Laws and Statistical Mechanics

Eric Winsberg†

This paper explores some connections between competing conceptions of scientific laws on the one hand, and a problem in the foundations of statistical mechanics on the other. I examine two proposals for understanding the time asymmetry of thermodynamic phenomenal: David Albert's recent proposal and a proposal that I outline based on Hans Reichenbach's "branch systems". I sketch an argument against the former, and mount a defense of the latter by showing how to accommodate statistical mechanics to recent developments in the philosophy of scientific laws.

1. Introduction. This paper will explore connections between competing conceptions of scientific laws on the one hand, and a puzzle in the foundations of statistical mechanics on the other. The puzzle is to reconcile the asymmetry of the second law of thermodynamics with the evident time symmetry of the laws of microscopic dynamics. Many solutions to this puzzle have been proposed, but one, a proposal based on Hans Reichenbach's "branch systems," has received insufficient attention because it is commonly thought to be inadequate.

In what follows, I try to salvage and reconstruct the branch-systems proposal. I demonstrate that there is a particular conception of scientific laws that motivates the arguments used to dismiss this proposal. I argue that the motivation for rejecting the branch-systems proposal out of hand is entirely grounded in what might be called the "possible worlds" conception of scientific laws. This is the view that laws are universally valid, and delimit the range of possible worlds. But there is an alternative view. On what I call the "framework" view of laws, the motivation for dismissing the branch-systems proposal disappears, and it becomes a viable solution to an important foundational problem. I take this to be grounds for a renewed interest in so-called "branch systems," as well as evidence for the strength of the framework view of laws.

†To contact the author, please write to: Department of Philosophy, University of South Florida, 4202 East Fowler Ave, FAO 226, Tampa, FL, 33620; e-mail: winsberg@chuma1.cas.usf.edu.

Philosophy of Science, 71 (December 2004) pp. 707–718. 0031-8248/2004/7105-0005$10.00
Copyright 2004 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/425234 Published online by Cambridge University Press

2. A Foundational Problem. The foundational problem I am concerned with is well known, and it can be articulated quite simply in general terms: On the one hand we have the laws of thermodynamics, which tell us that all macroscopic systems of a certain kind behave asymmetrically in time. On the other hand we have the fact that each of our candidate laws for describing the behavior of the microscopic constituents of those systems is time symmetric. How are we to reconcile this seeming incompatibility?

Suppose we have a thermodynamic system S, and set of laws M that govern the microscopic components of S. Thermodynamics tells us that if S is in a macrocondition A, then it will evolve over some period of time T into macrocondition B. To explain this, it is standard to begin by arguing that almost all of the volume of the region of phase space that is associated with A is occupied by microconditions that, when subjected to the laws M, will evolve over the period of time T into a region of phase space associated with the macrocondition B.[1] While this conclusion has never been shown to be true, it is safe to say that it has been made at least plausible by arguments going back to Boltzmann and Gibbs.

But that is only the first step. 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 volume of r—call this supposition "Boltzmann's postulate—then we could easily make the central prediction expected of statistical mechanics: that it is overwhelmingly probable that systems will evolve towards conditions of maximum entropy.

Since 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.

Unfortunately, all we have accomplished so far is to have put ourselves in a position to better articulate the real source of original problem, the problem pointed at by what is frequently referred to as the "reversibility objection." The problem is that if the second law of thermodynamics is a true law about the world, and our laws about the microworld are all time reversible, then the reversibility objection shows that the claim that Boltzmann's postulate holds for all systems at all times must be false.

1. This outline closely follows Albert 2000.

https://doi.org/10.1086/425234 Published online by Cambridge University Press

Suppose, for example, that our system S is an isolated glass containing water. Suppose that the present macrocondition of the system is that it contains some ice and some water. There is some region R of the phase space of that system that is compatible with a specific such macrocondition. And, if Boltzmann's postulate is true of that system at the present time, then it will certainly follow that it is overwhelming 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, it also follows, by the very same reasoning, that if Boltzmann's postulate 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. This retrodiction not only contradicts the second law of thermodynamics, it contradicts our everyday experiences. The only reasonable conclusion to draw from this is that, as plausible as it sounds, Boltzmann's postulate cannot possibly be true at all times.

3. Branch Systems. So what should we replace it with? David Albert (2000) has shown that all we need is an application of something like Boltzmann's postulate, but one that only holds at the beginnings of things. This, he makes plausible,[2] would get things right in predicting that macroconditions will evolve forward in time in accordance with the second law, without getting things wrong in our retrodictions. But here is the rub. How exactly can we suppose that to occur?

Here is one proposal. Inspired by Hans Reichenbach in *The Direction of Time* (1956), it is better articulated by Paul Davies. The idea is that

> all of our observations and experiments in thermodynamics refer to what Reichenbach calls *branch systems*, rather than permanently isolated systems. Branch systems are regions of the world which separate off from the main environment and exist thereafter as quasi-isolated systems, and usually merge once again with the wider environment after a sufficient time. Examples of this sort are countless, but one will suffice to remove any misunderstanding. When we take an ice cube, add it to a lukewarm drink, and watch the ice cube melt, the system + drink only comes into being *after* this event. It simply did not exist as a quasi-closed system before this event. (Davies 1977, 69)

2. The argument that the Boltzmann postulate need only be applied at the beginning of things is a bit complicated for the space available here. It relies centrally on the assumption that the abnormal subregions of a macroregion that result in antithermodynamic behavior are fibrillated, and randomly distributed. For more details, see Albert 2000, chapter 4. See also my discussion of what I call "Principle 2" in Winsberg 2004.

https://doi.org/10.1086/425234 Published online by Cambridge University Press

If we can identify the moment at which all thermodynamically relevant systems come into being, we need then only assume that Boltzmann's postulate holds at these, and only at these, moments. 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 equilibrium in the past in virtue of the simple fact that they *did not exist* in the past. The problem appears to be solved.

The branch-systems proposal has not been greeted with great acclaim in the literature. Presently, I will defend it against some criticisms. But first, I want to consider two arguments against versions of the branch-systems proposal, which I think are on the mark and which need to be accommodated. No defense should be mounted against these criticisms. Rather, I want to offer a version of the proposal that is immune to these criticisms.[3] The first of these critical remarks comes from Albert. Albert thinks that Davies *must* have an *epistemic* understanding of the statistical-mechanical probabilities involved (Albert 2000, 89n). That is, on Albert's construal of Davies' proposal, the claim that Boltzmann's postulate holds at a certain point in time is a fact about our collective ignorance of the exact state of affairs that obtain at that point in time. According to this picture, moreover, the claim is supposed to be "innocent" (Albert's word), that is, essentially without empirical content. It's supposed to be motivated simply by reasoning about we mean when we say that we know that a system is some macrostate. Albert has a very good argument that the epistemic argument provides inadequate justification for holding to the postulate (see Albert 2000, 64). Furthermore, if it is the case that Davies thinks it is my ignorance of the exact microconditions which somehow justifies me invoking Boltzmann's postulate, if the claim is "innocent," then it is a fair question to ask why this argument should only work at the times when these branch systems come into being, rather than the middles, or the ends (Albert 2000, 89n). And this, as we have seen, would be a disaster.

Since I take it that these are both good arguments, I in turn take it that it is a nonstarter to think that epistemic probabilities have anything at all to do with the proposal we should investigate. Thus, our game had better be to suppose that it is an empirical and contingent fact about the world that when branch systems come into being, the *objective* probability of their being in a particular region of phase space is proportional to the volume of space occupied by that region.

The second critical remark comes from Larry Sklar (1993). Sklar imagines that Davies and Reichenbach are playing the game of trying to ground

3. It is not my goal in this paper to provide an exegesis of what Reichenbach or Davies actually had in mind, but rather to articulate what I think is the most sensible position.

https://doi.org/10.1086/425234 Published online by Cambridge University Press

the very direction of time in terms of these considerations about branch systems. Sklar rightly argues that the posit about branch systems can not possibly accomplish this much. To accomplish that, we would need to sneak in two hidden assumptions: (1) that all branch systems are exhibiting time evolution in the same direction—presumably because (2) this direction of evolution is the same as that of causal connection (Sklar 1993, 328). Sklar notes that to make these two assumptions would be to illicitly smuggle a direction of time, the very thing we were supposedly trying to derive.

> The initial low entropy of a branch system is explained by its having been cut off from the main system. That is how it was "caused" to be brought into being . . . . But the final state, even if of low entropy, can't be explained by reference to the forthcoming reabsorption. (Sklar 1993, 328)

Exactly right. But of course the principal reason that the reabsorption can't cause the system to have any final state, or any initial state, or any state in between, is because the reabsorption occurs temporally after all of the states of the isolated branch system. As Sklar quite rightly points out, if what we are after is nothing less than an explanation of the origin of asymmetries in time, then there doesn't seem to be any non-question-begging way to explain why coming-into-beings should be at all privileged as compared to reabsorptions.

But if we are after significantly smaller game, that is, if all we are interested in doing is explaining how time-symmetric microlaws can be compatible with time-asymmetric macrolaws, then we are not doing anything question begging at all. We can readily admit that we are helping ourselves to some other arrow that is already out there (for example, in the causal structure). All we need to do is make the perfectly sensible and pragmatic assumption that time does indeed have a privileged direction, and that this is the direction in which causal influences propagate. We don't smuggle this assumption in illicitly, hoping that no one will notice. We simply state that just because all our candidate microlaws are time symmetric, this is no reason to deny the perfectly sensible claim that there is a privileged direction to time and that causal relations obey it. The asymmetry of the macrolaws thus does not appear mysteriously at all. It comes, as Sklar rightly points out, in part as a consequence of a basic assumption about causal orderings and temporality.

To review then, the proposal we should investigate after having reviewed the two forgoing critical points is the following: Whenever a portion of the world separates off from the main environment and becomes energetically isolated, we postulate that, as an empirical hypothesis, whatever macrocondition that system is in, it will always be in a subregion of the

https://doi.org/10.1086/425234 Published online by Cambridge University Press

region of microscopic phase space compatible with that macrocondition with an objective probability in proportion to the volume of that subregion.

4. Problems for Branch Systems. It is clear that *this* proposal about branch systems is immune to the two criticisms I have outlined above. They are not even directed at it. There are other criticisms, however, that cannot be dodged so easily. Albert, in particular, has some worries about how (and why), precisely, this proposal is to be carried out (2000, 88–89). He asks the following questions:

1. How are we to decide at exactly what moment a branch system comes into being?
2. Suppose our branch system is a glass with ice-water in it; why should we focus on the glass, and not the room in which the glass sits, or the building, or the city, etc.?
3. "Why in God's name *bother* with all this, when the uniform probability-distribution over the possible microconditions compatible with the macrocondition of the *world*, at the moment when *it* came into being, will very straightforwardly give us everything we need?" (Albert 2000, 89)

These are all good questions. I will try to answer them all. Clearly, however, it makes sense to start with the last question first. If indeed there is a straightforward way to get everything we need merely by assuming that the statistical postulate holds at the beginning of the universe (lets call this the "big-bang proposal"), then making the postulation over and over at the birth of each branch system would be, at best, superfluous. At the very least, Ockham's Razor would counsel strongly against it. Alas, recent reports in the literature of the death of this problem in the foundations of statistical mechanics have been somewhat exaggerated.[4] More precisely, if we attend carefully to exactly what should be included in the phrase "everything that we need," it will become difficult to see how in the world we can get all of that merely by postulating that there is a uniform probability distribution, on the standard measure, over those regions of phase space that are compatible with the macrocondition at the beginning of the world.

5. The Beginning of the Universe. Let's see how the big-bang proposal is supposed to work. Recall the problem. Suppose the universe at time T_0

4. There are a number of sources that suggest that what I call the "big-bang proposal" has closed the case on this issue. In addition to Albert 2000, see Callender 2001 and Lebowitz 1999.

https://doi.org/10.1086/425234 Published online by Cambridge University Press

is in macrocondition M_0 which is of less than maximum entropy. We wish to explain our ability to predict that at any future time T_n, the macrocondition of the system, M_n, will be of a higher entropy than M_0. The original Boltzmannian suggestion was to suppose that his postulate held at all times, including T_0. This will work. Of course, as we have seen, it will also have the unfortunate consequence of forcing us to predict that at all times earlier than T_0, the universe also had a higher entropy then it does at T_0. So, unless T_0 is the earliest instant in the history of the universe, this will not do. The solution offered by the big-bang proposal is to suppose that the Boltzmann postulate only holds at that one instant. Equivalently, the big-bang proposal supposes 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 macrocondition that held at the beginning of the universe.

We then add to the above what Albert calls the "past hypothesis." This 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"(96). From this we can predict (statistically) that for any time T_i in the history of the universe, the entropy of the universe will be greater than the entropy at time T_j, if and only if $j < i$ up until the moment when the universe reaches an equilibrium condition.

This much of the big-bang proposal may very well work. But if we hope to recover all of the predictions of thermodynamics from statistical mechanics, then 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 how we can predict that the entropy of every energetically isolated system will evolve over time.

Recall that it is precisely on this point that the big-bang proposal and the branch-systems proposal differ. The big-bang proposal wants to apply the statistical postulate only at the beginning of the universe, whereas the branch-systems proposal insists on applying it at the beginning of the life of every energetically isolated system. Thus, if Albert is going to claim that the big-bang proposal can do everything that the branch-systems proposal can do, then it would behoove him to explain how the big-bang proposal can predict the approach to equilibrium of individual isolated systems. Unfortunately, he does no such thing anywhere in the book. In fact, it is not obvious from reading the book what he thinks he has in mind.

Elsewhere, I argue for a stronger claim than that there is a missing argument. I argue, in fact, that the claim is false. It is false that by positing, at the earliest instant in the universe, a uniform probability distribution,

https://doi.org/10.1086/425234 Published online by Cambridge University Press

on the standard measure, over the microconditions compatible with the "past hypothesis," one can predict that individual isolated systems will behave in accordance with the second law. Here I do not have the space to offer more than a brief sketch of the argument. I would direct readers who would like to see the details of the argument to consult the other paper (Winsberg 2004).

Here is the sketch: Since Albert rejects the proposal that one should put the Boltzmann postulate in by hand at the beginning of the life of each isolated system, he clearly thinks that something akin to it can be gotten for free for each such system. But if it is plausible that I can get the postulate (conditioned on the "past hypothesis) for free at the beginning of the life of an energetic system, then it should be plausible that I can get it for free at any other time. Thus, he seems committed to believe that, for each such isolated system at every moment in its life, I can get for free that there is a uniform probability distribution over the region of microstates that is compatible both with the macrocondition of that system, and with the past hypothesis. But that simply cannot be the case.

Here is why: If there is a uniform probability distribution, on the standard measure, over the microconditions compatible with the present macrocondition of, say, a Coleman cooler, then, unless the cooler has been energetically isolated since the beginning of time, 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. 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 in any way we might imagine. And so the dynamics of what will end up being the contents of the cooler is completely unconstrained during the period prior to the lid being closed. 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.

Hence, adding the requirement that the microstates of the cooler be compatible with the initial macrostate of the universe doesn't screen off *any* of the possible microstates. It doesn't at all narrow down the volume of state space over which we apply the uniform distribution. And hence, the uniform probability distribution applied over this "restricted" volume of state space is precisely the same as the uniform probability distribution applied over the states compatible with the present macrocondition. And *that* distribution, as we have seen, forces us to retrodict in ways that are incompatible with the second law. Albert can't get for free anything akin to what the branch-systems proposal puts in by hand.

So, the branch-systems proposal seems to explain not only why the entropy of the universe is increasing, but also why the ice in my cooler

https://doi.org/10.1086/425234 Published online by Cambridge University Press

will melt and has been melting, and the big-bang proposal doesn't. And *that* is why we should bother with the branch-systems proposal.

6. The Autonomy of Laws. Let's now talk about those of Albert's worries that are specific to the branch-systems proposal. The first is that we would somehow have to make a difficult choice as to what would count as a branch system, say for example, the Coleman cooler in my kitchen, or the kitchen itself, or my house, or the state of Florida. Never mind that none of the latter are even close to energetically isolated, and suppose for a moment that I lived in a quasi-energetically-isolated house. There are no choices to be made. The house is a branch system with respect to the outside world, and the cooler is a branch system with respect to the house. On the branch-systems proposal, at the moment that they become energetically isolated both take on microconditions according to the probability distribution given by Boltzmann's postulate. I don't see why this should be a problem.

Recall, then, that the first question was this: How are we to decide at exactly what moment a branch system comes into being? From a methodological point of view, this worry seems to pack no punch at all. From the point of view of doing our physics, that is to say from the point of view of building models that are accurate enough for our needs, it seems entirely unnecessary to be able to pinpoint the precise moment at which a branch system comes into being. As in most modeling situations, an approximate value will do. At the risk of cartooning, the question that Albert really seems to want to ask is "How does the universe know the precise moment at which a branch system comes into being so that *it* can know to apply the Boltzmann postulate?"

So clearly, more is at stake in Albert's first two questions than simple methodological worries. When the question, "How are we to decide at exactly what moment a branch system comes into being?", is asked, we can hear the clear reverberations of a long-standing question from the foundations of quantum mechanics: "How does the universe know what counts as a measurement?" One way of thinking about the measurement problem, after all, is that quantum mechanics provides us with two different laws of evolution, the Schrödinger equation of state evolution and the projection postulate, and there doesn't seem to be any *principled* basis for choosing which one to apply. The worry here is very similar: there doesn't seem to be any principled way of deciding when the microlaws hold and the Boltzmann postulate will be applied.

Boltzmann's postulate, after all, is a postulate about microconditions. But once the microconditions for the world are specified for some particular time (say, at the beginning of the universe), isn't it the microlaws that are supposed to fix the microconditions of all states of affairs at all

https://doi.org/10.1086/425234 Published online by Cambridge University Press

future times? Isn't that what microlaws are supposed to do? Where does the branch-systems proposal get the temerity, then, to flagrantly disregard the clear and present authority of the microlaws and simply stipulate that the microconditions (or, at least, the probability distribution of a set of possible microconditions) *just are* such and such, as an *objective, empirical and contingent fact* about the world. But this worry, like all worries, has a presupposition.

7. Frameworks and Possible Worlds. The presupposition behind this worry is a supposition about the very nature of scientific laws. The supposition is that laws must give us necessary, universally applicable, complete, and exact descriptions of the way the world will necessarily be, given any physically possible set of initial conditions. There are many different ways in which philosophers have tried to spell out this kind of picture in detail. I will call the core of all such views about laws the "possible worlds" account of laws to emphasize the fact that on this picture, laws have universal (worldly) scope and that they, in conjunction with a set of physically possible initial conditions, delimit a set of possible ways that the world could present itself through the entire course of its temporal evolution.

There is, however, a strong, viable, and more pragmatic alternative to this conception of laws: call it the "framework" conception (Cartwright 1983, 1999; Giere 1988, 1999). On this conception, laws provide a framework for building models, schematizing experiments, and representing phenomena. Laws, moreover, have very broad, but not universal, domains of application. Rather than taking laws to be *universally true* and delimiting the character of all possible worlds, the proponent of the framework conception takes laws to be *broadly reliable* for a wide array of practical and epistemic tasks.

Can the final worry about the branch-systems proposal be motivated on the framework conception of laws? I don't think it can. If the worry is about the authority of the microlaws, the proponent of the framework conception can simply reply that it is a mistake to think that the microlaws need offer us a complete description of the universe. In fact, the best evidence from statistical mechanics is that our time-reversible microlaws are not complete in this sense. They need to be supplemented by Boltzmann's postulate whenever we want to model an energetically isolated system.[5] But this comes as no surprise to us framework types. We know that laws are for building models, and that all models have initial con-

5. Compare this proposal to Frisch's (2000) account of the origin of time asymmetry in electromagnetic radiation.

https://doi.org/10.1086/425234 Published online by Cambridge University Press

ditions and boundary conditions. Sometimes we specify these conditions explicitly, sometimes statistically.

This suggestion about how to think about the branch-systems proposal is closely aligned with Cartwright's approach to the measurement problem. Cartwright argues that since none of our physical models get things exactly right, it's just fine to rely on one model which gives us probabilities for outcomes, classically described, and then, when we find one of those outcomes, redescribe it again quantum mechanically. From this perspective, as long as we get a model accurate enough for our needs, the details of where we draw the line don't matter.

The possible worlds view demands that, though it may well be out of human reach, there is *one* model that accommodates all phenomena. In contrast the framework view endorses the logically weaker quantifier ordering. For each phenomenon there is at least one model that accommodates that phenomenon.[6] Given any of the possible energetically isolated systems of our scientific experience, we can provide a model in terms of the microlaws that perfectly explains the system's thermodynamic behavior. The possible worlds view demands of laws that they provide one model of all possible phenomena. In the domain of statistical mechanics, the world seems to resist this demand.

Once we dispense with the universal quantifiers and other such baggage in our conception of scientific laws, the worries that stand in the way of simple solutions to thermodynamic puzzles lose their bite. The possible worlds conception of laws leaves us with some troublingly intractable puzzles in the foundations of statistical mechanics. The framework conception, conversely, opens the doors to a relatively simple, straightforward solution to these puzzles. So much the worse for the possible worlds conception.

REFERENCES

Albert, David (2000), *Time and Chance*. Cambridge: Harvard University Press.
Callender, Craig (2001), "Taking Thermodynamics Too Seriously", *Studies in History and Philosophy of Modern Physics* 32(4): 539–553.
Cartwright, Nancy (1999), *The Dappled World: A Study of the Boundaries of Science*. Cambridge and New York: Cambridge University Press.
——— (1983), *How the Laws of Physics Lie.* Oxford: Oxford University Press.
Davies, P. C. W. (1977), *The Physics of Time Asymmetry*. Berkeley: University of California Press.
Frisch, Mathias (2000), "(Dis-)Solving the Puzzle of the Arrow of Radiation", *British Journal for the Philosophy of Science* 51: 381–410.
Giere, Ronald N. (1988), *Explaining Science: A Cognitive Approach.* Chicago: University of Chicago Press.

6. This discussion in terms of the order of quantifiers deliberately borrows from Paul Teller 2001, 143–144.

https://doi.org/10.1086/425234 Published online by Cambridge University Press

——— (1999), *Science without Laws*. Chicago: University of Chicago Press.
Lebowitz, Joel L. (1999), "Statistical Mechanics: A Selective Review of Two Central Issues", *Reviews of Modern Physics* 71: S346–357.
Reichenbach, Hans (1956), *The Direction of Time*. Berkeley: University of California Press.
Sklar, Lawrence (1993), *Physics and Chance: Philosophical Issues in the Foundations of Statistical Mechanics*. Cambridge and New York: Cambridge University Press.
Teller, Paul (2001), "Whither Constructive Empiricism", *Philosophical Studies* 106: 123–150.
Winsberg, Eric (2004), "Can Conditioning on the 'Past Hypothesis' Militate against the Reversibility Objections?" *Philosophy of Science* 71(4): 489–504.

https://doi.org/10.1086/425234 Published online by Cambridge University Press