Implications of quantum theory in the foundations of statistical mechanics

David Wallace*

August 15, 2001

Abstract

An investigation is made into how the foundations of statistical mechanics are affected once we treat classical mechanics as an approximation to quantum mechanics in certain domains rather than as a theory in its own right; this is necessary if we are to understand statistical-mechanical systems in our own world. Relevant structural and dynamical differences are identified between classical and quantum mechanics (partly through analysis of technical work on quantum chaos by other authors). These imply that quantum mechanics significantly affects a number of foundational questions, including the nature of statistical probability and the direction of time.

1 INTRODUCTION

> Classicality simply does not follow "as $\hbar \to 0$" in most *physically* interesting cases ... The Planck constant is $\hbar = 1.05459 \times 10^{-27}$ erg s and — *licentia mathematica* to vary it notwithstanding — it is a *constant.* (W. H. Zurek and J. P. Paz[(1)])

Why should we consider quantum issues when working in the foundations of statistical physics? The simple (too simple) answer is that classical physics is false. If our purpose, in doing foundational work, is to understand the actual world, it is necessary to use a theory which validly describes that world.

Of course, nonrelativistic quantum mechanics is strictly speaking also false. It is believed to be only a tractable limiting case of quantum field theory, which in turn is expected someday to be replaced by quantum gravity, and even that theory may not be the final word on physics.

This is not simply a foundational dilemma: we face it at all levels in physics. It is obviously impossible to work with theories which we do not yet know, and usually computationally intractable to work with our most fundamental known theories; yet if we do not do this then we are working with a theory which is known to be false — why, then, should we believe its conclusions?

From a theoretical point of view, we address this problem by finding subdomains of more fundamental theories in which they approximate our less fundamental, but more tractable theories — so general relativity reduces to special relativity in regions where matter densities are low and space is nearly flat, while special relativity in turn reduces to Newtonian mechanics when relative velocities are small.

*Centre for Quantum Computation, Oxford University, Oxford OX1 3PU, U.K.; email david.wallace@merton.ox.ac.uk.
I would like to thank Michael Bowler, Harvey Brown, David Deutsch, Artur Ekert and Simon Saunders for encouraging comments and constructive criticism on earlier versions of this work, Guido Bacciagaluppi, Hannah Barlow, Clare Horsman, Lev Vaidman, Lorenza Viola and Wojciech Zurek for useful conversations, an anonymous referee for extensive feedback, and in particular Jeremy Butterfield for extensive and detailed comments and advice.

Delineating these subdomains, however, is problematic. Although we can get a good general idea of their locations, it is generally not possible to predict perfectly at what level we can explain some natural phenomenon, and it is always possible to be surprised. (To take one example, it turns out that the color of gold is due to relativistic corrections to its electron orbits).

To see this approach in more detail, we can distinguish three theories. The first is classical mechanics (CM), treated as an exact theory of particles moving under Hamilton's equations. The second is quantum mechanics (QM), which is presumed to be approximated by CM in certain domains. Restricting QM to these domains gives something which we could call *classical-domain quantum mechanics*, or CDQM — this isn't really a new theory, of course, just a subtheory of QM.

In this language, the approximation of QM by CM amounts to some "approximate isomorphism" ι between CM and CDQM:[1]

$$CM \overset{\iota}{\leftrightsquigarrow} CDQM \subseteq QM. \tag{1}$$

ι is required to preserve most of the structural properties of CM, but we will not require it to preserve all of them — instead we take it as an open question as to what structural features can be preserved by ι.

As physicists studying a particular class of quantum phenomena, we start by tentatively deciding whether these phenomena lie in the domain of CDQM. If so, we try applying CM to the phenomena, but in doing so it is always necessary to remember the possibility that the particular features of CM which are doing explanatory or predictive work in our proposed account of the phenomena may not actually be preserved by the isomorphism ι. In this case, it is necessary make slight modifications to modify the structure of the part of CM being used so as to better approximate CDQM, or if necessary to abandon CM altogether and work entirely with QM.

Fleshing out the details of this analysis of theory approximation and approximate structure-preservation lies rather beyond the scope of this paper. See Wallace[(2)] for (some) further discussion of the idea of an "approximate isomorphism" between theories; see also the literature on structural realism (e. g. , Worrall,[(3)] Psillos,[(4)] Ladyman.[(5)])

Very much the same approach could be used to describe nonrelativistic mechanics (NRM) and special relativity (STR). NRM is a limit of STR as $c \to \infty$, but the value of c is fixed and finite so this limiting relationship cannot directly be the reason why NRM sometimes gives quite accurate results. Instead we consider domains of STR in which typical energies seem quite low, restrict STR to those domains (producing low-energy STR, or LSTR), and construct an approximate isomorphism τ between NRM and LSTR,

$$NRM \overset{\tau}{\leftrightsquigarrow} LSTR \subseteq STR. \tag{2}$$

To count as an approximate isomorphism τ should preserve most of the structure of NRM, but not necessarily all of it.

The movement of electrons in magnetic fields provides an example of this process (in the context of relativistic and non-relativistic QM, not CM). There, we begin by considering electrons with velocities far below c, which we suppose are covered by LSTR. However, on investigation we find that relativity predicts a coupling between the magnetic field and the internal spin of the electron, and to maintain the correspondence between NRM and LSTR in this domain we have to add such a coupling term to the non-relativistic Hamiltonian. If we were to allow the energy of the electron to rise too high, we would have to give up on even this modified form of NRM and work directly with STR, of course. In either case, there is not really any way internal to NRM for us to predict difficulties with it; such predictions have to come from STR itself.

Is this approach appropriate to foundational work? Yes and no, but mostly yes. There is genuine reason to be interested in the structure of old theories even when they are known not to apply to the world: trying to find the classical description of charged point particles, or examining the nature of spacetime as given by Newtonian physics, are interesting tasks in their own right and potentially might shed light on other problems. However, by and large we are interested in foundational problems

[1] If desired, ι may instead be understood as an approximate isomorphism between models of CM, and models of QM restricted to a certain domain.

— and, generally, in problems in physics and in philosophy — because of what their solutions would tell us about the world. In that case, clearly it is necessary to work with a theory which applies to the world, or at least to that part of it under study.

What of statistical mechanics? We might hope that at least a large portion of statistical mechanics could be understood using classical concepts — after all, classical statistical mechanics makes the assumption that classical mechanics is valid, and presumably there is a domain of phenomena in which that assumption is reasonable. There would be a separate set of foundational problems in purely *quantum* statistical mechanics — in the statistical mechanics of atomic nuclei, or of the quark-gluon plasma, for instance — but they could be addressed after the classical case was properly understood. However, this approach would depend upon the existence of domains in which statistical mechanics applies but in which quantum mechanics makes no appreciable difference to the underlying physics: in other words, domains of QM in which statistical mechanics applies, which are approximately classical in the sense that there exists an approximate isomorphism ι between the classical and quantum descriptions of the domains, and in which those features of classical mechanics relevant to statistical-mechanical explanation and prediction are preserved by ι.

It is the purpose of this paper to show that no such domains exist. Statements of classical statistical physics such as 'there exist isolated systems whose dynamics are chaotic and given by Hamilton's equations', or 'ordinary macroscopic systems have some unknown state which is approximately localized in momentum and position' turn out to be, not just only approximately true, but universally false.

As such, it would seem necessary to allow for at least certain quantum- mechanical phenomena in foundational work on statistical physics.[2] Specifically, our attitude to foundational problems in classical statistical mechanics should be that, if they are preserved by ι, they are to be taken seriously, but if instead ι does not preserve them then they are artifacts of CM which — whatever their intellectual interest — cannot be directly relevant to statistical mechanics as it applies to the actual world.

In the following I shall attempt to sketch out the ways in which quantum mechanics affects statistical physics, concentrating on those domains of quantum statistical mechanics in which it might appear that classical physics should be valid. In section 2 I shall review the classical description of statistical equilibrium, and attempt to show how problematic it is to extend this description to quantum mechanics — and hence, *a fortiori*, to CDQM; I shall then give some suggestions as to the correct description of quantum statistical equilibrium.

In section 3 I shall focus more specifically on CDQM, and analyze — with the help of the Wigner function — in exactly what sense certain quantum systems can be regarded as approximately classical; in passing I shall make some comments on the problem of justifying phase-space measures (in particular, the justification for neglecting regions of Liouville measure zero) in classical statistical mechanics.

A description of CDQM having been obtained, in section 4 I shall apply it to classically chaotic systems (systems describable by classical statistical mechanics generally being agreed to be of this type). I shall show (mostly by reference to existing work on chaos and decoherence) that such systems typically display relevantly different behavior when regarded as "approximately classical" quantum systems from when they are regarded as fully classical, and that as such there are properties of CM, highly relevant to foundational problems in classical statistical mechanics, which are not preserved by the approximate isomorphism between CM and CDQM. This has implications, in particular, for the problem of the direction of time, and I will briefly describe these implications. Section 5 is the conclusion.

[2]It may seem that there is a danger of infinite regress here: why not also allow for quantum field theory in foundational work, and so on? This danger is always present — in foundational work as much as in mainstream physics. All we can do is try to work out the correct level of theory at which to work, and be prepared to change our strategy if it is shown that our theory does not apply. In fact, I think there are good arguments that non-relativistic quantum statistical mechanics is not significantly affected by QFT, but a clear demonstration of a significant effect would indicate the falsity of any such arguments.

The approximation of quantum physics by classical physics is one of the most notoriously complex — and controversial — of the approximation schemes discussed above; it is extremely hard to discuss without taking a stance on the interpretation of quantum mechanics. In this paper I presume, wherever interpretational issues come up, that the wave-function of a closed system at all times evolves unitarily (i. e., in accordance with the Schrödinger equation). This is the viewpoint taken by Everett[(6)] and my treatment is most naturally compatible with decoherence-based Everett-type interpretations such as those proposed by Saunders,[(7,8)] Vaidman,[(9)] Zurek,[(10)] and myself,[(2,11)] but should also apply to other versions of the Everett interpretation such as the many-minds theories proposed by Lockwood[(12)] and Albert and Loewer,[(13)] and to some extent even to other unitary interpretations like the de Broglie-Bohm pilot-wave theory.[(14,15)]

Other theories will be less well described by the approach I take, although some fairly interpretation-independent comments will be made in section 4. It may be the case that other interpretations do in fact have different consequences for statistical mechanics, but if anything this would strengthen the arguments for considering quantum mechanics, since it would show that foundational issues in statistical mechanics are affected by interpretative issues in quantum theory.

2 DESCRIBING STATES AT EQUILIBRIUM

2.1 Equilibrium in classical systems

In classical equilibrium statistical mechanics, it is generally assumed that:

- The possible states of a classical system are given by the points in some phase space $\mathcal{P}$.
- At any given time t, the specific system under consideration has a determinate state given by a specific point in $\mathcal{P}$ — though this point is assumed not to be exactly known.
- At time t, the probability that this determinate state is in a given region of $\mathcal{P}$ is given by some probability distribution over $\mathcal{P}$.
- The time-evolution of the system is deterministic (given by Hamilton's equations) and so knowing the probability distribution at one time tells us what it is at all other times.
- A system is said to be *at equilibrium* when the probability distribution does not vary in time.

Fleshing out this program leads both to technical and to conceptual difficulties. On the technical side, we would like to establish that if some system's energy is required to be in some infinitesimal interval $(E, E+\mathrm{d}E)$, then the only possible equilibrium probability distribution is the *microcanonical* distribution, for which the probability of finding the system to have a state in some region $\mathcal{R}$ of phase space (where all states in $\mathcal{R}$ have allowed energies) is proportional to the Liouville volume[3] of $\mathcal{R}$. This result is guaranteed to hold whenever a system is *ergodic*, that is, when all dynamical trajectories through points of energy E pass arbitrarily close to all other points with the same energy.

Ergodicity (or weaker, but possibly still acceptable, analogues of ergodicity) is a plausible property of classical systems because phase space has no natural concept of distance — there is a mathematical notion of volume (the Liouville volume) which is conserved under all acceptable dynamics, but in general no conserved metric, so initially very close points in phase space can be taken by the dynamics to widely separated points. This means that classical mechanics is compatible with very chaotic dynamics, which in turn might plausibly lead to very wild motions in the phase space, and hence to the sort of mixing which implies ergodicity — although of course this is just a heuristic and not a proof.

On the conceptual side, there are problems in understanding the nature of the probability distribution which we have put on the phase space. It is often interpreted as being a relative-frequency distribution over an infinite ensemble of copies of the system — but what is the connection between

[3] Recall that the Liouville measure on phase space equals the Lebesque measure, if the coordinates used to define the latter are chosen to be canonical.

this ensemble (which obviously cannot physically exist in a finite classical universe) and the single systems which we in fact observe?

2.2 Equilibrium in isolated quantum systems

In discussions of quantum statistical mechanics one often[16–19] finds the following account of equilibrium: if the system is at energy E, and if there are n orthogonal eigenstates $|\phi_1\rangle, \ldots, |\phi_n\rangle$ all of energy E, then the system is determinately in one of these states, with probability $1/n$ for each. It is sometimes even claimed that quantum mechanics simplifies the classical description by replacing a continuum of classical states with finitely many quantum states.

It is easy to see that this description must be inadequate. Any system for which E is a degenerate eigenvalue of the Hamiltonian will have a continuum of eigenstates with energy E, and there is no reason to prefer one basis for the E-eigenspace over another. But if we say that $\{|\phi_i\rangle\}$ is such a basis, then we are claiming that the system is with certainty in state $|\phi_j\rangle$, for some j, and hence that it has no chance at all to be in superpositions such as

$$\frac{1}{\sqrt{2}}(|\phi_1\rangle + |\phi_2\rangle). \tag{3}$$

In any interpretation which treats the quantum wavefunction as real, this is implausible, as the choice of the $|\phi_i\rangle$ basis was made arbitrarily.

More careful treatments (see, e.g. , page 33 of Binney *et al.*[20] or page 99 of Garrod[21]) acknowledge this problem, and instead define the quantum microcanonical distribution in a manner very similar to that used in classical mechanics: a probability distribution is placed upon all states in the energy-E eigenspace, not just on a specific basis of that subspace. Since there is a continuum of such states, this requires a measure to be specified: the measure used is the volume measure on the unit sphere in Hilbert space. This process is superficially closely analogous to the procedure used in classical physics (the Hilbert-sphere measure, like the Liouville measure, is preserved under all allowed dynamics), and can be shown to give the same statistical results as the naive account.

However, it is easy to see that the analogy with classical mechanics can only be superficial. In the classical case a given energy-E state will move wildly around the energy-E subspace of $\mathcal{P}$ (this is what ergodicity is all about), but a quantum state of energy E is an eigenstate of energy, and hence will not change at all with time. Hence not just the canonical distribution, but any distribution at all within the quantum energy-E subspace will be preserved with time.

More generally, finding ergodic behavior in quantum systems — not just in energy eigenspaces — can reasonably be predicted to be much harder than in classical systems. The reason for this is that in quantum mechanics, unlike classical mechanics, there exists a natural distance measure — given by the Hilbert-space norm — which is preserved by the quantum dynamics. This means that states which begin close together (as defined by this measure) will forever after remain as close together. This property of quantum systems, though not ruling out a quantum analogue of ergodicity, clearly makes it less plausible as an assumption about a given system.

The conceptual questions regarding the probability distribution over states are also different in the quantum case. In quantum statistical mechanics we do not in fact work with probability measures over states, but with probability density operators — that is, with self-adjoint, positive, trace-one operators on the Hilbert space of the quantum system. This is of course possible because our ability to measure the state is much more restricted in quantum than in classical mechanics: in the latter case we can in principle measure the exact phase-space point describing the system, whereas in the former case it is a consequence of Gleason's theorem[22] that any probability distribution of observation outcomes on the system can be described by a single density operator.

This leads to an underdetermination of the probability distribution in quantum mechanics. The space of possible probability measures over an n-dimensional eigenspace is of course infinite-dimensional, whereas the space of density operators describing the same subspace is only (n^2-1)-(real)-dimensional. Hence there are a vast number of possible probability distributions on pure quantum states which give the same density operator — in other words, which are empirically equivalent.

We briefly recall why this is so: given a probability distribution $p(i)$ over some (not necessarily orthogonal) states $\{|i\rangle\}$, the density operator is

$$\rho = \sum_i p(i)\, |i\rangle \langle i| \,, \tag{4}$$

but this map is not 1 : 1 and in general a single density operator can be represented in vastly many ways as a probability distribution over non-orthogonal states. (Generically there is only one way to represent it as a distribution over orthogonal states, but we have already argued that restricting the probability distribution to a specific orthogonal family cannot be justified.)

2.3 Equilibrium in realistic quantum systems

The previous section showed that a statistical mechanics of pure quantum states would be disanalogous in a number of ways from the statistical mechanics of classical systems. However, even the concept of putting a probability distribution over the pure states is dubious, since a pure state is not the most general state for a quantum system.

To recall why this is, let the Hilbert space of the system under question be $\mathcal{H}_\mathcal{S}$. Unless our system is the entire Universe, there must also be some Hilbert space $\mathcal{H}_\mathcal{E}$ of environmental degrees of freedom, such that the Hilbert space of the Universe is[4]

$$\mathcal{H}_\mathcal{S} \otimes \mathcal{H}_\mathcal{E}. \tag{5}$$

Then even if the combined state of system-plus-environment is a pure state, it will not generically be a product state, but rather a highly entangled joint state of $\mathcal{H}_\mathcal{S} \otimes \mathcal{H}_\mathcal{E}$. In such a situation, our system can be assigned a state of its own (provided we avoid carrying out any joint operations on system and environment together) but this state will be a mixed state — i. e. a density operator over $\mathcal{H}_\mathcal{S}$ and not a pure state in $\mathcal{H}_\mathcal{S}$.

Furthermore, even if we do prepare the system in a pure state there is no reason actually to expect it to remain pure. Although the interaction between system and environment is by assumption weak, it has been repeatedly shown[23, 24] that even very weak connections lead to very rapid decoherence (i. e. entanglement with the environment). To get some insight into why this is, suppose that the system is initially described by a pure state made up of fairly localized particles. The uncertainty principle causes the wave-packets of the particles to spread out, and of course we may regard a spread-out packet as a superposition of differently-located, localized packets. These packets will exert a superposition of differently directed forces on the environment, causing it to evolve differently for different packets. In dynamical terms this difference may be tiny, leading to environmental states which differ by only nanometers. But all that is important for decoherence is that different states of the system are coupled with orthogonal states of the environment, and two non-overlapping wave-packets are just as orthogonal when separated by manometers as by parsecs.

All this means that if we wish to place a probability distribution over states of a quantum system it had better be a distribution over all states — pure and mixed — of the system, not just over the pure states. It is then rather unclear what that distribution should be, for while there is a measure on the space of density operators (generated by the inner product $\langle \widehat{A}, \widehat{B} \rangle = \mathsf{Tr}(\widehat{A}^\dagger \widehat{B})$ on the vector space of all linear operators) which is preserved under unitary transformations, the decoherence process which sends pure states to mixed ones is not unitary.

At this point one might object that the problem has been changed by sleight-of-hand. After all, in classical statistical mechanics it is generally assumed that a system is isolated from external influences, so should we not make the same assumption of isolation in quantum mechanics?

[4]This is something of an over-simplification, actually: if we really are considering the state of the Universe then we can expect it to break into wildly different branches, with the subsystem structure varying from branch to branch; hence, the tensor-product decomposition above is to be understood as applying only to the branch or set of branches in which the system actually exists.

There are three problems with this view. Firstly, even if the influence of an external environment is entirely blocked, a quantum system can still be entangled with that environment — so a mixed state is a perfectly valid state for that system, and prima facie should be considered when putting a probability distribution onto such states. (There is no classical analogue for this: even when a classical system is interacting with an environment, its state is still represented by a single point in phase space.)

Secondly, it is overwhelmingly difficult to isolate a quantum system even approximately from its environment — vastly more so than for a classical system. To understand why this is so, recall that for a classical system interference from the outside is generally in the form of energy transfer between system and environment, which can usually be neglected by making the characteristic energies within the system very large compared with the energy transfer in system-environment interactions. In quantum mechanics, by contrast, interference primarily takes the form of *information* transfer, and this cannot be eliminated in the same fashion. For instance, in considering the motion of the Earth around the sun, energy loss through damping from the microwave background radiation is utterly negligible — but if the Earth were in a coherent superposition of two macroscopically separated states, a single microwave photon would suffice to decohere them.

The third reason is that to recover classical mechanics from unitary quantum mechanics, it is generally agreed that we need decoherence to prevent the existence of unobserved coherent superpositions of macroscopically distinct states. As such, at least whilst we are considering domains of quantum mechanics which are in some sense "approximately" classical (and in particular, those which are macroscopic) we will be forced to include some decoherence. This point will be developed further in section 4.

So we return to the problem of finding a probability distribution over the pure and mixed states of the system — that is, over its density operators. Now, these density operators are conceptually speaking very different entities from the 'probability' density operators which were introduced in the previous section. Those were used to describe a probability distribution over states, to be used when the actual state was unknown. The density operators describing states of a system entangled with its environment, on the other hand, encode no classical ignorance at all, but rather reflect the impossibility of describing an entangled state in terms of its constituent systems alone. We might call these latter density operators "entanglement" density operators (they are also sometimes referred to as 'improper mixtures').

So applying the methodology of classical statistical mechanics would lead us to place a probability distribution over the $(n^2 - 1)$-dimensional space of such entanglement density operators. But here the underdetermination of the previous section returns with a vengeance, for any such distribution can itself be described by a *single* probability density operator. The map is of form

$$p(\rho) \longrightarrow \int \mathcal{D}\rho \; p(\rho)\, \rho \tag{6}$$

(where $p(\rho)$ is the given probability distribution over entanglement density operators ρ) and is obviously many-to-one.

So although the empirical success of quantum statistical mechanics leads us to expect this probability density operator to be proportional to the projection operator onto the energy-E subspace, it tells us nothing about the physical interpretation of this operator — at one extreme it might be a probability distribution over unknown pure states, at the other extreme the system might be known to be described by an entanglement density operator equal to the probability density operator, so that there is no classical ignorance or classical probability at all.

2.4 Failure of the classical description

In view of the time-independence of energy eigenstates, of the subsequent difficulty of using any ergodic results to establish uniqueness of equilibrium distributions, of the inappropriateness of describing a macroscopic quantum system as totally isolated from its environment, and of the great extent to which empirically accessible data underdetermines any probability distribution placed on quantum states, we conclude that there is no reason to think that the classical definitions and derivations of equilibrium apply to quantum-mechanical systems. It follows that the features of classical mechanics which allow such definitions and derivations cannot be well-approximated in any approximately classical domain of quantum theory. If it is accepted (as was argued for in section 1) that classical mechanics tells us about the world only insofar as it is approximated by quantum mechanics, it follows that the classical derivation of equilibrium cannot be used to justify the nature of equilibrium states in the world.

This leaves the canonical distribution in a rather precarious position: virtually all attempts to derive it begin with classical physics and attempt to extend the results to quantum physics by analogy, but the arguments of this section imply that

- this extension to quantum mechanics doesn't really work (in other words, the arguments used in the derivation don't work, even if the answer is correct);
- as such, the 'classical' account cannot be a valid account of the microcanonical distribution even in 'classical' regimes of quantum physics. If it can be made to work for genuinely classical systems (which is controversial, of course) then this must be an artifact of classical mechanics and not something which applies to the actual world.

This seems to leave an urgent need for a genuinely quantum account of equilibrium. Though fulfilling this need in any detail lies rather beyond this paper, some tentative ideas will be advanced in section 2.5.

2.5 The quantum interpretation of statistical probability

I would like to propose the following conjecture, which I shall refer to as the quantum interpretation of statistical probability (or QISP):

> 'Ignorance' probability, in the sense of a probability distribution over a space of many possible states of a system, one of which is actual, has no place in statistical mechanics. As such, the *probability* density operator should be banished from statistical mechanics. When a density operator is used to describe a statistical system, it is to be understood as the determinate — though highly non-pure — *entanglement* density operator which describes that specific system.

My reasons for recommending this conjecture are as follows. The first three reasons are basically conceptual; the other three are more dynamical, and probably more important.

1. It would solve — albeit by fiat — the problem of underdetermination of the probability distribution by the statistical facts. If no ignorance probability distribution is being introduced into the theory at all, then there is no reason to be embarrassed by the problems which would arise if such a distribution *were* to be introduced.
2. It would make the concept of 'ensemble' rather less problematic. In one sense, regarding the density operator as describing a single system removes the ensemble concept entirely; in another (particularly to a reader sympathetic to 'many-worlds' language) it makes the ensemble a physical entity rather than a theoretical abstraction.[5] This has consequences for the way we use the ensemble concept in statistical physics. In classical statistical mechanics we are used to a number of properties and descriptions applying not to the individual system (which is in some specific unknown microstate) but to an ensemble, a supposed infinite collection of such systems

[5] Of course, it is important not to confuse the ontological statement that an ensemble is physical, with the epistemological one that we have access to all of it: we are part of the environment of the system and as such will ourselves be entangled with it (I am grateful to an anonymous referee for this point). Also, if we are to regard the density operator as representing an ensemble of Everett worlds, in a sense we return to the underdetermination problem, for we may make a decomposition into worlds in many ways. Of course this is nothing more than the preferred-basis problem; see[(2)] for a discussion of why we need not let it get in the way of seeing a quantum state as genuinely describing a multiplicity of worlds.

whose microstates are distributed over those compatible with the macrostate according to some probability distribution. In particular, concepts like entropy are defined with respect to the ensemble, and the change with time of a macrostate is usually taken to apply to an ensemble rather than an individual system (to bypass the Poincaré recurrence theorem, amongst other reasons). A general discussion is given by Sklar.[(25)]

In quantum mechanics, if QISP holds then it makes sense to describe a single system as being in a macrostate (i. e. described by an entanglement density operator), and we should be able to assign macrostate properties such as entropy to that single system. This may make it more coherent to describe a unique system as having a certain probability distribution.

This redescription of single systems has relevance for the reduction of thermodynamics to statistical mechanics. In thermodynamics quantities such as entropy are properties of single systems, whereas in classical statistical mechanics they are often taken to be properties of ensembles, or of regions of phase space, rather than of individual systems. It appears that this does not hold in quantum statistical mechanics: entropy may be defined for a single system as in the thermodynamic case.

3. If QISP holds, then the (highly problematic[(25)]) probabilities of statistical mechanics are to a large extent removed from consideration, to be replaced with the probability intrinsic to quantum mechanics. Admittedly, the extent to which this is a good thing will depend upon how sanguine the reader is about our prospects for understanding quantum probability. In a pilot-wave theory the probabilities are actually introduced in virtually the same way as for classical statistical mechanics, so not much is gained; the topic of probability in Everett-type interpretations is much more controversial (for some recent proposals on how it is to be understood, see Deutsch,[(26)] Saunders,[(8)] Vaidman[(9)] and Zurek[(10)]). But at any rate the problems with probability in statistical mechanics and in quantum mechanics are prima facie two hard problems urgently in need of solution; there is perhaps something to be said for any strategy which replaces two problems with one.

4. QISP allows us to construct a 'transcendental' account of equilibrium — that is, a justification of the equilibrium state independent of any causal story as to how systems get into equilibrium in the first place — for quantum mechanics which is in some way analogous to the usual accounts in classical mechanics. Recall that in the latter, the microcanonical distribution is to be justified as the only probability distribution which is time-independent; heuristically we expect this to be so because the only obvious conserved quantity for the system is energy and we expect the dynamics of large systems to be unstable enough to explore all regions not expressly forbidden by conservation laws. But all the evidence from studies[6] of decoherence suggests that (in the absence of dissipation) the only density operators which are invariant under decoherence are projections (and sums of projectors) onto eigenspaces of the conserved quantities. For a system which has only energy as a conserved quantity, this is equivalent to saying that the only invariant density operators are microcanonical operators and their sums.

5. Although the classical arguments alluded to in (3) above seem plausible, it has generally been recognized in classical statistical mechanics that more technical work is needed to establish them rigorously — this is where ergodicity comes in. Much progress has been made on these technical questions (reviewed by Sklar[(25)]) and despite the arguments advanced above against the classical derivations of equilibrium, it is hard to deny that there is something very suggestive

[6]We can identify at least three separate research programs in decoherence which support these claims. The use of the predictability sieve[(27, 28)] as a method of picking out minimum-decoherence pure states tends to show that, although there exist states (usually coherent states) which are maximally resistant to decoherence, nonetheless all the pure states eventually decohere. (Admittedly this is only partial support, for it does not prove that there are no highly mixed states other than the microcanonical ones which don't decohere.) The algebraic approaches[(29, 30)] to studying decoherence in small (multi-qubit) systems suggest that only symmetries provide a method of stabilizing systems against decoherence. And the topic of quantum analogues of chaotic systems, which is perhaps the most suggestive for our purposes, will be discussed *in extenso* in section 4.1, and shown to give strong support to the claims in (4).

about some of the ergodic results — and thus something unsatisfactory in simply writing them off as an artifact of a superseded theory. It is then a virtue of the QISP approach that it rehabilitates this technical work, maintaining its connection with equilibrium even whilst to some extent abandoning the conceptual arguments usually used to establish the connection. This point cannot be developed further until we have discussed quantum mechanics on phase space; we will return to it in section 4.2.

6. The last point is perhaps the most telling: if we accept the plausibility of (4) then the microcanonical density operator (interpreted as an entanglement density operator) is the only state of the system (at given energy) which is a valid equilibrium state — all other states evolve to that state, so any probability distribution over any other states will not be an equilibrium distribution at all. Put another way: it may well be that QISP holds automatically at equilibrium, because the dynamics of the system force it upon us.

To illustrate these points — and in particular the last one — we now leave quantum mechanics in its full generality, and focus on those regimes of the theory in which it may be said to approximate classical physics.

3 CLASSICAL-DOMAIN QUANTUM MECHANICS

3.1 Quantum mechanics on phase space

Let us take stock. In section 1 it was argued that (despite its strict falsehood) we are able to use classical mechanics as a theory in some circumstances because quantum mechanics has regimes which are approximately isomorphic to regimes of classical physics, and that we should only be interested in the foundational problems of classical statistical mechanics *per se* insofar as they are taken by this approximate isomorphism into genuine foundational problems of quantum statistical mechanics.

Section 2 gave what might be called a 'non-constructive' proof that many such classical foundational problems are not in fact taken over to quantum mechanics. The argument can be summarized thus:

1. quantum mechanics is conceptually and structurally incompatible with the methods used to derive classical statistical mechanics;
2. hence those methods do not apply to quantum mechanics;
3. *a fortiori* they cannot apply to the classical domains of quantum mechanics;
4. hence problems with classical statistical mechanics which are formulated in terms of these classical methods cannot apply to classical-domain quantum mechanics.

In the remainder of the paper, we shall be concerned with the more 'constructive' question of why the approximate isomorphism between classical mechanics and classical-domain quantum mechanics fails to map the classical derivations of statistical mechanics into quantum mechanics. In doing so, we will shed some light on the QISP hypothesis of section 2.5, and restore to some extent the importance of the concept of classical ergodicity.

To do this, we need to spell out just what 'classical-domain quantum mechanics' actually looks like as a theory (more accurately, as a subtheory of quantum mechanics). The requirements on CDQM are fairly obvious: it should be a theory of quantum states approximately localized in both position and momentum, and those states should approximately speaking evolve in accordance with the classical equations of motion — or more precisely, in accordance with *some* classical equations of motion.

Of course, finding states which are *exactly* localized in both position and momentum is an impossible goal: position and momentum do not commute. But equally, there are some states which should obviously be described as effectively localized within a certain region both of position and momentum space: a Gaussian wave-packet is an obvious example. There is of course an extensive literature on approximate joint measurements of position and momentum, using the positive-operator-valued-measure (POVM) formalism (see Busch, Grabowski and Lahti[(31)] for a recent review). We shall not need the details of this formalism here, for it is enough to know that there exists a well-defined notion of approximate measurement of phase space position, and that we can use it to determine which states are approximately localized. Rather as we would expect, the POVM formalism predicts that it makes sense to describe states localized in phase-space volumes large in comparison to $\hbar^n$ (where $2n$ is the dimensionality of phase space) but that we cannot find a good concept of phase-space localization within significantly smaller volumes.

The phase-space dynamics of states is usually studied by the following method: we look for some representation of all the quantum states (not just the approximately localized ones) in some way which makes their associated phase-space probability distribution transparent; we then look at how that distribution evolves under the quantum dynamics. In general it is possible to represent both pure and mixed states in this fashion. We can then insist on two criteria which the system must satisfy in order to count as "approximately classical". Firstly, those pure states which are approximately phase-space-localized must remain approximately localized on relevant timescales: localized states must not be delocalized by the dynamics. (We are not restricting our attention to closed quantum systems, so we will permit localized pure states to be taken to entanglement density operators made up from localized pure states). Secondly, we require the states to obey approximately classical dynamics, which is to say that the localization centers of localized states are required to evolve in accordance with Hamilton's equations for some classical Hamiltonian (which should presumably be related in a relatively straightforward way to the quantum Hamiltonian).

3.2 The Wigner function

There is a considerable tradition of constructing phase-space representations of quantum states, going back to work of Wigner.[32] Wigner's own solution to the problem is still the most widely used and will suffice for our purposes: to every quantum state (pure or mixed) ρ, we will assign a real function $W_\rho(q,p)$ on phase space. This function (the *Wigner function*) has some of the formal properties of a classical probability distribution. In particular, it yields exactly the right 'marginals', i. e. the probability distributions over position and momentum separately:

$$\int \mathrm{d}p \, W_\rho(q,p) = \langle q| \, \rho \, |q\rangle \tag{7}$$

and similarly for momentum. However, it cannot be regarded even formally as a probability, as it is sometimes negative!

We always knew, though, that there was no prospect of quantum states yielding an exact phase-space probability distribution, because (a) the uncertainty principle tells us that no quantum state is perfectly localized in phase space, and (b) it is easy to exhibit pure states which are not remotely localized. As explained in the previous section, our goal is the much more modest one of providing a probability distribution over the outcomes of *approximate* joint measurements of q and p, and it can be shown[31] that if we "blur" the Wigner function over regions of volume $\sim \hbar^n$, we obtain precisely such a probability distribution for the state.

The Wigner function is far from the only function of this kind: many others (Husimi functions, P-distributions, Q-distributions, etc.) have been proposed at various times, and none of them avoid some pathologies when examined on scales small compared with $\hbar^n$; this is an inevitable (and provable) consequence of the uncertainty principle. It is tempting, but mistaken, to ask which function is the 'correct' representation on phase space. They are all representations of the same quantum state, and when observed at appropriate scales they all give the correct probability distribution with respect to approximate phase-space measurements. I have chosen to work with the Wigner function in this paper because it is the most commonly used in investigations of quantum phase-space dynamics and quantum chaos; this is in turn because of the relative simplicity of its definition and equations of motion.[7]

[7]The equation motion for the Wigner function is given by the Moyal bracket,[33] a generalization of the Poisson bracket. The Moyal bracket of two functions equals their Poisson bracket plus some higher terms suppressed by powers of $\hbar$.

With this technical machinery in hand, we can now state cleanly some necessary conditions for classical results to be treated as relevant to the actual quantum world. Results about the evolution of some classical probability measure on phase space apply to classical-domain quantum mechanics, and hence have some chance of applying to the actual world, only if:

1. There exists a Wigner function whose phase-space probability distribution approximately matches the classical distribution, and which is the Wigner function either of an approximately localized pure state, or of a mixed state made up of such pure states.
2. The quantum dynamics of the system cause the Wigner function's evolution in phase space approximately to match the evolution of the classical distribution.
3. These quantum dynamics also do not cause localized states to become delocalized, at least on the timescales relevant to the problem.
4. The classical results in which we are interested have not been lost as a consequence of all the 'approximately' caveats of the first three points!

It should be stressed that these are necessary but not sufficient conditions. In particular, all of the above could hold for some system and yet it might still be necessary to regard the Wigner function as being that of an entanglement density operator, and thus as having nothing to do with classical probability — unlike the case in the classical distribution being approximated.

We are now in a position to consider the dynamics of 'approximately classical' systems and compare them to genuinely classical ones. This will be the task of section 4, but first we digress onto another foundational problem of classical mechanics which does not apply to classical-domain quantum mechanics: the 'problem of measure zero'.

3.3 Justifying phase space measures

When doing classical statistical mechanics, to calculate probabilities it is necessary to have both a function on phase space and a measure on it, and it is customary to use the Liouville measure. There are various motivations for the naturalness of this choice: it is the only measure which has the same functional form in any canonical coordinates; it is the only measure preserved under all allowed (classical) phase-space evolutions; the equations of motion for probability distributions take on a particularly elegant form when Liouville measure is used, etc. Furthermore, in general our choice of measure can quite legitimately be made on grounds of convenience: given a probability distribution we can represent it with any measure[8] (provided we are willing to use singular functions like delta functions).

Unfortunately, there are some situations where the connection between the Liouville measure and probability must be made *a priori*: specifically, many classical dynamical results which we would like to hold for all phase-space points actually hold for all points *except for a set of Liouville measure zero*. We would like to exclude these points by saying that a given system has zero probability to be found at one of these points, but this requires us to justify why this should be the case — after all, there are plenty of measures which assign nonzero measure to these points. (This 'set of measure zero problem' is discussed, in the classical context, by Sklar.[(25)])

We shall see in section 4 that modifications to classical dynamics made by quantum mechanics mean that we should be quite cautious about the validity of classical dynamical arguments applied to ergodic systems, but there is a more fundamental reason why the set of measure zero problem may be alleviated by quantum mechanics — namely, the sets of measure zero do not appear to have any quantum analogue. In the Wigner function formalism, *every* state occupies a phase-space region of non-zero Liouville measure; individual points on the phase-space do not correspond to any physical states. Hence any dynamical result applying to all regions of nonzero measure would apply to all states.

[8] Given a function f, a measure μ, and a region $\mathcal{R}$ of phase space, the probability of finding the system in $\mathcal{R}$ is $\int_{\mathcal{R}} f \mathrm{d}\mu$; clearly, then, modifications of μ can be compensated for by modifications of f.

Of course, as mentioned earlier the Wigner function is not the only way to do quantum mechanics on phase space, and there are phase-space representations in which states are contained in regions of zero measure.[9] However, even in these cases only a special subclass of states have zero measure; the generic state has nonzero measure. Hence given any zero-measure state we can always find one arbitrarily close (as measured by the Hilbert-space norm) of non-zero measure, and since quantum dynamics are linear a sufficiently close quantum state will approximate the dynamics of the state in question to any desired degree of accuracy.

Another way of seeing why the measure-zero problem disappears in quantum mechanics is as follows: the problem assumes that arbitrarily small regions of phase space can legitimately be discussed, whereas all of the phase-space representations of quantum states only describe phase-space locality when we avoid probing them on scales smaller than $\hbar^n$. This is not a limitation of these quantum descriptions: it is not that we are giving up on describing smaller phase-space volumes, but rather that we are recognizing that quantum mechanics — and thus the actual world — has no concept of phase-space localization on such small scales. Thus the approximate isomorphism between classical mechanics proper and the classical domains of quantum mechanics will not preserve the concept of a measure-zero subset — so, by the fourth criterion of section 3.2, we should reject the measure-zero problem as a relevant foundational problem for statistical mechanics as it applies to the actual world.

4 QUANTUM DYNAMICS OF CLASSICALLY CHAOTIC SYSTEMS

4.1 Quantum chaos

In this section, it will be argued that 'classical' statistical-mechanical systems have their dynamical behavior profoundly modified by quantum theory. This will form the basis of my 'constructive' argument about the relevance of quantum effects to classical statistical mechanics, as well as providing the promised (in section 2.4 illustration of QISP and the arguments in favour of it.

At first sight it may seem trivially wrong that classical and quantum systems will have such profoundly different dynamics. After all, we *know* that classical mechanics is an effective limiting case of quantum mechanics, and furthermore can directly derive correspondence results: for instance, the Wigner function obeys dynamical laws equal to the classical dynamics together with some apparently negligible higher terms.(35)

The reason why quantum effects are after all important is linked to the nature of the dynamics of (classical) statistical-mechanical systems. To have any chance of showing the dynamical mixing properties which such systems need to approach equilibrium, their dynamics need to be highly chaotic: that is, trajectories initially close to one another must rapidly diverge, so that an initially localized probability distribution becomes spread across the whole phase space. In fact, one can characterize a chaotic system by the exponentially rapid divergence of its trajectories.

However, the concept of trajectory is foreign to quantum mechanics, as is the idea of arbitrarily close initial states ending up widely separated: after all, the Schrödinger equation is unitary. To see the implications for quantum states we need to consider the evolution of the whole classical probability distribution, and compare that to the evolution of the Wigner function or another phase-space representation of the quantum state.

The evolution of the classical distribution can be understood from two conflicting considerations: the trajectories are diverging exponentially so that initially close points are increasingly found very far away from one another in the phase space; but the overall volume of the region where the distribution is non-zero must be conserved, by Liouville's theorem.[10] This means that the distribution becomes

[9] Glauber's(34) "P-representation", for instance, assigns a delta function to coherent states; I am grateful to an anonymous referee for this reference.

[10] Of course, this was understood by Gibbs(36) more than a century ago, in the context of statistical mechanics; see Sklar(25) for a historical discussion.

highly distorted, extending long thin filaments which grow and develop fine structure at exponential rates. There is no dynamically preferred length scale in classical mechanics, hence nothing to stop this fine structure developing at arbitrarily small scales.

Quantum mechanics, however — and hence its classical limit — does have a preferred length scale. As Peres says:

> Any compact domain, obeying the Liouville equation of motion, is continuously distorted and tends to project increasingly long and thin filaments. As time passes, new, finer filaments emerge, whose volume is less than $\hbar^N$. The quantum density ρ (or the Wigner function ...) cannot reproduce these minute details and smoothes them away. We therefore expect the quantum dynamical evolution to be qualitatively *milder* than the classical one. ((35), p. 303)

How can we reconcile this with the observed accuracy of classical mechanics? The details are both complicated and somewhat controversial, but appear to be as follows (I follow Berry and co-workers[(37,38)] and Zurek and Paz;[(39,40)] however see Casati and Chirikov[(41)] for a criticism of Zurek and Paz's approach).

1. Suppose we begin with a quantum system in a pure state, and describe it by a Wigner function on phase space. To begin with (i. e. while the classical state has no fine structure on scales smaller than $\hbar^n$) the Wigner function accurately approximates the classical evolution.

2. After a time of approximately
$$\tau_c \ln(I/\hbar), \tag{8}$$
where τ_c is the timescale for exponential growth of classical filaments and I is the system's action,[(37)] fine structure has developed on scales too small for the Wigner function to follow; on a similar timescale,[(39)] the Wigner function will have become highly delocalized in position space in at least some directions. Note that (due to the logarithm in (8), which is in turn due to the exponential growth rate of fine structure) this timescale will not be comfortably large, even for macroscopic systems; it will instead be on roughly the same timescale at which the classical state spreads across phase space, and hence on the timescale of approach to equilibrium. (In fact, Zurek[(39)] calculates the timescale for Hyperion, Saturn's classically chaotic moon, to become delocalized across its entire orbit if allowed to evolve in isolation, and gets a figure of about twenty years. Zurek's paper, describing the quantum mechanics of planetary motion(!) is a salutary lesson in how careful we should be before deciding that quantum effects are negligible.)

3. If the system were genuinely isolated (and given our assumption that even macroscopic isolated systems have unitary dynamics) then the system will genuinely cease to behave classically. However, in any realistic case the presence of macroscopically delocalized states will by now have led to environment-induced decoherence, effectively collapsing the wave-function into a basis approximately localized in both position and momentum and changing it from a pure to a mixed state. The effect of this on the Wigner function is to prevent narrow filaments[(39,40)] from getting *too* narrow, allowing them to continue to track the classical evolution at the cost of no longer preserving the phase-space volume of the state. Since phase-space volume is related to entropy (via $S = k \ln(\text{volume})$, this increase in phase-space volume corresponds to an increase in the fine-grained entropy of the system.[11]

4. Of course, the fact that outside interference can increase the entropy of a system is also true in classical statistical mechanics. The difference here is that it seems reasonable to consider classically chaotic systems in isolation, whereas if we isolate chaotic quantum systems they develop, in short order, the pathology of macroscopically delocalized states. (Also phenomena which have a totally negligible effect on classical dynamics — the effect of the microwave background

[11] It should be acknowledged that we are assuming here that phase-space entropy is a good measure of the *quantum* entropy $-\mathsf{Tr}\rho \ln \rho$ of the mixed state; this seems plausible, and is motivated by Zurek,[(39)] but probably does not hold exactly.

radiation on the orbit of Jupiter, for instance — lead to significant levels of quantum decoherence. As was mentioned in section 2.3, this is because classical isolation is about limiting energy transfer, whilst quantum isolation is about the vastly more difficult task of limiting information transfer.)

This brief sketch cannot do justice to the technical work it attempts to describe, but the basic consequences can be fairly simply described — and moreover seem to a large extent to be interpretation-independent: a chaotic quantum system which is isolated and evolves unitarily will in fairly short order develop macroscopically delocalized states and stop obeying the classical equations of motion. To prevent this pathology we must invoke collapse of the wave-function (using whatever description of collapse our particular interpretation gives us; in the Everett interpretation it is environment-induced decoherence). The collapse continually converts the (initially pure) quantum state to a mixed state. This mixed state will spread out across the phase space at the same rate as the classical probability distribution, but will not develop the arbitrarily small-scale structure of the classical state. Instead, its filaments will have a minimum width, and as such its effective phase-space volume will increase exponentially — which implies that its entropy will increase linearly.

4.2 Foundational implications of quantum chaos

With an account of quantum chaos available, we are now able to make a number of foundational points, supporting and developing the arguments of section 2 (and in particular 2.5, in which the QISP was advocated).

1. The implications of these results for isolated systems are startling. Imagine taking a classically chaotic system (even a highly macroscopic one such as the Solar System), preparing it in an initially well-localized state, isolating it completely (*per impossibile*) from the rest of the Universe, and letting it evolve on timescales which are of the same order as the system's classical dynamical timescale. If the wave-function evolves unitarily, then the predictions made by classical physics for this system will not be correct, or even approximately correct: they will be totally wrong.

   Here we have possibly the most powerful reason for including decoherence (and thus, interaction with other systems) when considering statistical systems, as was advocated in section 2.3: there is just no such thing as a totally isolated, approximately classical system. Such systems are artifacts of classical mechanics proper: they do not exist in classical-domain quantum mechanics, and hence do not exist in the actual world.

2. Classically chaotic systems give us an example of the modification of CM which might be necessary to preserve its approximate isomorphism with CDQM, as mentioned in section 1. In chaotic systems CDQM is more closely isomorphic not to CM but to what we could call *irreversible classical mechanics*, or ICM. In ICM, there is negligible energy transfer into and out of the system, and when coarse-grained on scales large with respect to $\hbar^n$, the evolution of phase-space distributions is the same in ICM as in CM — but the fine-grained structure, produced on exponentially small scales in CM, is washed out in ICM, and phase-space volume, conserved in CM, increases exponentially in ICM.

   In terms of the criteria laid out at the end of section 3.2 for when classical results are to be treated as relevant, the situation is as follows: we are able to construct a Wigner function, and a quantum dynamics, which satisfies the first three criteria (that is, remains comprised of localized states, and approximately tracks the classical dynamics), but in the process, time-reversibility and conservation of Liouville volume are lost; thus, any classical result which depends upon time-reversibility and volume conservation fails the fourth criterion (that the relevant effects are not, in fact, lost due to the approximations made).

3. Note that quantum mechanics has converted classical unpredictability into quantum indeterminacy. Classically, in a chaotic system if we specify a state to any finite accuracy we will sooner or later lose our ability to predict where the state is. In quantum mechanics we can reliably

predict what the future state will be — but our prediction is that it will be a mixed state spread across the whole of the allowable region of phase space, and that it is in principle impossible (no matter how accurately we knew the initial state) to predict the result of a position or momentum measurement.

4. We are now in a position to illustrate point (6) in section 2.5: viz, that it may be dynamically impossible to avoid imposing QISP, i. e. replacing probability density operators with entanglement density operators. We shall show that QISP must hold for chaotic systems at equilibrium, as follows. Suppose we begin with a density operator describing a uniform distribution across phase space. If we interpret this operator as a *probability* density operator over rather well-localized pure states, then the actual state of the system will be relatively localized in phase space. Our ignorance about the results of position or momentum measurements will be largely epistemic: we don't know the state, but if we did then we could predict reasonably well the results of measurements.

   However, this pure state will not stay pure for long. The Wigner function of the state will begin to develop fine structure, which in conjunction with environmental entanglement will force a transition from a pure to a mixed state. On a timescale of order a few multiples of the timescale τ_c, our pure state will have been converted to a mixed state entirely delocalized in (the accessible region of) phase space: in other words, to a mixed state formally identical to the original probability density operator.

   As a consequence of this, *all* of the original unknown states will evolve into the *same* entanglement density operator. There is no longer any ignorance-type probability, so we can discard the probability distribution and just work with the entanglement density operator. However, since this is mathematically identical to the original probability density operator there is no practical change. The only change is in the nature of the probabilities, which now must be interpreted in the QISP sense.

5. We are also now able to fulfil the promissory note of (5) in section 2.5, and to rehabilitate the importance of the ergodicity concept. For the quantum-mechanical results of this section rely upon the fact that the classical evolution explores all of (a given energy surface in) phase space — in other words, upon classical ergodicity. If significant regions of phase space remained unexplored by the classical dynamics — as is implied, for instance, by the KAM theorem (of Kolmogorov, Arnold and Moser[42]) — then it seems less clear that the Wigner function would spread across all of phase space even given decoherence.

4.3 The arrow of time

The asymmetry between past and future remains one of the great puzzles of physics, and is not confined to statistical mechanics: in addition to the increase of entropy from past to future, 'arrows of time' are also defined by phenomena as disparate as the expansion of the universe and our own perception of temporal flow. It remains an open question as to what extent these different temporal asymmetries can be related, and (if they can be related) which should be viewed as fundamental.

Quantum mechanics seems to add another 'arrow': the collapse of the wave-function appears to be an explicitly time-asymmetric process. Furthermore, it has long been recognized (since Von Neumann;[43] see also[44]) that the pure-to-mixed-state transition associated with wave-function collapse is a process which increases entropy; this hints at some sort of connection between the entropic and quantum-mechanical arrows of time.

If the account of quantum chaos presented in section 4.1 is correct, the two are in fact linked in a fairly straightforward manner: chaotic systems will become macroscopically delocalized unless their wavefunction is collapsed, but adding this collapse into such systems causes an increase in entropy — hence the arrow of time defined by the collapse process gives the direction of entropic increase. (See[39] for further, and more technical, discussion.)

Of course, different interpretations of quantum mechanics give different explanations for the time asymmetry of wavefunction collapse. Theories of dynamical collapse (of which the most well-developed are due to Ghirardi, Rimini and Weber[45] and Pearle[46]) are time-asymmetric by *fiat*, of course, and presumably could cause entropic increase in even an isolated system (Albert[47] has developed ideas along these lines).

The Everett interpretation, on the other hand, incorporates no explicit time-asymmetry,[12] and so we have a problem analogous to that of statistical mechanics: given time-symmetric dynamics how do we explain observed asymmetries? Specifically, we need to explain why the branching structure of histories is indeed branching. When we phrase this question in terms of subsystems it becomes the question of why subsystems of the universe are less entangled in the past than in the future. Explanations of this in terms of a very special initial state of the universe begin to look indistinguishable from their counterparts in the foundations of statistical mechanics.

The moral seems to be that the quantum-mechanical and entropic arrows of time are closely linked. Depending on our interpretation we may solve one by *fiat* and thus solve the other; or we may recognize both as aspects of the same problem. But it seems unwise to attempt to solve them separately.

5 CONCLUSIONS

If the arguments of this paper are valid, then some rather basic concepts in classical statistical mechanics are significantly altered once quantum effects are taken into account. In terms of the conceptual structure of the theory, we have found that the existence of superposition and entanglement and the dual interpretations of the density operator cast serious doubt upon the validity of some classical derivations of equilibrium, and call into question the standard view of statistical mechanics' phase-space distributions as probability distributions over unknown but determinate microstates. On the dynamical side, the nature of quantum chaos requires us to replace CM with an irreversible cousin, drastically undermines the validity of treating systems as totally isolated, and implies a close link between the statistical-mechanical and quantum ('wave-packet-collapse') arrows of time.

These need not be seen as *negative* consequences, of course. Quantum considerations indicate ways to make progress on issues such as:

- the nature of ensembles;
- the reduction of thermodynamics to statistical mechanics;
- the concept of probability in statistical mechanics;
- the problem of measure zero;
- the connection between ergodicity and equilibrium.

Furthermore, statistical mechanics may offer a potential testing ground for rival interpretations of quantum mechanics. If the way in which we formulate statistical mechanics is dependent on our interpretation of quantum theory then it is *a priori* possible that two interpretations, identical in their predictions for any given experiment, will produce different statistical mechanics. If so we have a potential test between them: failure to reproduce the canonical ensemble is just as fatal for an interpretation as failure to reproduce the two-slit experiment.

But in any case, whether or not moving to quantum mechanics helps solve foundational problems in statistical physics, it is not an optional move. It is certainly no goal of this paper to undermine the superb and sustained work that has been done on understanding the foundations of classical statistical mechanics, nor to claim that this work is rendered useless by quantum considerations. But to see how it applies to statistical-mechanical systems in the actual world, it is necessary to see how — and to what extent — this classical work can be transferred across to quantum systems. If the approximate isomorphism between classical mechanics and classical-domain quantum mechanics was such as to transfer all our understanding of classical statistical mechanics unproblematically across to at least certain domains of quantum theory, there would have been no need to worry about quantum

[12] Neither does the de Broglie Bohm interpretation, of course.

effects. As the isomorphism is not in fact like this, we cannot justify the avoidance of quantum considerations.

References

[1] Wojciech H. Zurek and Juan Pablo Paz. Zurek and Paz reply [to[41]]. *Physical Review Letters*, 75(2):351, 1995.

[2] David Wallace. Worlds in the Everett interpretation. To appear in *Studies in the history and philosophy of modern physics*; available at http://www.arxiv.org/abs/quant-ph/0103092, 2001.

[3] John Worrall. Structural realism: the best of both worlds. *Dialectica*, 43:99–124, 1989. Reprinted in *The Philosophy of Science*, D. Papineau (Ed.), Oxford University Press, 1996, pp. 139–165.

[4] Stathis Psillos. Is structural realism the best of both worlds? *Dialectica*, 49:15–46, 1995.

[5] James Ladyman. What is structural realism. *Studies in the History and Philosophy of Science*, 29:409–424, 1998.

[6] H. Everett III. Relative state formulation of quantum mechanics. *Review of Modern Physics*, 29:454–462, 1957. Reprinted in *The many-worlds interpretation of quantum mechanics*, De Witt, B. and N. Graham, eds. (Princeton: Princeton University Press 1975).

[7] Simon Saunders. Time, decoherence and quantum mechanics. *Synthese*, 102(2):235–266, 1995.

[8] Simon Saunders. Time, quantum mechanics, and probability. *Synthese*, 114:373–404, 1998.

[9] Lev Vaidman. The many-worlds interpretation of quantum theory. *Stanford Encyclopedia of Philosophy*, available at http://www.tau.ac.il/∼vaidman/mwi/mwst1.html, 2000.

[10] Wojciech Zurek. Decoherence, einselection and the existential interpretation (the rough guide). *Philosophical Transactions of the Royal Society of London*, A356:1793–1820, 1998.

[11] David Wallace. Everett and structure. Forthcoming, 2001.

[12] Michael Lockwood. *Mind, Brain and the Quantum: the compound 'I'*. Blackwell Publishers, Oxford, 1989.

[13] David Albert and Barry Loewer. Interpreting the many worlds interpretation. *Synthese*, 77:195–213, 1988.

[14] D. Bohm. A suggested interpretation of quantum theory in terms of "hidden" variables. *Physical Review*, 85:166–193, 1952.

[15] Peter Holland. *The Quantum Theory of Motion*. Cambridge University Press, 1993.

[16] R. Baierlein. *Atoms and Information Theory*. Freeman, 1971.

[17] Walter Grandy, Jr. *Foundations of Statistical Mechanics, Vol. I, Equilibrium Theory*. Kluwer Academic Publishers, 1987.

[18] Richard Feynman. *Statistical Mechanics*. Addison Wesley, 1972.

[19] C. Kittel. *Thermal Physics*. Wiley, 1969.

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

[21] Claude Garrod. *Statistical Mechanics and Thermodynamics*. Oxford University Press, Oxford, 1995.

[22] A. M. Gleason. Measures on the closed subspaces of a Hilbert space. *Journal of Mathematics and Mechanics*, 6:885–893, 1957. Reprinted in.[48]

[23] E. Joos and H. D. Zeh. The emergence of classical properties through interaction with the environment. *Zeitschrift für Physik*, 59:223–243, 1985.

[24] Wojciech H. Zurek. Decoherence and the transition from quantum to classical. *Physics Today*, 44(10):36–44, 1991.

[25] Lawrence Sklar. *Physics and Chance: Philosophical issues in the foundations of statistical mechanics.* Cambridge University Press, 1993.

[26] David Deutsch. Quantum theory of probability and decisions. *Proceedings of the Royal Society of London*, A455:3129–3137, 1999. Available at http://www.arxiv.org/abs/quant-ph/9906015.

[27] W. H. Zurek. *Progress in Theoretical Physics*, 89:281–302, 1993.

[28] W. H. Zurek, S. Habib, and J.P Paz. Coherent states via decoherence. *Physical Review Letters*, 70(9):1187–1190, 1993.

[29] Paolo Zanardi. Stabilizing quantum information. *Physical Review A*, 63:12301, 2001.

[30] Emanuel Knill, Raymond Laflamme, and Lorenza Viola. Theory of quantum error correction for general noise. *Physical Review Letters*, 84(11):2525–2528, 2000.

[31] P. Busch, M. Grabowski, and P. Lahti. *Operational Quantum Physics.* Springer-Verlag, 1995.

[32] E. Wigner. On the quantum correction for thermodynamic equilibrium. *Physical Review*, 40:749, 1932.

[33] J. E. Moyal. Quantum mechanics as a statistical theory. *Proceedings of the Cambridge Philosophical Society*, 49:99–124, 1945.

[34] R. Glauber. Quantum theory of coherence. In S.M. Kay and A. Maitland, editors, *Quantum Optics: proceedings of the tenth session of the Scottish Universities' Summer School in Physics, 1969.* Academic Press, London, 1970.

[35] Asher Peres. *Quantum Theory: Concepts and Methods.* Kluwer Academic Publishers, 1993.

[36] J. Gibbs. *Elementary principles in statistical mechanics.* Dover, 1960.

[37] M. V. Berry and N. L. Balazs. Evolution of semiclassical quantum states in phase space. *Journal of Physics A*, 12(5):625–642, 1978.

[38] H. J. Korsch and M. V. Berry. Evolution of Wigner's phase space density under a nonintegrable quantum map. *Physica D*, 3:627–636, 1981.

[39] Wojciech H. Zurek. Decoherence, chaos, quantum-classical correspondence, and the algorithmic arrow of time. *Phys. Scripta*, T76:186–198, 1998. LANL preprint number quant-ph/9802054.

[40] Wojciech H. Zurek and Juan Pablo Paz. Decoherence, chaos and the second law. *Physical Review Letters*, 72(16):2508–2511, 1994.

[41] Giulio Casati and B. V. Chirikov. Comment on "Decoherence, chaos and the second law". *Physical Review Letters*, 75(2):350, 1995.

[42] V. Arnold and A Avez. *Ergodic problems of classical mechanics.* Benjamin, New York, 1968.

[43] John von Neumann. *Mathematical Foundations of Quantum Mechanics.* Princeton University Press, 1955.

[44] D. H. Zeh. *The Physical Basis of the Direction of Time.* Springer, Berlin, 1989.

[45] G. Ghirardi, A. Rimini, and T. Weber. Unified dynamics for micro and macro systems. *Physical Review D*, 34:470, 1986.

[46] P. Pearle. Combining stochastic dynamical state-vector reduction with spontaneous localization. *Physical Review A*, 39(5):2277–2289, 1989.

[47] David Albert. *Time and Chance.* Harvard University Press, 2000.

[48] C. A. Hooker, editor. *The Logico-Algebraic Approach to Quantum Mechanics*, volume 1. Reidel, Dordrecht, 1975.