

Do statistical mechanical probabilities need quantum shuffling? Albert's Boltzmannian Story and GRW collapsesⁱ

Jenann Ismael

*For a collection of newly commissioned papers on themes from David Albert's *Time and Chance* (Harvard, 2000), with replies by Albert, Loewer, B., Winsberg, E. and Weslake, B., eds., <currently untitled>, Harvard University Press.*

Introduction

In the last chapter of *Time and Chance*, Albert suggests that Ghirardi-Rimini-Weber theory (GRW) might provide a grounding for the statistical probabilities of thermodynamics. The proposal was that the dynamical laws of GRW provide general transition probabilities that generate the micro canonical distribution of statistical mechanics. In an otherwise favorable review of the book, Huw Price questioned the motivation for the proposal. He writes:

“According to the statistical approach, nothing forces a cup of coffee to cool. Everything depends on the initial arrangements of its microscopic constituents. In the rival view, he points out, the GRW mechanism introduces a kind of quantum shuffling, so that it doesn't matter how the coffee molecules are arranged initially—whatever the arrangement, the coffee is almost certain to get cold.

“But is there any work for this shuffling to do? To think so is to think that if we didn't have it, things would happen differently. Coffee wouldn't cool, or something of the kind. ... No one has yet shown that there is work for such a law to do.”

Price's claim is that there's nothing for the proposal in the last chapter to explain. It adds nothing to the Boltzmannian story.

The question of whether there is any explanatory work for shuffling raises some puzzles about the explanatory logic of the Boltzmannian Account. I will identify two ways in which the official version of the account seems to fall short of recovering thermodynamic phenomena. I will say how to reformulate it so that it hits the target. The revised formulation will open up an explanatory lacuna, which I will then suggest that the Shuffling hypothesis will fill. Although the discussion takes its departure from Price's question about the role of the Shuffling Hypothesis, it highlights two familiar difficulties, both of which run deep in the foundations of thermodynamics: (i) a tension between global and local accounts and (ii) a tendency to run together epistemic and statistical interpretations of probability.

The explanatory target

The worry that dominated the early history of foundational discussions in thermodynamics is that there was a strict incompatibility between temporal asymmetry of the second law and the underlying temporally symmetrical dynamical laws of the statistical theory. That worry was alleviated when it was observed that we don't have empirical evidence for the strict impossibility of anti-thermodynamic behavior, and the macroscopic asymmetry was formulated in probabilistic terms. If anti-thermodynamic behavior is not impossible, but simply highly *improbable*, there is no logical conflict between the macrophenomena and the microlaws, and the focus thus turned to understanding what facts, in addition to the laws, generate the macroscopic asymmetry.ⁱⁱ Since the laws are deterministic, the complete mechanical explanation of thermodynamic phenomena is “given laws, together with the initial conditions, which entail everything that happens in history in full specificity,” but that would be unilluminating because it wouldn't tell us what general features of initial conditions are responsible for generating the macroscopic asymmetry embodied in thermodynamic phenomena. For those purposes, we want to know what features of initial conditions are both necessary and sufficient for the production of thermodynamic phenomena. The aim of Albert's book is to provide an answer to this question, i.e., an account in microphysical terms of the facts about the world in virtue of which the generalizations embodied in the thermodynamic laws hold.

The first 6 chapters of *Time and Chance* are carried out in a Newtonian regime and the account is squarely in the Boltzmannian tradition. The theory is contained in three laws and a contingent hypothesis about the early history of the world.

1. The Newtonian laws of motion...
2. The Statistical Postulate...
3. The Past-Hypothesis."

All three postulates are to be applied at the global level. The Newtonian laws of motion are the familiar time symmetric laws that allow the calculation of the state of the world at one time from its state at any other.

The statistical postulate (SP) is a probability distribution over worlds.

The Past Hypothesis (PH) says that the universe started in a state of low enough entropy to make thermodynamic generalizations applicable for the roughly 15 billion years we think these generalizations have held.ⁱⁱⁱ The three postulates work together as follows: The Newtonian laws delimit the space of physically possible worlds. The PH knocks out all of those worlds that don't start in a low entropy state. The statistical postulate provides a probability distribution over remaining worlds that overwhelmingly favors those on an entropy increasing trajectory. The result is, in his words,

"That the most probable history of the universe is one wherein entropy rises."

The account has been very well picked over.^{iv} A good many of the questions concern technical matters, e.g., about whether the central postulates can be given mathematically rigorous formulation^v, or whether the mathematical conjectures on which the account rests are true. The questions that we will raise concern the basic explanatory logic of the account and arise before any of these delicate issues,^{vi} so we will follow the informal presentation and allow for whatever mathematical conjectures Albert needs. There is only one sense in which we will be somewhat stricter than Albert. He uses epistemic vocabulary freely, talking not of what *is* the case, but what we should *believe*, and speaking of theories often not as descriptions of the world, but as algorithms for deriving probability distributions for past and future events.^{vii} We are simply going to leave epistemology out of the picture. We will treat statistical mechanics entirely as a mechanical theory, i.e., a theory about the movements of particles whose purpose is to explain the macroscopic regularities embodied in phenomenological thermodynamics. We will talk of particle trajectories, the micromechanical laws that those trajectories obey. We will talk of the statistics of various kinds of trajectories in specified samples of various kinds, and we will allow talk of statistical probabilities, which will be understood explicitly as quantities characterized formally by the Kolmogorov axioms and interpreted physically by their relations to statistics.^{viii} But there will be no talk of belief, algorithms for setting credences, or epistemic probabilities as part of the explanation of thermodynamic behavior.

The first problem

There are two difficulties that seem to emerge within the explanatory logic of Albert's account. The first difficulty has to do with a mis-match between the explanatory target and what the 3 postulates achieve. If the goal is to articulate the facts about our world in virtue of which thermodynamic generalizations are true, then the conclusion that *it is highly probable* that the universe is on an entropy-increasing trajectory is too weak.^{ix} 1-3 might hold, and world might *not* be on an entropy-increasing trajectory, so 1-3 can't be a complete account of the facts about our world in virtue of which thermodynamic generalizations hold. That might seem like an easily fixable flaw. All that we need to do is strengthen the postulates so that they entail thermodynamic generalizations. The fact that the three postulates entail a very high probability that the universe is on an entropy-increasing trajectory tells us that we are close to the conclusion that we want. The problem is that as soon as we supplement the postulates by something strong enough to entail thermodynamic generalizations, the statistical postulate becomes irrelevant.

Interpreted statistically, the probability distribution over worlds tells us something about the composition of the *class* of worlds of which the actual world is a member, viz., that the proportion of those in an arbitrary

subselection of members of the class that are on an entropy increasing trajectory is large relative to the ones that aren't. But if we want to know what facts about the actual world make it the case that thermodynamic generalizations hold (or, as Albert says, what "more fundamental" facts), then the information about the composition of the class of physically possible worlds embodied in the probability distribution would seem to be extrinsic and irrelevant.^x The small bit of information that it contains about our world (viz., that it is a member of that class, and so doesn't have any features that are not possessed by any member of it) is screened off immediately by knowing what our world is actually like.

We can put this in the form of a dilemma: either thermodynamic generalizations hold at our world or they do not. If they don't hold, then all bets are off, and pointing out that they hold at worlds not like ours is a red herring. If they do hold, then they do so in virtue of some combination of laws and initial conditions (or, laws, initial conditions, and chancy departures from deterministic evolution along the way). In this case, the probability distribution – which gives us a relative measure of PH worlds in which they do and do not hold – is a red herring. So, in either case, the probability distribution is a red herring.

The problem that gives rise to the dilemma is that the probability distribution, if it is interpreted statistically, tells us something about the composition of the class of worlds that are physically possible according to our fundamental theory. But the explanandum concerns a particular case. We want to know what facts are responsible for the obtaining of thermodynamic generalizations at our world. What goes on at other worlds plays no role in the production of thermodynamic behavior at our world. What that shows is that the purely *modal* (or counterfactual) content of the distribution over worlds is irrelevant to the derivation of thermodynamic generalizations at our world. The only thing that is directly relevant is the very minimal content it holds about the actual world; content that is – once again – screened off by the addition of facts about our world sufficient to entail thermodynamic generalizations.^{xi} For the purposes of knowing the facts of our world in virtue of which thermodynamic generalizations hold, it matters naught whether the initial conditions at our world are chosen from a class that is dominated by those that lead to entropy increase, or dominated by those that lead to entropy decrease.

The logical situation would seem to be just like the following: suppose that, for a variety of reasons, 9 out of 10 of the children in a particular village die before the age of two. Suppose we have a good understanding of the various causes – from poor nutrition to accident- and that we also have a well-defined statistical probability distribution that an arbitrary child will survive or die from one of those causes. If we want to know why a particular child died, we want to know the specific facts about *her* physical state that led to her death, the probability of death for a child in her village – whether it is high or low – would seem to be irrelevant. Suppose you asked someone why this particular child died, and the response was "Oh, she died because 9 out of 10 children in this village die before age 2" would not be an adequate one. The general probability that a child in the village dies before age 2 might help us set our epistemic probabilities for the various factors that might explain her death, but her actual physical state is what explains her death. The *epistemic* relevance of the statistical probability is screened off immediately as soon as that state is known, and the statistical probability has no direct physical relevance to her death.

A distribution over worlds can have explanatory relevance in various ways, if we are not careful about *what* is being explained and because 'explanation' is a weasel word:

- (i) To say that P renders E highly probable is to say that the actual world is a member of a class in which the distribution of worlds on entropy-increasing and entropy-decreasing trajectories heavily favors entropy-increasing ones. That means that if I know nothing else about the actual world than that it is in a low entropy state, then my expectations ought to favor entropy increase.^{xii}
- (ii) It also signals that P is *****almost***** all I need to know about the actual world to know the more fundamental facts in virtue of which thermodynamic generalizations hold at our world. It doesn't tell me everything I need to know about the actual world, since there are P-worlds in which thermodynamic generalizations do not hold. But the fact that P renders E highly probable signals that the additional assumptions I need to make to derive thermodynamic generalizations are weak.^{xiii} This is even clearer if we link probability to typicality. When we say that some behavior is typical for systems in a certain class, we say that the behavior is explained in part by the laws that govern the members in the class, and in part by the accident of choice between the members. So when we say that 1-3 render saying that most of the explanation

is done by the laws, but part of it is an accident of the choice of initial conditions, i.e., that most or almost all choices of initial conditions would have resulted in the phenomenon, though some would have led to different behavior. In this capacity, it is playing a *metaexplanatory* role, functioning as a gauge of how much explanatory work is done by the laws and how much is an accident of initial conditions.^{xiv}

- (iii) Finally, the probability that a theory assigns to the evidence is a measure of its likelihood and likelihood is a measure of evidential support. So one very important way of showing that a theory is well supported by the evidence is to show that the theory renders the phenomena highly probable.^{xv}

A moment of reflection will show that none of these give the probability a role in the *physical production* of thermodynamic generalizations at our world. It is for this reason that I've been very explicit about what the account is supposed to achieve, and have avoided the term 'explanation'. If the account is supposed to provide a full account, in microscopic terms, of the facts about our world in virtue of which thermodynamic generalizations hold (rather than derive expectations about the actual world in the face of ignorance, gauge how much explanatory work is done by the initial macrostate and laws, or assess how well the Newtonian laws are confirmed by thermodynamic phenomena) information about the statistical distribution of initial conditions across the class of worlds, of which the actual world is a member, would seem to be irrelevant.

In sum, if (1) – (3) falls short of articulating the facts about our world in microphysical terms in virtue of which it is on an entropy-increasing trajectory. But as soon as we make the postulates strong enough to derive what needs to be explained, replacing PH with the slightly stronger PH*, which says that the world began in one of those low entropy, initial states that is on an entropy increasing trajectory, SP falls out of the picture. I see no role for a statistical postulate defined over worlds as wholes in articulating the microphysical facts that underlie the production of thermodynamic phenomena at our world.

If the probability in SP is not to be interpreted statistically and is not playing one of the peripheral roles i-iii, then we need some clarification about how it is to be interpreted, and what role it is playing. My own suspicion is that there is equivocation between the epistemic and statistical probability. We move between these so seamlessly in everyday reasoning that the distinction becomes invisible. This doesn't matter for everyday inferences, but it does matter here.

The second problem

I said that there were two problems with Albert's conclusion that the most probable history of the universe is one of increasing entropy. The first problem centered on the phrase 'most probable'. The second centers on the phrase 'the universe'. To say what this problem is, we need to be a little more precise about what we mean by 'thermodynamic generalization'. Phenomenological thermodynamics describes the behavior of relatively small, relatively short-lived subsystems in energetic isolation: gases in enclosed containers, an ice cube in a glass of water, a chunk of table salt in a test tube of hot water, a teaspoon of cream dropped in a cup of black coffee. These are standard examples of the kinds of systems to which thermodynamics is successfully applied. The standard procedures for deriving *predictions* is as follows. We begin by observing that such systems possess a large number of degrees of freedom, but that the gross constraints on the system consist of measured values of a relatively small number of macroscopic parameters. We introduce a probability distribution over the microstates states of the system compatible with the gross constraints, and we evolve it forward under the Newtonian laws. In retrodiction, we don't just reverse the procedure, evolving the same probability distribution backward. We impose a boundary condition in its past that ensures that it evolved from a state of yet lower entropy. The measure that is employed is the one that is uniform by the standard measure over the volume of phase space corresponding to the known macrostate of the system at the initial stage. It corresponds intuitively to the number of microstates compatible with the macrostate in question so that by this measure, the most probable states are those have the most fully mixed or random distribution of microstates across constituents (and in general that the more mixed they are, the more probable).^{xvi} These procedures work well to predict the macroscopic behavior of relatively small, relatively short-lived systems in adiabatic isolation - not just qualitatively predicting the march towards entropy, but quantitatively predicting relaxation times.

Let's see how well Albert's account does in explaining the facts about the world in virtue of which these procedures are successful. 1 -3 provide a global story that tells us that the universe is (very probably) on an

entropy-increasing trajectory. Let us put aside the ‘very probably’, and suppose our world *is* on an entropy increasing trajectory. What does that tell us about the behavior of the local, approximately adiabatically isolated subsystems to which thermodynamic generalizations apply? Does it follow that subsystems of the universe corresponding to cream in coffee cups, gases in enclosed containers, and ice cubes in glasses of water are also on entropy-increasing trajectories? The state-spaces for these systems are low-dimensional subspaces of that for the universe as a whole, obtained by effectively ignoring dimensions that correspond to the states of particles that are not part of the subsystem of interest. If the trajectory for the universe through its subspace is entropy increasing,^{xvii} do we have some assurance that the portion of the trajectory that goes through these subspaces is itself entropy increasing? Or do we have some assurance that the trajectory that the states of these systems follow through their subspaces will be a lower-dimensional projection of the whole?

The first thing to notice is that it is not generally true that there is this kind of match between the global trajectory and the trajectories of subsystems. It is easy to construct examples of systems, for example, whose global trajectory is opposite to that of most or even all of the local trajectories of its subsystems. Consider, for example, a very long train full of people clustered at the south end. The train moves south at a fixed velocity of .01 km/hr. The people on the train run north at .02 km/hr. ^{xviii} So we need to add something to reproduce the global trajectory at the local level. We know that the global trajectory is a product of the microphysical laws + an assumption about the low entropy initial state + SP. By hypothesis these systems obey the same microphysical laws, so to reproduce the global trajectory, we need analogues of PH and SP at the local level. Since the kinds of local, adiabatically isolated subsystems to which thermodynamics is applied – e.g., warm glasses of water with fresh ice cubes, or cups of coffee to which spoons of cream have been added – are typically far from equilibrium, a local PH might be granted as part of the definition of the class of systems that exhibit thermodynamic behavior.^{xix} So the pressure is on the question of whether 1-3 permit a derivation of a statistical postulate for such systems. Unless we have a derivation of SP for such systems, we don’t have a derivation of thermodynamic behavior. There is a tendency in foundational discussions, reproduced in Albert’s text, for the discussion to turn towards epistemic probabilities at this point. But we are leaving aside any discussion of belief. We are squarely focusing, here, on a physical phenomenon – the tendency of local, adiabatically isolated subsystems of the world to exhibit thermodynamic behavior – and we are looking for a mechanical (as opposed to epistemic) explanation.

In concrete terms, the principle to be derived to generate thermodynamic behavior is that there is a high statistical probability that the microstates of local, approximately adiabatically isolated subsystems of the world are normal, or – to put it another way – that abnormal microstates are so exceedingly rare under natural conditions that we should be shocked to see them, i.e., that when we sample such systems, we never find one arising naturally whose microstate is among those very special ones that produce anti-thermodynamic behavior. To get some concrete, intuitive feel for abnormal microstates - i.e., microstates whose forward dynamical evolution leads to all of the particles in a dispersed gas suddenly collected in one corner, for example, or cream dispersed through a cup of coffee spontaneously collecting to form a droplet - forward evolve a system that starts in a low entropy state to a state of higher entropy, and then reverse the states of its constituent particles keeping the system isolated.

Can we derive local SP is from a global SP? Callender describes what is needed in picturesque terms:

“Recall that any sub-system corresponds, in phase space, to a lower-dimensional subspace of the original phase space. The hope must be that when we sample the original approximately uniform distribution onto this subspace and renormalize we again find a distribution that is approximately uniform. Pictorially, imagine a plane and a thoroughly fibrillated set of points on this plane, a set so fibrillated that it corresponds to an approximately uniform measure. Now draw a line at random through this plane and color in the points on the line that intersect the fibrillated set. Are these colored points themselves approximately uniformly distributed? That is what the Boltzmannian needs for local thermodynamic systems, except with vastly many higher dimensions originally and much greater dimensional gaps between the phase space and subspaces. How are we to evaluate this reply?”

I have already said why I think the global SP has to be a red herring. In this case, it’s the non-probabilistic fact that our world *is* on an entropy-increasing trajectory that is physically relevant to the derivation of

thermodynamic behavior at our world. But even if we put aside this worry, and we allow Albert his version of PH, a global SP, if there is no logical entailment to the local SP, we have no derivation of thermodynamic behavior from statistical mechanical principles. Many others have noticed this lacuna: Winsberg (2004b) and Frigg (2008a) both express skepticism. Callender is willing to believe it, but writes: “Based on experience with many types of systems, some physicists don't balk at the thought of such fibrillation. They see it in some of the systems they deal with and in the absence of constraints ask why things shouldn't be so fibrillated. Others such as Winsberg (2004b) think it false. The unhappy truth, however, is that we simply have no idea how to evaluate this claim in general.”

What Albert needs to bridge this gap is a logically transparent derivation of a local SP from 1-3. Without such a derivation, we have a very incomplete understanding of the links that are supposed to bridge the gap between these global postulates and the phenomenological regularities embodied in thermodynamic generalizations. And as before, once it is filled, it would seem that we could dispense with the global SP, since it is the local one that is needed to recover thermodynamic generalizations. The upshot is the connection between global entropy and local entropy of any given subsystem is neither direct nor clear, and without some explicit physics that bridges the gap between a globally increasing trajectory and a local SP, we haven't got a derivation of thermodynamic generalizations.

Revising the Boltzmannian Story: moving SP to the local level

The problem with Albert's Boltzmannian theory is that it is formulated at the global level. All three postulates are to be applied at the global level and the consequence of the three postulates is stated in global terms. That is the source of both of the problems above. It is what makes the content of the probability distribution largely extrinsic to the actual world, and hence seemingly irrelevant to the derivation of thermodynamic generalizations. And it is what opens up the logical gap between his postulates and the local thermodynamic generalizations. There are other problems about a global formulation that I haven't mentioned. There are worries about the empirical meaning of a statistical distribution over worlds as wholes.^{xx} And there have been objections^{xxi} about whether the global version of the PH makes sense. Thermodynamics, however, wasn't formulated as a global theory and we have no direct evidence for thermodynamic generalizations at the global level. The procedure that Albert recommends – evolve forward a flat distribution over the volume of phase space defined by the initial macrostate of the world – is one that works at the *local* level in deriving predictions for *local* approximately adiabatically isolated subsystems.

The natural response to both problems is to trade the global probability postulate for a local one. After all, we need a local postulate to accomplish the explanatory task, and the global one is neither necessary nor sufficient for the derivation of the local one. It is not sufficient for reasons we just saw. And it's not necessary because if we have a statistical postulate that allows us to derive thermodynamic generalizations for local subsystems, we don't need one that applies at the level of the world as a whole. And indeed there is a long tradition stemming from Reichenbach (1956) of doing just that. The approach is developed in Paul Davies (1974) and discussed in Sklar (1993, 318- 332). Winsberg, Frigg, and Callender (2011) have all advocated the approach in opposition to Albert.

The proposal is that the isolated systems relevant to SM are what Reichenbach called 'branch systems'. These are systems that are created in low entropy states and exist as energetically isolated systems for a time before being destroyed or merging again with the environment. So, for example, a warm glass of water with an ice cube comes into existence when someone places the cube in the glass, and ceases to exist when the water is consumed or poured out. The proposal is to impose SP at (roughly) the first moment that such a system becomes suitably isolated. The low entropy starting state is explained by whatever procedures or historical accidents led to the creation of the system, and since it didn't exist as an adiabatically isolated subsystem before that moment, thermodynamic generalizations don't apply to it before that moment, and there is no question of whether the system has evolved into the current state from a higher entropy state.

So the proposal now being considered is to do away with this probability distribution over worlds as wholes. We retain the statistical postulate in a form that applies to local, adiabatically isolated subsystems (a random selection of systems in an X macrostate will be distributed among compatible microstates in a manner proportional to the volume of X-space occupied by those microstates), where it is treated as is an empirically

testable, statistical fact that describes the percentage of any not too small, arbitrarily selected set of such systems in a given macrostate that typically fall on abnormal phase points.^{xxii} This fact about the distribution of microstates in arbitrary samples of systems in a given macrostate can then be used with the Newtonian laws to generate the macrodynamics embodied in thermo-gens.

Albert's own attitude toward SP suggests some ambivalence between the local and global versions of the postulate. In the official presentation he talks of a distribution over initial conditions and is quite clear that he intends a global postulate,^{xxiii} however, at various places, the exposition waffles between global and local versions of postulate. And in some parts he clearly seems to have a local postulate in mind. Discussing the objectivity of the postulate, he writes:

"It does seem to be some sort of a fact – or at any rate it seems to yield correct predictions to suppose that it is some sort of a fact – that the percentage of any large collection of randomly selected *x*-systems whose microconditions lie within any particular subregion of the *X*-region of the phase space will be more or less proportional to the familiarly defined volume of that subregion."^{xxiv}

By *X*-systems, here, he must have in mind local subsystems of the world, for any talk of percentages in randomly selected *worlds* is not a fact of the kind about which we could have any experimental knowledge.

The flipping back and forth between the global story and the local one, which obscures the difficulties with the Boltzmannian story, is part of a tendency that we see more widely in foundational discussions. It has two effects, both of which are in evidence here (i) it conceals explanatory lacuna that need to be filled, and (ii) it allows us to uncritically extend notions that have a sensible interpretation at the local level to the global level, where they lose any empirical or operationalizable significance. The notions of a statistical probability distribution over worlds as wholes and the notion of global entropy are both examples of this. And the waffling between the local and global accounts mirrors the waffling that Albert himself rails against, between epistemic and statistical interpretations of SP. Global probabilities make perfect sense as epistemic probabilities, but little sense as statistical probabilities. Statistical probabilities make perfect sense as part of the explanation of mechanical phenomena, epistemic probabilities do not.

The move to a global formulation wasn't merely a whim. It was designed to overcome what would otherwise have been pressing questions for a Boltzmannian account. Resolving the inconsistency between the future and retrodictive probabilities is something it does by *fiat*: by imposing the statistical postulate once, far enough back in history to avoid multiple impositions of a past hypothesis to the same system at different times. It was a transition that Boltzman made in his famous 'lunge for totality'.^{xxv} The branch system proposal avoids the problem by restricting application of SP to the initial state of a branch system and remaining silent about the history of the particles that make up that system. Since those particles didn't form an adiabatically isolated subsystem before they branched off, thermodynamics is simply inapplicable to them and so no inconsistency arises. While they exist, and until they become energetically open on a macroscopic scale, thermodynamics works very well to predict their behavior.^{xxvi}

Albert, of course, is well aware of the approach. He discusses it briefly in chapter 4, lobbying a set of objections in quick succession: First, it is impossible to specify the precise moment at which a particular system comes into being; that is, we cannot specify the precise branching point. Second, there is no unambiguous way to individuate the system. Why does the system in question consist of the glass with ice, rather than the glass with ice and the table on which the glass stands, or the glass and ice and the table and the person watching it, or any of the indefinite number of subsystems of the world that include the glass and ice cube. The vagueness isn't particularly problematic since thermodynamics isn't a fundamental theory. And the obvious answer to the question of why thermodynamics applies to the glass of water with the ice-cubes, but not the other systems in which it is contained, is that the former but not the latter forms an (approximately) adiabatically isolated unit. Of course it is not impossible that there isn't some larger system in which the glass of water is embedded that is also adiabatically isolated, indeed, at the very least, the universe as a whole constitutes such a system, but then the question becomes about whether we have some assurance that an SP applied to the larger system won't conflict with an SP applied to these smaller systems embedded in it. And indeed, this seems to be the worry Albert has in mind when he writes:

"Serious questions would remain as to the logical consistency of all these statistical-hypotheses-applied-to-individual branch-systems with one another, and with the earlier histories of the branch systems those branch systems branched off from" (89)."

I don't know that I feel the force of this worry. It would be odd if we couldn't piece together local subsystems that obey SP into a larger system that also obeys it, and it is difficult to summon a positive reason for thinking that there would be an inconsistency here.

An explanatory lacuna opened up by the revision

Let's take stock. I have argued that thermodynamic phenomena involve generalizations about the behavior of local, approximately adiabatically isolated subsystems of the world. I have argued that if our goal is to make explicit the facts about our world in virtue of which thermodynamic phenomena arise, we have to write down facts about our world that are sufficient for the logical derivation of thermodynamic generalizations. I have argued that a global statistical postulate plays no role in the derivation of thermodynamic generalizations. I have argued that in order to derive thermodynamic generalizations, we need a local version of SP that applies to the kinds of subsystems that fall under thermodynamic generalizations. And I have argued that the local SP does not follow from the assertion that the universe as a whole is on an entropy increasing trajectory.^{xxvii} Now we are considering the proposal to revise the Boltzmannian story by taking the local SP on as an independent assumption, together with an explicit prescription about the systems to which it applies that guarantees that we won't have overlapping, inconsistent probabilities. Would that provide the micromechanical underpinning for thermodynamics that was sought? Callender (2011), for examples, suggests that we should just treat thermodynamics in the way that we treat other special sciences, treating SP as a boundary assumption that defines the conditions under which thermodynamics applies. On this view, thermodynamics 'kicks in' when the statistical postulate is satisfied. The complete explanation for thermodynamic behavior is provided by the microphysical laws and SP, but no explanation is provided for why SP holds when it does.

I am now going to argue that this isn't a stable resting place. Callender's suggestion that SP should be accepted as a boundary condition that defines a special science effectively says that whenever the local SP holds, we get thermodynamic behavior. But the fact that SP holds for all of these apparently unrelated subsystems isn't the sort of thing that can be accepted as an explanatory primitive in a global theory. The existence of a well-defined quantity $pr(\text{macrostate } m | \text{macrostate } M)$ for arbitrary m and M defined over local subsystems is not a given. It depends on a randomization over fine-grained variables relative to the macro-partition. And this effective randomization – if it is not accidental, but counterfactual supporting and lawlike – needs to be *derived*.^{xxviii} Suppose that ensembles with different distributions of black and white balls behave in very different ways and imagine a large, perhaps infinite urn of black and white balls. If you want to know how an arbitrary sequence of 10 or 100 balls drawn from the urn will behave, you have to know the distribution of black and white balls in the sequence, and there is no general *a priori* reason for thinking that the distribution in a 10 or 100 ball sequence is likely to be roughly similar to that in another, nor is it likely to reflect the distribution of black and white balls overall. That black balls are extremely rare overall, for example (or, in an infinite urn, over the first hundred gazillion years), is perfectly compatible with the claim that it is nothing but black for 100 years, then alternating black and white for 50, and later all white. The global facts tell us almost nothing about what we should expect of arbitrary subselections from the urn. There is an indefinite number of ways in which balls can be distributed. The same points hold for infinite urns even more clearly, and even more strongly. In that case, there will be no overall distribution, and an uncountable number of ways in which balls can be distributed. We can always find overall proportions, averages, or medians, ... but none of these quantities would be operationizable as probabilities.

Statistical methods will be successful if an arbitrary macroscopic ensemble of microparticles will contain particles in a random selection of microstates. What we need an understanding of, here, is not so much the particular value of the quantity $pr(m/M)$, but its *existence*.^{xxix} We need an explanation of the stabilization of relative frequencies across samples of systems in a given macrostate. That is a condition *sine qua non* of the assignment of well-defined probabilities. Statistical inferences succeed, to whatever extent they succeed, because there is a de facto randomization over the variables that determine microstate, relative to macroscopic constraints. In a deterministic setting, that de facto randomization is what needs to be explained.

The macrodynamics will give the dynamics for macroscopic ensembles of particles, and there has to be a reason for thinking that, in general, the microstates of particles in such ensembles will be evenly distributed across the volume of phase space associated with those microstates *reliably* and *stably* so from one to the next. And there is no *a priori* reason that there should be stability in the distribution of microstates of local subsystems in an arbitrarily selected set of systems in a given macrostate. It is not as though we can take any old partition P of phase space and end up with a well-defined probability that a system in a chosen cell of that partition is in a given microstate. Partition the set of persons according to their country of origin, for example, and ask whether there is a stable statistical distribution of country of origin across spatially localized ensembles of people. The answer is clearly not. The distribution will vary from one country to the next, being dominated by Norwegians in Norway, French people in France. Partition the set of persons according to age and ask yourself if there is a stable statistical age distribution in local ensembles? Again, clearly not. Some groups will have lots of young people and some will have lots of old people.

The objective dynamical facts embodied in thermodynamic generalizations are that for any macroscopic ensemble of Newtonian particles, if M_0 is a low entropy initial state, and we consider the macrostate at t_1 that results from dynamical evolution, the phase points that originate in M_0 evolve into, and appear to be uniformly spread over the region M_1 at t_1 . This will follow from an assumption to the effect that the distribution of microstates represented in any such ensemble is spread uniformly over the phase points in M_0 . But again, there is no *a priori* reason this assumption should hold. If we want to understand why thermodynamics applies to *this* glass of ice water, *that* gas in an enclosed container, *this* cup of freshly creamed coffee, it suffices to point to the distribution of initial microstates in the ensemble and the dynamics. But that *all* (or very nearly all) such glasses, gases, and cups satisfy SP – and hence that we can have macroscopic *generalizations* that apply to such systems independently of their microstates is something that needs to be explained in a global theory. This sort of statistical fact is a fact about the way that microstates are distributed in our world across samples of systems in a given macrostate, and it describes a regularity that should be explained in a global theory.

There is a general lesson here about the conditions of success for statistical methods in physics. By statistical methods, we mean coarse-graining the phase space, introducing a statistical distribution, and then deriving predictions evolving forward the distribution. Statistical methods will be successful if the probability distribution reflects the actual distribution of microstates in an arbitrary macroscopic ensemble of microparticles. We can't have a very informative coarse-grained method applying to collections of sea creatures drawn from arbitrary 50-gallon volumes of water, because samples drawn from different bodies of water, or different depths, have very different distributions of constituents. We need to know in more detail about their constituents before we have any generalizable conclusions. There just isn't the kind of de facto randomization in selection of constituents that would support an interesting dynamics at that level of coarse graining. One way that might get more detail about their constituents is to add new parameters so that there is effective randomization within the classes defined by those parameters. Or we might be able to find a level of coarse-graining at which we do get stable relative frequencies. But what those parameters are, and what that level of coarse-graining is, will depend on details of the microdynamics. Likewise the distribution of plants over a large-enough region of space, the number of threes in an arbitrarily chosen fifty digit sequence of the expansion of pi, and so on...

The program that tries to derive the local SP from an ergodic hypothesis is the most developed. The usual complaint about the ergodic program is that ergodicity is too weak. Because of its asymptotic character, accounting for the phenomena requires effective mixing to take place, over the relevant macrotime scale, for those phase space regions corresponding to the macrostates at issue. But it is not clear that this is so. The epsilon-ergodicity of Frigg and Werndl^{xxx}, for example, might solve this problem, but this a scientific question that can't be solved without the detailed, specific quantitative work that people involved in that program are doing. The derivation of thermodynamic phenomena depends on a dynamical story about why the local mixing assumption holds. We need some sort of randomization of factors that determine initial microstate within macroscopic constraints. The technical discussion is complex and ongoing, but if we can fill this explanatory lacuna, we will have a genuine dynamical understanding of why we can ignore abnormal phase points. Without a dynamical grounding for the mixing assumption, however, we don't have the counterfactual-supporting generalizations for local subsystems embodied in thermodynamic generalizations. There is nothing to keep the correlations in initial microstate that produce anti-thermodynamic behavior from arising at the local level.

The lesson is that extracting probabilities pertaining to local subsystems from global deterministic dynamics is nontrivial. It is a relatively simple matter to derive a coarse-grained *history* from a fine-grained one. And it is relatively simple to derive a coarse-grained dynamics from a fine-grained one if we are only interested in predicting what is *possible* given a coarse-grained history. But deriving a coarse-grained dynamics capable of making quantitative predictions about the relative probabilities of possible histories is much more difficult. There is a macrodynamics only if there are emergent macroprobabilities, and the existence of emergent macroprobabilities requires stable relative frequencies of microstates for local subsystems in a given macrostate. The existence of such stable relative frequencies is something that needs to be explained. ^{xxx}

A role for the Shuffling Hypothesis

Here, finally, we have found a role for the Shuffling Hypothesis. The Shuffling Hypothesis starts from the observation that abnormal phase points are not only rare but also scattered, which is to say that they are surrounded by normal phase-points. It adds to this a claim about the perturbing influences on the system, namely that these are the GRW collapses, probabilistically distributed as specified by that theory. The combination of these two features – the distribution of normal points, together with the perturbations – has the result that for *every* p , the probability is nearly 1 that the system, when started out in p , by a combination of the unperturbed dynamics of the system and the GRW perturbations, will, if it started out in a state of low entropy, end up in a higher one, and if it started out in equilibrium, remain there. In logical terms the Boltzmannian story and the Shuffling Hypothesis are differentiated by their modal implications. The Boltzmannian story tells us that abnormal states have virtually no chance of arising naturally, but that if we prepare systems in abnormal microstates we should observe entropy decrease. The Shuffling Hypothesis entails that no matter what the initial microstate, we would still be overwhelmingly unlikely to observe entropy decrease. It adds to the Boltzmannian story the instability of entropy-decreasing trajectories, making them not only unlikely to arise spontaneously but (virtually) impossible to produce.

Price complained that the hypothesis adds nothing to the Boltzmannian story in the first six chapters of *Time and Chance*, writing:

“It is hard to avoid the suspicion that Albert’s proposal is trying to sell us something we simply don’t need—the physical equivalent of an expensive gadget for choosing our Lotto numbers, which makes no difference at all to how often we win.” ^{xxxii}

In its global version, Price’s criticism is on point. In that case, the Shuffling Hypothesis adds nothing to the Boltzmannian story but a counterfactual robustness that is irrelevant to anything that actually happens in our world. In its local version it gives us a direct, intelligible story about why thermodynamic generalizations hold at the local level. The local dynamics introduces a probabilistic element into the evolution of every local subsystem that does two things. It gives us a mechanism for stabilizing frequencies across subensembles. And together with the instability of abnormal states under perturbations, it explains why no information about the initial microstate that goes beyond that given in the macrostate can be dynamically relevant to predicting subsequent macroscopically discernible behavior. The only way that additional microinformation could make a macroscopic difference is if the initial microstate was abnormal. But spontaneous collapses nudge systems that start out in abnormal states onto normal ones suppressing the kinds of microscopic differences that might otherwise show up. The continual agitation provided by GRW collapses nudges systems away from ordered states, and destroys the kinds of correlations in initial states that lead to anti-thermodynamic behavior. The provision of a dynamical mechanism driving systems towards stable macrostates explains why there is uniform macro behavior despite the fact that – as a purely topological matter – there are trajectories emerging from macrostates that lead to low entropy states that are clearly macroscopically discernible.

In the Boltzmannian story, where m_{ab} is an abnormal microstate and M is any macrostate, the low probability $pr(m_{ab}/M)$ was supposed to provide a rationale in calculations applied to local subsystems for ignoring the possibility of microstates on entropy decreasing trajectories. The problem with that story was that if those probabilities exist for local subsystems, they must follow from the global dynamics, and no reason for supposing they exist had been given. On the Shuffling Hypothesis, the probabilities come from the local dynamics and don’t need to be derived. The features of the Boltzmannian story that I have been contrasting

unfavorably with the Shuffling Hypothesis have to do with the fact that the Boltzmannian story is a theory of the whole universe; the local theory is hard to extract and remains largely conjectural. The Shuffling Hypothesis starts with the local dynamics and effectively *puts in by hand* what the Boltzmannian story has to *derive*. The local probabilities get generated internally by the dynamics of local subsystems. This keeps the probabilities firmly rooted in the local contexts in which they have practical and empirical significance. We might arrive at a distribution over physically possible worlds or initial conditions by extrapolating probabilities defined in the first instance over local subsystems, but such a distribution plays no explanatory role.

One interesting difference between the two is that there is a shift in the order of priority. The Shuffling Hypothesis takes the local laws as basic and the global story is extrapolated. The Boltzmannian story takes the global laws as basic and has to derive the local dynamics, so it owes an account of why the kind of ordered states that produce entropy decrease don't arise. The new story provides an intelligible mechanical account of thermodynamic behavior that, I have argued, we didn't find in Albert's version of the Boltzmannian account.

That concludes our discussion of Albert. I have used Price's question about a role for the Shuffling hypothesis to raise questions about the logic of Albert's account in the first six chapters. We found faults with that account, and the revision that was proposed opened up a space for the Shuffling hypothesis to contribute to a mechanical explanation of thermodynamic behavior (or better to contribute to the production of thermodynamic behavior).

Shuffling without GRW

The Shuffling Hypothesis rests on speculative physics. It stands or falls with GRW. But it provides clues for an entirely classical account. A proposal due to Leeds uses the clues in the Shuffling Hypothesis to fix up the classical branch system version of the Boltzmannian story.^{xxxiii} On this proposal, we give an account of local subsystems that makes the same assumptions about the rarity and fragility of abnormal states that were invoked in the Shuffling Hypothesis. But instead of an *internal* mechanism that introduces a probabilistic element to nudge systems on abnormal phase points onto normal ones, Leeds suggests that we exploit the fact that we are talking about *open* subsystems and look to the sorts of external perturbations that are bound to be present in any real situation. If the idea of a distribution over initial conditions of the world was problematic, there is no problem in talking about a distribution over external perturbations for open subsystems. That gives us an account of local subsystems that is very like the one that Shuffling Hypothesis gives us, except that we have offloaded the source of perturbations from the internal dynamics to the external environment. In this story, external perturbations, rather than internal GRW collapses drive the internal dynamics of local subsystems in the direction of entropy increasing states by nudging them onto normal phase points. We dispense with the need for an assumption about initial states of subensembles in favor of an assumption about the presence of disturbing influences from the environment. In this version of the Boltzmannian story, the strategy is to look for a dynamical explanation of the emergence of local probabilities, from an assumption about how things behave under random perturbations from outside, akin to the explanation of the emergence of probabilities for coins and gambling devices. Open subsystems that show regular statistical behavior can be thought of as physical mechanisms for turning microscopic noise into macroscopic information. A random distribution of initial microstates compatible with the macroscopic constraints will produce the familiar distribution of heads and tails. Consider the Newtonian model of coin tosses due to Diaconis.^{xxxiv} One might think that, since coins are deterministic systems, any stability in the output would have to derive from stability in the input. If that were the end of the story, we would have pushed stability in the distribution of heads and tails conditions in the sampling of initial states to produce stability in sampling of final states. A sampling of initial conditions will deterministically produce a dispersed sample of final conditions, and any quantifiable regularities in the outcome of tosses would be a product of the distribution of initial microstates produced by whatever selection procedure we used. But that's not how the story actually works. Even though the microdynamics is deterministic, we notice that edge-landings are unstable. Unless perturbations are perfectly balanced so that every push from one side is matched by an equal, opposite, and simultaneous push from the other, the coin will be pushed onto one or the other of the stable bands. For a fair coin, roughly half will terminate in heads and half will terminate in tails. In weighted coins, the distribution will be different. But in virtually no case would we expect tosses to terminate in edge landings. Two things stabilize probabilities of heads and tails here: the randomness of the

initial distribution relative to macroscopic constraints and the instability of all but flat landings. The local dynamics *adds* something that *produces* regularity that is not present in the environment, and that turns out to be tremendously useful, because it provides us with handy chance devices that can decide practical dilemmas. Should I have eggs or oatmeal? Ask Sandra or Sam? Do it now or later? A coin flip will decide.

The model explains the emergence of quantifiable regularities in the behavior of coins under tossing in a manner quite similar to the proposal that Leeds is making to fill the Boltzmannian lacuna (the analogy is his). We build dynamical information into the structure of the phase space and add perturbations, and we end with bands of states that are stable under perturbation. The same thing that makes edge-landings almost impossible makes it very difficult to sway the output of a fair coin away from a 50-50 distribution. If we could predict or control the perturbations, we could predict or control the outcome of individual tosses. But in practice, we have no more knowledge of, or control over the exogenous variables that determine coin trajectory than we do over the outcome of a GRW collapse. So this story rules out edge-landing and stabilizes probabilities for coins as surely as a genuine chance mechanism would, at least in a world in which there are no Maxwell-demon-type agents.^{xxxv} Even if we could control the initial microstate of the coin quite precisely, to control its final state, we would need to know in quite precise detail, what external perturbations would be present to control the outcome of a particular toss. In the absence of that kind of Maxwell-demon-style control of the trajectory determining-microvariables, we can expect a roughly 50-50 distribution in any appreciably large ensemble of tosses. The impossibility of controlling the outcome falls not on the imprecision of preparation of initial state, but the impossibility of controlling the external perturbations.

So how does this story run in the thermodynamic case? Again, we exploit the rarity and fragility of abnormal states, which played a large role in the Shuffling Hypothesis, but now we add uncontrollable external perturbations pushing trajectories onto stable states and filtering out trajectories that would otherwise lead to entropy decrease.^{xxxvi} The only difference is that the shuffling comes from the outside rather than the inside. Of course, the distribution of perturbations that was exogenous from the point of view of a local subsystem is endogenous in the global theory. But the need for a *distribution over initial conditions gets discharged in the global story* in favor an assumption that the perturbing influences on any given system are *not* perfectly balanced to preserve the sorts of correlations present in the initial state of an abnormal trajectory. The sort of coincidence would be required to preserve an entropy-decreasing trajectory would involve a collection of microscopic particles that would have to be perfectly aligned by accident to preserve the correlations present in an abnormal trajectory. The intuition is that the initial state of the universe would have had to been aligned just exactly right, carefully and delicately balanced to preserve the correlations present in an abnormal state. The only way for such a balancing act to arise in a Newtonian world is by the accidental confluence of a disparate set of dynamically unrelated events. There are no (non-coincidental) dynamical mechanisms for producing the balancing act.

This is just a sketch of potential line of development, but it provides an example of the *kind* of filling in that I argued is needed in the Boltzmannian account to turn it into an intelligible account of why our world behaves thermodynamically at the local level, i.e., to derive thermo-generalizations for local subsystems, but its feasibility will depend on the details. The important feature of the story is that a statistical postulate defined over worlds as wholes is no part of the global dynamics. A statistical postulate defined over local subsystems plays an ineliminable role in the local dynamics, but gets derived at the global level.

What the Leeds proposal and Shuffling hypothesis both do is provide mixing mechanisms to give us a much better physical understanding locally of why we don't observe anti-thermodynamic behavior. This kind of pervasive regularity can be traced to the denial of 'special initial conditions', but our intuitions about what initial conditions count as 'special' or 'ordered', I would suggest, is really derivative of our expectation that regularities among disconnected exogenous influences on local subsystems do not occur without explanation. That is an expectation that is rooted in our experience with such subsystems, not in a primitive measure over initial conditions.^{xxxvii}

I have argued that the statistical postulate applied at the global level doesn't derive what needs to be derived, which is thermodynamic generalizations in application to local approximately adiabatically isolated subsystems of the actual world. A local version of the statistical postulate does play a role in the derivation of thermodynamic generalizations, but it is not a free-standing postulate. If it is true at our world, it follows,

like every other truth, from some combination of law and accident. I suggested that this gives us an explanatory role for the Shuffling Hypothesis and an answer to Price's question: what is the Shuffling Hypothesis supposed to do? The Shuffling Hypothesis provides the internal randomizing mechanism that stabilizes relative frequencies at the local level, filling an explanatory lacuna in the Albert's account. Finally, I described an alternative way of filling the lacuna due to Leeds.

It is illuminating to look at the regularity from the point of view of the microtheory. From that point of view, since we know that there *are* initial microstates that would produce macroscopically different behavior if left to evolve without disturbance, what needs to be explained is why the sorts of differences in initial microstate that would produce macroscopically different behaviors in these kinds of systems don't manifest themselves. An unilluminating answer would be that the initial conditions of the universe were just set up so that they haven't arisen. That by itself would give us no reason for supposing they wouldn't arise in the future. A more illuminating answer – and one that would underwrite a continued expectation for thermodynamic behavior, would be that there are dynamical mechanisms in place that ensure (or come close to ensuring) that ensembles that form adiabatically isolated subsystems have a way of 'forgetting their microscopic histories', by which we mean that information about their histories encoded in correlations among the initial microstates gets wiped out. There is a kind of continual erasure of information. The GRW jumps, or the continuous exogenous mixing, would provide mechanisms for this erasure. Once the mechanism is incorporated, as Albert notes, the need for an SP disappears, and the past hypothesis plays a different role. In the Boltzmannian story, Albert thought that it was needed to keep the story from delivering disastrously wrong predictions about the past.^{xxxviii} The statistical interpretation doesn't work at the global level (or if it does, it is empirically empty. It tells us something about the composition of random subselections from the class of physically possible worlds. Whatever content you want to give that, it is mostly extrinsic to the actual world). If the probabilities are applied to ensembles and interpreted statistically, then again, but for a different reason, we have very little in the way of predictive content. If we look at a collection of particles that currently comprise a local adiabatically isolated subsystem (a gas in an enclosed container, a glass of water with some ice cubes in it), in the past, those particles didn't form an isolated subsystem. The 'initial conditions' of the predictive procedure are the final conditions of a retrodictive procedure, and the backward evolution is not just the temporal reverse of the predictive procedure because the forward evolution is the evolution of a (quasi-) isolated system. The backwards evolution is nothing of the sort. This is one of the things that can be obscured in moving seamlessly from the global to the local story. The set of particles that comprise the ensemble form a system in the past, but only nominally. Every one of those particles was following a different trajectory, subject to innumerable, unaccounted for influences that brought it to its present state. Without information about those influences, even if we knew the current microstate of the ensemble, we know very little about its past. And if the final *microstate* won't give us predictions about the past of the ensemble, neither will its final *macrostate*. Information about the *global* past falls outside the predictive scope of the theory, and it just fills in a part of our understanding of the history of the world that this part of the theory is silent on. The explanatory structure of the GRW story is well-behaved. It is an intelligible story, and it can be told entirely in mechanical terms in a way that I don't think the story in the first six chapters can.^{xxxix} It is a straightforward factual question what the real mechanical story is, and the answer can only be settled in scientific terms by experiment. But both of these stories provide intelligible mechanical explanations of thermodynamic phenomena.

Albert, comparing the branch systems approach with the GRW approach, says this: "whatever constitutes the environment of these two bodies evolves in accord with precisely the same sorts of deterministic dynamical laws as the constituents of the bodies themselves do, then whatever 'randomness' there is in the perturbations arising from interactions with that environment can only have gotten there in virtue of precisely the same sort of probability-distribution over that environment's initial conditions that we have been dealing with throughout this book. And so the whole exercise gets us nowhere."^{xl}

But here we might say 'well, no'. The randomness got there by some combination of the initial conditions at our world and global dynamics, and so what we really need to understand is what combination of facts about initial conditions and global dynamics would generate the kind of randomness we need in the environment of open systems. There are two ways to explain this that aren't entirely distinct: one places emphasis on the dynamics, showing that there are local mixing mechanisms that ensure or come close to ensuring that the exogenous influences on these kinds of subsystems is typically random by the standard measure, and one that looks at what initial conditions would have to be like to produce anti-thermodynamic behavior and appeals to

our sense that things would have had to have been set up very specially at the beginning of the world to orchestrate thermodynamic behavior at the level of the open subsystem. I suggested that the explanation that points to local mixing gives us a somewhat better physical understanding, by dispensing with a primitive sense of specialness over initial conditions and highlighting the local mechanisms that make anti-thermodynamic behavior atypical.

The question that worried Albert – i.e., the one that motivated the global move, and the one that bothered Callender – is why should SP work at one point in a system’s history and not earlier? Or, to put it a little differently, thermodynamic methods become applicable after you impose an SP. But what justifies its application when we do impose it? If we impose it at one point in the history of a system, what about before? As Callender put it “are the microstates suddenly supposed to scramble themselves when the particles come together to comprise a cup of coffee?”. The Shuffling Hypothesis and the Leeds proposal would both provide a good mechanical explanation at the local level for why we can safely impose SP to the local, adiabatically isolated subsystems at the start of their lives, i.e., why the microstates of such systems are bound to be distributed ‘evenly’ by the standard measure across those compatible with its macrostate, regardless of their microscopic history. The explanation is that when we are dealing with an adiabatically isolated subsystem, there are two dynamical rules that are working together: there is the ordinary deterministic Newtonian evolution from the initial microstate, but there is also a disturbing influence that is pushing the system off of abnormal phase points at every step of the evolution, and these work together to ensure thermodynamic behavior. Albert rejects the idea that local perturbations could do the work of the Shuffling Hypothesis. He writes:

“It has often been suggested in the literature ...that the sorts of perturbations we were talking about above are already all over the place, if one simply stops and looks in (say) the Newtonian picture of the world. The idea is that since none of the macroscopic two-body systems of which we have ever had any experience and none of the macroscopic two-body systems of which we ever will have any experience are genuinely *isolated* ones, the perturbations in question can be seen as arising simply from the interactions of the two-body system we’ve been talking about here with *its environment*. But if (as these authors always suppose) whatever constitutes the environment of these two bodies evolves in accord with precisely the same sorts of deterministic dynamical laws as the constituents of the bodies themselves do, then whatever ‘randomness’ there is in the perturbations arising from *interactions* with that environment can only have gotten there in virtue of precisely the same sort of probability distribution over that environment’s initial conditions that we have been dealing with *throughout this book*. And so the whole exercise gets us nowhere.” (p. 153)

Of course he is exactly right to say that in a deterministic setting there is a direct entailment from the initial conditions and laws to everything that happens, including thermodynamic phenomena. No matter what, the initial conditions are going to have to have been set up in such a way as to produce thermodynamic generalizations when evolved by the laws. What we are looking for is some illumination about the difference between those initial conditions that do and those initial conditions that don’t. And one way to do that is to look at the kinds of initial conditions that would have had to be in place to produce thermodynamic behavior. The best thing would be if we can get some local mechanical understanding of what is going on with local subsystems, since their behavior is a product of the same microlaws. But now the initial conditions are local states of parts of the world that arise all over the place across its history, so we need an understanding why the initial conditions of these systems obey SP. That can be a quite specific story about how the gross constraints that define this system were set up (how these ice cubes came to be dropped into this glass, how a partition was removed from the container of this gas... and so on). So what really needs to be explained is stabilization of the distribution of microstates across all of these scattered systems with their different histories. What we need to understand is not why this system evolves thermodynamically (something that is explained by its initial microstate and the laws) or that one or does. What we need to understand is why *all of these systems* – as diverse as they themselves and their histories are – exhibit thermodynamic behavior.

The Shuffling Hypothesis and the Leeds proposal both provide answers to this question. It is an entirely factual question which, if any of these, is correct, but they are good examples of the kind of purely mechanical understanding of thermodynamic phenomena from a microphysical perspective, something that I just don’t see how Albert’s story, with the global version of SP does.

Appendix: statistical and epistemic probability

What follows are some remarks about the interpretation of probability. We introduce statistical probability in a manner that is purged of epistemic association and defined in operational terms.

We distinguish general (or indefinite) from single-case (or definite) probabilities. General probabilities pertain to classes and the basic form is conditional. Single-case probabilities pertain to particular events. General probabilities are related to single case probabilities by the principle that the single case probability of an x that is randomly selected from the population of y 's is the general probability of x , given y . These definitions license an inference from a general probability to a single case probability where the selection procedure is random.

Statistical probabilities are general probabilities. What makes them *probabilities* is that they obey the Kolmogorov axioms.^{xli} What makes them *statistical* is a connection to statistics. The connection to statistics comes from the law of large numbers (in its weak or strong version) which says that the relative frequency of x in a class of randomly selected y 's approaches the general probability of x , given y , with increasing probability, as the class gets larger. This secures an evidential relationship between frequencies and probabilities, but it blocks reduction because it leaves open the possibility that the actual frequencies may diverge as far as you like from the probabilities.^{xlii} We might say that the connections between probabilities and statistics *interpret* the probabilities without reducing them.

In operational terms, statistical probabilities are elements in a theoretical matrix that mediate inference from observed statistics in local samples to others. Probabilistic notions are connected to actual observed frequencies by the body of operational procedures and norms by which we infer probabilities from collected statistics. The operational procedures are embodied in the canons of statistical inference. There is no compact simple rule or algorithm for statistical inference. Nothing short of the canons of statistical inference and all of the formal and informal rules that are part of knowing how to apply them will suffice to relate probabilities to statistics.

Assignment of statistical probabilities commits one to the expectation that the statistics for one sample will reflect those of others provided that no selection was exercised either in the collection of the sample or in those whose statistics it reflects. The expectation grows as the size of the sample increases, and is defeasible by the belief that the selection was biased, or that the selection process was not random. We don't reify probabilities. To say that there are statistical probabilities means nothing more than that there are stabilized relative frequencies with the right kind of subensemble structure.^{xliii}

Statistical probabilities contrast with epistemic probabilities. Where statistical probabilities are interpreted by a connection to frequencies, epistemic probabilities are interpreted by a connection to belief. What makes them *probabilities* is (as with statistical probabilities) that they obey the probability axioms. What makes them *epistemic* is a connection to belief. There is a descriptive form of epistemic probabilities according to which they represent the epistemic states (or degrees of belief) of agents. And there is a normative form, according to which they *guide* the beliefs of agents.

The connection between statistical and epistemic probabilities is mediated by a principle like Lewis' Principle Principal.^{xliv} The correct form and status of that principle is a matter of controversy, with some holding that it is a rational norm, and others giving a pragmatic justification that is mediated by a substantive epistemological theory.

References

Albert, D. Z. (2000), *Time and Chance*. Cambridge, Mass.: Harvard University Press.

Boltzmann, L. (1964), *Lectures on Gas Theory*. Berkeley: University of California Press.

Boltzmann, L. (1974), *Theoretical Physics and Philosophical Problems: Selected Writings*. Vol. 5. Dordrecht; Boston: Reidel Pub. Co.

Brown, H., Myrvold, W., and Uffink, J. (2009), Boltzmann's H-Theorem, its discontents, and The Birth of Statistical Mechanics, *Studies in the History and Philosophy of Modern Physics* 40, pp. 174 – 191.

Callender, C. (2011): The Past Histories of Molecules, in Claus Beisbart and Stephan Hartmann (eds) *Probabilities in Physics*, Oxford University Press, 2011.

Callender, C. (2010), The Past Hypothesis Meets Gravity. Forthcoming in Gerhard Ernst and Andreas Hüttemann (eds.): *Time, Chance and Reduction*. Philosophical Aspects of Statistical Mechanics. Cambridge: Cambridge University Press.

Callender, C. (2004), Measures, Explanations, and the Past: Should 'Special' Initial Conditions be Explained? *British Journal for the Philosophy of Science* 55, 2:195-217.

Davies, P. C. W. (1974), *The Physics of Time Asymmetry*. Berkeley: University of California Press.

Diaconis, P. (1998) "A Place for Philosophy? The Rise of Modeling in Statistical Science." *Quarterly of Applied Mathematics* 56: 797-805.

Earman, J. (2006), The Past Hypothesis: Not Even False. *Studies in History and Philosophy of Modern Physics* 37, 3:399-430.

Emch, G. G., and Liu, C. (2002), *The Logic of Thermostatistical Physics*. Heidelberg: Springer.

Frigg, R. (2008a), A Field Guide to Recent Work on the Foundations of Thermodynamics and Statistical Mechanics. In *The Ashgate Companion to the New Philosophy of Physics*, edited by D. Rickles. London: Ashgate.

Frigg, R. (2008b). Typicality and the approach to equilibrium in Boltzmannian Statistical mechanics. Available online at <http://philsci-archive.pitt.edu>.

Frigg, R. and Werndl, C (2013), <http://arxiv.org/abs/1310.1717>

Hájek, Alan, "Interpretations of Probability", *The Stanford Encyclopedia of Philosophy* (Winter 2012 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/win2012/entries/probability-interpret/>>.

Hall, Ned (1994). "Correcting the Guide to Objective Chance". *Mind* 103:504–517.

Leeds, S. (2003), Foundations of Statistical Mechanics—Two Approaches. *Philosophy of Science* 70,126–144.

Lewis, D. [1980]: 'A Subjectivist's Guide to Objective Chance', in R. C. Jeffrey (ed.),

Studies in Inductive Logic and Probability, Vol II, Berkeley, University of California Press, pp. 263-93,

Lewis, D. [1994]: 'Humean Supervenience Debugged', *Mind*, 103, pp. 473-90.

Price, Huw review of Albert, <http://www.usyd.edu.au/time/price/preprints/albertreview.html>

Reichenbach, H. (1999), *The Direction of Time*. Edited by M. Reichenbach. Mineola, N.Y.: Dover.

Sklar, L. (2006), Why does the Standard Measure Work in Statistical Mechanics? In *Interactions. Mathematics, Physics and Philosophy, 1860–1930*, edited by V. F. Hendricks, K. F. Jorgensen, J. Lützen and S. A. Pedersen. Dordrecht, Holland: Kluwer.

Thau, Michael (1994). “Undermining and Admissibility”. *Mind*

103:491–503.

Wallace, D. (2010), Gravity, Entropy, and Cosmology: In Search of Clarity, forthcoming in *British Journal for the Philosophy of Science*.

Wallace, D. (?), “The Logic of the Past Hypothesis”, this volume.

Winsberg, E. (2004a), Laws and Statistical Mechanics. *Philosophy of Science* 71, 5:707–718.

Winsberg, E. (2004b), Can Conditioning on the Past Hypothesis Militate Against the Reversibility Objections? *Philosophy of Science* 71, 4:489–504.

ⁱ I’d like to thank Toby Handfield and Alastair Wilson for organizing a small workshop on these topics at Monash University, with a special debt to Alastair Wilson for very pressing comments on an early version of this paper.

ⁱⁱ Put crudely, thermodynamic phenomena are compatible with the laws, we know that there are some ways that the initial conditions could be do generate thermodynamic phenomena and some that don’t. And since our world is one in which thermodynamic phenomena arise, the initial conditions at our world must be among the latter.

ⁱⁱⁱ A good deal of critical attention has gone to the status of the Past Hypothesis. I will remain uncritical about that component of the story.

^{iv} See this volume and references.

^v Earman (2006), Frigg (2008a).

^{vi} The norm in discussions in foundations of statistical mechanics is to focus on theorem proving and barrage the reader with technical results. There is a tendency for logical clarity to get buried by mathematical detail and Albert’s departure from the norm is part of his uncompromising insistence on logical clarity. Theorem proving only does any good if the pieces of the account can be assembled into an intelligible explanation of thermodynamic behavior. As David Wallace makes a nice remark in this context that the job of philosopher of physics is not to prove theorems but to say which theorems are worth proving.

^{vii} The use of the word ‘prediction’ can conceal this, since we ordinarily think of a theory’s predictions (and retrodictions) as part of the empirical content of the theory. In this case, however, the ‘predictions’ are probability distributions which don’t represent what does happen – which is always one thing – but theoretically recommended credences over a range of possibilities.

^{viii} Although the interpretation of probability in other areas of philosophy is generally contested, this is *statistical* mechanics; SP is introduced as an explicitly *statistical* postulate. See the Appendix for a few more remarks on the interpretation of probability.

^{ix} There are other motivations one might have of course. One might want to establish the *compatibility* of thermodynamic generalizations with underlying theory. Or one might aim to show that thermodynamic laws, or the ontology of thermodynamics reduces to microphysics. But the compatibility of thermodynamics with Newtonian mechanics (which furnishes the relevant microtheory in the first six chapters) is not at issue and Albert disavows reduction in any stronger sense than reproducing thermodynamic experience’.

^x One might think that this isn’t exactly true, because if the class contains no object that is F, I can conclude that the member I am interested in is not F (and conversely, if all members of the class are H’s, I can conclude that the member I am interested in is H). And if the class contains a single member, I know that it does, and I know that it contains a J, then I know that the member I am interested in is a J. And so on. So one might think that the statistical probability that the initial conditions at our world are at least among those that have a non-zero statistical probability. But even that is not true, since having a zero statistical probability is entirely compatible with occurring.

^{xi} We can restate it in terms of a distribution over initial microstates of our world: either the initial microstate of our world is typical, or it is not. If it is, then thermodynamic generalizations are explained by those conditions together with the laws. If they are not, there is nothing to explain. Either way, what is true of other worlds is a red herring.

^{xii} Here we go from a statistical postulate (which tells us something about the distribution of trajectories in the class of physically possible worlds) to an epistemic probability for the trajectory of our world. That move – which we make seamlessly in everyday life – is part of what I suggest lends illusory explanatory import to the probability.

^{xiii} This is the sense in which that the probability distribution plays a metaexplanatory role acting as a gauge of how much explanatory work is done by the different elements in a derivation of thermodynamic generalizations (how much is explained by the laws and how much is explained by the assumptions we have to make about initial conditions to rule out entropy decreasing trajectories locally, given P).

^{xiv} One way of putting the issue is that a purely statistical interpretation of SP is possible if it is applied locally. Applied at the global level, the statistical interpretation makes the content of the probability distribution mostly extrinsic to the actual world. An epistemic interpretation of the probability postulate makes sense. On this interpretation the probability is expressing partial information about the initial conditions at our world. That is why its import is screened off by knowing what the initial conditions are. But clearly the epistemic interpretation cannot do any work in the mechanical explanation of thermodynamic behavior. That work is done by the initial conditions about which it encodes information.

^{xv} It can't be seriously taken to be assessing how well Newtonian dynamics is confirmed by thermodynamic phenomena, for then, the relevant probability is not that of thermodynamic generalizations, conditional on PH, but thermodynamic generalizations, unconditionally (since PH is a contingent hypothesis that is necessary, but not sufficient, to generate thermodynamic generalizations). By that measure thermodynamic phenomena assign Newtonian Mechanics a very *low* likelihood.

^{xvi} The numbers are important for deriving quantitative results, but for our purposes, it will do simply to connect the probability assigned by the standard measure to the randomness of the distribution of microstates across constituents.

^{xvii} When we say the trajectory that a system follows through its state space; we mean the trajectory that is traced by the time evolution of its state through its state space.

^{xviii} It is true that if the train is finitely long, the runners will eventually hit the end and begin moving to the south with the train, but there is no limit to how long that could take, and no reason that it could not take longer than the life of the system.

^{xix} We would need an historical account of how each of these systems came to be, but that story would be highly individual in each case. It would involve the story of how Gina put a particular tray of water in a freezer, and later took the cubes from the freezer and put them in a glass of warm water, or of how Gus prepared a gas in an enclosed container and then removed a barrier allowing it to expand, or these are approximately adiabatically isolated systems in low entropy macrostates: fresh ice cubes in warm glasses of water, gases in closed containers from which barriers have been removed, and the like.

^{xx} One can challenge the intelligibility of a probability distribution over worlds, if what one means by 'worlds' is not components of an actual multiverse, but a set of universes, only one of which is actual, and if this distribution is not simply an extrapolation from a distribution over local subsystems of the world obtained by mixing, or obtained from dynamical symmetries (or an ergodic hypothesis or something that ties it logically to observed regularities in the actual world), but is rather presented as a logically independent hypothesis that serves as an explanatory primitive, then I don't know what it means. Physicists employ distributions over worlds in a number of ways, most of which can be rendered empirically intelligible by one of the means above.

^{xxi} Earman ().

^{xxii} Empirically testable, of course, means that it receives evidential support from observed frequencies, though not definitive proof. Such is the epistemology of statistical claims. The difficulty of verifying a statistically probability claim is of an entirely different order than the difficulty of verifying a probability distribution over worlds. In that case, there is no frequency. We have only a single case.

^{xxiii} Here, for example, and in numerous other places: "Why in God's name bother with all this when the uniform probability-distribution over the possible microconditions compatible with the macroconditions of the world, at the moment when it came into being, will very straightforwardly give us everything we need?" , *ibid*, p. 89.

^{xxiv} *Ibid*, p. 65.

^{xxv} Callender (2011) has an excellent discussion of the history, and describes the global move thus: "We have gained consistency with the mechanics by pushing SP back to the beginning of the universe. For the global state of the universe now, Let the Past State (the state the Past Hypothesis stipulates) be at time t_0 . The new measure for some later time t , where $t > t_0$, is the old SP except changed by conditionalizing on both the current macrostate and the Past State. That is, consider all the microstates at time t compatible with the current macrostate and the Past State macrostate. These microstates form a set, a proper subset of the set of microstates compatible with the current macrostate. Intuitively put, these are the microstates compatible with the current macrostate that allegedly don't give us reversibility headaches. Put the uniform (Lebesgue) measure on this restricted set. Then the chances that any microstate is in region A is determined the usual way, i.e., by the proportion of the measure of A with respect to the measure of this restricted set. Let's call this new probability principle SP*. Stepping back from the details, we seem to have achieved our goal. We wanted to know how SP could be compatible with the dynamics. The answer is twofold: one, push back the static probabilities to the beginning of the universe, and two, replace the static probabilities used at later times with a new probability, that encoded in SP*, one that makes essentially the same probabilistic predictions for the future as the old static probability (presumably, since there is no "Future Hypothesis")."

^{xxvi} Frigg (2008) and Winsberg (2004a)

^{xxvii} And if it doesn't follow from that claim, nor does it follow from the weaker claim that there is a high probability that the universe is on entropy increasing trajectory.

^{xxviii} Like every fact in a deterministic world, this kind of regularity could be derived from initial conditions plus laws. But that would leave a set of unanswered questions: How robust is the regularity under differences in initial conditions? Do we have any reason for thinking it will persist, or might we start seeing lots of anti-thermodynamic behavior tomorrow? We are looking for that sweet spot, where we see that there are general mechanisms in place that are bound to produce thermodynamic behavior no matter what the initial conditions, or if not that, then under almost all choices in initial conditions.

^{xxxix} I assume that the quantity exists whenever there is a distribution of m 's among arbitrarily chosen subselections from the class of M 's with the right kind of stability. For precise quantitative characterization of 'the right kind of stability', see Hajek (2012).

^{xxx} Frigg and Werndl (2013)

^{xxxi} While questions about what needs explanation and why are heavily contested in philosophical circles, in practice, there is wide agreement in practice among scientists that an emergent, law-like statistical regularity across the initial states of unrelated systems needs a dynamical underpinning.

^{xxxii} <http://www.usyd.edu.au/time/price/preprints/albertreview.html>

^{xxxiii} Leeds (2003)

^{xxxiv} Diaconis (1998).

^{xxxv} Or at least not in a world in which they don't arise naturally and spontaneously or with enough regularity mess things up.

^{xxxvi} And explain departures from the expected distribution in heads and tails by departures from randomness in the initial distribution, while also explaining the incapacity to control the initial distribution in a manner that will produce departures.

^{xxxvii} See Wallace (The Logic of the Past Hypothesis". His account points to the specialness of the initial conditions that would produce anti-thermodynamic behavior for any macrostate.

^{xxxviii} Using 'prediction' here generically to mean entailment of theory. Again, here, there is a slippage between the epistemic and statistical notions of probability. The retrodictive procedure without PH gives us *expectations* that are disappointed, but it doesn't give us *predictions* that are falsified. A prediction is falsified when a theory says 'x did happen', when in fact x didn't. In this case, the theory just tells us that something unlikely (or unexpected) happened. The difference here is the difference between a statement about the world and a recommendation for belief. The two are logically distinct, but so closely connected the slippage has to be carefully guarded against

^{xxxix} What has to be true of the GRW jumps quantitatively for them to play this role is that "the probability of abnormality that follows from every single individual one of the $P_A(B)$'s is roughly equal to the probability of abnormality that follows from the standard statistical-mechanical measure over the *entirety of the macrocondition* within which the A in question happens to *fall*." (p. 156) The same is true of the probabilities derived from exogenous influences.

^{xl} Albert (2000), p. 153.

^{xli} http://en.wikipedia.org/wiki/Probability_axioms

^{xlii} Relative frequencies, of course, are a kind of statistics. There is another reason that Bernoulli's theorem doesn't reduce probabilities to frequencies. It doesn't connect probabilities to frequencies simpliciter, but between probabilities and frequencies in *randomly selected* subclasses from the class of y 's and 'random' is an undefined, explicitly probabilistic primitive. What keeps the probability-frequency connection from collapsing into triviality is that even though we can't define 'random' non-probabilistically we give some empirical content to 'random selection' by operationalizing selection procedures. Defeasibility conditions for random selection are tacitly understood. We mostly know a non-random selection when we see one and it doesn't matter in practice that we can't give an explicit recipe or necessary and sufficient conditions.

^{xliii} It is a virtue of this account of probabilities that we can regard questions about what the subensemble structure is like as dictated pragmatically by what kinds of inferences we regard an assignment of a probability as licensing.

^{xliv} Hall (1994), Lewis (1980,1994), Thau (1994).