From the holdings of the Department of Special Collections , Princeton University Library, Princeton, NJ

These documents are either in the public domain, or can only be used for educational and research purposes (Fair use) as per U.S. Copyright law (text below). By accessing this file, all users agree that their use falls in the public domain or within fair use as defined by the copyrightlaw. We do not charge any permission or usage fees for the publication of images of material inour collections, including those provided by Princeton University Library, via our website, catalog, or directly from staff.For more information regarding copyright, credit, and citation please review the information onthe Copyright, Credit and Citation Guidelines page of our website.
https://library.princeton.edu/special-collections/policies/copyright-credit-and-citation-guidelines. If you have any questions, including how to order higher-quality images from the collection, contact Special Collections Public Services via our Ask Us! Form. https://library.princeton.edu/special-collections/ask-us

---

Doctoral Dissertations, 2011-2026 Philosophy

2014

Entropy and the Asymmetry of Time

Hershey, Joshua

Princeton, NJ : Princeton University

---

http://arks.princeton.edu/ark:/88435/dsp01v692t6340

U.S. Copyright law: The copyright law of the United States (Title 17, §108, United States Code) governs the making of photocopies or other reproductions of copyrighted material. Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the photocopy or reproduction is not to be used for any purpose other than private study, scholarship, or research. If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of fair use that user may be liable for copyright infringement. This institution reserves the right to refuse to accept a copying order if, in its judgment, fulfillment of the order would involve a violation of copyright law. https://www.copyright.gov/title17/title17.pdf

ENTROPY AND THE ASYMMETRY OF TIME

Joshua Nathan Hershey

A DISSERTATION

PRESENTED TO THE FACULTY

OF PRINCETON UNIVERSITY

IN CANDIDACY FOR THE DEGREE

OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE

BY THE DEPARTMENT OF

PHILOSOPHY

Advisors: Adam Elga (primary), Hans Halvorson

January 2014

© Copyright by Joshua Nathan Hershey, 2013. All rights reserved.

Abstract

Many thermodynamic processes are irreversible: ice melts in warm water but never spontaneously freezes out of warm water, air leaks out of a pressurized bicycle tire but never leaks back in, and so on. The Second Law of Thermodynamics provides a high-level explanation for these temporal asymmetries, but a satisfactory explanation in terms of more fundamental laws has been notoriously elusive. I defend a cosmological explanation, according to which the temporal asymmetry of thermodynamics derives ultimately from an asymmetry in the boundary conditions of the universe—namely, the fact that the distribution of mass was quite different at the beginning of the universe than it will be at the end.

Recently, three formidable objections have been raised against this cosmological explanation of thermodynamics. Two of these objections challenge the idea that the specified boundary conditions of the universe place significant constraints on the behaviors of mundane thermodynamic systems. A third objection calls into question the applicability of certain thermodynamic and statistical mechanical concepts to the universe as a whole. Building upon suggestions due to Hans Reichenbach and (more recently) Huw Price, I show how the specified boundary conditions can indeed constrain thermodynamic behavior in the requisite way; and I argue that when the relationship between the universal boundary conditions and thermodynamic phenomena is properly understood, the third objection can be sidestepped altogether.

Table of Contents

Abstract iii

Acknowledgements vi

Chapter 1. Statistical mechanics and the temporal asymmetry problem 1

Prologue: The mystery of the teapot 1

Section 1. The Puzzle of temporal asymmetry 3

Section 2. Statistical mechanics and the reversibility problem 6

Section 3. Boltzmann's cosmological conjecture & the Past Hypothesis 14

Section 4. Objections to the cosmological explanation 20

Section 5. Conclusion 27

Chapter 2. The Principle of the parallelism of entropy increase 29

Section 1. Entropic asymmetry and the Past Hypothesis 29

Section 2. The Principle of Parallelism 33

Section 3. The Reichenbach-Davies argument 38

Section 4. A proof of POP for the urn model 43

Section 5. A Boltzmannian explanation of the Principle of Parallelism 54

Chapter 3. Replies to Earman and Callender 61

Section 1. The Definability worry and the time scales objection 61

Section 2. Gravity and cosmic entropy 63

Section 3. The time scale problem 68

Section 4. How to sidestep the definability worry 79

Section 5. Conclusion 84

Epilogue: Sherlock's reply.....87

Appendix. Boltzmann's combinatorial argument.....88

References.....91

Acknowledgements

I am indebted to Adam Elga for serving as my primary advisor, for offering very helpful guidance as this dissertation took shape, and for teaching the graduate seminar that first provoked my interest in the topic. I am grateful also to Hans Halvorson for serving as my secondary advisor, and for providing much valuable guidance not only concerning my research but concerning important professional and personal decisions along the way. Many thanks to Noel Swanson, Cory Nichols, Rachel Parsons, and two anonymous reviewers, for providing extensive written comments on earlier drafts of a paper that has been incorporated into this dissertation. Thanks also to the other members of the Department of Philosophy, both faculty and graduate students, who have inspired and contributed to my work in more ways than I can name.

Chapter 1. Statistical mechanics and the temporal asymmetry problem

Prologue: The mystery of the teapot

Finding a comfortable chair in a quiet corner of the parlor, Sherlock settles in for tea. As he tips the steaming teapot to pour a nice hot cup, he notices something peculiar: floating in the pot are several chunks of ice! The ice must have been dropped into the teapot very recently, he muses; for otherwise, it would have melted and cooled the tea. The detective glances about to see who might have been responsible for the misdeed, and is confronted by his nemesis, Professor Sklariarty.[1]

"I suppose you think someone has tried to ruin your tea."

"Well," the detective replies, "since its contents are not in thermal equilibrium, I suspect someone has tampered with the teapot."

The professor smiles condescendingly. "My dear fellow, I suppose you are unaware of the fact that the dynamical laws of fundamental physics are temporally symmetric.[2] If it is possible for ice to melt in hot tea, then the temporal reverse of that process must also be possible—it must be possible, that is, for lukewarm tea to evolve into ice and hot tea, quite independently of any tampering."

"I am aware that it is physically possible for a system to randomly fluctuate out of equilibrium," Sherlock responds with a conciliatory nod. "Nonetheless, random fluctuations are more likely to increase the entropy of a system than to lower it. Since

[1] Sklariarty, unlike the villainous Professor Moriarty of the famed novels by Sir Arthur Conan Doyle, is really more of a witty prankster than a reprobate criminal.

[2] In some recently discovered special cases, the time-reversal of a process requires that matter be replaced with anti-matter, but these rare exceptions are not plausibly responsible for the ubiquitous temporal asymmetry of thermodynamic processes. See Wallace (2013, p. 4 and pp. 10-12) for further discussion.

equilibrium is the condition of maximum entropy, a system in equilibrium will (in all likelihood) stay in equilibrium until it is disturbed by an interaction with its environment. The entropy of the teapot is now very low, so its present condition is almost certainly not the result of a mere fluctuation. It is far more probable that the ice cubes were dropped into hot tea than that they spontaneously froze out of lukewarm tea." Picking up a teaspoon, he calmly scoops out the ice.

"That is very clever reasoning, my good man." Sklariarty's brow rimples into a pensive frown. "I'll grant that the teapot was in a low-entropy condition because of an interaction with its environment. But why do you assume that the interaction responsible for that low-entropy condition must have occurred *prior* to the condition in question? Given the temporal symmetry of the laws, that assumption is unwarranted. Just now, you returned the teapot to equilibrium by removing the ice; but viewed in reverse, that is like disturbing the equilibrium condition by putting ice into the teapot. Perhaps the teapot initially contained lukewarm tea, which gradually evolved into hot tea with ice, and finally you removed the ice. Viewed in reverse, that whole process is just the same as a scenario in which ice is added to hot tea and then melts away. The laws do not distinguish these two possibilities; one is just the time-reversal of the other. So I don't see how you can dismiss the former possibility as unlikely. If you think the teapot would be likely to approach equilibrium after a disturbance in which ice is added, then by parity of reasoning you ought to conclude that it evolves toward equilibrium in the reverse time direction from a disturbance in which ice is removed."

Sherlock sets down his teacup and shifts in his chair, which is beginning to feel a lot less cozy than it had at first. “Perhaps we can continue this conversation after I’ve had a chance to finish my tea.”

1. The Puzzle of temporal asymmetry

Drop an ice cube into a steaming cup of coffee and it will quickly melt, leaving you with a cup of lukewarm, diluted coffee. Let a cup of lukewarm diluted coffee sit around as long as you like; it will never spontaneously separate into an ice cube floating in hot coffee. This temporal asymmetry is explained by the Second Law of Thermodynamics, which says that the entropy of an energetically isolated system does not decrease. (A precise definition of entropy is provided in the appendix; for now, we can think of entropy simply as a theoretical quantity that increases as a system evolves toward a condition of equilibrium.) The entropy of the coffee-cup system increases as the ice melts, then levels off to a constant value as the contents of your cup reach thermal equilibrium. The temporal reverse of that process would violate the Second Law, so we have a thermodynamic explanation for the temporally asymmetric behavior of ice in coffee. However, the laws of thermodynamics do not explain why changes in entropy are temporally asymmetric in this way. A deeper explanation of this temporal asymmetry, which I’ll call the *entropic asymmetry*, is desirable.

The coffee-cup system is comprised of molecules, which are comprised of atoms, and so on. If an asymmetry could be found in the dynamical laws governing elementary particles, that asymmetry might provide a way of explaining the entropic asymmetry described by the Second Law. Interestingly, however, the fundamental microphysical

laws appear to be almost perfectly time-symmetric.[3] In other words, whatever happens can also happen in reverse, so far as the microphysical laws are concerned. If an ice cube melts in hot coffee, then it is possible—according to the most fundamental laws of physics—for an ice cube to spontaneously freeze out of lukewarm coffee!

Asymmetries in the microphysical laws may yet be discovered, of course. It is sometimes suggested that an asymmetry of quantum collapse (e.g., in a GRW-style interpretation of quantum mechanics) might yield the right sort of temporal asymmetries at the macroscopic level.[4] On the other hand, if the fundamental microphysical laws really are time-symmetric, as they appear to be, then the entropic asymmetry must come from temporal asymmetries in the boundary conditions of thermodynamic systems.[5] But where could the latter asymmetries come from?

It is widely believed that the entropic asymmetry ultimately derives from an asymmetry in the boundary conditions of the universe as a whole.[6] Cosmologists tell us

---

[3] See footnote 2, above.

[4] Penrose (1989) and Albert (1994 and 2000) have made suggestions along these lines. See Wallace (2013, pp. 11-12) for further discussion.

[5] See North (2011) for a helpful survey of the various approaches to solving the asymmetry puzzle. As she notes, in order to explain the asymmetry of thermodynamics, "we will need a temporal asymmetry somewhere in our fundamental theory. … The question is where. Answers can generally be divided into one of two camps: an asymmetry in boundary conditions or an asymmetry in the dynamics." (p. 17) Price (2004) refers to the latter as the "two-asymmetry" approach, since it still relies on the assumption of a low-entropy boundary condition for the universe. However, the one-asymmetry approach is not necessarily simpler or explanatorily preferable to the two-asymmetry approach, which may have other explanatory advantages. For instance, Albert (2000) suggests that the quantum collapse approach (see footnote 4, above) might eliminate the need for an independent statistical posit: statistical mechanical probabilities might instead be grounded in quantum mechanical chances (p. 152).

[6] Depending on one's view of physical laws, this idea may seem wrong-headed. Winsberg (2009) argues that statistical mechanics makes better sense if one adopts a "framework conception" of laws, à la Cartwright and Giere. On this view, "laws provide a framework for building models, schematizing experiments, and representing
Phenomena," but do not have universal domains of application and are not universally true (p. 716). Such a view may provide a way of resisting the idea that the temporal asymmetry of thermodynamics must be grounded either in fundamental laws or in universal boundary

that the distribution of mass and energy shortly after the Big Bang was extremely smooth. The distribution of mass is becoming much less smooth as matter clumps together under the force of gravity, however, and eventually much of the universe's mass will be concentrated in black holes.[7] So the distribution of mass and energy across the universe was very different at the beginning than it will be at the end. It has frequently been suggested that this *global* asymmetry of the mass-energy distribution might somehow be responsible for the temporally asymmetric behaviors of *local* (ordinary-sized) thermodynamic systems.[8] Some have been so optimistic as to suggest that the global asymmetry might explain not only the entropic asymmetry but nearly all of the temporal asymmetries of our everyday experience: it could explain why causes precede their effects, why we remember the past but not the future, why we can influence the future but not the past, why things break but never spontaneously reassemble, and so on.[9]

Unfortunately, it is far from obvious how an asymmetry in the global mass-energy distribution could explain the thermodynamic behavior of a cup of coffee, let alone the myriad other temporal asymmetries of our mundane experiences. Numerous forceful objections—due to Eric Winsberg, John Earman, Craig Callender, and others—have recently been pressed against the idea that the global asymmetry places significant constraints on local thermodynamic systems. The aim of this dissertation is to answer

---

conditions. Although this issue is worth addressing, I don't want my arguments to hinge on any particular view of laws, so I will sidestep the issue for now.

[7] It is hypothesized that black holes might eventually evaporate through Hawking radiation, so that all of the matter in the universe will ultimately be converted into electromagnetic radiation.

[8] The idea that the entropic asymmetry might have a cosmological explanation was first suggested by Boltzmann, whose cosmological hypothesis will be discussed in section 3, below. The idea that local temporal asymmetries might be grounded in asymmetric global boundary conditions has been suggested in various ways by Reichenbach (1956), Feynman (1965), Penrose (1989; 2005), Lebowitz (1993a; 1993b), Bricmont (1995), Price (1996), Albert (2000), Goldstein (2001), Callender (2004a; 2004b), Wald (2006), and numerous others.

[9] Albert is among the most optimistic. See chapters 5 and 6 of Albert (2000).

three of those objections. I will argue that the asymmetry of thermodynamics can indeed be explained in terms of the global asymmetry, or at least that such an explanation remains quite plausible despite the objections raised by Winsberg, Earman, and Callender. The other asymmetries (of causation, memory, etc.) will be left aside, though I share the optimism that a wide variety of temporal asymmetries might be explained in terms of the universe's boundary conditions.

The following section of this first chapter provides a brief introduction to Boltzmannian statistical mechanics, which provides a simple but useful framework for investigating the puzzle of entropic asymmetry. Section 3 discusses several attempts to provide a cosmological solution to the puzzle, beginning with one of Boltzmann's own ideas. The modern cosmological account will also be presented, in two slightly different forms. In section 4, I consider several objections against the very idea of explaining the local entropic asymmetry in terms of the global asymmetry. These objections will then be taken up in chapters 2 and 3 of this dissertation.

2. Statistical mechanics and the reversibility problem

As noted above, the fundamental microphysical laws appear to be *time-symmetric* in the sense that anything that's allowed to happen forwards is also allowed to happen backwards. Since it is possible for an isolated system of ice and coffee to evolve into lukewarm diluted coffee, for example, it must likewise be possible for lukewarm diluted coffee to evolve into hot coffee with ice. Thus, the microphysical laws imply that the Second Law of Thermodynamics can be violated: the entropy of an isolated system can

decrease.[10] Such violations of the Second Law are rarely (if ever) observed, of course: ice *just doesn't* spontaneously freeze out of lukewarm coffee. Why not? Statistical mechanics is supposed to provide an answer. That is, statistical mechanics is supposed to provide a way of explaining the thermodynamic behavior of macroscopic systems in terms of the laws governing their microscopic components.

One might reasonably doubt that this ambitious explanatory feat can be accomplished by assuming that the universe consists of point-sized particles that obey the laws of classical mechanics—an outdated theory that gives the wrong predictions on both the quantum and the cosmic scales. Nevertheless, most of the recent philosophical literature on the temporal asymmetry problem has been framed in terms of classical Hamiltonian[11] mechanics, and this paper will be no exception. Despite the fact that real particles are not point-sized and do not obey the classical laws, imagining that the world is a classical system helps to simplify the discussion considerably and may provide insights that can later be generalized or extended to the actual physics.[12] Henceforward, let us pretend that the world is made up of point-sized particles whose motions obey the laws of classical Hamiltonian mechanics.

---

[10] This point is often acknowledged by formulating the Second Law as a statistical or probabilistic principle: the entropy of an isolated system is very *unlikely* to decrease.

[11] Hamilton Mechanics is simply a reformulation of the classical Newtonian laws. The details of that reformulation won't matter for the present discussion.

[12] There is more to be said in favor of this simplifying assumption. Sometimes the best explanations deliberately misconstrue the causal story behind some system's behavior, in order to emphasize which aspects of that behavior are relevant to the explanandum. If the quantum and relativistic aspects of the world make no difference to the macroscopic behavior of ice in coffee, for example, then it is appropriate to explain this macroscopic behavior in classical terms. (For a brief but insightful discussion of this point, see Strevens (unpublished); for a more detailed exposition of his view, see Strevens (2008)). It may well be that non-classical features of reality *do* make a difference, but it seems worthwhile at least to explore the possibility of explaining thermodynamic phenomena classically. Even if classical statistical mechanics ultimately fails, perhaps the way in which it fails will illuminate which aspects of the actual physics are responsible for thermodynamic phenomena.

Imagine a classical system comprised of n particles of unit mass in a three-dimensional physical space. The precise state of such a system at any instant can be represented by a single point in an abstract $6n$-dimensional *phase space*, with three coordinates representing the position and three representing the momentum of each particle. The Hamiltonian laws determine a set of possible trajectories through this phase space, representing the precise ways in which the system (or other systems comprised of equally many particles) can evolve over time. The laws are deterministic, in the sense that each point in phase space lies on exactly one Hamiltonian trajectory. (Figure 1.1)

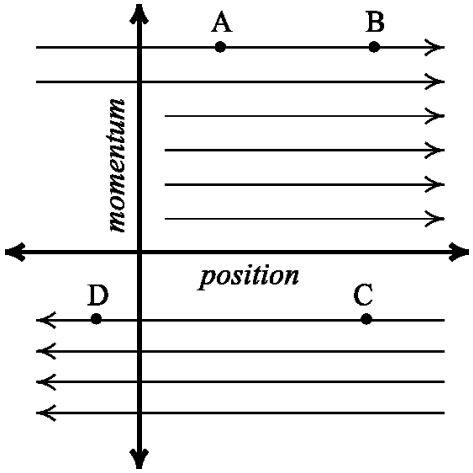


Figure 1.1
Consider a single particle confined to just one dimension of physical space. The phase space of such a system has only two dimensions: one to represent the particle's position and one to represent its momentum. Provided there are no forces acting on it, the particle will move with constant momentum, so the Hamiltonian trajectories will be straight lines. Each point on this diagram represents a possible microstate, and the thin horizontal arrows represent Hamiltonian trajectories. If the system's actual microstate is presently at point A, for instance, the particle will move quickly to the right and the system will eventually pass through microstate B. If the system's actual microstate is presently at point C, on the other hand, the particle will move slowly to the left and the system will eventually pass through microstate D. For systems consisting of multiple interacting particles, the trajectories will not be straight lines; ordinarily the Hamiltonian trajectories will meander haphazardly due to the changing forces exerted between particles.

Certain macroscopic properties of a system can be represented in phase space as well. In particular, the thermodynamic condition (defined by macroscopic variables like temperature, pressure, and volume) of a system at any given time corresponds to a contiguous region of phase space rather than a single point. The exact state of a system at a time—in full and precise microscopic detail—is often called its *microstate*, while the

term *macrocondition* refers to its thermodynamic condition.[13] With a slight abuse of terminology, I will use these words to denote also the corresponding point or region in the system's phase space.

In the late 19th century, Ludwig Boltzmann found a simple way to define entropy in terms of phase space rather than in terms of thermodynamic variables. His notion of entropy, called *Boltzmann entropy*, is not exactly the same as thermodynamic entropy; but it is closely analogous to thermodynamic entropy in that it increases as a system evolves toward equilibrium. For most of what follows, the differences between Boltzmann entropy and thermodynamic entropy will not matter. When they do matter, I will make the necessary clarifications. To avoid confusion at the outset, however, it is worth mentioning that there are other types of entropy besides thermodynamic entropy and Boltzmann entropy.[14] In fact, Boltzmann entropy is not the type of entropy defined in most statistical physics textbooks. Boltzmann framed his statistical considerations in a way that only strictly applies to ideal gasses. His contemporary J. Willard Gibbs developed a more versatile framework, which is usually employed by physicists today.

For the purposes of the present discussion, however, there are at least two good reasons for preferring the simpler Boltzmannian framework. First, Boltzmann entropy is determined by the state of the physical system independently of any probabilistic assumptions, whereas Gibbs entropy is defined in terms of a statistical posit. The statistical posit itself will come under scrutiny in what follows, so we'll need a concept of

[13] At least, that is how I will use the term. Sometimes the term 'macrocondition' or 'macrostate' is used, somewhat misleadingly, to refer to the level of coarse-graining associated with Boltzmann entropy. See footnote 16, below.

[14] Other types of entropy include *Shannon entropy* (a concept used in information theory), *Kolmogorov-Sinai entropy* (used in measure theory), and *von Neumann entropy* (used in quantum statistical mechanics), to name just a few. For a discussion of various types of entropy, see Frigg (2011a).

entropy that does not rely on such a posit. Secondly, Boltzmann entropy is more closely analogous to thermodynamic entropy: it increases as a system evolves toward equilibrium. In contrast, Gibbs entropy is calculated only for equilibrium conditions. For systems that are evolving from one equilibrium condition to another (due to the removal of an external constraint, for example), Gibbs entropy is regarded as a constant during the system's evolution. For this reason, the Boltzmannian framework lends itself to a more perspicuous explanation of the asymmetry of thermodynamics.[15]

Boltzmann's crucial insight was that high-entropy macroconditions take up *vastly* greater volumes of phase space than low-entropy ones.[16] (This is the upshot of his famous combinatorial argument, which is explained in the appendix to this dissertation.) Indeed, for systems whose phase space is bounded—i.e., systems of finite energy confined to a finite region of physical space—the macrocondition representing thermodynamic equilibrium fills nearly all of the available phase space.[17] Given fairly uncontroversial assumptions about the shape of the Hamiltonian trajectories,[18] this suggests that any low-

[15] See Goldstein (2001) for further discussion of this point.

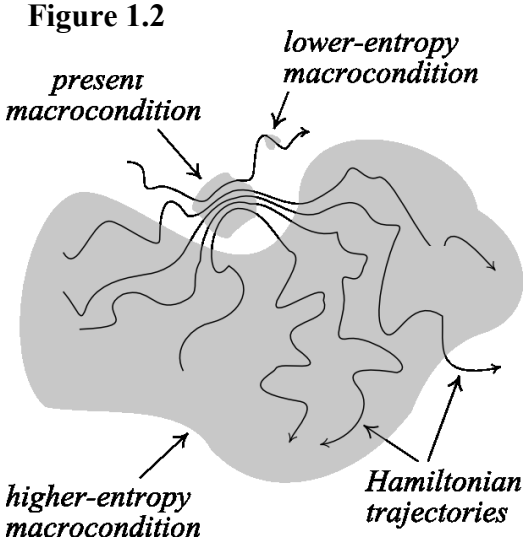
[16] By 'volume' here, I mean the Lebesgue measure. For the meticulous reader, a few further points of clarification are in order. The Boltzmann entropy of a system is associated with a much more fine-grained 'macrostate' than the thermodynamic macrocondition. (See appendix.) Boltzmann entropy is proportional to the log of the volume of the Boltzmannian macrostate, so the volume increases exponentially as the entropy increases. Something similar is presumed to be true for thermodynamic macroconditions and their corresponding thermodynamic entropies, so the differences between thermodynamic and Boltzmann entropy make little difference for the present discussion. Both kinds of entropy increase monotonically as the system evolves toward equilibrium, and both kinds of macrostates take up vastly greater volumes of phase space as entropy increases.

[17] This claim is true only for the thermodynamic macrocondition, not for the so-called 'macrostate' that is associated with Boltzmann entropy (see appendix). The thermodynamic macrocondition includes very many macroscopically indistinguishable Boltzmannian macrostates, since the latter notion is comparatively fine-grained.

[18] Though it is often assumed that the trajectories are ergodic, such a strong assumption is not required for Boltzmann's argument; his reasoning will hold for almost any plausible hypothesis about the shape of the trajectories. (A trajectory is said to be *ergodic* if, in the limit as the trajectory is extended in both temporal directions from any point, the fraction of time it spends in

entropy macrocondition is a local entropy minimum for nearly all of the trajectories passing through it. To see why, imagine a set of trajectories going through one of the tiny regions of phase space that correspond to a low-entropy macrocondition. As the trajectories wander out of that macrocondition in various directions, most of them are sure to run into higher-entropy macroconditions, since higher-entropy macroconditions fill such a large proportion of the surrounding phase space. In contrast, lower-entropy macroconditions are nearly impossible targets to hit, because they are so small. (Figure 1.2) Thus, the overwhelming majority of trajectories through any low-entropy macrocondition will evolve more or less monotonically toward equilibrium (maximum entropy) in both temporal directions.

Figure 1.2



Since most of the possible low-entropy microstates (i.e., the microstates in which a system could be, given its low-entropy macrocondition) lie on trajectories that lead to greater entropy, it seems almost inevitable that systems in low-entropy states will evolve toward equilibrium. Thus, it is tempting to think that Boltzmann's insight explains why low-entropy systems that we encounter in the real world always—or at any rate, almost always—obey the Second Law of Thermodynamics. Things are not quite that straightforward, however. In the first place, the foregoing considerations seem to suggest that most systems in low-entropy states should evolve toward equilibrium in both

---

any particular region is equal to the proportion of the accessible phase space occupied by that region.)

temporal directions. That's deeply problematic, as we'll see in a moment. Secondly, the fact that higher-entropy macroconditions take up vastly greater volumes of phase space implies nothing at all about the behavior of actual physical systems. Phase space is just a space of physically *possible* microstates, and provides no further information about which microstates actually occur in real physical systems. For all that phase space tells us, the exceptional microstates that lead to anti-thermodynamic behavior might be instantiated more frequently. Of course that would be surprising, given what we know about the world, but the crucial point here is that phase space by itself has no statistical or probabilistic implications: it implies nothing at all about what most real systems will do, or about what any particular system is objectively likely to do. Though Boltzmann's insight seems to go a long way toward explaining thermodynamic behavior, something further is needed.

To explain the approach to equilibrium in terms of Boltzmann's insight, some sort of probability distribution must be introduced into the phase space.[19] It is far from obvious how that is to be done. Suppose we assume that a system is equally likely to be in any of the myriad microstates that are compatible with its present macroscopic properties. That is, we posit a probability density function that is uniform over the present

[19] Lebowitz (1993a), (1993b); Goldstein (2001) and others have suggested that a probability distribution is not required to explain why systems tend toward equilibrium. Instead, they argue, systems can be understood as evolving away from *atypical* microstates and toward *typical* microstates: equilibrium microstates are typical, insofar as the overwhelming majority of possible microstates are in equilibrium; whereas non-equilibrium microstates are atypical in that they occupy only a tiny proportion of phase space. Thus, the approach to equilibrium can be explained without appealing to notions of probability. See Frigg (2009), (2011b) for discussion and criticism of the 'typicality' account.

macrocondition of the system, and zero elsewhere.[20] I'll call this the *Naïve Statistical Postulate*:

> **Naïve Statistical Postulate:** The probability that an isolated system's *present* microstate lies within any given subregion of its *present* macrocondition is proportional to the volume (Lebesgue measure) of that subregion.

For any isolated system that is not already in equilibrium, this probability distribution will imply that the system is overwhelmingly likely to evolve toward equilibrium in the future, as we'd expect. However, since the present macrocondition is a local entropy minimum for most of the trajectories passing through it, the Naïve Statistical Postulate also makes it very likely that the system has evolved *from* a higher-entropy condition in the *past*. In other words, the Naïve Statistical Postulate implies that the system has almost certainly been violating the Second Law of Thermodynamics up to the present moment (but from now on it will obey the Second Law).

To illustrate the problem, suppose that several minutes ago you dropped a few ice cubes into a mug of hot coffee. By the present moment, most of the ice has melted. Now, a uniform probability density over the present macrocondition of the mug will imply, wrongly, that the remaining bits of ice in your coffee most likely got there by spontaneously freezing out of lukewarm coffee, rather than coming from larger ice cubes that have been melting away into hot coffee. Although the Naïve Statistical Postulate correctly predicts the system's future entropic behavior, it gets the past dramatically

[20] A probability density function is uniform when any two regions of equal volume (Lebesgue measure) have the same probability.

wrong. This problem, which was first pointed out by Boltzmann's colleague and former teacher Johann Loschmidt in 1876, has been dubbed the *reversibility problem*.[21]

3. Boltzmann's cosmological conjecture and the Past Hypothesis

In an ingenious attempt to solve the reversibility problem, Boltzmann conjectured that the universe might be vast enough or old enough to make large-scale fluctuations from equilibrium inevitable.[22] Although fluctuations away from equilibrium are extremely unlikely to occur in any particular place at any particular time, there is a good chance that low-entropy conditions will randomly occur sometimes in at least some places if the universe is large enough or lasts long enough. To see why, imagine a trajectory meandering aimlessly around in phase space. Although it spends the vast majority of its time at or near equilibrium, sooner or later it's sure to wander briefly through some low-entropy macrocondition or other. (Figure 1.3) In fact, one of Boltzmann's contemporaries, Henri Poincaré, proved that Hamiltonian trajectories through phase space will eventually bring almost any isolated system arbitrarily close to its initial state, even if that initial state is one of very low entropy.[23]

---

[21] The reversibility problem was originally posed as an objection to Boltzmann's famous H-theorem, which showed that given an assumption similar to the Naïve Statistical Postulate, a certain quantity 'H' (which is inversely proportional to Boltzmann entropy) would not decrease after the moment specified in the assumption. A closely related objection to the H-theorem, called the recurrence objection, was raised by Ernst Zermelo in 1896. (See footnote 23, below.) Although neither of these objections was explicitly aimed at the Naïve Statistical Postulate, both show that the Naïve Statistical Postulate is inadequate for explaining temporally asymmetric phenomena in terms of the temporally symmetric Hamiltonian laws.

[22] See Sklar (1993), pp. 43-44.

[23] The Poincaré recurrence theorem was published in 1890. Zermelo's recurrence objection, mentioned in a previous footnote (fn 21), made use of the Poincaré recurrence theorem to show that the quantity 'H' of Boltzmann's famous H-theorem could increase, and therefore that entropy could decrease. See Sklar (1993), pp. 36-37.

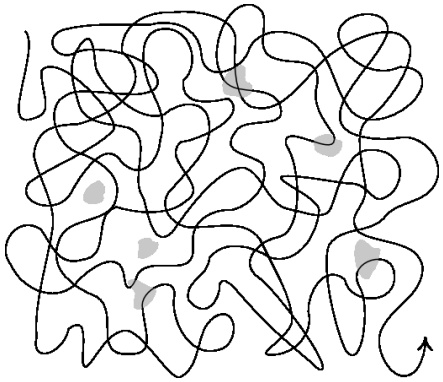


Figure 1.3

The tangled line represents a single Hamiltonian trajectory, wandering aimlessly through phase space. (Actual Hamiltonian trajectories do not intersect, but it is difficult to draw a tangled mess without intersecting lines.) The unshaded region represents the equilibrium macrocondition, while the small grey regions represent low-entropy macroconditions.

Perhaps the universe is almost always in a state of equilibrium, Boltzmann suggested, but on rare occasions some parts of the universe randomly fluctuate out of equilibrium and back again. It so happens that at this particular moment, in this particular region of the universe, we are experiencing one of those extraordinarily rare fluctuations. Since complex life can only exist where there is a steep entropy gradient, the fact that we find ourselves in such a low-entropy region of time and space might be explained by the observer selection effect known in more recent literature as the 'anthropic principle.'[24] The fact that we are in such a region could, in turn, explain why the second law of thermodynamics holds for the systems that we observe, *assuming* the entropy gradients of local isolated systems tend to parallel the entropy gradient of the region as a whole. This last assumption is an important one, and we will return to it in a moment.

Boltzmann's radical hypothesis did not meet with widespread acceptance in his own day, and later discoveries in astronomy and cosmology have shown that the universe is neither large enough nor old enough to make large-scale fluctuations from equilibrium

[24] According to the anthropic principle, we should expect to find ourselves in a region of space and time that can sustain complex life (given that we exist), even if most of spacetime is inhospitable to life.

probable. Boltzmann's proposal also suffers from rather severe philosophical defects.[25] Nevertheless, the idea that cosmology might provide an explanation for entropic asymmetry has not been abandoned. A new cosmological explanation has recently gained widespread acceptance among philosophers and physicists alike.[26] This popular view appeals to a remarkable discovery about the early universe, namely, the discovery that the distribution of matter shortly after the Big Bang was extremely smooth. Since gravity is an attractive force, it is suggested, this smooth distribution of matter constitutes (or is at least analogous to) a very low-entropy condition.[27] The idea that the universe began in an extremely low-entropy macrocondition (or something closely analogous to such a condition) has been called the *Past Hypothesis*.[28]

Given the Past Hypothesis, it is unsurprising that the entropy of the universe as a whole has been steadily increasing. To see why, recall Boltzmann's insight that nearly all of the possible microstates within any low-entropy macrocondition lie on trajectories that evolve toward higher entropy in both temporal directions (as previously shown in figure 1.2). Applying this insight to the macrocondition of the early universe, we can infer that

---

[25] The most serious philosophical objection to Boltzmann's cosmological hypothesis is the so-called 'Boltzmann's brain paradox.' The problem is that small fluctuations from equilibrium are likely to occur much more frequently than large ones. Hence, if the universe were vast enough or old enough to make large fluctuations probable, an overwhelmingly higher number of small fluctuations would also occur—so many, in fact, that conscious brain-like arrangements of matter would spontaneously occur more often than low-entropy regions of the size required for complex life to evolve. So, if Boltzmann's hypothesis were correct, then we should conclude that our memories and present conscious experiences of the world are merely illusory, and that we are merely brain-like arrangements of matter that have spontaneously formed in an otherwise high-entropy region of the universe! Interestingly, the same problem arises for some recent cosmological models that posit multiple universes to explain the early condition of our universe. For further discussion of the Boltzmann's brain paradox, and a novel cosmological hypothesis that escapes it, see Carroll and Chen (2004).

[26] See footnote 8, above, for references.

[27] The reason will be briefly explained in section 4 of this chapter; more details will be given in chapter 3.

[28] This sobriquet originated with Albert (2000), but is now widely used.

almost any of the low-entropy microstates in which the universe could have begun must be microstates that evolve toward higher entropy. This explains why the universe as a whole has a steep entropy gradient sloping upward in the forward time direction. Thus, we have an explanation for the entropic asymmetry of the universe as a whole.[29]

But does the Past Hypothesis also provide an explanation for the entropic asymmetry of local thermodynamic systems? It is far from obvious how the distribution of mass in the early universe could explain the thermodynamic behavior of a cup of coffee billions of years later. One might be tempted to reason as follows: the Past Hypothesis explains the entropic asymmetry of the universe as a whole, and cups of coffee are just subsystems of the universe, so an explanation of the global asymmetry is tantamount to an explanation of all the local asymmetries. Unfortunately, the situation isn't quite so simple. From the fact that the entropy gradient of the universe as a whole slopes in one direction, it does not follow that the entropy gradients of most local systems will slope in the same direction. In fact, it could happen that nearly all local isolated subsystems decrease in entropy while the entropy of the universe as a whole nevertheless increases.[30]

It may turn out that the entropy gradients of local thermodynamic systems are probabilistically correlated with that of the surrounding environment, but this remains to be shown. In his cosmological explanation of thermodynamics, Boltzmann simply assumed that the entropy gradients of local thermodynamic systems would tend to

[29] There is some dispute concerning the explanatory status of the Past Hypothesis itself. Price (2002; 2004) contends that the low-entropy condition of the early universe is itself in need of explanation, while Callender (2004a; 2004b) argues (in agreement with Albert) that the Past Hypothesis should be treated as a fundamental non-dynamical law that is not in need of explanation.

[30] This will be illustrated with a specific example in the next chapter.

parallel the entropy gradient of the surrounding environment. This assumption, which Hans Reichenbach later dubbed the "principle of the parallelism of entropy increase," has been the subject of some controversy, as we'll see in the next chapter. For now, let's simply take note that without a careful justification for that assumption, we cannot explain the local entropic asymmetry in terms of the global asymmetry.

It is sometimes suggested, however, that the Past Hypothesis provides a more direct solution to the reversibility problem. For instance, in *Time and Chance*, Albert suggests that the problem can be solved simply by conditioning on the Past Hypothesis. That is, we replace the Naïve Statistical Postulate with a uniform probability distribution over just those regions of phase space that are compatible with the Past Hypothesis. I'll call this revised postulate the Conditional Statistical Postulate:

> **Conditional Statistical Postulate:** The probability that an isolated system's microstate lies within any given subregion of its macrocondition is proportional to the volume (Lebesgue measure) of that subregion, excluding any points that are incompatible with the Past Hypothesis.

Supposedly, the Conditional Statistical Postulate will assign zero probability to microstates that lie on trajectories whose entropy was higher in the past, since those trajectories are inconsistent (it is claimed) with the Past Hypothesis. Whether that claim is true remains to be shown, and in fact there are good reasons to doubt the claim, as Eric Winsberg has argued. We'll consider his argument in the next section.

The modern solution to the reversibility problem is often presented in a slightly different way, a way that conceals the need for an explanatory link between the entropy gradient of the universe as a whole and the entropy gradients of local thermodynamic systems. Although in *Time and Chance* Albert endorses something like the Conditional Statistical Postulate, in more recent papers he and numerous other authors have suggested that all statistical mechanical probabilities are to be derived ultimately from a uniform probability distribution over the initial macrocondition of the universe.[31] I'll call this the Global Statistical Postulate:

> **Global Statistical Postulate:** The probability that the *universe's* initial microstate lies within any given subregion of its initial macrocondition is proportional to the volume (Lebesgue measure) of that subregion.

To understand the behavior of a local system like a cup of coffee, on this view, we must regard the system as merely a subsystem of the universe, and explain its behavior in terms of the probabilities given by the Global Statistical Postulate. (We can—in

[31] Loewer (2007) endorses such a view, and provides a Humean/Lewisian argument for regarding the Global SP as a fundamental law of physics. Wallace (2013) argues that the Global SP is "conceptually preferable, as it makes clear that we are imposing two independent conditions on the initial state of the universe: one on its macrostate [viz., the Past Hypothesis], one on its microstate (that it be given by a reasonably-smooth probability distribution over microstates compatible with some macrostate or other)." (p. 272) Albert (2009) endorses the Global SP as well. In fact, this view tacitly appears even in *Time and Chance*, or at least, one of Albert's arguments there seems to commit him to such a view. On p. 89, Albert criticizes the "branch systems" account of Reichenbach and Davies on the grounds that logical inconsistencies may arise from applying statistical posits to branch systems that contain smaller branch systems. (For further discussion of the objection, see Winsberg (2004b); Frigg, R. (2008).) However, the same objection could be raised against the Conditional SP. When one isolated system contains another, applying the Conditional SP to both systems can yield logically inconsistent probability distributions. Hence, Albert's objection against Reichenbach and Davies seems to commit him to the Global SP rather than the Conditional SP.

principle—still represent the local system in its own phase space, provided we use the Global Statistical Postulate to generate a probability distribution in the local system's phase space. An "in principle" method for doing so will be described below.)

It is claimed that the Global Statistical Postulate makes it very likely that local isolated systems will obey the Second Law of Thermodynamics. No rigorous arguments have been offered for this claim, however, and it is not easy to see why it should be true—unless we take for granted the assumption that local entropy gradients are likely to parallel that of the universe. Hence, the modern cosmological explanation of entropic asymmetry, like Boltzmann's original hypothesis, appears to rest on a hitherto unjustified assumption. That assumption, again, is what Reichenbach called "the principle of the parallelism of entropy increase": the idea that entropy gradients of isolated subsystems should be likely (in terms of statistical mechanical probabilities) to parallel the entropy gradient of the main system. In chapter 2, I will show that the assumption holds for a simple model, and argue that it likewise holds for classical Hamiltonian systems and probably for real physical systems as well. Before defending the cosmological explanation of entropic asymmetry, however, it will be worthwhile to consider a few objections against the very idea of explaining thermodynamics in terms of the Past Hypothesis.

4. Objections to the cosmological explanation

Several formidable objections have recently been raised against the idea that the thermodynamic behavior of teapots and the like might be explained by appealing to the condition of the early universe. The first of these objections was mentioned already:

Winsberg argues that conditioning on the Past Hypothesis does not solve the reversibility problem. His argument will be considered presently; I will then turn to some further objections raised by John Earman and Craig Callender (and in slightly different ways by Hans Reichenbach and Adolf Grunbaum).

Winsberg's argument is directed specifically against Albert's putative solution to the reversibility problem as presented in *Time and Chance*, though I will argue that the objection applies to both the Conditional Statistical Postulate and the Global Statistical Postulate. Recall the problem with the Naïve Statistical Postulate, discussed in section 2, above: although it predicts future behavior correctly, it wrongly predicts that the entropy of a system was higher in the past. Albert contends that the problem can be avoided simply by conditioning on the Past Hypothesis. That is, we assign zero probability to any possible microstate that is incompatible (in a sense to be explained below) with the Past Hypothesis, and distribute probability uniformly over the remaining microstates in the system's macrocondition. When the uniform probability density given by the Naïve Statistical Postulate is conditioned on the Past Hypothesis in this way, supposedly, most of the possible microstates within a system's present macrocondition will be ruled out—i.e., they will be assigned zero probability. (This is the claim that Winsberg contests, as we'll see in a moment.) All of the probability will be distributed over a tiny minority of microstates, namely, those microstates that evolve toward lower entropy conditions in the past, since those are the microstates that are compatible with the Past Hypothesis.

To illustrate, suppose you've added ice to a cup of hot coffee a minute ago. If we apply the Conditional Statistical Postulate to the present macrocondition of the cup, most of the points (possible microstates) within that macrocondition will receive zero

probability, allegedly, since most of those points lie on trajectories whose entropy was higher in the past, and that supposedly is inconsistent with the Past Hypothesis. Thus, we avoid the unwelcome consequence of the Naïve Statistical Postulate, which implied that the ice had been less melted a minute ago. Of the possible microstates that have not been ruled out, on the other hand, most lie on trajectories whose entropy increases toward the future; so we get correct predictions about both past and future thermodynamic behavior. Thus, the reversibility problem is solved… or so we're told.

Winsberg contends, however, that the Past Hypothesis is compatible with all or nearly all points in any local system's phase space. The fact that the universe began in a certain macrocondition, he argues, does not rule out the possibility of the local system coming into existence in almost any microstate whatsoever: "any microstate that is compatible with the present state of the [local system] is also one that is in principle compatible with the past hypothesis."[32] Thus, conditioning on the Past Hypothesis has little or no effect on the probability distribution given by the Naïve Statistical Postulate, and the reversibility problem remains unsolved.[33]

To assess Winsberg's argument, we must consider in what sense the Past Hypothesis might "rule out" or be incompatible with a point in a local system's phase space. While the Past Hypothesis certainly does not rule out the *epistemic* possibility of a local system coming into existence in any microstate whatsoever, in the present context *nomological* possibility is what matters. Let's return to the example of a cup containing hot coffee and ice. Each point in your coffee cup's phase space represents a (nomologically and epistemically) possible microstate, i.e., a possible way for the

[32] Winsberg (2004a), p. 503
[33] Winsberg's objection is also discussed by Frigg (2008), p. 128 and by Earman (2006), fn. 38. Earman adds further considerations to bolster the objection, as we'll see.

individual fundamental particles to be arranged at the most precise microscopic level. Assuming for simplicity that the world is governed by classical Hamiltonian dynamics, any such microstate will be nomologically compatible with the Past Hypothesis only if the Hamiltonian laws permit such a microstate to occur while the Past Hypothesis is true. In other words, a point in your coffee cup's phase space is compatible with the Past Hypothesis only if there is a Hamiltonian trajectory in the *universal* phase space that begins in the condition specified by the Past Hypothesis and that later passes through a microstate which includes (as a subsystem) a coffee cup in the precise microstate represented—in your coffee cup's phase space—by the point in question.

Now, I have no idea how to determine whether much of your coffee cup's phase space is compatible (in the sense just described) with the Past Hypothesis. I suspect that neither Winsberg nor Albert does either. For this reason, I am not convinced by Winsberg's claim that the Past Hypothesis rules out few if any possible microstates for local systems[34]; but by the same token, Albert has given us no compelling reason to believe that the Past Hypothesis rules out precisely those trajectories on which your cup's entropy was higher in the past. It's anyone's guess whether conditioning on the Past Hypothesis will have much effect on the uniform probability distribution given by the Naïve Statistical Postulate. The Past Hypothesis might be compatible with nearly every possible configuration of the molecules in your coffee mug, or it might be compatible

[34] I've been assuming that the PH specifies some *particular* macrocondition (viz., "whatever particular low-entropy highly condensed big-bang sort of macrocondition it is that the normal inferential procedures of cosmology will eventually present to us", as Albert says on p. 96). Winsberg's suggestion may be more plausible if we construe the PH instead as saying merely that the universe began in *some* low-entropy macrostate or other. Even on that interpretation, however, it is far from obvious that the PH will be compatible with most of the local system's possible initial microstates as Winsberg claims.

with very few. In any case, we have been given little reason to suppose that conditioning on the Past Hypothesis yields correct predictions concerning the behavior of ice in coffee.

Even if we assume, *pace* Winsberg, that the Past Hypothesis rules out quite a lot of the local system's phase space, it is still very far from obvious whether the remaining trajectories will be precisely the ones which obey the second law of thermodynamics. The Past Hypothesis certainly cannot rule out anti-thermodynamic behavior forever. As Boltzmann recognized, spontaneous fluctuations away from equilibrium are bound to occur at least some times, in at least some places, if the universe lasts long enough. In fact, there is a sense in which the Past Hypothesis may be compatible with almost anything in the long run. If trajectories in the universal phase space are ergodic,[35] as they are often assumed to be, any given trajectory will eventually fluctuate arbitrarily close to any possible microstate. For instance, nearly every trajectory will eventually bring the universe to a state in which one of its subsystems closely resembles your coffee cup. For most trajectories, this will happen long after the universe has reached a condition of equilibrium, and will involve a relatively small fluctuation. The coffee-cup-like subsystem will be a lonely aberration in a world that is otherwise in equilibrium. Such lonely fluctuations will not exhibit the entropic asymmetry: the aberrant subsystems will have higher entropy at both temporal ends of the fluctuation. So, if all of this is correct, the Past Hypothesis is compatible—at least in the long run—with a wide range of possible anti-thermodynamic behaviors.

The crucial question, of course, is whether the Past Hypothesis is similarly permissive in the short run, that is, in the relatively brief period (billions of years) before

[35] A trajectory is said to be *ergodic* if, in the limit as the trajectory is extended in both temporal directions from any point, the fraction of time it spends in any particular region is equal to the proportion of the accessible phase space occupied by that region.

the universe reaches equilibrium. Given that the Past Hypothesis does not rule out anti-thermodynamic behavior for all time, why assume that it rules out anti-thermodynamic behavior for the systems we observe today? Without a satisfactory answer to that question, Albert's solution to the reversibility problem seems at best incomplete. We need additional or independent reasons for thinking that the Past Hypothesis places the right sorts of constraints on local thermodynamic systems.

Though Winsberg's argument was aimed specifically at the Conditional Statistical Postulate, the foregoing considerations likewise call into question the success of the Global Statistical Postulate. Since it is far from obvious what sorts of constraints are placed on local systems by the condition of the early universe, positing a uniform probability density function over the macrocondition of the early universe won't solve the reversibility problem—at least, not without some further explanatory link between the entropy gradient of the universe and the entropy gradients of local thermodynamic systems. In chapter 2, I will argue that Reichenbach's "principle of the parallelism of entropy increase" can provide the requisite explanatory link. Before turning to that, however, some further objections are worth considering.

Above, I introduced the Past Hypothesis as the claim that the universe began in an extremely low-entropy macrocondition, or something closely *analogous* to such a condition: namely, a very smooth distribution of matter. When repulsive forces are dominant, for example in the free expansion of a gas, matter tends to disperse and become smoothly distributed; but on large scales where the attractive force of gravity rules the day, matter tends to clump together. So in mundane thermodynamic systems, smooth distributions of matter correspond to high-entropy states, whereas smooth

distributions of matter are analogous to low-entropy states in a gravitating system. Unfortunately, the gravitational contribution to entropy is not fully understood, and the Boltzmannian framework is not well equipped to make sense of gravitational entropy—if there is any such thing. In fact, as we'll see in chapter 3, neither thermodynamic entropy nor Boltzmann entropy is well-defined for the universe as a whole. Since the entropy of the universe as a whole is not even a well-defined notion, the prospects of explaining the entropic asymmetry in terms of the Past Hypothesis may seem to be doomed at the start. This objection has been expressed most recently (and most poignantly) by John Earman and Craig Callender, though similar worries were raised previously by Reichenbach and Grunbaum.[36]

Earman also adds a further objection. Even if we can define a notion of entropy that accommodates gravity, he points out, the gravitational contribution to entropy "is irrelevant for the kinds of subsystems of interest to us." Gravitational attractions between the molecules in your coffee cup, for instance, are negligible. Moreover, he continues:

> Nor is it plausible that the entropy increases we perceive in typical thermodynamic systems is the result of an increase in gravitational entropy that is transferred to the non-gravitational degrees of freedom; for on the time scales of typical thermodynamic observations and inferences, the values of the variables that characterize the gravitational degrees of freedom are effectively constant.[37]

[36] See Reichenbach (1956), p. 114; Grunbaum (1963); Earman (2006); Callender (2010).
[37] Earman (2006), p. 419.

If I understand Earman's point correctly, his claim is that the entropy gradients of small subsystems cannot be expected to parallel that of the universe, because the universal entropy gradient is primarily a matter of gravitational entropy, which is effectively constant over the course of mundane thermodynamic processes. I will answer this objection in chapter 3.

5. Conclusion

Let's take stock of the arguments discussed above. Assuming that the fundamental microphysical laws really are time-symmetric, as they appear to be, the entropic asymmetry of thermodynamic systems must be explained by an asymmetry in the boundary conditions of those systems, and ultimately by some *global* asymmetry—i.e., an asymmetry in the boundary conditions of the universe as a whole. A plausible candidate for such a global asymmetry is the fact that the distribution of mass and energy was very different at the beginning of the universe than it will be at the end. Since the smooth distribution of mass in the early universe is thought to constitute an extremely low-entropy condition (or is at least analogous to such a condition), this global asymmetry is widely believed to provide an explanation for the entropic asymmetry of thermodynamics.

Three formidable objections have been raised against the cosmological explanation of entropic asymmetry. One of these is the objection that cosmic entropy is not a well-defined notion. Near the end of the third chapter of this dissertation, I will suggest a way to reformulate the cosmological explanation so that it need not invoke the notion of cosmic entropy, thereby sidestepping the definability issue. Before doing that,

however, it will be helpful to see how the cosmological explanation might work if the concept of entropy were applicable on a cosmic scale. To that end, I will address the other two objections raised above, namely, the objections concerning the explanatory connection between the global asymmetry and the entropic asymmetry of local thermodynamic systems. One of these 'subsystems objections,' as they have been called,[38] is Winsberg's argument that the Past Hypothesis places little if any constraint on the behaviors of mundane systems. The second is Earman's worry that increases in the gravitational entropy of the universe need not be correlated with increases in thermodynamic entropy, since the relevant gravitational variables are effectively fixed over the time scales of typical thermodynamic observations.

Both of the latter objections will be answered by arguing that the entropic behavior of an isolated subsystem is likely to parallel that of the main system. As mentioned above, Reichenbach called this claim "the principle of the parallelism of entropy increase." I'll call it the Principle of Parallelism (POP) for short. A careful defense of POP is offered in the next chapter, where I show that it holds for a simple model and draw an analogy between that model and the actual case. Chapter 2 also provides a response to Winsberg's version of the subsystems objection. The objections raised by Earman and Callender are then tacked in chapter 3. Finally, a brief epilogue brings the fictitious conversation between Detective Sherlock and Professor Sklariarty to a more satisfying conclusion.

[38] I borrow this name from Callender (2011b), (2011c), who regards these problems as "among the most pressing facing the Globalist solution." (2011b), p. 100.

Chapter 2: The Principle of the parallelism of entropy increase

1. Entropic asymmetry and the Past Hypothesis

Many commonplace thermodynamic processes never happen in reverse: ice melts in hot tea but never spontaneously freezes out of lukewarm tea, air leaks out of a pressurized bicycle tire but never leaks back in, and so on. The Second Law of Thermodynamics provides a high-level explanation—i.e., an explanation in terms of non-fundamental laws—for these temporal asymmetries. A thermodynamic process is irreversible if it involves a net entropy increase, in other words, if the total entropy of the system and its environment is raised in the process. The Second Law forbids the total entropy to decrease,[1] and that is why such processes do not occur backwards.

However, the laws of thermodynamics describe macroscopic regularities, and are themselves amenable to explanation in terms of more fundamental laws. Surprisingly, the most fundamental laws of physics appear to be almost perfectly time-symmetric.[2] This suggests that the temporal asymmetry of thermodynamics, which I'll call the *entropic asymmetry*, derives from an asymmetry not in the laws but in the boundary conditions of thermodynamic systems, and ultimately from an asymmetry in the boundary conditions

[1] The Second Law can be formulated in different ways; sometimes it is expressed in a probabilistic form to accommodate the statistical mechanical insight that it is possible, albeit extremely unlikely, for the entropy of an isolated system to decrease. In the previous chapter, the Second Law was formulated as the claim that the entropy of an energetically isolated system cannot decrease. Since thermodynamic entropy is an *extensive* quantity (i.e. the total is the sum of the parts), this implies that a non-isolated subsystem can only decrease in entropy when it interacts with its environment in a way that raises the entropy of the environment by at least as much, so that the total change in entropy is non-negative.

[2] In some recently discovered special cases, the time-reversal of a process requires that matter be replaced with anti-matter (see footnote 2 from the previous chapter). It is also worth mentioning that some regard the Past Hypothesis and/or the Global Statistical Postulate as fundamental non-dynamical laws rather than boundary conditions (see e.g. Loewer for a defense of this view), and of course both of those are explicitly intended to introduce a temporal asymmetry into fundamental physics. Nonetheless, the fundamental *dynamical* laws appear to be time-symmetric.

of the universe as a whole. A widely accepted view holds that the entropic asymmetry stems ultimately from a condition of extremely low entropy in the early universe. The idea that the universe began in such a condition has been called the *Past Hypothesis*.

Several worries have been raised for the view that entropic asymmetry can be explained in terms of the Past Hypothesis. For one, it is unclear whether the Past Hypothesis even makes sense. The ordinary thermodynamic and statistical mechanical definitions of entropy cannot be straightforwardly applied to gravitating systems like stars and galaxies, much less to the universe as a whole.[3] Thus, the Past Hypothesis is technically ill-defined, unless a broader concept of entropy can be devised and extended to cases where gravity plays more than a negligible role. Moreover, even if there is such a quantity as the entropy of the entire universe, further worries have been raised concerning the putative explanatory link between the Past Hypothesis and the behaviors of relatively small, isolated or quasi-isolated[4] thermodynamic systems. For instance, the entropy of the early universe appears to have been low primarily because of gravity; but as John Earman points out, the gravitational contribution to entropy seems irrelevant to the behaviors of ordinary thermodynamic systems. This worry will be revisited in the next chapter, but let us set it aside for the moment. The present chapter is concerned with yet another worry about the explanatory link between the Past Hypothesis and the temporal asymmetry of thermodynamics. The worry is simply this: it is unclear how the Past Hypothesis could have significant implications for very small, short-lived subsystems of the universe. Even if it is true that the universe began in a condition of extremely low entropy, how could

[3] As mentioned in chapter 1, Boltzmann entropy is the appropriate statistical mechanical notion of entropy in this context, though in fact *none* of the ordinary definitions of entropy can be unproblematically applied at the cosmic scale. See Frigg (2011a) for a discussion of the various types of entropy.

[4] I'll explain what is meant by 'quasi-isolation' in section 2, below.

that fact have significant implications for the behavior of ice in hot tea or of air in bicycle tires?

In the preceding chapter, we examined Eric Winsberg's objection against the idea that the Past Hypothesis has significant implications for thermodynamics. (John Earman also presses this objection, and adds a further consideration to bolster it. His version of the objection will be examined in chapter 3.) To briefly recap, Winsberg contends that conditioning on the Past Hypothesis makes no difference at all to the statistical mechanical probabilities associated with local thermodynamic systems. This claim was perhaps too strong, I suggested, but at the very least Winsberg has made it clear that the explanatory link between the Past Hypothesis and local thermodynamic phenomena is far from trivial. Simple conditionalization won't do the trick, because it is not obvious what sorts of constraints, if any, the Past Hypothesis places on local isolated systems.

Nor will it suffice to posit a uniform probably density over the low-entropy macrocondition of the early universe, I argued, because again it is far from obvious whether this *Global Statistical Postulate* (as I called it) has the right sorts of implications for local systems. At best, the Global Statistical Postulate provides an explanation for the steep entropy gradient of the universe as a whole; but we still need an explanatory link between the universal entropy gradient and the entropy gradients of local thermodynamic systems. Craig Callender poses the issue this way: "What follows for isolated subsystems from the fact that the global entropy is likely to increase? To me, this question is among the most pressing facing the Globalist solution. … It's a mistake to rely on the intuition

that if the system's entropy is increasing then 'most' of the subsystem's entropies are increasing too."[5]

From the fact that the universe's entropy is increasing, it does not follow that entropy is increasing in most of the subsystems as well. (We'll see why in section 2, below.) On the other hand, it *might* be the case that the entropy gradient of a subsystem is *likely*—in terms of the probabilities given by the Global Statistical Postulate—to parallel the universal entropy gradient. That is what some proponents of the so-called 'globalist solution' have suggested, and what Winsberg denies.[6] As I argued in the preceding chapter, however, Winsberg's objection seems inconclusive. It is entirely unclear what sorts of constraints, if any, the Past Hypothesis places on ordinary thermodynamic systems. It's no good to rely simply on intuitions one way or the other. In this chapter, I will defend the globalist solution with a careful argument for the claim that the direction of entropy change for an isolated subsystem is probabilistically correlated with that of the main system. I call this crucial claim the Principle of Parallelism (POP).

The following section gives a brief overview of the contemporary cosmological explanation for entropic asymmetry, highlighting the way in which POP could provide the requisite link between the Past Hypothesis and local thermodynamic behavior. Section 3 provides some historical context for my defense of POP. Hans Reichenbach and Paul Davies both gave arguments for that principle, but Lawrence Sklar later pointed out a subtle way in which their arguments appear to beg the question. To avoid making the same mistake, it is important to see how and why the Reichenbach-Davies argument

[5] Callender (2011b), p. 100.
[6] Perhaps I should say that Winsberg *implicitly* denies this, as his arguments are directed explicitly against the Conditional Statistical Postulate rather than the Global Statistical Postulate. As I pointed out in the first chapter, however, his argument can likewise be directed against the Global Statistical Postulate.

is unsuccessful. In the fourth section, I show that POP holds for a toy model similar to one famously used by the Ehrenfests to illustrate the process of thermodynamic equilibration. Then, in section 5, the argument is recast directly in terms of Boltzmannian statistical mechanics.

2. The Principle of Parallelism

Given the Past Hypothesis, which says that the universe had extremely low entropy at its past temporal boundary, it is unsurprising that the universe has a steep entropy gradient sloping upward in the opposite temporal direction, i.e., toward the future. (For further explanation, refer again to sections 2 and 3 of the preceding chapter.) Plausibly, this *global* entropy gradient makes it extremely improbable that small bits of the universe—teapots and such—have entropy gradients sloping the other way.[7] This could explain why ice melts in hot tea but never spontaneously freezes out of lukewarm tea, and more generally, why thermodynamic entropy is often observed to increase but rarely (if ever) observed to decrease. Unfortunately, however, there is a significant lacuna in this cosmological explanation for the temporal asymmetry of thermodynamics.[8] From the fact that the entropy of the universe as a whole is rapidly increasing, it does not obviously

[7] This probabilistic implication is only supposed to hold for subsystems that are isolated or at least "quasi-isolated," as I will explain below.

[8] Actually, there are a number of lacunae in this explanatory project, some of which will be addressed in this chapter and the next. Here are a few additional problems that will not be addressed in this dissertation. First, the low-entropy initial condition may itself be in need of explanation, so the explanation of thermodynamic asymmetry may be incomplete. Price has made this point numerous times; see for instance (1995), (2002), and (2004). (For an opposing view, see Callender 2004a.) Three closely related further issues concern the status of the initial conditions: 1) should they be regarded as law-like or as merely contingent facts? 2) the justification for assuming a uniform probability density in phase space, and 3) the question of how to interpret the probabilities. See Albert (2000), pp. 85-96; Winsberg (2004b), (2009); Myrvold (unpublished) for further discussion of these issues.

follow that the entropy of a relatively small and more-or-less isolated subsystem, like a teapot, is unlikely to decrease.

It is tempting to think that the entropy gradients of most subsystems must parallel that of the universe, since thermodynamic entropy is an *extensive* quantity.[9] That is, the total thermodynamic entropy of a system is just the sum of the entropies of its partitioning subsystems; so if the entropy of the universe increases, it follows trivially that the average change in entropy for any partitioning set of subsystems must be positive. For instance, if we divide the entire universe up into teapot-sized parts, the average change in entropy among those parts must go in the same direction as the total change in the entropy of the universe as a whole. (This is true regardless of whether the parts are isolated from each other.) That trivial observation, however, does not entail that the entropy of most subsystems must change in the same direction as that of the main system. To see why, imagine a system comprised of several subsystems, most of which undergo a very slight entropy decrease while the remaining few undergo a very dramatic entropy increase. The total change in entropy may be positive, even though most of the subsystems decrease in entropy. (A specific example is given in figure 2.1.) Thus, it could happen that entropy decreases almost everywhere, while the total entropy of the universe nevertheless increases.[10]

[9] Callender (2011b) points out that Boltzmann entropy is not extensive (p. 100). Nevertheless, thermodynamic entropy is extensive, so one might still be tempted by this idea.
[10] Sklar (1993) makes essentially this point (pp. 330-331).

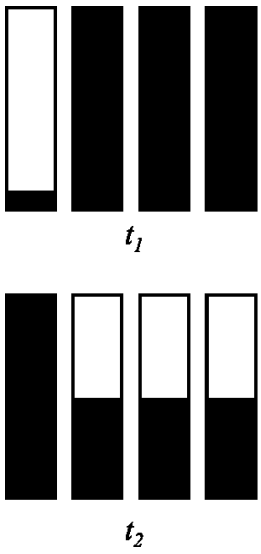


Figure 2.1

Consider a system comprised of four boxes, each containing one mole of ideal gas. Suppose the gas in the first subsystem is initially confined to the bottom $1/16^{th}$ of the box by a partition that is suddenly removed at time t_1. The gas undergoes Joule expansion, reaching a new equilibrium condition at t_2. Suppose furthermore that the remaining three subsystems undergo the temporal reverse of Joule expansion: they are filled with gas at t_1, but somehow all of the molecules in each of those boxes just happen to wind up in the bottom half at t_2. The change in entropy for each box is given by $R\ln(V_f/V_i)$, where R is the ideal gas constant and V_i and V_f are the initial and final volumes, respectively, to which the gas is confined. The total entropy of the four-box system *increases* by $R\ln(16) - 3R\ln(2) = R\ln(2)$, even though three of the four subsystems *decrease* in entropy.

To bridge this explanatory gap, let us turn to a controversial assumption concerning what Reichenbach called *branch systems*—systems that are temporarily isolated from their environments, or at least *quasi-isolated* in the sense that interactions with the environment make only a negligible impact on the system's entropy. (In what follows I will ignore the distinction between perfect isolation and quasi-isolation, as it makes no difference to the arguments that I will be presenting.) Reichenbach referred to the assumption in question as "the principle of the parallelism of entropy increase."[11] In deference to historical precedent, I've adopted his phrase as the title of this chapter; but to make a more pronounceable acronym I will call it simply the *Principle of Parallelism* (POP):

[11] Reichenbach (1956) expressed the assumption thus: "In the vast majority of branch systems, the directions toward higher entropy are parallel to one another and to that of the main system." (pp. 136-7)

(POP) The entropy of a branch system is very unlikely, in terms of statistical mechanical probability, to change in the opposite direction from that of the main system (provided the main system is likewise isolated).[12]

Branch systems are precisely the systems to which the Second Law of Thermodynamics applies, since the Second Law says that the entropy of a system does not decrease *while the system is isolated*.[13] Hence, an argument for POP could provide (in conjunction with the Past Hypothesis) an explanation for the asymmetry of thermodynamics.

Reichenbach showed that POP could be derived from statistical mechanical assumptions that appeared, at first, to be unproblematic. (A simplified version of the derivation, due to Davies, will be explained below.) However, Sklar later pointed out a subtle way in which Reichenbach's argument seems to beg the question, or in any case, fails to provide a non-question-begging explanation for the fact that branch systems have entropy gradients parallel to that of the universe.[14] Sklar finds similar shortcomings in Davies' version of the argument, and also dispenses with a few other possible ways in which one might attempt to justify the parallelism assumption. He concludes that "the existing arguments attempting to derive this branch system asymmetry from the temporal asymmetry of cosmic entropy increase are all fatally flawed."[15]

---

[12] This parenthetical condition is needed because the entropy of an isolated subsystem is not affected when the total entropy of the main system is changed by an interaction with the main system's environment. I will return to this point at the end of section 3.

[13] The Second Law can be formulated in different ways, but this simple formulation will suffice for the present discussion.

[14] Sklar (1992), pp. 144-5; (1993), pp. 319-31. In fairness, Sklar mentions that Reichenbach's argument may not have been intended to explain entropic parallelism (though Reichenbach clearly does think that his argument is explanatory in some respect). I will not attempt to weigh in on this exegetical subtlety.

[15] Sklar (1993), p. 418.

The failure of historical arguments for POP doesn't bode well for the cosmological account of entropic asymmetry, especially given that these failures were brought to light roughly two decades ago and no new arguments have been forthcoming. Nevertheless, it is a widely held opinion that the thermodynamic asymmetry should be explainable *somehow* in terms of the cosmic entropy gradient, even though no one has successfully spelled out the details of such an explanation.[16] Here is a representative statement by Huw Price, who also provides a helpful suggestion for how one might proceed to construct the desired explanation:

> The crucial point is that in terms of the Boltzmann probabilities, the 'cheapest' (most probable) way to reach a specified low-entropy boundary is to get there monotonically (or almost monotonically—the expected fluctuations from monotonicity are very small compared to the overall gradient). The most likely universe with a single low-entropy boundary will be one in which all entropy gradients in quasi-isolated subsystems slope down towards that boundary. … Of course, all of this needs to be backed up by some numerical arguments. A good place to start would be with elaborations of the Ehrenfest Urn models.[17]

In the fourth section of this paper, I follow Price's suggestion and show how POP can be derived for an urn model similar to the Ehrenfest models. Then, in section 5, the

[16] Sklar (2009) remarks: "One can argue, in a manner tracing back to H. Reichenbach, that the entropic increase of the universe as a whole will lead, again using the usual time symmetric probabilistic posits, to a high probability that a random 'branch system' will show entropic increase parallel to that of the universe and parallel to that of other branch systems. Most of the arguments in the literature that this will be so are flawed, but the inference is reasonable nonetheless."

[17] Price (2002), p. 113.

argument is recast directly in terms of Boltzmannian statistical mechanics. Before presenting my arguments, though, it will be worthwhile to briefly rehearse the unsuccessful Reichenbachian argument and Sklar's objection to it. It is not obvious just what is required for a successful explanation of entropic parallelism, and Sklar's critique of the Reichenbach-Davies argument will help to illuminate what is at issue here.

3. The Reichenbach-Davies argument for the Principle of Parallelism

To understand the Reichenbach-Davies argument and what it was intended to achieve, let us begin by situating the argument in its historical context. The Principle of Parallelism (POP) assumption originated almost as early as statistical mechanics itself, decades before cosmology provided much information about the history of the universe. In an ingenious attempt to explain the asymmetry of thermodynamics, Ludwig Boltzmann conjectured that the world might be vast enough or old enough to make large-scale fluctuations from equilibrium inevitable. Although random fluctuations from equilibrium are extremely unlikely to occur in any particular place at any particular time, if a system is large enough or lasts long enough there may be a good chance that low-entropy conditions occur sometimes in at least some regions. Since complex life can only exist where there is a steep entropy gradient, Boltzmann suggested, the fact that we find ourselves in a low-entropy region might be explained by an observer selection effect known in more recent literature as the 'anthropic principle.'[18] The fact that we are in such a region could, in turn, explain why the second law of thermodynamics holds for the

[18] According to the so-called 'anthropic principle,' we should expect (given that we exist) to find ourselves in a region of space and time that can sustain complex intelligent life, even if most of spacetime is inhospitable to life.

systems that we observe, assuming the Principle of Parallelism—i.e., assuming that the entropy gradients of branch systems tend to parallel that of the region as a whole.

In his posthumously published book *The Direction of Time*, Hans Reichenbach made a sophisticated effort to work out the details of Boltzmann's proposal. The contemporary cosmological account that I will be defending is in large part a continuation of that explanatory project. The hypothesis that the universe randomly fluctuates out of equilibrium and back again is no longer tenable, of course; the low-entropy condition of the early universe is more plausibly explained by something like the inflation model.[19] Nonetheless, the modern account shares with Boltzmann's the idea that the asymmetries of thermodynamics (and perhaps other temporal asymmetries as well) ultimately stem from a cosmic entropy gradient, via entropic parallelism.

One of the most important attractions of the cosmological account is that it promises to explain thermodynamic behavior without appealing to any temporal asymmetries apart from the fact that the universe has a low-entropy boundary condition at one temporal end and not the other. In order for the account to succeed in providing such an explanation, however, the crucial assumption on which it relies—namely, POP—must be justified in a way that does not presuppose any further temporal asymmetries. In a fascinating discussion of Boltzmann's cosmological hypothesis, Reichenbach provided a statistical mechanical derivation of POP that seemed to provide the right sort of justification. A simpler reformulation of his argument, due to Paul Davies, goes roughly as follows.

[19] According to the inflation model, the extremely smooth distribution of matter in the early universe (see chapter 1) resulted from a period of exponential expansion that occurred a split second after the Big Bang. Whether this model succeeds in providing the right sort of explanation has been disputed. See Price (1995), pp. 71-72.

Suppose the exact microscopic state, or *microstate*, of a system at time t is determined "at random," in the sense that the position and velocity of each molecule is about as likely to be in any small range as in any neighboring range of similar size.[20] Then, according to Boltzmann's combinatorial argument,[21] the system's entropy is unlikely to decrease if the system remains isolated as it evolves in either temporal direction from t. In other words, the system is likely to approach equilibrium (maximum entropy) as it evolves either forward or backward in time from a randomly-determined microstate, for as long as the isolation continues.

Now, here's the trick: Davies assumes that when a branch system is *formed*—that is, when a thermodynamic system becomes isolated—its *initial* microstate is determined at random in the sense just described. It follows that the entropy of the system is very unlikely to decrease as the system evolves in the forward temporal direction from that initial time t_i, as long as the system remains isolated. (The same would be true of the reverse temporal direction from t_i, if the system had already been isolated beforehand. But the system was not isolated prior to t_i, by hypothesis, since t_i is the moment when the branch system was formed.) Assuming that all branch systems are formed in this way, it follows trivially that their entropy gradients are likely to parallel one another. Thus,

[20] For readers familiar with the technical jargon of statistical mechanics, the idea here is that the probability density is approximately uniform (as judged by the Lebesgue measure on phase space) over the system's macrostate at t.

[21] This claim is the upshot of Boltzmann's famous combinatorial argument (which I explain in the appendix) together with plausible assumptions about the dynamics of the system. As far as the dynamics are concerned, neither ergodicity nor mixing are required: the entropy of system will be very unlikely to decrease on almost *any* plausible assumption about the dynamics. See Goldstein (2001) for an argument against the relevance of ergodicity and mixing.

according to Davies: "It is the asymmetry regarding the formation of branch systems which brings about the parallel increase in all (nearly) branch system entropies."[22]

Suspiciously, the foregoing argument never invokes the direction of entropy increase for the main system. Sklar points out that the Reichenbach-Davies argument relies instead on a temporally asymmetric assumption about the boundary conditions for branch systems: the *initial* microstate of a branch system is assumed to be determined at random, but no similar assumption is made about the *final* microstate, i.e., the microstate of the branch system at the moment when its isolation ends. [23] Such temporally asymmetric assumptions are reasonable if all we're after is predictive accuracy, Sklar notes; but in order to provide a cosmological explanation for thermodynamic asymmetry we need something more:

> Systems do (in the probabilistic sense) have their entropies increase in a time direction parallel to each other and to the main system. That isn't the point. The point is that we wanted an *explanation* of that fact. We thought we were going to get one from the time direction of entropy increase of the main system. We didn't. We simply got a rule that generates a prediction of parallel behavior by positing it and avoids predicting anti-parallel behavior by prohibiting the other way of applying the posit of random micro-state. But parity of reasoning considerations would lead us to think that the posit is legitimate applied both ways. It isn't, of course, but we still don't know why.[24]

---

[22] Davies (1977), p. 72.
[23] See Sklar (1992), pp. 144 -5 for a brief discussion of Reichenbach's argument and the problem with it; see (1993), pp. 318-31 for more detailed discussion.
[24] Sklar (1993), p. 327.

If our explanatory goal is ultimately to give an account of thermodynamic asymmetry in terms of the cosmic entropy gradient, then our statistical mechanical derivation of entropic parallelism had better not rely on a further temporal asymmetry that is just as mysterious as the one we set out to explain. In order for the asymmetry of cosmic entropy increase to explain the entropic asymmetry of branch systems, it must be shown that the Principle of Parallelism (POP) can be derived from genuinely time-symmetric assumptions.

The prospects for a genuinely time-symmetric justification of POP may not seem especially good. The arguments of Winsberg and Earman, discussed in the preceding chapter, give reasons to doubt that the entropy of the universe as a whole is relevant to the behaviors of branch systems.[25] Here is a further worry that, to my awareness, has not been mentioned elsewhere. It is perfectly commonplace for the entropy of a system to decrease when the system exchanges matter or energy with its environment. Turning on the air conditioner, for instance, lowers the entropy of a room through the operation of a heat pump. Yet the behavior of isolated subsystems in the room does not change when the total entropy decreases in this way; ice continues to melt in hot tea, and so on. The Principle of Parallelism does not imply otherwise, since it applies only to isolated systems. POP does imply, however, that if the total entropy of an *isolated* room were to decrease (as the result of a random fluctuation), then any ice in that room would be obliged *not* to melt in hot tea. It is difficult to see why that should be the case. How could it make a difference to the behavior of an isolated subsystem whether the total entropy of

[25] Winsberg (2004a); also Earman (2006), p. 419. These arguments are also discussed by Frigg (2008), p. 128; Frisch (2010); North (2011); Callender (2011b), p. 100.

the room decreases due to an interaction or a random fluctuation? This question will be revisited at the end of the next section.

Despite all of these worries, I will argue that POP can be derived from well-motivated assumptions that do not presuppose any temporal asymmetry. The following section shows how this can be done for a simple model of branch systems. A closely analogous argument, formulated within the framework of Boltzmannian statistical mechanics, is then developed in section 5.

4. A proof of POP for the urn model

As noted above, Price advises us to think about the Principle of Parallelism in terms of Ehrenfest Urn models. The aim of this section is to construct an urn model for which POP holds independently of any temporally asymmetric assumptions. This will require a bit of caution because, as I shall explain presently, the dynamics of Ehrenfest Urn models are not time-symmetric.

On a simple urn model, a number of balls are initially distributed arbitrarily between two urns, *A* and *B*. A ball is selected at random, in such a way that all the balls have an equal probability of being chosen, and the selected ball is moved to the opposite urn. The random selection process is repeated many times, always moving the chosen ball to the other urn. This dynamical algorithm is not time-symmetric, as can been seen from an example. Imagine that there are ten balls, of different colors, and suppose one of them—the red one, say—is in urn *A* after three trials. Its probability of moving to urn *B* on the *next* (fourth) trial is 1/10. However, the probability that the red ball was in urn *B* just *before* the third trial depends on which urn it was in at the start. If the red ball began

in urn *A*, for instance, then its probability of being in urn *B* just before the third trial (given that it is in urn *A* after the third trial) is lower than 1/10. If the red ball began in urn *B*, on the other hand, then its probability of being in urn *B* just before the third trial (given that it is in urn *A* after the third trial) is higher than 1/10.[26] In neither case is the red ball's probability of moving from urn *A* to urn *B* on the fourth trial the same as its probability of having moved from urn *B* to urn *A* on the third trial. In that sense, the dynamics of the urn model are temporally asymmetric.

However, there is a sense in which the dynamics are approximately time-symmetric *in the long run*. Since the aim of this section is to construct a proof of POP that does not presuppose any temporal asymmetries, I will make use of that long-run symmetry in what follows. It will therefore be worthwhile to carefully explain the sense in which the long-run dynamics of the urn system are temporally symmetric. Let *n* be the total number of balls in the urns. If the random selection process is repeated very many times, always moving the selected ball to the opposite urn, it is extremely probable that both urns will eventually contain about *n*/2 balls, regardless of how the balls were initially distributed. In fact, if the number of trials is many times more than *n*, each ball will be about equally likely to end up in either urn. For instance, if the urns contain ten balls between them, then the probability that a particular ball will be in urn *A* after a million trials is practically the same as the probability that it will end up in *B*, regardless of which urn it was in at the start. Thus, the initial distribution of balls "washes out" in the long run.

Since the initial distribution washes out in the long run, the dynamics of the urn system are approximately time-symmetric after very many trials. To see this, consider the

[26] The exact probabilities are 1/42 and 41/122, respectively.

illustration used above, in which ten balls of different colors are initially distributed in some arbitrary way between urns *A* and *B*. Now, instead of supposing that the red ball is in urn *A* after just three trials, let us suppose it is in urn *A* after a million trials. Given this new assumption, what is the probability that the red ball will be in urn *B* just before the millionth trial? In other words, what is the conditional probability that the red ball will be in urn *B* after the 999,999th trial, *given* that it will be in urn *A* after the millionth trial? By Bayes' theorem:

$$p(\mathrm{B}|\mathrm{A}) = p(\mathrm{A}|\mathrm{B})p(\mathrm{B})/p(\mathrm{A}),$$

where $p(\mathrm{B})$ is the prior probability that the red ball will be in urn *B* just *before* the millionth trial and $p(\mathrm{A})$ is the prior probability that the red ball will be in urn *A* just *after* the millionth trial. Those two probabilities will be almost equivalent, to an extremely close approximation, regardless of how the balls were distributed before the very first trial. It follows that the conditional probabilities must likewise be almost equivalent: $p(\mathrm{B}|\mathrm{A}) \approx p(\mathrm{A}|\mathrm{B})$. In words, the probability that the red ball will be in *B* just *before* the millionth trial, given that it will be in *A afterwards*, is almost identical to the probability that it will be in *A afterwards* if it is in *B* just *beforehand*.

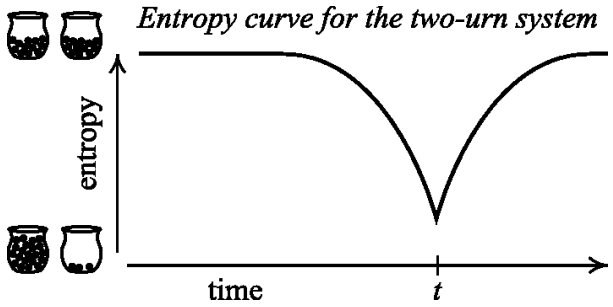
This last conditional probability, $p(\mathrm{A}|\mathrm{B})$, is directly specified by the dynamics. Since each ball has a 1/10 chance of being selected and moved to the opposite urn on any given trial, the probability that the red one will be selected on the millionth trial is one in ten. The same argument will go through, of course, for each of the other balls in the system. What we have just seen is that the long-run dynamics of the system are almost

perfectly time-symmetric: given that a particular ball is in a certain urn at time *t*, its probability of *having been* in the other urn at *t*–1 is practically the same as its probability of *moving to* the other urn at *t*+1 (where the time units are in 'trials' and *t* is very large). This will be important in what follows.

Now that we see how the dynamics of the urn system are temporally symmetric in the long run, let us see how the so-called "reversibility problem" manifests itself in this model.[27] Suppose that a low-entropy condition occurs after a long run of trials. Though low-entropy conditions are unlikely to occur at any particular time, in a sufficiently long sequence of trials the system will occasionally fluctuate to a low-entropy state (where one urn contains significantly more balls than the other) and then back to the "equilibrium" condition in which the urns contain roughly the same number of balls. On the assumption that the system is in a low-entropy state at *t*, where *t* occurs late in a long run of trials, it is very likely that *t* will be preceded by a (mostly) monotonic fluctuation away from equilibrium and followed by a (mostly) monotonic evolution back to equilibrium again. (Figure 2.2)

To see why, consider first the forward temporal direction from *t*. Since the system is, by assumption, in a low-entropy state at *t*, there must be many more balls in one urn than in the other at that time. Subsequent repetitions of the dynamical algorithm are very likely to take the system more-or-less monotonically toward equilibrium, since it is more likely that a ball will

Figure 2.2



[27] See chapter 1, section 2 for an explanation of the reversibility problem.

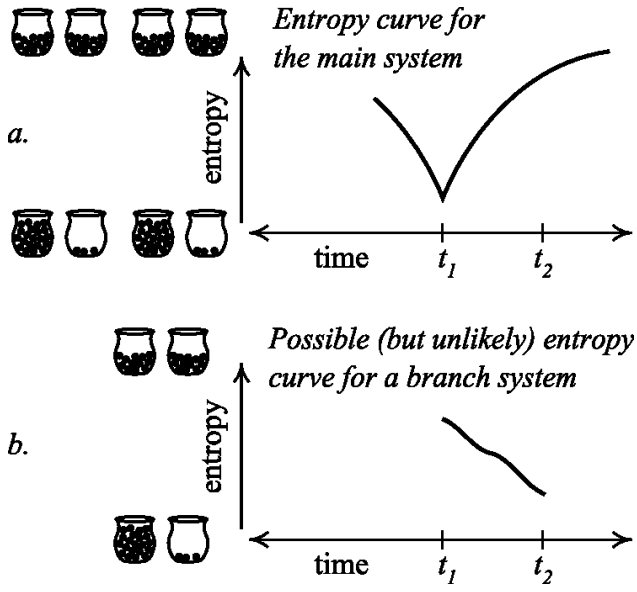
be chosen from the crowded urn than from the sparsely populated urn on any given trial. For the reverse temporal direction, recall that the long-run dynamics of the system are temporally symmetric with respect to any particular specification of which balls are in which urns. If the system is likely to approach equilibrium in one temporal direction from the specified state, then it is likely to approach equilibrium in the opposite temporal direction as well. In other words, since the system is likely to evolve toward equilibrium after t, it is equally likely that the system will be in equilibrium at some earlier time and evolve more-or-less monotonically away from equilibrium as time t approaches.

We are now in a position to derive the Principal of Parallelism without relying on any temporally asymmetric assumptions. Of course, our model will have to be embellished a bit, since the simple two-urn system is not complex enough to represent branch systems (i.e., subsystems that are temporarily isolated from the main system). Before constructing a model that is complicated enough to represent the formation and reabsorption of branch systems, though, let's try a model comprised of subsystems that are permanently isolated from each other. This model will be referred to later, so to give it name, I'll call it "the warm-up exercise":

> *Warm-up exercise*. Suppose the main system is comprised of very many pairs of urns (A_i, B_i). Each pair is governed by the same dynamics as the pair of urns considered above: whenever a ball is chosen from urn A_i it moves to the corresponding urn B_i and vice versa. Let t_1 and t_2 be times that occur long after the initial distribution has washed out, and suppose the total entropy of the main system at t_1 is low. (Figure 2.3a)

Given these assumptions, how likely is it that any particular pair of urns will have significantly lower entropy at t_2 than at t_1? (Figure 2.3b) A moment's reflection will show that it is very *unlikely* that any given subsystem will have lower entropy at t_2, especially if t_2 is a significant number of trials away from t_1 (in either temporal direction). To see this, notice that there is no temporal asymmetry with respect to t_1. The probability of any particular ball being selected and moved to the opposite urn is the same whether we run the experiment forward from t_1, or trace the history of the system backwards from t_1. It makes no difference whether t_2 occurs before or after t_1, therefore; but it may be heuristically useful to imagine that t_2 occurs after t_1. In other words, we can suppose that the main system *began* in a low-entropy state at t_1, and evolved toward another state at t_2.

Figure 2.3



When viewed in this way, the solution to our warm-up exercise is easy to see. Some pairs of urns must have low entropy at t_1, given that the main system has low entropy then. The selection algorithm makes it extremely likely that *those* subsystems—which I'll call the *contributors*, since they contribute to the main system's having low

entropy at t_1—will evolve toward equilibrium in the "future" of time t_1.[28] (Of course, contributors will tend toward equilibrium in the "past" direction from t_1 as well, since the dynamics are time-symmetric with respect to t_1.) If a pair of urns has low entropy at t_2, on the other hand, then it is *not* likely to have come from a state of higher entropy in the "past." Why not? Well, if the pair is not in equilibrium at t_2, then it is very likely to be a contributor, as I will argue below.[29] And, as just noted, contributors tend to have higher entropy at t_2 than at t_1.

So long as the main system does not reach equilibrium at any time between t_1 and t_2, any pair of urns that is not in equilibrium at t_2 is almost certainly a contributor, as the following considerations show. First, note that it is extremely unlikely that a pair of urns will fluctuate to a low-entropy state from a state of higher entropy at t_1; but if the urns were already far from equilibrium at t_1, there is a good chance that their entropy will still be low at t_2 (though perhaps not as low as before). So, the probability that any given subsystem has low entropy at t_2 is very much greater on the assumption that it is a contributor than on the assumption that it is not:

$$p(\mathrm{L}|\mathrm{C}) >> p(\mathrm{L}|\neg\mathrm{C}),$$

where L says that the subsystem is in a low-entropy state at t_2, and C says that it is a contributor (i.e., it has low entropy at t_1). This implies that for some extremely large

[28] It may not be terribly unlikely for the entropy of a contributing subsystem to be *very slightly* lower at t_2, especially if t_1 and t_2 are only a few trials apart or if the subsystem in question contains only a small proportion of the total number of balls. But it is *extremely* unlikely that the entropy at t_2 will be *significantly* lower than at t_1, for the same reason that random fluctuations away from equilibrium are unlikely in general.

[29] A caveat is needed: the subsystem is very likely to be a contributing one *unless* the main system reached equilibrium at some time between t_1 and t_2.

positive constant k, $p(L|C) = kp(L|\neg C)$. It follows by the probability calculus that $p(C|L)$ differs from $p(\neg C|L)$ by a factor of $kp(C)/p(\neg C)$.[30]

Now, it is important recognize that the probability function p is already conditionalized on the stipulation that the main system has low entropy at t_1. Thus, $p(C)$ is the probability that a particular pair of urns will have low entropy at t_1 *given* that the main system randomly fluctuates to a low-entropy state at that time. This probability cannot be terribly small, since at least some pairs of urns must have low entropy when the main system has low entropy. So $kp(C)/p(\neg C)$ is still very large,[31] and hence:

$$p(C|L) >> p(\neg C|L).$$

That is, the subsystem is much more likely to be a contributor than a non-contributor, given that it has low entropy at t_2. Therefore, since (as argued above) contributors are

---

[30] By Bayes' theorem, the left side of the equation is equivalent to $p(C|L)p(L)/p(C)$, and the right side is equivalent to $kp(\neg C|L)p(L)/p(\neg C)$. So we have:
$p(C|L)p(L)/p(C) = kp(\neg C|L)p(L)/p(\neg C)$
$p(C|L)/p(C) = kp(\neg C|L)/p(\neg C)$
$p(C|L) = kp(\neg C|L)p(C)/p(\neg C)$
$p(C|L) = p(\neg C|L)kp(C)/p(\neg C)$

[31] One might worry that there is an element of arbitrariness at this step of the argument, since the truth of this claim depends on just what is meant by 'low' entropy. If the main system consists of extremely many pairs of urns and is only very slightly out of equilibrium at t_1, then an overwhelming majority of the subsystems may be in (or very near) equilibrium at t_1, so $p(C)$ may be rather small. Moreover, if the main system is close enough to equilibrium then k need not be very large at all, since $p(L|C)$ may be only slightly greater than $p(L|\neg C)$. To allay these worries, we must remember that the aim of my argument is, ultimately, to lend plausibility to the idea that the Principle of Parallelism holds between ordinary thermodynamic systems and the universe as a whole. The worrisome instances of the urn model have little bearing on that issue, since the universe is extremely far from equilibrium. In order for the urn model to provide an instructive representation of the thermodynamic case, let us stipulate that 'low' entropy means that a significant proportion of the total number of balls would have to switch urns for the main system to be in equilibrium. Similarly, we may stipulate that each subsystem contains a large number of balls, and that a pair of urns is in a low-entropy state when a significant proportion of the balls would have to switch urns for the pair to be in equilibrium. Given these stipulations, k will be unimaginably enormous, while $p(C)$ will not be too small; hence, $kp(C)/p(\neg C)$ will be quite large.

very likely to evolve toward equilibrium in both temporal directions from t_1, it is very unlikely that the pair of urns in question has significantly lower entropy at t_2 than at t_1.

In the "warm-up exercise" just described, it was assumed that all the subsystems were isolated for the duration of the experiment. But we can easily generalize the argument for cases in which a subsystem *branches off* (becomes isolated) from the main system for some shorter interval of time. Imagine that we have a collection of very many urns, with the following dynamics. On each trial, a ball is selected at random, as before; but now it is moved with equal probability to *any* of the other urns—with a special exception for the time interval $[t_1, t_2]$. During this interval, any ball selected from a certain subset of the urns can *only* be moved to another member of that subset, and no balls from outside the subset are allowed in. The subset is a *branch system* during the specified interval: it is isolated from the rest of the system.

To demonstrate that the Principle of Parallelism (POP) holds in this more general case, we need to show that the branch system is very unlikely to have an entropy gradient that slopes in the opposite direction from that of the main system, regardless of whether t_1 occurs before t_2 or *vice versa*. Let t_0 be some time that is closer to t_1 than to t_2 (the order of these times doesn't matter[32]), and assume all three times occur long after the initial distribution has been washed out. Now, suppose that at time t_0 the main system is in a low-entropy state—that is, some of the urns contain significantly more balls than others.[33]

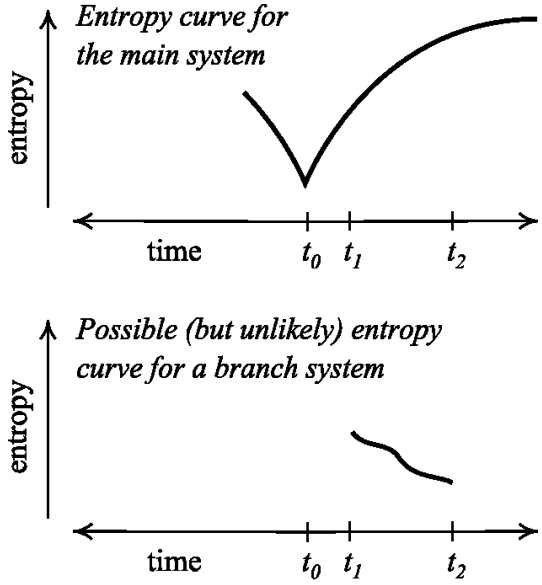
[32] In fact, the argument will go through even if t_0 occurs *between* t_1 and t_2, so long as it is closer to t_1 than to t_2.

[33] If the main system consists of a large enough number of urns as compared to the number of balls, it is likely that a few urns will have significantly more or fewer balls than the average at any given time, even when the system is in equilibrium. To keep things simple, let us stipulate that the number of urns is relatively small compared to the number of balls.

The crucial question is this: how likely is it that the branch system has significantly lower entropy at t_2 than at t_1? (Figure 2.4)

The reasoning employed in the previous example (the "warm-up exercise") can be straightforwardly extended to this case. Here, let us say that a branch system is a *contributor* if and only if its urns did *not* all contain about the same number of balls at the main system's low-entropy boundary t_0. Provided t_1 and t_2 are not extremely far apart, the probability that a given branch system will be in a low-entropy state at any time between t_1 and t_2 is very low if the system is not a contributor, i.e., if all of its urns contained about the same number of balls at t_0. On the assumption that the branch system *is* a contributor, on the other hand, the probability that it has low entropy at some time between t_1 and t_2 is much greater (assuming also that the main system does not reach equilibrium at any time between t_0 and t_1). Thus, as in the previous case, $p(L|C) >> p(L|\neg C)$, where L says that the subsystem is in a low-entropy state at some time between t_1 and t_2, and C says that it is a contributor. Since $p(C)$ cannot be too small,[34] it follows again by Bayes' theorem that $p(C|L) >> p(\neg C|L)$. That is, the branch system is very likely to be a contributor given that it has low entropy at any time during the period of its isolation, provided the main system does not reach equilibrium between t_0 and the time in question. Notice that nothing

Figure 2.4



[34] See footnote 31, above.

whatsoever has been assumed about the temporal order of t_0, t_1, and t_2; we have assumed only that the time at which the main system is stipulated to be in a low-entropy state, viz. t_0, is closer to t_1 than to t_2.

Now, the dynamics of the system make it overwhelmingly likely that any two urns containing significantly different numbers of balls will eventually contain about the same number, provided the urns are not isolated from each other. This is true regardless of whether the two urns belong to a branch system for part of their history, so long as they belong to the *same* branch system. Thus, if any two urns contain significantly different numbers of balls at any time before or after t_0, it is overwhelmingly unlikely that they contained roughly equal numbers of balls at t_0, provided the two urns have not been isolated from each other.[35] And in general, any two urns that differ at t_0 are likely to approach each other more or less monotonically in both temporal directions from t_0. Since t_2 is further (i.e., more trials away) from t_0 than t_1 is, it follows that a branch system is very unlikely to have lower entropy at t_2 than at t_1. So the Principle of Parallelism (POP) holds in general for the long-run dynamics of the urn model.

In the next section, a closely analogous argument will be developed within the framework of classical statistical mechanics. Before proceeding to reformulate the above argument in statistical mechanical terms, though, let us revisit the worry raised at the end of section 3. The urn model provides a simple way of illustrating why entropic parallelism fails when the main system is not isolated. On the assumption that the main system is in a very low-entropy state at t_0, the entropy gradient of a branch system

[35] The need for this proviso can be seen by considering the following example. Two urns might initially contain the same number of balls, but later become trapped in separate branches. If the average number of balls in one branch differs significantly from the average of the other, the two urns will diverge as they approach the averages of their respective branch systems.

(temporarily isolated subsystem) is overwhelmingly likely to increase in the direction leading away from t_0, as we have already seen. What happens if we add an additional assumption, namely, that the entropy of the main system is changed by an interaction with its environment during the interval $[t_1, t_2]$, when the branch system is isolated? This new set of assumptions is still temporally symmetric with respect to t_0, and it still follows from the dynamics that any *isolated* subset of the urns will tend to equilibrate with one another as the system evolves in either temporal direction from t_0. Whether the urns outside the branch system are disturbed—by dumping the contents of one of those urns into another, for instance—is just irrelevant to the behavior of the isolated urns, as intuitively it ought to be.

5. A Boltzmannian explanation of the Principle of Parallelism

One might suspect that the derivation of POP given in section 4 is peculiar to the urn model, or that it depends essentially on the fact that the dynamics of the urn model are stochastic. In classical statistical mechanics, the dynamical laws are deterministic—that is, the microstate of an isolated system at any given time, in conjunction with the laws, determines the microstate of the system at any other time. It is sometimes suggested that the temporal asymmetry of thermodynamics depends in some way on stochastic processes that occur in quantum interactions.[36] According to another hypothesis, sometimes called 'interventionism,' the asymmetry of thermodynamics depends at least partly on the fact that no real physical systems are perfectly energetically isolated. All systems interact at least to some small degree with their environments, and this is said to

[36] Albert (2000) suggests that such a view may be correct, though he also defends the purely cosmological account in both (2000) and (2009).

have a randomizing effect on the evolution of a thermodynamic system.[37] The aim of this concluding section is to show that neither of those hypotheses are required to derive the Principle of Parallelism. POP follows straightforwardly from the standard time-symmetric probability posit within the Boltzmannian framework. My argument in this section will involve a few more technical notions than the preceding sections. Readers familiar with the Boltzmannian approach to statistical mechanics will notice that my way of characterizing these technical notions is neither rigorous nor as detailed as it could be. Please excuse these shortcomings, as my aim is to make the argument more accessible to a general audience.

In Boltzmannian statistical mechanics, thermodynamic entropy is replaced by a different notion, Boltzmann entropy, which is associated not with the system's thermodynamic condition but with its 'macrostate.'[38] For present purposes, however, these differences need not concern us. The macrostate of a system can be regarded as equivalent to the system's thermodynamic condition, and Boltzmann entropy is closely analogous to thermodynamic entropy: it increases as the system approaches equilibrium.

The macrostate of a system at any given time is represented as a region in a high-dimensional abstract *phase space*, the space of all the possible microstates (exact microscopic states) that the system could be in. Boltzmann's crucial insight was that the Lebesgue measure (intuitively, the "volume") of the phase space region representing a system's macrostate increases exponentially as entropy increases.[39] In other words, high-

[37] For a brief discussion of this idea and the problems with it, see Sklar (1993), pp. 250-4.
[38] The Boltzmann entropy of a system is associated with a more fine-grained 'macrostate' than the thermodynamic macrocondition (see appendix), but this point is often ignored even in very technical presentations of the relation between statistical mechanics and thermodynamics.
[39] Boltzmann entropy is proportional to the log of the volume of the Boltzmannian macrostate, so the volume increases exponentially as the entropy increases (see appendix). Something similar is

entropy macrostates correspond to vastly greater volumes of phase space than low-entropy ones. In fact, the equilibrium macrostate takes up almost all of the available phase space.

If we regard the Lebesgue measure as a measure of probability, it follows (given relatively uncontroversial assumptions about the dynamical laws)[40] that isolated systems are overwhelmingly likely to evolve toward states of higher entropy *in both temporal directions*. That is, they are likely to have evolved *from* higher entropy in the past, and will probably evolve *toward* higher entropy in the future. Therefore, if we assume a uniform probability density over the system's present macrostate (intuitively, this is like assuming that the system is equally likely to be in any of the possible microstates that are consistent with its macrostate), it follows that the system will almost certainly evolve toward equilibrium in both temporal directions. So, treating the Lebesgue measure as a measure of probability results in a temporal symmetry closely analogous to the symmetry of the long-run dynamics of the urn models. If the entropy of the system is low at any given time t, the uniform probability density assumption makes it overwhelmingly likely that the system will evolve monotonically toward equilibrium in both temporal directions from t, just as in the case of the urn models. (Cf. Figure 2.2)

Now, suppose the main system is in a low-entropy state at t_0, and assume a uniform probability density over its macrostate at that time. Furthermore, suppose a

---

presumed to be true for thermodynamic macroconditions and their corresponding thermodynamic entropies, so the differences between thermodynamic and Boltzmann entropy make little difference for the present discussion. Both kinds of macrostates take up vastly greater volumes of phase space as the corresponding kinds of entropy increase.

[40] It is often assumed that the trajectories are ergodic, but such a strong assumption is not required for Boltzmann's argument; his reasoning will hold for almost any plausible hypothesis about the shape of the trajectories. (A trajectory is said to be *ergodic* if, in the limit as the trajectory is extended in both temporal directions from any point, the fraction of time it spends in any particular region is equal to the proportion of the accessible phase space occupied by that region.)

subsystem becomes isolated during the interval $[t_1, t_2]$, where t_1 is closer to t_0 than t_2 is. (Cf. figure 4.) What is the probability that the entropy of this branch system will be lower at t_2 than at t_1? I will argue that the entropy of the branch system at t_2 is very unlikely to be lower than the entropy at t_1; to the contrary, the entropy gradient of the branch system is likely to parallel that of the main system.

In the case of the urn models, I provided an argument for the claim that if the branch system is in a low-entropy state at any time between t_1 and t_2, it is much more likely to be a contributor than a non-contributor. (Recall that a *contributor* is a subsystem that "contributes" to the main system's having low entropy at a specified time.) Here, I will simply take an analogous claim for granted, since it is not in dispute in the literature. Sklar is happy to grant this much:

> Now it is certainly correct that the universe being in such an extraordinarily low-entropy state, we ought generally to infer, if we find a system isolated for a finite time, that its low entropy is a result of its interaction with the main system and not an improbable fluctuation from equilibrium. But, as far as I can see, that is all that cosmology by itself can really give us.[41]

That is to say, any branch system with low entropy is almost certainly a *contributor*, in a sense analogous to the notion introduced in section 4.

In the Boltzmannian version of my argument for POP, the uniform probability density, conditionalized on the low-entropy state of the main system at t_0, will play a role similar to that of the stochastic dynamics of the urn models. Given the uniform density

[41] Sklar (1993), pp. 325-6.

assumption, I shall argue, contributors are very likely to evolve to a state of higher entropy in both temporal directions from the specified low-entropy boundary of the main system. To see why this must be so, consider first the case of a branch system that is already isolated at t_0, the low-entropy boundary of the main system. If we begin with the uniform probability density conditionalized on the main system's macrostate at t_0, then restrict that density function to the dimensions of the main system's phase space that represent the branch system and renormalize, the resulting probability density will still be uniform. And, as previously noted, a uniform probability density over the macrostate of a system at t_0 implies that the system will almost certainly evolve monotonically toward equilibrium in both temporal directions from t_0.

Next, consider the more general case of a branch system that is isolated during the interval $[t_1, t_2]$, where t_1 is closer to t_0 than t_2 is. Evolving the main system's probability distribution forward (or backward, as the case may be), via the dynamical laws, will yield probability densities for the main system at t_1 and t_2. The evolved probability densities will almost certainly be non-uniform over the macrostate of the main system at those times. Nonetheless, those distributions will continue to make it very probable that the main system evolves toward higher entropy in the temporal direction that leads away from t_0.

Suppose we restrict that probability density function to the dimensions of phase space that represent the branch system and renormalize, as before. It is not immediately obvious whether this restricted density function will make it probable that the branch system evolves toward equilibrium in the same temporal direction as the main system.[42] Nevertheless, there are at least two compelling reasons to think that it will. First, it is

[42] Callender (2011b) makes this point (p. 100).

difficult to imagine how the fact that the system becomes isolated at some time other than t_0 could be likely to affect whether its entropy gradient parallels that of the main system; but as we have just seen, parallelism will hold in the special case where the branch system is already isolated at t_0.

Secondly, although the argument I gave for POP in section 4 did appeal to the stochastic dynamics of the urn model, the fact that the dynamics were non-deterministic did not play any direct role in the argument. All that mattered, so far as the dynamics were concerned, was that any two urns would tend to equilibrate with one another as the main system evolved in either temporal direction from the low-entropy boundary condition. A closely analogous claim will surely be true for typical Boltzmannian models—e.g., for classical gas particles bouncing around in a system of boxes with removable partitions. Given a uniform probability density over the macrostate of such a system at time t_0, any two boxes that have unequal pressures at t_0 will tend to equilibrate with one another in both temporal directions; and it seems implausible that this process could be reversed simply by isolating the two boxes in question from the *other* boxes. So the argument given in section 4 can be easily reinterpreted in terms of a classical deterministic system.

If all of this is correct, then the Principle of Parallelism follows from the very simplest way of assigning probabilities to phase space. It does not require any temporally asymmetric assumptions at all, nor does it require stochastic dynamical laws, nor again does it rely on the fact that real branch systems are only approximately isolated. Perhaps the temporal asymmetry of thermodynamics can, after all, be explained by the Past Hypothesis. The entropy gradient of the universe as a whole can be explained by the low-

entropy condition of the early universe, and thermodynamic irreversibility can in turn be explained by that global entropy gradient together with the Principle of Parallelism. Of course a few other wrinkles in this cosmological story have yet to be ironed out, as we'll see in the next chapter. Nonetheless, I hope that my defense of the Principle of Parallelism represents a step in the right direction.

Chapter 3: Replies to Earman and Callender

1. The Definability worry and the time scales objection

According to the so-called *Past Hypothesis*, the universe was in an extremely low-entropy condition just after the Big Bang. From this cosmological posit together with a very plausible probabilistic assumption, a satisfactory explanation for the asymmetry of thermodynamics can supposedly be constructed, at least in principle.[1] The devil is in the details. There is reason to doubt, in the first place, whether claims about the entropy of the universe as a whole even make sense. This objection originated with empiricists Reichenbach and Grunbaum, who worried that the concept of entropy might not be applicable to the cosmos.[2] More recently, John Earman has pointed out that Boltzmann entropy is a technically ill-defined concept in the context of General Relativity. In a similar vein, Craig Callender argues that there are difficulties applying the Boltzmannian framework to the universe even if we assume a classical (non-relativistic) theory of gravity. Suffice it to say that the gravitational contribution to entropy is not yet fully understood, and none of the standard statistical mechanical definitions of entropy—nor the thermodynamic definition, for that matter—can be unproblematically applied to gravitating systems, much less to the universe as a whole. Strictly speaking, claims about

[1] The idea that the asymmetry of thermodynamics might be grounded in an asymmetry of cosmic boundary conditions was first suggested by Boltzmann, and has been developed in various ways by Reichenbach (1956), Feynman (1965), Penrose (1989; 2005), Lebowitz (1993a; 1993b), Bricmont (1995), Price (1996), Albert (2000), Goldstein (2001), Callender (2004a; 2004b), and Wald (2006), among others.

[2] Reichenbach (1956), p. 114; Grunbaum (1963).

the entropy of the universe are nonsensical; hence the Past Hypothesis is "not even false," as Earman puts it.[3] This objection has been called the *definability worry*.[4]

Moreover, even if the concept of entropy is reworked to incorporate gravity, so that it is applicable at the cosmic scale, there are further problems concerning the explanatory link between the entropic behavior of the cosmos and the entropic behaviors of mundane thermodynamic systems. One such problem was addressed in the previous chapter: Winsberg contends that the Past Hypothesis, even if true, places little constraint on the behaviors of local thermodynamic systems. I have answered that objection by defending the Principle of Parallelism (POP)—the claim that the entropic behavior of an isolated subsystem is likely to parallel that of the main system. However, Earman poses an additional challenge along similar lines:

> The Boltzmann entropy of the initial state of the universe is low primarily because of the gravitational contribution, whereas that contribution is irrelevant for the kinds of subsystems of interest to us. … Nor is it plausible that the entropy increases we perceive in typical thermodynamic systems is the result of an increase in gravitational entropy that is transferred to the non-gravitational degrees of freedom; for on the time scales of typical thermodynamic observations and inferences, the values of the variables that characterize the gravitational degrees of freedom are effectively constant.[5]

[3] Earman levels this charge on p. 400; it also appears in the title of his paper, "The 'Past Hypothesis': Not even false."
[4] I borrow this name from Callender (2011b).
[5] Earman (2006), p. 419

As I will argue below, this objection—which I'll call the *time scales objection*—is best understood as a challenge to the claim that POP holds between the universe and its thermodynamic subsystems.

Before responding to these objections, it will be helpful first to examine the role that gravity is supposed to play in determining the entropy of cosmic systems. In the following section, I will give an intuitive characterization of the way in which gravitational variables might contribute to the entropy of a system, and explain why it is difficult to make sense of gravitating systems within the Boltzmannian framework. In section 3, I address the time scales objection, arguing that the problem can be solved if our cosmological explanation of thermodynamic irreversibility proceeds in two stages: the Past Hypothesis explains the entropy gradient of a certain large system that I call our 'cosmic neighborhood,' and that gradient in turn explains—via POP—the asymmetry of local thermodynamic systems. The definability worry is then addressed in section 4, where I propose a way of avoiding the issue altogether. I conclude by sketching an explanation of thermodynamic irreversibility that escapes all of the aforementioned objections.

2. Gravity and cosmic entropy

Assuming that the Second Law of Thermodynamics holds at the cosmic scale, the Past Hypothesis seems almost trivial. The Second Law implies that the entropy of a system only decreases when the entropy of its environment increases by at least as much, so the total entropy of a system and its environment can only change in the positive direction. Since the total entropy very often does change in the positive direction (but never in the

negative direction), of course the entropy of the universe must have been much lower in the beginning than it is now.

On the other hand, if we think of entropy solely in terms of ordinary thermodynamic variables like temperature and pressure, the idea that the universe began in a low-entropy state appears to have been refuted by cosmology. Late in the twentieth century, cosmologists discovered that the distribution of mass and energy shortly after the Big Bang was extremely smooth. In other words, the universe's temperature and pressure was very uniform, much like the equilibrium condition of air in a balloon. The temperature and pressure became less uniform as stars and galaxies formed through the process of gravitational collapse. In typical thermodynamic systems, such a transition from uniform to non-uniform temperature and pressure ordinarily represents a *decrease* in entropy. Thus, it may seem as though the universe has been evolving away from equilibrium and toward a condition of lower entropy. This point has been raised by physicist Robert Wald:

> The above claim that the entropy of the very early universe must have been extremely low might appear to blatantly contradict the ''standard model'' of cosmology: there is overwhelmingly strong reason to believe that in the early universe matter was (very nearly) uniformly distributed and (very nearly) in thermal equilibrium at uniform temperature. Does not this correspond to a state of (very nearly) maximum entropy, not a state of low entropy? [6]

[6] Wald (2006), p. 395.

Nevertheless, proponents of the Past Hypothesis have argued that the smooth distribution of mass and energy in the early universe in fact constitutes a very low-entropy condition. This claim is controversial, since physicists have not yet formulated a precise definition of entropy that applies to gravitating systems; but at the very least, the smooth condition of the early universe is closely analogous to a low-entropy condition, in a way that I will explain presently. At the cosmic scale, gravity is the dominant force,[7] and gravity is attractive. In contrast, the behavior of air in a balloon is dominated by the repulsive electromagnetic forces acting between the gas molecules. When repulsive forces are dominant, the entropy of a system increases as its temperature and pressure become more uniform; but this is not the case when the dominant force is attractive.

To illustrate, suppose we fill a container not with gas molecules but with hundreds of tiny magnets. Provided the magnets are allowed to move freely (without much friction against the container), the equilibrium condition will be one in which the magnets are clumped together in clusters or perhaps in one large cluster. The equilibrium condition of the system of magnets obviously will not be a condition in which the magnets are uniformly distributed throughout the container. Similarly, in cosmic systems where gravity plays the salient role, smoothly distributed matter is very far from equilibrium. In that respect, the condition of the early universe is at least analogous to a condition of very low entropy.

[7] Two caveats are in order. The most successful theory of gravity to date is General Relativity, which does not regard gravity as a force at all, but as curvature of spacetime. Moreover, recent cosmological research has revealed that the expansion of the universe is accelerating due to a mysterious force called 'dark energy'—a sort of anti-gravity which appears to be gaining hegemony over the attractive 'force' of gravity. I won't consider the possible role that dark energy might play in determining the entropic behavior of the universe.

Unfortunately, the entropy of a gravitating system cannot be straightforwardly characterized in terms of Boltzmannian statistical mechanics, since it is difficult (Earman suggests it may be impossible) to extend the Boltzmannian framework to General Relativity—our most successful theory of gravity. Earman considers an idealized general relativistic model, and points out that the measures most naturally used to calculate phase space volumes are non-normalizable. This prevents us from treating phase space volume as a measure of probability, making it impossible to give statistical mechanical explanations in the usual Boltzmannian way, at least within the idealized model.

What is worse, it turns out that Boltzmann entropy is a problematic concept even within the classical (non-relativistic) theory of gravity, as Craig Callender has recently argued. Callender suggests that the arguments ordinarily used to justify treating Boltzmann entropy as a physically significant quantity (analogous to thermodynamic entropy) might not be applicable in the case of gravitating systems, even in classical mechanics. After all, Boltzmann's original arguments for treating phase space volume as a measure of probability were given in the context of ideal gas models, and it is far from obvious whether similar arguments can be given in the gravitational context:

> Boltzmann's argument can plausibly be extended to some systems for which it was not originally intended, and new arguments mimicking Boltzmann's can show that the Boltzmann entropy for some non-dilute gases have physical significance … However, in the gravitational case we know we in general cannot use

Boltzmann's argument and there is not much reason to hope anything like it will help.[8]

This is not necessarily the death of the Past Hypothesis, fortunately. Callender responds to his own objection by pointing out that for at least some of the kinetic equations used in astrophysics, including the highly successful Fokker-Planck equation, it is possible to provide an argument analogous to Boltzmann's. Moreover, "one can show that the stationary or equilibrium distribution of [the Fokker-Planck equation] is equivalent to the solution one obtains from maximizing the Boltzmann entropy in the presence of an external potential."[9] Of course, this approach only works for subsystems of the universe than can be accurately described by a gravitational kinetic equation like the Fokker-Planck equation, which does not hold for the universe as a whole. To really make use of the Past Hypothesis, then, we'll have to do better.

Elsewhere, Callender suggests that a proponent of the cosmological account "might advocate rolling up one's sleeves and finding a way to write the Past Hypothesis in terms of the most fundamental physics. That Earman found it impossible in the known physics of some ideal systems is a long way from a no-go theorem implying it's impossible."[10] I would like to propose an easier strategy. Although the Past Hypothesis is typically construed as a claim about the Boltzmann entropy of the early universe, we might instead formulate the Past Hypothesis in some other way, perhaps a way that does not involve statistical mechanics at all. Indeed, Earman himself advocates such an approach:

[8] Callender (2010), pp. 48-9.
[9] Callender (2010), p. 53
[10] Callender (2011b), p. 102

> Needless to say, some sort of 'Past Hypothesis' in the form of a restriction on initial conditions will have to be used. But it is to be hoped the restriction can be justified without having to appeal to dubious and perhaps ill-defined notions such as the Boltzmann entropy of the universe."[11]

This idea will be revisited in section 4, where I argue that the asymmetry of thermodynamics can be explained in terms of the smooth initial condition of the universe, regardless of whether that condition can be regarded as one of low entropy.

3. The time scale problem

Even if entropy could be defined for general relativistic systems, Earman raises a further worry about the role that gravity is supposed to play in the cosmological explanation of thermodynamic irreversibility. The Boltzmann entropy of the early universe—if there is such a quantity—is low "primarily because of the gravitational contribution, whereas that contribution is irrelevant for the kinds of subsystems of interest to us." Mundane thermodynamic systems, cups of hot tea and the like, are too small for self-gravitation to be significant.

One might think that the entropy gradients of such systems will nonetheless tend to parallel that of the universe, as I argued in chapter 2. But Earman disagrees:

> Nor is it plausible that the entropy increases we perceive in typical thermodynamic systems is the result of an increase in gravitational entropy that is

[11] Earman (2006), p. 422

> transferred to the non-gravitational degrees of freedom; for on the time scales of typical thermodynamic observations and inferences, the values of the variables that characterize the gravitational degrees of freedom are effectively constant.[12]

Before responding to this argument, it is important to clarify just what the target of the argument is supposed to be. One may have the impression, from the quotation excerpted above, that Earman is refuting a view according to which ordinary thermodynamic processes are causally driven by gravitational collapse. He says that "the entropy increases we perceive" in ordinary thermodynamic processes cannot be "*the result of* an increase in gravitational entropy" (emphasis mine). That may be an infelicitous way of putting the point. Proponents of the Past Hypothesis surely are not committed to the dubious claim that increases in gravitational entropy somehow cause or drive the entropy increases that occur in ordinary thermodynamic processes. The Past Hypothesis is not intended to explain why entropy tends to increase in the future time direction; it is intended to explain why entropy does *not* increase toward the past as well. Recall the *reversibility problem*, which was explained in Chapter 1: assuming that an isolated system is equally likely to be in any of the microstates compatible with its present thermodynamic condition,[13] we get the prediction that entropy will increase in both time directions. The Past Hypothesis is supposed to solve this problem, by providing an explanation for the fact that isolated thermodynamic systems tend to have their lowest entropies at their past temporal boundaries. Given that an isolated system has low entropy

[12] Earman (2006), p. 419

[13] More precisely, we posit a probability density function that is uniform over the present macrocondition of the system, and zero elsewhere. I referred to this as the 'Naïve Statistical Postulate.'

at one temporal boundary, the fact that its entropy increases from there is readily explained by the standard statistical mechanical considerations, without further appeal to cosmology. Presumably Earman is not simply making this uncontroversial point.

Nor is Earman denying that entropy *deficits* (i.e. disturbances from equilibrium) can be transferred, over sufficiently long time scales, from gravitational to non-gravitational degrees of freedom. That sort of transfer obviously does happen when, for instance, a nebula (interstellar cloud of dust and gas) gravitationally collapses on itself and becomes much hotter than the surrounding environment as its pressure increases. Rather, he is challenging the idea that the entropy gradients of small, short-lived subsystems can reasonably be expected to parallel that of the surrounding cosmic system, whichever direction the latter's gradient goes. His target, in other words, is the Principle of Parallelism (POP), which I defended in chapter 2.

If I understand Earman correctly, his claim is that POP need not hold between the universe and local thermodynamic systems, because the universal entropy gradient—if there is such a thing—is primarily a matter of gravitational collapse. Gravitational collapse occurs so slowly that the entropy gradient of the universe is essentially flat over the time scales of ordinary thermodynamic observations. The universal entropy gradient will be so nearly flat over the lifespan of a cup of tea, for instance, that the tea won't know (so to speak) which way its own entropy gradient ought to go. This seems to me the most charitable interpretation of Earman's time scales objection, and it has some intuitive force, at least initially.

However, we must not ignore the fact that the gravitational contribution to entropy is quite significant over long periods, e.g. millions or billions of years. Earman's

objection presents no challenge to the claim that entropy gradients of long-lived isolated subsystems should parallel that of the universe as a whole. Moreover, if Earman is willing to grant that POP holds for long-lived subsystems, then his point about short-lived subsystems can be answered simply by noting that the entropy gradient of a small subsystem should parallel that of the larger and longer-lived subsystem that contains it, which in turn should parallel that of a still larger subsystem, and so on up to the universe as a whole.

The probability that the smallest subsystem parallels that of the universe will be equal to the product of all of these probabilities, of course, and one might worry that this product may be rather low: the more times POP has to be applied, the lower the resulting probability. However, I suspect it won't take many applications of POP to recover the predictions of mundane thermodynamics from the Past Hypothesis. Perhaps just two applications of POP will do the trick, as I will suggest with the following example.

Our solar system formed about 4.6 billion years ago from a process of gravitational collapse, probably from the debris of a prior supernova. If POP holds for long-lived cosmic subsystems, then we can expect the entropy gradient of our solar system to parallel that of the universe. (Actually, we need a slightly more general version of POP, since the solar system is not energetically isolated. We'll return to this important caveat in a moment.) But the entropy gradient of our solar system is not primarily a matter of gravitational changes, since the process of collapse that formed our solar system is essentially complete.[14] Moreover, the non-gravitational degrees of freedom responsible for our solar system's entropy gradient are not effectively fixed during the course of

[14] Since the planetary orbits are for the most part stable, very little gravitational collapse is occurring within our solar system. When the hydrogen in the sun is depleted, the sun's inner layers will collapse further, but that won't happen for another 5 billion years or so.

mundane thermodynamic processes. In the few minutes that it takes for ice to melt in your tea, for instance, the Sun dumps unfathomably many joules of energy outward toward the colder extremities of the solar system.[15] So far as Earman's argument is concerned, then, we can expect POP to hold between a cup of tea and the solar system, and in turn between the solar system and the universe. Since parallelism is obviously transitive (with the exception noted in the preceding paragraph), we can therefore expect the entropy gradient of the teacup to parallel that of the universe.

The foregoing example should suffice to convey the basic idea behind my reply to Earman, but it is not quite accurate as it stands. Recall that the Principle of Parallelism (POP), as defined in the previous chapter, only applies when both the main system and the subsystem are at least approximately isolated. As you may recall from my example of the air-conditioned room, we don't expect parallelism when either the main system or the subsystem has its entropy gradient reversed by an interaction with the environment. Turning on the air conditioner may lower the entropy of a room by making one side colder than the other; but we still expect ice to melt in hot tea, even when the air conditioner is running.[16] The room is not isolated, because it is exchanging energy with its environment via the air conditioner's electrical cord and heat pump. This interaction with the room's environment temporarily reverses the direction of the room's entropy gradient, but of course it does not affect the entropy gradients of isolated subsystems within the room. Likewise, the entropy gradients of non-isolated subsystems within the

[15] According to NASA's "Sun Fact Sheet" <http://nssdc.gsfc.nasa.gov/planetary/factsheet/sunfact.html>, the sun radiates energy at a rate of approximately 3.8×10^{26} joules per second.
[16] This example was given in chapter 2 (near the end of section 3).

room can be reversed. For example, the entropy of a refrigerator in the room can be reversed when it is plugged in.

Let us say that an interaction with the environment is *gradient-reversing* if and only if the system's entropy gradient would have sloped the other way, had the system been isolated during the period in question.[17] From that definition, it follows trivially that if POP holds for isolated systems, then non-isolated systems will exhibit similar parallelism so long as they do not undergo any gradient-reversing interactions. Thus, POP can be reformulated as the following principle, which I'll call "POP star":[18]

> (POP*) The entropy of a subsystem (whether isolated or not) is very unlikely to change in the opposite direction from that of the main system, so long as neither the main system nor the subsystem undergoes a gradient-reversing interaction with its environment.

Let us see how this works in the urn model from chapter 2. Gradient-reversing interactions are unlikely to occur with the dynamics of the urn model as previously constructed, but we can easily modify the dynamics to simulate a gradient-reversing interaction, e.g. by slowly dumping the contents of one urn into another while the usual dynamics continue for all the other urns. If we do something like that, parallelism obviously will fail. On the other hand, however, it is easy to show that parallelism holds

[17] There may be some ambiguity or arbitrariness here regarding what constitutes the same non-isolated system over time, since non-isolated systems by definition are exchanging matter and/or energy with their environments. For this reason, there may be cases in which there is no fact of the matter how "the system" would have behaved had it been isolated. Fortunately, the cases that will matter for my argument are not like that, as I will try to show with some examples in what follows.

[18] Please accept my apologies for the bad pun.

for both isolated and non-isolated urn systems whenever we don't mess with the dynamics. In fact, I've already given an argument to that effect, which I will review presently.

In section 4 of chapter 2, we supposed that the main system (a large collection of urns) is in a low-entropy state (i.e. some urns contain many more balls than others) at t_0, a time long after the initial distribution has "washed out"—that is, where the long-run dynamics are approximately time-symmetric in the sense I explained on pages 44-46. Then we considered a subsystem—a subset of the urns—that has low entropy at some other time between the interval t_1 and t_2 (where t_1 is closest to t_0). I showed that if the subsystem has low entropy at any time between t_1 and t_2, then it almost certainly has low entropy at t_0 as well—that is, the urns comprising the subsystem contain significantly different numbers of balls at t_0. I then argued that, given the dynamics of the urn model, any two urns that differ at t_0 are likely to approach each other more or less monotonically in both temporal directions from t_0. Therefore, I concluded, the subsystem is likely to approach equilibrium in the direction from t_1 to t_2 (since the latter is furthest from t_0). In other words, its entropy gradient "parallels" that of the main system.

Note that the foregoing argument did not hinge on the assumption that the subsystem is isolated between t_1 and t_2. I did use that assumption in chapter 2, but only for another purpose, namely to show that the Principle of Parallelism holds in particular for isolated systems or "branch systems," the same systems to which the Second Law of Thermodynamics applies. But in fact, I argued, any two urns will tend to equilibrate with one another as the system evolves in either time direction from t_0, provided the two urns are not isolated *from each other*. Whether they are isolated from the main system, I

pointed out, makes no difference. (See page 53.) So parallelism holds for all subsystems in the urn model, whether isolated or not. Apart from special interventions, gradient-reversing interactions are unlikely to happen in the urn model at all.

Of course, gradient-reversing interactions often do happen in real-life thermodynamics. So far as my reply to Earman's time scales objection is concerned, the crucial question is whether our solar system is presently undergoing a gradient-reversing interaction. If our solar system is not presently undergoing such an interaction with its environment, my reply to Earman will succeed: we can expect the entropy gradients of isolated thermodynamic systems to parallel that of the solar system, which in turn can be expected to parallel that of the universe, despite the fact that the universal entropy gradient is essentially flat over the course of typical thermodynamic observations. On the other hand, if the solar system *is* presently undergoing a gradient-reversing interaction, then I'll need to cast my argument in terms of some larger subsystem of the universe that is not presently undergoing such an interaction.

Alas, it seems that our solar system is indeed undergoing a gradient-reversing interaction, as I will explain. The idea that our solar system is undergoing a gradient-reversing interaction may seem counterintuitive, since this means that its entropy gradient would slope the other way if it were isolated. If the solar system were somehow isolated—if one could trap the entire solar system within a giant insulated thermos, say—surely its entropy would continue to increase, right? But this intuition mistakenly assumes that the entropy of the solar system is in fact increasing rather than decreasing. That assumption can and should be contested. Since the solar system radiates away more energy than it absorbs from its environment, the total *energy* of the solar system is

steadily decreasing. The fact that its energy is decreasing does not immediately imply that its *entropy* is decreasing too, of course, but significant decreases in energy tend to be accompanied by decreases in entropy.

To illustrate, imagine a balloon floating out in the cold darkness of deep space. As the gas inside the balloon radiates away some of its heat energy, the gas cools and the balloon shrinks. As a result, the thermodynamic entropy of the balloon decreases. The Boltzmann entropy of the balloon also decreases, because as the balloon's temperature and volume decrease, the space of possible microstates available to the gas shrinks as well. In other words, the volume of the region of phase space representing the balloon's macrocondition shrinks as the balloon loses energy. Boltzmann entropy is proportional to the log of that phase space volume; so if the volume decreases, so does the entropy.

Now our solar system is more or less stable in size, unlike the shrinking balloon; and thanks to the energy released from the sun, the solar system isn't cooling significantly either. Nonetheless, the entropy of the solar system is likely decreasing, due to the way in which energy is being lost. In terms of phase space, the dimensions corresponding to degrees of freedom at the subatomic level are shrinking as the sun uses up its nuclear fuel. The temperature of the sun isn't increasing (since the released energy is being radiated away), so the result is a net loss in both energy and entropy. Perhaps a dramatic increase in entropy elsewhere in the solar system could compensate for the sun's entropy loss, but that seems improbable since the remainder of the solar system is relatively stable. Thus, it appears that the entropy of the solar system is decreasing as it radiates away much of the sun's energy.

If the solar system were isolated within a giant insulated thermos, on the other hand, then its entropy would increase as the temperature would become more uniform throughout the solar system. Therefore, the loss of the sun's energy is a gradient-reversing interaction, as defined above: the solar system's entropy gradient slopes in one direction (down), but it would slope the other way (up) if the solar system were isolated. If all of this is correct, then neither POP nor POP* apply to our solar system, so I can't appeal to the solar system's entropy gradient in my reply to Earman.

Fortunately for me, my argument can be easily revised to avoid this problem. In place of the solar system, let us substitute what I'll call our *cosmic neighborhood*. By this, I mean the smallest region of space surrounding the solar system that is not undergoing a gradient-reversing interaction. Such a region surely exists, as I will argue in what follows. Imagine drawing a large sphere around our solar system. If we make the sphere very large, so that the sun is a tiny speck in the middle, the outflow of the sun's energy will be miniscule compared to the inflow of energy from stars outside the sphere. Unless there is a lot of non-luminous matter within the sphere to absorb it, most of that inbound radiation energy will pass right through the sphere and out the other side, so that the total energy within the sphere will still be decreasing. If we make the sphere large enough, however, eventually the inflow of energy will be roughly matched by the outflow of energy.

Provided the inflow and outflow of radiation energy are roughly equal, surely it makes little difference whether the inbound energy comes from the environment or is simply reflected back by the walls of a giant insulated thermos. Systems within the sphere should behave in more or less the same way in either scenario. Therefore, I

contend, such a sphere would not be undergoing a gradient-reversing interaction; it would behave more or less like an isolated system.

Our cosmic neighborhood is the smallest region like that. Surely there are very many regions around us through which the outflow and inflow of energy are roughly equal (though most of them are probably not spherical in shape), and surely not all of them are undergoing gradient-reversing interactions. Pick the smallest one that is not interacting with its environment in a gradient-reversing way: that's what I'm referring to as our cosmic neighborhood.

To keep track of our cosmic neighborhood over time, I will stipulate its persistence conditions in the same way: at any given moment, our cosmic neighborhood is the smallest region around us that is not undergoing a gradient-reversing interaction. The size and shape of our cosmic neighborhood may fluctuate dramatically over time; its borders may look arbitrary and gerrymandered. None of that matters, so far as my argument is concerned. What matters is simply that we're in *some* neighborhood that is not undergoing a gradient-reversing interaction with its environment.

Now, my reply to Earman can be recast in terms of our cosmic neighborhood rather than in terms of our solar system. Here is the argument, in brief. The gravitational contribution to the universe's entropy gradient is significant (*not* essentially flat) over long periods, and our cosmic neighborhood exists for a long period. Thus, we can expect the entropy gradient of our cosmic neighborhood to parallel that of the universe as a whole, Earman's objection notwithstanding. Moreover, I contend, the entropy gradient of our cosmic neighborhood is not primarily a matter of gravitational entropy. After all, there isn't much gravitational collapse going on in our little corner of the universe; but

there is quite a bit of nuclear fusion occurring in our sun and other nearby stars. So the entropy gradient of our cosmic neighborhood primarily involves the non-gravitational degrees of freedom, which are not effectively fixed over the course of mundane thermodynamic processes. For all Earman has said, therefore, we can still expect POP* to hold between mundane thermodynamic systems and our cosmic neighborhood, and in turn between our cosmic neighborhood and the universe as a whole.

I have argued that only two applications of POP* are required to recover the predictions of thermodynamics from the Past Hypothesis. The Past Hypothesis provides an explanation for the entropy gradient of the universe as a whole, which in turn explains the entropy gradient of our cosmic neighborhood (via POP*), and the gradient of our cosmic neighborhood then explains mundane thermodynamics (via a second application of POP*). Of course, all of this assumes that the notion of entropy itself is applicable at the cosmic scale. As we have already seen, Earman contests that claim as well, since it is not obvious how to make sense of the gravitational contribution to entropy. To that issue I now turn.

4. How to sidestep the definability worry

In my solution to the time scales problem, I used two applications of the reformulated Principle of Parallelism, which I've been calling POP*:

> (POP*) The entropy of a subsystem is very unlikely to change in the opposite direction from that of the main system, so long as neither the main system nor the subsystem undergoes a gradient-reversing interaction with its environment.

I have claimed that POP* is true. It follows trivially from POP (which I defended in chapter 2) together with my definition of a gradient-reversing interaction. However, even if POP* holds for systems with definable entropies, it cannot be legitimately applied to the universe as a whole if—as Earman contends, and as Callender worries—there might be no such thing as the entropy of the universe. My aim in this section is to avoid that problem, by arguing that the smooth distribution of matter near the Big Bang can still play the requisite explanatory role, regardless of whether it is possible to define a notion of entropy that applies to the universe. More specifically, I will argue that the entropy gradient of our cosmic neighborhood is dependent upon the gravitational dynamics of the universe as a whole. If this argument succeeds, then the first of my two applications of POP* will not be needed: the smooth condition of the early universe will provide a more direct explanation for the entropy gradient of our cosmic neighborhood (without relying on POP*), and the entropy gradient of our cosmic neighborhood in turn will explain local thermodynamic behavior (via a single application of POP*).

This new explanatory strategy will still require that the concept of entropy applies to our cosmic neighborhood. The idea that our cosmic neighborhood has well-defined entropy may be disputed, so let me offer a few words in defense of that idea. As noted above, the entropy gradient of our cosmic neighborhood primarily involves non-gravitational changes—in particular, the release and dispersal of energy from the nuclear fusion reactions occurring in our sun and other nearby stars. There was a time, billions of years ago, when gravitational changes did contribute to the entropy gradient (or something closely analogous to an entropy gradient) in this vicinity; but if entropy is ill-

defined for systems undergoing gravitational collapse, then my definition of 'cosmic neighborhood' is likewise inapplicable to such systems. For this reason, we should regard our cosmic neighborhood as existing only for as long as the region around and including our solar system has been gravitationally stable, or at least approximately so.[19] Then we can think of its entropy solely in terms of ordinary thermodynamic variables like temperature and pressure.

Since our cosmic neighborhood is a region through which the outflow of energy is matched by the inflow, its total energy is constant. Its entropy, however, is increasing as energy from the sun and other stars is distributed to colder areas.[20] So the entropy gradient of our neighborhood slopes upward in the future time direction, and downward in the past direction. For the purposes of my argument, though, it is not enough to simply point out that the entropy gradient slopes the appropriate way in the actual case; I need to show that it is counterfactually dependent upon the gravitational dynamics of the universe. In other words, I need to show that the entropy gradient of our cosmic neighborhood would slope in the opposite direction if the universe were evolving toward (instead of away from) a smooth condition like that of the early universe. For the cosmological explanation to work, it must turn out that our cosmic neighborhood will

[19] By 'gravitationally stable,' I mean that significant gravitational collapse is not occurring. More to the point, gravitational changes are making only a negligible difference to the entropy of our neighborhood.

[20] In particular, the amount of background radiation is increasing. At the same time, that background radiation is being diluted by the expansion of the universe, so in order for the total energy of our cosmic neighborhood to remain constant, it may be necessary for our neighborhood to grow larger over time. This complicates matters, because if the cosmic neighborhood increases in size, then so do the dimensions of phase space corresponding to its spatial degrees of freedom; and that will introduce an additional increase in Boltzmann entropy that results not from physical processes but from the way in which 'cosmic neighborhood' was defined. To avoid this complication, we can either discount the additional entropy increase or make the simplifying assumption that the size of the universe is constant.

have its lowest entropy at the temporal boundary nearest the smooth condition of the universe, regardless of whether that smooth condition occurs at the beginning or the end of time.

It is a bit difficult to imagine how the universe could evolve from its present state to a condition like that of the early universe. Some extraordinary coincidences would have to occur. Huw Price pictures it this way:

> Suppose, for a moment, that in the past thirty years or so, physics had discovered that the matter in the universe is collapsing towards a Big Crunch, fifteen billion years or so in our future—and that as it does so, something very, very extraordinary is happening. The motions of the individual pieces of matter in the universe are somehow conspiring to defeat gravity's overwhelming tendency to pull things together. Somehow, by some unimaginably intricate balancing act, the various forces are balancing out, so that by the time of the Big Crunch, matter will have spread itself out with great uniformity. A butterfly—nay, a molecule—out of place, and the whole house of cards would surely collapse![21]

As hard as it may be to imagine that scenario, it is still less easy to envisage how our cosmic neighborhood might behave if such a scenario were to occur.

To give our intuitions a boost, let's simplify the story, imagining that the universe consists of a single cloud of gas particles that do not emit radiation and do not undergo nuclear fusion. The cloud collapses under its own gravitational pull, eventually stabilizing as a sphere with a certain temperature and pressure—like a star or planet, but

[21] Price (2002), p. 115

one that neither emits nor absorbs energy from radiation. The sphere reaches equilibrium, and stays there for a very long time. But every now and then (after extremely long intervals of time) it randomly fluctuates out of equilibrium, upsetting the gravitational stability: certain parts become hotter, and as a result expand outward, while other parts become colder and collapse inward.

In the simplified scenario we've just imagined, the gravitational stability is upset by random anti-thermodynamic fluctuations: the temperature and pressure of the gas become less uniform, causing it to spread out in some places. Whereas ordinary thermodynamic behavior allows the gas to maintain gravitational stability in the shape of a sphere, anti-thermodynamic behavior allows the gas to "defeat gravity's overwhelming tendency to pull things together," as Price puts it. Similarly, in the more complex scenario that Price describes, the only plausible way in which matter could conspire to defeat gravity is by exhibiting dramatic anti-thermodynamic behavior on a cosmic scale. How else could the opposite of gravitational collapse occur? Consider the fact that when an interstellar cloud of ordinary matter (i.e. a nebula) collapses under its own gravitational pull, the increasing pressure raises the temperature, and some of the heat energy is radiated away. In order for the nebula to un-collapse back to its original state, then, the nebula must absorb energy from the surrounding space, even though the nebula is already hotter than its environment. Energy would have to flow from cold to hot! This suggests that in order for the universe to return to a smooth distribution of matter, ordinary thermodynamic processes would have to occur in reverse, at least for cosmic systems. Therefore, I claim, the entropy of our cosmic neighborhood (and of every other

cosmic neighborhood) would probably be decreasing rather than increasing if the universe were headed toward a future condition like that of the early universe.

I have argued that the entropy gradient of our cosmic neighborhood is dependent upon the gravitational dynamics of the universe, in the following way: if the smooth distribution of matter were to occur at the *end* of the universe, then our cosmic neighborhood would have its lowest entropy in the *future*; but since the universe in fact *began* with a smooth distribution of matter, our cosmic neighborhood has its lowest entropy in the *past*. Thus, the smooth distribution of matter in the early universe explains why our neighborhood's entropy gradient slopes the way it does, quite independently of whether the concept of entropy applies to the universe as a whole. This allows us to explain the asymmetry of thermodynamics in terms of the global gravitational asymmetry, even if the concept of entropy does not strictly apply on a global scale. That explanation will be outlined in the concluding section, below.

5. Conclusion

In this dissertation, I have defended a cosmological explanation for the temporal asymmetry of thermodynamics. As noted in the first chapter, attempts to provide such an explanation began with Ludwig Boltzmann, one of the earliest and most important contributors to our modern understanding of the entropic asymmetry. The modern account builds upon Boltzmann's crucial idea that the asymmetry of thermodynamics stems ultimately from an asymmetry in our cosmic environment, rather than from an asymmetry in the dynamical laws that govern physical processes. Whereas Boltzmann attributed that cosmic asymmetry to an extremely rare random fluctuation, however, the

modern account simply takes for granted that the universe was in a low-entropy condition—or something closely analogous to such a condition—at its past temporal boundary. This claim about the early universe has been called the Past Hypothesis.

The preceding two chapters of this dissertation have addressed the worry that the Past Hypothesis, even if true, places little constraint on the behaviors of mundane thermodynamic systems. (This was Winsberg's objection.) This objection was answered by appealing to the Principle of Parellelism (POP), which I defended in chapter 2. The Principle of Parallelism also played a crucial role in the present chapter, which answers two further objections to the modern cosmological account: the definability worry and the time scales objection.

To briefly recapitulate, the so-called 'definability worry' is the objection (raised by Earman, Callender, and others) that the concept of entropy cannot be meaningfully applied to the universe as a whole. Even if an adequate definition of entropy could be devised, moreover, the entropy gradient of the universe is primarily a matter of gravitational collapse. This poses a further problem, which I've called the 'time scales objection.' Since gravitational collapse occurs very slowly, the universal entropy gradient is essentially flat over the course of mundane thermodynamic processes. Therefore, Earman argues, it is not plausible that ordinary thermodynamic systems take their cues from the universal gradient. I believe that both of these further objections can be avoided by embellishing the cosmological account in the ways that I have suggested. I will conclude with an outline of the cosmological explanation of thermodynamics, as I conceive it.

The cosmological explanation of thermodynamic irreversibility proceeds in two stages. First, the entropy gradient of our cosmic neighborhood is explained in terms of the smooth condition of the early universe. As I argued in section 4, that explanation can be given without applying the concept of entropy to the entire universe. The concept of entropy must apply to our cosmic neighborhood, of course; but that seems unproblematic, since the entropy gradient of our neighborhood is not primarily a matter of gravitational changes (see pages 80-81, above). Moreover, the lifespan of our cosmic neighborhood is long enough to avoid the time scales objection: the universe's gravitational variables are not effectively constant over the course of our cosmic neighborhood's evolution.

The second stage of the explanation involves an application of the Principle of Parallelism. If POP holds for isolated systems, I pointed out, it follows trivially that non-isolated systems will exhibit similar parallelism so long as they do not undergo any gradient-reversing interactions. Our cosmic neighborhood is not undergoing a gradient-reversing interaction, by definition, since it was defined as the smallest region around us that is not undergoing such an interaction. Therefore, the reformulated version of the principle (POP*) holds for our cosmic neighborhood and its subsystems, which include local thermodynamic systems. This explains why mundane thermodynamic processes, like the melting of ice in hot tea, occur in one time direction but not the other: the entropy gradients of local systems slope in the same direction as the gradient of our cosmic neighborhood.

Epilogue: Sherlock's reply

Carefully, so as to avoid spilling any of the dregs from the tea leaves into his cup, the clever detective pours the last few drops from the teapot. Having had time to ponder Sklariarty's challenge, Sherlock had gone from badly flustered to robustly confident. With the faintest smile of calm satisfaction, he enjoys the last swig and clears his throat. "Given that the distribution of matter is very smooth at one temporal end of the universe but not the other, it very plausibly follows that ice is likely to melt in hot tea but not to spontaneously freeze out of lukewarm tea. What do you say to that, Professor?"

Appendix: Boltzmann's combinatorial argument

The following is a very brief (though not brief enough for a footnote) explanation of Boltzmann's famous argument for the claim that high-entropy macroconditions correspond to vastly greater volumes of phase space than low-entropy ones. A simple version of his argument, often referred to as the *combinatorial* or *permutational argument*, goes as follows.[22] Instead of representing a system's microstate as a single point in phase space, we can represent the microstate as a collection of points in *μ-space*—the phase space of a single particle. Now, if we partition the μ-space into a number of disjoint cells of equal measure, we can represent the state of the system in a coarse-grained way by specifying how many particles are in each cell. This coarse-grained state is sometimes referred to as the system's 'macrostate,' but we must be careful not to confuse it with the system's thermodynamic *macrocondition*. Thermodynamic macroconditions are specified in terms of "macroscopic" variables like temperature, pressure, and volume; whereas the coarse-grained states associated with Boltzmann entropy are not individuated by macroscopic variables. Moving a single particle from one μ-space cell to another will not in general make any macroscopic difference, for example; yet the Boltzmannian coarse-grained states are sensitive to such minute changes. To avoid confusion, I will refer to these coarse-grained states as *distributions* rather than 'macrostates.'

The volume of a distribution in *phase space* (not μ-space) is proportional to the number of permutations of particles that can be made without changing the distribution.

[22] For a more technical presentation of the argument, see Uffink (2007), (2008). A very simple version of the Boltzmannian argument is presented by Albert (2000), pp. 43-45. On a spectrum from technical and precise to more accessible and less precise, my rendition of Boltzmann's argument lies somewhere between those two presentations.

In other words, if we think of the particles haecceitistically, then the volume of a distribution is proportional to the number of compatible *arrangements*—where an arrangement is a specification of which particles are in which cells.[23] This fact about the volume of a distribution does not depend on any haecceitistic assumptions, by the way;[24] but the haecceitistic notion of an 'arrangement' makes the combinatorial argument a bit easier to follow, so I will continue to employ the notion in what follows. The crucial observation behind Boltzmann's combinatorial argument is, very roughly, that dispersed distributions correspond to more arrangements (and hence take up larger volumes of phase space) than concentrated distributions do.

For a simple illustration, suppose we have a system comprised of just 8 particles, which are free to move within a finite interval of just one spatial dimension. To simplify further, let's ignore the momentum of these particles, and regard a specification of their positions alone to be a complete specification of the system's microstate. The μ-space of this system is then just a one-dimensional interval. If we partition this interval into 4 equal cells, there will be very many possible distributions (165, to be precise). For instance, all 8 particles could be located in the leftmost partition. This distribution corresponds to only 1 arrangement, and hence takes up a minimal volume of phase space. On the other hand, the distribution in which each cell contains 2 particles corresponds to 2520 arrangements, and hence takes up a dramatically larger volume of phase space.

For ideal gasses, the Boltzmann entropy S of a distribution is defined as:

[23] I borrow the terms 'distribution' and 'arrangement' from Albert's presentation of the combinatorial argument.
[24] See Huggett (1999); also Albert (2000), pp. 45-7.

$$S = k \cdot \log W,$$

where k is Boltzmann's constant[25] and W is the Lebesgue measure (i.e., the volume) of the corresponding region in phase space. Since S is proportional to the log of the volume in phase space, conditions of high Boltzmann entropy correspond to exponentially greater volumes of phase space compared to conditions of low Boltzmann entropy. Something similar is presumed to be true in general for thermodynamic entropy: macroconditions of high thermodynamic entropy correspond to more dispersed distributions than low-entropy macroconditions, and thus correspond to larger regions of phase space. For example, consider a case in which gas molecules are confined by a partition to one side of a container. The thermodynamic entropy of the system will increase if the partition is removed and the gas is allowed to disperse throughout the whole container. Thus, conditions of high entropy in either sense (thermodynamic entropy or Boltzmann entropy) correspond to much greater volumes of phase space than conditions of low entropy.

[25] Boltzmann's constant is approximately 1.38×10^{-23} J/K, and is equal to the ideal gas constant divided by the Avogadro constant.

References

Albert, D. (1994), "The Foundations of Quantum Mechanics and the Approach to Thermodynamic Equilibrium," *British Journal for the Philosophy of Science* 45, 669-77.

Albert, D. (2000), *Time and Chance*, Cambridge, MA: Harvard UP.

Albert, D. (2009), "Physics and Chance," presented at the Foundations of Statistical Mechanics workshop held on May 13-14, 2009. URL = <http://fas-philosophy.rutgers.edu/hdemarest/StatMech.html>.

Bricmont, J. (1995). "Science of Chaos or Chaos in Science?" *Physicalia* 17, 159–208.

Callender, C. (2004a). "There is no puzzle about the low-entropy past," in C. Hitchcock (Ed.), *Contemporary Debates in Philosophy of Science*, Blackwell, 240-57.

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

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

Callender, C. (2011b), "The Past Histories of Molecules," in Claus Beisbart and Stephan Hartmann (eds), *Probabilities in Physics*, Oxford UP.

Callender, C. (2011c), "Thermodynamic Asymmetry in Time," *Stanford Encyclopedia of Philosophy* (Fall 2011 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/fall2011/entries/time-thermo/>.

Carroll, S. and J. Chen (2004), "Spontaneous Inflation and the Origin of the Arrow of Time," URL = <http://arxiv.org/PS_cache/hep-th/pdf/0410/0410270v1.pdf>

Davies, P. (1977), *The physics of time asymmetry*, University of California Press.

Earman, J. (2006), "The 'Past Hypothesis': Not even false," *Studies in History and Philosophy of Modern Physics*, 37, 399-430.

Feynman, R. (1965), *The Character of Physical Law*, Cambridge, Mass.: MIT Press.

Frigg, R. (2008), "A Field Guide to Recent Work on the Foundations of Statistical Mechanics", in *The Ashgate Companion to Contemporary Philosophy of Physics*, ed. Dean Rickles.

Frigg, R. (2009). "Typicality and the Approach to Equilibrium in Boltzmannian Statistical Mechanics," *Philosophy of Science*, 76, 997–1008.

Frigg, R. (2011a), "Entropy," in Claus Beisbart and Stephan Hartmann (eds), *Probabilities in Physics*, Oxford UP.

Frigg, R. (2011b), "Why Typicality Does Not Explain the Approach to Equilibrium," in *Probabilities, Causes and Propensities in Physics*, Mauricio Suárez (ed.), Kluwer, Dordrecht, pp. 77-93.

Frisch, M. (2010), "Does a low-entropy constraint prevent us from influencing the past?" in G. Ernst and A. Huttemann (eds.) *Time, Chance, and Reduction: Philosophical Aspects of Statistical Mechanics*, Cambridge UP, 13-33.

Goldstein, S. (2001), "Boltzmann's approach to statistical mechanics," in J. Bricmont, D. Durr, M. Galavotti, F. Petruccione, and N. Zanghi (Eds.), *Chance in Physics: Foundations and Perspectives*, Berlin: Springer.

Grünbaum, A. (1963), *Philosophical Problems of Space and Time*, New York: Knopf.

Huggett, N. (1999), "Atomic Metaphysics," *Journal of Philosophy*, 96, 5-24.

Lebowitz, J. (1993a), "Boltzmann's Entropy and Time's Arrow," *Physics Today* 46, 32–8.

Lebowitz, J. (1993b), "Macroscopic Laws, Microscopic Dynamics, Time's Arrow and Boltzmann's Entropy," *Physica A* 194, 1–27.

Loewer, B. (2007), "Counterfactuals and the Second Law," in Huw Price and Richord Corry (eds.), *Causation, Physics, and the Constitution of Reality*. New York: Oxford UP.

Myrvold, W. (unpublished), "Probabilities in Statistical Mechanics: Objective, Subjective, or a Bit of Both?", URL = <http://philsci-archive.pitt.edu/8642/>.

North, J. (2011), "Time in Thermodynamics," in C. Callender (ed.) *The Oxford Handbook of Philosophy of Time*, Oxford: Oxford UP, 312-52.

Penrose, R. (1989), *The Emperor's New Mind: concerning computers, brains and the laws of physics*. Oxford: Oxford UP.

Penrose, R. (2005), *The Road to Reality*, New York: Knopf.

Price, H. (1995), "Cosmology, time's arrow and that old double standard," in Savitt, S. (ed.), *Time's Arrows Today*. Cambridge UP, 66-94.

Price, H. (1996), *Time's Arrow and Archimedes' Point: New Directions for the Physics of Time*, Oxford: Oxford UP.

Price, H. (2002), "Boltzmann's Time Bomb," *British Journal for the Philosophy of Science*, 53, 83-119.

Price, H. (2004), "On the origins of the arrow of time: why there is still a puzzle about the low entropy past," in C. Hitchcock, ed., *Contemporary Debates in the Philosophy of Science*. Blackwell, 219-39.

Reichenbach, H. (1956), *The Direction of Time*. Berkeley: U of California Press.

Sklar, L. (1992), *Philosophy of Physics*. Westview.

Sklar, L. (1993), *Physics and Chance: Philosophical Issues in the Foundations of Statistical Mechanics*. Cambridge: Cambridge UP.

Sklar, L. (2009), "Philosophy of Statistical Mechanics," *The Stanford Encyclopedia of Philosophy (Summer 2009 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/sum2009/entries/statphys-statmech/>.

Strevens, M. (2008), *Depth: An Account of Scientific Explanation*. Cambridge: Harvard UP.

Strevens, M. (upublished), "Why Explanations Lie: Idealization in Explanation." URL = <http://www.strevens.org/research/expln/Idealization.pdf>

Uffink, J. (2007), "Compendium of the Foundations of Classical Statistical Physics," in *Philosophy of Physics (Handbook of the Philosophy of Science)*, J. Butterfield and J. Earman (eds.), North Holland, Amsterdam, 923-1047.

Uffink, J. (2008), "Boltzmann's Work in Statistical Physics," *The Stanford Encyclopedia of Philosophy (Winter 2008 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/win2008/entries/statphys-Boltzmann/>.

Wallace, D. (2013), "The Arrow of Time in Physics," in A. Bardon and H. Dyke (eds.), *A Companion to the Philosophy of Time*, West Sussex: Wiley.
Wald, R.M. (2006), "The Arrow of Time and the Initial Conditions of the Universe," *Studies in History and Philosophy of Modern Physics* 37 (3), 394-398.
Winsberg, E. (2004a), "Can Conditioning on the "Past Hypothesis" Militate Against the Reversibility Objections?" *Philosophy of Science*, 71, pp. 489–504.
Winsberg, E. (2004b), "Laws and Statistical Mechanics," *Philosophy of Science*, 71, 707-718.
Winsberg, E. (2009) "Laws and Chances in Statistical Mechanics," presented at the Foundations of SM workshop, URL = <http://fas-philosophy.rutgers.edu/hdemarest/StatMech.html>.