

Reality as a Vector in Hilbert Space

Sean M. Carroll¹

*Walter Burke Institute for Theoretical Physics
California Institute of Technology, Pasadena, CA 91125
and Santa Fe Institute, Santa Fe, NM 87501*

Abstract

I defend the extremist position that the fundamental ontology of the world consists of a vector in Hilbert space evolving according to the Schrödinger equation. The laws of physics are determined solely by the energy eigenspectrum of the Hamiltonian. The structure of our observed world, including space and fields living within it, should arise as a higher-level emergent description. I sketch how this might come about, although much work remains to be done.

Invited contribution to the volume *Quantum Mechanics and Fundamentality: Naturalizing Quantum Theory Between Scientific Realism and Ontological Indeterminacy*; Valia Allori (ed.).

¹e-mail: seancarroll@gmail.com

[I]t would be naive in the extreme to be a ‘Hilbert-space-vector realist’: to reify Hilbert space, and take it as analogous to physical space. – Wallace (2017)

In Hilbert space, nobody can hear you scream. – Aharonov & Rohrlich (2005)

Almost a century after the 1927 Solvay Conference, the question of the ultimate ontology of quantum mechanics remains unsettled. Essentially all formulations of quantum theory rely on the use of a wave function or state vector (or mathematically equivalent structures). But researchers do not agree on whether the state vector is a complete and exact representation of reality, whether it represents part of reality but needs to be augmented by additional variables to be complete, or whether it is an epistemic tool rather than a representation of reality at all. And they further do not agree on whether the state vector should be thought of as purely an element of some abstract Hilbert space, or whether there is some fundamental ontological status to a particular representation of that vector in terms of something more directly physical, such as configuration space or particles or fields in honest three-dimensional “space.”

Here I want to argue for the plausibility of an extreme position among these alternatives, that the fundamental ontology of the world is completely and exactly represented by a vector in an abstract Hilbert space, evolving in time according to unitary Schrödinger dynamics. Everything else, from particles and fields to space itself, is rightly thought of as emergent from that austere set of ingredients. This approach has been called “Mad-Dog Everettianism” (Carroll & Singh, 2019) although “Hilbert Space Fundamentalism” would be equally accurate.

Let’s see how one might end up seduced by an ideology that is so radically different from our direct experience of the world. When we are first taught quantum mechanics, we are shown how to construct quantum theories by taking classical models and quantizing them. Imagine we have a classical precursor theory defined on some phase space, expressed mathematically as a symplectic manifold Γ , with evolution determined by some Hamiltonian function $H : \Gamma \rightarrow \mathbb{R}$. We choose a “polarization” on phase space, which amounts to coordinatizing it in terms of canonical coordinates Q (defining “configuration space”) and corresponding canonical momenta P , where each symbol might stand for multiple dimensions. This is a fairly general setup; for N point particles moving in D -dimensional Euclidean space, configuration space is isomorphic to \mathbb{R}^{DN} , but we could also consider field theory, for which the coordinates are simply the values of the fields throughout space.

One way to construct a corresponding quantum theory is to introduce complex-valued wave functions of the coordinates alone, $\Psi(Q) \in \mathbb{C}$. Wave functions must be normalizable, in the sense that they are square-integrable, $\int \Psi^* \Psi dQ < \infty$, where Ψ^* is the complex conjugate of Ψ . Momenta are now represented by linear operators \hat{P} , whose form can be derived from the canonical commutation relations $[\hat{Q}, \hat{P}] = i\hbar$ (where the operator \hat{Q} is simply multiplication by Q). This lets us promote the classical Hamiltonian to a self-adjoint operator $\hat{H}(\hat{Q}, \hat{P})$ (up to potential operator-ordering ambiguities). We then posit that the wave function evolves according

to the Schrödinger equation,

$$\hat{H}\Psi = i\hbar\frac{\partial}{\partial t}\Psi. \tag{1}$$

This form of the Schrödinger equation is perfectly general, and applies to relativistic theories as well as non-relativistic ones, as long as one uses an appropriate Hamiltonian.

This procedure gives us the beginnings of a quantum theory. For Everettians, it gives us the complete theory; a unitarily-evolving quantum state describes the entire ontology at a fundamental level. Other approaches require additional dynamical rules, physical structures, or some combination thereof. Here our interest is in seeing how far we can get from a minimal starting point, so the Everettian approach is appropriate. [For more on structure in the Everett interpretation, see (Saunders, 2021; Wallace, 2003).]

But how precisely should we think about this ontology? It might seem to include at least original configuration space coordinatized by $\{Q\}$, the set of wave functions $\Psi(Q)$, and the Hamiltonian \hat{H} . This has led Albert (1996) to suggest “wave function realism” – the idea that it is the wave function $\Psi(Q)$ that describes reality, and that the wave function is defined on configuration space, so that configuration space ($3N$ dimensional for N particles in 3 dimensions, infinite-dimensional for a field theory) is where the wave function “really lives” [see also (Ney, 2013; North, 2013)]. This point of view has been criticized by Wallace (2017), who points out that configuration space is simply one possible way of *representing* the wave function; it could also be represented in momentum space (as a function of P), or an infinite number of other choices. [Wallace also has other criticisms, and other authors criticize the idea for other reasons (Allori, 2013; Myrvold, 2015).]

The ambiguity between defining wave functions in position versus momentum space is an example of a broader issue. “Quantization” does not give us a one-to-one map from classical theories to quantum ones, or even a well-defined map at all. A single classical precursor may correspond to multiple quantum theories, due to operator-ordering ambiguities. Moreover, distinct classical models may have identical quantizations, as in the dualities of quantum field theory. This makes it difficult to uniquely pinpoint what a quantum theory is supposed to be a theory “of.” In the duality between the massive Thirring model and sine-Gordon theory, a single quantum theory can arise from classical precursors describing either fundamental bosons or fundamental fermions (Coleman, 1975). In the AdS/CFT correspondence, a nongravitational field theory in D spacetime dimensions is dual to a quantum gravity theory in $D + 1$ dimensions (Maldacena, 1999). So from a quantum theory alone, it might be impossible to say what kinds of fields the model describes, or even the number of dimensions they live in.

But the complicated relationship between classical and quantum theories is a problem for physicists, not for physics. Nature simply is quantum from the start, and the classical world arises as an emergent approximation in some appropriate limit. If our interest is in fundamental ontology, rather than focusing on quantum theories derived by quantizing classical precursors, it would make sense to consider the inverse problem: given a quantum theory, what kind of

classical limits might arise within it?

So let us think carefully about what it means to be given a quantum theory (at least from an Everettian perspective). Wave functions, as von Neumann noted long ago, can be added and scaled by complex numbers, and have a natural inner product defined by $(\Psi, \Phi) = \int \Psi^* \Phi$. They therefore describe a complex, normed vector space, called Hilbert space, and the fact that the Schrödinger equation is linear means that it respects this structure. In vector-space language, the choice between expressing the wave function in configuration space or momentum space is simply a change of basis in Hilbert space, which presumably has no physical importance whatsoever. The physical quantum state – as distinguished from its representation in some particular basis – is simply a vector in Hilbert space, sometimes called the “state vector,” and written in Dirac notation as $|\Psi\rangle$.

It would therefore seem natural, if our goal is to take a nature’s-eye view of things and specify the correct quantum theory of the world in its own right, to define that theory as a set of state vectors $|\Psi\rangle$ in a Hilbert space \mathcal{H} , evolving under the Schrödinger equation via a specified Hamiltonian \hat{H} . The problem is that this is very little structure indeed. Hilbert space itself is featureless; a particular choice of Hilbert space is completely specified by its dimension $d = \dim \mathcal{H}$. A vector in Hilbert space contains no direct specification of what the physical content of such a state is supposed to be; there is no mention of space, configuration space, particles, fields, or any such familiar notions. Presumably all of that is going to have to somehow emerge from the dynamics, which seems like a tall order.

One might imagine that we can somehow read off the physical structure being described from the explicit form of the Hamiltonian. For example, if we were handed

$$\hat{H} = -\frac{\hbar^2}{2m} \frac{\partial^2}{\partial x^2} + \frac{1}{2} m \omega^2 \hat{x}^2, \quad (2)$$

we would quickly surmise that we had a simple harmonic oscillator on our hands. But that’s only because we have conveniently been given the Hamiltonian in a particularly useful basis, in this case the position basis $|x\rangle$. That is not part of the specification of the theory itself. Hilbert space does not come equipped with a preferred basis; we should be able to deduce that the position basis (or some other one) is useful to our analysis of the system, rather than assuming we have been given it from the start.

Given that the only data that comes along with Hilbert space is its dimension d , the only other thing we have to work with is the Hamiltonian considered as an abstract operator. That does pick out one particular basis: that of the energy eigenstates, vectors satisfying $\hat{H} |n\rangle = E_n |n\rangle$. (For simplicity we assume non-degenerate energy eigenvalues, so that the states $\{|n\rangle\}$ define a unique basis.) There is no information contained in the specification of the energy eigenstates themselves; they are just a collection of orthonormal vectors. The information about the Hamiltonian is contained entirely in its set of eigenvalues $\{E_n\}$, called the “spectrum” of the Hamiltonian. A specification of a quantum theory consists entirely, therefore, of this list of

real numbers, the energy eigenspectrum.

There is an important caveat to this statement. If Hilbert space is non-separable (infinite non-countable dimension), there can be unitarily inequivalent representations of the canonical commutation relations (Haag, 1955). It is therefore necessary to give additional information to define the theory; typically this would amount to specifying an algebra of observables.

We have good reason, however, to expect that the Hilbert space for the real world is finite-dimensional, at least if we restrict our attention to our observable universe or any other finite region of space (Banks, 2001; Bao et al., 2017; Bekenstein, 1981; Bousso, 1999; Jacobson, 2012). The reason is gravity. In the presence of gravity, to make a long story short, the highest-entropy configuration we can construct in a spherical region \mathbf{R} of radius r is a black hole, and black holes have a finite entropy $S_{\text{BH}} = A/4G = \pi r^2/G$, where A is the area of the event horizon and G is Newton’s constant (and we set $\hbar = c = 1$). If we decompose Hilbert space into a tensor product of a factor describing that region and one describing the rest of the world,

$$\mathcal{H} = \mathcal{H}_{\mathbf{R}} \otimes \mathcal{H}_{\mathbf{E}}, \tag{3}$$

it follows that we have an upper bound on the dimensionality of $\mathcal{H}_{\mathbf{R}}$, given by

$$\dim \mathcal{H}_{\mathbf{R}} \leq \exp(e^{\pi r^2/G}). \tag{4}$$

If \mathbf{R} represents our observable universe with current values of the cosmological parameters, this works out to approximately $\dim \mathcal{H}_{\mathbf{R}} \leq e^{e^{123}}$, which is large but still smaller than infinity.¹

We don’t know whether the dimension of the full Hilbert space \mathcal{H} is finite or infinite, other than the indirect consideration that infinite time evolution in a finite-dimensional Hilbert space would lead to a proliferation of Boltzmann Brains (Carroll, 2008; Dyson et al., 2002). But for our current purposes, it suffices that operations confined to our observable universe effectively act on a finite-dimensional part of Hilbert space. In that case, the set of in-principle observables is simply all Hermitian operators on $\mathcal{H}_{\mathbf{R}}$. What can be observed in practice will depend on how we split further split Hilbert space into systems and observers, but the theory is completely specified by the Hilbert space and Hamiltonian, or in other words, by the discrete list of numbers $\{E_n\}$ in the energy spectrum. By which we mean the spectrum of the Hamiltonian acting on $\mathcal{H}_{\mathbf{R}}$ itself, assuming that interactions with degrees of freedom in $\mathcal{H}_{\mathbf{E}}$ can be neglected for practical purposes. Henceforth we will speak as if the finite-dimensional factor $\mathcal{H}_{\mathbf{R}}$ is the relevant part of Hilbert space, and our world can be described by a unitarily-evolving state within it. Technically it is more likely to be a mixed state described by a density operator, but that can always be purified by adding a finite-dimensional auxiliary factor to Hilbert space, so we won’t worry about such details.

¹There are a number of nuances here. In gauge theories, we cannot precisely decompose Hilbert space into factors representing regions of space. And we are somewhat cheating by invoking “regions of space” at all, although this will be a sensible notion on individual semiclassical branches of the universal quantum state. For elaboration see (Bao et al., 2017).

The challenge facing such an approach should be clear. The world of our experience doesn't *seem like* a vector in Hilbert space, evolving according to a list of energy eigenvalues. It seems like there is space, and objects located in space, and those objects interact with each other, and so forth. How in the world is all of that supposed to come from a description as abstract and featureless as a vector evolving through Hilbert space? To make matters seemingly worse, the actual evolution is pretty trivial; in the energy eigenbasis, an exact solution to the Schrödinger equation from an initial state $|\Psi(0)\rangle = \sum_n \psi_n |n\rangle$ is

$$|\Psi(t)\rangle = \sum_n \psi_n e^{-iE_n t} |n\rangle. \quad (5)$$

Each component simply evolves via an energy-dependent phase factor. It seems like a long way from the complicated nonlinear dynamics of our world. This challenge accounts for the epigraphs at the beginning of this paper.

Everettians tend to think that the right strategy for understanding the fundamental nature of reality is not necessarily to start with what the world seems like and to construct an ontology that hews as closely as possible to that. Rather, we should start with some proposed ontology and ask what it would seem like to observers (if any such exist) described by it. Clearly, “observers” are not represented directly by a vector in Hilbert space, nor is the world that they observe. What we can instead ask is whether there could be a higher-level description, emergent from our ontology, that can successfully account for our world. Such a description is not forced on us at the God’s-eye (or Laplace’s-Demon’s eye) view of the world. It would always be possible to say that reality is a vector in Hilbert space, evolving through time, and stop at that. But that’s not the only thing we’re allowed to say. The search for emergent levels is precisely the search for higher-level, non-fundamental descriptions that approximately capture some of the relevant dynamics, perhaps on the basis of incomplete information about the fundamental state. The question is whether we *can* recover the patterns and phenomena of our experience (space, objects, interactions) from the behavior of our fundamental ontology.

To get our bearings, consider the classic case of N massive particles moving in three-dimensional space under the rules of classical Newtonian gravity. The state of the system is specified by one point in a $6N$ -dimensional phase space. Yet there is an overwhelming temptation to say that the system “really lives” in three-dimensional space, not the $6N$ -dimensional phase space. Can we account for where that temptation comes from without postulating any *a priori* metaphysical essence to three-dimensional space?

There are two features of the description as N particles that make it seem more natural than that featuring a single point in phase space, even though they are mathematically equivalent. The first is that the *internal dynamics* of the system are more easily interpreted in the N -particle language. For example, it is immediately clear that two particles will strongly affect each other when they are nearby and the others are relatively far away. This kind of partial and approximate understanding of the dynamics is transparent in the N -particle description,

and obscured in the point-in-phase-space description. The second is that the system *looks like* N particles. That is, in the real-world analogues of this toy model, when we observe the system by interacting with it as a separate physical system ourselves, what we immediately see are N particles. There can be multiple equivalent ways of describing the internal dynamics of a system, but the one we think of as “natural” or describing what “really exists” is often predicated on how that system interacts with the outside world.

Similar considerations apply to familiar examples of emergence, such as treating a box of many atoms as a fluid. In this case the two descriptions are not equivalent – the emergent fluid description is an approximation obtained by coarse-graining – but the same principles apply. The internal dynamics of the particles in the box are more easily apprehended in the fluid description (fewer variables and equations are required, given some short-distance coarse-graining scale), and we can measure the fluid properties directly (using thermometers and barometers and so on), while the states of each individual atom are inaccessible to us. Emergent structures that accurately describe the dynamics of part of a system using only information accessible within the emergent description itself have every right to be thought of as “real,” even if they are not “fundamental” (Dennett, 1991; Wallace, 2012).

These two considerations (internal dynamics and what we see) work in tandem: we want the information we gather by observing a system to be sufficient for us to predict its subsequent behavior. The ontology that seems most natural for us to ascribe to a physical subsystem thus depends on how that subsystem interacts with the rest of the world (Zanardi et al., 2004). For our program of Hilbert-space fundamentalism, this suggests that we should look for emergent descriptions by considering ways to factorize \mathcal{H} into a tensor product of factors representing different subsystems, and ask how those subsystems interact with each other. [For certain situations we might also consider direct-sum structures (Kabernik et al., 2020).]

We are therefore interested in *quantum mereology*: how to decompose the whole of Hilbert space into parts such that individual subsystems have simple internal dynamics, and those dynamics are readily observed via interactions with other subsystems, given nothing but the spectrum of the Hamiltonian [(Carroll & Singh, 2019); for related work see (Brun & Hartle, 1999; Hartle, 2011; Ney, 2020a; Tegmark, 2015)]. This seems ambitious, but turns out to be surprisingly tractable.

Consider the most basic thing we might want to describe by such a factorization, the distinction between a “system” representing a macroscopic object exhibiting quasi-classical behavior and an “environment” describing degrees of freedom that passively monitor the system and lead to decoherence. This corresponds to expressing Hilbert space as the tensor product

$$\mathcal{H} = \mathcal{H}_S \otimes \mathcal{H}_E. \tag{6}$$

Fixing the dimensions of \mathcal{H}_S and \mathcal{H}_E , different factorizations are related by unitary transformations that mix the two factors together. For any specific factorization, we automatically

get a decomposition of the Hamiltonian into a self-Hamiltonian for the system, another self-Hamiltonian for the environment, and an interaction term:

$$\hat{H} = \hat{H}_S \otimes \mathbb{I}_E + \mathbb{I}_S \otimes \hat{H}_E + \hat{H}_{\text{int}}. \quad (7)$$

We can now ask, within the set of all possible factorizations, which ones lead to system dynamics that can be described by approximate classical behavior in appropriate circumstances?

In this decomposition, internal dynamics of the system are governed by \hat{H}_S and how it is observed by the environment is governed by \hat{H}_{int} . (We assume that $\dim \mathcal{H}_E \gg \dim \mathcal{H}_S$ and that the environment’s internal dynamics are largely irrelevant.) To recover classical behavior when the system is macroscopic, we want localized wave packets in \mathcal{H}_S to stay relatively localized and follow classical equations of motion under \hat{H}_S , but also for unentangled states in $\mathcal{H}_S \otimes \mathcal{H}_E$ to remain relatively unentangled; the environment is supposed to passively monitor the system, not rapidly reach maximal entanglement with it. It was argued by Carroll & Singh (2019) that generic Hamiltonians, as defined by their spectra, have neither of these features in any factorization; the Hamiltonian of the real world is apparently non-generic, to nobody’s surprise.

When the Hamiltonian and the system/environment split allow for classical behavior, the density operator for the system rapidly diagonalizes in the dynamically-preferred *pointer basis* $\{|\phi_n\rangle\}$, such that the corresponding pointer states are robust under environmental monitoring (Zurek, 1981). The pointer basis defines a pointer observable in the system’s Hilbert space, $\hat{Q}_S = \sum |\phi_n\rangle \langle \phi_n|$. Pointer states are the ones that appear classical, and a standard pointer observable is the position in space of a system; a macroscopically coherent cat is described by a pointer state, but a superposition of cats in different physical configurations is not (Joos & Zeh, 1985; Riedel et al., 2012; Zurek, 1993). In the Quantum Measurement Limit, where \hat{H}_{int} dominates over \hat{H}_S , the pointer observable satisfies Zurek’s commutativity criterion (Zurek, 1981),

$$[\hat{H}_{\text{int}}, \hat{Q}_S \otimes \mathbb{I}_E] \approx 0, \quad (8)$$

reflecting the fact that the system in a pointer state does not continually entangle with the environment.

These ideas about pointer states and decoherence are usually discussed within the context of a known factorization, but they also suggest a criterion for determining the factorization that allows for classical behavior (Carroll & Singh, 2021). Given \mathcal{H} and \hat{H} , we search through all possible factorizations, and for each one we define a candidate pointer observable as the one that comes closest to achieving (8). This candidate pointer observable embodies the idea that external observers “see” certain features of the system, and that those features evolve classically. We can then calculate, from an initially localized and unentangled system state, the rate of delocalization and entanglement growth. The correct factorization is the one that minimizes both. Simple numerical examples verify that this criterion picks out what we usually

think of as the standard system/environment split. Again, the sought-after behavior is non-generic; it doesn't occur for random Hamiltonians, nor for random factorizations with any given Hamiltonian.

The data contained within the spectrum of \hat{H} therefore seems to provide the requisite information to pinpoint an appropriate emergent classical description, when one exists. Thus reassured, let us be somewhat more ambitious, and train our sights on the ontology that seems to accurately describe our low-energy world (Carroll, 2021): an effective quantum field theory that includes both gravity and the Standard Model of particle physics, defined in four-dimensional spacetime.

We can again take guidance from our experience of the real world. In quantum field theory (modulo previously-mentioned nuances), we can think about degrees of freedom as being associated with regions of space. If we partition space into regions indexed by α , there is a corresponding decomposition of Hilbert space into a tensor product of local factors,

$$\mathcal{H} = \bigotimes_{\alpha} \mathcal{H}_{\alpha}. \quad (9)$$

This lets us expand the Hamiltonian as a sum of self-Hamiltonians for each region, plus interactions between pairs of regions, triplets of regions, and so on:

$$\hat{H} = \sum_a h_a \hat{O}_a^{(\text{self})} + \sum_{ab} h_{ab} \hat{O}_{ab}^{(2\text{-pt})} + \sum_{abc} h_{abc} \hat{O}_{abc}^{(3\text{-pt})} + \dots, \quad (10)$$

where the h_{\dots} are numerical parameters.

A necessary requirement for the emergence of a structure recognizable as “space,” with local dynamics therein, is that the series (10) does not continue indefinitely when the factorization (9) corresponds to local regions. Degrees of freedom only interact with a small number of nearest neighbors, not with regions arbitrarily far away, which means that we only need interactions between a small number of Hilbert-space factors to capture the entire Hamiltonian. Happily, this requirement is also essentially sufficient. As shown by Cotler et al. (2019), generic Hamiltonians admit no local factorization at all, and when such a factorization exists, it is unique up to irrelevant internal transformations (and some technicalities that we won't go into here). Therefore, the spectrum of the Hamiltonian is enough to pick out the correct notion of an emergent spatial structure when one exists.

A natural next step would be to define an emergent dynamical metric on spacetime, as in general relativity. Here results are much less definitive, but initial auguries are promising. [For alternative perspectives see (Carlip, 2014; Giddings, 2015; Hu, 2009; Huggett & Wüthrich, 2013; Ney, 2020b; Nielsen & Kleppe, 2013; Raasakka, 2017).] In the local decomposition of Hilbert space (9), given an overall state $|\Psi\rangle$ we can calculate the reduced density matrix $\hat{\rho}_{\alpha}$ for each factor or collection of factors, and the corresponding von Neumann entropy $S_{\alpha} = -\text{Tr} \hat{\rho} \log \hat{\rho}$.

The amount of entanglement between two factors can be measured by the quantum mutual information,

$$I(\alpha : \beta) = S_\alpha + S_\beta - S_{\alpha\beta}. \quad (11)$$

In the vacuum state of a quantum field theory, we know that the entanglement between two regions decreases monotonically (exponentially in theories with massive fields, as a power law in conformal theories) with the distance between them. We can therefore imagine turning this around, and *defining* a distance metric depending inversely on the mutual information.² If there exists an emergent best-fit smooth geometry, that can be uniquely determined using techniques such as classical multidimensional scaling. Thus, both the dimensionality and geometry of space can be defined using the entanglement information contained in a state $|\Psi\rangle$, at least if that state is near to the vacuum (Cao et al., 2017).

We can go further, and show that under certain optimistic assumptions (most notably, the eventual emergence of approximate Lorentz invariance), the resulting geometry obeys Einstein’s equation of general relativity in the weak-field limit (Cao & Carroll, 2018). The basic idea, following Jacobson (1995, 2016), is to posit that perturbations in the quantum state lead to a change in the emergent area of a codimension-one surface that is proportional to the change in entanglement entropy across that surface,

$$\delta\mathcal{A} \propto \delta S. \quad (12)$$

The area perturbation, being a geometric quantity, can then be related to the Einstein tensor, while in the long-distance limit the entropy can be related to the stress-energy tensor. Assuming emergent Lorentz symmetry, the resulting dynamical equation for a perturbation of the vacuum (representing flat spacetime) is

$$\delta G_{\mu\nu} \propto \delta T_{\mu\nu}, \quad (13)$$

just as Einstein leads us to expect.

This is a provocative result, but one that shouldn’t be over-interpreted. The accomplishment is not that we recover specifically Einstein’s equation; there aren’t that many other locally Lorentz-invariant equations one could imagine for a dynamical spacetime metric. What matters is that the necessary information required for such a description to emerge can be found in an abstract quantum state evolving according to a Hamiltonian specified purely by its spectrum, without any additional ontological assistance. Moreover, while we can use entanglement to define *an* emergent metric, we haven’t shown that it is *the* emergent metric, the one whose

²The resulting metric is defined on the same spacetime structure in which entanglement exists. This is in contrast to the emergence of spacetime from entanglement in the AdS/CFT context (Faulkner et al., 2014; Maldacena & Susskind, 2013; Swingle, 2012; van Raamsdonk, 2010), where entanglement is on the boundary and geometry is in the bulk. There is no inconsistency, as the procedure described here applies to weak-field situations far from any horizons, while the AdS/CFT construction extends over a cosmological spacetime where large-scale curvature is a central part of the description.

geodesics define the motion of localized test particles. We can imagine how this might happen; as just one piece of the puzzle, the Lieb-Robinson bound in quantum information theory (Lieb & Robinson, 1972) provides a natural mechanism whereby light-cone structures can emerge from quantum information theory (Hamma et al., 2009).

What strictly emerges from entanglement in this story is the metric on space, not on space-time; a reconstruction of the latter depends on a procedure for stitching together spatial surfaces evolving over time. Since our entire analysis is based on the eigenvalues of the Hamiltonian appearing in the Schrödinger equation $\hat{H} |\Psi\rangle = i\partial_t |\Psi\rangle$, we need to imagine that time itself is fundamental, rather than emergent. The appearance of an explicit time parameter does not imply that this parameter is in some sense preferred. Any Lorentz-invariant theory can be written in Hamiltonian form by choosing a frame; such a form will not be manifestly Lorentz invariant, but the symmetry is still there, so there is no obstacle to the emergent theory being approximately Lorentz- and diffeomorphism-invariant. It would be natural if Lorentz invariance is indeed only approximate, since we are working in a finite-dimensional factor of Hilbert space, and there are no nontrivial representations of non-compact symmetry groups on finite-dimensional vector spaces. This raises the intriguing possibility of potential experimental signatures of these ideas, as Lorentz invariance can be tested to high precision Liberati (2013). One might also worry about compatibility with the Wheeler-DeWitt equation of quantum gravity, which takes the form $\hat{H} |\Psi\rangle = 0$. In this case the eigenvalues of \hat{H} are seemingly irrelevant, since the world is fully described by a zero-energy eigenstate [cf. (Albrecht & Iglesias, 2008)]. But there is also no fundamental time evolution; that is the well-known “problem of time” (Anderson, 2010). A standard solution is to imagine that time is emergent, which amounts to writing the full Hamiltonian as

$$\hat{H} = \hat{H}_{\text{eff}} - i \frac{d}{d\tau}, \quad (14)$$

where τ is the emergent time parameter. In that case everything we have said thus far still goes through, only using the eigenvalues of the effective Hamiltonian \hat{H}_{eff} .

In addition to spacetime, we still have to show how local quantum fields can emerge in the same sense as the spacetime metric. Less explicit progress has been made in reconstructing approximate quantum field theories from the spectrum of the Hamiltonian, but it’s not unreasonable to hope that this task is more straightforward than reconstructing spacetime itself. One promising route is via “string net condensates,” which have been argued to lead naturally to emergent gauge bosons and fermions (Levin & Wen, 2005).

Nothing in this perspective implies that we should think of spacetime or quantum fields as illusory. They are emergent, but none the less real for that. As mentioned, we we may not be forced to invoke these concepts within our most fundamental picture, but the fact that they play a role in an emergent description is highly non-trivial. (Most Hamiltonians admit no local decomposition, most factorizations admit no classical limit, etc.) It is precisely this non-generic characteristic of the specific features of the world of our experience that makes it possible to

contemplate uniquely defining them in terms of the austere ingredients of the deeper theory. They should therefore be thought of as equally real as tables and chairs.

This has been an overly concise discussion of an ambitious research program (and one that may ultimately fail). But the lesson for fundamental ontology is hopefully clear. Thinking of the world as represented by simply a vector in Hilbert space, evolving unitarily according to the Schrödinger equation governed by a Hamiltonian specified only by its energy eigenvalues, seems at first hopelessly far away from the warm, welcoming, richly-structured ontology we are used to thinking about in physics. But recognizing that the latter is plausibly a higher-level emergent description, and contemplating the possibility that the more fundamental vocabulary is the one straightforwardly suggested by our simplest construal of the rules of quantum theory, leads to a reconstruction program that appears remarkably plausible. By taking the prospect of emergence seriously, and acknowledging that our fondness for attributing metaphysical fundamentality to the spatial arena is more a matter of convenience and convention than one of principle, it is possible to see how the basic ingredients of the world might be boiled down to a list of energy eigenvalues and the components of a vector in Hilbert space. If it did succeed, this project would represent a triumph of unification and simplification, and is worth taking seriously for that reason alone.

Acknowledgments

It is a pleasure to thank Ashmeet Singh, Charles Cao, Spiros Mikhalkis, and Ning Bao for collaboration on some of the work described here. This research is funded in part by the Walter Burke Institute for Theoretical Physics at Caltech, by DOE grant DE-SC0011632, and by the Foundational Questions Institute.

References

- Aharonov, Y., & Rohrlich, D. (2005). *Quantum Paradoxes*. Wiley, New York.
- Albert, D. (1996). Elementary quantum metaphysics. In J. Cushing, A. Fine, & S. Goldstein (Eds.) *Bohmian Mechanics and Quantum Theory: An Appraisal*, (pp. 277–284). Kluwer, Dordrecht.
- Albrecht, A., & Iglesias, A. (2008). The Clock ambiguity and the emergence of physical laws. *Phys. Rev.*, *D77*, 063506.
- Allori, V. (2013). Primitive ontology and the structure of fundamental physical theories. In A. Ney, & D. Albert (Eds.) *The Wave Function: Essays on the Metaphysics of Quantum Mechanics*, (pp. 58–75). Oxford University Press, Oxford.

- Anderson, E. (2010). The Problem of Time in Quantum Gravity.
URL <https://arxiv.org/abs/1009.2157>
- Banks, T. (2001). Cosmological breaking of supersymmetry? *Int. J. Mod. Phys., A16*, 910–921.
- Bao, N., Carroll, S. M., & Singh, A. (2017). The Hilbert Space of Quantum Gravity Is Locally Finite-Dimensional. *Int. J. Mod. Phys., D26*(12), 1743013.
- Bekenstein, J. D. (1981). A Universal Upper Bound on the Entropy to Energy Ratio for Bounded Systems. *Phys. Rev., D23*, 287.
- Bousso, R. (1999). A Covariant entropy conjecture. *JHEP*, 07, 004.
- Brun, T. A., & Hartle, J. B. (1999). Classical dynamics of the quantum harmonic chain. *Physical Review D*, 60(12), 123503.
- Cao, C., & Carroll, S. M. (2018). Bulk entanglement gravity without a boundary: Towards finding Einstein’s equation in Hilbert space. *Phys. Rev. D*, 97(8), 086003.
- Cao, C., Carroll, S. M., & Michalakis, S. (2017). Space from Hilbert Space: Recovering Geometry from Bulk Entanglement. *Phys. Rev.*, D95(2), 024031.
- Carlip, S. (2014). Challenges for Emergent Gravity. *Stud. Hist. Phil. Sci. B*, 46, 200–208.
- Carroll, S. M. (2008). What if Time Really Exists?
URL <https://arxiv.org/abs/0811.3772>
- Carroll, S. M. (2021). The Quantum Field Theory on Which the Everyday World Supervenes.
URL <https://arxiv.org/abs/2101.07884>
- Carroll, S. M., & Singh, A. (2019). Mad-Dog Everettianism: Quantum Mechanics at Its Most Minimal. In A. Aguirre, B. Foster, & Z. Merali (Eds.) *What Is Fundamental?*, (pp. 95–104). Springer.
- Carroll, S. M., & Singh, A. (2021). Quantum mereology: Factorizing Hilbert space into subsystems with quasiclassical dynamics. *Phys. Rev. A*, 103(2), 022213.
- Coleman, S. (1975). Quantum sine-gordon equation as the massive thirring model. *Phys. Rev. D*, 11, 2088–2097.
- Cotler, J. S., Penington, G. R., & Ranard, D. H. (2019). Locality from the Spectrum. *Commun. Math. Phys.*, 368(3), 1267–1296.
- Dennett, D. C. (1991). Real patterns. *The Journal of Philosophy*, 88(1), 27–51.

- Dyson, L., Kleban, M., & Susskind, L. (2002). Disturbing implications of a cosmological constant. *JHEP*, *10*, 011.
- Faulkner, T., Guica, M., Hartman, T., Myers, R. C., & Van Raamsdonk, M. (2014). Gravitation from Entanglement in Holographic CFTs. *JHEP*, *03*, 051.
- Giddings, S. B. (2015). Hilbert space structure in quantum gravity: an algebraic perspective. *JHEP*, *12*, 099.
- Haag, R. (1955). On Quantum Field Theories. *Matematisk-fysiske Meddelelser*, *29*(12), 1–37.
- Hamma, A., Markopoulou, F., Premont-Schwarz, I., & Severini, S. (2009). Lieb-Robinson bounds and the speed of light from topological order. *Phys. Rev. Lett.*, *102*, 017204.
- Hartle, J. B. (2011). The Quasiclassical realms of this quantum universe. *Found. Phys.*, *41*, 982–1006.
- Hu, B.-L. (2009). Emergent/quantum gravity: macro/micro structures of spacetime. In *Journal of Physics: Conference Series*, vol. 174, (p. 012015). IOP Publishing.
- Huggett, N., & Wüthrich, C. (2013). Emergent spacetime and empirical (in)coherence. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, *44*(3), 276–285.
- Jacobson, T. (1995). Thermodynamics of Spacetime: The Einstein Equation of State. *Physical Review Letters*, *75*, 1260–1263.
- Jacobson, T. (2012). Gravitation and vacuum entanglement entropy. *Int. J. Mod. Phys., D21*, 1242006.
- Jacobson, T. (2016). Entanglement Equilibrium and the Einstein Equation. *Phys. Rev. Lett.*, *116*(20), 201101.
- Joos, E., & Zeh, H. D. (1985). The Emergence of classical properties through interaction with the environment. *Z. Phys. B*, *59*, 223–243.
- Kabernik, O., Pollack, J., & Singh, A. (2020). Quantum State Reduction: Generalized Bipartitions from Algebras of Observables. *Phys. Rev. A*, *101*(3), 032303.
- Levin, M. A., & Wen, X.-G. (2005). String net condensation: A Physical mechanism for topological phases. *Phys. Rev. B*, *71*, 045110.
- Liberati, S. (2013). Tests of Lorentz invariance: a 2013 update. *Class. Quant. Grav.*, *30*, 133001.

- Lieb, E. H., & Robinson, D. W. (1972). The finite group velocity of quantum spin systems. *Commun. Math. Phys.*, *28*, 251–257.
- Maldacena, J., & Susskind, L. (2013). Cool horizons for entangled black holes. *Fortschritte der Physik*, *61*, 781–811.
- Maldacena, J. M. (1999). The Large N limit of superconformal field theories and supergravity. *Int. J. Theor. Phys.*, *38*, 1113–1133.
- Myrvold, W. C. (2015). What is a wavefunction? *Synthese*, *192*(10), 3247–3274.
- Ney, A. (2013). Introduction. In A. Ney, & D. Albert (Eds.) *The Wave Function: Essays on the Metaphysics of Quantum Mechanics*, (pp. 1–51). Oxford University Press, Oxford.
- Ney, A. (2020a). Finding the world in the wave function: Some strategies for solving the macro-object problem. *Synthese*, *197*(10), 4227–4249.
- Ney, A. (2020b). From quantum entanglement to spatiotemporal distance. In C. Wüthrich, B. Le Bihan, & N. Huggett (Eds.) *Philosophy Beyond Spacetime: Implications from Quantum Gravity*. Oxford University Press, Oxford.
- Nielsen, H. B., & Kleppe, A. (2013). Towards a Derivation of Space. *Bled Workshops Phys.*, *14*(2), 171–196.
- North, J. (2013). The structure of a quantum world. In A. Ney, & D. Albert (Eds.) *The Wave Function: Essays on the Metaphysics of Quantum Mechanics*, (pp. 184–202). Oxford University Press, Oxford.
- Raasakka, M. (2017). Spacetime-Free Approach to Quantum Theory and Effective Spacetime Structure. *SIGMA*, *13*, 006.
- Riedel, J. C., Zurek, W. H., & Zwolak, M. (2012). The rise and fall of redundancy in decoherence and quantum Darwinism. *New Journal of Physics*, *14*(8), 083010.
- Saunders, S. (2021). The everett interpretation: structure. To be published in *The Routledge Companion to Philosophy of Physics*, E. Knox and A. Wilson (eds.).
URL <http://philsci-archive.pitt.edu/18772/>
- Swingle, B. (2012). Entanglement renormalization and holography. *Phys. Rev. D*, *86*(6), 065007.
- Tegmark, M. (2015). Consciousness as a State of Matter. *Chaos Solitons Fractals*, *76*, 238–270.
- van Raamsdonk, M. (2010). Building up spacetime with quantum entanglement. *General Relativity and Gravitation*, *42*, 2323–2329.

- Wallace, D. (2003). Everett and structure. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 34(1), 87–105.
- Wallace, D. (2012). *The emergent multiverse: Quantum theory according to the Everett interpretation*. Oxford University Press.
- Wallace, D. (2017). Against wavefunction realism. In S. Dasgupta, R. Dotan, & B. Weslake (Eds.) *Current Controversies in Philosophy of Science*, Current Controversies in Philosophy, (pp. 63–74). Routledge, New York.
- Zanardi, P., Lidar, D. A., & Lloyd, S. (2004). Quantum tensor product structures are observable induced. *Phys. Rev. Lett.*, 92, 060402.
- Zurek, W. H. (1981). Pointer Basis of Quantum Apparatus: Into What Mixture Does the Wave Packet Collapse? *Phys. Rev.*, D24, 1516–1525.
- Zurek, W. H. (1993). Preferred states, predictability, classicality and the environment-induced decoherence. *Progress of Theoretical Physics*, 89(2), 281–312.