

Recurrence Theorems: a unified account

David Wallace*

Balliol College, University of Oxford

(Dated: June 17, 2013)

I discuss classical and quantum recurrence theorems in a unified manner, treating both as generalisations of the fact that a system with a finite state space only has so many places to go. Along the way I prove versions of the recurrence theorem applicable to dynamics on linear and metric spaces, and make some comments about applications of the classical recurrence theorem in the foundations of statistical mechanics.

I. INTRODUCTION

The Poincaré recurrence theorem plays an important role in the foundations of statistical mechanics, dating back to Zermelo's original objection to Boltzmann's H theorem.[11] The theorem exists in both classical and quantum forms, but only the classical version is widely discussed. This is unfortunate, as there are foundationally important differences between the two. Furthermore, in proofs of the classical theorem measure-theoretic technicalities can obscure the very elementary idea underpinning the theorem, and standard proofs of the quantum version obscure the links between classical and quantum versions.

In this paper I provide a unified treatment of recurrence, beginning with the toy example of a system with discretely many states and based around the intuition that if a system only has finitely many places to go it will eventually enter one such place twice. In the discrete case this is literally how the theorem works; I generalise it to dynamics on spaces equipped with notions of conserved volume, of conserved inner product, and (in the Appendix) of conserved length, showing that in each case the conserved quantity extends the basic idea from discrete to continuous state spaces. I illustrate how recurrence theorems for classical and quantum mechanics, respectively, are instances of the first two extensions. I stress the distinction between recurrence (every point recurs after some time interval) and *uniform* recurrence (there is some time interval after which every point recurs). On the basis of these results, I observe that two properties of recurrence that have been frequently discussed in the literature are just artifacts of classical mechanics.

II. WARM-UP: RECURRENCE IN FINITE SYSTEMS

At a very abstract level, we can define a *dynamical system* as specified by:

- A set \mathcal{S} of *states*

- An evolution rule U , which maps \mathcal{S} to itself.

The idea is that U maps each state to the state into which it will evolve after 1 time-step. We can then define a time evolution operator $U(n)$ as follows: $U(1) = U$, $U(2) = U \cdot U$, $U(3) = U \cdot U \cdot U$, etc. (or, more formally: $U(1) = U$, $U(n+1) = U \cdot U(n)$), satisfying $U(n+m) = U(n)U(m)$. Although in this model of dynamics time is discrete, this is no real limitation: given a continuous-time evolution operator $V(t)$, we can pick some arbitrarily short time τ and define $U = V(\tau)$; then $U(n) =_{df} V(n\tau)$.

The idea of recurrence is that every state, as it evolves forward in time, in some sense eventually returns to its original state. If the state space \mathcal{S} has finitely many points, we can interpret this entirely literally. Suppose we say that:

- A dynamical system is exactly recurrent if for every $s \in \mathcal{S}$, there is some $n_s > 0$ such that $U(n_s)s = s$.
- A dynamical system is invertible if U is a one-to-one map: that is, if for every $s \in \mathcal{S}$ there is a unique $t \in \mathcal{S}$ such that $U(t) = s$. (From this it follows that $U(n)$ can be extended uniquely to negative n while preserving $U(n+m) = U(n)U(m)$.)
- A dynamical system is finite if its state space \mathcal{S} contains finitely many points.

Then we can easily prove the

Theorem 1 (*Finite Recurrence Theorem*) *Any finite invertible system is exactly recurrent.*

Proof Let \mathcal{S} contain N members. Then for any given s , not all of

$$s, U(1)s, U(2)s, \dots, U(N)s$$

can be distinct, so there must be distinct n, m with $n < m$ such that $U(n)s = U(m)s$. Then $s = U(-n)U(n)s = U(-n)U(m)s = U(m-n)s$. \square

Informally: since there are only finitely many states, the evolution of s must eventually pass through the same state twice; because the dynamics is given by a one-to-one map, the *first* state to be entered twice must be the original state, so that no two states are mapped to the same state.

*Electronic address: david.wallace@balliol.ox.ac.uk

Note that both requirements are needed. If the system is not finite, it can continue exploring new states forever; if it is not invertible, it can get stuck at some state or set of states and never get out again.

Now consider the space $\mathcal{U}(\mathcal{S})$ of *all* maps from \mathcal{S} to itself. Given some evolution operator U , we can define a function L_U from $\mathcal{U}(\mathcal{S})$ to itself by

$$L_U(V) = U \cdot V. \quad (1)$$

Formally, we can regard $\mathcal{U}(\mathcal{S})$ as a state space, and L_U as an evolution operator on that state space, so as to create a new dynamical system. If \mathcal{S} is finite, so is $\mathcal{U}(\mathcal{S})$ (there are N^N maps from an N -member set to itself); if U is invertible, so is L_U ($L_{U^{-1}} \cdot L_U(V) = U^{-1} \cdot U \cdot V = V$). So in fact, this dynamical system is also recurrent. In particular, if 1 is the identity map, then there exists some n such that

$$(L_U)^n(1) \equiv U(n) \cdot 1 = U(n) = 1. \quad (2)$$

But if $U(n) = 1$, then $U(n)s = s$ for all $s \in \mathcal{S}$: that is, there is some *fixed* time n such that after n time-steps, every state has been time-evolved back to itself. Let us call this stronger property *uniform recurrence*: formally, a dynamical system is *uniformly exactly recurrent* if for some fixed $n > 0$, *all* states $s \in \mathcal{S}$ satisfy $U(n)s = s$.

We have now proved the

Theorem 2 (Finite Uniform Recurrence Theorem) *Any finite invertible system is uniformly exactly recurrent.*

In essence, all recurrence results we will consider have this form: recurrence occurs just because the system only has so many places to go, and runs out of them, and because its one-to-one dynamics means that this can only happen if it returns to its original state. The problem is that in most cases of physical interest, the state space has infinitely many points in it, so simple arguments based on state-counting will not work. We need instead to find some precise sense in which the system has only finitely many *interestingly different* places to go, so that it returns to some state *virtually the same as* its initial state. This will be our task for the rest of the paper.

III. RECURRENCE IN VOLUME

Suppose that the state space \mathcal{S} of a given dynamical system is equipped with a *measure*, a rule that defines a notion of volume on the subsets of \mathcal{S} . (Informally, a measure μ is a rule that associates to a subset $S \subset \mathcal{S}$ a value $\mu(S)$ which is a nonnegative real number or ∞ , such that the volume of a union of disjoint sets is the sum of the volumes of the individual sets and the volume of the empty set is zero; for a more formal definition, see [6]. Not all subsets can be assigned a measure (there are pathological ‘non-measurable’ sets) but any remotely simple set is measurable.) We define a *measurable dynamical system*

as a dynamical system equipped with a measure in this way. We can now use this notion to define an appropriate generalisation of recurrence.

- A measurable dynamical system is *recurrent in volume* iff for any measurable $S \subset \mathcal{S}$, the set $T \subset S$ of points in S that do not eventually re-enter S under dynamical evolution has measure zero.

To put this more informally: pick any state, and choose any non-zero-volume region containing that state. If the system is recurrent in volume, virtually all points in that region will eventually return to it — “virtually” in the sense that the set of those which do not has measure zero. By taking the region arbitrarily small, we can ensure that (in some sense) almost all states arbitrarily close to our original state will return arbitrarily close to it.

To prove recurrence in volume, we need assumptions analogous to finiteness and invertibility. They are:

- A measurable dynamical system is *finite in measure* iff $\mu(\mathcal{S}) < \infty$.
- A measurable dynamical system is *volume-preserving* iff $\mu(U(S)) = \mu(S)$ for all measurable S .

We can then prove the

Theorem 3 (Volume Recurrence Theorem) *Any measurable dynamical system that is finite and volume-preserving is recurrent in volume.*

Proof Pick a measurable subset $S \subset \mathcal{S}$, and let $T \subset S$ be the set of points in S that never return to S , i.e. the points $s \in \mathcal{S}$ which satisfy $U(n)s \notin S$ for all $n > 0$. Now define $T_n = U(n)T$. Since the dynamics is measure-preserving, for any n, m (with $n < m$) we have $\mu(T_n \cap T_m) = \mu(T \cap T_{m-n}) = 0$, i.e. the overlap of any two T_n has measure zero. So if we define T'_m as the set of points in T_m but not in any of the T_n with $n < m$, then $\mu(T'_n) = \mu(T_n)$.

The T'_n are mutually disjoint, so if $\mathcal{T} = \cup_n T'_n$ (and so also $\mathcal{T} = \cup_n T_n$), $\mu(\mathcal{T}) = \mu(T'_1) + \mu(T'_2) + \dots$.

But since the dynamics is volume-preserving, $\mu(T) = \mu(T_i) = \mu(T'_i)$ for all i . There are two ways of satisfying this requirement:

- $\mu(\mathcal{T}) = \infty$.
- $\mu(T) = 0$.

But (i) is ruled out by the assumption that the system is finite in measure, so (ii) is the only possibility. \square

How can we define a volume-based version of uniform recurrence? It’s not entirely clear, but it’s moot in any case, because no such result is true. Consider, for instance, the following dynamical system: the state space is a cylinder of radius 1 and height 1, so that states are labelled by pairs (z, θ) with $z \in [0, 1]$ and $\theta \in [0, 2\pi)$. If the dynamical rule is $(z, \theta) \rightarrow (z, [\theta + z] \text{ modulo } 2\pi)$

(i.e., if it rotates each point an amount proportional to z) then the ordinary volume of this cylinder provides a preserved measure, so that the system is a finite, invertible, measure-preserving dynamical system. The volume recurrence theorem holds, and indeed this is easy to see informally: for any z , and any ϵ , we can find n so that nz is within ϵ of being an integer multiple of 2π .

If the system were *uniformly* recurrent, though, we would expect (e.g.) the line of points $(z, 0)$ for arbitrary z to return arbitrarily close to itself. But it's clear that this won't happen: the line will wind more and more tightly around the cylinder, and never unwind.

To get further mathematical insight into why this happens, consider that the space of measure-preserving functions on a (non-finite) space is typically very large — so large, in fact, that there is no natural way to define a volume measure on it. So the analogue of our finite-case strategy, where we considered the space of maps itself as a dynamical system, will not work in this case.

We will shortly see that there are alternative (and physically more relevant) generalisations of recurrence in which uniform recurrence *does* hold. Firstly, though, we should stop and consider the most important application of volume recurrence.

IV. APPLICATION: CLASSICAL MECHANICS

Classical mechanics in the Hamiltonian formalism (for given time-step τ , and assuming finitely many degrees of freedom) is a dynamical system: the state space is phase space, the evolution rule is evolution for time τ under Hamilton's equations, and the volume is given by[12]

$$\mu(V) = \int_V dq^1 dp_1 \cdots dq^n dp_n \quad (3)$$

By Liouville's theorem, this is conserved under Hamiltonian evolution; as such, classical mechanics is an invertible, measure-preserving, measureable dynamical system.

It is not a finite system (phase space, in general, has infinite volume, since momentum can increase without limit). However, Hamiltonian dynamics conserves energy, and it is entirely possible for the subset of phase space corresponding to energies between E and δE (for arbitrarily small δE) to have finite volume.[13] This gives us a form of the

Theorem 4 (*Poincaré recurrence theorem*) *For any classical-mechanical system with finite-dimensional phase space, if the subset of phase space with energies lying in some band $(E, E + \delta E)$ has finite Liouville volume then the restriction of the dynamics to that subset is recurrent in volume.*

As an important example of this, consider a system whose Hamiltonian has the form

$$H(q^1, p_1, \dots, q^n, p_n) = \sum_{i,j} C^{ij} p_i p_j + V(q^1, \dots, q^n), \quad (4)$$

where V is bounded below, C^{ij} is positive definite, and the positions are confined to a finite region. Then the momenta will also be confined to a finite region for given energy, and the Poincaré recurrence theorem will hold. (For a case where the positions are not confined, consider a particle moving freely in empty space (clearly this will not recur); for a case where V is not bounded below, consider Newtonian gravity.)

V. RECURRENCE IN LINEAR DYNAMICAL THEORIES

We say that a dynamical theory is *linear* if its state space is some subspace of an inner product space (whose inner product we write $\langle \cdot, \cdot \rangle$). Provided the space is of some finite dimension N , we can use this inner product to define a volume on the space, informally[14] just by defining the volume of a cube of side λ as λ^N . Also given the inner product, we can define the distance $d(x, y)$ between two states x, y as the length of the vector between them: that is, $d(x, y) = \sqrt{\langle x - y, x - y \rangle}$. This function is sometimes called the *metric*.

We now define a linear dynamical theory as

- *bounded* if for some fixed D , $d(s, t) < D$ for all states s, t .
- *inner-product-preserving* if for all states s, t , $\langle U(s), U(t) \rangle = \langle s, t \rangle$.

We then have

Theorem 5 (*Linear recurrence theorem, preliminary form*) *Any linear dynamical theory which is finite-dimensional, bounded, invertible and inner-product-preserving is recurrent in volume with respect to the volume function defined by the inner product.*

Proof If the system is inner-product-preserving, its dynamics preserves the volume measure. If the system is bounded, (say, with $d(s, t) < D$ for some D and all states s, t), then the state space is contained within a cube of side D and hence volume D^N , and so has finite volume. The result is now a corollary of the volume recurrence theorem. \square

Since quantum-mechanical systems are linear systems (taking the state space to be Hilbert-space states of norm 1), this result applies in particular to finite-dimensional quantum systems, and indeed can be extended to infinite-dimensional ones under certain conditions (as we will see). But this route to quantum recurrence, although it has the virtue of hewing closely to the classical model, understates the strength of recurrence that actually occurs in quantum theory — essentially because preservation of an inner product is a much stronger condition than preservation of a volume.

Indeed, the mere presence of an inner product allows us to define an alternative concept of recurrence: we can

consider a system to be recurrent if any state returns arbitrarily close to its starting point with respect to this metric. More precisely,

- A linear dynamical system is *recurrent in metric* iff for any state s , and any $\epsilon > 0$, there is some n_s such that $d(U(n_s)s, s) < \epsilon$.
- A linear dynamical system is *uniformly recurrent in metric* iff for any $\epsilon > 0$, there is some n such that for all states s , $d(U(n)s, s) < \epsilon$.

We can then prove a stronger result:

Theorem 6 (Linear recurrence theorem) *Any linear dynamical theory which is finite-dimensional, bounded, and inner-product-preserving is uniformly recurrent in metric.*

Proof We begin by proving ordinary recurrence. Suppose that the dynamical theory has a state space \mathcal{S} that is a subset of an N -dimensional vector space \mathcal{V} , and pick some orthonormal set of vectors v_1, \dots, v_n . For any N -tuple of real[15] numbers x^1, \dots, x^n , and any $\epsilon > 0$, we can construct the ball

$$B_\epsilon(x^1, \dots, x^n) = \{x \in \mathcal{V} : d(x, \sum_i x^i v_i) < \epsilon/2\}. \quad (5)$$

If \mathcal{B} is the set of such balls defined by N -tuples $n_1\epsilon/2\sqrt{N}, \dots, n_N\epsilon/2\sqrt{N}$ for arbitrary (positive or negative) integers n_1, \dots, n_N , it is easy to verify that

- every point in \mathcal{V} is in some element of \mathcal{B} ;
- if \mathcal{S} is bounded, there is some *finite* subset B_1, \dots, B_M such that every point in \mathcal{V} is in one of the B_i .

Now pick an arbitrary state s , and consider its dynamical evolutes

$$U(1)s, U(2)s, \dots$$

Each one must lie in (at least) one of the B_i (we can pick one arbitrarily if it lies in more than one) and so, since there are only finitely many of the B_i , we must be able to find n, m ($n < m$) such that $U(n)s$ and $U(m)s$ both lie within the same B_i .

But if so, $d(U(n)s, U(m)s) < \epsilon$ (B_i has radius $\epsilon/2$, so diameter ϵ) and so, since the dynamics preserves d , $d(s, U(m-n)s) < \epsilon$: that is, the system is recurrent in metric.

We prove *uniform* recurrence by the same trick used in the finite case: we construe the space of evolution operators as a dynamical system in its own right. Specifically: consider the space $\mathcal{U}(\mathcal{V})$ of linear transformations of \mathcal{V} . This too is a vector space, and it can be equipped with an inner product as follows:

$$\langle U, V \rangle = \text{Tr}(U^\dagger V). \quad (6)$$

Any norm-preserving V satisfies $\langle V, V \rangle = \text{Tr}(V^\dagger V) = \text{Tr}(1) = N$, so that the set of all such U is a bounded subset of $\mathcal{U}(\mathcal{V})$. If we define a dynamics on that set by

$$L_U(V) = UV \quad (7)$$

then we can readily verify that it is norm-preserving. Applying our existing result tells us that this dynamics is recurrent in metric, and so for any ϵ' , there is some n such that $\epsilon' > \text{Tr}((U(n) - 1, U(n) - 1)$.

Now, for any operator A and any vector v , $\text{Tr}(A^\dagger A)\langle v, v \rangle \geq \langle Av, Av \rangle$. And because the original state space is bounded, there is some D such that $\langle s, s \rangle < D$ for any state s . So if we take $\epsilon' = \epsilon^2/D$, we have

$$\begin{aligned} \epsilon^2 > \text{Tr}(U(n)-1, U(n)-1)D > \text{Tr}(U(n)-1, U(n)-1)\langle s, s \rangle \\ &\geq \langle U(n)s - s, U(n)s - s \rangle = d(U(n)s, s)^2. \end{aligned} \quad (8)$$

□

In fact, the result rests only on the existence of a preserved distance function, and on the fact that the space can be covered, for any ϵ , with finitely many balls of radius ϵ . (This latter property is called *total boundedness* in the theory of metric spaces.) Since the details are mildly fiddly, I postpone them to the appendix.

VI. APPLICATION: QUANTUM MECHANICS

In the quantum theory of pure states, quantum states are taken to be vectors of norm 1; in the theory of mixed states, they are taken to be self-adjoint operators of trace 1 all of whose eigenvalues are nonnegative. In both cases the states form a bounded subset of a linear space. So whenever the quantum theory's Hilbert space is finite-dimensional, the dynamics are uniformly recurrent.

In many cases of physical interest, though, we take the Hilbert space to be infinite-dimensional. In some cases the recurrence theorem can be extended to these situations also. Specifically, we can prove the

Theorem 7 (Quantum recurrence theorem) *Given a quantum system with Hamiltonian \hat{H} , if the set of eigenvalues of \hat{H} (a) is bounded below, (b) is discrete, and (c) has only finitely many elements within any bounded interval, then under the evolution operator $\hat{U} = \exp(-i\hat{H}\tau)$ (for arbitrary $\tau > 0$),*

- The dynamics are recurrent in metric;
- For any E , the restriction of the dynamics to states of expected energy $< E$ is uniformly recurrent in metric.

Proof To begin with, let the energy eigenstates be given by $|\varphi_i\rangle$, with $\hat{H}|\varphi_i\rangle = E_i|\varphi_i\rangle$, and with the energies in increasing order, so that if $i > j$, $E_i \geq E_j$. (That they

can be so listed is a consequence of (a)-(c.) Write \mathcal{H}_N for the Hilbert space spanned by states $\{|\varphi_1\rangle, \dots, |\varphi_N\rangle\}$; this space has dimension N . \widehat{U} leaves \mathcal{H}_N invariant and so defines a dynamics on it; since \mathcal{H}_N is finite-dimensional, this dynamics is uniformly recurrent.

Next, note that any given state can be expressed as a sum of energy eigenstates: $|\psi\rangle = \sum_i^\infty \alpha_i |\varphi_i\rangle$, where $\sum_i^\infty |\alpha_i|^2 = 1$. We can usefully define the partial sums $|\psi_N\rangle = \sum_i^N \alpha_i |\varphi_i\rangle$ (that is, as the projection of $|\psi\rangle$ onto \mathcal{H}_N , so that $\langle\psi|\psi_N\rangle = \sum_i^N |\alpha_i|^2$). Using the linear space definition of distance, we can readily calculate

$$d(|\psi\rangle, |\psi_N\rangle) = \sqrt{2 \left(1 - \sum_{i=0}^N |\alpha_i|^2\right)} = \sqrt{2 \sum_{i=N+1}^\infty |\alpha_i|^2}.$$

To prove (i), pick $\epsilon > 0$. There will exist some N such that $d(|\psi\rangle, |\psi_N\rangle) < \epsilon/3$. Since the dynamics on \mathcal{H}_N are recurrent, there will also exist m such that $d(\widehat{U}(m)|\psi_N\rangle, |\psi_N\rangle) < \epsilon/3$. By the triangle inequality,

$$\begin{aligned} d(\widehat{U}(m)|\psi\rangle, |\psi\rangle) &< d(\widehat{U}(m)|\psi\rangle, \widehat{U}(m)|\psi_N\rangle) \\ &+ d(\widehat{U}(m)|\psi_N\rangle, |\psi_N\rangle) + d(|\psi_N\rangle, |\psi\rangle). \end{aligned} \quad (9)$$

Since \widehat{U} preserves Hilbert-space distance, the first and third terms are equal, and each is less than $\epsilon/3$; the result now follows.

To prove (ii), assume that all the energy eigenvalues are nonnegative (if this is not the case, we can just add a constant to \widehat{H} without affecting the dynamics) and that there are eigenvalues of arbitrarily high energy (if not, then by (b) and (c), the system is finite-dimensional and so already known to be uniformly recurrent). Now pick $E > 0$. For any $\lambda > 0$, there will exist some energy eigenvalue E_N with $E_N > E/\lambda$. For any given state $|\psi\rangle$ with expected energy $< E$, we can decompose it as $|\psi\rangle = |\psi_N\rangle + |\delta\psi\rangle$, and we have

$$E \geq \langle\psi|\widehat{H}|\psi\rangle \quad (10)$$

$$= \langle\psi_N|\widehat{H}|\psi_N\rangle + \langle\delta\psi|\widehat{H}|\delta\psi\rangle \quad (11)$$

$$\geq \langle\delta\psi|\widehat{H}|\delta\psi\rangle \quad (12)$$

$$= \sum_{i=N+1}^\infty E_i |\alpha_i|^2 \quad (13)$$

$$\geq E_N \sum_{i=N+1}^\infty |\alpha_i|^2 \quad (14)$$

$$> (E/\lambda) \sum_{i=N+1}^\infty |\alpha_i|^2 \quad (15)$$

$$(16)$$

Rearranging,

$$2\lambda > d(|\psi\rangle, |\psi_N\rangle)^2. \quad (17)$$

Since λ was chosen arbitrarily, we can in particular choose N so that

$$d(|\psi\rangle, |\psi_N\rangle) < \epsilon/3. \quad (18)$$

(In other words, all states with expected energy below E are within distance $\epsilon/3$ of \mathcal{H}_N). Since the dynamics are uniformly recurrent on \mathcal{H}_N , we can find m such that for all states $|\psi\rangle$, $d(\widehat{U}(m)|\psi_N\rangle, |\psi_N\rangle) < \epsilon/3$. Repeating the argument in the proof of (i) tells us that $d(\widehat{U}(m)|\psi\rangle, |\psi\rangle) < \epsilon$. \square

When will this somewhat abstract condition be realised? On heuristic grounds we would expect it to hold in nonrelativistic QM whenever the system is spatially bounded and has a potential bounded below: in this situation, a solution to the Schrödinger equation can only have discretely many nodes, and any increase in nodes increases the average curvature of the solution and so the kinetic energy. In fact this can be formally proved in functional analysis (see [3, p.1330]). More generally, [1] has advanced arguments to the effect that even in quantum field theory, there are only finitely many energy eigenstates of any spatially bounded theory below a given total energy.

If the theory is not spatially bounded, the spectrum may be continuous, and there is no guarantee of recurrence, just as in the classical case (the free particle is a trivial counterexample).

VII. ANOTHER ROUTE TO LINEAR RECURRENCE

The above argument is not in fact how the quantum recurrence theorem was originally proved. In this section I sketch the standard proof, which has the advantage of making more transparent the way in which recurrence actually plays out, at the cost of blurring the continuity with the classical and discrete cases.

Firstly, suppose that the system's Hilbert space is finite-dimensional, and that the Hamiltonian has eigenvalues E_1, \dots, E_N and eigenstates $|\varphi_1\rangle, \dots, |\varphi_N\rangle$. Then any state $|\psi\rangle$ may be expanded in energy eigenstates, and we have

$$\widehat{U}(t)|\psi\rangle = \exp(-i\widehat{H}t)|\psi\rangle = \sum_{i=1}^N \alpha_i e^{-iE_i t} |\varphi_i\rangle. \quad (19)$$

Now suppose that all of the energy eigenvalues E_1, \dots, E_N are rational numbers. Then for some time t , $E_i t = 2n_i \pi$ for some integers n_i . It follows that $\widehat{U}(t)|\psi\rangle = |\psi\rangle$: that is, the theory is uniformly recurrent.

But any real number is arbitrarily well approximated by a rational number! So this suggests that this result ought to hold (to an arbitrarily good approximation) for arbitrary E_1, \dots, E_N . This leads us to the theory of *almost periodic functions*, and outside the scope of this article. But an elementary sketch proof (using our existing

recurrence results!) would go as follows: for given τ , the effect of $U(\tau)$ is to increase the phase of the i th eigenstate by $E_i\tau$. These phases take values between 0 and 2π ; that is, the set of all phases lives on an N -dimensional torus. If we define a distance or volume function on that torus in the usual way, we can see that these phase increases preserve both. Applying the volume or metric versions of the recurrence theorem then tells us that to any degree of approximation, and for some N , applying the phase rotation N times gets us back where we started to that degree of approximation. (This argument essentially follows [7].)

VIII. TWO ARTEFACTS OF CLASSICAL MECHANICS

In the literature on the foundations of statistical mechanics, two subtleties of the classical (Poincaré) recurrence theorem are often stressed:[16]

1. Although the theorem establishes (roughly) that the states that do not recur have collective measure zero, nothing rules out the existence of anomalous states that never recur.
2. The theorem establishes recurrence, but not *uniform* recurrence.

The first point (often called the ‘problem of measure zero’), and its cousin in the theory of ergodicity, poses a supposed problem to any attempt to give a dynamical justification to the standard measures used in statistical mechanics: only by making measure-theoretic assumptions can we even get the recurrence theorem to tell us about actual systems. The second point is advanced as a supposed justification for the move from the Boltzmannian description of thermodynamical systems as represented by phase-space *points*, to the Gibbsian description which represents them by phase-space *distributions*. The idea here is that recurrence poses a threat to the monotonicity of entropy increase — but if entropy is a function of distributions rather than individual states, there is no objection to its increasing indefinitely.

I make no comment on the *conceptual* status of these concerns — but we can see that technically speaking, they rely on features of classical dynamics which do not occur in quantum theory. The quantum recurrence theorem demonstrates *uniform* recurrence, and it tells us that *all* states return arbitrarily close to their starting point. So insofar as a given classical system fails to demonstrate these features, all we can learn is that the classical description must at some point fail to be a good approximation to any underlying quantum system. And since our only (non-historical) reason to describe a system via classical dynamics is because it is assumed to be a good approximation to some underlying quantum dynamics, these subtleties can be dismissed as classical artefacts, without conceptual significance.

(To be fair, at least the second point is sometimes noted. [8], in particular, notes (pp.200-1) that the Gibbsian approach does indeed fail in finite-volume quantum systems. But he suggests that the solution is to go to the thermodynamic limit: that is, to infinite-volume (but finite-density) systems. However, there is no reason to expect even *non-uniform* recurrence to apply in that context, and indeed, a move to the infinite-volume case would have saved classical *Boltzmannian* statistical mechanics from any problem of recurrence. So it is hard to see what is really intended here.)

It is interesting to ask just what features of classical mechanics really prevent uniform recurrence. On possibility is the absence of a dynamically conserved distance function (let alone a dynamically preserved linear structure) on phase space. Classical dynamics preserves size but not shape[17] and so allows a wider variety of possible dynamics.

However, the space of probability distributions over phase space does have a linear structure, and the dynamics on that space is likewise linear: the Liouville equation, $\dot{\rho} = \{\rho, H\}$, is a linear differential equation. Furthermore, if we restrict our attention to a finite-volume region of phase space and define an inner product on the space of density operators by

$$\langle \rho, \sigma \rangle = \int dx \rho(x)\sigma(x) \quad (20)$$

(where the integral is taken with respect to Liouville measure) then it is easy to show that this inner product is preserved by the Liouville equation. In this case, the reason that classical dynamics fails to abide by the linear recurrence theorem is that the distribution space is infinite-dimensional, even if phase space has finite volume: distributions can have structure on arbitrarily short scales.

Formally, quantum theory can also be done on the space of phase-space distributions: the Wigner function formalism[18] maps each quantum state to a real function on phase space (with the proviso that such functions are not reliably nonnegative) and the resultant mapping of the Schrödinger equation — the *Moyal equation* — is equal to Liouville’s equation to leading order in \hbar . But Wigner functions of a given energy cannot vary on arbitrarily short lengthscales, essentially because of the uncertainty principle. We can again see the infinite dimensionality of the space of classical distributions as an artefact of classical physics.

IX. CONCLUSION

Systems undergo recurrence, essentially, because there are only so many places they can go, and eventually they have to go back on themselves, and because if they have 1:1 dynamics, the only way they can do this is by returning to their starting point. Generalising this to state

spaces with infinitely many states can be done via a conserved volume, or a conserved inner product, or a conserved metric. The first of these is most commonly seen, and applies to classical mechanics; the others, however, are more powerful, implying *exceptionless, uniform* recurrence, and are applicable to quantum mechanics. The relative weakness of the classical recurrence theorem is therefore a curiosity rather than a point of conceptual significance.

APPENDIX: RECURRENCE IN METRIC SPACES

To begin with some definitions: a *pre-metric space* is a set \mathcal{S} equipped with a function $d(x, y)$ from pairs of elements x, y of the set to the nonnegative real numbers, such that for any $x, y, z \in \mathcal{S}$:

1. $d(x, x) = 0$.
2. $d(x, y) = d(y, x)$;
3. $d(x, y, z) \leq d(x, y) + d(y, z)$ (the triangle inequality).

(A *metric space* satisfies in addition that if $d(x, y) = 0$ then $x = y$, but we will not need this condition.) A metric dynamical theory is then just a dynamical theory whose state space is equipped with a metric.

We now repeat some of the definitions used for linear spaces:

- Given a point x , and $\epsilon > 0$, the *ball of radius ϵ* , $B_\epsilon(x)$ is the set $\{y \in \mathcal{S} : d(x, y) < \epsilon\}$.
- A premetric space is *totally bounded* if for any $\epsilon > 0$, there are finitely many points x_1, \dots, x_N such that

$$B_\epsilon(x_1) \cup \dots \cup B_\epsilon(x_N) = \mathcal{S}.$$

- A map $U : \mathcal{S} \rightarrow \mathcal{S}$ is *metric-preserving* if for any $x, y \in \mathcal{S}$, $d(x, y) = d(U(x), U(y))$.

The proof that any metric-preserving evolution map is recurrent in metric proceeds the same way as for linear spaces:

Theorem 8 (Metric recurrence theorem) *Any metric dynamical theory which is totally bounded and metric-preserving is recurrent in metric.*

Proof Fix $\epsilon > 0$ and $x \in \mathcal{S}$ (where \mathcal{S} is the state space), and let $\{x_1, \dots, x_N\}$ be such that the union of the $B_{\epsilon/2}(x_i)$ is \mathcal{S} . Since there are only finitely many x_i , we can find $n, m (n < m)$ and i such that $U(n)x$ and $U(m)x$ are both in $B_{\epsilon/2}(x_i)$, and so by the triangle inequality, $d(U(n)x, U(m)x) < \epsilon$. Since U is metric-preserving, $d(x, U(m-n)x) < \epsilon$.

To extend this to uniform recurrence, we can use the usual trick of applying recurrence to dynamics on the space of metric-preserving maps. To do so, we need to make this space into a totally bounded metric space.

To do so, first note that a totally bounded space is also bounded: that is, for some $D > 0$, $d(x, y) < D$ for all x and y . (If the balls of radius (say) 1 around $x_1 \dots x_N$ contain all points in the space, then for some $D' > 0$ and all i, j , $d(x_i, x_j) < D'$, and so by the triangle inequality, $d(x, y) < D' + 2$.) If $\mathcal{U}(\mathcal{S}, \mathcal{T})$ is the space of metric-preserving functions from \mathcal{S} to \mathcal{T} , and \mathcal{S} and \mathcal{T} are totally bounded metric spaces, we can then make $\mathcal{U}(\mathcal{S}, \mathcal{T})$ into a metric space by defining

$$d(f, g) = \max_x d(f(x), g(x)).$$

We can prove the following lemma:

Lemma: If \mathcal{S} and \mathcal{T} are totally bounded premetric spaces (with metrics d_S, d_T), so is $\mathcal{U}(\mathcal{S}, \mathcal{T})$ (with the metric d defined above).

Proof Pick $\epsilon > 0$, and pick $x_1 \dots x_N \in \mathcal{S}$, $y_1, \dots, y_M \in \mathcal{T}$ such that

$$B_{\epsilon/4}(x_1) \cup \dots \cup B_{\epsilon/4}(x_N) = \mathcal{S}$$

$$B_{\epsilon/4}(y_1) \cup \dots \cup B_{\epsilon/4}(y_M) = \mathcal{T}.$$

Now let Π be the set of functions from $\{x_1, \dots, x_N\}$ to $\{y_1, \dots, y_M\}$, and for each $\pi \in \Pi$, define $F_\pi \in \mathcal{U}(\mathcal{S}, \mathcal{T})$ to be the set of metric-preserving functions such that $f(x_i) \in B_{\epsilon/4}(\pi(x_i))$ for each i . Suppose f and g are both elements of some F_π , and that $x \in B_{\epsilon/4}(\pi(x_j))$. Then by the triangle inequality,

$$d_T(f(x), g(x)) \leq d_T(f(x), f(x_i)) +$$

$$d_T(f(x_i), \pi(x_i)) + d_T(\pi(x_i), g(x_i)) + d_T(g(x_i), g(x)). \quad (21)$$

Since f and g are metric-preserving, $d_T(f(x), f(x_i)) = d_S(x, x_i) < \epsilon/4$; since $f \in F_\pi$, $d_T(f(x_i), \pi(x_i)) < \epsilon/4$; the same holds for g . So we conclude that $d_T(f(x), g(x)) < \epsilon$.

Now for each non-empty F_π , pick an arbitrary $f_\pi \in F_\pi$. $F_\pi \subset B_\epsilon(f_\pi)$, so it follows that the union of all the $B_\epsilon(f_\pi)$ is $\mathcal{U}(\mathcal{S}, \mathcal{T})$. Since there are only finitely many of these, $\mathcal{U}(\mathcal{S}, \mathcal{T})$ is totally bounded.

In particular, $\mathcal{U}(\mathcal{S}, \mathcal{S})$ is totally bounded. The dynamics defined by $L_U(f) = U \cdot f$ is a metric-preserving evolution on this space, and so by theorem 8, for any ϵ we can find n such that $d(U(n)1, 1) < \epsilon$, and hence $d_S(U(n)x, x) < \epsilon$ for any x . That is, we have proved

Theorem 9 (Uniform metric recurrence theorem) *Any metric dynamical theory which is totally bounded and metric-preserving is uniformly recurrent in metric.*

For readers more familiar with the topological notion of *compactness*, note that any compact metric space is totally bounded[19] (for given $\epsilon > 0$, consider the open cover consisting of balls of radius ϵ around every point in

the space) and so, as a corollary of theorem 8, any metric-preserving dynamical theory with a compact state space is uniformly recurrent in metric.

-
- [1] Bekenstein, J. (1981). Universal upper bound of the entropy-to-energy ratio for bounded systems. *Physical Review D* 23, 287–298.
- [2] Brown, H. R., W. Myrvold, and J. Uffink (2009). Boltzmann’s H-theorem, its discontents, and the birth of statistical mechanics. *Studies in History and Philosophy of Modern Physics* 40, 174–191.
- [3] Dunford, N. and J. Schwartz (1988). *Linear Operators Part II: Spectral Theory*. John Wil.
- [4] Hillery, M., R. F. O’Connell, M. O. Scully, and E. P. Wigner (1984). Distribution functions in physics: Fundamentals. *Physics Reports* 106, 121–167.
- [5] Moyal, J. E. (1949). Quantum mechanics as a statistical theory. *Mathematical Proceedings of the Cambridge Philosophical Society* 45, 99–124.
- [6] Rudin, W. (1987). *Real and Complex Analysis* (3rd ed.). New York: McGraw-Hill.
- [7] Schulman, L. (1978). Note on the quantum recurrence theorem. *Physical Review A* 18, 2379–2380.
- [8] Sklar, L. (1993). *Physics and Chance: Philosophical Issues in the Foundations of Statistical Mechanics*. Cambridge: Cambridge University Press.
- [9] Sutherland, W. (1975). *Introduction to Metric and Topological Spaces*. Oxford Univer.
- [10] Wigner, E. (1932). On the quantum correction for thermodynamic equilibrium. *Physical Review* 40, 749–759.
- [11] See [2], and references therein
- [12] More formally, this is the volume defined by the symplectic structure ω on phase space: $\mu(V) = \int_V \omega \wedge \omega \cdots \wedge \omega$, where $\omega \wedge \omega \cdots \wedge \omega$ is the n -fold exterior product of ω with itself.
- [13] With only slightly more care, we can prove a version of the recurrence theorem that applies if *any given energy hypersurface* is finite in volume. However, as there can be no physically interesting distinction between the two cases, even this slight increase in care seems unnecessary.
- [14] Formally (at least if the state space is an open subset of the inner product space) pick an orthonormal basis; use this basis to identify the state space with a subset of R^N ; use this identification to carry the Lebesgue measure over to the state space.
- [15] We’re not really assuming that the space is a *real* rather than complex space here; an N -dimensional complex vector space is also a $2N$ -dimensional real vector space.
- [16] See, e. g. , [8, pp.182–7,199–202], and references therein.
- [17] That is not to say that it preserves *only* size: by no means are all volume-preserving maps symplectic maps.
- [18] The Wigner function was first introduced in [10] and explored further by [5]; see Hillery *et al* 1984 for a review of its properties.
- [19] The converse is not true, but any *complete* totally bounded metric space is compact. For an elementary proof, see [9, p.141].