The Metaphysical Foundations of Statistical Mechanics: on the status of PROB and PH

Eric Winsberg

November 21, 2010

What if some day or night a demon were to steal after you into your loneliest loneliness and say to you: "This life as you now live it and have lived it, you will have to live once more and innumerable times more" ... Would you not throw yourself down and gnash your teeth and curse the demon who spoke thus? Or have you once experienced a tremendous moment when you would have answered him: "You are a god and never have I heard anything more divine." [Nietzsche, *The Gay Science*, §341]

Contents

1 Introduction

One of the central aims of *Time and Chance* (2000) was to bring all of our experience of the macroscopic world under the umbrella of a small package of fundamental physical postulates. For many of the central features of our experience–particularly the especially puzzling temporally irreversible ones: like the thermodynamic behavior of isolated systems, the asymmetry of knowledge of the past and future, and the temporal asymmetry of causal interventions–the goal of *Time and Chance* was to show how these features of our experience follow from a package of laws consisting of a set of time-reversible, deterministic, dynamical microlaws; a hypothesis about the macroscopic state of the early universe (the "past hypothesis"; PH); and a statistical postulate about the distribution of microstates compatible with that macrostate (PROB).

Other than the claim that they are to be understood as laws, however, *Time and Chance* did not have very much to say about the metaphysical underpinnings of PH and PROB. This was despite the fact that PH and PROB appear to be somewhat unusual from the point of view of most ordinary understandings of laws of nature. As for the probabilities postulated by PROB, the book contains some brief arguments against construing them as degrees of belief, but for the most part, it is relatively quiet about a potentially puzzling question: how are we to understand the claim that the initial physical state of the universe involves objective chances? There is, after all, only one universe, and it began only once.

In subsequent developments of David Albert's program by Albert and Barry Loewer, however, the scientific program of *Time and Chance* has been given a more rigorous metaphysical underpinning. They have done this by embedding Albert's Boltzmannian scientific project into a Humean metaphysical project; one where PH and PROB have taken on the status of fundamental laws of a so called "best system" variety associated with John S. Mill, Frank Ramsey and David Lewis, and where the probabilities in PROB are taken to be objective chances of the kind that appear in the extension of the project by Lewis.

In this chapter, I want to attempt to clarify Albert and Loewer's metaphysical project, and subject it to critical scrutiny.[1] The metaphysical project is based, in part, on David Lewis' "Humean" conception of laws and chances; but substantial changes to Lewis' system are required in order to try to make it work for statistical mechanics. One of the first things I will want to do, therefore, is to try to clarify exactly what these modifications need to consist in.

Once it is clear what the modifications need to be, and we have a clear view of what Albert and Loewer's version of Humean laws and chances looks like, we can begin to ask how attractive this metaphysical project looks from a variety of points view. In particular, I will be interested in examining how attractive it looks from the point of view of underwriting what I view as a particularly strong form of "fundamentalism" that Albert and Loewer endorse, and which partially motivates the project in the first place. More will need to be said in due course, but the strong version of fundamentalism I have in mind insists not only that the truths of the special sciences be recoverable from the fundamental laws, but that lawfulness of their laws must be as well. Ultimately, what I will argue is that the metaphysical picture offered by Albert and Loewer is unsuccessful, and that the version of fundamentalism that motivates it is too strong.

2 The basic Humean picture

Let us begin by giving a first pass at the metaphysical picture. How are PROB and PH, which, intuitively, inform us about the initial conditions of the universe, meant to be understood as laws? According to the basic picture of laws offered to us by Lewis (1986, 1994), what is fundamental in the world is the Humean mosaic: a distribution of points in space-time with their natural intrinsic properties. The "laws" of such a world are whatever regularities occur in the axiomatic system that "best" summarizes the facts about the mosaic. The "best system" is the system that maximizes simplicity, informativeness, and fit. Whatever system of axioms, in other words, that offers the best compromise between being simple, being informative about the mosaic, and assigning the highest probability to the actual course of events ("fit"), is the best system. Its regularities are the laws of the world. And any probabilities that are featured in those laws are the objective chances.

Unfortunately, Lewis' conceptions of laws and chances, in the precise form he originally formulated it, is ill-suited to understanding PH and PROB. For one thing, neither PH nor PROB are regularities. More importantly, perhaps, Lewis did not paint a nuanced enough picture of the role of initial conditions in assessing the informativeness and simplicity of a "system" for his account to straighforwardly underwrite the claim that PH and PROB are laws. PH and PROB are, after all, coarse-grained summaries of the initial conditions of the universe.

Albert often illustrates that point by asking us to imagine a scenario in which a physicist is having a conversation with God. The physi-

[1] In what follows, I will take the success of what I am calling the scientific project for granted. I take it for granted, that is, PH and PROB, however we are meant to construe them, are sufficient, in conjuction with whatever the dynamical laws are, to underwrite, at least, the thermodynamic behavior of isolated systems. Doubts about this have been expressed in Winsberg (2004), Earman (2006), Callender (forthcoming), and elsewhere; but I will be ignoring those here.

cist, quite naturally, asks God to tell her what the universe is like. God begins by telling the physicist what the dynamic laws of the universe are. But of course this is not enough. Because she wants to be able to make predictions, the physicist asks to know in what particular initial condition the universe began. To which God replies: "The universe has 1050 particles, how much time do you have?" Since the answer is obviously "not enough", God suggest the following alternative: "Pretend as if there were a uniform probability measure over the set of microconditions that are compatible with whatever macrocondition in which the universe began. From that fact and whatever macroscopic observations you make, you should have enough information to make any practical macroscopic predictions that you like. That should save us some time." The trick to having PH and PROB come out as laws, in other words, is to imagine that a succinct summary of the initial conditions of the universe ought to be regarded as part of the best system. This idea, though, for whatever it is worth, was not on Lewis' radar. Indeed, it is not entirely clear that Lewis thought through the role initial conditions ought to play in his game of best system analysis carefully enough for us to evaluate the claim that PH and PROB should be part of one.

3 Measuring informativeness and simplicity: the role of initial conditions

So let us begin to flesh out the ways in which Albert and Loewer's conception of a best system analysis might differ from Lewis' by discussing the role that initial conditions might be imagined to play in evaluating the informativeness and simplicity of a system. Informativeness, obviously, involves the comparison of two things: the system, on the one hand, and the Humean mosaic—the distribution of fundamental facts in the world whatever they happen to be—on the other. The question, then, is what role the sorts of things that we–or at least working scientists–ordinarily call initial conditions play in such a comparison. There are three ways we can imagine this could go.

1. We could imagine that in evaluating the informativeness of a system, what we ordinarily think of as initial conditions can be appealed to in measuring the informativeness of the system, but they are not part of the system for the purposes of measuring the system's simplicity—and importantly, do not count as one of its laws. Imagine, for example, a world consisting of hard spheres bouncing around in what we would ordinarily think of as in accord with the laws of Newtonian mechanics. Think, then, of two books. In one book would be recorded the distribution of fundamental facts about the world—in this case, presumably, the distribution of the spheres across space-time. In the second book would be written the laws of Newtonian mechanics (the laws of motion and whatever force laws "govern" the interaction of the spheres.) But nothing about the positions and velocities of the spheres at any particular time would be recorded in the second book. On this conception of informativeness, then, the informativeness of the second book is measured by how informative it is about about what is written in the first book, given some specification of the positions and velocities of the spheres at some instant of time. The initial conditions get plugged into the calculations used when evaluating the informativeness of the second book, but they do not count against its simplicity.

2. As a second possibility, we could imagine that when evaluating the informativeness of a system, all we do is compare the first book with the second book, straight up. In such a conception, the informativeness of the second book in the example above would be rather low–if it has any at all. But by adding some description of the initial state of the

world to the second book, we would make it much more informative—while decreasing, obviously, its simplicity.

Presumably, on this second conception, whether or not the second book contains "initial conditions" is a matter of balance, and whether or not they get included in the version of book two that ends up being the "best system" will depend on whether the reduced simplicity that comes from adding them does or does not balance out the increased informativeness. Of course, it may turn out that a book without any "initial conditions" is never adequately informative, and so it may turn out that any best system, on this conception, will always include what we usually think of as "initial conditions". Or it may turn out that partial description of the intial coditions (something like in the God story above) wins out. It may even turn out that, for the real world with all its complexity, the inclusion of a precise specification of initial conditions never counts as adequately "simple." More on this later.

3. A third possibility is that we think of informativeness in a much more pragmatic way. Rather than thinking of informativeness as a direct, God's eye, comparison between the system and the human mosaic, informativeness is something that is *for us*. So, if I want to compare the informativeness of two competing systems, what I do is I try to use each of those two systems, *in combination with other facts I happen to know*, in order to make predictions about things that I can observe. And then I decide which of the two systems is more useful to me in making those sorts of prediction. The one that, in combination with other things I know, enables me to make better predictions, is the the more informative system. If system one makes better such predictions than system two without a sacrifice in simplicity, than it is the "better" system.

Here's one thing that should be clear: on the first conception, initial conditions will never be part of the package. On the second and third conceptions, it will be an interesting and open question whether or not they are, and what form they should take (i.e., should they they include precise specifications, or coarse grained descriptions like PH and PROB). It should also be clear that on the third conception (unlike the other two), how well a system rates for informativeness will depend on the epistemic agent using it. That is because on the third conception, it might very well turn out that a system without initial conditions is very informative to an agent who already knows certain other facts, but much less informative to another soft of agent. How informative a system is, on the third conception, will depend, to a great extent, on whatever other things you happen to know about the world, on what you are hoping to predict, and on other related factors.

Having discussed each of these three ways of thinking about informativeness, we should now ask: when Loewer says "Adding PH and PROB to the dynamical laws results in a system that is only a little less simple but is vastly more informative than is the system consisting only of the dynamical laws" (2007, p. 305) which conception does he have in mind? More importantly, on which, if any, of three conceptions above does this claim turn to be plausible?

Obviously, the first conception is a non-starter. On the first conception, you get to use the precise initial conditions for free when evaluating a system that does not contain them. So, on that conception, adding information about the initial conditions to a system will never increase its informativeness.[2] But what about the second and

[2] That is one way of putting the point. A virtually equivalent thing to say is that a precise specification of the initial conditions is always part of the package. These are virtually equivalent because in either case, no two systems will ever differ in either their informativeness or their simplicity as a result of a difference with respect to this part of the package. Regardless of how we characterize the point, PROB and PH will not come out as part of the system since they would be redundant given a precise specification of the initial conditions of the universe.

third conception? What I want to argue here is that, for various reasons, neither the second nor the third conception support the idea that PROB should be considered a fundamental law in the way that Albert and Loewer would like it to be.

Let us begin by looking at the second conception.

First, there are reasons to think that the second conception is not a coherent way to think about what a law is. This has been argued by Mathias Frisch. Frisch (fortcoming) gives the example of a simple world consisting of nothing but a smallish collection of objects 'governed' by Newtonian gravitation. In such a world, a Lewisian system consisting of Newton's laws would be made vastly more informative, and not much less simple, by including a precise specification of the world's initial condition. But if such a thing is included in the set of laws, then every fact about that world becomes nomically necessary–and the distinction between laws and nomically contingent facts would collapse. This is an unappealing result, and raises doubts that the second conception is a coherent one.

Of course, the world Frisch imagines is not our world. And it is a premise of Albert and Loewer's proposal (made vivid by the God story), that in our world, a precise specification of the initial conditions is too costly in terms of simplicity to be a candidate for being part of a best system. Perhaps this is so. So perhaps we ought only to conclude that the distinction Frisch is worried about collapses only for other worlds not like ours. So, even though Frisch's example clearly reduces the attractiveness of the second conception of informativeness, we might still want to entertain it as a possibilitiy, and ask instead whether or not its plausible that, on the second conception, PROB comes out as part the best system. I do not think it is, for two reasons.

The first reason has been articulated by Cohen and Callender (2009). While my primary concern in this paper is with the status of PROB, it should be recalled that Albert and Loewer's argument for the nomic status of PROB also depends on PH being a law. And it is precisely this that Cohen and Callender have called into question. What they point out is that whether or not PH counts as a simple addition to the package depends on what language it is expressed in. The idea is that PH presumably has some relatively simple expression in a thermodynamic language.[3] But if we try to express PH in the language of the microphysics, then its expression becomes almost as complex as writing down the precise microscopic initial condition of the universe.

A second worry is the following:

Lets call the the ordinary dynamical laws that we expect physics to deliver to us at the end of inquiry the P-laws. In most discussions of Boltzmannian statistical mechanics, we assume that the P-laws are ordinary Newtonian mechanics along with some ordinary conservative Hamiltonian. It doesn't really matter, here, what the P-laws turn out turn out to be–but the kind of thing I have in mind is that the P-laws are whatever future development of QFT, or quantum gravity, or whatever, physics has in store for us.[4] By P-laws, in other words, I mean the sort of things that working physicists will expect to find in their best textbooks in the library. What Albert and Loewer seem to be committed to is that, at the end of the day, when we do the best system analysis of our world, what we will get is the P-laws–the ordinary laws of physics we would find in our textbooks–plus PROB, plus PH.

But is that true? Or is it even plausibly true (on the second conception)? I would argue that it is not, and it is only a confusion, or a slipping back and forth, between the first conception of informativeness and the second one that makes

[3] Even this might be open to doubt. Notice that Albert formulates PH as the claim that "the universe began in whatever low-entorpy macrocondtion the normal inferential procudes of cosmology" tell us it did. In this form, PH is more of a law-schema than a candidate law. Its hard to know if PH is precisely formulable, even in the thermodynamic language.

[4] Unless, of course, the P-laws turn out to be indeterministic, or non-reversible, in just the right sort of way to make PROB unnecessary. I will not be discussing that possibility here.

it seem like it is plausibly true. Here is what I mean:

Let us first of all suppose that a basic intuition of Lewis' was correct. Let us suppose that, at the end of the day, the kind of laws, call them the L-laws, that one would get if we did a best system analysis, with informativeness defined as in the first conception above (the one I am assuming Lewis himself, and most of his readers, had in mind), then we would end up with the P-laws. Let us assume, in other words, that the P-laws are the L-laws. This is presumably something which Lewis argued for reasonably persuasively, and about which many philosophers of a Lewisean/Humean bent are convinced.

Let us then also suppose that the following intuition of Loewer's is correct:

"Adding PH and PROB to the dynamical laws results in a system that is only a little less simple but is vastly more informative than is the system consisting only of the dynamical laws." (p. 305)

I will assume here that when Loewer talks here of "the dynamical laws" what he means is what I am calling the P-laws.[5] (Let us call the dynamical portion of the package of laws Albert and Loewer want to contstruct the "ALD-laws"). So, suppose that this claim is true. And continue to suppose, as above, that the L-laws are the P-laws. How does it follow from this that the P-laws, plus PROB, plus PH, wil make up the best system? It follows only if we think that that there is some antecendent reason to be committed to the proposition that the P-laws necessarily make up a subset of whatever that best system will be, on the second conception. But why should we think that? We are only tempted to think, that, I would argue, if follow some mistaken reasoning.

Here is how the mistaken reasoning goes: First, rely on conception 1 of informativeness to decide what the dynamical laws are. Next, chose to believe, as most Humeans will do, that Lewis' intuitions were right, and these "L-laws" (defined in terms of the first conception) will turn out to be the P-laws. Then, switch to conception two of informativeness, and figure out if any summary of initial conditions, added to those L-laws/P-laws, add to informativeness without overly sacrificing simplicity. If we play this funny sort of game, then–plausibly– we end up with P-laws, PROB and PH as the best system. And then we end up with the P-laws being the ALD-laws. Indeed, this seems to be the reasoning behind the inference from

P. "Adding PH and PROB to [the P-laws] results in a system that is only a little less simple but is vastly more informative than is the system consisting only of the dynamical laws." (p.305)

to

C. The P-laws, PROB and PH make up the best system.

Unless we are following the mistaken reason, why would we care what adding something to the P-laws does to informativeness and simplicity?

Of course, this kind of reasoning is not legitimate. No one should think that the laws of nature are the sorts of things that come out of this kind of two step analysis. But this seems to be the reasoning that Loewer has in mind. If its not, then the claim that adding PH and PROB to the P-laws results in more informative and only slightly less simple system than just the P-laws is a red herring. Unless we have antecedently assumed that the P-laws will be part of the package, then the claim that adding PH and PROB will be more informative and only a little less simple is irrelevant to determining what is the very best system (as opposed to the best system that is *assumed to contain* the P-laws).

Suppose, then, that we don't antecedently assume that the P-laws will be part of the package. How would be convince ourselves that the P-laws would come out to be part of the best system on the 2nd conception of informativeness. Its hard to see how we would. After all, physicists dont test their putative laws by adding them to PH and PROB. They test them by combining them with detailed microscopic and macroscopic initial conditions. The P-laws are chosen because of their ability, in other words, to make predictions

[5] He can't mean anything like the de dicto best system laws or the sentence becomes strangely circular.

when combined with microscopic initial conditions. They are not chosen (by working physicists) for their ability to work well with PROB and PH.

If we are really trying to guess what will be the best system, where "best" is defined purely in terms of the second conception of informativeness, why not consider the possibility that some much simpler set of dynamical laws than the ones physicists will offer us–added to PH and PROB, will be just as informative, or close, as the desired package. After all, quite plausibly, any set of dynamical laws that is "friendly" to Boltzmannian reasoning, (i.e. that takes almost all of small regions of state space to large regions, that fibrillates volumes of state space, etc), will be just as informative, or close, on conception two, as the genuine P-laws. And surely there are candidate laws that are Boltzmann friendly that are simpler than the real P-laws, whatever they turn out to be. Or maybe the simplest most informative system will be one without any microscopic dynamical laws at all–i.e. one that just gives thermodynamic laws.

So, the second conception of informativeness is not going to work. It either forces us to rely on an ad hoc, two step process for determining what is the best system, or else it makes it implausible that the best system analysis will contain the P-laws as the dynamical portion of the best package.

What, then, of conception three of informativeness? Recall that the third proposal regarding informativeness is significantly different than the first two in that it does not involve a God's eye comparison between the world, on the one hand, and some fixed system of sentences on the other. Rather, the informativeness of a system is evaluated by an active scientist/participant in the world. This sort of agent has certain pragmatic goals–things that she knows and things that she wants to predict and retrodict. Such an agent will compare the informativeness of two systems by using each of them, in combination with whatever microscopic and macroscopic information she happens to come by in the pursuit of those pragmatic goals, to predict and retrodict whatever interests her. The most informative system, on the third conception, is whichever system is best at doing that.

In one respect, this is an attractive way of thinking about informativeness for grounding the notion of PROB as a fundamental law. It is radically naturalistic and Humean. And it dovetails well with Albert's now famous line that PROB is "the right probability distribution to use for making inferences about the past and future." (2000, p.96) And it might be genuinely plausible, I would at least grant here for the sake of argument, that P-laws+PROB+PH will come out on top on this conception of informativeness. That's because an agent like the one we envision here will want to make both microscopic predictions from the bits of microscopic knowledge she can infer to, and will want to make macroscopic predictions about macroscopic objects about which she has only macroknowledge. Doing the former will require that the microlaws in her package be the real P-laws, rather than any old Boltzmann friendly set, and doing the latter will require that she add PROB and PH to her macroknowledge. These considerations make the third conception attractive.

What makes the third conception much less attractive, in my opinion, is the strong tension that exists between the pragmatic nature of the proposal, on the one hand, and a very un-pragmatic feature of the rest of what is going on in the overall scheme of things.

Recall, first of all, that what we are after here is a system that will be *constitutive* of the fundamental laws of the universe. Making such a system dependent on the pragmatic goals of a particular agent, or set of agents, strikes me as unattractive. Though I would think someone like Lewis would share such a worry, I admit that perhaps having the worry is not entirely Humean . Perhaps a genuine Humean wants to relativise the fundamental laws to particular kinds of epistemic agents.

If so, however, then the following consideration becomes all the more forceful: there is an ambiguity in the claim that PROB is the "right probability disribution to use...". In the sense that is most in harmony with the pragmatic spirit of the third conception of informativenss, PROB is not a probability distribution that can be used *at all*, for making predictions and retrodictions. That's because, pragmatically, we lack[6] the ability to compute predictions using the P-laws at the scale of the entire universe. The only sense in which PROB, along with PH, is the "the right probability distribution to use" is this:

If we encounter some macroscopic system, say a cocktail with ice in it, and we want to predict what it will be like in ten minutes, then here is what we do: we use macroscopic laws that we hold to be correct on the basis of macroscopic inductive evidence to predict what it will be like in ten minutes. Then, if we like, we can perhaps convince ourselves that it is plausible, given the correctness of those macroscopic laws, that *had we been able to make the calculation*, it would have come out that that outcome was overwhelmingly likely on the basis of what we knew, the P-laws, PH, and PROB. But that is a very very long way from the claim that, as a pragmatic matter, we could have predicted that outcome from what we knew, the P-laws, PH, and PROB.

A short example illustrates this point rather well[7]. Famously, in *Time and Chance* Albert makes the claim that the package of the dynamical laws, PH, and PROB not only enable us to predict things about the behavior of traditional thermodynamic systems, it also enables us to predict, from our knowledge of the design of an apartment and the fact that it contains a spatula, that the spatula is very likely to be found in the kitchen. In various debates about the book, the objection has been raised that unless we know whether or not martians have visited the apartment, moved around all the household items, and subsequently erased all macro traces of their visit, then we dont know what the probability, from that package, of the spatula being in the kitchen is. Albert's reply has been that the probability of such a martian visit, given what we know, and the package, is very low–and so conditional on what we know, the probability of the spatula being in the kitchen is still high. Others are skeptical that the low probability of the martian hypothesis follows from the package. I take no stand on that issue here. I merely want to point out that all we have here is a clash of intuitions about how much confidence to have in the package. We have nothing even remotely like the ability to even crudely estimate those probabilities. The point, in short, is that from a genuinely *pragmatic* point of view, the package of the P-laws, PH, and PROB is *entirely uninformative.*

Together, I find that these two worries about the pragmatic nature of the third conception of informativeness combine to make it an unattractive way of grounding the lawfulness of PROB and PH. In sum, then, none of three conceptions of informativeness provide a clear, coherent and attractive way of understanding how PROB and PH could be fundamental laws of a Lewisian flavor.

4 What are PROB and PH after all?

What then, of the arguments that PROB and PH must be laws? Recall that part of the point of this whole enterprise, according to *Time and Chance*, was to understand how the probabilities in statistical mechanics could be objective chances–the kinds of things that are involved "*in bringing it about*, in *making it the case*, that (say) milk dissolves in coffee," rather than being about "mere degrees of belief." So one argument in favor of PROB being a law was that that was only way it could deliver probabilities that were not mere degrees of belief, and that the probabilities in statistical mechanics could not be degrees of belief.

[6] And given the physical constraints on computation, we will probably always lack them.

[7] This example was raised by Stephen Leeds in his (2003).

Let us accept, for the moment, that if PROB is not a fundamental law, then it cannot provide objective chances. So what status does it have? And what kind of probabilities does it deliver? I would argue that we ought to understand the probabilities in PROB as objective epistemic degrees of belief, and that we ought to regard PROB itself as a nomically contingent empirical hypothesis.

Elsewhere (Winsberg, 2008), I have given independent arguments for doubting that the probabilities inPROB should be understood as objective chances. I also articulated in more detail the idea that the probabilities in PROB could be understood as representing degrees of belief, and I defended this possibility against the worries, expressed in *Time and Chance*, that degrees of belief cannot explain why the ice melts in my cocktail. Let me just reiterate here that by saying PROB is about epistemic probabilities I do *not* mean to say that they are innocent, or a priori, or reflect uniform ignorance–and I believe it is only against those positions that the arguments in *Time and Chance* apply. Rather, I think the probabilities in PROB are (objectively correct) epistemic probabilities precisely because they are the right probabilities to use given our particular epistemic situation. I also think that it is the job of PROB to *underwrite the legitimacy of our inferences* about what will happen to the ice in my cocktail. It is not the job of the probabilities in PROB to "*make it the case*" that the ice in my cocktail melts. Only the dynamical laws, in conjunction witth the specific initial condition of the universe, can do that.

5 Underwriting the laws of the special sciences

Here, I would like to take the space to address the argument, put forward by Loewer, that : "PH and PROB *can bestow lawfulness* on [the higher level laws] only *if they themselves are laws...*" (Loewer, 2007, pp. 304–305, my emphasis.)

Loewer's argument is rather simple: it appears to be a genuine law, for example, as opposed to an accidental regularity, that ice in a glass of lukewarm water will melt. Therefore, according to Loewer, this fact needs to be underwritten by the lawfullness of PROB and PH. And so the underlying premise is clear: since thermodynamic laws are not fundamental laws–they are laws of a special science–not only their truth, but also their lawfulness, must follow from fundamental physics.

Loewer's argument, then, depends on a fairly strong variety of what we might call fundamentalism. "Fundamentalism" is word coined by Nancy Cartwright, and it is normally understood as the doctrine that there exist fundamental laws of physics whose scope is universal (Cartwright, 1999). According to this version of fundamentalism, then, any truth discovered by a special science ought to be, at least in principle, recoverable from those fundamental laws. But the kind of fundamentalism being appealed to by Loewer here is even stronger. Not only, according to the reasoning above, must the truth of the laws of the special sciences be recoverable from the fundamental laws, but their lawfulness must be as well.

The question, then, is: does what we know about the relationship between what we take to be the laws of the special sciences and the fundamental laws support confidence in this very strong form of fundamentalism? I believe the answer to this question is "no".

To see why this is so, we need only ask ourselves why Loewer insists that both PROB and PH are required to underwrite the lawfulness of the special science laws. The reason that PH is required, of course, is that the laws of thermodyamics have force in both temporal directions. I can not only predict, using the laws of thermodynamics, that the ice in my cocktail will melt; I can only also retrodict that, at some time in the recent past, the ice in my cocktail was less melted than it is now. That this is so will not follow from the dynamical laws and PROB alone–indeed, from just them, the opposite follows. It also requires PH. Famously so.

Unfortunately, it does not follow from PROB and PH (plus the dynamical laws) alone either. The reason for this has to do with Poincaré's recurrence theorem. Roughly, Poincaré's recurrence theorem (PR) tells us that any system governed by a conservative Hamiltonian, will, after a sufficiently long period of time, return to a state arbitrarily close to any state that it occupied in the past. If the requirements for PR are met, then for any presently observable macrostate of the universe, we should expect a future of maximum entropy followed endlessly recurring thermal fluctuations back to that macrostate.

This is state of affairs that leads to what Sean Carroll and others have dubbed the "Boltzmann's Brain" (BB) paradox.[8] To understand the BB, it helps to revisit a suggestion once made by Boltzman himself. One of Boltzmann's initial reactions to Poincaré's theorem (when it was put forward as a "reversibility objection" to Boltzmann's attempt to reduce the 2nd laws to microphysics) was to postulate that the universe was undergoing a long history of thermal fluctations, and that we just happened to be on a local fluctuation, with the forward direction of time just being, by definition, whatever direction was locally headed toward equilibrium. Famously, the problem with this scenario is that it leads to a skeptical paradox. If we suppose that the universe is undergoing thermal fluctuations around equilibrium, then no matter what macrocondition we observe, it becomes overwhelmingly likely that the macrostate we think we are observing is in fact a spontaneous fluctuation out of chaos, rather than a macrostate with a normal macrohistory (one which runs in accord with the second law.) And so it becomes overwhelmingly likely that any apparent "records" are nothing but spontaneous fluctuations, and hence are not veridical. If we assume that the only real records we have genuine access to are our own mental states, then the assumption that we are in a universe undergoing fluctuations around equilibrium leads to the conclusion that it is overwhelmingly likely that we are nothing but BBs–disembodied brains surrounded by a thermal equilibrium, whose memories and perceptions are not veridical: they are nothing but spontaneous fluctuations out of chaos. And it really does not matter what we take to the observable macrostate; the conclusion is more general. No matter what we take to be the directly surveyable macrostate, we should conclude that that it is nothing but a thermal fluctuation. If you take youself to be veridically observing that you are in a room with a desk with this book on it, you should conclude that it is overwhelmingly likely that the room with the book is a fluctuation, surrounded–outside of the reach of your eyesight–by thermal chaos, and preceded in time by more of the same.

Strangely, Carroll and others take this result to be "a direct disagreement between theory and observation", (2010, 314). Presumably, what they mean is that any set of laws that gives rise to a BB scenario is empirically disconfirmed by what we see. But this is incorrect. In fact, no observation is ever inconsistent with the hypothesis that I am a Boltzmann Brain. In fact, *that* is the problem. Of course, this is no different from the fact that no observation is ever inconsistent with the possibility that the universe is presently in the lowest entropy condition it will or ever has been in. That is why, in order to avoid skeptical paradox, we need to postulate PH. PH, we might say, is a condition for the possibility of our having knowledge of the past. As we will see in what folows, something else, something akin to PH, but different, will also be required to block BB from giving rise to similar intractable worries.

What, after all, follows from the fact that we

[8] See, for example, (Dyson et al, 2002), (Albrecht and Sorbo, 2004), (Linde, 2007) and (Carroll, 2010). More precisely, what I have given is the conditions of a "classical" BB paradox. Carroll and others are more concerned with a more modern, quantum-gravitational version of the paradox. On this version, we are to take the recent evidence of the presence of dark energy in the universe to be reason to believe that we are headed toward a future universe that is a de Sitter space. In such a universe, vacuum fluctuations would give rise to behavior that was, for all intents and purposes, just like classical Poincare recurrence. And so if fact the details do not really matter.

have reason to believe that the universe will eventually reach a state of equilibrium, followed endless recurrences to macroscopic states arbitrarily close to the one we are in now. It follows that if I reason in the way that Albert proposed that I do, that is, if I try to make predictions and retrodictions using a package consisting of only PH, PROB, and DL[9], then it is overwhelimingly likely that the macrocondition I am presently "observing" is not the one that evolved directly out of the big bang, but is one of the infinitely many fluctuations out of equilibrium that will occur after the universe relaxes. And so it follows that it is overwhelmingly likely that I am a BB and no one will ever read this paper, which exists only in my mind. More importantly, it is overwhelmingly likely that all the records (including memories) I have of the experiments that made me believe in modern physics in the first place are not veridical. It is also, then, overwhelmingly likely that none of my records are veridical. And it is also the case–and here is the rub–that none of the laws of the special sciences will hold, since none of hold for Bolzmann Brains.

And so rather than being in disagreement with experience , the BB scenario is epistemologically unstable–and denying one of its premises is precondition for the possibility of veridical records. And so denying one of its premises is a precondition for the possibility of believing in the body of evidence on which we base all of physical theories in the first place, and indeed of reasoning statistical mechanically in the first place.

6 Conclusion

So, just as we need to postulate PH in order to avoid the skeptical paradoxes discussed in *Time and Chance*, we also need to assume one more thing in order to avoid the BB scenario (and a resulting skeptical paradox that is just as bad) as the outcome of our package: we need to assume a postulate, call it the Near Past postulate (NP) which says that our present state lies somewhere between the time of the PH and time when the universe first relaxes . The epistemological status of such a principle would be exactly like the one that Albert attributes to PH in *Time and Chance*–it is a principle that we reason to trancendentally, as it were–it is a condition for the possibility of our having the very knowledge base that forms the empirical evidence base for the physics we were doing in the first place.[10]

The prolem for Albert and Loewer's fundamentalism, however, is that, unlike PROB and PH, NP is not even a prima facie contender for being a law. It can not be a law because it does not describe a fact about the world; it describes a fact about our present location of the world. To put this point most starkly, imagine supposing that it is a law of nature that the universe is 14 billion years old. This is absurd. (Can laws have expiration dates?). But to suppose that NP is a law is akin to this.

So NP cannot be a law. But NP plays exactly the same role as PH (and a role not unlike that of PROB) in "underwriting" the laws of the special sciences. So something must have gone wrong in the inference to the conclusion that everything in the package to which PH and PROB belong must be fundamental laws; it turns out to be impossible to derive the laws of the special science from any observable set of conditions plus a package containing only laws.

And here I agree wholehartedly with Albert (and not incidentally, with Boltzmann, Einstein, and Feynman) that this is indeed the correct epistemological status of PH. And hence I think that NP has exactly the same epistemological status. And hence, *pace* Carroll and others, I think it is a mistake to think that a theoretical system that results in a BB scenario, but for which that scenario can be blocked with the ad-

[9]Assuming these dynamical laws are friendly to PR.

[10]And whether you want to think about that classically–in terms of reaching a state of thermal equilibrium, or in terms of some fancy quantum-gravitation version of entropy–presumably de Sitter space–doesn't seem to me to make any difference at all. And hence I do not really understand the worry, expressed by some, that the BB scenario is a special worry that arises in the context of recent discoveries/proposals in cosmology.

dition of NP, is "in disagreement with observation".

So if there is a simple informative package the underwrites the truth of the laws of the special sciences–a package which, when it is added to what we can observe, tells us, e.g. that ice will melt going forward and unmelt going backward–that package must contain NP. And NP is not a law. So it can not be that the things that are required to be in the package must be laws. So it cannot be that PROB and PH are *required* to be laws so that they can underwrite the lawfulness of the laws of the special sciences. But anyway: the various possible ways in which one could have made sense of the claim that PH and PROB are laws were unattractive. What should we conclude? I would argue that we should conclude that the strong form of fundamentalism that motivated the desire to elevate them to the status of laws in the first place was ill-advised. The lawfulness of the laws of the special sciences must be autonomous features of their respective levels of description.

References

[1] Albert, D. (2000). *Time and chance*. Cambridge: Harvard University Press.

[2] Albrecht, A., and Sorbo, L. *(2004)*"Can the universe afford inflation?", *Physical Review D, 063528*

[3] Callender, C. (forthcoming) "The Past Histories of Molecules", in Claus Beisbart and Stephan Hartmann's (eds) *Probabilities in Physics*, Oxford University Press.

[4] Cartwright, Nancy (1999), *The Dappled World: A Study of the Boundaries of Science*, Cambridge University Press.

[5] Cohen, Jonathan and Craig Callender. (2009) "A better best system account of lawhood". *Philosophical Studies* 145:1, 1-34

[6] Dyson, L., Kleban, M., and Susskind, L. (2002) "Disturbing Implications of a Cosmological Constant", *Journal of High Energy Physics*, 0210 011

[7] Earman, J (2006) "The 'Past Hypothesis': Not Even False," *Studies in History and Philosophy of Science* 37: 399-430

[8] Frisch, M. (forthcoming) "From Arbuthnot to Boltzmann: The Past-hypothesis, the Best System, and the Special Sciences."

[9] Leeds, S. (2003). Foundations of Statistical Mechanics—Two Approaches. *Philosophy of Science* 70 (1):126-144

[10] Lewis, D. (1986). A subjectivist's guide to objective chance, *Philosophical Papers*, Vol. II. Oxford: Oxford University Press.

[11] Lewis, D. (1994). "Humean supervenience debugged." *Mind*, 103, 473–489.

[12] Linde, A. (2007), "Sinks in the Landscape, Boltzmann Brains, and the Cosmological Constant Problem." *Journal of Cosmology and Astroparticle Physics*, 0701, 022

[13] Loewer, B. (2007). "Counterfactuals and the second law." In H. Price, & R. Corry (Eds.), *Causation, physics, and the constitution of reality: Russell's republic revisited*. Oxford: Oxford University Press.

[14] Winsberg, E. (2004). "Can conditioning on the 'past hypothesis' militate against the reversibility objections?" *Philosophy of Science*, 71, 489–504.

[15] Winsberg, E. (2008). "Laws and Chances in Statistical Mechanics." *Studies In History and Philosophy of Science Part B: Studies In History and Philosophy of Modern Physics*, 39, 872-888