

Can Conditioning on the “Past Hypothesis” Militate Against the Reversibility Objections?*

Eric Winsberg^{†‡}

In his recent book, *Time and Chance*, David Albert claims that by positing that there is a uniform probability distribution defined, on the standard measure, over the space of microscopic states that are compatible with both the current macrocondition of the world, and with what he calls the “past hypothesis”, we can explain the time asymmetry of all of the thermodynamic behavior in the world. The principal purpose of this paper is to dispute this claim. I argue that Albert’s proposal fails in his stated goal—to show how to use the time-reversible dynamics of Newtonian physics to “underwrite the actual content of our thermodynamic experience” (Albert 2000, 159). Albert’s proposal can satisfactorily explain why the overall entropy of the universe as a whole is increasing, but it does not and cannot explain the increasing entropy of relatively small, relatively short-lived systems in energetic isolation without making use of a principle that leads to reversibility objections.

1. Introduction. One of the oldest foundational puzzles in the philosophy of physics concerns the origin of the time asymmetry of the second law of thermodynamics. The laws of thermodynamics have a temporal arrow associated with them, but none of the candidates for the dynamical laws from which we suppose them to arise have such an arrow. That such a puzzle exists has been well known since the so called “reversibility objections” to Boltzmann’s theory of statistical mechanics were articulated by Loschmidt more than a hundred years ago.¹

*Received August 2003; revised January 2004.

[†]To contact the author write to Department of Philosophy, The University of South Florida, FAO 226, 4202 Fowler Ave., Tampa FL 33620; e-mail: winsberg@cas.usf.edu.

[‡]I thank Mathias Frisch, Paul Teller, David Albert and three anonymous reviewers for helpful discussions, comments, and criticisms. Special thanks to Bailey Quillen of the Pinellas County Center for the Arts at Gibbs High School, Florida, for help with the illustrations.

1. For a thorough discussion of the history of the various “reversibility” and “recurrence” objections to Boltzmann’s ideas, see Chapter 1 of Lawrence Sklar’s (1995).

Philosophy of Science, 71 (October 2004) pp. 489–504. 0031-8248/2004/7104-0004\$10.00
Copyright 2004 by the Philosophy of Science Association. All rights reserved.

In his remarkable recent book, *Time and Chance*, David Albert (2000) outlined a proposal² that he argues can finally close the book on the century-old foundational puzzle. The proposal will have to be outlined with greater care in due course, but roughly, Albert claims that by positing that there is a uniform probability distribution defined, on the standard measure, over the space of microscopic states that are compatible with the current macrocondition of the world, conditionalized on what he calls the “past hypothesis,” we can explain the time asymmetry of all of the thermodynamic behavior in the world.

The principal purpose of this paper is to dispute this now widely held claim.³ Specifically, I will argue in this paper that while Albert’s proposal does contain one very powerful insight, it nevertheless fails in its stated goal—to show how to use the time-reversible dynamics of Newtonian physics to “underwrite the actual content of our thermodynamic experience.” (159). While the details will be complex, the reason is simple: Albert’s proposal can satisfactorily explain why the overall entropy of the universe as a whole is increasing, but I will argue that it cannot explain the increasing entropy of relatively small, relatively short-lived⁴ systems in energetic isolation without succumbing to the same problems that beset Boltzmann.

Let’s begin by outlining the source of the problem, as it has been more or less understood from the beginning. Suppose we have a thermodynamic system S , and a macroscopic law L that describes systems like S . We also have a set of laws M that govern the microscopic components of S . Suppose furthermore that L tells us that if S is in a macrocondition A , then it will evolve over some period of time into macrocondition B . The goal is to provide an account of how it is that our macroscopic laws, whatever they may be.

2. Reversibility. If the macroscopic law L is the second law of thermodynamics, then we almost certainly must begin by making plausible something like the following principle. And indeed, most students of statistical mechanics believe that such a principle, which we will presently label

2. In the book, Albert does not make it altogether clear what the origin of this proposal is—that is, to whom the various ideas are due. I will not concern myself with these questions here. When I refer to “Albert’s proposal” or say, “Albert argues . . .” consider it to be for brevity only. I make only one remark. The view is often attributed to Richard Feynman in his (1967). As far as I can tell, however, Feynman only makes the claim that positing the “past hypothesis” is a necessary condition to explaining the second law, but not that it is sufficient.

3. In addition to (Albert 2000), see also (Callender 2001), and (Lebowitz 1999).

4. With respect to the size and age of the universe.

Principle 1, can be made plausible by the kind reasoning made famous by Boltzmann and Gibbs.

We begin by representing the complete state of an isolated system made up of N classical particles by a point \mathbf{X} in a phase space. The point \mathbf{X} is denoted by $(\mathbf{x}_1, \mathbf{p}_1, \mathbf{x}_2, \mathbf{p}_2, \dots, \mathbf{x}_n, \mathbf{p}_n)$, where \mathbf{x}_i and \mathbf{p}_i are the position and momentum, respectively, of the i th particle. If an isolated system of particles is in some macrocondition A , then we define the N dimensional “volume” of the subregions r of the entire region R compatible with macrocondition A to be (Lebowitz, 1999):

$$|\Gamma_r| = \int_{\Gamma_r} \prod_{i=1}^N \partial x_i \partial p_i$$

normalized over $|R|$. We can now state the principle, call it Principle 1 (see Figure 1).

Principle 1 (P1). If condition A is a condition of low entropy, and condition B is a condition of maximum entropy (the macrostate whose constitutive microstates occupy the largest volume), then the overwhelming majority of the volume of the region compatible with macrocondition A —all but a tiny volume of “abnormal regions”—evolves into the region of microstates compatible with B .

It’s important to be clear about the status of P1. P1 has not been shown to be true by mathematical reasoning or theoretical demonstration, but neither is it simply being assumed as a postulate. Rather, most people believe that it has been made plausible by a variety of considerations—that is, most believe it is likely a consequence of a minimal set of assumptions about the microlaws M , even if that consequential relation cannot be rigorously demonstrated.

So far so good, but that is only the first step; we need more. So now suppose that we are given a system like S , and only the information that it has the property of being in macrocondition A . We know that the microcondition of this system lies within some region R of phase space; the region that is compatible with A . If, as our next step, we were to assume that, given any system S in condition A , the probability of the microcondition of S being in some tiny region of R , (call it r) is proportional to the phase space volume of r , then we would be off to the races.

Since, by P1, almost all of the microconditions compatible with A (in the sense of phase space volume) evolve into microconditions compatible with B , if the probability of the microcondition of S being in r is proportional to the volume of r , then the probability that condition A will evolve into condition B is overwhelming high. And this is exactly what we want to show.

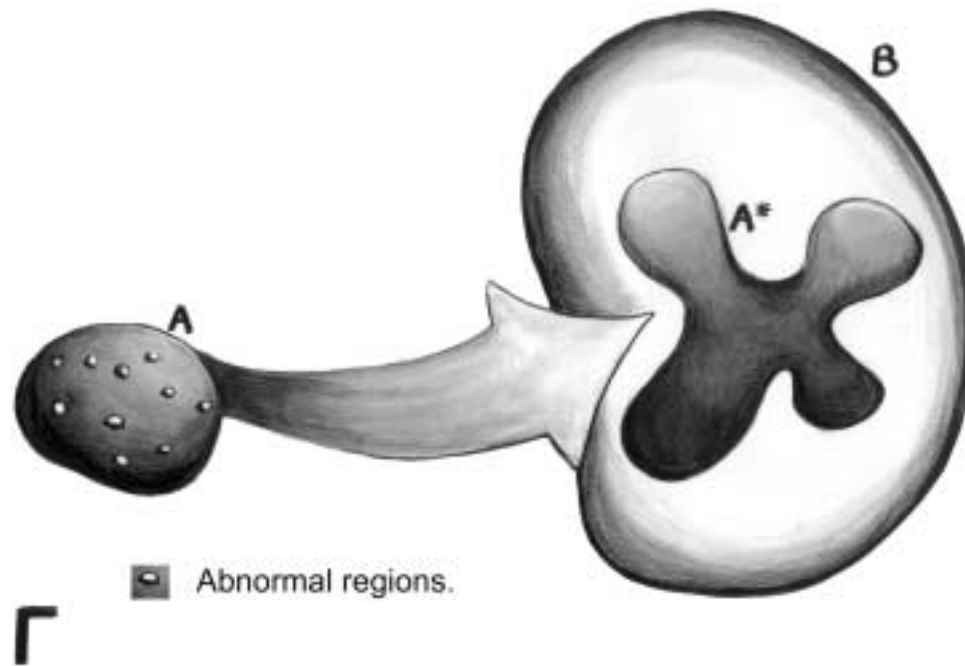


Figure 1. Principle 1. If A is region of lower entropy, and B is the region of maximum entropy, then the overwhelming majority of the volume associated with A evolves into A^* , inside of B . The tiny remainder of A , which does not evolve into B , we call the "abnormal regions."

Unfortunately, all we have accomplished so far is to have put ourselves in a position to better articulate the real source of the original problem. Let's examine the supposition that "given any system S in macrocondition A , the probability of the microcondition of S being in some tiny region r of R (the region compatible with A), is proportional to the volume of r (on the standard measure)." Let's call this Boltzmann's Postulate (BP) for brevity (see Figure 2).

The problem with BP arises because we are inclined to believe that the laws that govern the transitions from microstate to microstate are, in the relevant sense, time reversible. More precisely, we take our microlaws, at least on the classical picture, to have the following property: The dynamics specified by those laws link a microstate at time t_0 , $\mathbf{X}(t_0)$, to microstates $\mathbf{X}(t)$ at all times t . Now, take $\mathbf{X}(t_0)$ and $\mathbf{X}(t_0 + 1)$, for any $t > 0$. If we reverse all of the velocities at time $t_0 + 1$, we obtain a new microstate. If we now follow the evolution for another interval from this new microstate, "reversibility" tells us that we will find the new microstate at time $t_0 + 2$ is just $R\mathbf{X}(t_0)$, the microstate $\mathbf{X}(t_0)$ with all velocities reversed, $R\mathbf{X} = (\mathbf{x}_1, -\mathbf{p}_1, \mathbf{x}_2, -\mathbf{p}_2, \dots, \mathbf{x}_N, -\mathbf{p}_N)$ (Lebowitz, 1999). If reversibility, so defined, is indeed a property of our microlaws,⁵ then the reversibility objection shows that the claim that BP holds for all systems at all times must be false.

To see why this is so, suppose that our system S is an energetically isolated glass containing a quantity of water. And suppose that the present macrocondition of the system is that it contains some ice and some lukewarm water. There is some region R of the phase space of that system that is compatible with a specific such macrocondition. And, if the postulate about statistics is true of that system at the present time, then it will certainly follow that it is overwhelmingly probable that at some time in the future, the glass will contain only water at some uniform temperature. So far, this sounds good. But since our microlaws are time symmetric, if P1 is true, then its inverse must be true too. It follows, therefore, by the very same reasoning, that if the postulate about statistics holds at the present time, then it is overwhelmingly likely that at some time in the past, the glass contained only water at some uniform temperature. But of course this retrodiction not only contradicts the second law of thermodynamics, it contradicts our everyday experiences. The only reasonable

5. Reversibility, so defined, certainly is a property of classical Hamiltonian mechanics. The idea, however, that it is necessarily a property of the relevant laws for undergirding statistical mechanics has important critics, including Albert himself in chapter 7 of (2001). See also a discussion of this possibility in Sklar (1995). For the purpose of this paper, I take the reversibility of the microlaws for granted.




Figure 2. Boltzmann's postulate. Given any system S in macrocondition A , the probability of the microcondition of S being in some tiny region r of R (the region compatible with A), is proportional to the volume of r (on the standard measure).

conclusion to draw from this is that, as plausible as it sounds, Boltzmann's postulate cannot possibly be a universally true fact about the world.

3. Two Proposals. The question is: what should we replace it with? What appears to be called for is some application of something like this postulate, but one which will somehow shatter the time symmetry that is a fundamental feature of our microscopic laws. This is where Albert's powerful insight comes in. The insight is as follows: Following Albert, call the regions of the phase space of a system that are compatible with that system's macroconditions " M -regions". P1 tells us that the volume of the subregion of any M -region that is occupied by microstates that lead to decreases in entropy toward the future, (the "abnormal subregions") are overwhelmingly small. Albert's insight is to remark that they are not only small, but also extremely scattered, "in unimaginably tiny clusters, more or less at random, all over the place" (82). To introduce a term, these abnormal regions are "fibrillated". And, of course, the same thing can be said of the subregion of any M -region which is taken up by microstates that, when the dynamics are run in reverse, lead to decreases in entropy towards the past. Even more important is the insight that "there is patently no reason at all that those two subregions of any particular M -region of the phase space of any particular thermodynamic system should have any tendency whatsoever to be aligned or to be correlated, or to be otherwise matched up with each other" (82).

In other words, Albert is here introducing a second principle:

Principle 2 (see Figure 3). If M is an M -region and M^F is the subregion of M which is taken up by the (abnormal) microstates that lead to lower entropy going forward, and M^B is the subregion of M taken up by the (abnormal) microstates that lead to lower entropy going backward then, because of their fibrillation, it is reasonable to conclude that

$$P(M^F|M, M^B) = P(M^F|M) \text{ and } P(M^B|M, M^F) = P(M^B|M)$$

(where $P(M^F|M, M^B)$ means the probability of M^F given M and M^B , etc.)

P2 has exactly the same status in Albert's argument as P1. It has neither been shown, nor is it being introduced as a postulate. Rather, it is being argued that the principle is plausible to believe, given what we know about our microlaws. The beauty of P2, for those who are willing to accept it, is that allows us to use a severely restricted form of BP, one that might help us to get the predictions we want without forcing us to make bad retrodictions.

Albert uses P2 to show that, in order to underwrite the time asymmetry




Figure 3. Principle 2. The abnormal regions M^F , which lead to higher entropy in the future, and the abnormal regions M^B , which lead to higher entropy in the past, are scattered more or less at random, all over the place. Moreover, they have no tendency to be aligned or coordinated. And so the probability, given that one is in one of the abnormal regions, of finding oneself in the intersection of the two sorts of regions, is still always very small

of the second law for a particular system, one need not make the obviously false claim that BP holds of the system at all time. One need only assume that the postulate held at one moment in the past. And so if we want BP to help us get things right in predicting that macroconditions will evolve forward in time in accordance with the second law, without helping us get to get things wrong in our retrodictions, then we had better somehow make it such that BP only holds at the beginnings of things.

Prima facie, there are two ways I might go about doing this—that is, there are two ways I might interpret what it means to say that the postulate holds *at the beginning of things*. To see why, imagine that we have a glass of ice water, sitting alone in the universe, at the beginning of time. If I encounter the glass at that moment, and I apply BP—that is, if I assume that there is a uniform probability distribution over all the microstates compatible with the macrostate of the glass at the beginning of time, then I can be sure that at some later time T , it is overwhelmingly likely that the ice water will be in a state of higher entropy.

If I encounter the glass at time T , what I do is that I assume that the glass was in its low entropy macrostate at the beginning of time, and that at that time, BP held. This is equivalent to assuming that, at time T , there is a uniform probability distribution over the set of states that are compatible both with the current macrostate of the ice-water, and with the fact that the ice-water began in the low entropy macrostate that it did. Since, by P2, the probability of being in a region of microstates that leads to higher entropy in the future is not correlated with being in a region of microstates that come from the low entropy original state of the ice-water, I can confidently predict both that the ice-water will be in a higher state of entropy in the future, and that it was in a lower state of entropy in the past.

But the universe does not consist entirely of one solitary glass of water. The universe as a whole is indeed one thermodynamic system, but it also contains many occasionally energetically isolated sub-systems that are of thermodynamic interest. So, there are, at least apparently, two different ways that we could make use of the apparent benefits of P2. One proposal, which Albert rejects, I will call the "branch systems" proposal. The other proposal, the one which Albert endorses, I will call the "big bang" proposal. For the discussion that follows, it will be helpful to outline not only Albert's proposal, but also the alternative, in order to draw some contrasts.⁶

6. It is not my goal in this paper to defend the branch systems proposal. I offer it here only as a foil for better understanding what Albert's proposal is and what he expects it to accomplish. For a defense of the branch system proposal, or more precisely, for an attempt to understand what a good defense of the branch systems proposal

3.1. The “Branch Systems” Proposal. The branch systems proposal is, in a sense, the simpler of the two. It goes something like this: If we can identify the moment at which all thermodynamically relevant systems come into being, then we need then only assume that the postulate about statistics holds at these, and only at these, moments. The idea would be that something like the act of preparing the system as an energetically isolated system brings it about that there is, as a matter of objective probability, a uniform probability distribution over the region of microstates that are compatible with whatever macrostate the system is in at the moment that it is prepared.

If P2 is true, then we can then expect that systems will evolve forwards in time towards equilibrium, but we are free from the worry that they should have been expected to have evolved from a higher entropy state in the past (relative to that moment) in virtue of the simple fact that they *did not exist* in the past. If we encounter an isolated system in a particular, say, medium entropy macrocondition, we can even retrodict that is overwhelmingly likely that the system has evolved from a lower entropy state, because, were anything else the case, it would have been overwhelmingly unlikely to have evolved into its present macrocondition. The problem appears to be solved.

But Albert rejects accounts of this type. He does so, as I understand him, for two reasons.

1. He thinks that there is no principled way to specify the precise moment at which a particular system has become energetically isolated—the moment, that is, at which we might say of it that it has been prepared (89).
2. He thinks it is entirely unnecessary, and indeed illegitimate, to postulate BP over and over again as an unexplained primitive for each branch system when he thinks he can get away with postulating it only once, at the beginning of the universe.

3.2. The “Big Bang” Proposal. This later idea, that we should apply BP only once—at the beginning of the universe, is what I call the “big bang” proposal. Here, the idea is to suppose that the universe began in a very low entropy state, and that, at that moment, a uniform probability distribution obtained over all the microstates compatible with that macrostate.

Equivalently, and a bit more precisely, the solution offered by the big

would likely require from us in the way of a metaphysical understanding of scientific laws, see Winsberg (forthcoming). For the original presentations of these ideas, see Reichenbach (1956) and Davies (1977).

bang proposal is to suppose that the relevant system is the universe as a whole, and that there is a uniform probability distribution, on the standard measure, over the region of microconditions that are compatible with the present macrocondition, but further restricted to those microconditions that are compatible with the "past hypothesis." The past hypothesis is the supposition that the universe began "in whatever particular low-entropy highly condensed big-bang sort of macro-condition it is that the normal inferential procedures of cosmology will eventually present to us" (Albert 2000, 96). From this, and using P1 and P2 in the same reasoning we applied to the glass of ice water, we can then predict that for any time T_i in the history of the universe up until the expiration of the relaxation time, the entropy of the universe will be greater than the entropy at time T_j , if and only if $j < i$. After that, the universe will remain in a state of maximum entropy.

So far so good; but now we will begin to run into problems. In addition to showing how we predict that the entropy of the universe as a whole will increase over time, we also need to show that we can predict how the entropy of every energetically isolated subsystem will evolve over time. Otherwise, we will have failed to "underwrite the actual content of our thermodynamic experience" (Albert 2000, 159).

Suppose, for example, I were to right now go down to my kitchen, grab my Coleman cooler, fill it half way with lukewarm water, dump in the contents of the ice tray from my freezer, and shut the lid. Presumably, I could then quite confidently predict that after a few hours time, if were to open up the lid, what I would find is cold water and no ice. The question is how these sorts of predictions can be underwritten.

4. Working with Isolated Sub-Systems. Let's be more careful.

Let S be the separation time, the time when you dump the ice and water in the cooler and shut the lid. We assume that after S , the ice water is an isolated system.

Let P be some time a little bit after shutting the lid, when the ice has just started melting.

Let T be 10 minutes after P when the ice is about half melted.

Let $T+$ be some time after the ice is all melted.

Suppose I come upon the (closed) cooler at time P and watch it for a while, so I know nothing is molesting whatever is in the cooler, and then at T , I open the cooler and see melting ice and cold water. I make the inference that at time P , there was more ice in the cooler than at T . I

also infer that if I close the cooler back up, more of the ice will melt by time T^+ .⁷

The trick now is to explain, statistically, why these are good inferences.

One possibility would be to use what we have already established about the universe as a whole, and rely on the fact that the average entropy of the universe is increasing from time P to time T^+ . But just because the average entropy of the entire universe is increasing from P to T^+ , this should give me no confidence that the entropy of my local system, which is tiny in comparison with the universe, will also increase in entropy. A decrease in the entropy of my cooler could easily be offset by a more than average increase in some other part of the universe.

So how will a defender of Albert's proposal try to underpin these inferences? Unfortunately, Albert does not spell out how this kind of example is supposed to work, in detail, in the book. Nevertheless, I think the following is a pretty good reconstruction of the reasoning Albert would hope to apply to this case. The reasoning would almost certainly have to go something like this:⁸

We know that at time S , the microstate of the universe occupies a very special region of the phase-space compatible with the universe's macrostate at S , namely the thin fibrillated region of states that are compatible with the past hypothesis. But we also know something else: that besides being constrained to be in this special region, there are no other constraints that restrict where in this fibrillated region the actual microstate is. All subregions of the fibrillated region are (so far as nature will let us know) equally likely microstates. More specifically, P2 tells me that almost all of this fibrillated region's microstates lead to greater entropy in the future, for the universe as a whole.

Our interest for the moment, however, is not the universe, but the cooler. So let's think about the reasoning that the proponent of the big bang hypothesis *must* expect us to use with regard to the ice water: At time S , no matter what the macrostate of the universe, the microstate of the universe is confined to the small fibrillated region that is compatible with the past hypothesis. Now restrict attention to the subspace of the universe's statespace that represents only the particles that will be trapped inside the cooler. The macrostate of the cooler will be confined to the region that is compatible with the past hypothesis. *Since it is overwhelming likely that the universe is in a region that will lead in the future to steadily higher entropy, it must be overwhelmingly likely that the microstate of the cooler is in some subregion of the fibrillation that will lead in the future to*

7. Thanks to an anonymous referee for suggesting this way of clarifying the cooler example.

8. Thanks again to the same referee for spelling out how this argument should go.

the cooler's steadily higher entropy. Conversely, it is overwhelmingly unlikely that the microstate of the cooler-contents is in the extremely small subregion that will lead to its having decreasing entropy from time P to time T . Thus, I know that the entropy at time P must be lower than at time T , which is in turn lower than at time $T+$.

This is the sort of argument that Albert must give. But clearly, the only way that any such argument is going to go through is if something like the kind of inference that occurs in the italicized sentence is a valid one. But what might assure us that such an inference is valid? We know that at time P , it is highly likely that the universe is in a region of microstates that leads to a steadily higher entropy. But how do we conclude that the cooler is highly likely to be in such a region? We can only reach this conclusion if we believe that we can move effortlessly from this property of the universe as a whole to a corresponding property of the cooler. More specifically, we could only reach this conclusion by applying another “principle”, call it Principle 3:

Suppose the universe is in some macrostate M . By application of the past hypothesis and BP, we know that the universe is the small region of microstates (call it M^*) compatible with M and with the past hypothesis, and that there is a uniform probability distribution over that small, fibrillated region.

Principle 3. Given any such M^* , if M_S^* is constructed by restricting that M^* to any subspace of M^* corresponding to a system made up of a subset of the particles in the universe, *then* (because there is a uniform probability distribution over M^*) there will also be a uniform probability distribution over M_S^* .

Only if we believe P3 can we make the inference in the italicized sentence.

5. The Status of Principle 3. Certainly, if P3 is true, then the argument above goes, and all is well with cooler, and all is well with Albert's proposal—the actual content of our thermodynamic experience has been underwritten.

So the question is: what is the status of P3? Clearly, P3 is not a simple theorem of measure theory. Since M_S^* is constructed by restricting that M^* to a subspace of M^* corresponding to vastly fewer particles than are represented in M^* , M_S^* has many fewer dimensions.⁹ Just because there is a uniform distribution of over some space, it does not follow mathe-

9. Recall that the phase space of a system has a number of dimensions equal to six times the number of particles it represents.

matically that there is a uniform distribution of every (lower-dimensional) subspace of that space.

But maybe, like principles P1 and P2, P3 can be motivated as a plausible consequence of the microlaws. If we were willing to accept P1 and P2 on the grounds that they had been “made plausible,” why not P3?

The problem is that we can easily see that P3, at least as written, is clearly false. There is a relatively simple argument to see that it must be. It follows: In the argument I attributed to Albert above, we applied P3 to the subspace representing the cooler at time S , and we got great results. But let’s see what happens if we apply the principle to the subspace representing the contents of the cooler at time T :

Recall that at time T , we know that the universe is in a microcondition that is compatible not only with its current macrocondition, but also with the past hypothesis. We also know that there is a uniform probability distribution over all such microconditions. From these facts, we can easily conclude that the cooler is in a microstate that is compatible with its own macrocondition and with the past hypothesis. P3 allows us to infer, furthermore, that there is a uniform probability distribution over those states. This is a big problem! Here is why:

Suppose for the sake of argument I were able to show that there was a uniform probability distribution over *all* of the microstates compatible with the macrostate of the cooler at time T (that is to say, not just with the ones that are also compatible with the past hypothesis.) Clearly, that would be a disaster, since this would enable me to conclude that it was overwhelmingly likely that, at time P , there was less ice in the cooler than there was at time T . And this is exactly the opposite of the inference I hope to underwrite!

“But hold on,” you say, “what we were able to show using P3 was not that there is a uniform distribution over the set of microconditions compatible with the macrocondition of the cooler at time T , but rather with that set restricted to those microstates compatible with the past hypothesis.”

But here is the rub: *these are the same two sets!* Restricting the set to those microconditions that are compatible with the past hypothesis does nothing because the cooler-content’s previous interaction with the rest of the universe effectively randomizes the microconfiguration of the cooler-content.

If I apply a uniform probability distribution over the microconditions compatible with the present macrocondition of the cooler, then, unless the cooler has been energetically isolated since the beginning of time, then adding the further requirement that these microconditions be compatible with the low entropy state of the beginning of the universe is not adding any substantive further requirement at all. The reason is this: in the time

between the beginning of the universe and the time when the cooler lid gets shut, outside influences from the rest of the universe have been free to interfere with the state transitions of the contents of the cooler in any way we might imagine. The dynamics is completely unconstrained. Consequently, any microstate that is compatible with the present state of the cooler is also one that is in principle compatible with the past hypothesis. Since the contents of the cooler were around long before the lid was closed, any of the possible current microstates of the contents of the cooler (those compatible with M) are dynamically compatible with any of the possible microstates of the contents of the cooler at the beginning of the universe.

In sum, if P3 holds, then it can be applied to the cooler at time T , and hence by the above argument, there is a uniform distribution over all microstates compatible with the macrostate of the cooler, and therefore, it is overwhelmingly likely that there was less ice in the cooler at P than there is at T . Clearly, not only is P3 poorly motivated, it is false!

6. Can Principle 3 Be Restricted? If P3 is demonstrably false, then the only reply left open to a defender of the big-bang proposal is to deny the *universal* applicability of P3, but try to hang on to some limited version of it. Looking back at the arguments that I claim Albert needs to rely on, it's clear that the arguments only rely on applying principle 3 at time S , and not at time T . But it's only applying P3 at time T that gets us into trouble. Why not reason, therefore that we are safe in assuming that there is some time which we might as well call S , when the contents were interacting with the rest of the environment, and that it was up to then, and only up to then that something like P3 applied?

Why not, in other words, suggest that P3 only applies to subspaces that represent subsystems at time when they are not energetically isolated? The problem is that this reply treads on very thin ice. It treads on especially thin ice for someone like Albert—someone who has argued against the branch systems proposal on the grounds that he has.

Recall his two principle objections to the branch systems proposal:

1. He thinks that there is no principled way to specify the precise moment at which a particular system has become energetically isolated—the moment, that is, at which we might say of it that it has been prepared (89).
2. He thinks it is entirely unnecessary, and perhaps illegitimate, to postulate over and over BP as an unexplained primitive for each branch system when he thinks he can get away with postulating it only once, at the beginning of the universe.

Let's apply the reasoning in these two objections to the proposed reply that principle 3 only applies to non-energetically isolated systems.

1. If there is no principled way to specify the precise moment at which a subsystem of the universe becomes energetically isolated then how could there be a physical principle that applies before that precise moment, and ceases to apply after it?
2. Remember what the intended nature of these “principles” (as opposed to “postulates”) is supposed to be. These are not assumptions that are intended to be added as unexplained primitives. That, according to Albert, is the second sin committed by the branch systems proposal. The principles are meant to be offered as claims that can be plausibly argued for as consequences of the dynamics—of the microlaws. But what argument could possibly make P3 seem plausible prior to S , and not so plausible after S —even if there were a principled way of saying precisely when S occurred? To suggest P3 should apply at S but not at T smacks heavily of using P3 as an unexplained primitive, not as a plausibly motivated principle. And if it is being used as an unexplained primitive, then the big bang proposal simply reduces to the branch systems proposal.

In sum, there is no way that Albert can avail himself of P3 at time S and not also be forced to apply it at time T without committing precisely the same sins as those of which he accuses proponents of the branch systems proposal. So P3 is false, and using a suitably restricted form of P3 is illegitimate. But without P3, applying BP only at the beginning of the universe will not allow us to make the standard predictions about ordinary sized thermodynamic systems. One application of the BP, conditionalized on the past hypothesis, will not “underwrite the actual content of our thermodynamic experience.”

REFERENCES

- Albert, David Z. (2000), *Time and Chance*. Cambridge, MA: Harvard University Press.
- Callender, Craig (2001), “Taking Thermodynamics too Seriously”, *Studies in History and Philosophy of Modern Physics* 32: 539–553.
- Davies, Paul C. W. (1977), *The Physics of Time Asymmetry*. Berkeley: University of California Press.
- Feynman, Richard (1967), *The Character of Physical Law*. Boston: MIT Press.
- Lebowitz, J. L. (1999), “Statistical Mechanics: A Selective Review of Two Central Issues”, *Reviews of Modern Physics* 71: S346–57.
- Reichenbach, Hans (1956), *The Direction of Time*. Berkeley: University of California Press.
- Sklar, Lawrence (1995), *Physics and Chance: Philosophical Issues in the Foundations of Statistical Mechanics*. London: Cambridge University Press.
- Winsberg, Eric (forthcoming), “Laws and Statistical Mechanics”, *Philosophy of Science*.