The Arrow of Time in Physics

David Wallace

April 5, 2012

1 Introduction

Essentially any process in physics can be described as a sequence of physical states of the system in question, indexed by times. Some such processes, even at the macroscopic, everyday level, do not appear in themselves to pick out any difference between past and future: the sequence run backwards is as legitimate a physical process as the original. Consider, for instance,a highly elastic ball bouncing back and forth, or the orbits of the planets. If these processes — in isolation — were filmed and played to an audience both forwards and backwards, there would be no way for the audience to know which was which.

But *most* physical processes are not like that. The decay of radioactive elements, the melting of ice in a glass of water, the emission of light by a hot object, the slowing down by friction of a moving object, nuclear or chemical reactions . . . each of these processes seems to have a direction to it. None of them, run backwards, is a physical process that we observe in nature; any of them, if filmed and played backwards, could easily be identified as incorrect.[1] In the standard terminology, these processes define *arrows of time*.

So: some physical processes are undirected in time, but most are directed. So what? The problem is that among the physical processes that seem to be undirected in time (at least, so it seems) are essentially all the processes that govern fundamental physics. Yet we have strong reasons to think that the equations that govern larger-scale physical processes are somehow derivative on, or determined by, those that govern fundamental physics. So we have a puzzle at least, a paradox at worst: if there is no fundamental directedness in fundamental physics, how does it come to be present in other areas of physics? The purpose of this essay is first to sharpen and make more precise this dilemma, and then to review the main strategies for solving or dissolving it.

Many processes beyond physics are also directed. Indeed, virtually every process studied in any science other than physics defines an arrow of time — to say nothing for the directedness of the processes of causation, inference, memory, control, counterfactual dependence and the like that occur in everyday

[1] And indeed, even our examples of everyday processes without a direction of time actually experience small asymmetric effects — friction, emission of gravitational radiation, etc, — so that sufficiently careful observation would pick out a direction of time there too.

life. But for reasons of space and of clarity, I confine my attention here to the arrow of time as it occurs in physics. (For broader accounts, see (e.g.) Albert (2000), Price and Weslake (2009), Price (1996), Callender (2011), and references therein). For similar reasons, I set aside the metaphysical question of whether time itself has some intrinsic directedness, given (say) by some flow of time, and whether this could in any way influence our understanding of the direction of time in physics. (See Maudlin (2007, ch.4) for considerations along these lines.) My question is more modest but perhaps thereby more manageable: can we understand the directedness of time within physics using only resources that are themselves from physics?

The structure is as follows. In section 2 I briefly discuss those features of microscopic physics which seem to conflict with time asymmetry. In section 3 I explain just how this conflict plays out in the important context of thermodynamics and statistical mechanics. In section 4 I review the main strategies available for resolving the conflict between reversible microdynamics and irreversible macrodynamics, working within the context of the quantitative, time-irreversible evolution laws that seem to apply to large-scale phenomena. In section 5 I broaden the discussion to cover other arrows of time within physics. Most of the physics I discuss is standard textbook fare; where this is so, I make no attempt to provide detailed references.

2 Features of underlying microphysics: reversal and recurrence

Many physical theories of the microscopic world have at various points been advanced. No such theory is currently believed by physicists to be *true*; the closest we have to one is the Standard Model of particle physics, which is inadequate in (at least) the gravitational domain. And in fact, when studying (say) melting ice, or friction, it is rare in the extreme to use the Standard Model, mostly because of its complexity; a variety of simpler microscopic models are normally used.

All such models in contemporary physics, however, fall under two frameworks: they are examples either of classical, or of quantum, mechanics, with the former framework generally regarded (by physicists, at any rate) as relevant only because it approximates the latter in certain circumstances. And in either framework, the particular physical theory can be discussed in very abstract terms: by specifying a *state space'* of possible states of the system at a given time, and an *evolution law*, which constrains the system's future and past states given its present state. In the classical case, the intepretation of the state space is unproblematic and uncontroversial. We can do classical mechanics on *phase space*, in which case the individual states unambiguously represent possible physical states of the system; or we can do it on the space of distributions over phase space, in which case the individual states unambiguously represent probability distributions over possible physical states of the system. In quantum mechanics, things are not so simple due to the infamous *quantum measurement problem*: the quantum state appears to behave for some purposes like a physical state, for others like a probability distribution over such states.[2] (I return to this issue later.)

But for our purposes, the interpretation of the state space can be set aside for now. The microphysics, stated in this abstract way — and quite independent of the interpretation of the state — has in general three properties that will be crucial to my discussion.

Invertibility

In both classical and quantum mechanics, the evolution law takes the form of a first-order differential equation (called Hamilton's equation in the classical case, Schrödinger's equation in the quantum). It is a characteristic property of such equations that the present state of the system completely determines both its past and future states, for all times. Leaving aside certain singularities that can occur in classical physics, this is a universal feature of the dynamical systems we will be considering.

Time reversal

There is a one-to-one map — $\mathcal{T}$, the *time reversal map* — of states to states such that

(i) if $f(t)$ is a continuous sequence of states (one for each time $t \in (-\infty, \infty)$) satisfying the dynamical equations of the theory (and so determining a possible history of the system being studied), then so is $\mathcal{T}f(-t)$.

(ii) $\mathcal{T}$ leaves alone the "macroscopically relevant" properties, so that $\mathcal{T}x$ has the same macro-level characterisation as x.

Time reversal captures the sense in which a dynamically possible process, run backwards in some appropriate sense, is itself dynamically possible.

The second condition is admittedly a little vague. It would be straightforward to interpret in certain (so-called *second-order*) formulations of classical mechanics, in which a system's state *and the rate of change of that state* are required to determine its past and future evolution: in that formalism, in general $\mathcal{T}$ acts trivially, and it is literally true that if $f(t)$ is a possible sequence of states, so is $f(-t)$. But in the first-order formalism we are adopting, it is easy to show that this strong form of time reversal is only possible if the dynamics are trivial, with $f(t)$ =constant. Furthermore, in quantum theory there does not appear to be a natural second-order formulation. For that reason, Albert (2000, ch.1) has argued that most physical theories are not truly "time-reversal-invariant" (to use the physicists' phrase); his remarks have sparked considerable

[2] This is a slightly heterodox way of expressing the quantum measurement problem; for a defence and exposition, see Wallace (2011).

discussion (see *inter alia* Earman (), Arntzenius and Greaves (2009), Callender (2000), and Malament (2004)).

For the purposes of understanding the arrow of time, however, this is all a red herring (as Albert himself freely acknowledges). For most theories of relevance to us, we can find a time reversal operator in the above sense which leaves invariant a system's energy, pressure, mass density, and essentially every other property of relevance to macroscopic, observable properties of a system. I will say that systems like this have the *time reversal property*, leaving to others the question of whether that is "true" time reversal invariance.

I said previously "most" theories have the time reversal property; are there exceptions? Within contemporary quantum theory, the full Standard Model very nearly has the property; that is, it has the property if we neglect certain extremely weak interactions between quarks and leptons. Furthermore, even the full standard model *does* have an operator — the CPT[3]symmetry operator — which satisfies condition (i) of time reversal, but does so at the cost of transforming particles to their antiparticles and carrying out a mirror reflection on those particles (turning left-handed gloves into right-handed ones, and so forth). Whether this CPT transformation satisfies (ii) — whether it leaves "macroscopically relevant" properties invariant — is somewhat context-dependent, and in particular depends upon whether the system in question is interacting with an environment. (A box full of anti-hydrogen gas has the same pressure, density, energy distributions as an equivalent box of hydrogen gas, but heaven help the surrounding landscape if the box contains antiparticles but its surroundings are made of "normal" particles.)

Recurrence

A system exhibits recurrence if a given state, if left to evolve under the dynamics for a sufficiently long time, returns to its original state (or more precisely, to a state arbitrarily close to its original state). It is a rather striking fact about both classical and quantum physics that, under reasonable conditions, recurrence (or something close to it) can be proved to occur. Specifically, in quantum physics,*any system* confined to a finite volume of physical space will undergo recurrence. The precise statement in classical physics is a little more complicated (see, e.g., Albert (2000) or Sklar (1993) for discussion), but essentially the same is true (in any case, classical mechanics is of interest to us in underpinning observed phenomena only insofar as it is a good approximation to quantum mechanics.)

Recurrence is a very striking feature of classical and quantum mechanics, but three points should be stressed:

(i) The timescales on which recurrence occurs are *incredibly long* for macroscopic systems. Very crudely, recurrence happens because a system just runs out of distinct places in phase space to explore, and so has to revisit

[3]CPT, for **C**harge, **P**arity, **T**ime. Also known as the PCT symmetry, the TCP symmetry, etc.

old haunts. But for macroscopically large systems, there are extraordinarily many places to try exploring first. The recurrence timescale for a gas of $\sim 10^{23}$ particles is approximately $10^{10^{23}}$, a time absurdly longer than the age of the universe.[4]

(ii) Recurrence is a property of the mathematical framework we use to describe isolated physical systems, but we should not expect it to apply to any real physical system in the universe. The reason is simply that the recurrence timescales are *so* long as to exceed any reasonable timescale on which the system can be expected to behave as if isolated.

(iii) We have no solid reason to think that the universe *as a whole* will undergo recurrence. In classical cosmology, the theory used to describe the large-scale features of the universe is general relativity, a theory which does *not* have recurrence as a property (basically because of the formation of singularities; see the appendix of Tipler (1994) for further discussion). In the absence of a fully worked out quantum theory of gravity, it is an open question whether we should expect recurrence in that theory.

So much for microphysics. We now turn to the derived theories of large(r)-scale physics, to see just how their properties (appear to) clash with what we have discussed above.

3 Irreversibility of macrophysics

There are two general kinds of irreversibility in physics: *qualitative principles* that are time-asymmetric, and *quantitative dynamical laws* that are time-asymmetric; I discuss both in turn.

Beginning with the qualitative principles, by far the most important are the principles of *thermodynamics*, the 19th-century theory that is still a key part of contemporary physics , chemistry, and engineering. Very loosely speaking, thermodynamics makes two main sets of predictions:

1. Large-scale systems, left to themselves, evolve into an "equilibrium" state characterised only by the system's conserved quantities and externally imposed parameters, *and stay there.* A litre of gas, for instance, in general requires $\sim 10^{27}$ real numbers to describe its state, because we need to specify the position and velocity of every single particle. Thermodynamics predicts that the gas will evolve into an equilibrium state characterised *uniquely* by its total energy and the volume of the box in which it is confined.

[4]What are the units? Nanoseconds. Or possibly, millenia. It doesn't matter. Suppose we want to convert $10^{10^{23}}$ nanoseconds into millenia. There are $\sim 10^{20}$ nanoseconds in a millenium, so we need to divide by 10^{20}. But $10^{10^{23}}/10^{20}$ is $10^{10^{23}-20}$, and $10^{23} - 20$ is as close to 10^{23} as makes no odds. This should give some impression of how ridiculously big the recurrence timescale really is.

2. Two such systems at equilibrium, allowed to interact, move into a joint equilibrium state; the reverse of this process does not spontaneously occur.

We can now see that thermodynamics appears to be in sharp conflict with the underlying microphysics: to be precise, it clashes both with time reversal, and with recurrence.

Conflict with reversal: according to thermodynamics, *every* state evolves to equilibrium. But if $f(t)$ is a sequence of states evolving to equilbrium that satisfies the evolution laws of the underlying physical system, then there is a sequence $\mathcal{T}f(-t)$ that also satisfies those laws. And since $f(t)$ and $\mathcal{T}f(t)$ are macroscopically indistinguishable, *this* sequence is evolving away from equilibrium. So it cannot be true that all states evolve to equilibrium and stay there.

Conflict with recurrence: Thermodynamic systems are normally confined to finite volumes; therefore, we would expect that their underlying dynamics undergoes recurrence. But in that case, it cannot be true that *any* system not already at equilibrium evolves to equilibrium and stays there forever: eventually, that system must return arbitrarily close to its initial, non-equilibrium, state. Recall, though, that the recurrence timescale is very long indeed, longer by far than the timescales on which thermodynamics is applied, and that there is in any case no *physical* prospect of a macroscopic system undergoing recurrence. The real significance of recurrence is this: there can be no valid mathematical derivation of the principle of approach to equilibrium that presupposes that the underlying dynamics are of a sort to produce recurrence.

However, in this part of the discussion my main focus is on the quantitative equations of macrophysics. Here it will be useful to consider a specific example (not often discussed in philosophy of statistical mechanics): the *Pauli master equation*. The Pauli equation[5] is applicable to quantum-mechanical systems where there is some small interaction that perturbs the otherwise interaction-free dynamics. In the absence of that interaction, the system can exist in any of a number of states of definite energy (labelled by some parameter k), or in quantum-mechanical superpositions thereof; the interaction permits transitions between such states.

The Pauli equation is a probability equation: it governs the evolution of the quantities $\rho_k(t)$, each of which is the probability of finding the system in energy state k at time t (the interpretation of this probability depends on one's preferred approach to quantum mechanics). The equation itself is

$$\frac{\mathrm{d}}{\mathrm{d}t}\rho_k(t) = \sum_{l\neq k} A(k \to l)(\rho_l(t) - \rho_k(t)) \tag{1}$$

[5]Not to be confused with the Schrödinger-Pauli equation that governs nonrelativistic particles with spin, interacting with a magnetic field.

where $A(k \to l) = 0$ unless states k and l have (very nearly) the same energy (according to the unperturbed dynamics). It has a reasonably intuitive interpretation: $A(k \to l)$ indicates the transition probability per unit time of a transition between state k and state l; the probability in some small time of the system leaving state k and ending up in state l is proportional to this transition probability times the probability of it being in state k in the first place; the probability of it leaving state l and ending up in state k is similarly proportional to this transition probability (assuming the probability is the same in both directions), times the probability of it being in state l in the first place. Combining these, and summing over all states l, we get the Pauli equation. It is widely and successfully used in physics, notably in the theory of atomic energy-level transitions.

The Pauli equation has an *equilibrium state*, or, more properly: an equilibrium probability distribution: if there are N energy states all with the same energy, then the probability distribution $\rho_l(t) = 1/N$ is a solution to the equation. Furthermore, for generic values of A (specifically, for any choice of A such that $A(k \to l)$ is never exactly zero for states of the same energy), arbitrary probability distributions converge on this solution. So, for a system governed by the Pauli equation, no matter what the appropriate probability distribution is at early times, at sufficiently late times it will be given, near enough by this steady-state distribution. In classical terms (which are slightly dubious in this quantum example): whatever energy state the system starts in, if we wait sufficiently long it will be equally likely to be found (according to the Pauli equation) in any state of the same energy. In quantum terms: if we prepare the system in any quantum state of definite unperturbed energy, it will in due course evolve (according to the Pauli equation) into a state which assigns equal probability to each eigenstate of unperturbed energy of the same energy eigenvalue.

This seems to mesh quite naturally with thermodynamic concepts of equilibrium.[6], but it is once again in direct conflict with the microdynamics. The microdynamics ought to obey time reversal, in which case for every state that will evolve *to* the steady-state distribution, there should be one that evolves away from it — but the approach to steady state of the Pauli equation is universal. And the microdynamics ought to obey recurrence, in which case with probability 1 any probability distribution ought to reconstitute itself — but the Pauli equation dictates that the steady-state probability distribution, once reached, is never left.

The Pauli equation is just one example of a very wide class of such irreversible equations; there are equations governing dissipation, quantum decoherence, radiative transitions and radioactive decay, nuclear and chemical reactions, and gas kinetics.[7] Typically, they are equations for the evolution of a certain subset

[6] Whether it *does* mesh is highly controversial, and turns on the choice between "Gibbsian" and "Boltzmannian" conceptions of equilibrium; this issue lies beyond the scope of this article. See Frigg (2007), Albert (2000), or Sklar (1993) for further discussion.

[7] Traditionally, one particular subclass of such equations — those governing classical electromagnetic radiation — are singled out and used to define a *radiation arrow of time*, to be regarded as distinct from the thermodynamical/statistical-mechanical arrow of time discussed

of the features of the overall system's state (in the Pauli equation, for instance, the probability distribution over states of definite unperturbed energy is tracked, but the terms that indicate quantum interference are ignored). In each case the same story plays out: these equations typically have attractor states (normally, attractor probability distributions) such that arbitrary initial conditions evolve towards those attractors.

In each case, then, the conflict with microdynamics persists.

But there is a crucial extra part of the story, not emphasised enough. These equations are highly successful at describing the phenomena. But in general they are not first formulated from observations of the phenomena. Instead, they are typically "derived" from the very microphysics with which they conflict. In the case of the Pauli equation, for instance, we can give an explicit formula for the transition probabilities $A(k \rightarrow l)$ in terms of the underlying quantum physics.[8]

How can this be? Of course, the supposed "derivations" are not literally such. Typically they have a very standard form.[9] Recall that each such equation tracks the evolution of only a subset of the total information specifying the system. We can think of a state of the system as decomposed into "relevant" degrees of freedom, tracked by the macro-level equation, and "irrelevant" degrees of freedom, ignored by that equation. (This terminology is due to Zwanzig (1961).) We then select, for a given value of the relevant degrees of freedom (written schematically as x) a state $\rho(x)$, in which the irrelevant degrees of freedom have a particular form — normally, a combination of discarding quantum entanglement and interference, and discarding classical correlation.[10] We can then define what Zwanzig calls the *projection map* on the space of states, which takes a given state with relevant information x to the state $\rho(x)$.

With this in hand, the "derivation" is now as follows:

1. Apply the projection map.

2. Evolve the system forward for a short time under the microdynamics.

3. Repeat.

It is fairly easy to see how this recipe could generate time-irreversible evolution, and permit attractor states. It is equally easy to see that the assumption that the true state evolves this way is in flat contradiction with the microdynamics. It is *not* easy to see — but is no less true for that — that it nonetheless very

---

here. I am sceptical that this distinction marks anything fundamental (especially as many of the cases I mention, including the Pauli equation, are not self-evidently thermodynamical, but space does not allow for a more detailed discussion. See Frisch (2006), North (2003) (and references therein) for more on the radiation arrow of time.

[8]Specifically: $A(k \rightarrow l) = \delta(E_k - E_l)(2\pi/\hbar) \langle k| \widehat{V} |l\rangle$, where the total-system Hamiltonian is $\widehat{H} = \widehat{H}_0 + \widehat{V}$ and where we have chosen $\widehat{H}_0$ and $\widehat{V}$ so that $\langle k| \widehat{V} |k\rangle = 0$. (The formula holds strictly only in the limit of infinite volume; for large but finite volumes we replace the delta function with a very sharply peaked function.)

[9]For more detail, see Zeh (2007, ch.3) or Wallace (2010).

[10]In the case of the Pauli equation, the state space is the space of mixed states, and the projection map discards the off-diagonal elements of the state in the unperturbed-energy basis.

reliably generates empirically adequate equations, when sufficient care is taken in the choice of projection. (Quite how one makes that choice is an issue to which I will return later.)

The message from this, I take it, is that it is insufficient for philosophers of statistical mechanics to explain how macro-level irreversibility is *compatible* with reversible microdynamics. They also need to explain the success of the apparently-contradictory framework physicists use to *derive*— or, better, to *construct* — irreversible macrodynamics from reversible underpinnings. I call this the *construction constraint*; we shall see more of it below, as we turn to the various strategies that have been advanced to resolve the apparent paradox.

4 Reconciling reversible and irreversible dynamics

As a matter of logic, if isolated systems confined to a finite volume display both time reversal and recurrence, we cannot validly derive any claim to the effect that such systems invariably display irreversible behaviour. The observed irreversibility in Nature — and the success of the irreversible equations we use to describe it — must therefore be for one (or more) of the following reasons:

(i) The systems in question cannot in fact be treated as isolated, finite-volume systems

(ii) The correct microdynamics for the systems in question is not in fact time-reversal-invariant and recurrent.

(iii) The observed behaviour is not in fact invariant, but requires additional constraints on the initial (or final, or other) conditions of the system.

Each has been suggested; I will discuss each briefly below.

Real systems are not isolated and finite-volume?

We can dispense quickly with the possibility that allowing for infinite volumes will be of help. It solves the recurrence issue (infinite-volume systems do not in general experience recurrence) but has no effect at all on the reversibility issue. (See Zeh (2007, p.91) for more on this point.)

The suggestion that real systems are not isolated at first sight looks more attractive. To be sure, real systems are *not* isolated, and this actually matters for the physics of those systems. In particular, even a very small interaction with an external system can serve, in effect, to induce the kind of erasure of entanglement and correlation that characterises the Zwanzig projection map. Interaction with an external environment, then, appears to provide a mechanism by which that map is physically implemented, justifying the construction of the Pauli equation and its kin and satisfying the construction constraint.

Appearences, however, can be deceptive. "Interaction-based" solutions suffer from three severe problems. Firstly, we can always expand the "system" to include as much of its environment as we like. Of course, typically even this larger super-system is not fully isolated from its environment, so we can play the same trick again ... and respond the same way again ... and so on. Ultimately, the issue turns on a key issue of quantum cosmology and the philosophy of cosmology: whether or not it is legitimate to treat the Universe as a whole as a physical system.[11]

Secondly, even if we allow environments that cannot simply be analysed as systems themselves, it is in a sense cheating. To reproduce the required effects, the environment's effect on the system must be "random" or "uncontrolled" or somesuch. But this is really a time-dependent notion: a series of interactions with random noise is in general the time reverse of an exquisitely delicate set of precisely aligned correlations. So by postulating that the environmental interaction has this form, we are postulating precisely what it is that we wish to explain. (This is an example of what Price (1996) correctly identifies as a pervasive double standard in discussions of the direction of time: assumptions that seem simple or reasonable or natural are often the time reverse of assumptions that seem no such thing, so that those assumptions build in precisely what they seek to explain.)

Thirdly, there seem to be real-life examples where systems are provably isolated from their environment and yet display (what looks awfully like) irreversible behaviour. In the so-called *spin-echo* experiment, a system behaves for an extended time just as the equations of irreversible dynamics predict — and then a carefully tailored magnetic pulse physically performs the time reversal operation, and the system smoothly returns to its initial state, demonstrating in the process that it was isolated from its environment. So irreversibility (of a kind) here seems to coexist with isolation. (For extensive discussion of the spin-echo experiment and of this argument, including a response to the initial reaction that the experiment shows precisely that isolated systems *do not* display irreversible behaviour, see Sklar (1993).)

4.1 The fundamental laws of physics are not time-reversal-invariant?

As I noted previously, it is generally said that contemporary physics is not invariant under time reversal, due to some subtle violations of time-reversal symmetry in the weak interaction between leptons and quarks. Occasionally (though rarely seriously) one hears this proposed as an explanation of irreversibility in macrophysics.

Problems abound. The interactions that violate time invariance are *extremely weak*, and there is no obvious mechanism by which they could lead to the ubiquitous irreversibility we come across in macrophysics. But more fundamentally, even contemporary physics is invariant under the CPT symmetry

[11] The question is under-explored in the philosophy literature (though arguably it is closely tied to questions of realism vs anti-realism; it divides physicists, with the divide roughly corresponding to the divide between Everettians and neo-operationalists; cf (Deutsch 1985) Carroll (2010) for the former view, Peres (1993), Fuchs and Peres (2000) for the latter.

(which, recall, reverses both spatial and temporal directions and replaces matter with antimatter). So if these exotica of the weak interaction are the cause of irreversibility, we should expect antimatter to display time-reversed anti-thermodynamic behaviour. So far as we can tell, it doesn't.[12] In addition, this strategy violates the construction constraint: physicists routinely construct irreversible macro-equations from models of the microphysics that are time-reversal invariant (and, indeed, which make no commitment as to whether they govern matter or antimatter). I conclude that the strategy is hopeless.

A more intereresting possibility arises from the foundations of quantum mechanics. Recall that in many textbook presentations of quantum mechanics — and in Dirac's (1930) and von Neumann's (1955) original codifications of quantum mechanics — the collapse of the wave-function is taken as a literal, physical process. This process is explicitly time-asymmetric; it might be said to define a "quantum-mechanical arrow of time". It is also physically highly unsatisfactory, since the point of collapse is defined in terms of higher-level concepts like 'measurement', 'observation', or even 'consciousness'. (Indeed, the unsatisfactory nature of wave-function collapse is frequently taken as synonymous with the quantum measurement problem; see Wallace (2011) for criticism of this.)

Most solutions to the quantum measurement problem (notably the Everett interpretation and Bohm's pilot-wave theory, in which the wave-function is physical but always evolves unitarily, and neo-operationalist or neo-Copenhagen approaches, in which it is not physical) do not assume a physical collapse. But an alternative strategy is to accept collapse but to formulate it in a physically satisfactory way: the so-called 'dynamical collapse' strategy. This amounts to replacing the Schrödinger equation with a new dynamical equation that is neither reversible, nor time-symmetric nor recurrent. (The classic example, in the non-relativistic regime, is the GRWP theory developed by Ghirardi, Rimini and Weber (1986) and developed further by Pearle (1989); see Bassi and Ghirardi (2003) for a review.) Both Penrose (1989) and Albert (2000) have proposed such a link between dynamical collapse and the arrow of time.

From the perspective of this article, the most attractive feature of this strategy is that it amounts to a physical modification of the dynamics so as to implement something like the Zwanzig projection: the kind of collapse mechanism needed to solve the quantum measurement problem pretty reliably ends up implementing the kind of discard-correlations, remove-interference transformation that is made in the construction of irreversible macrodynamics. The construction constraint is thus thoroughly satisfied.[13]

The downsides are three-fold. Firstly, there is no evidence at all that dy-

[12]To be sure, we don't generally manufacture antimatter in the lab in large quantities. But we do manufacture it in small quantities, and in those situations it obeys the normal laws of radioactive decay; cosmologists and high-energy astrophysicists also need to consider antimatter in bulk, and they make empirical predictions using the normal rules for constructing irreversible laws.

[13]I should note in passing that Prigogine (Prigogine (1984); see Bishop (2004) for philosophical discussion) has proposed modifying *classical* mechanics to achieve essentially the same effect. I refrain from further discussion of this proposal due to a lack of clarity in how it helps resolve the problems of *quantum* statistical mechanics.

namical collapse occurs,[14] and so far only very limited success in generalising dynamical-collapse theory beyond the restricted regime of nonrelativistic particle mechanics (the state of the art is Tumulka (2006), which generalises the GRWP theory to relativistic physics under the fairly draconian restriction that there are no interactions). Secondly, the spin-echo objection made previously applies here too: the spin-echo system seems to display thermodynamic behaviour but cannot have undergone any truly irreversible evolution. Thirdly, the dynamical collapse has the same general form as the Zwanzig projection, but this does not of itself guarantee that it can play the role of the Zwanzig projection.

4.2 A need for special conditions?

Proposals to explain the arrow of time based on a particular initial state of the universe have been espoused frequently in recent philosophy of physics (see, for instance, Lebowitz (2007), Goldstein (2001), Callender (2009), Penrose (1994), or Albert (2000)) and mostly take one particular form, which I will briefly expound here. First, it is observed that for a typical macroscopic system, the overwhelming majority of the volume of its state space corresponds to the same macroscopic state of the system (this is provable in the case of an ideal gas, highly plausible in the case of a weakly interacting gas, and open to debate in general[15]). It is then proposed that, for rather generic forms of the dynamics (those that are not "ridiculously special", as Goldstein puts it), we should expect the system's to make its way into the set of equilibrium states and stay there for a very long time (eventually it will recur, but only on timescales far longer than those that are empirically relevant).

Now, this cannot be true universally, because of time reversal invariance. But it can be true almost-universally: true, that is, for all but a very small minority of states ('small' with respect to some appropriate measure of state-space volume, not with respect to cardinality). So advocates of this approach postulate some constraint on the state of the system. The most straightforward such constraint is probabilistic: we can postulate that the microstate of the system is drawn randomly from some appropriately smooth probability distribution over all states compatible with the present *macroscopic* conditions of the system. (These proposals are almost invariably made on the assumption that classical physics holds, so I refer to probability distributions and not, more properly, to possibly-mixed quantum states.)

There is a snag with this postulate, though. By time reversal invariance, if it is sufficient to guarantee that the future evolution of the system will be towards

[14]This is controversial. Advocates of dynamical collapse tend to claim that our observations of macroscopic definiteness are themselves direct observational evidence of collapse. This seems question-begging to me, but further discussion would take us beyond the scope of this article.

[15]It is explicilty false for an important class of systems — those with spontaneous symmetry breaking (cf Binney et al (1992) or, for philosophical discussion, Reutsche (2011) or Batterman (2001) — but it would be fairly straightforward to fix the account to avoid this problem.

equilibrium, it is equally sufficient to guarantee that its *past* evolution is towards equilibrium: after all, the argument given above is itself invariant under time reversal. This is solved by an additional postulate: that the macroparameters which describe the initial state of the universe (or at any rate some very early state) correspond to some particular, far-from-equilibrium state — a 'low-entropy' state in the jargon of statistical mechanics (where 'entropy' refers to Boltzmann entropy, a measure of the state-space volume of microstates compatible with a given macrostate). This postulate is generally called the *past hypothesis* (following Albert); here I call it the Low-Entropy Past Hypothesis, or LEPH.

Given the LEPH, the probability postulate needs to be reformulated in one of two equivalent ways:

1. The probability distribution over microstates appropriate for the *present* time is some appropriately smooth distribution over all microstates, conditionalised both on the present-day macro-state and on the LEPH.

2. The probability distribution over microscates appropriate for the *initial state of the Universe* is some appropriately smooth distribution over all microstates, conditionalised on the LEPH.

The second is, in my view, conceptually preferable, as it makes clear that we are imposing two independent conditions on the initial state of the universe: one on its macrostate (the LEPH), one on its microstate (that it be given by a reasonably-smooth probability distribution over microstates compatible with some macrostate or other). Call this latter hypothesis the *smooth past hypothesis*, or SPH (following Wallace (2010), 'S' could also stand for 'simple'). The Past Hypothesis strategy, then, consists of LEPH+SPH, with LEPH receiving most of the attention in discussion (and most of the criticism; see, in particular, Earman (2006)).

Does it work? As stated, it seems to violate the construction constraint: it gives a rather qualitative argument that systems must evolve to equilibrium, without providing an underpinning for the equations that actually govern that evolution (and which, in some cases, tell us that the evolution of the system will not in fact go to equilibrium, or not for an extremely long time). But this is too quick: there are, in fact, good technical reasons to think that if SPH holds at some time t_0, the evolution of the system's relevant degrees of freedom at subsequent times $t > t_0$ such that $|t < t_0|$ is much less than the recurrence time will satisfy the equations constructed by the Zwanzig method, This seems to be exactly what the Past Hypothesis strategy requires. (For details, see Wallace (2010) or Zeh (2007).) Of course, if this is correct, then parity of reasoning tells us the *time reverse* of the macro-equations will hold at times $t < t_0$. But since in the Past Hypothesis strategy we choose t_0 to be the time of the begiinning of the Universe, this does not in fact lead to problems (at least within classical cosmology.[16])

[16]For (speculative) discussion of extending thermodynamics and statistical mechanics before the Big Bang, see Carroll (2010).

But notice that the LEPH component of the Past Hypothesis strategy does no work whatsoever in this accoun, and so apparently can be dropped entirely. The Past Hypothesis strategy reduces to a condition on the micro-level structure of the initial state of the Universe. (For a more extended version of this argument, see Wallace (2010).)

Before moving on, I should mention that the status of any initial condition requirement (contingent fact? physical law? something else?) is unclear and controversial. For opposing views, see Callender (2004) and Price (2004).

5 The qualitative arrows revisited

The discussion in the previous section dealt almost exclusively with the quantitative arrows of time: the relation between time-irreversible dynamics at one scale and time-reversible dynamics at another. Let us now return to qualitative distinctions in physics between past and future.

5.1 The thermodynamic arrow

As I noted previously, the most important such distinctions — certainly those most discussed — are induced by thermodynamics: the tendency of systems to go to equilibrium states and stay there, the irreversible constraints on which operations on such states are permitted. Indeed, my account of the arrow of time is somewhat anomalous in the literature: *most* discussions try to account directly for the asymmetries of equilibrium thermodynamics, without much digression into, or appeal to, the details of non-equilibrium physics. (Perhaps most famously, the old idea of *ergodicity*, along with many modern variants, is often appealed to for this purpose; see, e.g.,Malament and Zabell (1980) and Frigg and Werndl (2011).)

This is, I think, a mistake. We (that is: physicists) have a very good, quantitative, detailed understanding of how, when and if a given system evolves to equilibrium, applicable to a wide variety of systems ranging from thermonuclear reactions through to the mixing of gases or the cooling of hot bodies.[17] To be sure, there are deep philosophical problems involved in understanding how the myriad models and equations used in this analysis — this has been the main topic of my discussion so far. But given a solution to those problems, the asymmetries of thermodynamics do not obviously pose additional philosophical conundra: they emerge, rather, as dynamical consequences of the quantitative equations that govern the macro-world. Conversely, though, it is not obvious that a general argument that a given system must *at some point* reach equilibrium will help us understand the microphysical underpinnings of the actual models that govern its rate of approach to equilibrium.

[17]Lest there be doubt, I am not claiming that we have such a quantitative understanding in all or even most cases. But the field is marked by steady progress.

5.2 The computational arrow

There is a longstanding tradition in physics of claiming that certain computation processes — in particular (in contemporary accounts) the erasure of memory — are necessarily aligned with the thermodynamica arrow of time. This tradition has come under sustained attack by philosophers of physics in the last decade — notable criticisms include Earman and Norton (1998, 1999), Norton (2005), and Maroney (2005, 2010); relevant defences include Bub (2002), Bennett (2003), and Ladyman, Presnell, Short, and Groisman (2007) — but the details quickly get technical and lie beyond the scope of this article. (For an anthology of classic readings on the subject in the physics literature, see Leff and Rex (2002).)

5.3 The singularity arrow

For a less-frequently discussed example of a qualitative constraint, consider the singularity structure of the universe according to general relativity (readers unfamiliar with general relativity may wish to skip this paragraph). We seem to live in a universe with many final singularities (formed within black holes) but only one initial singularity (the Big Bang). Were this not to be the case, there would be general failures of determinism in cosmology, because subsequent initial singularities would require new data to be specified in some way that GR does not constrain. Hawking's "cosmic censorship principle" elevates this quasi-empirical observation to a general principle: such subsequent initial singularities must either fail to exist or at any rate must be sealed within event horizons. But this is explicitly time asymmetric; from the time-reversed point of view, the final singularities within black holes become naked initial singularities, and each time-reversed black hole becomes a new failure point for determinism. (For further discussion, see Earman (1995) on singularities in general, or Penrose (1989) on their implications for time reversal symmetry.)

Having said this, the actual processes by which black holes form in our universe seem reasonably clear — albeit our analysis of these processes rely crucially on a variety of time-asymmetric quantitative equations of the macro-world — and it is not as if there are plausible dynamical processes by which naked singularities could form. So it might well be that a solution to the problem of time in quantitative macrophysics might also help here. But the topic is murky — and made more so by the links between black holes and thermodynamic entropy, links which remain tantalizing but little-understood.[18]

5.4 The quantum-mechanical arrow of time

Although the normal dynamical equations of quantum theory respect time reversal invariance, the notorious "collapse of the wave-function" on measurement

[18] The canonical papers are Bekenstein (1973) and Hawking (1976); for philosophical discussion, see Curiel and Bokulich (2009) and references therein. See also Penrose (1989) for deep — but speculative — proposals connecting black hole entropy with the arrow of time (and the quantum measurement problem, and the problem of quantum gravity, and the mind-body problem).

does not. This prima facie suggests that quantum measurements define an additional arrow of time. Whether this is actually the case depends on ones preferred solution to the measurement problem:

- In no-collapse approaches, like the Everett interpretation[19] or the de Broglie-Bohm theory,[20] the underlying dynamics remains time-reversal invariant, and "wave-function collapse" is an effective, macro-level process, the irreversibility of which is just a special case of the dynamical irreversibilities we have already considered.
- In dynamical-collapse approaches, like the already-mentioned GRWP theory, collapse is a fundamental physical process which breaks time-reversal invariance by fiat. As we have already seen, there is some prospect that this fundamental irreversibility could ground other irreversibilities.
- In the traditional Copenhagen interpretation, and its various modern successors, external concepts like 'observation' or 'measurement' play an irreducible role in quantum theory and cannot be analysed in terms of closed physical processes; in these approaches, therefore, it is not really possible to discuss the arrow of time in a non-question-begging way.

5.5 The cosmological arrow

Sometimes one can get carried away in looking for arrows of time. The universe is expanding, right? So it's bigger in the future than in the past, right? So there must be an asymmetry between past and future in cosmology!

Well, yes. But there is not anything obviously mysterious about it *per se*. General relativity — itself a perfectly time-symmetric theory — predicts that cosmological solutions to its equations have the form of expansion from an initial singularity — possibly followed by collapse back to a final singularity, possibly — and, so current data suggests, actually) — proceeding indefinitely. In the latter case there is no obvious problem to explain. Even in the latter case, it's fairly easy to come up with anthropic grounds to expect that we live in the expanding phase, rather than the contracting phase; it's also not clear that this is the kind of thing that needs explanation.

Having said all this, the time reverse of expansion is contraction, so it *is* reasonable to ask why we describe the current state of affairs as an expansion rather than a contraction. The most direct answer would be in terms of our perceptions and memories — but therefore lies outside the scope of this article. Internally to physics, though, we can ask why the "cosmological arrow" is aligned as it is with the arrow defined by thermodynamics and by the irreversible equations of macrophysics. Put another way, why is the direction of time defined by the approach to equilibrium the same as the direction defined by the expansion of the Universe?

[19] Everett (1957); see Wallace (2012) or Saunders, Barrett, Kent, and Wallace (2010) for discussion.

[20] Bohm (1952); see Cushing, Fine, and Goldstein (1996) for discussion.

The answer will depend on our preferred understanding of the arrow of time in quantitative macrophysics. But at least for the solutions which postulate a particular boundary condition — macro or micro — there is no deep puzzle. If the Universe expands forever, then (presumably) there is only one end at which to impose that boundary condition. Even if it re-contracts, we have to impose it at one end or another, and the aforementioned anthropic considerations give us fairly good reason (mostly to do with how long stars shine for) to expect to be in the period of the Universe's history when systems go to equilibrium in the expansion direction rather than the contraction direction.

5.6 The "arrow of increasing entropy"

Entropy is, loosely, a measure of (microscopic) disorder; more precisely, it is a quantity defined in both thermodynamics and in the quantitative equations of statistical physics. Each of these theories has as a central result that entropy is non-decreasing over time; indeed, the claim that entropy never goes down — often called the *second law of thermodynamics*[21] — is sometimes taken as *defining* the thermodynamic arrow of time.

That might suggest that there is no residual arrow of entropy over and above that given by any successful explanation of thermodynamics. But thermodynamics predicts only that entropy does not go *down*: it — and the irreversible dynamics that governs approach to equilibrium — are entirely compatible with the idea that the Universe just started off in a state of maximal entropy, and stayed there.[22] In fact, the entropy of the early Universe was far lower than its present-day value, so that a robust direction of time is defined by the direction of increasing entropy.

It has been argued that this fact stands in need of explanation. But it is unclear whether this is so, and it is further unclear whether there is any *special* fact about the low entropy initial state of the universe — as opposed to any of a number of other facts about the early universe, such as its geometry, or its admixture of particles and antiparticles — which creates a particular puzzle. Arguments that there *is* something special about the low entropy of the early universe generally turn on the idea that "almost all" possible initial conditions have very high entropy. Penrose (1989) uses a vivid analogy: if God chose the initial state of the Universe by throwing a dart at random into the space of initial conditions, it would be stupendously unlikely that he would pick the actual, low-entropy, initial condition. Penrose (and others) have suggested that some new physics is required to make sense of the mystery.

Now, there are continuum many initial conditions, so this sort of objection only makes sense with respect to some kind of probability measure over the space of conditions — and it's not clear what would justify a particular measure. But more importantly, why should *any* measure have any particular significance?

[21] Often, but inaccurately; see Uffink (2001).

[22] Our theories of classical cosmology are *not* so compatible — they don't have non-trivial time-invariant solutions. This seems to create another problem for this idea, one not often discussed (though see (Earman 2006)).

To put the matter crudely: yes, the "God threw a dart into the space of initial conditions" theory, absent new physics, looks incompatible with the data. But that particular account of Creation didn't look all that impressive in any case.

6 Epilogue

If there is a general moral to this article, it is that we do best to approach the problems of the arrows of time in physics in small steps and with attention to the detail. It is tempting to set the problem up directly as a conflict between our microphysics and thermodynamics — or even as a conflict between our microphysics and our observations. Put that way, it can easily seem as if the problem is insurmountable. But in the first instance the problem shows up as an incompatibility between certain equations of macrophysics and the very pieces of microphysics that are used to construct them. If we cannot understand that apparent inconsistency, there is no hope of understanding contemporary physical practice; conversely, if we *can* understand it then we approach other problems of time asymmetry in physics with the inestimably useful tool of time-asymmetric dynamical equations, whose own time asymmetry is (ex hypothesi) understood.

I promised an article focussed exclusively upon physics, but I will bend that promise in this final paragraph, to suggest that something of the same moral might be applicable to our attempts to ground some of the more "philosophical" asymmetries — inference, prediction, memory, control, causation, and so forth. It's tempting to assume that such attempts (insofar as they attempt to be naturalistic) should work inside a framework of time-symmetric physics, and to ground the target asymmetry in some asymmetry of boundary conditions. But if we can understand why our world is describable in terms of (admittedly non-fundamental) time-asymmetric physics, we have access to a major resource to use in giving a grounding to the philosophical asymmetry in question. Granted, no metaphysically fundamental account will be possible this way. But metaphysical fundamentality is overrated: for many purposes, an account that is at the level of emergent, macro-level physics ought to do everything we need of it.

Acknowledgements

I have benefitted greatly in my understanding of this topic from conversations with David Albert, Katherine Brading, Harvey Brown, Tim Maudlin, Simon Saunders, and Wayne Myrvold. I wish to acknowledge financial support from the John J. Reilley Center at Notre Dame University and the Arts and Humanities Research Council of the UK.

References

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

Arntzenius, F. and H. Greaves (2009). Time reversal in classical electromagnetism. *British Journal for the Philosophy of Science 60*, 557–584.

Bassi, A. and G. Ghirardi (2003). Dynamical reduction models. *Physics Reports 379*, 257.

Batterman, R. W. (2001). *The Devil in the Details: Asymptotic Reasoning in Explanation, Reduction and Emergence*. Oxford: Oxford University Press.

Bekenstein, J. D. (1973). Black holes and entropy. *Physical Review D 7*, 2333–2346.

Bennett, C. H. (2003). Notes on landauer's principle, reversible computation, and maxwell's demon. *Studies in the History and Philosophy of Modern Physics 34*, 501–510.

Binney, J. J., N. J. Dowrick, A. J. Fisher, and M. E. J. Newman (1992). *The Theory of Critical Phenomena : an introduction to the renormalisation group*. Oxford: Oxford University Press.

Bishop, R. C. (2004). Nonequilibrium statistical mechanics brusselsaustin style. *Studies in the History and Philosophy of Modern Physics 35*, 1–30.

Bohm, D. (1952). A Suggested Interpretation of Quantum Theory in Terms of "Hidden" Variables. *Physical Review 85*, 166–193.

Bub, J. (2002). Maxwells demon and the thermodynamics of computation. *Studies in the History and Philosophy of Modern Physics 32*, 569–579.

Callender, C. (2000). Is time 'handed' in a quantum world. *Proceedings of the Aristotelian Society 100*, 247–269.

Callender, C. (2004). There is no puzzle about the low-entropy past. In C. Hitchcock (Ed.), *Contemporary Debates in the Philosophy of Science*, pp. 240–257. Oxford: Blackwell.

Callender, C. (2009). The past hypothesis meets gravity. In G. Ernst and A. Hütteman (Eds.), *Time, Chance and Reduction: Philosophical Aspects of Statistical Mechanics*, Cambridge. Cambridge University Press. Available online at http://philsci-archive.pitt.edu/archive/00004261.

Callender, C. (2011). Thermodynamic asymmetry in time. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy (Fall 2011 Edition)*. http://plato.stanford.edu/archives/fall2011/entries/time-thermo/.

Carroll, S. (2010). *From Eternity to Here: The Quest for the Ultimate Theory of Time*. Dutton.

Curiel, E. and P. Bokulich (2009). Singularities and black holes. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy (Fall 2009*

Edition). http://plato.stanford.edu/archives/fall2009/entries/spacetime-singularities/.

Cushing, J. T., A. Fine, and S. Goldstein (Eds.) (1996). *Bohmian Mechanics and Quantum Theory: An Appraisal*, Dordrecht. Kluwer Academic Publishers.

Deutsch, D. (1985). Quantum Theory as a Universal Physical Theory. *International Journal of Theoretical Physics 24*(1), 1–41.

DeWitt, B. and N. Graham (Eds.) (1973). *The many-worlds interpretation of quantum mechanics*. Princeton: Princeton University Press.

Dirac, P. (1930). *The Principles of Quantum Mechanics*. Oxford University Press.

Earman, J. What time reversal invariance is and why it matters. *International Studies in the Philosophy of Science 16*.

Earman, J. (1995). *Bangs, Crunches, Whimpers, and Shrieks: singularities and acausalities in relativistic spacetimes*. Oxford: Oxford University Press.

Earman, J. (2006). The 'past hypothesis': Not even false. *Studies in the History and Philosophy of Modern Physics 37*, 399–430.

Earman, J. and J. Norton (1998). EXORCIST XIV: The wrath of Maxwell's demon. part I. from Maxwell to Szilard. *Studies in the History and Philosophy of Modern Physics 29*, 435–471.

Earman, J. and J. Norton (1999). EXORCIST XIV: The wrath of Maxwell's demon. part II. from Szilard to Landauer and beyond. *Studies in the History and Philosophy of Modern Physics 30*, 1–40.

Everett, H. I. (1957). Relative State Formulation of Quantum Mechanics. *Review of Modern Physics 29*, 454–462. Reprinted in DeWitt and Graham (1973).

Frigg, R. (2007). A field guide to recent work on the foundations of thermodynamics and statistical mechanics. In D. Rickles (Ed.), *The Ashgate Companion to the New Philosophy of Physics*, pp. 99–196. London: Ashgate.

Frigg, R. and C. Werndl (2011). Explaining thermodynamic-like behaviour in terms of epsilon-ergodicity. *Philosophy of Science 78*, 628–652.

Frisch, M. (2006). A tale of two arrows. *Studies in the History and Philosophy of Modern Physics 37*, 542–558.

Fuchs, C. and A. Peres (2000). Quantum theory needs no "interpretation". *Physics Today 53*(3), 70–71.

Ghirardi, G., A. Rimini, and T. Weber (1986). Unified Dynamics for Micro and Macro Systems. *Physical Review D 34*, 470–491.

Goldstein, S. (2001). Boltzmann's approach to statistical mechanics. In J. Bricmont, D. Dürr, M. Galavotti, F. Petruccione, and N. Zanghi (Eds.), *In: Chance in Physics: Foundations and Perspectives*, Berlin, pp. 39. Springer. Available online at http://arxiv.org/abs/cond-mat/0105242.

Hawking, S. W. (1976). Black holes and thermodynamics. *Physical Review D 13*, 191–197.

Ladyman, J., S. Presnell, A. Short, and B. Groisman (2007). The connection between logical and thermodynamic irreversibility. *Studies in History and Philosophy of Modern Physics 38*(1), 58–79.

Lebowitz, J. (2007). From time-symmetric microscopic dynamics to time-asymmetric macroscopic behavior: An overview. Available online at http://arxiv.org/abs/0709.0724.

Leff, H. and A. F. Rex (2002). *Maxwell's Demon: Entropy, Information, Computing* (2nd ed.). Institute of Physics Publishing.

Malament, D. (2004). On the time reversal invariance of classical electromagnetic theory. *Studies in the History and Philosophy of Modern Physics 35*, 295–315.

Malament, D. and S. L. Zabell (1980). Why Gibbs phase space averages work: the role of ergodic theory. *Philosophy of Science 47*, 339–349.

Maroney, O. J. E. (2005). The (absence of a) relationship between thermodynamic and logical reversibility. *Studies in the History and Philosophy of Modern Physics 36*, 355–374.

Maroney, O. J. E. (2010). Does a computer have an arrow of time? *Foundations of Physics 40*, 205–238.

Maudlin, T. (2007). *The Metaphysics within Physics*. Oxford: Oxford University Press.

North, J. (2003). Understanding the time-asymmetry of radiation. *Philosophy of Science 70*, 1086–1097.

Norton, J. D. (2005). Eaters of the lotus: Landauers principle and the return of maxwells demon. *Studies in the History and Philosophy of Modern Physics 36*, 375–411.

Pearle, P. (1989). Combining Stochastic Dynamical State-Vector Reduction with Spontaneous Localization. *Physical Review A 39*(5), 2277–2289.

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

Penrose, R. (1994). On the second law of thermodynamics. *Journal of Statistical Physics 77*, 217–221.

Peres, A. (1993). *Quantum Theory: Concepts and Methods*. Dordrecht: Kluwer Academic Publishers.

Price, H. (1996). *Time's Arrow and Archimedes' Point*. Oxford: Oxford University Press.

Price, H. (2004). Why there is still a puzzle about the low-entropy past. In C. Hitchcock (Ed.), *Contemporary Debates in the Philosophy of Science*, pp. 219–239. Oxford: Blackwell.

Price, H. and B. Weslake (2009). The time-asymmetry of causation. In H. Beebee, C. Hitchcock, and P. Menzies (Eds.), *The Oxford Handbook of Causation*. Oxford: Oxford University Press. Available online at phisci-archive.pitt.edu.

Prigogine, I. (1984). *Order out of Chaos*. Bantam Books.

Reutsche, L. (2011). *Interpreting Quantum Theories*. Oxford: Oxford University Press.

Saunders, S., J. Barrett, A. Kent, and D. Wallace (Eds.) (2010). *Many Worlds? Everett, Quantum Theory, and Reality*, Oxford. Oxford University Press.

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

Tipler, F. (1994). *The Physics of Immortality: Modern Cosmology, God and the Resurrection of the Dead*. New York: Doubleday.

Tumulka, R. (2006). Collapse and relativity. In A. Bassi, T. Weber, and N. Zanghi (Eds.), *Quantum Mechanics: Are There Quantum Jumps? and on the Present Status of Quantum Mechanics*, pp. 340. American Institute of Physics Conference Proceedings. Available online at http://arxiv.org/abs/quant-ph/0602208.

Uffink, J. (2001). Bluff your way in the second law of thermodynamics. *Studies in the History and Philosophy of Modern Physics 32*, 305–394.

von Neumann, J. (1955). *Mathematical Foundations of Quantum Mechanics*. Princeton: Princeton University Press.

Wallace, D. (2010). The logic of the past hypothesis. Available online at http://users.ox.ac.uk/ mert0130/papers.shtml.

Wallace, D. (2011). Decoherence and its role in the modern measurement problem. Forthcoming in *Philosophical Transactions of the Royal Society A*; available online at http://arxiv.org/abs/1111.2187.

Wallace, D. (2012). *The Emergent Multiverse: Quantum Theory according to the Everett Interpretation*. Oxford: Oxford University Press.

Zeh, H. D. (2007). *The Physical Basis of the Direction of Time* (5th ed.). Berlin: Springer.

Zwanzig, R. (1961). Memory effects in irreversible thermodynamics. *Physical Review 124*, 983.