Burbury's Last Case: The Mystery of the Entropic Arrow

Huw Price

Does not the theory of a general tendency of entropy to diminish [sic¹] take too much for granted? To a certain extent it is supported by experimental evidence. We must accept such evidence as far as it goes and no further. We have no right to supplement it by a large draft of the scientific imagination. (Burbury 1904, 49)

1. Introduction

Samuel Hawksley Burbury (1831–1911) was an English barrister and mathematician, who favoured the latter profession as loss of hearing increasingly curtailed the former. The Bar's loss was Science's gain, for Burbury played a significant and perhaps still under-rated part in discussions in the 1890s and 1900s about the nature and origins of the Second Law of Thermodynamics. One contemporary commentator, reviewing Burbury's *The Kinetic Theory of Gases* for *Science* in 1899, describes his role in these terms:

[I]n that very interesting discussion of the Kinetic Theory which was begun at the Oxford meeting of the British Association in 1894 and continued for months afterwards in *Nature*, Mr. Burbury took a conspicuous part, appearing as the expounder and defender of Boltzmann's H-theorem in answer to the question which so many [had] asked in secret, and which Mr Culverwell asked in print, '*What is the H-Theorem and what does it prove?*' Thanks to this discussion, and to the more recent publication of Boltzmann's *Vorlesungen über Gas-theorie*, and finally to this treatise by Burbury, the question is not so difficult to answer as it was a few years ago. (Hall 1899, 685)

In my view, however, it is at best half-right to describe Burbury as a defender of the H-Theorem. In some respects, he was the leading advocate for the prosecution. The crucial issue arising from Culverwell's (1890a, 1890b, 1894) enquiry was the source of the *time-asymmetry* of the H-Theorem, and while Burbury put his finger on the argument's time-asymmetric premise, he himself fingered it—in true forensic spirit—as an object of considerable suspicion. A decade later, as the quotation with which we began

¹By 'entropy' here, Burbury seems to mean the quantity H of H-Theorem fame, which diminishes as a gas approaches equilibrium.

indicates, he still wasn't convinced that we are *entitled* to assume it; and in their important survey article of 1912, Ehrenfest and Ehrenfest note his continuing dissent on the matter (1959, 42; 95, n. 168). Indeed, in reading Burbury's work from the late 1890s and 1900s, one gets the sense that the issue of the basis and justification for the Second Law remained *the* great intellectual puzzle of his later life.

Later in the paper I want to return to Burbury's contribution to the debate of the 1890s. I want to endorse his scepticism about our present entitlement to assume that entropy will continue to increase, and I want to show that in one important respect, his contribution to the debate has been systematically misunderstood. (In uncovering this misunderstanding, I shall argue, we undermine one widespread conception of what it would take to account for the Second Law.)

However, the main task of the paper is to call attention to two kinds of objection to some well-known strategies for explaining the temporal asymmetry of thermodynamics. Neither of these objections is particularly novel, but I think that despite the long history of the problem, neither has received the prominence it deserves. As a result, in my view, the strategies concerned continue to enjoy more popularity than they merit.

In a more constructive vein, I want to locate the strategies I am criticising within a kind of taxonomy of possible approaches to the problem of the thermodynamic asymmetry. In this way, I want to suggest that they are mistaken not only for the reasons identified by the two objections in question, but also for a more basic reason. In effect, they are trying to answer the wrong question. This may sound doubly critical, but the double negative yields a positive—given that the proposed strategies don't work, it turns out to be good news that we don't need them to work, in order to understand the thermodynamic asymmetry. In this respect, the present paper deals with a topic I have examined in more detail in other recent work (Price 2002): the issue as to the precise nature of the time-asymmetric explanandum presented to us by thermodynamic phenomenon.

1.1 Origins of the problem

The time-asymmetry of thermodynamics is associated with the Second Law. According to this principle, non-equilibrium systems progress monotonically towards equilibrium. Temperature differences decrease, energy concentrations dissipate, and entropy increases monotonically until equilibrium is reached.

In the latter half of the nineteenth century, having recognised and described this tendency as a phenomenological principle, physics sought to explain it in statistical mechanical terms. The Second Law turned out to be different from more familiar laws of physics in at least two ways. First, it was probabilistic rather than strictly universal in nature—exceptions were possible, though very unlikely. Secondly, and more interestingly for our present purposes, it was time-asymmetric—the phenomena described by the Second Law showed a clear temporal preference. Late in the nineteenth century, and especially in the debate mentioned above, physics began to see that the latter fact is rather puzzling, in the light of the apparent time-symmetry of the laws of mechanics. How could symmetric underlying laws give rise to such a strikingly timeasymmetric range of phenomena as those described by the Second Law?

More than a century later, there is surprisingly little consensus as to how this question should be answered. Late in twentieth century, indeed, a leading authority on the conceptual foundations of statistical mechanics could still refer to the puzzle of the time-asymmetry of the Second Law as the elusive object of desire. (Sklar 1995) In my view, the fact that a solution has remained elusive rests in part on some confusion about what we are actually looking for. It is unusually difficult to be clear about what precisely needs to be explained about the time-asymmetry of the Second Law. There are several competing conceptions of what the problem is, with the result that proponents of rival approaches tend unwittingly to be talking at cross-purposes.

In order to try to clarify matters, I proceed as follows. I focus first on the nature of the time-asymmetric phenomena that are the source of the problem. Here, I think, a few simple remarks do a great deal to help us keep the true object of desire in view, and to avoid issues which are not directly relevant to the puzzle of time-asymmetry. I then distinguish two major competing approaches to the explanation of these timeasymmetric phenomena—two approaches which differ markedly in their conception of what needs to be done to solve the puzzle.

Despite their differences, the two approaches do agree about one part of the puzzle—they both hold that an important contributing factor to the observed thermodynamic time-asymmetry is that entropy was much lower in the past than it is now. However, whereas one approach argues that this asymmetric boundary condition is the *sole* time-asymmetric source of the observed asymmetry of thermodynamic phenomena, the other approach is committed to the existence of a *second* time-asymmetry—a time-asymmetric lawlike generalisation. So far as I know, the distinction between these two approaches has not been drawn explicitly by other writers. Without it, it is not easy to appreciate the possibility that many familiar attempts to explain the time-asymmetry of thermodynamics might be not *mistaken* so much as *misconceived*—addressed to the wrong problem, in looking for time-asymmetry in the wrong place.

In my view, the 'two-asymmetry' approaches are misdirected in just this way. In looking for a basis for a time-asymmetric generalisation, they are looking for something the explanation of the phenomena in question neither needs nor wants. However, my main aim is not to argue that these approaches address the problem of the thermodynamic asymmetry in the wrong terms, but to show that they don't succeed in addressing it in their own terms—they fail by their own lights. Specifically, I shall raise two objections to the two-asymmetry approach. Neither objection applies to all versions of this approach. However, I shall note two characteristics, one or other of which seems a feature of any version of the two-asymmetry view. Each characteristic leaves the view in question open to one of my two objections. My claim is thus that any version of the twoasymmetry approach is subject to at least one of the two objections. Some, as we shall see, are subject to both.

As I said, however, I take these negative conclusions to be good news for the project of trying to understand the thermodynamic asymmetry itself. For they bolster the case—already strong, in my view, on simplicity grounds—for taking the oneasymmetry view to provide the better conception of what actually needs to be explained.

2. Three Preliminary Clarifications

I begin with three preliminary remarks about the nature of thermodynamic time asymmetry. The purpose of these remarks is to 'frame' the relevant discussion, as I see it, and explicitly to set aside some issues I take not to be of immediate relevance. I am aware, of course, that these judgements of relevance themselves are not uncontroversial, and that some readers will feel that I am ignoring the interesting part of the subject. But in territory thick with the criss-crossed tracks of previous expeditions, it is useful to all parties to mark out one's own course as clearly as possible. What follows are three assumptions that I shall take for granted in the remainder of the paper.

2.1 Numerical asymmetry not practical reversibility

It is common to characterise the time-asymmetry of thermodynamic phenomena in terms of the irreversibility of the processes by which matter tends to thermodynamic equilibrium. However, the term 'irreversible' is ambiguous. In particular, some writers interpret it in a very practical sense. The time-asymmetry of thermodynamics is hence thought to be tied to the practical difficulty of 'reversing the motions' in real systems. (Ridderbos and Redhead 1998 appear to take this view, for example, focussing on the contrast between most ordinary systems and those involved in the so-called spin-echo experiments, which do permit reversal of motions, at least to an unusual degree.)

In my view, this focus on practical reversibility mislocates the important timeasymmetry of thermodynamic phenomena. Consider a parity analogy. In a world containing handed structures of a certain kind, we may distinguish two sorts of question: (i) Can a left-handed example of such a structure be 'reversed' into a right-handed version, and vice versa? (ii) Is there a numerical imbalance in nature between the number of left-handed and right-handed examples? These two questions are logically independent—one can easily imagine worlds with any of the four possible combinations of answers. Similarly in the temporal case, I think. The issue of the practical reversibility of a time-oriented phenomenon is logically independent of that of the numerical imbalance in nature between examples of the phenomenon in question with the two possible temporal orientations. Clearly, our world shows a vast numerical imbalance of this kind, for the kind of phenomena described by the Second Law. It is this numerical imbalance that is the primary puzzle, in my view, not the issue of the practical reversibility of individual systems.

2.2 A monotonic entropy gradient, not an increase or decrease

When we say that entropy always *increases*, or that entropy change is always *non-negative*, we presuppose a convention as to which is to count as the 'positive' direction on the temporal axis. If we reverse this convention, an increase is redescribed as a decrease, a non-negative change as a non-positive change. In one rather uninteresting sense, this means that the Second Law is itself a conventional matter—in its standard formulations, it depends on a convention concerning the labelling of the temporal axis.

Some people may feel that they can make sense of the view that the choice of labelling is not merely conventional—that one or other labelling is objectively correct, and that time itself is objectively 'directed' in this way. The point I wish to emphasise is that these are separate issues from that of the thermodynamic asymmetry. In particular, the objectivity of the thermodynamic asymmetry does not depend on the view that time itself is objectively directed (whatever that might mean).

Indeed, the thermodynamic asymmetry is easily characterised in a way which avoids the conventional choice of labelling altogether. We simply need to say that the entropy gradients of non-equilibrium systems are all aligned in the same direction, leaving it unspecified (and a conventional matter) whether they all increase or all decrease. (By way of analogy, we might describe a universe containing nothing other than a single hand as objectively 'handed', while regarding it as a conventional matter whether it contains a left hand or a right hand.)

2.3 'Entropy' is inessential

One source of complexity in discussions of the Second Law is the existence a variety of competing definitions of entropy. In view of this complexity, it is helpful to keep in mind that if necessary, the puzzling time-asymmetry of thermodynamic phenomena can be characterised without using the notion of entropy—at least if we are prepared to tolerate some loss of generality. We can describe the puzzle by being more specific—by listing some of the actual kinds of physical phenomena that exhibit a temporal preference; that occur in nature with one temporal orientation, but not the other. Warm objects cool down rather than heat up in a tub of ice, pressurised gas flows out from but not into a bottle, so on. (In all these cases the description presupposes the ordinary temporal labelling, in the way noted in Section 2.1.)

The notion of entropy may turn out to provide a useful way of generalising over this class of phenomena, but doing without it wouldn't deprive us of a way of talking about the temporal bias displayed by real systems—which, after all, is the real source of the puzzle. I shall continue to use the term in this paper, but take it to be merely a placeholder for more specific descriptions of the relevant properties of time-asymmetric thermodynamic systems.

3. Two Models of the Origin of the Thermodynamic Asymmetry

In this section I distinguish two different explanatory models for the thermodynamic asymmetry. In my view, insufficient attention to this distinction does a great deal to explain the striking lack of consensus about this topic—in particular, it explains the major respect in which various participants in the debate have been talking at cross purposes.

3.1 The two-asymmetry view

As noted above, the nineteenth century attempt to reduce thermodynamics to statistical mechanics led to two realisations about the Second Law: first, that it is probabilistic rather than strictly universal in nature; and second, that it is time-asymmetric. As a result, the Second Law is commonly seen as a time-asymmetric probabilistic generalisation, with a lawlike or quasi-lawlike character—a general constraint on the

behaviour of matter, preventing or at least discouraging (in a probabilistic sense) entropy from decreasing, in the ordinary time sense.

What does 'lawlike' mean in this context? This question deserves more attention than I can give it here, but for present purposes the crucial feature seems to be *projectibility*. At the very least, for two-asymmetry views, the Second Law retains the status of a time-asymmetric generalisation *on which we are entitled to rely*, in forming expectations about the future behaviour of matter. (As we shall see, it cannot have this status for one-asymmetry views.)

The nomological character of two-asymmetry approaches sometimes also shows up as the view that matter has a certain time-asymmetric *property*, or *disposition*—that of being 'thermodynamic' rather than 'anti-thermodynamic'. Thus Richard Feynman, in comments attributed to 'Mr X' in discussion reproduced in (Gold 1963), refers to 'the assumption that matter is thermodynamically "one-sided", in the ordinary sense that it damps when you try to shake it.' (Gold 1963, 17) Here it is the *property* which is thought of as projectible—the time-asymmetric disposition to absorb heat, disperse energy, and the like.

According to this picture, then, a large part of the task of explaining the observed thermodynamic asymmetry is a matter of finding a basis for a time-asymmetric nomological generalisation or disposition—finding something that 'makes' entropy increase towards the future, which 'makes' matter thermodynamic, in the above sense. Whatever this basis is, it needs to be time-asymmetric itself. Otherwise, the generalisation would hold in both temporal directions, and entropy could only be constant. (Damping becomes anti-damping under time reversal, so matter that damped in both time directions would be incompatible with the existence of shakers.)

However, it is crucial to note that even a time-asymmetric nomological constraint of this kind would not give rise to the observed thermodynamic asymmetry unless entropy were initially lower than its maximum possible value—unless there were something around to do the shaking, in Feynman's example. Otherwise, the effect of the constraint would simply be to maintain a state of equilibrium, with no observable entropy gradient, and no asymmetry. So in order to explain the observed timeasymmetry, according to this conception of the Second Law, we actually need *two* timeasymmetric ingredients: the asymmetry nomological tendency or generalisation, and an asymmetric boundary condition, to the effect that entropy is low at some point in the past (again, in the ordinary time sense).

Diagrammatically, the observed time-asymmetry thus arises as follows, according to the two-asymmetry view:

Asymmetric boundary condition-entropy low in the past

+ Asymmetric lawlike tendency—entropy *constrained* to increase

Observed asymmetry.

As we shall see, these two-asymmetry views share the need for the first component—the asymmetric boundary condition—with their 'one-asymmetry' rivals. Their distinguishing feature is thus the attempt to find a basis for the second component—for a nomological time-asymmetry. (In practice, this is often seen as very much the more important aspect of the problem.)

In my view, the search for such a time-asymmetric principle or property is both misconceived and unsuccessful. It is unsuccessful in the light of one or both of the two objections I want to raise below to different versions of this two-asymmetry approach to the explanation of the thermodynamic asymmetry. And it is misconceived because there is a preferable alternative view of the origins of the observed time-asymmetry of thermodynamic phenomena, which—because it does not involve a time-asymmetric generalisation in the first place—simply does not face the problem of a basis or explanation for such a generalisation.²

²In Price 2002 I call the two-asymmetry view the 'causal-general' approach, and the one-asymmetry view the 'acausal-particular' approach.

3.2 The one-asymmetry view

We noted that a nomological time-asymmetry is at best half of what we need to explain the observed asymmetry of thermodynamic phenomena. A constraint preventing entropy from decreasing does not yield an asymmetry unless entropy is low to start with. This low entropy past boundary condition is the second of the two-asymmetry approach's two asymmetries.

The essence of the alternative approach is that this low entropy boundary condition is actually the only asymmetry we need. There need be no additional *time-asymmetric* generalisation at work in nature (at least in its thermodynamic manifestations), in order to account for the observed phenomena. This one-asymmetry approach originates in the late 1870s, in Boltzmann's response to Loschmidt's 'reversibility objections'. In essence, I think—although he himself does not present it in these terms—what Boltzmann offers is an alternative to his own famous H-Theorem. The H-Theorem offers a dynamical argument that the entropy of a non-equilibrium system must increase over time, as a result of collisions between its constituent molecules. As Burbury made clear in the 1890s (see Section 7), the time-asymmetry of the argument stems from a time-asymmetric independence assumption. Roughly, it is assumed that the velocities of colliding particles are independent *before* they interact.

The statistical approach does away with this dynamical argument altogether. In its place it offers us a simple statistical consideration. In a natural measure on the space of possible microstates of a physical system, microstates which are such that the system approaches equilibrium over time vastly outnumber those in which it behaves in other way in the set of all microstates compatible with a give macrostate). In a sense, this is not so much an explanation for the behaviour of the system as a reason why no special explanation is necessary—a reason for thinking the behaviour unexceptional.

As Boltzmann himself notes (1877, 193), the statistical considerations involved in this argument are time-symmetric. For a system in a given non-equilibrium macrostate, most microstates compatible with that macrostate are such that the system equilibrates towards the past, as well as towards the future. In practice, of course, this is not our experience. The low entropy systems with which we are familiar typically arise from systems of even lower entropy in the past. For example, the temperature difference between a cup of tea and its environment arises from the greater temperature difference between the boiling kettle and its environment. Entropy seems to decrease towards the past, and Boltzmann's time-symmetric statistics render this behaviour exceptionally puzzling—in making it unsurprising that the tea cools down, they make it puzzling why there was ever such a concentration of heat in the first place.

Taken seriously, Boltzmann's statistical approach thus directs our attention on the fact that entropy was very low at some point in the past. This 'boundary condition' is time-asymmetric, as far as we know, but this is the only time-asymmetry in play, on this view. The observed time-asymmetry of thermodynamic phenomena is thus taken to arise from the imposition of an asymmetric boundary condition³ on the time-symmetric probabilities of the Boltzmann measure. Diagrammatically:

Asymmetric boundary condition-entropy low in the past

+ Symmetric default condition—entropy likely to be high, ceteris paribus

Observed asymmetry.

3.3 Advantages of the one-asymmetry view

The immediate advantage of the one-asymmetry view is its relative simplicity, or theoretical economy. If it works, it simply does more (or as much) with less. The best response to such a claimed advantage would be to show that rival views do something more—and something worth doing, of course, for otherwise the extra expenditure is for nought. In this area, unfortunately, one's conception of what needs doing depends very much on very issue at stake. Proponents of two-asymmetry views take it that there is a time-asymmetric nomological generalisation evident in thermodynamic phenomena, and hence take explanation of this generalisation to be something worth paying for (in the

³The use of the term 'boundary condition' here is intended to reflect the role the proposition in question plays in the explanation of the thermodynamic asymmetry, relative to 'local' dynamical constraints on the behaviour of matter. Thus, it reflects the view that the condition does not flow from those constraints, but needs to be given independently. It is not intended to exclude the possibility that the condition in question might come to be treated as lawlike by some future physics.

coin of theoretical complexity). But this consideration will not move their oneasymmetry rivals, who deny that there is any such generalisation at work.

Disputes of this kind are difficult to settle, because the two sides have such different views of what counts as winning. The best dialectical strategy is often to exploit one's opponents' own conception of what counts as losing—in other words, to show that their view fails by their own lights. This is the strategy I shall follow below, in arguing against the two-asymmetry view. First, however, I want to respond to what may seem a objection to the one-asymmetry approach.

3.4 A surprising consequence of the one-asymmetry view?

The probabilities involved in the statistical picture are time-symmetric. They imply that entropy is very likely to be high in the past, as well as in the future. Yet entropy seems to have been extremely low in the past. Doesn't this amount to a strong disconfirmation of the view, by any reasonable standards? In other words, doesn't it give us good grounds for thinking that the Boltzmann measure is inapplicable to the real world?

It seems to me that there are two possible strategies available to one-asymmetry view, in response to this objection:

(i) *The no asymmetry strategy.* This involves saying that the Boltzmann measure is not erroneous, because the past low entropy is just a fluke—the kind of possible though extremely unlikely outcome explicitly permitted by the Boltzmann probabilities. This sounds like simply thumbing one's nose at the canons of confirmation theory, but Boltzmann himself suggests a way of making this option less unappealing than initially it seems. In an infinite universe, even very low entropy states may be expected to occur occasionally, simply by random fluctuation. If creatures like us can only exist in appropriate proximity to such fluctuations, then it is not surprising that we find ourselves in an otherwise unlikely kind of world.⁴ (Since this approach does away altogether with the need for an additional time-asymmetric restriction on the Boltzmann probability

⁴The closest Boltzmann seems to come to making this anthropic proposal explicit is in his (1895, 415). I am grateful to Jos Uffink for pointing out to me that it is not nearly as obvious that Boltzmann actually had this point in mind as tradition has tended to assume.

measure, it is genuinely a 'no asymmetry' approach: on the global scale, it tells us, there is no thermodynamic asymmetry!)

(ii) *The ceteris paribus strategy.* The second approach takes the Boltzmann probabilities to be merely 'default' expectation values, to be used in the absence of overriding factors— and the past low entropy boundary condition, as yet not fully understood, seems to be one such factor. In other words, the Boltzmann measure tells us what is likely to be the case, other things being equal—and we know of one relevant respect in which other things are not equal, namely, the low entropy past.

Note that the one-asymmetry approach cannot avail itself of a third escape strategy, that of claiming that the Boltzmann probabilities are reliable only in one direction. If the probabilities become asymmetric in this way, we no longer have a one-asymmetry view.

Thus for the one-asymmetry approach the choice seems to be between (i) and (ii). Option (i) is now widely regarded as unworkable. There are two major problems. The first is that if this suggestion had been true, we should not have expected to find any more order in the universe than was already known to exist, for that was certainly enough to support us. The second is that it is much easier for a fluctuation to 'fake' historical records, by producing them from scratch, than by producing the real state of affairs (of even lower entropy) of which they purport to be records. On this view, then, historical records such as memories are almost certainly misleading.⁵

Option (ii) has a surprising consequence. Given that the Boltzmann probabilities are time-symmetric, it implies that they have must the same 'default' status towards the future as towards the past, and hence that they cannot give us grounds for confidence that entropy will not decrease in the future. All we can reasonably say is that it is very unlikely to decrease, other things being equal—i.e., *in the absence of the kind of overriding*

⁵ For more details, see Price 1996, ch. 2, and Albert 2000, ch. 4. I believe that the point was made originally by C. F. von Weizsäcker 1939, in a paper which appears in English translation as §II.2 in von Weizsäcker 1980. Von Weizsäcker notes that "improbable states can count as documents [i.e., records of the past] only if we presuppose that still less probable states preceded them." He concludes that "the most probable situation by far would be that the present moment represents the entropy minimum, while the past, which we infer from the available documents, is an illusion" (1980, 144–145).

factors which make it decrease towards the past. In other words, statistical arguments alone do not give us good grounds for confidence that the Second Law will continue to hold. The statistics cannot rule out (or even give us strong grounds for doubting) the possibility that there might be a low entropy boundary condition at some point in the future, as there seems to be in the past. (In order to clarify this possibility, we need to understand more about the reasons for past boundary condition.)

In other words, the one-asymmetry view leads us to the kind of open-mindedness about the future thermodynamic behaviour of the universe that Burbury himself seemed to recommend (for somewhat different reasons, as we'll see), in the passage with which we began. In my view this is a major 'sleeping' consequence of Boltzmann's statistical approach to the thermodynamic asymmetry—'Boltzmann's Time Bomb', as I call it in Price 2002—which has been triggered by the recent discovery of the cosmological origins of the low entropy past.

Some people may feel that this consequence of the one-asymmetry view is a reason for preferring its rival—better time-asymmetric probabilities than probabilities whose predictions can't always be trusted. In my view, this reaction is a case of shooting the messenger. If this is what our best theory of the origins of the thermodynamic asymmetry tells us, shouldn't we take it seriously, rather casting around for an alternative theory? (It is not as though we have any good *independent* reason to reject the conclusion.)

Why is the one-asymmetry approach the best theory? In part because it is simpler, and in part because the alternatives are fatally flawed. Justifying the latter claim brings me to the main task of the paper, which is to lay out some objections to the twoasymmetry approach.

4. Two Kinds of Two-asymmetry View

I want to raise two objections to the two-asymmetry view. Neither objection applies to all versions of the two-asymmetry view, but all versions are subject to one objection or other, and some versions are subject to both. In order to make things clear, I need a simple taxonomy of possible views. For completeness, I include the one-asymmetry view. I'll present the taxonomy as a decision tree, as in Figure 1. Answers to a few simple questions lead to one to one or other of the four possible positions I want to distinguish.



Figure 1: Four possible views of the origin of the thermodynamic arrow.

The first issue is whether we need a second asymmetry at all—whether the explanation of thermodynamic phenomena requires a time-asymmetric nomological generalisation, in addition to low entropy past boundary condition. One-asymmetry approaches deny this, of course, and hence occupy position **A** on the tree.

Next, two-asymmetry approaches may usefully divided into two broad (though non-exclusive) categories, depending on their view of the source of the second asymmetry. One category comprises a group of approaches which seek a source for the second asymmetry in some asymmetric 'boundary condition', involve randomness, independence, or 'lack of correlation', in the initial motions of the microscopic constituents of matter. More precisely, these approaches take the second asymmetry to rest on the fact the fact that there is such randomness or lack of correlations in the *initial* but not the corresponding *final* conditions. Approaches of this kind may be further subdivided according to their answer to the following question. Is the 'internal' randomness of the microscopic motions of the constituents of a system sufficient (position **B**), or do environmental influences play a crucial role, as held by interventionist approaches (position **C**)?

The second broad category of two-asymmetry approaches comprises views that seek a dynamical *cause* for the second asymmetry—in other words, some identifiable factor dynamical without which entropy would not be constrained to increase. Again, views of this kind may usefully be sub-divided according to their view of the 'location' of this cause. For some views it essentially involves an external influence (position C). For others it is a feature of the *internal* dynamics of normal systems—i.e., it would normally be present even in a completely isolated system (position D).

Interventionism thus turns up in both broad categories. On the one hand, the interventionist's external influences comprise an identifiable *cause*, without which (according to the interventionist) entropy would not necessarily increase. On the other, it turns out to be crucial that these external influences are suitably *random*, uncorrelated with the internal motions of the systems on which they exert an influence.

Interventionism thus provides the obvious overlap between the two broad secondlevel categories of two-asymmetry approaches identified in Figure 1. In principle, these categories might also turn out to be non-exclusive for a more subtle reason, namely that 'initial randomness' itself might be held to be a dynamical cause for the second asymmetry. This possibility will be clarified below. For the time being, I want to discuss causal approaches and initial randomness approaches separately. As I have said, I want to raise two objections to two-asymmetry approaches. The first objection is to causal approaches, and hence applies to options C and D. The second is to the initial randomness approach, and hence applies to options B and C. Interventionism at least thus turns out to be doubly at fault, by my lights, being vulnerable to both objections. And if there is a further region of overlap between the causal and initial randomness approaches, the same will be true there.

5 The Counterfactual Confinement Problem

Causal approaches seek a dynamical factor responsible for the general tendency of entropy to increase—some factor without which entropy would not increase, at least with the observed regularity. As we noted, this factor needs to be time-asymmetric, for otherwise entropy would be constrained to be constant (non-decreasing in both directions).

One version of such an approach is interventionism, which takes the cause to be provided by influences from the external environment 'coming at haphazard' (in the words of Burbury (1894, 320), who seems to have been one of the first to suggest this idea). For the moment, let's ignore the role of 'haphazardness', or randomness, and focus on the external nature of these influences. I want to call attention to the implied counterfactual claim:

(5a) If there were no such external influences, the observed phenomena would be different—entropy would not increase monotonically, in the observed fashion..

This counterfactual seems part and parcel of what it means to say that the observed increase in entropy is *caused* by such external influences.

If we accept that we need a dynamical cause for observed monotonic increase in entropy, but reject interventionism, then the alternative is to locate the cause in some asymmetric feature of the internal dynamics of matter. A recent example of such a view is David Albert's (1994, 1998, 2000) suggestion that the collapse mechanism in the GRW interpretation of quantum theory provides such an asymmetry. The details of this proposal need not concern us here, but again I want to call attention to the implied counterfactual: (5b) If there were no such asymmetric mechanism, the observed phenomena would be different.

In his 1994 paper, Albert accepts this counterfactual claim, at least implicitly, in taking seriously the objection that if his suggestion were correct, entropy would not increase in systems containing too few constituents to allow the GRW mechanism to have its effect:

[T]he collapse-driven statistical mechanics ... will entail that an extraordinarily tiny and extraordinarily compressed and *absolutely isolated* gas will have *no lawlike tendency* whatever to spread out.

It can hardly be denied, therefore, that runs strongly counter to our intuitions. [sic]

What it does *not* run counter to, however (and *this* is what has presumably got to be important, in the long run) is our empirical experience. (1994, 677)

In a later paper, Albert accepts the counterfactual claim explicitly, saying that it is

perfectly right ... that anybody who claims that one or another causal mechanism called M is what actually underlies the tendencies of the entropies of thermodynamic systems to increase must also be claiming that if that mechanism were not *operating* then would *be* no such tendencies. (1998, 16)

Let us think about what this means. In a deterministic world, on this view, thermodynamic systems would not behave as they normally behave. Deterministic coffee would not grow cold, Newtonian cider would not need bottling. More precisely, if it were true that the observed thermodynamic behaviour is *caused* by some asymmetric indeterministic mechanism in the dynamics, then a means for turning off this mechanism would be a means for turning off the general tendency to equilibration. Useful as this might prove in practice, our present concern is with plausibility. Do we have good grounds for thinking that such counterfactuals are true? Similarly in the interventionist case. Is it really plausible that temperatures would not equalise in a completely isolated laboratory?

In the remarks quoted above, Albert says that these counterfactuals do not 'run counter to ... our empirical experience', and on the assumption that the mechanism in question (GRW, or random external influences) does obtain in the world of our empirical experience, this is quite true. But it is true in the trivial sense that in virtue of having an antecedent which is contrary to fact, *every* strictly *counter*factual conditional does not run counter to empirical experience. For strict counterfactuals especially, then, not running counter to experience is a very long way from being supported by experience (and, *pace* Albert, the latter is what has *really* got to be important, in the long run).

Albert also suggests that the issue is whether a gas has a 'lawlike tendency ... to spread out', and this seems to me to be a little misleading. The lawlike character of what we observe is disputable. Indeed, it is precisely the point at issue. As I have noted, oneasymmetry approaches to the thermodynamic arrow differ from their two-asymmetry rivals mainly in denying that there is any asymmetric *nomological* generalisation or disposition of matter to be explained. The real issue is not whether the gas would have a lawlike tendency to spread out, but simply whether it would spread out. It is not enough for these causal views to claim that a world without the cause in question would differ *solely* in modal respects.

These approaches thus require that the causal factor 'makes a difference'—that it alters the behaviour of matter in some way. Note especially the burden of proof. It is the proponents of these causal views who need a justification for the counterfactual claims (5a) and (5b). Their opponents need only say that they see no reason to accept such counterfactuals. Hence it is not sufficient for the causal view to object that its opponents have no basis for a contrary counterfactual claim. Thus Albert 1998 responds to the present objection as follows: 'If the GRW-theory is right, then there simply *are* no [Boltzmann-like probability distributions over initial conditions]; and so (of course) there would be no such distributions to fall back on in the event that one were to entertain a counterfactual sort of GRW-theory with the spontaneous localizations removed.' However, my point is that it is Albert who needs to justify the counterfactual (5b), not his opponents who need to refute it.

This objection seems to me to be very powerful, and it has a long history, at least as applied to interventionism. I'm not sure who first noted that it seemed implausible to claim that temperatures would not equalise, or gas not disperse, in a genuinely isolated system, but the point is certainly well known. Yet interventionism continues to attract supporters. This sort of situation is common in science and philosophy, of course, but it is often a sign that the two contending theories are operating with different conceptions of the nature of the problem. One way forward is to try to clarify the common task—to step back from the various proposed solutions, and ask what the problem actually is.

The debate about the asymmetry of thermodynamic phenomena seems to me very much in need of this kind of clarification. The approaches I have criticised in the section—interventionism, and the asymmetric internal dynamics view—seem to me to be take for granted that thermodynamic phenomena present us with a distinctive *lawlike* regularity, for which we need to find a basis. Once this is taken for granted (and it is accepted that the Second Law of Thermodynamics cannot simply be a primitive law, inexplicable in terms of any more basic feature of matter), it implies that there is *something* more basic, without which the regularity in question would not obtain. The task then is just to discover what that 'something' actually is. Against this background, the present objection is bound to seem more of an issue of detail than an objection of principle.

However, the background assumption is not compulsory. There is an alternative view of the phenomena, according to which equilibration towards the future is not a nomological regularity. On the contrary, it is the product of a weaker, *time-symmetric*, default condition on the behaviour of matter, in combination with a single time-asymmetric constraint (the low entropy past boundary condition). On this view—the one-asymmetry view—what happens towards the future is simply *not* the manifestation of a time-asymmetric nomological constraint.

Of course, what happens in the other direction—towards the past—is not in accordance with the time-symmetric default condition. Hence it does call for explanation, on this view. Thus there is a time-asymmetry in the phenomena, something puzzling in the behaviour of matter—but the puzzle is the way it behaves towards the past, not the way it behaves towards the future. In other words, the one-asymmetry approach takes interventionism and the asymmetric internal dynamics view to be trying to answer a question which doesn't need addressing.

It would be naive to hope that this kind of second-order point will be suddenly convincing, where the first-order objection has failed for so long. However, it does seem to me important to clarify the debate in this way. To a considerable extent, rival approaches to the thermodynamic asymmetry are not offering different answers to the same question, but answers to different questions. In order to make progress, we need first to recognise these competing conceptions of the nature of the problem—competing views of what an answer to the puzzle of the thermodynamic time-asymmetry would look like. In particular, we need to distinguish between the one-asymmetry and twoasymmetry conceptions of the nature of the task.

6. The Use and Abuse of Initial Randomness

I now turn to those versions of the two-asymmetry approach which attribute the general tendency of entropy to increase to a characteristic of the initial microstates of interacting matter—to 'initial randomness', or 'lack of initial correlations'. (In Price 1996 I called it the 'Principle of the Independence of Incoming Interactions', or PI^3 .) Again, this assumption needs to be time-asymmetric, applying to initial conditions but not to final conditions. Otherwise, it yields no asymmetry.

Figure 1 gave us two kinds of example of such views. One was interventionism, which relies on the assumption that the incoming influences exerted on a thermodynamic system by the external environment are uncorrelated with the microstates of the system itself. The incoming influences are assumed to 'come at haphazard'. In this case the time-asymmetry is held to reside in the fact that the

outgoing connections *to* the external environment do not 'go at haphazard', for otherwise, as just noted, the argument yields no asymmetry. In other words, it needs to be held that the system is correlated with its environment after their interaction, in a way in which it is not correlated beforehand.

The second version of this initial randomness approach is the one exemplified by the grandfather of would-be derivations of the Second Law, the H-Theorem itself. This differs from interventionism in that the initial randomness invoked is essentially an 'internal' matter, not dependent on a contribution from the environment. Burbury's great contribution to the discussion of the H-Theorem in the 1890s was to put his finger on the time-asymmetric assumption in question. It is the famous 'assumption of molecular chaos', or the *stoßzahlansatz*—roughly, the principle that the velocities of interacting molecules are independent *before* they collide. As we shall see in a moment, it was Burbury who pointed out why the corresponding condition could not be expected to hold *after* collisions.

My discussion of these initial randomness approaches is in two parts. I begin (Section 7) by examining the historical and theoretical basis of some assumptions typically associated with this approach, from which it derives much of its appeal. This examination reveals a surprising foundational lacuna, apparently obscured from view, at least in some quarters, by some longstanding errors of interpretation of early arguments in the field (including especially those of Burbury himself).

In the second stage, in Section 8, I want to treat the initial randomness approach in a more abstract way, paying particular attention to issues of explanatory and epistemological structure. I distinguish several possible forms the approach may take, and argue that all forms fall victim to one or other of several sub-problems. The main distinction is between (i) those versions of the approach which claim that our grounds for believing the entropy will continue to increase are that we have prior justification for believing in initial randomness (in some appropriate form); and (ii) those which take the epistemology to run the other way—which take the relevant initial randomness principle to be known by inference from an independently justifiable belief that entropy will continue to increase. I argue that in different ways, both versions suffer from the problem that their assumed epistemology is ungrounded—the required primary beliefs are simply not justifiable, by ordinary scientific standards. The latter approach also suffers from another problem. The initial randomness assumption it invokes turns out to come in weaker and stronger forms. In weaker forms the proposed explanation of entropy increase turns out to be question-begging in its epistemology, and to introduce no genuine second time-asymmetry, of the kind the two-asymmetry view requires. In stronger forms it runs out to be vulnerable to the counterfactual containment problem.

These objections may seem surprising, for the initial randomness assumptions have an apparent naturalness and plausibility. This is strikingly exemplified in what for more than a century has often been seen as one of the key problems in the thermodynamics of non-equilibrium systems—that of dealing with objections to the H-Theorem which turn on the claim that the assumption of molecular chaos cannot continue to hold over time, because the molecules of a gas become correlated as they interact. On this view, as Ridderbos puts it,

the central problem in non-equilibrium statistical mechanics [is that] equilibrium can only be obtained and entropy can only increase for those systems for which an appropriate analogue of the Stoßzahlansatz can be shown to hold. That is, a necessary condition for the approach to equilibrium is that the system has to get rid of correlations which are continually being built up dynamically by the interactions between the constituents of the system. (1997, 477)

A number of authors—e.g., Bergmann and Lebowitz (1953), Blatt (1959), Ridderbos and Redhead (1998)—have suggested that interventionism provides a solution to this 'central problem'. Where does the time asymmetry come from, on this view? From the fact, it is claimed, that the incoming influences 'from the environment' are uncorrelated, whereas the corresponding 'outgoing influences' are correlated, *in virtue of the interaction between system and environment*. As Ridderbos and Redhead put it: This is why the argument cannot be applied in the reverse time direction to argue that equilibrium will be approached into the past; in the ordinary time direction the 'incoming' influences are the influences from the environment on the system, and these are uncorrelated, but in the reversed time direction the 'incoming' influences are the influences the system exerts on its environment and these will be correlated. (1998, 1261)

This tradition thus seems committed to the following two principles:

(6a) Particle motions are independent *before* interactions—i.e., if the particles concerned have not interact in their common past.

(6b) Interactions give rise to (or in general increase) correlations.

These principles are logically independent, in the sense that neither implies the other. Yet each is clearly time-asymmetric, and necessarily so, if both are to be maintained. For each principle is incompatible with the *time-inverse* of the other.

On the face of it, moreover, neither of these principles depends on the low entropy past. Thus, remarkably, this tradition seems to be committed to the existence of *three* independent time-asymmetries: (6a), (6b), and the low entropy past boundary condition. And the combination is of doubtful coherency. After all, if there is a dynamical reason for the accumulation of correlations in one direction due to interaction, and (as typically assumed in these discussions) the dynamics is timesymmetric, then how could it fail to be the case that correlations accumulate in both directions? Puzzling assumptions, then, but as we have seen, also very natural ones to contemporary intuitions.

How did the tradition get to this point? Where did this peculiar combination of assumptions come from? It turns out that in order to answer these questions, we need to return to Samuel Burbury's contribution to the debate about the H-Theorem in the 1890s.

7. Origins of an Orthodoxy

In his initial (1894, 1895) contributions the debate of the 1890s, Burbury points out that Boltzmann's proof of the H-Theorem relies on what Burbury calls 'Condition A' (later 'Assumption A'). This amounts to the assumption that the velocities of colliding particles are independent, before collisions. Burbury also points out that if this condition applied also to reverse motions, the H-Theorem would imply that the gas in question would also be approaching equilibrium in reverse, and this could only be true if the gas were already at equilibrium. So in a case in which the gas is not in equilibrium, and in which Assumption A does hold initially, the assumption is not true of the reversed motions at the end of the interval in question.

After Burbury's initial presentation of the point in his (1894) and (1895) letters to *Nature*, the argument is reproduced in many places both by Burbury himself (e.g., Burbury 1899, §39) and by others. In particular, it is reproduced by the Ehrenfests in their important survey article in 1912 (Ehrenfest and Ehrenfest 1959, 40; 85, n. 65).

Roughly stated, Burbury's observation might seem to amount to the principle that particle motions are not uncorrelated after collisions—i.e., to principle (6b). It seems to be interpreted this way by Ridderbos, among others:

It was the Ehrenfests in their famous 1912 article [Ehrenfest and Ehrenfest 1959] who in a careful analysis of the Stoßzahlansatz pointed out the asymmetry which it contains; in general a distribution of molecules in a gas which satisfies the Stoßzahlansatz at an initial moment will fail to do so at a later moment *as a result of the correlations between the molecules which are built up dynamically as a result of the collisions*. (1997, 526, emphasis added.)

This interpretation is simply mistaken, however. The Burbury, Ehrenfest and Ehrenfest (BEE) argument that the *stofszahlansatz* (or Assumption A) does not hold after a period of interaction relies on two assumptions: (i) that the gas in question is not in equilibrium, and (ii) that the condition does obtain at the beginning of the time interval in question. Given these assumptions, the H-Theorem implies that entropy increases over the time interval in the forward time sense, and hence that it decreases in the reverse time sense.

From the latter fact, a second application of the H-Theorem, and *modus tollens*, we may infer that Assumption A does not obtain at the end of the time interval.

Note the crucial role of the assumption of initial non-equilibrium. Far from revealing a general tendency for correlations to accumulate *as a result of collisions*, the BEE argument turns on nothing more than this: Assumption A fails in the reverse motion *because* entropy was low in the past. Not only is this not a *result* of interactions; it seems independent of whether there actually are any interactions!

The BEE argument thus supports neither (6a) nor (6b). It provides no basis for the view that there is an asymmetrical tendency for interaction to give rise to correlations in the 'forward' time sense, as (6b) asserts. And gives us no a priori justification for assuming that Assumption A does hold at any given 'initial' moment, as (6a) maintains. On the contrary, in fact. Since our reason for thinking that Assumption A does not hold later is that we know that entropy was low earlier, we might think—by symmetry, as it were—that we are not justified in assuming with (6a) that Assumption A holds earlier, until we know whether entropy is low later. Clearly, this thought has the potential to undermine any proposal to make (6a) our *reason* for thinking that entropy will not decrease in the future. (More on this in a moment.)

It is worth noting that Burbury himself seems to regard Assumption A as unrealistic, and indeed, in the remark I quoted at the beginning, offers what seems a lone voice of caution about our right to assume that entropy will continue to increase. However, even the astute Burbury does not seem to doubt that we are justified in assuming (6a) for the case of incoming influences from an external environment. His caution about Assumption A seems to be based on the intuition that interactions within all but an ideally rarefied gas will induce violations of Assumption A (due to the production of 'velocity streams', for example). He does not seem to see that there is a more basic puzzle about the origin of the asymmetry in the assumption of independence for initial but not for final motions. The BEE argument is not the only reason for thinking that the *stoßzahlansatz* cannot continue to hold over an extended time period. A more basic reason is first noted, so far as I am aware, by J. H. Jeans:

The effect of this assumption [i.e., Assumption A] is to enable us to regard certain probabilities at any given instant as independent, and we then assume not only that the probabilities at a later instant are *inter se* independent, but also that they are independent of the events which took place at any earlier instant. This assumption cannot be logically reconciled with the fact that the motion of the system is continuous in time, *i.e.*, that the events which occur at any instant depend on those which occurred at a previous instant. (1903, 598)

Jeans calls this an a priori argument that Assumption A cannot continue to hold over a time interval. He then goes on to give a second argument, which he terms a posteriori. This latter argument is essentially the BEE argument. And Burbury, in turn, immediately adopts Jeans' a priori argument, endorsing it, for example, in (1903, 530) and (1904).

Does Jeans' a priori argument support (6a) and (6b)? Not at all, for it is not timeasymmetric. Nor incidentally, does it depend on the assumption that the constituents of the system in question actually interact. (As I noted, this is also true of the BEE argument.) In other words, Jeans' logical point cannot be the basis of a one-way tendency for correlations to be built up dynamically as a result of collisions.

I know of one more argument which might at first sight seem to support (6a) and (6b). Consider two particles, of initial momenta p and p'. If they interact with each other (and with nothing else), conservation of momentum ensures that their combined momentum after collision is p + p'. Doesn't this imply that the two particles are now correlated—that their momenta cannot vary independently?

However, a moment's reflection reveals that this simple argument cannot provide a basis for (6b) and (6a). For one thing, there is nothing time-asymmetric here. The argument applies equally well in reverse. Indeed, it applies not only at any earlier or later time, whether or not the particles collide; but also, in a sense, at the initial time itself. If we are given the total momentum, and told that there are only two particles, then, *relative to this information* the two momenta are not independent. The diachronic content of this argument is essentially that of Jeans' argument, viz., that in virtue of the determinism of the dynamics, a choice of the dynamical variables at any one time fixes those values at all other times (so long as the system remains isolated, of course).

Thus it seems to me that in so far as the tradition maintains that there is some time-asymmetric sense in which correlations accumulate as a result of collisions, it is simply mistaken. There is no sense in which the dynamics produces correlations in a forward time sense in which it does not also do so in the reverse time sense. As a result, the idea that the asymmetry of (6b) could bolster that of (6a)—could show why (6a) doesn't hold in reverse—is also without foundation. On the contrary, in fact. Our only relevant reason for thinking that there are correlations in *final* conditions are that we know entropy was low in the *past*. Far from showing us that there is a past–future asymmetry, this suggests that we cannot know whether there is such an asymmetry—whether, as (6a) claims, there are no corresponding correlations in *initial* conditions—until we know whether entropy is low in the *future*.

Perhaps even more damagingly, this brief excursion into the history of the initial randomness approach suggests that it provides no legitimate basis whatsoever for the nomological *second asymmetry* required by the two-asymmetry approach. At best, the asymmetry of micro-correlations is simply the *first asymmetry* re-described—the asymmetry of macroscopic boundary conditions, characterised in another way.

In order properly to evaluate these objections, however, it is necessary to bring their target more sharply into focus. It turns out that there are several subtle variants of the initial randomness approach, differently sensitive to these points. In the next section I want to look at the approach in abstract terms, paying close attention to matters of explanatory and epistemological structure—that is, in particular, to the issue of what is being taken as providing grounds for believing what. It turns out, I think, that while there are possible versions of the initial randomness approach which are not guilty of question begging, and which do genuinely introduce a second time-asymmetry, these versions have other problems. For one thing, they are subject to the counterfactual containment problem.

8 The Formal Structure of the Appeal to Initial Randomness

Let's begin with some terminology, to be read in conjunction with Figure 2. The propositions $HiEnt_p$ and $HiEnt_f$ say that there is a macroscopic condition of high entropy in the past and the future, respectively. **Ran**_f and **Ran**_p represent the corresponding microscopic conditions at the other end of time. Figure 2 depicts the combinations of these propositions and their negations we ordinarily assume to be true of the actual universe: \neg HiEnt_p and \neg Ran_f, and HiEnt_f and Ran_p.



Figure 2. Macroscopic entropy and microscopic correlation.

Thus the two shorter double-headed arrows in Figure 2 represent logical implications (modulo the laws). Given the dynamical laws, a high entropy past (HiEnt_p) would imply and be implied by the corresponding kind of randomness in the future (Ran_f); and a high entropy future (HiEnt_f) implies and is implied by the corresponding kind of randomness in the past (Ran_p). (In particular, therefore, Ran_p is logically the *weakest* assumption that implies HiEnt_f.)

In the most abstract terms, then, the initial randomness approach is committed to the idea that a high entropy future is a consequence of, and is to be explained by, the fact that Ran_p —that is, by the randomness of the earlier microstates. Perhaps the true explanans is some stronger principle of which Ran_p is a consequence—more on this option below—but at any rate the later high-entropy macrostate is thought of as a consequence of the earlier microstate, and Ran_p is by definition the weakest condition which will do the trick.

8.1 Explanation and epistemology

The initial randomness approach thus takes Ran_{p} to explain HiEnt_{f} . It also seems natural to take $\neg\operatorname{HiEnt}_{p}$ to explain $\neg\operatorname{Ran}_{f}$. Note that the 'intuitive' direction of explanation is past-to-future, in both cases. But this alignment in our explanatory intuitions should not be allowed to mask an important difference in the epistemological structures of the two claims, and hence a reason to be suspicious of the former. The grounds for suspicion turn on the issue of the justification the initial randomness approach takes us to have to accept that HiEnt_{f} .

In the case of the claim that \neg HiEnt_p explains \neg Ran_f, the epistemology is relatively unproblematic. Our grounds for thinking that \neg Ran_f—i.e. that there are microscopic correlations in the future—are simply that we have observed that the universe is highly ordered now and in the past. Epistemological inference thus follows the explanatory arrow, and all is well because the past state of affairs—the assumed explanans—is an observable matter.⁶

In the other case, however, the epistemology turns out to be problematic. To see why, let us note first that there are two broad possibilities. The first is that the epistemology does follow the explanatory arrow (taken to be past-to-future), so that our grounds for thinking that **HiEnt**f are that we have independent reason for thinking that the required microscopic randomness (**Ran**_p) obtains now (or in the past). The second is the converse possibility—our grounds for thinking that **Ran**_p obtains now and in the past are that we have independent reason to believe that entropy does not decrease in the future (from which past randomness follows as a logical or at least abductive consequence). Let us call these possibilities *past-to-future inference* and *future-to-past*

⁶For present purposes I ignore the sceptical issue concerning our entitlement to believe our apparent evidence for a low entropy past, stemming from von Weizsäcker's observation (see Section 3.4 above) that if Boltzmann's probabilities are our guide, then it is much easier to produce fake records and memories, than to produce the real events of which they purport to *be* records. Albert 2000, ch. 4, provides an excellent recent account of this sceptical point.

inference, respectively. I want to show that both are problematic, though for different reasons.

8.2 Past-to-future inference?

This version of the initial randomness approach requires that we have grounds for believing that **Ran**_p obtains, other than by inference from the fact that it is required for **HiEntf**. In what could these grounds consist?

The first option is that we might observe **Ran**_p directly. However, it is not difficult to see that this is not possible, at least in general. Just as it becomes impossible, in practice, to detect the correlations associated a system's low-entropy past—to tell, for example, whether a particular equilibrium sample of gas had a specified non-equilibrium state at a given earlier time—so it is impossible to make such determinations with respect to the future. (Among the factors which make this kind of observation impossible, in real systems, is the fact that measurements would be required across an entire spacelike hypersurface of the past light cone of the relevant segment of future spacetime.) Thus there is no prospect whatsoever of detecting directly the 'hidden' correlations which would be required for entropy to decrease in the future—nor, for the same reasons, of excluding them on observational grounds.

If our acceptance of Ran_p does not rest on direct observational grounds, then on what? Could there be some direct but non-observational reason for accepting Ran_p ? There seem to be two kinds of argument on offer at this point. One tries to appeal to the intuition above, that correlations are typically produced by interactions, and 'hence' that none are to be expected before interactions—incoming influences can be expected to be independent. However, we have seen that the history of the subject offers us no sound basis whatsoever for this appeal, or for the asymmetric principle about interactions on which it relies.

The second kind of argument appeals to statistical considerations. It claims that failure of Ran_p would require highly improbable coordination among the microscopic motions of matter, in order to give rise to $\neg HiEnt_f$. However, in the absence of an independent reason for thinking that the statistics are time-asymmetric, we have no

reason to think that the correlations required to give rise to \neg HiEnt_f are any more unlikely than those required to give rise to \neg HiEnt_p. Since the latter obtains *despite* these statistical considerations, we have very good reason to doubt the general reliability of statistical reasoning in this context. (After all, it fails in one case out of a possible two.)

It might be objected that we do have good reason—very good reason—to think that the relevant statistical arguments are reliable past-to-future, namely, that they perform so well in predicting the past-to-future behaviour of many real systems. However, in so far as it claims to offer us a reason to believe **Ran**_p, this argument involves an inductive step. In effect, it relies on the claim that we are justified in expecting this statistical success to continue. As such, the claimed inference to **Ran**_p thus becomes indirect, and future-to-past (on which more in a moment).

Thus it appears that the epistemology of the claimed explanation of $HiEnt_f$ by Ran_p cannot follow the explanatory arrow. Unlike in the case of the explanation of $\neg Ran_f$ by $\neg HiEnt_p$, our grounds for accepting the future explanandum cannot be that we have independent grounds to accept the past explanans. If we have good reason to believe the supposed explanandum at all, then, it must be on some basis more direct. Let us now turn to that possibility.

8.3 Future-to-past inference—a direct case for HiEntf?

In this version, the initial randomness approach is the view that we are led to some sort of initial randomness principle by abductive inference from the prior discovery that entropy will continue to increase. It is the view that in some form, initial randomness is the best explanation for this independently established fact about the physical world.

I want to raise two kinds of objection to this proposal. The first echoes Burbury's point, from the quotation with which we began. We simply *don't* have particularly strong reasons for believing that entropy will continue to increase, at least in the distant future. However, although the sentiment is Burbury's, the present justification for it depends on considerations of which Burbury could not have been aware. It depends on our current understanding of the nature of the past low entropy boundary condition, and hence on modern cosmology.

Within the last thirty years or so, cosmology has offered us a plausible hypothesis as to the origins of the observed low entropy. Briefly, everything seems to turn on the fact that matter was very smoothly distributed, soon after the Big Bang. This is a very low ordered condition for a system dominated by gravity, and seems to supply a lowentropy 'reservoir' which supplies the entire thermodynamic gradient we observe in our region. (This story is very well told by Penrose 1989. On its links to the present issue, see Price 1996, ch. 2, Price 2002.)

Given this reason to think that the low entropy past has cosmological origins, we also have some basis for thinking that the issue of future low entropy also turns on cosmological issues. In particular, if we are entitled to believe that entropy will not decrease in the future, then it must be because we are entitled to believe—on cosmological grounds—that there is no future low entropy condition, of the same kind as the past condition. There seem to be two possible routes to such a belief. One would be theoretical, and turn on the question as to whether our best cosmological models excluded the possibility (or at least the likelihood) of such a future condition. The other would be observational, at least in part, and turn on the issue of what *present* manifestations there might be of such a future boundary condition. At present, both routes are best very inconclusive. (I shall say more about the latter possibility in a moment.) As it stands, then, Burbury's cautious agnosticism remains the appropriate attitude. In other words, it is very far from clear that there is an explanandum (HiEntf) in need of the kind of explanans (Ran_p) this approach takes initial randomness to offer.

I shall also argue that even if we did have good reason for thinking that entropy will continue to increase, this wouldn't count in favour of the initial randomness approach, which would still fall victim to the following dilemma. Depending on what is meant by initial randomness, it is either (i) a condition too weak to support a non-trivial two-asymmetry approach to the thermodynamic asymmetry; or (ii) a condition that, in being stronger than necessary, is subject to an analogue of the counterfactual containment problem.

8.3.1 Why HiEntf isn't presently observable.

If entropy were to decrease on a large scale at some time in the future, would that fact be detectable now? If we assume determinism, then in one sense, of course, it would be observable in principle to a Laplacian ideal observer. As we have already noted, however, there is no prospect whatsoever of this sort of microstate-based observation being possible for real observers in real systems. If it were possible, moreover, then it would count as direct observation of **Ran**_p, and hence as a basis for a past-to-future inference. The possibility we have now to consider is different. It is that there might be macroscopic evidence now as to whether entropy continues to increase in the future.

This possibility is not as bizarre as it might at first sound. There seem to be quite good grounds for thinking that certain kinds of future low entropy boundary conditions would have observable consequences now—and hence that at least to some extent, we can have observational evidence that those boundary conditions do not obtain in the future. The kind of evidence concerned is very much the kind of evidence we have of the low entropy past. Essentially, the latter evidence consists in low entropy 'remnants'—in systems that had not yet had time to reach equilibrium, following the low entropy condition in the past. (Stars and galaxies are striking examples of such remnants.)⁷

In an exactly analogous way—and proceeding on the same assumptions, applied in reverse—future low entropy conditions might well be expected to have manifestations now. These would be systems that, in the reverse time sense, have likewise not had sufficient time to equilibrate.

This possibility has been discussed in the literature in connection with a proposal made by Thomas Gold 1962, that the universe might be globally symmetric in entropy terms, with a low entropy future endpoint. The question of the advanced effects of such a boundary condition has been addressed in variations of the Ehrenfest two-urn models, with two-time boundary conditions. (See, for example, Cocke 1967 and the references in Schulman 1997, 156.) The upshot of these investigations is that such a boundary condition at a time T_f is not expected to be detectable before $T_f - t_{relax}$, where t_{relax} is

⁷Again, I am ignoring von Weizsäcker's sceptical difficulties about inference to the past.

the relaxation time of the relevant physical processes. Beyond that relaxation time, the system behaves in a manner macroscopically indistinguishable from one which lacks the future boundary condition. (The predictions assume a statistical measure that is symmetric, before application of boundary conditions. However, the conclusion concerning detectability outside the relaxation time seems likely to apply even in asymmetric background measures.)

Thus it seems that even a very large low entropy future boundary condition, comparable in magnitude to that in our past, might be undetectable now, so long as it is sufficiently far away. How far away is sufficiently far away? The issue depends on the relaxation time of the relevant real physical processes. This question has been discussed a little in recent years (e.g., in Gell-Mann and Hartle 1994; see also Price 1996, ch. 4), but the relevant issues remain open. At any rate, it is clear that at present there is no strong observational case for saying that there isn't a low entropy boundary condition in the future.

It might be thought that there is easier observational route to the conclusion that HiEntf. Can't we simply get there by induction from the observed present behaviour of matter? We observe that the Second Law holds now, with great consistency. Isn't that a reason to believe that it will continue to hold in the future? In other words, doesn't induction give us our explanandum (and abduction then our explanans)?

The problem is that induction depends on the assumption that the future will be like the past in the relevant respect. It is therefore question-begging to try to use it to exclude the contrary hypothesis. In other words, induction is powerless to exclude a theoretically well-motivated hypothesis to the effect that the future will not be like the past. (This is simply a more localised version of Hume's point, viz., that we can't use inductive methods to justify induction.) Indeed, as I noted in Section 3.1, the issue between the one-asymmetry and two-asymmetry views seems to come down to that as to whether there is a *projectible* (or induction-supporting) generalisation to the effect that entropy will continue to increase. There is no prospect of settling this issue by induction. 8.3.2 Why initial randomness would be in trouble, even if HiEntf were observable

Suppose for a moment that we did have good theoretical or observational grounds for believing HiEnt_f. Would we be justified in making an abductive inference to Ran_p? Yes, at least on some conceptions of explanation. Since Ran_p is by definition equivalent to HiEnt_f, given the relevant laws, it follows that Ran_p is an initial condition sufficient to imply HiEnt_f, given the laws. If this is held to be sufficient for explanation, then Ran_p can indeed explain HiEnt_f.

Unfortunately for the two-asymmetry approach, however, such an explanation does not conform to the two-asymmetry model. Ran_p implies $HiEnt_f$, but the implication is logical, not nomological. So long as $HiEnt_f$ is a factlike matter, so too is Ran_p . After all, the one-asymmetry view also accepts that Ran_p . On this view, the default expectation is for both Ran_p and Ran_f (or equivalently, both $HiEnt_f$ and $HiEnt_p$); and the single asymmetry is the constraint which supplies $\neg HiEnt_p$ (or equivalently $\neg Ran_f$).

Thus if abductive inference is to get us anywhere useful from HiEnt_f, from the two-asymmetry approach's point of view, it has to get us to something stronger than **Ran**_p. Call such a stronger proposition ***Ran**_p. In effect, the initial randomness approach will then be committed to the counterfactual

(8a) If not *Ran_p, then (probably) not HiEnt_f.

For if $HiEnt_f$ were held to be likely to occur anyway, even in the absence of * Ran_p , the supposition that * Ran_p could hardly be required to *explain* HiEnt_f.

But why should we accept (8a)? We know it is *logically* possible that Ran_p (and hence HiEntf) should obtain without *Ran_p. (By definition, Ran_p is sufficient for HiEntf, and *Ran_p is logically stronger than Ran_p, so failure of *Ran_p does not entail failure of Ran_p.) This version of the initial randomness approach will therefore need to argue on non-logical grounds that if *Ran_p failed, so too would Ran_p—the result would not simply be a world in which Ran_p obtained anyway.

The dialectical position of the initial randomness approach here is the same as that of causal versions of the two-asymmetry theory. As we saw in Section 5, those approaches are committed to the claim that if the causal mechanism in question did not operate—if a system were truly isolated from its environment, if there were no GRW collapse, or whatever—thermodynamic systems would not behave as they are observed to behave. In effect, this amounts to the idea that entropy needs to be *prevented* from decreasing—that the initial condition of the universe is such that entropy will decrease, unless something intervenes. The one-asymmetry approach challenges this idea, arguing that the observed phenomena can be construed as the product of a single timeasymmetric constraint (the low entropy past) on an otherwise symmetric space of possibilities. Unless we are given some reason to rule out this conception of the origin of the phenomena in question, the causal approach is not entitled to its counterfactuals. While the one-asymmetry model remains a viable possibility, in other words, we cannot be justified in claiming that without the second asymmetry, the relevant phenomena would have been different.

All this transfers to the present case. Given the supposition that HiEntf, the oneasymmetry approach will propose that the resulting past–future asymmetry—the fact that HiEntf but not HiEntp—reflects only one-time asymmetry, that of the low entropy past boundary condition. In order to make a case for (8a), and hence for a second asymmetry in the form of *Ranp, the initial randomness approach needs a reason to disallow this one-asymmetry proposal. No such reason seems to have been forthcoming.

8.3.3 Summary: The case against nomological initial randomness

The argument against an appeal to a 'strong' form of initial randomness may be summarised in this way. If *Ran_p were to be justified on theoretical grounds, these grounds would need to be time-asymmetric (since otherwise they would equally give us *Ran_f, and hence HiEnt_p). A time-asymmetric theoretical assumption of this kind might be justified if we already knew HiEnt_f, and therefore sought an explanation of this known fact. But (i) this is not our situation—on the contrary, we are still looking for some reason to believe HiEnt_f. And (ii) even if we did have such reason, *Ran_p wouldn't be necessary, since we have no reason to think that situation is not merely the result of symmetric Boltzmann probabilities, subject to a single time-asymmetric boundary condition.

8.3.4 Is the argument too strong?

It might be objected that the last objection proves too much. If it worked, wouldn't it show that nomological explanations are always too strong? For isn't the possibility always open that the same phenomena would have occurred without the law in question?

This is an interesting point, which raises issues which go far beyond the scope of this paper. In my view, the counterfactual containment argument does pose a general problem for some strongly realist conceptions of the nomological realm. For present purposes, however, what matters is something that seems to distinguish the use of argument in the thermodynamic context from possible general uses.

In the present context, we have a relatively clear alternative to the two-asymmetry approach's conception of the origins of the relevant thermodynamic phenomena. That is, we have clearly in view a means by which the same phenomena could have arisen, without a nomological asymmetry. Assuming a deterministic dynamics, all it takes is for the initial microstate of the universe to be 'normal', in the sense of Boltzmann's measure, in the space of possibilities compatible with the initial macrostate. Hence the onus is on the proponent of a two-asymmetry view to convince us that without the nomological asymmetry supposedly in question—GRW collapse, external influence, *Ran_p, or whatever—the initial microstate would not have been of that kind. (As I noted in Section 5, there is no corresponding onus on opponents of the two-asymmetry approach to *defend* the Boltzmann measure. Proponents of the two-asymmetry approach are obliged to justify the relevant counterfactual, and hence to *exclude* the Boltzmann measure, but opponents can simply afford to be agnostic.)

A familiar argument for realism about nomological necessity is that without it, observed regularities become incredible coincidences. The implicit counterfactual claim is that without real necessities, the regularities in question would probably fail. Whatever we think about the merits of this argument, it is better than the analogous defence of nomological asymmetry in the thermodynamic case. There may perhaps be a prima facie case that regularities in general would fail without laws. Post-Boltzmann, there isn't even a prima facie case that thermodynamic systems would equilibrate differently without a nomological constraint. Boltzmann showed us that there is at least one plausible measure on the space of possible initial microstates under which the observed phenomena are highly likely, even in the absence of a nomological asymmetry. To make their view stick—to return the ball—defenders of nomological asymmetry. To make their view stick—to return the ball—defenders of that view need a reason to rule out Boltzmann's measure. So far as I know, there has been no serious, let alone *successful*, attempt to do so. (Once again, the point is not that Boltzmann's measure is a priori *right*, but simply that it is a serious contender, which has not been shown to be *wrong*.)

8.4 Conclusion

When disambiguated, the initial randomness approach turns out to be unable to provide a satisfactory explanation of the second asymmetry, of the kind the two-asymmetry approach requires. The faults of the approach depend on the precise form in which it is presented to us.

If it is offered in the spirit of a justification for believing that entropy will continue to increase, then it fails because there is no good reason to believe that the required condition of initial randomness does obtain. (It is not observable, and statistical arguments are unreliable, by parity of reasoning.)

If it is offered to us in the spirit of an inferred explanation of something independently determinable, then again the approach fails. In this case, the problem is both that the explanandum is not independently justifiable, and that even if it were, the proposed explanans would be either (i) too weak for the two-asymmetry approach, in not supplying a nomological second asymmetry, or (ii) unnecessary.

9. Starting in the Wrong Place

The failings of the two-asymmetry approach seem to boil down to two points. The first is an epistemological failing, a vulnerability to a modern version of Samuel Burbury's objection. It simply hasn't been adequately established that on a global scale, the relevant phenomena are as the two-asymmetry approach takes them to be. That is, it hasn't been established that entropy will continue to increase. We saw that what has been added since Burbury's time is the cosmological case for regarding this as an open question. Burbury was rightly suspicious of the crucial assumption (the assumption that he himself had first identified) of the leading argument for thinking that entropy will continue to increase. The cosmological connection supplies an independent reason for holding open the possibility that it will not do so.

The second failing is methodological. Even if we grant the relevant phenomena, the two-asymmetry approach has failed to establish that the explanation of those phenomena requires something either nomological or time-asymmetrical (in addition, that is, to the low entropy past boundary condition). To make this case, the twoasymmetry approach needs to defend the kind of counterfactual we identified at (5a), (5b) and (8a), and thus to exclude the alternative one-asymmetry conception of the origin of the asymmetric phenomena in question. No such defence has been offered.

In my view, these two failings are symptoms of a more basic failing, that of asking the wrong question to start with. The two-asymmetry approach begins with the question: Why does entropy go up towards the future—why is matter engaged in this 'uphill' journey? The one-asymmetry approach begins instead with the question: Why does entropy go down towards the past—why does the matter begin its journey at such a 'low' spot? It is true that by one-asymmetry lights these are in a sense the same question: the journey is uphill *because* it starts in a low place. All the same, the latter version puts the emphasis in the right place—it brings to the foreground the crucial puzzle.

By two-asymmetry lights, the questions are not the same. A reason why matter cannot go 'downhill' is not automatically a reason why its journey must start at a low place. Hence the need for two asymmetries, in this picture—and so much the worse for this approach, as we have seen.

In my view, a great but still under-appreciated consequence of Boltzmann's contribution to our understanding of the thermodynamic asymmetry is that it directs our attention 'backwards' in this way. It show us that if our interest is in the timeasymmetry, the proper focus of our attention is the condition of the universe in what we regard as its past, rather than some intrinsically time-asymmetric characteristic of its current journey.

Thus the crucial question is why entropy is so low in the past, and as noted in Section 8.3, modern cosmology has given this issue a remarkably concrete form. But this question arises from an even more basic one, which isn't essentially time-asymmetric: Why isn't entropy almost always high, as the time-symmetric Boltzmann measure would lead us to expect? We'll still need to answer this latter question, even if—as we currently have no very strong reason to disbelieve, in my view—entropy turns out to decrease in the distant future, and the 'end' of the universe is as peculiar as its 'beginning'.⁸

Department of Philosophy University of Edinburgh David Hume Tower Edinburgh EH8 9JX

and

Department of Philosophy University of Sydney NSW Australia 2006

Huw.Price@ed.ac.uk

⁸Early versions of some of this material were presented at a conference in Groningen in September, 1999, at the Royal Institute of Philosophy Conference at LSE in September, 2000, and in talks in Utrecht, College Park and Tucson in April, 2001. I am grateful to participants for discussions on those occasions; and indebted to David Atkinson and Craig Callender for comments on earlier drafts, amongst much else.

- Albert, D. 1994: 'The Foundations of Quantum Mechanics and the Approach to Thermodynamic Equilibrium', *British Journal for the Philosophy of Science*, 45, 669–677.
- Bergmann, P.G. and Lebowitz, J.L. 1955: 'New Approach to Non-equilibrium Processes', *Physical Review*, **99**, 578–587.
- Blatt, J.M. 1959: 'An Alternative Approach to the Ergodic Problem', *Progress in Theoretical Physics*, 22, 745–756.
- Boltzmann, L. 1877: 'Über die Beziehung zwischen des zweiten Hauptsatze der mechanischen der Wärmetheorie" ('On the Relation of a General Mechanical Theorem to the Second Law of Thermodynamics"), *Sitzungsberichte, K. Akademie der Wissenschaften in Wien, Math.-Naturwiss*, 75, 67–73 (reprinted in Brush 1966).
- ————1964: Lectures on Gas Theory, Berkeley: University of California Press.
- Brush, S. 1966: Kinetic Theory. Volume 2: Irreversible Processes, Oxford: Pergamon Press.
- Burbury, S. H., 1894: 'Boltzmann's Minimum Function', Nature, 51, 78.

- ——1904: 'On the Theory of Diminishing Entropy', *Philosophical Magazine, Series 6*, 8, 4349.
- Cocke, W. 1967: 'Statistical Time Symmetry and Two-Time boundary Conditions in Physics and Cosmology', *Physical Review*, **160**, 1165–1170.

- Culverwell, E. 1890a: 'Note on Boltzmann's Kinetic Theory of Gases, and on Sir W. Thomson's Address to Section A, British Association, 1884', *Philosophical Magazine*, **30**, 95–99.

Distributions', Nature, 50, 617.

Ehrenfest, P. and Ehrenfest, T. 1959: *The Conceptual Foundations of the Statistical Approach in Mechanics.* English translation by M.J. Moravcsik (Ithaca, NY: Cornell University Press). Dover edn. 1990 (NewYork: Dover Publications).

Gell-Mann, M. and Hartle, J. 1994: 'Time Symmetry and Asymmetry in Quantum Mechanics and Quantum Cosmology,' in Halliwell, Perez-Mercader, and Zurek (1994), pp. 311–345.

- Gold, T. 1962: 'The Arrow of Time,' American Journal of Physics, 30 403-410.
- Gold, T. (ed.) 1963: The Nature of Time. Ithaca: Cornell University Press.
- Hall, E. H. 1899: Review of S. H. Burbury, *The Kinetic Theory of Gases* (Cambridge: Cambridge University Press, 1899), *Science, New Series*, **10**, 685-688.
- Halliwell, J., Perez-Mercader, J. and Zurek, W. (eds.), 1994: *Physical Origins of Time Asymmetry*, Cambridge: Cambridge University Press.
- Jeans, J. H. 1903: 'The Kinetic Theory of Gases developed from a New Standpoint', *Philosophical Magazine, Series 6*, 5, 587-620.
- Lebowitz, J. 1993: 'Boltzmann's Entropy and Time's Arrow', Physics Today, 9:93, 32-38.
- Price, H. 1996: Time's Arrow and Archimedes' Point: New Directions for the Physics of Time, New York: Oxford University Press.
- Ridderbos, T. M. 1997: 'The Wheeler-Feynman Absorber Theory: A Reinterpretation?', Foundations of Physics Letters, 10, 473–486.

- Ridderbos, T. M. and Redhead, M. 1998: 'The Spin-echo Experiments and the Second Law of Thermodynamics', *Foundations of Physics*, 28, 1237–1270.
- Schulman, L. 1997: *Time's Arrows and Quantum Measurement*, Cambridge: Cambridge University Press.
- Sklar, L. 1995: 'The Elusive Object of Desire: in Pursuit of the Kinetic Equations and the Second Law', in Savitt, S., ed., *Time's Arrows Today.*, Cambridge: Cambridge University Press, 191–216. Originally published in Fine, A. and Machamer, P., eds., *PSA 1986: Proceedings of the 1986 Biennial Meeting of the Philosophy of Science Association*, vol. 2.
- von Weizsäcker, C. 1939: "Der zweite Hauptsatz und der Unterschied von der Vergangenheit und Zukunft," *Annalen der Physik (5 Folge)*, 36, 275–283.
 ——1980: *The Unity of Nature*, New York: Farrar Straus Giroux.