The need for a long sequence also explains why Reichenbach requires there to be large numbers of branch systems with entropy gradients (Posits 2 and 4).

¹⁶ There are other approaches to probability in statistical mechanics, such as typicality approaches (Maudlin 2007b), as well as different ways of thinking about what a probability postulate amounts to, including Gibbsian ensemble approaches (Sklar 1993, Ch. 5) and approaches based on indifference reasoning. I have focused on two plausible options that fit well within the Boltzmannian program.

presuppose conditionalising on the initial state, it is more general and is my preferred option.¹⁷

While a probability postulate does not fit with Reichenbach's views about probability, it is not in tension with frequentist approaches more generally. Loewer (2004) and Albert (2015, Ch. 1), for example, argue that the Statistical Postulate is derivable from a sophisticated frequentism, such as a Lewisian Best Systems Analysis of probability. Moreover, a probability postulate is compatible with empiricism, even if one is not a frequentist. The probability postulate need not be a postulate of logic or based on indifference reasoning but may be an empirical posit about real objective modal features of the world. See Fernandes (Forth. a.) for my preferred non-Humean account of objective probability.

By using Posit 3*, the behaviour of branch systems can be derived from the probabilistic behaviour of the universe. Albert (2000, Chs. 3–4) argues that, given the type of macrodynamics usual at our world, when subsystems of the universe become isolated, there is an overwhelmingly high probability derivable from the Lebesgue Postulate that they will rise in entropy over time (if they are isolated at non-maximal entropy) or that they will stay at or near maximal entropy (if they are isolated at equilibrium). The regions of phase space which correspond to entropy decreasing or low-entropy maintaining macroscopic behaviour towards the future are incredibly tiny and are randomly distributed throughout the phase space of the universe such that, when a subsystem is isolated from the main system, it is overwhelmingly likely not to be on one of these low-entropy trajectories. There may be very unusual cases where macroscopic dynamics are highly sensitive to the microstates of the system—so entropy decreasing behaviour is likely to occur.¹⁸ But entropy decreasing behaviour is highly improbable for the usual macroscopic dynamics we encounter, which are insensitive to a system's particular microstate.

While there are aspects of Albert's account that require further elaboration and defence, it provides a plausible starting point for solving the derivation and parallelism problems.

¹⁷ These postulates will, of course, do equivalent work when one *does* conditionalise on the initial macrostate.

¹⁸ For example, there may be Maxwell-demon cases where the demon sets up an isolated system that then spontaneously decreases in entropy—see Albert (2000, Ch. 5).

Once we conditionalise on there being many branch systems (Posit 2), with entropy gradients (Posit 4), and the universe being in the middle of an entropic slope (Posit 1), it is overwhelmingly likely that the entropy gradients of isolated systems are parallel to one another and to that of the main system. A branch system having an entropy gradient implies it either starts or ends in a low-entropy state. Conditional on a branch system starting out in a low-entropy state, it is overwhelmingly likely that the system rises in entropy towards the future, rather than decreases—there are gradients in one temporal direction and not the other. A branch system may also rejoin the main system in a low-entropy state. But, conditionalising on these states and the entropy slope of the universe, it is much more probable that such a system was isolated in an even *lower* entropy state in the past than that it was isolated in a higher entropy state and evolved to a low-entropy state. Altogether, gradients in parallel with the main system are much more probable than gradients in the opposite direction. So, parallelism is not presupposed, but derived. The asymmetric posit that does the work here is Posit 1—the entropy slope of the universe—rather than a posit about the particular state that the universe began in.

3.3 Entropy

So far, I have argued that Reichenbach’s account needs to be supplemented with a probability postulate. Next, I will argue that a Reichenbachian account, thus supplemented, has an important advantage over its most prominent Boltzmannian rival—Albert’s account. Only the Reichenbachian account correctly identifies the minimal temporally asymmetric posit required to derive the entropy rise of isolated subsystems.

To explain the entropy rise of isolated subsystems, Reichenbach uses the fact that the universe is currently on an entropy slope.¹⁹ Albert, in contrast, uses the ‘Past Hypothesis’—the claim that the universe started out in the *particular* low-entropy macrostate that it did—sometimes combined with the claim that there is no analogous ‘Future Hypothesis’ (2000 p. 119).²⁰ Both Reichenbach and Albert use global conditions to explain the probabilistic behaviour of isolated subsystems which in turn explain why

¹⁹ One could replace the entropy posit with a similarly general posit with similar consequences, such as assuming micro-randomness of initial conditions combined with micro-order of final conditions. For reasons of space, I won’t pursue such alternatives here.

²⁰ I take a Future Hypothesis to be a later macrostate that is highly improbable given the Past Hypothesis, the dynamical laws and the Lebesgue Postulate. Albert suggests a Future Hypothesis would have ‘some simple macroscopic organization’ (2000, p. 119). I discuss problems with that view below (Section 5.1).

there are vastly more isolated subsystems that have entropy gradients oriented towards the future (rather than towards the past). The posit of a slope and the posit of *an* initial low-entropy state (and no analogously restrictive later or final state) both imply the universe moves from a low-entropy state to a higher-entropy state.

In specifying what kind of entropy slope the universe is on (or what restrictive later or final states are absent) one may need to introduce further restrictions to ensure that the entropy of the universes increases at an appropriate *rate*. I will introduce one such restriction here because it will be useful when explaining the record asymmetry. The probability for later macrostates of the universe given its earlier macrostates should be reasonably high (relative to other possible later macrostates), where the relevant probabilities are derived from the Lebesgue postulate and the fundamental dynamical laws. This restriction ensures that the entropy rise will be at an appropriate rate and that the macroevolution of the universe forwards in time will not be too improbable. This proposal might appear to introduce two independent asymmetries: a low entropy condition (or slope) and the requirement that the universe's forwards evolution (but not its backwards evolution) is reasonably probable. But appearances are misleading. Either one uses just the low-entropy past state and the lack of any analogously restrictive future state (a state that implies that the macroevolution of the universe forwards in time is improbable, given the Lebesgue postulate and the dynamical laws), which implies an asymmetry in which evolutions are reasonably probable. Or, one posits that the universe's forwards evolution but not its backwards evolution is reasonably probable—which will imply there is a low-entropy past state and no analogously restrictive future state.

There are a number of differences between Reichenbach's and Albert's posits. Let me put aside two that are less important. Firstly, Reichenbach is open to the possibility that entropy rises merely within our temporal region of the universe (pp. 125–9), while Albert takes it that the entropy of the universe rises from its initial state until now.²¹ But this difference does not affect the form of their explanations, since the temporally asymmetric phenomena to be explained are all within our temporal region.

²¹ An exception is if the concept of entropy does not apply at the start of the universe (Albert 2000, p. 85; Earman 2006, p. 412).

Secondly, it might seem that only Reichenbach has to assume that the universe is currently on a *long* entropic upgrade, while Albert explains this fact. It might seem, therefore, that Albert and Reichenbach have different explanandum in mind. However, appearances are deceiving. If Albert is to *derive* the probabilistic behaviour of isolated systems, he needs to assume the Past Hypothesis, as well as the fact that the universe lacks a Future Hypothesis anytime soon—the Past Hypothesis alone is not sufficient to derive subsystem or global entropy slopes.²² Moreover, Albert also has to assume that we are currently located near enough to the initial low-entropy state.²³ In other words, Albert has to assume posits that are *logically* sufficient to establish that the universe is currently on a long entropic upgrade, independently of anything about the dynamics of the system or the laws of nature. For this reason, Albert’s account does not offer a scientific explanation (or scientific derivation) of the global entropy slope.

The more important difference between Reichenbach’s and Albert’s accounts is that Reichenbach uses a *general* entropy posit, while Albert uses a posit about a *particular* low-entropy macrostate. Albert explains the general entropy rise of isolated subsystems using the Past Hypothesis—the claim that the universe started out in the particular macrostate that it did. Albert is explicit that a posit about a particular state is required to recover the derivations of empirical behaviour we need fundamental posits to recover, including the generalisations captured by the Second Law of Thermodynamics (2000, p. 96; 2015, p. 5).²⁴ For now, I want to focus on what *temporally asymmetric* posit is needed to explain the entropy rise of isolated subsystems—the posit that introduces a temporally asymmetric component that, combined with temporally neutral posits, accounts for the explanandum. Both Albert and Reichenbach employ a single posit to introduce a temporal asymmetry—they are example of ‘one asymmetry’ approaches, in Price’s terminology (2002a, 2002b). I will argue that, when it comes to identifying this temporally asymmetric posit, no particular macrostate is required.

Here’s the argument:

²² See Price (2002a, 2002b) for a version of this argument—the presence of one boundary condition in no way rules out the presence of others.

²³ My thanks to Mike Hicks for discussion.

²⁴ Part of Albert’s motivation also appears to be that a particular state is required to explain how we reason towards the past. I’ll consider our reasoning towards the past below (Sections 4–5).

1. Explanation requires identifying minimal explanatory posits—features of the world that are required to derive the explanandum, given other suitable explanatory posits.
2. The temporally asymmetric posit needed to derive the general entropy rise of isolated systems (when combined with time symmetric posits concerning probabilities and laws) is an appropriate entropic slope or a posit with that minimal implication—not the particular state.
3. The correct temporally asymmetric explanatory posit is an appropriate entropic slope or a posit with that minimal implication—not the Past Hypothesis.

Premise 1 presumes an account of explanation in which explanation consists in deriving phenomena (or deriving probabilities for phenomena, or phenomena with high probability) using dynamical laws, (perhaps objective probabilities,) and other premises concerning contingent states. Explanation requires using minimal explanatory posits, at least one of which expresses a law, from which the phenomenon can be logically derived.²⁵ There is a restriction to relevant information—posits should be minimal and should not contain information irrelevant to the derivation of the explanandum (given the other posits). While we often won't be able to give a complete explanation, we aim to identify posits that would feature in such an explanation.

My claim is that a posit about the appropriate entropic upgrade (or a posit with that minimal implication) features as a temporally asymmetric posit in a complete explanation of the entropy rise of isolated subsystems but that the Past Hypothesis does not. The Past Hypothesis contains irrelevant information: information about the particular state, beyond its being low-entropy. This information is not required to derive the entropy rise of isolated subsystems: the universe starting out in a different low-entropy macrostate would imply the entropy rise of isolated subsystems, given the other explanatory posits. If the universe is currently situated on an appropriate entropy slope (Posit 1), then, given

²⁵ Models of explanation that would satisfy these requirements include Reichenbach's (pp. 136, 143) and Hempel's covering law model of explanation (Hempel and Oppenheim 1948), each when combined with a restriction to relevant information. While Hempel's account has been much criticised (Salmon 1992), many objections concern whether it can account for *all* scientific explanations—not its use in particular contexts. One influential objection was that Hempel's account didn't capture the temporal asymmetry characteristic of causal explanation (ibid.). However, this feature is a strength in this context. Using a temporally or causally asymmetric account of explanation would presuppose a temporal or causal asymmetry in the *form* of the explanation. Whereas, when explaining temporal or causal asymmetries, we may prefer a model of explanation that doesn't import a temporal or causal asymmetry. For related concerns with using multiple asymmetries when explaining temporally asymmetric phenomena, see Price (2002a, 2002b).

the posits about probability (Posit 3*) and the existence of many branch systems (Posit 2) with entropy gradients (Posit 4), the entropy of these branch systems will, with overwhelmingly high probability, rise in entropy in a temporal direction parallel to that of the entropy rise of the universe. No particular low-entropy state is required.

One might object as follows. Perhaps, given the macrostate of our universe now, the Past Hypothesis is the *only* macrostate the universe could have begun in that would be consistent with the entropy slope of the universe satisfying the restrictions mentioned above—being an ‘appropriate’ slope. So, perhaps the Past Hypothesis is a minimal explanatory posit. The current macrostate of the universe, however, now does *not* feature as an explanatory posit in Albert’s account. Nor does it seem it should—the current macrostate of the universe is ever-changing and contains large amounts of information irrelevant to the explanandum.

One might argue that, even if the Past Hypothesis is not needed to derive the *temporal asymmetry* of the relevant phenomena, it is required to derive the fact that these temporally asymmetric phenomena *exist*—such as the existence of branch systems at all. However, more general posits, of the kind Reichenbach employs, are enough to derive the existence of temporally asymmetric phenomena. In the case of the entropy rise of isolated systems, for example, the Reichenbachian account allows us to derive the fact that there *are* local isolated systems with entropy gradients parallel to that of the main system—not merely the claim that, if there were branch systems, their temporal direction would be parallel to that of the main system. Such a derivation is possible because the Reichenbachian account includes posits about there being local isolated subsystems (Posits 2) and their having entropy gradients (Posits 4). Even if one attempts to derive the *existence* of particular temporally asymmetric phenomena, this doesn’t license including the Past Hypothesis as an explanatory posit. The minimal explanatory posits are structural information that abstracts away from the particular details of the macrostate—such as Posits 2 and 4. Moreover, the *temporally asymmetric* minimal explanatory posit remains the entropy slope—not any particular state.

Albert could give up the claim that the Past Hypothesis refers to a particular state and take it to instead only claim the universe started out in *a* suitable low-entropy state. But the Past Hypothesis being true, and no analogous Future Hypothesis being true, is

logically equivalent to the entropy upgrade posited by Reichenbach—Albert would still not be giving a scientific explanation of the global entropy slope. Moreover, what is distinctive about Albert’s approach would be lost. Albert would also need some other account of how we reason to the past (Section 5.1).

To summarise so far; if we are to explain the general entropy rise of isolated subsystems, we need a Reichenbachian account. This account uses a general entropy posit and temporally neutral posits (as Reichenbach does) but combined with a probability postulate of the kind Albert employs (see Table 1). One might think that a posit about a particular state is required to derive other temporal asymmetries—such as the record asymmetry. If so, perhaps Albert’s account is to be preferred because it can derive a range of temporally asymmetric phenomena using the same resources. The work of the second half of this paper will be to argue that no particular state is required to explain the record asymmetry either.

<i>Reichenbach</i>	<i>Albert</i>	<i>Reichenbachian</i>
Lattice of mixtures (Posit 3)	Statistical Postulate	Probability Postulate (Posit 3*)
Entropic upgrade (Posit 1)	Past Hypothesis	Entropic upgrade (Posit 1)
Time-symmetric laws	Time-symmetric laws	Time-symmetric laws
Branch systems (Posits 2, 4)		Branch systems (Posits 2, 4)

Table 1: Posits used in the Reichenbachian account of the entropy rise of isolated subsystems (and other temporally asymmetric phenomena) in comparison with those used by Reichenbach and Albert. The Reichenbachian account retains most of Reichenbach’s posits, but includes a probability postulate of the kind Albert employs.

4. The Record Asymmetry

We’ve seen how a Reichenbachian approach can explain the entropy rise of isolated systems. I’ll now consider how Reichenbach explains the record asymmetry—an asymmetry that is interesting in itself and plausibly linked to our idea of the direction of time and to the temporal asymmetry of causation (see Section 1).

The record asymmetry is, roughly, the fact that there are local states now that serve as reliable indicators of past states, but not, in the same way, of future states. To explain the record asymmetry, we need to identify more precisely what records are and explain why there are records of the past and not the future. Identifying what records are requires identifying what ‘inferential mechanism’ is used when reasoning using records—what ‘in the same way’ amounts to in the rough formulation. The inferential mechanism need not be something we apply consciously, but instead indicates what worldly structure *underlies* our use of records—what has to be in place for a given local state to *count* as a record. What records are should be specified in temporally neutral terms so that their temporal asymmetry is traced back to worldly features and not to how records are defined. Keep in mind, I’m defining records as reliable indicators. We may attempt to infer ways that look like reasoning using records, but that are not reliable. These won’t be cases of reasoning using records. In this section, I consider Reichenbach’s account, before comparing it to Albert’s and presenting a Reichenbachian alternative (Section 5).

Reichenbach (pp. 108–125) identifies records as states of quasi-isolated entropy-increasing local subsystems that are improbable given the local system’s dynamics when isolated, but highly probable given the state of the main system at the point of interaction. The inferential mechanism involved is reasoning to states of systems at other times that render their otherwise improbable states now probable. Consider a box of gas currently in a mid-entropy state: most of the gas particles are on the left-hand side. Assume we know that this box is in the middle of a period of isolation. The probability of the box spontaneously evolving into this state, given its dynamics when isolated, is very small. So, when we observe the box in this state, we reason that it reached this state by being in an even lower entropy state in the past when it was isolated, after which it evolved to its present mid-entropy state. The state of the box now is a record of its own past state, as well as the state of the main system when the box became isolated—the point of ‘interaction’.

We have records of the past and not the future because, in the overwhelming majority of cases, recorded states that render the otherwise improbable record state probable lie in the past of that state. Future states don’t render the record state probable. So, there are records of the past and not the future. The record asymmetry follows from the fact that, in the overwhelming majority of cases, branch systems that have entropy gradients are on

entropy gradients oriented towards the future that take them from improbable states (given the system's dynamics when isolated) to more probable states (given the system's dynamics when isolated). So, we can reliably infer to the past in a way we can't reliably infer to the future. While we could attempt to reason in a 'record-like' fashion in either temporal direction, only towards the past is the method reliable, implying there are only records of the past.

With some modifications, Reichenbach takes his explanation to apply at the level of 'macrostatistics'—where probabilities derive from the system's macrodynamics rather than their microdynamics (pp. 145–54).²⁶ For simplicity, I will keep to the case of microstatistics. 'Entropy' will always refer to Boltzmannian entropy at the level of microstatistics.

Before considering competitors to Reichenbach's account, let me clear up some confusion regarding its scope. It's tempting to think that, on Reichenbach's account, states we would identify as records must be entropy increasing and must be lower in entropy than relevant alternatives. Neither assumption is correct.

Consider an objection from Sklar (1993, p. 398), based on Earman (1974)—see also Horwich (1987, p. 83). Say cans of soup are usually arrayed neatly on shelves in a supermarket and are kept that way by efficient shelf-stockers. Cans of soup in disarray on the floor will be a record of a past event—such as a shopper colliding with the cans. But the cans of soup are not changing significantly in entropy. Moreover, the scattered cans of soup are more *disordered* than the relevant alternative—being stacked neatly on the shelf. So, it seems the cans *cannot* be a record, on Reichenbach's account.

Sklar's objection is based on a misinterpretation of Reichenbach's account. On Reichenbach's account, states we would *typically* identify as records may be higher in entropy than relevant alternatives. But a record, properly speaking, is a state of the *entire recording system*. What Reichenbach's account requires is that, on the relevant time scales,

²⁶ Reichenbach's extension to the case of macrostatistics will only work if there is a consilience between macro and microstatistics. In the Reichenbachian approach I defend (Section 3.3), the consilience is achieved by deriving the macrodynamics of a system from its micro-dynamics and a probability postulate. One might argue that Reichenbach's analysis of records will only work at the level of macrostatistics. Reichenbach himself reserves the term 'record' for states at this level. I'll say something about what justifies my analysis of records in Section 5.2.

the state of the *entire recording system* (the record) must be improbable, given the system's dynamics when isolated—not that the *subpart* we would typically identify as a record must be improbable. In a supermarket with efficient shelf stockers, states of the supermarket where the cans are highly ordered are highly *probable* (given the supermarket's dynamics when isolated). So, states of the *supermarket* where the cans are disordered are unlikely states—so, they *can* be records. The mistake was to look to the cans alone to determine which states were probable, rather than looking to the states of the entire isolated system (cans, shelves and stockers). For the same reason, it won't matter if the cans themselves aren't increasing significantly in entropy—they are part of a larger isolated system whose entropy is significantly increasing.

Here is a second concern. Horwich (1987, p. 83) and Albert (2000, p. 124) argue that disordered states can be records of an *absence* of interactions—such as when your apartment remaining messy is a record that no one has broken in and tidied it up. On Reichenbach's account, it seems probable disordered states cannot be records. In response, it is no great revision to Reichenbach's account to allow that the absence of records may indicate the absence of interactions. Highly probable states of isolated systems can indicate the absence of interactions, when such interactions would be expected to leave records. These might not be records in the strict sense, but they are also temporally asymmetric—the absence of records of the future cannot indicate an absence of future interactions.

5. How to Explain the Record Asymmetry

I'll now argue that Reichenbach's account, supplemented with the probability postulate introduced above (Section 3), does better in identifying what records are and why they are temporally asymmetric, compared to its most prominent Boltzmannian rival—Albert's account (2000, 2015).

5.1. Albert's Account

Albert intends to explain the record asymmetry (2000, Ch. 6; 2015, Ch. 2). Yet it is not entirely clear what Albert takes records to be. Recall, explaining the record asymmetry requires identifying more precisely what records are (in temporally neutral terms) and explaining why they are temporally asymmetric. I'll focus on the most plausible interpretation of Albert and consider alternatives below.

Albert ends his 2000 account (pp. 122–3) with a precise statement of an epistemic asymmetry that holds at our world. In giving this statement, he introduces a method of reasoning. This method is *not* designed to capture our knowledge of the past. Instead, the method is designed to highlight how inferior our knowledge about the past would be *if* it relied on this method—and to argue that the Past Hypothesis is what allows us to know vastly more about the past (but not the future) than is available on this method. I’ll call the method Albert introduces the ‘contrast method’—it is meant to contrast unfavourably with the inferential mechanism that does capture our knowledge of the past. Readers may be more familiar with Albert’s terminology of prediction and retrodiction. Prediction/retrodiction involves conditionalising on any known macrostates in the *present*, applying the Statistical Postulate (Section 3.2), and using the dynamical laws to reason about states in the future/past.²⁷ While prediction/retrodiction doesn’t feature in Albert’s final statement of the record asymmetry, it serves a similar function—Albert argues we know vastly more about the past than is knowable through retrodiction, and it is the Past Hypothesis that allows us to know this additional information. The reason that retrodiction/prediction doesn’t feature in Albert’s final statement of the record asymmetry is that we also have more knowledge of the *future* than is available through prediction (2000, pp. 119–22)—so the prediction/retrodiction method does not allow the formulation of a sharp epistemic asymmetry. Albert’s final statement of the record asymmetry is meant to remedy this defect.

Take the ‘contrast method’ (my term) to be conditionalising on any known macrostates in the *present and in one temporal direction*, applying the Statistical Postulate (Section 3.2), and using the dynamical laws to reason about states in the other temporal direction. According to Albert, the contrast method allows us to derive pretty much everything we take ourselves to know about the future. For example, we can reason that the entropy of isolated systems at non-maximal entropy will be higher in the future. But, he argues, the contrast method allows us to derive almost nothing we take ourselves to know about the past. When we reason towards the past, the contrast method would lead us to infer, incorrectly, that the entropy of isolated systems at non-maximal entropy was *higher* in the past.

²⁷ Albert sometimes speaks of conditionalising on ‘directly surveyable’ states in the present, which includes macrostates and any microstates we have direct epistemic access to now (2000, p. 114). This latter detail won’t matter to my arguments and doesn’t feature in Albert’s explicit statements of the record asymmetry.

Albert argues that the Past Hypothesis is what allows us to have more accurate and precise knowledge of the past than is available using the contrast method. Say we reason towards the past using the contrast method, *but also conditionalise on the Past Hypothesis*. (In Albert’s 2015 terminology, this amounts to using statistical mechanical probabilities generated by the ‘Mentaculus’ and conditionalising on any known present (or future) macrostates.) When we reason towards the past by conditionalising on the Past Hypothesis, we reason to a state that is between two known states—the known Past Hypothesis and any known states in the present. Albert calls inferences of this form ‘measurement’ (2000, pp. 117–8; 2015 pp. 36–8). Measurement involves inferring *from* the state of a system at two times to the state of a system (itself or another) at a time in between. The recorded state lies temporally between the ‘ready’ state of the recording device and its ‘record-bearing’ state (2000, p. 117). I will call the latter the ‘record’. Albert argues that measurement is typically much more *informative* than reasoning from the state of a system at a single time. The notion of ‘informativeness’ that Albert presumably has in mind, and one I will employ elsewhere, is that a restriction is more informative than another when it restricts the possible phase space of a given system at a given time more than another (using a flat measure), given that the possibilities included or excluded can be specified in relatively simple macroscopic language.²⁸ Since the Past Hypothesis specifies the *particular* macrostate the universe begun in, and is known, it can function as a ready state when reasoning towards the past. It is the ‘mother of all ready conditions’ (2000, p. 118) because it is prior in time to all other past states we might infer to. But the Past Hypothesis does not allow us to ‘measure’ the future, since, when inferring towards the future and conditionalising on the Past Hypothesis, we are not inferring to an unknown state at a time between two known states. So, conditionalising on the Past Hypothesis does not provide a more informative way of reasoning about the future, compared to the contrast method.

Ultimately, we have records of the past and not the future because conditionalising on the Past Hypothesis is more informative when reasoning towards the past, but not the future, in comparison with the contrast method—in virtue of the fact that conditionalising on the Past Hypothesis when reasoning towards the past (but not the

²⁸ For the need to only consider possibilities specifiable in relatively simple macroscopic language, see Albert (2015, pp. 40–1).

future) constitutes a case of measurement. On Albert's account, records are local states that, using the contrast method and conditionalising on the Past Hypothesis, provide more precise or accurate information about the recorded states than could be had using the contrast method alone.²⁹

Note that, at this point, Albert doesn't give an account of what features of the Past Hypothesis allow it to play its role, beyond the fact that it is a known particular state. The argument regarding measurement, for example, does not make use of the fact that the Past Hypothesis is a low-entropy condition. While its entropic features may explain how we come to know the Past Hypothesis, the argument that we can have records of the past and not the future relies merely on the fact that conditionalising on the Past Hypothesis allows us to measure the past and not the future—and the fact that we can measure the past and not the future relies merely on the Past Hypothesis being a known particular state. This is partly why, when discussing the possibility of the world having a 'Future Hypothesis', Albert doesn't think this would be 'a hypothesis to the effect that the far-future state is a low entropy one, but that it is characterized by some simple macroscopic organization' (2000, p. 119). A Future Hypothesis would be a known particular state—not a low-entropy state. The only constraint that might be suggested is that the Past Hypothesis cannot be something we can reason to using the contrast method—or else conditionalising again on the Past Hypothesis would not be more informative than the contrast method. But even this feature is not discussed in Albert's argument.

Later I'll suggest that the entropic features of the Past Hypothesis can be used to provide an account of how we reason to it and to ultimately underwrite the record asymmetry (Section 5.2). Indeed, the Reichenbachian account I defend could be reached by generalising Albert's account and considering what entropic features of the Past Hypothesis are relevant for explaining the record asymmetry. But this would amount to abandoning Albert's argument based on measurement, which relies only on the Past

²⁹ Let me briefly note and dismiss one alternative interpretation. One might take Albert's account of records to simply be his account of measurement—there are records whenever states of systems at two times are reliable indicators of states at times in between. The problem is this account doesn't rule out there being records of the future. For example, the current non-equilibrium state of an isolated system and its future state at equilibrium do reliably indicate its near future state. While, in this case, the system's present state is *by itself* a reliable indicator of its future state, this point only makes the problem clearer: the account needs to say something about how the inferential mechanism behind records compares to other inferential mechanisms—not merely that records involve reasoning from states at two times.

Hypothesis being a known particular state. For now, I want to focus on the measurement-based argument that Albert actually gives for why we have records of the past and not the future.

I have two initial concerns concerning the argument that Albert gives. Firstly, Albert's argument risks triviality. When we reason about the future using the contrast method, we *inevitably* conditionalise on the Past Hypothesis, as long as the Past Hypothesis is a known past state. Albert accepts that the Past Hypothesis will have to be known, at least implicitly, if we are to reason using records (2000, p. 119; 2015, pp. 38–9). So, of course, conditionalising *again* on the Past Hypothesis makes no difference to our reasoning about the future—and there can't be records of the future. When reasoning towards the past, conversely, the Past Hypothesis *cannot* be conditionalised on when using the contrast method, simply because it is about the past. It is then no surprise that conditionalising on new information, the Past Hypothesis, leads to more accurate and precise inferences (provided, as noted above, the Past Hypothesis cannot be reasoned to using the contrast method). The reason there *can* be records of the past and that there *can't* be records of the future comes down to the fact that the Past Hypothesis is a known past state, known by some means other than the contrast method—and not to any of its (other) physical features. This is a worry. For, as hinted at above, it seems like the entropic features of the Past Hypothesis should be relevant for how it can function as a vastly informative ready condition. For this reason, we also have no account of why there isn't an analogous Future Hypothesis—a known future state—that *would* give us records of the future.

There might be other aims that Albert's account aims to achieve. For example, one might want to show that the analysis of records is *justified* because it delivers all the records we take there to be.³⁰ To clarify, the triviality concern arises only regarding Albert's aim to explain the *temporal asymmetry* of records—that is, why there are local states now that serve as reliable indicators of past states, but not, in the same way, of future states. It is not addressed to other conditions records may be required to satisfy.

³⁰ I discuss some such conditions at the end of Section 5.2. An anonymous reviewer suggested that perhaps one needs to show that conditionalising on the Past Hypothesis and the universe's present macrostate will deliver the entire past macrohistory of the world with high probability. However, this is not an aim that either Albert or I subscribe to. Albert (2015, pp. 46–8) accepts the possibility of large scale macroevents in the past that don't leave *any* records in the present and that would be assigned a very low probability by conditionalising on the Past Hypothesis and the universe's present macrostate.

Secondly, and relatedly, Albert's explanation risks incompleteness. While the explanation rests on the Past Hypothesis being a *known* past state, it doesn't provide an account of *how* we come to know the Past Hypothesis. As Albert emphasises, we certainly don't reason to the Past Hypothesis using records, since the Past Hypothesis is not prior in time to the Past Hypothesis (2000, p. 118; 2015, p. 38). Nor do we reason to it using the contrast method. So, there must necessarily be accurate methods of reasoning towards the past that aren't cases of records—and that, presumably, don't work towards the future (since the contrast method is supposed to deliver everything we take ourselves to know of the future). While Albert offers an independent account of how we come to know the Past Hypothesis (see Section 5.2), that account is not part of his explanation of the record asymmetry and does not appeal to the entropic features of the Past Hypothesis. It is not provided as part of a general account of how we come to know the past. It seems Albert's account does not give us the most general account available of why we have reliable modes of reasoning towards the past that aren't available towards the future.³¹

One might attempt to avoid the triviality concern (but not the incompleteness concern), by offering a different interpretation of Albert's account. Recall, prediction/retrodiction involves using the Statistical Postulate and the dynamical laws to reason about other states (as before), but also conditionalising on known directly macrostates at a *single* time—rather than directly known states at a single time *and in one temporal direction*. Perhaps reasoning using records involves using prediction/retrodiction but also conditionalising on the Past Hypothesis, in a way that provides more information than using prediction/retrodiction alone. This account avoids the triviality concern, since, when reasoning about the future, we don't inevitably conditionalise on the Past Hypothesis. But, as I mentioned above, the problem is that such an account *doesn't* imply a record asymmetry—as Albert is aware (2000, pp. 94–5). When we reason using prediction, conditionalising on the Past Hypothesis typically makes a great difference to our reasoning about the future. Say you observe a lion's fresh footprint. You can reason that the lion was here in the recent past and *will be somewhere nearby in the near future*. According to Albert, the latter inference relies on the footprint being a *record* of the lion's

³¹ For related concerns (that we have ways of reasoning towards the past more generally that also apply to the Past Hypothesis), see Frisch (2007, pp. 372–3). Incompleteness is also a concern if we accept Albert's definition of reasoning using records as *all* ways of reasoning that aren't cases of 'prediction' or 'retrodiction' (2000, p. 116).

location in the past and so can only be made by conditionalising on the Past Hypothesis. So, we do use the Past Hypothesis to reason more accurately about the future, compared to prediction/retrodiction, and there are records of the future.³²

Keeping to my original interpretation, one might attempt to respond to half of the triviality concern by arguing that, even if it is trivial that there *can* be records of the past, Albert's account of measurement nevertheless explains why the Past Hypothesis can serve as an *informative* restriction towards the past. Recall, records must be more informative than the contrast method. The Past Hypothesis is informative because it is a (known) particular past state that cannot be reasoned to using the contrast method. So, the account gives a substantial explanation of why there actually *are* records of the past—even if it remains trivial that there *can* be such records. However, this response leaves untouched the second half of the triviality concern. There is no explanation for why there aren't *future* particular states that would be more informative than the contrast method regarding the future. Indeed, in general, there are such states. Knowing the state of the universe 10 minutes from now will typically be much more informative for reasoning about the future, compared to the contrast method. Knowledge of near future states will tend to be very informative, for example, in cases where we don't know the full macrostate the universe at a single time. The problem is, we don't think we could have knowledge of such informative future states. Once again, knowability is a key feature. But it is not part of Albert's explanation of the record asymmetry.

Here is a diagnosis of what has gone wrong. When Albert explains why the Past Hypothesis is informative when reasoning towards the past, his focus is on the fact that the Past Hypothesis specifies the *particular* macrostate that the universe began in (2000, p. 96; 2015, p. 5). Measurement, recall, is about reasoning from (particular) known states at two times to a state at a time in between. But the feature the account turns out to rely on

³² A final alternative interpretation. Perhaps reasoning using records is to be compared to a method where we conditionalise on *known* states in the *present*—whether known directly or by other means. In the case of the lion, we conditionalise on the lion's location now. While this interpretation rules out their being records of the future in these cases, something close to triviality remains. In the case of the lion, the reason we know about the lion's present location *is*, Albert thinks, because we conditionalise on the Past Hypothesis. Given how much of our knowledge of the present is based on the Past Hypothesis, we still effectively conditionalise on the Past Hypothesis in most of our reasoning towards the future—so, no wonder conditionalising again on the Past Hypothesis makes no difference. The real puzzle is why conditionalising on the Past Hypothesis still makes a difference in our reasoning towards the past, even once we have detailed knowledge of the present. But Albert's account doesn't address this puzzle or why there isn't an analogous Future Hypothesis that makes a difference to our reasoning about the future. So, the account doesn't address the more general concerns I raise below.

is that the Past Hypothesis is known by means other than records or the contrast method. The Past Hypothesis being a particular state contributes nothing to explaining this feature. The universe will also be in a particular state 10 minutes in the future. Yet this doesn't imply that such a state can be known by means other than records or the contrast method. Nor does it imply that we can have records of the future.

The concern underlying all this is that Albert's account is not sufficiently general. It does not specify what is about the Past Hypothesis that makes it knowable and so why there isn't an analogous Future Hypotheses that would be similarly knowable. Nor does it explain how the Past Hypothesis is informative in ways that future states are not. Altogether, while Albert correctly identifies an epistemic asymmetry that holds, his account fails to explain *why* this asymmetry holds. So, it doesn't adequately explain the record asymmetry.

5.2 A Reichenbachian Account

I'll now present a *Reichenbachian* account of the record asymmetry, which combines the Reichenbachian account of temporal asymmetries (Section 3) with Reichenbach's approach to records (Section 4). I'll argue that this Reichenbachian account gets the right level of grain for explaining *what records are* and *why there are records of the past and not the future*. It provides the generality missing from Albert's account. The Reichenbachian account succeeds because it can explain a) how we come to know the Past Hypothesis, and b) why there isn't a similarly knowable analogous Future Hypothesis that *would* be informative towards the future—that is, would restrict the phase space of the universe (given possibilities specified in relatively simple macroscopic language), more than the contrast method.

Regarding b), the Reichenbachian account explains why the Past Hypothesis is a very *broad* informative restriction, and why there are no similarly broad informative restrictions in the future. For any isolated system heading towards disequilibrium towards the past, provided our knowledge of the present and future is something less than knowledge of its complete *microstate* at a time, knowing its past macrostate will always be more informative about the near past than the contrast method—*no matter how much knowledge of present and future macrostates we have*. Not so towards the future. If an isolated system is heading towards equilibrium towards the future, and if we know the full

present and past macrostates, knowledge of the future macrostate will *not* typically be more informative about near future states than the contrast method.³³ Reichenbach applied probability-based reasoning only to isolated subparts of the universe. I've argued (Section 3.2) that probabilities should also be applied to the universe as a whole. If so, we can say that the Past Hypothesis is a broad informative restriction because the universe is heading towards *improbable* states towards the past. But, because the universe is always heading towards probable states towards the future, there can be no Future Hypothesis as broadly informative as the Past Hypothesis.

However, while the probabilistic entropic gradient of the universe explains why there is no Future Hypothesis as *broad* as the Past Hypothesis, it doesn't yet explain why there is no Future Hypothesis at all. Recall (Section 5.1), if we know less than the full macrostate of the universe at any given time, information about future macrostates will typically be very informative. To explain why there is no analogous Future Hypothesis at all, we need to address a)—how we come to know the Past Hypothesis. If our knowledge of the future is limited to what we can infer to using the contrast method, informative restrictions on the future will not be knowable. However, if we have ways of reasoning to the past that are more informative than the contrast method, we will have a way of reasoning to past states, such as the Past Hypothesis, that are more informative than the contrast method. While it will be impossible to show that there are no informative ways of reasoning towards the future other than the contrast method, showing that there is a way of reasoning to the Past Hypothesis that is not available towards the future provides a more general and adequate account of the record asymmetry.

Albert does have things to say about how we come to know the Past Hypothesis—but he doesn't specify an inferential mechanism that operates towards the past and not the future and doesn't identify what entropic features of the Past Hypothesis are relevant to its epistemic accessibility. Albert suggests we reason to the Past Hypothesis via 'the sorts of things that typically justify our beliefs in *laws*' (2000, p. 119)—namely that they are empirically successful in helping us make good predictions and bring about coherence in our beliefs. The Past Hypothesis is particular helpful in achieving coherence in our beliefs about the veracity of our records and our beliefs about the fundamental

³³ Knowledge of future macrostates may still be more informative if the macrodynamics are stochastic or if we lack relevant knowledge of the system's past and present microstates.

dynamical and statistical laws—for, without the Past Hypothesis, we would be led by our beliefs in the time-symmetric fundamental dynamical laws and probability postulate to infer that the past was high entropy and our records were unreliable.

Albert also speaks of an implicit knowledge of the Past Hypothesis being hard-wired into us by evolution, experience and explicit study, as part of a package with fundamental dynamical and statistical laws (2015, p. 39). Indeed, Albert counts the Past Hypothesis as one of the fundamental laws (2015, p. 5). As noted above, Albert briefly discusses what it would be for the world to have a Future Hypothesis (as well as a Past Hypothesis), suggesting a Future Hypothesis wouldn't be a low-entropy condition but one characterised 'by some simple macroscopic organisation' (2000, p. 119). Because we reason to Hypotheses as we reason to (other) laws, being simple is part of how we reason to it. What I suspect Albert has in mind is a broadly subjective Bayesian account, where simplicity is built into the prior probability of a hypothesis. I also suspect he is influenced by his Humean 'Best Systems' account of laws (2015, Ch. 1), in which laws are axioms of *simple* and general systems that summarise non-modal events. If the Past Hypothesis is to be a fundamental law, it must be simple.

One might attempt to leverage Albert's brief remarks into a more general account of how we come to know past states by means other than (Albert's) records or the contrast method. That is, one might attempt to identify a general inferential mechanism (in the sense I began with) by which we reason to past states and not future states. (My thanks to an anonymous reviewer for pressing me on this point.) As mentioned above, Albert thinks we reason to the Past Hypothesis partly because it brings about coherence in our beliefs about the reliability of our records and the (other) laws, by blocking us inferring to a high-entropy past. For the Past Hypothesis to play this role, it must be a low-entropy state. However, there are problems with relying too heavily on this aspect of Albert's explanation. Firstly, it's not obvious that the requirement to avoid incoherence or the epistemic undermining of records secures positive belief in the Past Hypothesis—one could equally revise one's beliefs in the laws or statistical postulate, for example. Secondly, it's not obvious this argument buys Albert what he would want—a mode of reasoning that takes you to a single simple state (the Past Hypothesis). Instead, it looks like it might lead to many (mostly quite complicated) intermediate mid-entropy states placed in whatever point in time is required to underwrite the records we actually have at

a given time. One might try appealing to simplicity at this point to select a single simple state. But it's not obvious why simplicity should play a role at all. It seems that the temporal asymmetries we observe would still hold, even if the Past Hypothesis turned out to be a low-entropy *complicated* state. Thirdly, if we reason to the Past Hypothesis using the same method we use to reason to (other) laws, since those laws concern the future, the method can't be temporally asymmetric—it can't be a method that is reliable towards the past and not the future.

Ultimately, one could enrich Albert's discussion in order to identify a particular inferential mechanism by which we reason to past states like the Past Hypothesis (and not to future states), one that would rely on the entropic features of the Past Hypothesis.³⁴ These same entropic features could also explain in what way conditionalising on the Past Hypothesis is more informative than the contrast method—a point I consider below. But these amount to revisions of Albert's account, which relies much more closely on the fact that the Past Hypothesis is a simple particular state. The revised account would move much closer to the more general and entropy-based account given by Reichenbach. At the end of the day, it won't matter whether one considers the account I offer a variant of Albert's or Reichenbach's—it combines features of each to give a single explanation of how we reason to the Past Hypothesis as well as to past states more generally.

According to Reichenbach, when we reason using records, we reason to states of systems at other times that render their otherwise improbable states now probable. This method is unreliable towards the future because systems head towards equilibrium towards the future—there aren't future states that would render their otherwise improbable states now probable—so, there aren't records of the future. If we accept universe wide probabilities (Section 3.2), the Reichenbachian account explains how we can come to know the Past Hypothesis (or at least as much of the Past Hypothesis as we have epistemic access to). We reason to the past macrostate of the universe that renders its otherwise improbable low-entropy state now probable. So, we have a way of reasoning towards informative states towards the past that is not available towards the future, one underwritten by the probabilistic entropic gradient of the universe.

³⁴ I read Frisch (2007, pp. 372–3) as suggesting something along these lines.

This Reichenbachian account provides the generality missing from Albert's account. As noted, Albert emphasises the fact that the Past Hypothesis is a 'particular' state (2000, p. 96; 2015, p. 5). However, what records are, and what explains why we have records of the past and not the future, is *not* the fact that the Past Hypothesis is a particular state—but the fact that it is a knowable informative restriction.

Frisch (2007, pp. 375–7) criticism of Albert's account on this point does not apply to the Reichenbachian account I defend. As noted above, Albert's measurement-based argument requires the Past Hypothesis to be a known particular state. Frisch argues that merely knowing the past was low entropy will not be enough for Albert's argument to apply. Unless knowledge of the Past Hypothesis amounts to knowledge of the complete initial macrostate of the universe, Frisch argues, it can't function as the kind of ready state that would provide us precise knowledge of the past. I agree with Frisch that measurement requires knowledge of particular states. But the Reichenbachian account does not use Albert's account of measurement to explain how we have knowledge of the past. Instead, the method involved is reasoning to past states that render later states probable. The Reichenbachian account may imply we have less precise knowledge of the past than would be available under measurement (if we knew the complete initial macrostate of the universe). But the account is not committed to us having this kind of precise knowledge.

The Reichenbachian account takes the informativeness and knowability of the Past Hypothesis to be deeply tied. The most informative restrictions are those that go counter to the probabilistically expected behaviour of isolated system, given its dynamics when isolated. What allows such informative restrictions to be knowable is the fact that they can render otherwise improbable states now probable.

One might worry that the Reichenbachian account implies a vicious circularity in our reasoning. We reason that the universe is on a long entropic upgrade by reasoning to states of systems at other times that render their otherwise improbable states now probable. We apply this method only when reasoning towards the past and only towards the past is this method reliable. What *explains* why this method is reliable when applied to the past is the fact that the universe is on a long probabilistically appropriate entropic upgrade. I claim there is no vicious circularity here. One side speaks to how we reason

(applying an inferential method only towards the past), the other to the physical facts that externally justify such reasoning. This kind of circularity is to be expected when we use posits we reason *to* in order to explain *how* we reason. The aim of this kind of philosophical account is not to provide posits that are both epistemically and physically foundational. Instead, the aim is to specify how we can reason to fundamental physical posits (using methods we may take as epistemically basic) and also how those same posits can justify why our methods of reasoning are appropriate. We may still want to give further and perhaps internalist justifications for how we reason to these relatively basic epistemic and physically fundamental posits.³⁵ But it is still an important constraint on our physical theorising that this kind of benignly circular justification is available.³⁶

I'll end with some observations about the Reichenbachian account of records and responses to objections. Firstly, not all our reasoning about the past need take the form of Reichenbachian records. There may be methods that apply equally towards the past and future, such as when we reason about systems whose macroevolution is deterministic. We still might think of the present as 'recording' the past in these cases. I suggest this is because we generalise from real and apparent temporal asymmetries and take the temporal asymmetry of records to be more pervasive than it is. For discussion of such generalising, see Reichenbach (p. 151, 56) and Fernandes (2021).

Secondly, Reichenbachian records need not always be states of small isolated systems. They can include states of much larger systems, including the universe as a whole. As I argued above (Section 4), what we typically call a 'record', such as a footprint, may be part of a larger isolated record system—such as a windswept beach. I also suspect that when we reason about the past using Reichenbachian records, we often do so by making assumptions about what the rest of the universe is like. For example, when we reason that the footprint-shaped indentation is a record of someone having walked on the sand, we use assumptions about what objects or events are likely to have been around to produce such an indentation. If we reasoned absent assumptions about what the rest of the universe was like, we could not infer anything so precise as a person walking. This 'global' aspect of our reasoning using records is something Albert's account rightly draws

³⁵ One option is to adopt a broadly subjective Bayesian account. I suspect such an account should still grant less of a role to simplicity than Albert does.

³⁶ For examples of theories that don't satisfy this minimal requirement, see discussions of Boltzmann brains (Carroll 2017) and the bare theory in Quantum mechanics (Barrett 1996).

attention to—we reason about local aspects of the past using much broader assumptions about the past and present.³⁷

Does the need to conditionalise on further relevant information suggest that there is nothing unique about reasoning using records or that reasoning using records is simply induction—an inferential mechanisms that applies equally well to the future (Earman 1974)? In response, first note that the use of conditionalising is not unique to the Reichenbachian account. On Albert’s account one must also conditionalise on other parts of the universe’s ‘present physical situation’ (2000, p. 94) when reasoning from local records such as Napoleon’s left boot. Second, building in conditionalising does not rule out there being temporally asymmetric inferential mechanisms—and so sharp asymmetries in the physical world that license us treating our reasoning towards the past as underwritten in a different way from our reasoning towards the future, *even once we conditionalise on other relevant information*. Third, there being some inferential mechanisms that operate towards the past and future (such as induction) does not rule out there being others which are temporally asymmetric. One might remain skeptical there are any temporally asymmetric inferential mechanisms. The best response I know is to put forward plausible accounts of them.

Thirdly, while the Reichenbachian account focuses on our reasoning about the past, knowing the universe is on an entropy gradient can also assist our reasoning about the future—at least we can know that isolated systems are heading towards equilibrium. But note that this fact doesn’t undermine the claim that there are records of the past and not the future. Recall, reasoning using records was defined to be a *reliable* inferential mechanism. Reasoning using records involves reasoning to a state of the system at another time that renders its otherwise improbable state now (given the dynamics of the system when isolated) probable. Applying such a method towards the future would lead one to infer *incorrectly* to a low-entropy future state. So, the state of the system now is not a record of its future state. If we ask why the method does not work reliably towards the future, the answer is given by the probabilistic entropy gradient of the universe (and isolated subsystems).

³⁷ For a more detailed account of how assumptions about both the past and present feature in reasoning about the past, see Fernandes (Forth b.).

Fourthly, there are certainly aspects of our reasoning using records that one might want to hear more about. For example, one might want to know why records typically provide more, or more reliable, knowledge, compared to our knowledge of the future. While I won't attempt a full account here, I want to suggest that the Reichenbachian account provides a plausible starting point. Firstly, if our knowledge of the future is limited to what may be had by the contrast method, we will be inferring towards the future to states that are more probable (given the relevant system's dynamics when isolated). When inferring using records towards the past, we will be inferring to states that are less probable. Plausibly, knowing the system will be in a more probable state is typically less informative than knowing it will be in a less probable state. If so, records of the past will typically be more informative than contrast-based reasoning about the future. Secondly, the Reichenbachian account explains some of the ways in which records can be more reliable. Records can be insulated from the disruptive or complex dynamics governing the main system. Reichenbach speaks of records as 'frozen order' (p. 151)—they are unlikely states of systems that may change much more slowly than the systems whose states they record. In situations where the dynamics of the main system make past states difficult to infer, records that are insulated from these dynamics may still reliably indicate past states of that system.

Finally, there is a further debate to be had about whether 'records', in the Reichenbachian sense I've defined, are close enough to the ordinary use of the term 'records' to warrant the term. Here are two points of concern. As discussed above (Section 4) I haven't made use of ordered 'macroarrangements' (such as packs of cards with all red suits in the top half) when defining records. Instead, records are improbable macrostates given the system's dynamics when isolated. There may be records that aren't ordered macroarrangements—should these also count as records? One might also worry that what we might typically call a record (soups cans in disarray) are actually only a *subpart* of the Reichenbachian record (the state of the whole supermarket) (Section 4). In response to the second concern, one could relativize local records to systems of which they are subparts—and only call the subpart the record. But my preferred response to both concerns is to keep the account of records as general as possible and then allow for more particular debates about whether only a certain subset (or subparts) of these warrant the ordinary term. The Reichenbachian account would still offer necessary conditions on ordinary-language 'records' (or systems in which they're embedded), and

provide a precise account of what records are that can be used to explain other temporal asymmetric phenomena.

6. Conclusion

There is a promising Reichenbachian account of the entropic rise of isolated subsystems and the asymmetry of records. This account combines the structural features of Reichenbach's approach with a probabilistic postulate of the kind used by Albert. While I haven't attempted a full defence of the Reichenbachian account, I have shown that it has advantages over its most prominent Boltzmannian rival—Albert's account. The Reichenbachian account correctly identifies the relevant temporally asymmetric posit that accounts for the entropy rise of isolated systems and the record asymmetry—it is the probabilistically appropriate entropy upgrade of the universe. The Reichenbachian account also has the right level of generality for explaining why there are records of the past and not the future, precisely because it relates the record asymmetry to the entropic upgrade. While there is much work still to be done, the Reichenbachian account provides the right foundations for understanding two phenomena that are arguably central to making sense of the direction of time and of temporally asymmetric phenomena more generally.

Acknowledgements: My warm thanks to David Albert for several helpful discussions and comments on this paper, Michael Hicks, who undertook a two-person reading group with me on Reichenbach's *The Direction of Time* that prompted this work, Erik Curiel and Flavia Padovani for the opportunity to present at the *All Things Reichenbach Conference* at the Munich Center for Mathematical Philosophy and their work organising that conference and this volume, and audiences at Munich and Columbia University.

The author has no conflicts of interest to declare.

References

- Albert, David Z. 2000. *Time and Chance*. Cambridge, Mass.: Harvard University Press.
- . 2015 *After Physics*. Cambridge, Mass.: Harvard University Press.
- Barrett, Jeffrey. 1996. Empirical Adequacy and the Availability of Reliable Records in Quantum Mechanics. *Philosophy of Science* 63: 49–64.

- Callender, Craig. 2017. *What Makes Time Special?* New York: Oxford University Press.
- . 2021. Thermodynamic Asymmetry in Time. In Edward N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy* (Summer 2021 Edition), URL = <<https://plato.stanford.edu/archives/sum2021/entries/time-thermo/>>.
- Carroll, Sean. 2017. Why Boltzmann Brains Are Bad. <<https://arxiv.org/abs/1702.00850>> Accessed 20 July 2021.
- Davies, Paul. 1974. *The Physics of Time Asymmetry*. Berkeley: UC Press.
- Earman, John. 1974. An attempt to Add and Little Direction to ‘The Problem of the Direction of Time.’ *Philosophy of Science*. 41: 15–47.
- . 2006. The “Past Hypothesis”: Not Even False. *Studies in History and Philosophy of Modern Physics* 37: 399–430.
- Eberhardt, Frederick. and Glymour, Clark. 2011. Hans Reichenbach’s Probability Logic. In D. M. Gabbay, J. Woods and S. Hartmann (eds.), *Handbook of the History of Logic*, Vol. 10, Amsterdam: Elsevier. 357–389.
- Elga, Adam. 2007. Isolation and Folk Physics. In Huw Price and Richard Corry (eds.), *Causation, Physics, and the Constitution of Reality*, Oxford: Oxford University Press. 106–119.
- Fernandes, Alison. 2017. A Deliberative Approach to Causation. *Philosophy and Phenomenological Research* 95(3): 686–708.
- . 2021. Does the Temporal Asymmetry of Value Support a Tensed Metaphysics?. *Synthese*. 198(5): 3999–4016.
- . 2022. Caring for Our Future Selves. In Christoph Hoerl, Teresa McCormack and Alison Fernandes (eds.) *Temporal Asymmetries in Philosophy and Psychology*. Oxford: Oxford University Press.
- . Forth. a. Naturalism, Functionalism and Chance: Not a Best Fit for the Humean. In Michael Hicks, Siegrfried Jaag and Christian Loew (eds.), *Humean Laws for Human Agents*, Oxford: Oxford University Press.
- . Forth. b. Functionalist Accounts of Laws, Chances and Counterfactuals.
- Frisch, Mathias. 2007. Causation, Counterfactuals, and Entropy. In Huw Price and Richard Corry (eds.), *Causation, Physics, and the Constitution of Reality*, Oxford: Oxford University Press. 351–395.
- . 2014. *Causal Reasoning in Physics*. Cambridge: Cambridge University Press.
- Frigg, Roman. 2008. A Field Guide to Recent Work on the Foundations of Statistical Mechanics. In Dean Rickles (ed.) *The Ashgate Companion to Contemporary Philosophy of Physics*. London: Ashgate. 99–196.

- Glymour, Clark and Eberhardt, Frederick. 2016. Hans Reichenbach. *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/win2016/entries/reichenbach/>>.
- Grünbaum, Adolf. 1973. *Philosophical Problems of Space and Time*. Dordrecht, Holland: D. Reidel.
- Hájek, Alan. 2009. Fifteen Arguments Against Hypothetical Frequentism. *Erkenntnis* 70(2): 211–235.
- Hempel, Carl G., and Paul Oppenheim. 1948. Studies in the Logic of Explanation. *Philosophy of Science* 15: 135–75.
- Hitchcock, Christopher. 2018. Probabilistic Causation. *The Stanford Encyclopedia of Philosophy* (Fall 2018 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/fall2018/entries/causation-probabilistic/>>.
- Horwich, Paul. 1987. *Asymmetries in Time*. MIT Press: Cambridge, Mass.
- Lee, Ruth, Shardlow, Jack, Hoerl, Christoph O'Connor, Patrick A., Fernandes, Alison, McCormack, Teresa. Forth. Towards an account of intuitive time. *Cognitive Science*.
- Leeds, Stephen. 2003. Foundations of Statistical Mechanics—Two Approaches. *Philosophy of Science* 70(1): 126–144.
- Loewer, Barry. 2004. David Lewis's Humean Theory of Objective Chance. *Philosophy of Science* 71(5): 1115–1125.
- . 2007. Counterfactuals and the Second Law. In Huw Price and Richard Corry (eds.) *Causation, Physics, and the Constitution of Reality*. Oxford: Oxford University Press, 293–326.
- Maudlin, Tim. 2007a. *The Metaphysics within Physics*. Oxford: Oxford University Press.
- . 2007b. What could be objective about probabilities? *Studies in History and Philosophy of Modern Physics* 38: 275–291.
- North, Jill. 2008. Two Views on Time Reversal. *Philosophy of Science* 75(2): 201–223.
- Price, Huw. 1996. *Time's Arrow and Archimedes' Point*. Oxford: Oxford University Press.
- . 2002a. Boltzmann's Time Bomb. *British Journal for the Philosophy of Science*. 53(1): 83–119.
- . 2002b. Burbury's Last Case. In Craig Callender (ed.), *Time, Reality and Experience*. Cambridge: Cambridge University Press. 19–56.
- Reichenbach, Hans. 1949. *The Theory of Probability. An Inquiry into the Logical and Mathematical Foundations of the Calculus of Probability*, Transl. by E.H. Hutten and M. Reichenbach, Berkeley-Los Angeles: University of California Press.
- . 1956. *The Direction of Time*. Berkeley: University of California Press.

- Schrödinger, Erwin. 1950. Irreversibility. *Proceedings of the Royal Irish Academy: Section A: Mathematical and Physical Sciences*. 53: 189–195.
- Salmon, Wesley C. 1992. Scientific Explanation. In Merrilee H. Salmon, John Earman, Clark Glymour, James G. Lennox, Peter Machamer, J.E. McGuire, John D. Norton, Wesley C. Salmon and Kenneth F. Schaffner (eds.), *Introduction to the Philosophy of Science*. Indianapolis/Cambridge: Hackett. pp. 7–41.
- Sklar, Lawrence. 1993. *Physics and Chance*. Cambridge: Cambridge University Press.
- Smart, J. J. C. 1967. Time. In Paul Edwards (ed.), *Encyclopedia of Philosophy*, New York: Macmillan.
- Winsberg, Eric. 2004a. Can Conditioning on the Past Hypothesis Militate Against the Reversibility Objections? *Philosophy of Science* 71(4): 489–504.
- . 2004b. Laws and statistical mechanics. *Philosophy of Science* 71(5): 707–18.