

The Minimal Modal Interpretation of Quantum Theory

Jacob A. Barandes^{1,*} and David Kagan^{2,†}

¹*Jefferson Physical Laboratory, Harvard University, Cambridge, MA 02138*

²*Department of Physics, University of Massachusetts Dartmouth, North Dartmouth, MA 02747*

(Dated: May 5, 2017)

We introduce a realist, unextravagant interpretation of quantum theory that builds on the existing physical structure of the theory and allows experiments to have definite outcomes but leaves the theory's basic dynamical content essentially intact. Much as classical systems have specific states that evolve along definite trajectories through configuration spaces, the traditional formulation of quantum theory permits assuming that closed quantum systems have specific states that evolve unitarily along definite trajectories through Hilbert spaces, and our interpretation extends this intuitive picture of states and Hilbert-space trajectories to the more realistic case of open quantum systems despite the generic development of entanglement. We provide independent justification for the partial-trace operation for density matrices, reformulate wave-function collapse in terms of an underlying interpolating dynamics, derive the Born rule from deeper principles, resolve several open questions regarding ontological stability and dynamics, address a number of familiar no-go theorems, and argue that our interpretation is ultimately compatible with Lorentz invariance. Along the way, we also investigate a number of unexplored features of quantum theory, including an interesting geometrical structure—which we call subsystem space—that we believe merits further study. We conclude with a summary, a list of criteria for future work on quantum foundations, and further research directions. We include an appendix that briefly reviews the traditional Copenhagen interpretation and the measurement problem of quantum theory, as well as the instrumentalist approach and a collection of foundational theorems not otherwise discussed in the main text.

* barandes@physics.harvard.edu

† dkagan@umassd.edu

CONTENTS

| | |
|---------------------------------------|----|
| I. Introduction | 2 |
| II. Preliminary Concepts | 6 |
| III. The Minimal Modal Interpretation | 16 |
| IV. The Measurement Process | 45 |
| V. Lorentz Invariance and Locality | 53 |
| VI. Conclusion | 61 |
| Acknowledgments | 75 |
| Appendix | 76 |
| References | 84 |

I. INTRODUCTION

A. Why Do We Need a New Interpretation?

1. The Copenhagen Interpretation

Any mathematical-physical framework like quantum theory¹ requires an interpretation, by which we mean *some* asserted connection between the theory’s abstract formalism and the real world’s physical, conceptual, logical, or empirical features.

The textbook formulation of the Copenhagen interpretation, with its axiomatic Born rule for computing empirical outcome probabilities and its notion of wave-function collapse for establishing the persistence of measurement outcomes, works quite well in most practical circumstances.² At least according to some surveys [271, 305], the Copenhagen interpretation is still the most popular interpretation today, despite the absence of a clear consensus on all the interpretation’s precise metaphysical commitments.

Unfortunately, the Copenhagen interpretation also suffers from a number of serious drawbacks. Most significantly, the definition of a measurement according to the Copenhagen interpretation relies on a physically questionable demarcation, known as the Heisenberg cut (*Heisenbergscher Schnitt*) [211, 322], between the large classical systems that carry out measurements and the small quantum systems that they measure. (See Figure 1.) This ill-defined Heisenberg cut has never been identified in any experiment to date and must be put into the interpretation by hand. Attempting to reframe the Heisenberg cut as being nothing more than an arbitrary quantum-classical boundary chosen for convenience conflicts with allowing quantum state vectors to represent objective features of reality.

An associated issue is the Copenhagen interpretation’s assumption of wave-function collapse, known more formally as the Von Neumann-Lüders projection postulate [227, 323]. By our use of the term “wave-function collapse” here, we are referring to the supposedly instantaneous, discontinuous change in a quantum system immediately following a measurement by a classical system, in stark contrast to the smooth time evolution that governs dynamically closed systems.

¹ We use the term “quantum theory” here in its broadest sense of referring to the general theoretical framework consisting of Hilbert spaces, operator algebras of observables, and state vectors that encompasses models as diverse as nonrelativistic point particles and quantum field theories. We are not referring specifically to the nonrelativistic models of quantum-mechanical point particles that dominated the subject in its early days.

² See Appendix 1 for a detailed treatment of the textbook definition of the Copenhagen interpretation, as well as a description of the famous measurement problem of quantum theory and a systematic classification of attempts to solve it according to the various prominent interpretations of the theory.

The Copenhagen interpretation is also unclear as to the ultimate meaning of the state vector of a system: Does a system’s state vector merely represent the experimenter’s knowledge, is it some sort of objective probability distribution,³ or is it an irreducible ingredient of reality like the state of a classical system [172]? For that matter, what constitutes an observer, and can we meaningfully talk about the state of an observer within the formalism of quantum theory? Given that no realistic system is ever perfectly free of quantum entanglement with other systems, and thus no realistic system can ever truly be assigned a specific state vector in the first place, what becomes of the popular depiction of quantum theory in which every particle is purportedly described by a specific wave function propagating in three-dimensional space? The Copenhagen interpretation leads to additional trouble when trying to make sense of thought experiments like Schrödinger’s cat, Wigner’s friend, and the quantum Zeno paradox.⁴

2. An Ideal Interpretation

Physicists and philosophers have expended much effort over many decades on the search for an alternative interpretation of quantum theory that could resolve these problems. Ideally, such an interpretation would eliminate the need for an *ad hoc* Heisenberg cut, thereby demoting measurements to an ordinary kind of interaction and allowing quantum theory to be a complete theory that seamlessly encompasses *all* systems in Nature, including observers as physical systems with quantum states of their own. Moreover, an acceptable interpretation should *fundamentally* (even if not always *superficially*) be consistent with all experimental data and other reliably known features of Nature, including relativity, and should be general enough to accommodate the large variety of both presently known and hypothetical physical systems. Such an interpretation should also address the key no-go theorems developed over the years by researchers working on the foundations of quantum theory; should not depend on concepts or quantities whose definitions require a physically unrealistic measure-zero level of sharpness in a sense that we’ll develop in this paper; and should be insensitive to potentially unknowable features of reality, such as whether we can sensibly define “the universe as a whole” as being a closed or open system.

3. Instrumentalism

In principle, a straightforward solution to this project is always available: One could instead simply insist upon instrumentalism, known in some quarters as the “shut-up-and-calculate approach” [234, 236]. Instrumentalism is a rudimentary interpretation in which one regards the mathematical formalism of quantum theory merely as being a calculational recipe or algorithm for predicting real-world measurement results and empirical outcome probabilities obtained by the kinds of physical systems (agents) that can self-consistently act as observers. Moreover, instrumentalism refrains from making further definitive metaphysical claims about any underlying reality.⁵

On the other hand, if the physics community had unquestioningly accepted instrumentalism from the very beginning of the history of quantum theory, then we might have missed out on the many important spin-offs from the search for a better interpretation: Decoherence, quantum information, quantum computing, quantum cryptography, and the black-hole information paradox are just a few of the far-reaching ideas that ultimately owe their origins to people thinking seriously about the meaning of quantum theory.

4. The Overdetermination Problem in Interpreting Quantum Theory

There now exists a large variety of candidate interpretations of quantum theory, and, indeed, it has been argued that quantum theory suffers from an interpretational underdetermination problem [189]. However, precisely because all these candidate interpretations suffer from serious flaws and shortcomings—some of which we’ll examine in detail in this paper—we argue that the real risk is that the nontrivial conceptual structure of quantum theory actually leads to an *overdetermination* problem that prevents the theory from admitting any complete, self-consistent, fully general, realist interpretation. One of our chief goals in presenting a new realist interpretation of quantum theory is, at a minimum, to provide a proof of principle that such an overdetermination problem doesn’t exist.

³ Recent work [37, 87, 264] casts considerable doubt on assertions that state vectors are nothing more than probability distributions over more fundamental ingredients of reality.

⁴ We will discuss all these thought experiments in Section IV D.

⁵ “Whereof one cannot speak, thereof one must be silent.” [344] Schrödinger described instrumentalism in the following way: “Reality resists imitation through a model. So one lets go of naïve realism and leans directly on the indubitable proposition that *actually* (for the physicist) after all is said and done there is only observation, measurement. Then all our physical thinking thenceforth has as sole basis and as sole object the results of measurements which can in principle be carried out, for we must now explicitly *not* relate our thinking any longer to any other kind of reality or to a model. All numbers arising in our physical calculations must be interpreted as measurement results.” [273] We provide a more extensive description of the instrumentalist approach in Appendix 1 e.

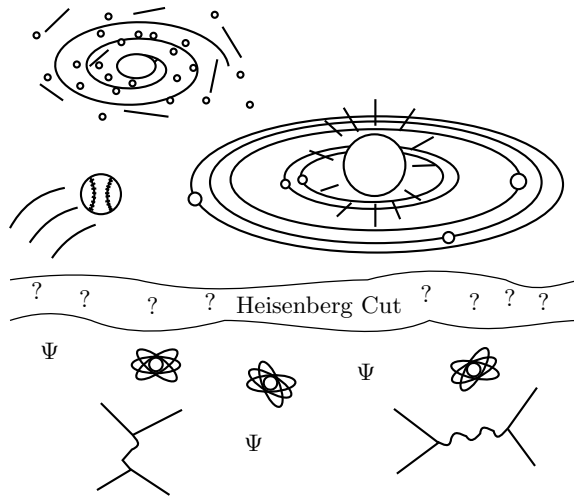


Figure 1. The Heisenberg cut.

B. Our Interpretation

In this paper and in a brief companion letter [36], we present a realist interpretation of quantum theory that hews closely to the basic structure of the theory in its widely accepted current form. Our primary goal is to move beyond instrumentalism and describe an actual reality that lies behind the mathematical formalism of quantum theory. We also intend to provide new hope to those who find themselves disappointed with the pace of progress on making sense of the theory’s foundations [333, 336].

Our interpretation is fully quantum in nature. However, for purposes of motivation, consider the basic theoretical structure of classical physics: A classical system has a specific state that evolves in time through the system’s configuration space according to some dynamical rule that may or may not be stochastic, and this dynamical rule exists whether or not the system’s state lies beneath a nontrivial evolving probability distribution on the system’s configuration space; moreover, the dynamical rule for the system’s underlying state is consistent with the overall evolution of the system’s probability distribution, in the sense that if we consider a probabilistic ensemble over the system’s initial underlying state and apply the dynamical rule to each underlying state in the ensemble, then we correctly obtain the overall evolution of the system’s probability distribution as a whole.

In particular, note the *insufficiency* of specifying a dynamical rule *solely* for the evolution of the system’s overall probability distribution but *not* for the system’s underlying state itself, because then the system’s underlying state would be free to fluctuate wildly and discontinuously between macroscopically distinct configurations. For example, even in the simple case in which a classical system’s probability distribution describes constant probabilities p_1 and p_2 for the system to be in respective macroscopically distinct states q_1 or q_2 , there would be nothing preventing the system’s state from hopping discontinuously between q_1 and q_2 with respective frequency ratios p_1 and p_2 over arbitrarily short time intervals. Essentially, by imposing a dynamical rule on the system’s underlying state, we can provide a “smoothness condition” for the system’s hidden physical configuration over time and thus eliminate these kinds of instabilities. (“Hidden variables need hidden dynamics.”)

In quantum theory, a system that is *exactly* closed and that is *exactly* in a pure state (both conditions that are unphysical idealizations) evolves along a well-defined trajectory through the system’s Hilbert space according to a well-known dynamical rule, namely, the Schrödinger equation. However, in traditional formulations of quantum theory, an *open* quantum system that must be described by a density matrix due to entanglement with other systems—that is, a system in a so-called improper mixture—does not have a specific underlying state vector, let alone a Hilbert-space trajectory or a dynamical rule governing the time evolution of such an underlying state vector and consistent with the overall evolution of the system’s density matrix.

It is a chief goal of our interpretation of quantum theory to provide these missing ingredients—in large part by assigning an explicit meaning to improper mixtures. In a sense that we will make much more precise, our interpretation of quantum theory asserts that systems have actual states that evolve along kinematical trajectories through their state spaces, and that those trajectories are governed by specific (if approximate) dynamical rules.

C. Conceptual Summary

We present a technical summary of our interpretation in Section [VI A](#). Section [VI B](#) contains tables in which we compare our interpretation with the implicit axiomatic content of classical realism as well as with both classical and quantum instrumentalism.

In short, rather than invoking the Born rule together with a collapse postulate that converts *improper* mixtures into *proper* mixtures—that is, into classical probability distributions over sets of definite outcomes—we instead attach an interpretation directly to improper mixtures: For a quantum system in a fully improper mixture, our interpretation identifies the *eigenstates* of the system’s density matrix with the *possible states* of the system in reality and identifies the *eigenvalues* of that density matrix with the *probabilities* that one of those possible states is *actually* occupied.⁶ Our interpretation introduces just enough minimal structure beyond that simple picture to provide a dynamical rule for underlying state vectors as they evolve along Hilbert-space trajectories and to evade criticisms made in the past regarding similar interpretations. This minimal additional structure consists of a new class of conditional probabilities amounting essentially to a series of smoothness conditions that kinematically relate the states of parent systems to the states of their subsystems, as well as dynamically relate the states of a single system to each other over time.

D. Comparison with Other Interpretations of Quantum Theory

1. Differences from Other Approaches

Our interpretation, which builds on the work of many others, is general, model-independent, and encompasses relativistic systems, but is also conservative and unextravagant: It includes only metaphysical objects that are either already a standard part of quantum theory or that have counterparts in classical physics. We do not posit the existence of exotic “many worlds” [[68](#), [93](#), [94](#), [96](#), [111–113](#), [325](#), [326](#), [337](#)], physical “pilot waves” [[59](#), [60](#), [63](#), [91](#)], or any fundamental GRW-type dynamical-collapse or spontaneous-localization modifications to quantum theory [[8](#), [39](#), [144](#), [249](#), [252](#), [334](#)]. Indeed, our interpretation leaves the widely accepted mathematical structure of quantum theory essentially intact.⁷ At the same time, we will argue that our interpretation is ultimately compatible with Lorentz invariance and is nonlocal only in the mild sense familiar from the framework of classical gauge theories.

Furthermore, we make no assumptions about as-yet-unknown aspects of reality, such as the fundamental discreteness or continuity of time or the dimensionality of the ultimate Hilbert space of Nature. Nor does our interpretation rely in any crucial way upon the existence of a well-defined maximal parent system that encompasses all other systems and is dynamically closed in the sense of having a so-called cosmic pure state or universal wave function that *precisely* obeys the Schrödinger equation; by contrast, this sort of unsubstantiated cosmic assumption is a necessarily *exact* ingredient in the traditional formulations of the de Broglie-Bohm pilot-wave interpretation [[59](#), [60](#), [63](#), [91](#)] and the Everett-DeWitt many-worlds interpretation [[68](#), [93](#), [94](#), [96](#), [111–113](#), [325](#), [326](#), [337](#)]. (Given their stature among interpretations of quantum theory, we will have more to say about the de Broglie-Bohm interpretation in Section [III A 6](#) and the Everett-DeWitt interpretation in Sections [III A 9](#) and [VI F 4](#).)

2. The Russian-Doll Problem

Indeed, by considering merely the *possibility* that our observable universe is but a small region of an eternally inflating cosmos of indeterminate spatial size and age [[17](#), [135](#), [137](#), [138](#), [162–164](#), [220](#), [221](#), [297](#), [318](#)], it becomes clear that the idea of a biggest closed system containing all other systems (“the universe as a whole”) may not generally be a sensible or empirically verifiable concept to begin with, let alone an axiom on which a robust interpretation of quantum theory can safely rely. Our interpretation certainly allows for the existence of a biggest closed system in an objectively pure state, but is also fully able to accommodate the alternative circumstance that if we were to imagine gradually enlarging our scope to parent systems of increasing physical size, then we might well find that the hierarchical sequence

⁶ In the language of [[333](#)], our interpretation therefore comports “with the idea that the state of a physical system is described by a vector in [a] Hilbert space rather than by numerical values of the positions and momenta of all the particles in the system,” and is not an interpretation “with no description of physical states at all.” That is, in contrast to instrumentalism, our interpretation is not “only an algorithm for calculating probabilities.”

⁷ Moreover, our interpretation does not introduce any new violations of time-reversal symmetry and gives no fundamental role to relative states [[111](#)]; a cosmic multiverse or self-locating uncertainty [[9](#)]; coarse-grained histories or decoherence functionals [[140](#), [141](#), [173](#), [174](#)]; decision theory [[94](#), [324](#), [327](#)]; Dutch-book coherence, SIC-POVMs, or *urgleichungs* [[19](#), [77](#), [78](#), [125–131](#)]; circular frequentist arguments involving unphysical “limits” of infinitely many copies of measurement devices [[9](#), [72](#), [79](#), [114](#)]; infinite imaginary ensembles [[34](#), [35](#)]; quantum reference systems or perspectivalism [[49–54](#)]; relational or non-global quantum states [[56](#), [181](#), [183](#), [266](#), [285](#)]; many-minds states [[12](#)]; mirror states [[181](#), [183](#)]; faux-Boolean algebras [[97](#), [98](#), [316](#)]; “atomic” subsystems [[28](#), [29](#)]; algebraic quantum field theory [[106](#)]; secret classical superdeterminism or fundamental information loss [[299–301](#), [303](#)]; cellular automata [[302](#), [304](#)]; classical matrix degrees of freedom or trace dynamics [[6](#), [7](#)]; or discrete Hilbert spaces or appeals to unknown Planck-scale physics [[71](#), [72](#)].

never terminates at any maximal, dynamically closed system, but may instead lead to an unending “Russian-doll” succession of ever-more-open parent systems in objectively mixed states representing improper mixtures.⁸

To avoid this problem, one might try to argue that one can always *formally* define a closed maximal parent system—presumably in an objectively pure state—just to be “the system containing all systems.” Whatever logicians might say about such a construction, we run into the more prosaic issue that if we cannot construct this closed maximal parent system via a well-defined succession of parent systems of incrementally increasing size, then it becomes unclear mathematically how we can generally define any human-scale system as a subsystem of the maximal parent system and thereby define the partial-trace operation, to be described in detail in Section III C 1. Furthermore, if our observable cosmic region is indeed an open system, then its own time evolution may not be exactly linear—an important feature of generic open systems that is rarely acknowledged in the literature [251] but that we will describe in detail later—in which case it is far from obvious that we can safely and rigorously embed that open-system time evolution into the hypothetical unitary dynamics of any conceivable closed parent system.

E. Outline of this Paper

In Section II, we lay down the conceptual groundwork for our interpretation and review features of classical physics whose quantum counterparts will play an important role. In Section III, we define our interpretation of quantum theory in precise detail and compare it with some of the other prominent interpretations, as well as identify an interesting geometrical structure, herein called subsystem space, that has been lurking in quantum theory all along.

Next, in Section IV, we describe how our interpretation makes sense of the measurement process and provide a first-principles derivation of the Born rule for computing empirical outcome probabilities. We also discuss possible corrections to the naïve Born rule that are otherwise invisible in the traditional Copenhagen interpretation, according to which, in contrast to our own interpretation, the Born rule is taken axiomatically to be an exact statement about reality. We then revisit several familiar “paradoxes” of quantum theory, including Schrödinger’s cat, Wigner’s friend, and the quantum Zeno paradox.

In Section V, we study issues of locality and Lorentz invariance and consider several well-known thought experiments and no-go theorems. Additionally, we show that our interpretation evades claims [56, 99, 106, 241, 242] that interpretations similar to our own are incompatible with Lorentz invariance and necessarily depend on the existence of a universal “preferred” inertial reference frame. We also address the question of nonlocality more generally, and argue that the picture of quantum theory that emerges from our interpretation is no more nonlocal than are classical gauge theories.

We conclude in Section VI, which contains a concise summary of our interpretation together with a discussion of its falsifiability, future research directions, and relevant philosophical issues.

In our appendix, we present a brief summary of the Copenhagen interpretation and the measurement problem of quantum theory—including a systematic analysis of the various ways in which the prominent interpretations have attempted to solve the measurement problem—as well as a description of the instrumentalist approach to quantum theory and a summary of several important foundational theorems not covered in the main text.

II. PRELIMINARY CONCEPTS

A. Ontology and Epistemology

Our aim is to present our interpretation of quantum theory in a language familiar to physicists. We make an allowance, however, for a small amount of widely used, model-independent philosophical terminology that will turn out to be very helpful.

For our purposes, we will employ the term “ontology” (adjective “ontic”) to refer to a state of being or objective physical existence—that is, the way things, including ourselves as physical observers, *really are*.⁹ This language is crucial for being able to talk about *realist* interpretations of quantum theory such as our own.

⁸ Schrödinger appears to allude to this troublesome hierarchical sequence of increasingly entangled improper mixtures when he refers in [274] to a “sinister” and “threatening” “*regressus in infinitum*” inherent in the measurement process.

⁹ Bell coined the term “beables” (that is, “be-ables”) to refer to ontic states of *being* or *reality*, as opposed to the mere “observables” (“observe-ables”) that appear in experiments. Bell explained this definition further in [45]: “The beables of the theory are those elements which might correspond to elements of reality, to things which exist. Their existence does not depend on ‘observation.’” Today, the term “beables” is often used in the quantum foundations community to refer more specifically to a *fixed* and *unchanging* set of elements of reality, such as the coordinate basis in the original nonrelativistic formulation of the de Broglie-Bohm interpretation of quantum theory. In this sense, beables are therefore conceptually distinct from the kind of ontic ingredients that appear in our own interpretation of quantum theory.

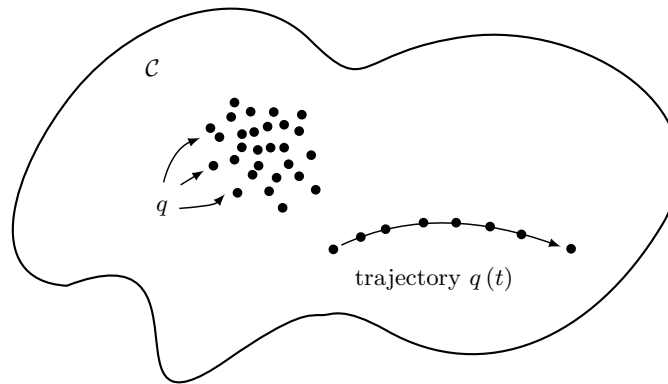


Figure 2. A schematic picture of a classical configuration space \mathcal{C} , with examples of allowed ontic states q and an example of an ontic-state trajectory $q(t)$.

Meanwhile, we will use the term “epistemology” (adjective “epistemic”) to refer to knowledge or information regarding a particular ontic piece of reality. We can further subdivide epistemology into subjective and objective parts: The subjective component refers to information that a particular observer possesses about an ontic piece of reality, whereas the objective component refers to the information that an ontic piece of reality reveals about itself to the rest of the world beyond it—meaning the *most complete* information that any observer could possibly possess about that ontic piece of reality without disturbing it.

All successful scientific theories and models make predictions about things that we can directly or indirectly observe, but, without a sufficiently extensive interpretation, we can’t really talk of an underlying ontology or its associated epistemology. To say that one’s interpretation of a mathematical theory of physics adds an ontology and an associated epistemology to the theory is to say that one is establishing some sort of connection between, on the one hand, the mathematical objects of the theory, and, on the other hand, the ontic states of things as they really exist as well as what epistemic information those ontic ingredients of reality make known about themselves and can be known by observers.

We regard an interpretation of a mathematical theory of physics as being realist if it asserts an underlying ontology of some kind—say, with some notion of systems and some notion that systems have states of some kind—and anti-realist if it asserts otherwise. We call an interpretation agnostic if it refrains from making any definitive claims about an ontology, either whether one exists at all or merely whether we can say anything specific about it. In the sense of these definitions, instrumentalism is agnostic, whereas the interpretation that we introduce in this paper is realist.

B. Classical Theories

Before laying out our interpretation of quantum theory in detail, we consider the salient features of the classical story, presented as generally as possible and in a manner intended to clarify the conceptual parallels and distinctions with the key ingredients that we’ll need in the quantum case.

1. Classical Kinematics and Ontology

The kinematical structure of a classical theory is conceptually straightforward and admits an intuitively simple ontology: An ontic state of a classical system, meaning the state of the system as it could truly exist in reality, corresponds at each instant in time t to some element q of a configuration space \mathcal{C} whose elements are by definition *mutually exclusive* possibilities. A full sequence $q(t)$ of such elements over time t constitutes an ontic-state trajectory of the system. (See Figure 2.)

2. Classical Epistemology

From the mutual exclusivity of all the system’s allowed ontic states, we also naturally obtain the associated epistemology of classical physics: In the language of probability theory, this mutual exclusivity allows us to regard the system’s configuration space as being a *single* sample space $\mathcal{C} = \Omega$, meaning that if we don’t know the system’s *actual* ontic state precisely, then we are free to describe the system in terms of an epistemic state defined to be a probability distribution $p : \mathcal{C} \rightarrow [0, 1]$ on $\mathcal{C} = \Omega$, where $p(q) \in [0, 1]$ is the epistemic probability that the actual ontic state is truly q and definitely not any other allowed ontic state $q' \neq q$.¹⁰ Equivalently, we can regard the system’s epistemic state p as the collection $\{(p(q), q)\}_{q \in \mathcal{C}}$ of ordered pairs that each identify one of the system’s *possible* ontic states $q \in \mathcal{C}$ together with the corresponding probability $p(q) \in [0, 1]$ with which q is the system’s *actual* ontic state. In analogy with ontic-state trajectories $q(t)$, we will sometimes say that a full sequence $p(t)$ of a system’s epistemic states over time constitutes an epistemic-state trajectory of the system.

Epistemic states and epistemic probabilities describing kinematical configurations in classical physics are always ultimately subjective, in the sense that they depend on the particular observer in question and are nontrivial only due to prosaic reasons of ignorance on the part of that observer. For instance, the system’s ontic state might be determined by a hidden pseudorandom number generator or by the roll of an unseen die, or the observer’s recording precision might be limited by technological constraints.

In the classical context of our present discussion, an optimal observer has no subjective uncertainty and knows the classical system’s present ontic state precisely; in this idealized case, the probability distribution $p(q)$ singles out that one ontic state with unit probability and we say that the system’s epistemic state is pure. In the more general and realistic case in which the observer’s knowledge is incomplete and the probability distribution $p(q)$ is nontrivial, we say that the system’s epistemic state is mixed.

3. Formal Epistemic States

The subjectivity of classical epistemic states means that, if we wish, we could work with *formal* epistemic probability distributions over possibilities that are *not* mutually exclusive. As an example, suppose that we fill a large container with marbles, 72% of which are known to come from a box containing only red marbles and 28% of which are known to come from a box containing marbles that are either red or blue according to an unknown red:blue ratio. If we single out one marble from the container at random, then we can formally describe the epistemic state of the marble in terms of the non-exclusive possible statements $x = \text{“red”}$ and $y = \text{“red or blue”}$ as $\{(0.72, x), (0.28, y)\}$. Similar reasoning would apply in describing the epistemic state of a marble obtained from a dispenser that stochastically releases marbles in the state x at a frequency of 72% and in the state y at a frequency of 28%.

However, to avoid ambiguities in our discussion ahead—especially when employing the notion of entropy to quantify our level of uncertainty—we will generally assume that all our classical epistemic states always encode *logically rigorous* probability distributions, by which we mean that they involve only *mutually exclusive* possibilities.

4. Classical Surprisal and Entropy

Given a system’s epistemic state p , the smooth function

$$\log \frac{1}{p(q)} \in [0, \infty], \quad (1)$$

called the surprisal [310], captures our “surprise” at learning that the system’s actual ontic state happens to be q . Our “average level of surprise,” also called the Shannon (or Gibbs) entropy of the epistemic state p , is defined by

$$S \equiv \left\langle \log \frac{1}{p} \right\rangle = - \sum_q p(q) \log p(q) \in [0, \log(\#\text{states in } \mathcal{C})] \quad (2)$$

and provides an overall measure of how much we currently *do not* yet know about the system’s ontology [192, 193, 280, 281]. That is, if we think of an epistemic state p as characterizing our information about a system’s underlying

¹⁰ We do not attempt to wade into the centuries-old philosophical debate regarding the ultimate meaning of probability in terms of anything more fundamental, as encapsulated wonderfully by Russell in 1929 (quoted in [191]): “Probability is the most important concept in modern science, especially as nobody has the slightest notion what it means.” For our purposes, we treat probability as a primitive, irreducible concept, just as we treat concepts like ontology and epistemology. Interestingly, Russell’s quotation parallels Feynman’s famous quip in 1961 that “I think I can safely say that nobody understands quantum mechanics.” [115]




Figure 3. The subject system Q and the measurement apparatus A (a) before the measurement and (b) after the measurement.

ontic state, then the entropy S defined in (2) provides us with a quantitative measure of how much information that we *lack* and that is therefore still hiding in the system.

Indeed, in the limit of a pure epistemic state, meaning that we now know the actual ontic state of the system with certainty and thus $p(q) \rightarrow 1$ for a single value of q and $p(q') \rightarrow 0$ for all $q' \neq q$, the entropy goes to zero, $S \rightarrow 0$. On the other hand, if we know nothing about the actual ontic state of the system, so that all the probabilities $p(q)$ become essentially equal, then the entropy approaches its maximally permitted value $S \rightarrow \log(\#\text{states in } \mathcal{C})$.

5. Classical Measurements and Signals

Combining classical kinematics with the definition of entropy in (2) as a total measure of the information that we lack about a system’s underlying ontic state, we can give more precise meanings to the notions of classical measurements, signals, and information transfer.

Suppose that we wish to study a system Q whose epistemic state is p_Q —that is, consisting of individual probabilities $p_Q(q)$ for each of the system’s possible ontic states q —and whose entropy $S_Q > 0$ represents the amount of information that we currently lack about the system’s actual ontic state. We proceed by sending over a measurement apparatus A in an initially known ontic state “ \emptyset ,” by which we mean that the dial on the apparatus initially reads “blank” and the initial entropy of the apparatus is $S_A = 0$. (See Figure 3a.)

When the apparatus A comes into local contact with the subject system Q and performs a measurement to determine the state of Q , the state of the measurement apparatus becomes correlated with the state of the subject system; for example, if the actual ontic state of the subject system Q happens to be q , then the ontic state of the apparatus A evolves from “ \emptyset ” to an ontic state in which the dial on the apparatus now shows the value “ q .” As a consequence, the measurement apparatus, which is still far away from us, develops a nontrivial epistemic state p'_A essentially identical to that of the subject system—meaning that $p'_A(\text{“}q\text{”}) = p_Q(q)$ —and thus the apparatus develops a nonzero final correlational entropy $S'_A = S_Q$. (See Figure 3b.)

Correlational entropy of this kind therefore indicates that there has been the transmission of a signal—that is, the exchange of *observable* information, which, as we know from relativistic causality, is constrained by the vacuum speed of light c . In the present example, this signal is communicated from the subject system Q to the measurement apparatus A , and, in particular, neither A nor Q is a closed system during the communication of this signal. We therefore also obtain a criterion for determining that *no* signal has passed between two systems, namely, that their respective epistemic states evolve independently of one another and thus neither system develops any correlational entropy.

With these concepts of information, signaling, and correlational entropy in hand, we can obtain a measure for the precision of a system’s information in terms of its correlational entropy. Consider a system—say, a measurement apparatus—that measures some previously unknown numerical property of another system using a resolution based on n distinguishable apparatus states. Then the correlational entropy of the apparatus grows by at least an amount $\Delta S \sim \log n$, and so the minimum measurement blurriness or error of the apparatus satisfies the error-entropy bound

$$\text{minimum error} \sim 1/n \sim e^{-\Delta S} \geq e^{-S}, \quad (3)$$

where $S \geq \Delta S$ is the *total* final entropy (2) of the apparatus. Hence, the precision with which a measurement apparatus can specify any numerical quantity is bounded from below by $\exp(-S)$. We will see that this same sort of error crops up in the analogous case of quantum measurements when we study them in Section IV A.¹¹

In particular, for the case of a physically realistic measurement apparatus with a finite maximum possible entropy $S_{\max} < \infty$, we find a fundamental limitation $\exp(-S_{\max})$ on the sharpness of its measurement precision. On the other

¹¹ Curiously, $\exp(-S)$ effects also seem to play an important role in the famous black-hole information paradox [147, 148, 175, 176, 263], because the mixed-state final density matrix computed semiclassically appears to differ by off-diagonal entries of size $\sim \exp(-S)$ [32, 186] from that of the pure-state final density matrix that would be required by information conservation. See also footnote 45.

hand, we could regard the abstract “agent” of instrumentalist probability theory—that is, instrumentalism’s ideal external observer—as being a formal measurement apparatus having infinite information capacity and thus an infinite maximum entropy $S_{\max} \rightarrow \infty$, in which case the minimum error would be zero, $\exp(-S_{\max}) \rightarrow 0$. The large size of S_{\max} for a typical human-scale measurement apparatus explains why one can successfully take an instrumentalist point of view toward classical probability theory and regard measurements as though they are performed by abstract agents.

6. The First and Second Laws of Thermodynamics

The present discussion provides an excellent opportunity to mention a simple set of arguments due to Jaynes [192, 194] by which one can derive the first and second laws of thermodynamics from first principles. We begin by letting $E(q)$ denote the energy of a classical system in an ontic state q and then define the system’s average energy by

$$U \equiv \langle E \rangle = \sum_q p(q) E(q). \quad (4)$$

Under a small change $dp(q)$ in the system’s epistemic state $p(q)$ and a small change $dE(q)$ in the system’s energy function $E(q)$, we see immediately that the corresponding change in the system’s average energy U is given according to the first law of thermodynamics,

$$dU = \delta Q + \delta W, \quad (5)$$

where the heat δQ consists of contributions due to statistical changes $dp(q)$ in the system’s internal state and the work δW consists of contributions due to alterations $dE(q)$ in the external dynamical parameters of the system:

$$\delta Q \equiv \sum_q dp(q) E(q), \quad (6)$$

$$\delta W \equiv \sum_q p(q) dE(q). \quad (7)$$

Jaynes argued in [192] that if all we know about a system is its average energy, then the correct epistemic state $p(q)$ that we should use to describe the system is the epistemic state that otherwise maximizes our ignorance, meaning the epistemic state that maximizes the Shannon (or Gibbs) entropy S as defined by (2). The result of this constrained maximization problem is the familiar canonical ensemble or Boltzmann distribution

$$p(q) = \frac{1}{Z} e^{-E(q)/k_B T}, \quad (8)$$

where k_B is Boltzmann’s constant, T is the system’s temperature, and the partition function Z serves as the required normalization factor:

$$Z \equiv \sum_q e^{-E(q)/k_B T}. \quad (9)$$

An elementary calculation then shows that the heat δQ is related to the system’s temperature T and the corresponding change in the system’s entropy dS according to the Clausius entropy formula $dS = \delta Q/T$, and thus the first law of thermodynamics (5) takes the textbook form

$$dU = TdS + \delta W. \quad (10)$$

A corollary to this discussion is that the maximum entropy of a system at a given temperature is equivalent to the conventional experimental (or Clausius) definition of the system’s thermodynamic entropy:

$$S = S_{\max} = S_{\text{exp}} \text{ at thermal equilibrium at a given temperature.} \quad (11)$$

As Jaynes noted [194], this observation leads to a simple proof of the traditional second law of thermodynamics. To see why, consider a system in thermal equilibrium, in which case the system’s initial Shannon entropy is maximal

$S_{\text{initial}} = S_{\text{initial,max}}$ at the given temperature and therefore agrees with the system's initial experimental entropy $S_{\text{exp,initial}}$, as in (11):

$$S_{\text{initial}} = S_{\text{initial,max}} = S_{\text{exp,initial}}. \quad (12)$$

If the system is subsequently allowed to evolve as a closed system (meaning no exchanges of signals with the outside world), then, whether or not the evolution is irreversible, the system's Shannon entropy (2) cannot change with time, and so the system's *initial* Shannon entropy is equal to its *final* Shannon entropy:

$$S_{\text{initial}} = S_{\text{final}}. \quad (13)$$

This final Shannon entropy S_{final} sets a lower bound on the system's final *maximal* possible Shannon entropy $S_{\text{final,max}}$ and thus also a lower bound on the system's *later* experimental entropy $S_{\text{exp,final}} = S_{\text{final,max}}$, which we therefore see must be greater than or equal to the system's *initial* experimental entropy $S_{\text{exp,initial}}$:

$$S_{\text{exp,final}} = S_{\text{final,max}} \geq S_{\text{final}} = S_{\text{initial}} = S_{\text{exp,initial}}. \quad (14)$$

We therefore obtain the traditional version of the second law of thermodynamics, namely, that in evolving reversibly or irreversibly from thermal equilibrium, the change in experimental entropy $\Delta S_{\text{exp}} \equiv S_{\text{exp,final}} - S_{\text{exp,initial}}$ of a closed system is non-negative:

$$\Delta S_{\text{exp}} \geq 0. \quad (15)$$

7. Classical Dynamics

A classical system has well-defined dynamics if final ontic states of the system can be predicted (probabilistically at least) based on given data characterizing initial ontic states of that *same* system. More precisely, a system has dynamics if for any choice of final time t' there exists a rule for picking prescribed initial times $t_1 \geq t_2 \geq \dots$ and there exists a mapping $p(\cdot; t' | (\cdot; t_1), (\cdot; t_2), \dots)$ that takes *arbitrary* initial ontic states q_1 at t_1 , q_2 at t_2 , \dots and yields corresponding conditional probabilities $p(q'; t' | (q_1; t_1), (q_2; t_2), \dots) \in [0, 1]$ that the system's final ontic state at t' is q' :

$$p(\cdot; t' | (\cdot; t_1), (\cdot; t_2), \dots): \underbrace{(q_1; t_1), (q_2; t_2), \dots}_{\text{initial data}}, \underbrace{(q'; t')}_{\text{final data}} \mapsto \underbrace{p(q'; t' | (q_1; t_1), (q_2; t_2), \dots)}_{\text{conditional probabilities}}. \quad (16)$$

The initial times t_1, t_2, \dots that are necessary to define the mapping could, in principle, be only infinitesimally separated, and their total number, which could be infinite, is called the order of the dynamics. (As an example, Markovian dynamics, to be defined shortly, requires initial data at only a *single* initial time, and is therefore first order.)

Just as was the case for the *epistemic* probabilities characterizing an observer's knowledge of a given system's ontic state at a single moment in time, the *conditional* probabilities that define a classical system's dynamics can be subjective in the sense that they are merely a matter of the observer's prosaic ignorance regarding details of the given system. However, these conditional probabilities can also (or instead) be objective in the sense that they cannot be trivialized based solely on knowledge about the given system itself; for example, nontrivial conditional probabilities may arise from the observer's ignorance about the properties or dynamics of a larger environment in which the given system resides and that we have implicitly averaged over—an environment that could, for instance, be infinitely big and old—or nontrivial conditional probabilities may emerge in an effective sense from chaos, to be described later on.

We call the dynamics deterministic in the idealized case in which the conditional probabilities $p(q'; t' | (q_1; t_1), (q_2; t_2), \dots)$ specify a particular ontic state q' at t' with unit probability for each choice of initial data $(q_1; t_1), (q_2; t_2), \dots$.¹² Otherwise, whether subjective, objective, or a combination of both, we say that the dynamics is stochastic.

By assumption, the conditional probabilities appearing in (16) are determined *solely* by the initial ontic data $(q_1; t_1), (q_2; t_2), \dots$ together with a specification of the final ontic state q' at t' . In particular, the conditional probabilities are required to be independent of the system's evolving epistemic state p at *all* intermediate times, and thus

¹² We address the definition and status of “superdeterminism,” a stronger notion of determinism, in Section VIF1.

$p(\cdot; t' | (\cdot; t_1), (\cdot; t_2), \dots)$ naturally lifts to a multilinear dynamical mapping relating the system's evolving epistemic state at each of the initial times $t_1 \geq t_2 \geq \dots$ to the system's epistemic state at the final time t' ,

$$p(\cdot; t' | (\cdot; t_1), (\cdot; t_2), \dots): \underbrace{p(q'; t')}_{\text{epistemic state at } t'} = \sum_{q_1, q_2, \dots} \underbrace{p(q'; t' | (q_1; t_1), (q_2; t_2), \dots)}_{\text{conditional probabilities (independent of epistemic states)}} \underbrace{p(q_1; t_1) p(q_2; t_2) \dots}_{\text{epistemic states at } t_1, t_2, \dots}, \quad (17)$$

which we can think of as a dynamical Bayesian propagation formula.

If a system has well-defined dynamics (16), then we see from (17) that the dynamics actually operates at *two* levels—namely, at the level of ontic states and also at the level of epistemic states—and consistency requires that these two levels of dynamics must be related by $p(\cdot; t' | (\cdot; t_1), (\cdot; t_2), \dots)$. Indeed, we can regard $p(\cdot; t' | (\cdot; t_1), (\cdot; t_2), \dots)$ equivalently as either describing the ontic-level dynamics in the sense of a mapping (16) from initial ontic states to the conditional probabilities for final ontic states, or as describing epistemic-level dynamics in the sense of a multilinear mapping (17) between initial and final epistemic states. This equivalence means that not only does the existence of ontic-level dynamics naturally define epistemic-level dynamics, but any given epistemic-level dynamics also determines the system's ontic-level dynamics because we can just read off the ontic-level dynamical mapping (16) from the matrix elements (17) of the multilinear dynamical mapping.

Although it's possible to conceive of an alternative world in which dynamics and trajectories are specified *only* at the *epistemic level* and *do not* determine *ontic-level* dynamics or trajectories, it's important to recognize the *inadequacy* of such a description: If all we knew were the rules dictating how epistemic states p evolve in time, then there would be nothing to ensure that a system's underlying ontic state should evolve in a sensible manner and thereby avoid fluctuating wildly between macroscopically distinct configurations. Quantum theory, as traditionally formulated, finds itself in just such a predicament: A quantum system's dynamics (when it exists) is specified in terms of density matrices (or, in the idealized case of a perfectly closed and information-conserving system, in terms of state vectors) and not directly as a multilinear dynamical mapping (17) acting on epistemic states themselves, which we distinguish conceptually from density matrices, as we'll explain later. As part of our interpretation of quantum theory, we fill in this missing ingredient by providing an explicit translation of density-matrix dynamics into dynamical mappings for ontic and epistemic states, and in such a way that we avoid macroscopic instabilities (such as eigenstate swaps, to be defined in Section III B 3) that have presented problems for other interpretations.

Returning to our classical story, the assumed independence of the conditional probabilities $p(q'; t' | (q_1; t_1), \dots)$ from the system's evolving epistemic state p also implies that any conditional probabilities depending on ontic states at *additional* initial times $\tilde{t}_1, \tilde{t}_2, \dots \notin \{t_1, t_2, \dots\}$ must (if they exist) be equal to $p(q'; t' | (q_1; t_1), \dots)$,

$$p(q'; t' | (q_1; t_1), \dots, (\tilde{q}_1, \tilde{t}_1), (\tilde{q}_2, \tilde{t}_2), \dots) = p(q'; t' | (q_1; t_1), \dots), \quad (18)$$

because otherwise the dynamical Bayesian propagation formula

$$p(q'; t' | (q_1; t_1), \dots) = \sum_{\tilde{q}_1, \dots} p(q'; t' | (q_1; t_1), \dots, (\tilde{q}_1, \tilde{t}_1), (\tilde{q}_2, \tilde{t}_2), \dots) p(\tilde{q}_1; \tilde{t}_1) p(\tilde{q}_2; \tilde{t}_2) \dots$$

would imply that the conditional probabilities $p(q'; t' | (q_1; t_1), \dots)$ have a disallowed implicit dependence on the epistemic states $p(\tilde{q}_1; \tilde{t}_1), p(\tilde{q}_2; \tilde{t}_2), \dots$ of the system at the additional times $\tilde{t}_1, \tilde{t}_2, \dots$. As an example, in the case of a system whose dynamics is Markovian [231], meaning first order, the mapping $p(\cdot; t' | \cdot; t)$ defined in (16) requires the specification of just a *single* initial time t , and any additional mappings (16) involving multiple initial times $t \geq t_1 \geq t_2 \geq \dots$ must be equal to the single-initial-time mapping $p(\cdot; t' | \cdot; t)$:

$$p(\cdot; t' | (\cdot; t), (\cdot; t_1), (\cdot; t_2), \dots) = p(\cdot; t' | \cdot; t). \quad (19)$$

In other words, a system whose dynamics is Markovian retains no “memory” of its ontic states prior to its most recent ontic state, apart from whatever memory is directly encoded in that most recent ontic state.

We call a system closed in the idealized case in which the system does not interact with or exchange information with its environment. Unless we allow for literal destruction of information, closed classical systems are always governed *fundamentally* by deterministic dynamics. (Subtleties arise for chaotic systems, to be discussed shortly.) More realistically, however, systems are never exactly closed and are instead said to be open, in which case information can “leak out.” The existence of dynamics for an open system is a delicate question because conditional probabilities inherited from an enclosing parent system may contain a nontrivial dependence on the parent system's own epistemic state, but any such open-system dynamics, if it exists, is generically stochastic.

8. Classical Systems with Continuous Configuration Spaces

Many familiar classical systems have *continuous* configuration spaces that we can regard as manifolds of some dimension $N \geq 1$. When working with a classical system of this kind, we can parameterize the ontic states $q = (q_1, \dots, q_N)$ of the configuration space \mathcal{C} in terms of N continuously valued degrees of freedom q_α ($\alpha = 1, \dots, N$), which play the role of coordinates for the configuration-space manifold. It is then more natural to describe the system's epistemic states in terms of probability *densities* $\rho(q) \equiv \rho(q_1, \dots, q_N) \equiv dp(q)/d^Nq$ rather than in terms of probabilities $p(q)$ *per se*.

If we can treat time as being a continuous parameter in our description of a classical system having a continuous configuration space, then, purely at the level of kinematics, the ontic states $q(t) = (q_1(t), \dots, q_N(t))$ of the system at infinitesimally separated times $t, t + dt, t + 2dt, \dots$ are independent kinematical variables, and thus the instantaneous coordinate values $q_\alpha(t)$, instantaneous velocities $\dot{q}_\alpha(t) \equiv dq_\alpha(t)/dt \equiv (q_\alpha(t + dt) - q_\alpha(t))/dt$, instantaneous accelerations $\ddot{q}_\alpha(t) \equiv d^2q_\alpha(t)/dt^2$, and so forth, are all likewise independent kinematical variables. We are therefore free to extend the notion of the system's epistemic state to describe the joint probabilities $p(q, \dot{q}, \ddot{q}, \dots)$ that the system's ontic state is $q \equiv (q_1, \dots, q_N)$, that its instantaneous velocities have values given by $\dot{q} \equiv (\dot{q}_1, \dots, \dot{q}_N)$, and so forth.

This generalization is particularly convenient when examining a system whose dynamical mapping (16) is *second order* in time—meaning that it involves initial ontic states q at a *pair* of times—and calls for infinitesimally separated pairs of times—namely, t and $t + dt$. We can then equivalently express the dynamics in terms of a mapping involving the system's ontic state $q = (q_1, \dots, q_N)$ and its instantaneous velocities $\dot{q} = (\dot{q}_1, \dots, \dot{q}_N)$ at the *single* initial time t . Hence, if we now think of the ordered pair (q, \dot{q}) as describing a point in a generalized kind of configuration space (mathematically speaking, the tangent bundle of the configuration-space manifold), then we can treat the dynamics as though it were Markovian—that is, depending only on data at a *single* initial time. When the dynamics is, moreover, deterministic, and we can express the dynamical mapping from initial data $(q, \dot{q}; t)$ to final data $(q', \dot{q}'; t')$ as a collection of equations, known as the system's equations of motion, then we enter the familiar terrain of textbook classical physics, with its language of Lagrangians $L(q, \dot{q})$, action functionals $S[q] \equiv \int dt L(q(t), \dot{q}(t); t)$, canonical momenta p_α , phase spaces (q, p) (mathematically speaking, the *cotangent* bundle of the system's configuration-space manifold), Hamiltonians $H(q, p)$, and Poisson brackets

$$\{f, g\}_{\text{PB}} \equiv \sum_{\alpha} \left[\frac{\partial f}{\partial q_{\alpha}} \frac{\partial g}{\partial p_{\alpha}} - \frac{\partial g}{\partial q_{\alpha}} \frac{\partial f}{\partial p_{\alpha}} \right]. \quad (20)$$

For classical systems of this type, we can formulate the second-order equations of motion as the so-called canonical equations of motion

$$\dot{q}_{\alpha} = \frac{\partial H}{\partial p_{\alpha}} = -\{H, q_{\alpha}\}_{\text{PB}}, \quad \dot{p}_{\alpha} = -\frac{\partial H}{\partial q_{\alpha}} = -\{H, p_{\alpha}\}_{\text{PB}}. \quad (21)$$

For consistency, the system's generalized epistemic state $\rho(q, p)$ on its phase space must then satisfy the classical Liouville equation

$$\frac{\partial \rho}{\partial t} = \{H, \rho\}_{\text{PB}}. \quad (22)$$

A familiar corollary is the Liouville theorem: The *total* time derivative of the probability distribution $\rho(q(t), p(t); t)$ vanishes on any trajectory that solves the equations of motion,

$$\frac{d\rho}{dt} = 0, \quad (23)$$

meaning that the ensemble of phase-space trajectories described by $\rho(q(t), p(t); t)$ behaves like an incompressible fluid.

9. Chaos

Classical systems with continuous configuration spaces can exhibit an important kind of dynamics known as chaos. While chaotic dynamics is *formally* deterministic according to the classification scheme described earlier—initial ontic states are mapped to unique final ontic states with unit probability—the predicted values of final ontic states are exponentially sensitive to minute changes in the initial data, and thus the decimal precision needed to describe the




Figure 4. A Venn diagram representing a quantum system’s configuration space, with an orthonormal basis for the system’s Hilbert space corresponding to a partitioning $\{\Psi_i\}_i$ by mutually exclusive ontic states. An ontic state Ψ' not orthogonal to all the members of the orthonormal basis is also displayed, and is not mutually exclusive with all the ontic states in $\{\Psi_i\}_i$.

system’s final epistemic state is exponentially greater than the decimal precision needed to specify the system’s initial data. Indeed, our relative predictive power, which we could define as being the ratio of our output precision to our input precision, typically goes to zero as we try to approach the idealized limit of infinitely sharp input precision. Because any realistic recording device is limited to a fixed and finite decimal precision, chaotic dynamics is therefore *effectively* stochastic even if the system of interest is closed off from its larger environment.

C. Quantum Theories

1. Quantum Kinematics

The picture that our interpretation paints for quantum theory is surprisingly similar to the classical story, the key kinematical difference being that a quantum system’s configuration space is “too large” to be a *single* sample space of mutually exclusive states. More precisely, our interpretation asserts that the configuration space of a quantum system corresponds (up to meaningless overall normalization factors) to a *vector space* \mathcal{H} ,¹³ called the system’s Hilbert space, for which the notion of mutual exclusivity between states is defined by the vanishing of their inner product, much in the same way that nondegenerate eigenstates of the Hermitian operators representing observables in the traditional formulation of quantum theory are always orthogonal.

As a consequence, every quantum system exhibits a *continuous infinity* of distinct sample spaces, each corresponding to a particular orthonormal basis for the system’s Hilbert space. It follows that an epistemic state makes *logically rigorous* sense only as a probability distribution over an orthonormal basis, although, just as we saw in Section II B 3 for the classical case, it is sometimes useful to consider *formal* subjective epistemic states over possibilities that are not mutually exclusive.

Just as a helpful visual analogy, one can loosely imagine all of a system’s allowed ontic states as corresponding to finite-size regions collectively covering an abstract Venn diagram. Then any orthonormal basis for the system’s Hilbert space corresponds to a partition of the Venn diagram into disjoint regions, whereas any two non-orthogonal ontic states describe overlapping regions. (See Figure 4.) Keep in mind, however, that this picture is entirely metaphorical: We regard quantum states as being irreducible sample-space elements, and thus we are not *literally* identifying the Venn diagram’s *individual points* as elements of the system’s sample space.

This new feature of quantum theory opens up a possibility not available to classical systems, namely, that a system’s sample space is fundamentally contextual, changing over time from one orthonormal basis to another as the system

¹³ It is not a goal of this paper to provide deeper principles underlying this fact, which we take on as an axiom; for a detailed discussion of the complex vector-space structure of Hilbert spaces in quantum theory, and attempts to derive it from more basic principles and generalize it in new directions, see Section VI E 1. It’s interesting to note the historical parallel between the way that quantum theory modifies classical ontology by replacing classical configuration spaces (even those that are discrete) with smoothly interpolating continuous vector spaces, and, similarly, the way that probability theory itself modified epistemology centuries ago by replacing binary true/false assignments with a smoothly interpolating continuum running from 0 (false) to 1 (true).




Figure 5. A schematic depiction of our postulated relationship between a system’s density matrix $\hat{\rho}$ and its associated epistemic state $\{(p_1, \Psi_1), (p_2, \Psi_2), (p_3, \Psi_3), \dots\}$. The latter consists of epistemic probabilities p_1, p_2, p_3, \dots (the eigenvalues of the density matrix) and possible ontic states $\Psi_1, \Psi_2, \Psi_3, \dots$ (represented by the eigenstates of the density matrix), where one of those *possible* ontic states (in this example, Ψ_2) is the system’s *actual* ontic state.

interacts with other systems, such as with measurement devices or the system’s larger environment. That is, in contrast to classical physics, a quantum system’s configuration space—essentially, the system’s Hilbert space, up to meaningless phase factors—is not in one-to-one correspondence with the system’s sample space—which corresponds to a particular orthonormal basis for the Hilbert space—and, indeed, interactions with the outside world can cause the system to change from one sample space to another sample space incompatible with the first but nonetheless lying in the *same* configuration space.

Hence, our interpretation of quantum theory builds in the notion of quantum contextuality [287–289] from the very beginning, closely in keeping with the Kochen-Specker theorem [202] described in Appendix 2. Many interpretations of quantum theory, including our own, take advantage of this general feature of the theory in order to resolve the measurement problem through the well-known quantum phenomenon of decoherence [58, 67, 196, 197, 269, 270, 349].

2. Density Matrices

As we will review in Section III B, the formalism of quantum theory naturally furnishes a mathematical object $\hat{\rho}$, called a density matrix [210, 321, 323], that our interpretation places in a central role for defining a system’s epistemic state and its underlying ontology, especially in the context of relating parent systems with their subsystems in the presence of quantum entanglement. A density matrix, which we regard as being related to but conceptually distinct from the system’s epistemic state, is a Hermitian operator on the system’s Hilbert space whose eigenvalues are all nonnegative and sum to unity.

For the case of a density matrix $\hat{\rho}$ whose impurity arises entirely due to external quantum entanglement—that is, for a system in a totally improper mixture—our interpretation identifies the eigenvalues p_1, p_2, p_3, \dots of the density matrix as collectively describing a probability distribution—namely, the system’s epistemic state $\{(p_1, \Psi_1), (p_2, \Psi_2), (p_3, \Psi_3), \dots\}$ —over a set of possible ontic states $\Psi_1, \Psi_2, \Psi_3, \dots$ that are represented by the orthonormal eigenstates $|\Psi_1\rangle, |\Psi_2\rangle, |\Psi_3\rangle, \dots$ of the density matrix, where one of those *possible* ontic states is the system’s *actual* ontic state. (See Figure 5.) In the very rough sense in which observables in the traditional formulation of quantum theory correspond to Hermitian operators whose eigenvalues represent measurement outcomes and whose eigenstates represent their associated state vectors, our interpretation therefore identifies a system’s density matrix as the Hermitian operator corresponding to the system’s “probability observable.”

3. Quantum Dynamics

As is well known, density matrices describing *closed* systems evolve in time according to a *quantum* version of the Liouville equation, known as the Von Neumann equation, which bears a striking resemblance¹⁴ to the classical

¹⁴ This resemblance becomes even closer after an appropriate change to complex phase-space variables, as we discuss in Section VI E 2.

Liouville equation (22) and is given by

$$\frac{\partial \hat{\rho}}{\partial t} = -\frac{i}{\hbar} [\hat{H}, \hat{\rho}]. \quad (24)$$

This equation for the time evolution of density matrices has been generalized [67, 196, 219, 275] to cases involving open systems losing information to their larger environments, but a missing ingredient has been a general specification of the dynamics that governs the system at the level of its *ontic and epistemic states*, where, again, we distinguish density matrices $\hat{\rho}$ from epistemic states $\{(p_1, \Psi_1), (p_2, \Psi_2), (p_3, \Psi_3), \dots\}$.

Our interpretation supplies this key ingredient by translating the evolution law governing density matrices into a Markovian dynamical mapping in the sense of (16) that satisfies the key consistency requirement (17) relating the dynamics of ontic states with the dynamics of epistemic states. This mapping belongs to a new class of conditional probabilities that both dynamically relate the ontic states of a single quantum system at initial and final times and also define the kinematical relationship between the ontic states of parent systems with those of their subsystems. In addition, the mapping serves as a kind of smoothness condition that sews together ontic states in a sensible manner and avoids some of the ontic-level instabilities that arise in other interpretations of quantum theory.

III. THE MINIMAL MODAL INTERPRETATION

A. Basic Ingredients

Having presented the motivation, conceptual foundations, and rough outlines of our interpretation of quantum theory in the preceding sections, we now commence our precise exposition. The axiomatic content of our interpretation, which we summarize in detail in Section VIA, is intended to be minimal and consists of

1. definitions provided in this section of what we mean by ontic and epistemic states in quantum theory, as well as a dichotomy between subjective (“classical”) and objective epistemic states;
2. a postulated relationship detailed in Section IIIB between objective epistemic states and density matrices, and, as a consequence, a relationship also between ontic states and density matrices;
3. a well-known rule (the partial-trace operation) described in Section IIIC for calculating density matrices for subsystems given density matrices for their parent systems, and, as an immediate corollary, also a rule for relating the objective epistemic states of parent systems with those of their subsystems;
4. and, finally, in Section IIID, a general formula defining a class of quantum conditional probabilities that serve both for *kinematically* relating the ontic states of parent systems with those of their subsystems, as well as for *dynamically* relating ontic states to each other over time and objective epistemic states to each other over time.

Note that apart from the references to density matrices, all these rules and metaphysical entities have necessary (if often implicit) counterparts in the description of classical systems detailed in Section IIB: For classical systems as well as for quantum systems, we must define what we mean by ontic and epistemic states, establish relationships between epistemic states of parent systems and of their subsystems as well as between ontic states of parent systems and of their subsystems (the latter relationship being trivial in the classical case, but crucial for making sense of measurements in the quantum case), and dynamical relationships between ontic states over time and between epistemic states over time.

There are two chief differences between the classical and quantum cases, however, the first difference being that in the quantum case, we need to invoke density matrices, ultimately because quantum entanglement leads to the existence of fundamentally objective epistemic states and makes density matrices a crucial link between parent systems and their subsystems. The second difference is that we can’t afford to take as many of the other ingredients for granted in the quantum case as we could in the classical case—quantum theory forces us to be much more explicit about our assumptions.

We claim that all the familiar features of quantum theory follow as consequences when we take all these basic principles to their logical conclusions, including the Born rule for computing empirical outcome probabilities and also consistency with the various no-go theorems, as we will discuss starting in Section IV. We treat several additional no-go theorems in Appendix 2.

1. Ontic States

Quantum physics first enters our story with the condition that the distinct ontic states Ψ_i that make up our system’s configuration space are each represented by a particular unit vector $|\Psi_i\rangle$ (up to overall phase) in a complex inner-product vector space \mathcal{H} , called the system’s Hilbert space:

$$\boxed{\Psi_i \leftrightarrow |\Psi_i\rangle \in \mathcal{H} \text{ (up to overall phase)}}. \quad (25)$$

That is, we identify the system’s configuration space as the quotient space \mathcal{H}/\sim , where \sim represents equivalence of vectors in the Hilbert space \mathcal{H} up to overall normalization and phase. Note that we regard (25) as merely a *correspondence* between an ontic object and a mathematical object, and so we will not use the terms “ontic state” and “state vector” synonymously in this paper, in contrast to the practice of some authors.

2. Epistemic States

As part of the definition of our interpretation, and in close parallel with the classical case presented in Section II B, we postulate that a quantum system’s evolving epistemic state consists at each moment in time t of a collection $\{(p_i, \Psi_i)\}_i$ of ordered pairs that for each i identify one of the system’s *possible* ontic states $\Psi_i = \Psi_i(t)$ together with the corresponding epistemic probability $p_i = p_i(t) \in [0, 1]$ for which Ψ_i is the system’s *actual* ontic state at the current time t .¹⁵ As in the classical case, an epistemic state is called pure in the idealized case in which its probability distribution is trivial—that is, if one epistemic probability p_i is equal to unity and all the others vanish—and is called mixed if the probability distribution is nontrivial.

Note that a system’s epistemic probabilities p_i at a particular moment in time should not generally be confused with the empirical outcome probabilities that are used to quantify the results of experiments and that are computed via the Born rule. Whereas epistemic probabilities describe a system’s *present* state of affairs, we will ultimately show that empirical outcome probabilities can be identified with a measurement device’s predicted *future* epistemic probabilities conditioned on the hypothetical assumption that the device will have performed a particular measurement on a given subject system.

3. Modal Interpretations and Minimalism

Our use of the modifiers “possible and “actual,” together known formally as modalities, identifies our interpretation of quantum theory as belonging to the general class of modal interpretations originally introduced by Krips in 1969 [205–207] and then independently developed by van Fraassen (whose early formulations involved the fusion of modal logic [86, 216, 340] with quantum logic [57, 146]), Dieks, Vermaas, and others [21, 28, 29, 70, 75, 224, 225, 313, 314, 316, 317].

The modal interpretations are now understood to encompass a very large set of interpretations of quantum theory, including most interpretations that fall between the “many worlds” of the Everett-DeWitt approach and the “no worlds” of the instrumentalist approaches. Generally speaking, in a modal interpretation, one singles out some preferred basis for each system’s Hilbert space and then regards the elements of that basis as being the system’s *possible* ontic states—one of which is the system’s *actual* ontic state—much in keeping with how we think conceptually about classical probability distributions. For example, as we will explain more fully in Section III A 6, and as emphasized in [316], the traditional de Broglie-Bohm pilot-wave interpretation can be regarded as being a special kind of modal interpretation in which the preferred basis is permanently fixed *for all* systems at a universal choice. Other modal interpretations, such as our own, instead allow the preferred basis for a given system to change—in our case by choosing the preferred basis to be the evolving eigenbasis of that system’s density matrix. However, we claim that no existing modal interpretation captures the one that we introduce in this paper.

A central guiding principle of our interpretation is *minimalism*: By intention, we make no changes to the way quantum theory should be used in practice, and the only requirements that we impose are those that are absolutely mandated by the need for our notions of ontology and epistemology to account for the observable predictions of

¹⁵ Throughout this paper, we will always work in the so-called Schrödinger picture, in which a system’s time evolution is carried by states and not by operators. Indeed, because most of the systems that we consider in this paper aren’t dynamically closed and thus don’t evolve according to unitary or even linear dynamics—unitary dynamics usually being a good approximation only for systems that are microscopic and therefore easy to isolate from their environments and linear dynamics being a good approximation only for systems whose environmental correlations wash away sufficiently quickly with time—none of the other familiar pictures (such as the Heisenberg picture) will generally be well-defined anyway.

quantum theory—no more, no less. We are also metaphysically conservative in the sense that we include only ingredients that are either already a part of the standard formalism of quantum theory or that have counterparts in classical physics. We therefore naturally call our interpretation the *minimal modal interpretation* of quantum theory.

Our reasons for this minimalism go beyond philosophically satisfying notions of axiomatic simplicity and parsimony. Based on an abundance of historical examples, we know that when trying to add an ontology to quantum theory, including (even implicitly) more features than are strictly necessary is not just a metaphysical extravagance, but also usually leads to trouble. This trouble may take the form of unacceptable ontological instabilities, conflicts with various no-go theorems, an uncontrollable profusion of “epicycles,” or just an overwhelming structural ornateness.

By contrast, we will show that our minimal modal interpretation evades a recent no-go theorem [56, 99, 106, 241, 242] asserting that other members of the class of modal interpretations are incompatible with Lorentz invariance at an ontological level. We will also argue that the same minimalism makes it possible to provide first-principles derivations of familiar aspects of quantum theory, including the Born rule (and possible corrections to it) for computing the empirical outcome probabilities that emerge from experiments.

4. Hidden Variables and the Irreducibility of Ontic States

To the extent that our interpretation of quantum theory involves hidden variables, the actual ontic states underlying the epistemic states of systems play that role. However, one could also argue that calling them hidden variables is just a matter of semantics because they are on the same metaphysical footing as both the *traditional* notion of quantum states as well as the actual ontic states of *classical* systems.

In any event, it is important to note that our interpretation includes *no other* hidden variables: Just as in the classical case, we regard ontic states as being *irreducible* objects, and, in keeping with this interpretation, we *do not* regard a system’s ontic state itself as being an epistemic probability distribution—much less a “pilot wave”—over a set of more basic hidden variables. In a rough sense, our interpretation *unifies* the de Broglie-Bohm interpretation’s pilot wave and hidden variables into a single ontological entity that we call an ontic state.

In particular, we do not attach an epistemic probability interpretation to the *components* of a vector representing a system’s ontic state, nor do we assume *a priori* the Born rule, which we will ultimately *derive* as a means of computing empirical outcome probabilities. Otherwise, we would need to introduce an unnecessary *additional* level of probabilities into our interpretation and thereby reduce its axiomatic parsimony and explanatory power. In a sense that we’ll make clear later, the components of a vector in a system’s Hilbert space represent “potentialities” [5] that, after certain environmental interactions or an appropriate measurement process, can eventually *become* the eigenvalues of the system’s density matrix and thus actualized as true epistemic probabilities.

Via the phenomenon of environmental decoherence, our interpretation ensures that the evolving ontic state of a sufficiently macroscopic system—with significant energy and in contact with a larger environment—is highly likely to be represented by a temporal sequence of state vectors that presumably approximate coherent states, as described in Section [VIE 2](#), and whose labels evolve in time according to recognizable semiclassical equations of motion. For microscopic, isolated systems, by contrast, we simply accept that the ontic state vector may not usually have an intuitively familiar classical description.

5. Constraints from the Kolmogorov Axioms

The logically rigid *Kolmogorov axioms* [203], which are satisfied by any well-defined epistemic probability distribution, require that our epistemic probabilities $p_i = p_i(t)$ at any moment in time t can never sum to a value greater than unity. In the idealized limit in which the system cannot decay, we can require the stronger condition that the epistemic probabilities always sum exactly to unity, $\sum_i p_i = 1$. By contrast, for systems having $\sum_i p_i < 1$, we naturally interpret the discrepancy $(1 - \sum_i p_i) \in [0, 1]$ as being the system’s probability of no longer existing, as we will explain in greater detail in Section [III C 6](#).

In keeping with the Kolmogorov axioms, we additionally require that the system’s corresponding possible ontic states $\Psi_i = \Psi_i(t)$ at any single moment in time t are mutually exclusive. We translate this requirement into quantum language as the condition that no member $|\Psi_i\rangle$ of the system’s set of possible ontic state vectors at the time t can be expressed nontrivially as a superposition involving any of the others; none is a “blend” involving any of the others. This condition is equivalent to requiring that the system’s possible ontic state vectors at the time t must all be mutually orthogonal,

$$\langle \Psi_i | \Psi_j \rangle = 0 \text{ for } i \neq j, \tag{26}$$

much in the familiar way that eigenstates corresponding to nondegenerate eigenvalues of the Hermitian operators representing observables in the traditional formulation of quantum theory are always orthogonal. Hence, *any* orthonormal basis for the system’s Hilbert space defines a distinct allowed sample space admitting logically rigorous probability distributions. However, sets of non-orthogonal state vectors (such as the system’s Hilbert space *as a whole*) do not, strictly speaking, admit logically rigorous probability distributions, although, in keeping with our discussion in Section II B 3, it is occasionally useful to define *formal, subjective* probability distributions even in those cases.

6. The de Broglie-Bohm Pilot-Wave Interpretation of Quantum Theory

In contrast to regarding *any* orthonormal basis as a potentially valid sample space, the so-called “fixed” modal interpretations [70, 316]—which include the traditional formulation of the de Broglie-Bohm pilot-wave interpretation as a special case—require that the sample space remains permanently fixed forever at some universally preferred choice, which is then a matter for the interpretation to try to justify once and for all. A generic state vector is then regarded as both being an epistemic probability distribution over this fixed sample space of ontic states and also as being an ontological entity in its own right, namely, a physical “pilot wave” that guides the evolution of the system’s hidden ontic state through this fixed sample space while being conceptually distinct from that hidden ontic state.

In its original nonrelativistic formulation, the de Broglie-Bohm interpretation [59, 60, 63, 91] takes its fixed orthonormal basis to be the position eigenbasis of point particles, and offers what might appear to be a straightforward solution to questions concerning measurements in quantum theory [233]. Unfortunately, in the *relativistic* regime, spatial position ceases to exist as an orthonormal basis: The Hilbert-space inner product of two would-be particle-position eigenstates is nonvanishing, although it is exponentially suppressed in the particle’s Compton wavelength $\lambda_{\text{Compton}} = h/mc$ and therefore indeed goes to zero in the nonrelativistic limit $c \rightarrow \infty$ [257]. (Preserving causality in the presence of this spontaneous superluminal propagation famously necessitates the existence of antiparticles [117, 257, 330].)

Attempts to recast the de Broglie-Bohm interpretation in a relativistic context require giving up the elegance and axiomatic parsimony of the interpretation’s nonrelativistic formulation. Some approaches involve replacing the nonrelativistic position eigenbasis with bases that remain orthonormal in the relativistic regime, such as the field-amplitude eigenbasis for bosonic fields [60, 62]. However, fermionic fields present a serious problem, because their “field amplitudes” are inherently non-classical, anticommuting, nilpotent Grassmann numbers; although Grassmann numbers provide a convenient *formal* device for expressing fermionic scattering amplitudes in terms of Berezin path integrals [55, 257], taking Grassmann numbers seriously as physical ingredients in quantum Hilbert spaces would lead to nonsense “probabilities” that are not ordinary numbers. Trying instead to choose the fermion-number eigenbasis [45] brings back the question of ill-definiteness of spatial position, leading some advocates to drop orthonormal bases altogether in favor of a positive-operator-valued measure (POVM) [104, 105, 293], whose elements are *not* mutually exclusive. Even then, the claimed *universal* preferred POVM singles out a correspondingly preferred *global* Lorentz frame, and it’s unclear how Nature singles out any such frame. Relativity therefore casts doubt on hopes of finding a safe choice of fixed orthonormal basis or POVM to provide the de Broglie-Bohm interpretation with its foundation: At best, there’s no canonical choice to fix once and for all, and, at worst, there is no choice that’s consistent or sensible.

Whether or not the de Broglie-Bohm interpretation’s proponents ultimately find a satisfactory fixed orthonormal basis or POVM to define their sample space,¹⁶ the interpretation still runs into other troubles as well, including its inability to accommodate the non-classical changes of particle spectrum that can arise in quantum field theories and the non-classical changes in configuration space that can emerge from quantum dualities, both of which play a central role in much of modern physics.¹⁷

The de Broglie-Bohm interpretation also suffers from a somewhat more metaphysical problem: Because a pilot wave has an ontological existence *over and above* that of the hidden ontic state of its corresponding physical system, we run into the well-known difficulty [68] of making sense of the ontological status of all its many branches. Indeed, the pilot wave’s branches behave precisely as the “many worlds” of the Everett-DeWitt interpretation of quantum theory in all their complexity, despite the twin assertions that only one of those branches is supposedly “occupied” by the system’s hidden ontic state and that the rest of the branches are “empty worlds” filled with ghostly people living out presumably ghostly lives.

¹⁶ Even for a nonrelativistic system, it’s not clear why the coordinate basis should be favored in the de Broglie-Bohm interpretation, as opposed to, say, an orthonormal basis approximating the system’s far more classical-looking coherent states [152, 272], as defined in Section VI E 2.

¹⁷ Non-classical changes in particle spectrum in quantum field theories include prosaic examples like transitions from tachyonic particle modes to massive radial “Higgs” modes and massless Nambu-Goldstone modes after the spontaneous symmetry breaking of a continuous symmetry [155, 156, 243], as well as more exotic examples like Skyrmions [256, 282–284] and bosonization [88, 229]. For a wide-ranging survey of dualities, see, for example, [259]. Prominent examples of dualities include generalizations of electric-magnetic duality in certain supersymmetric gauge theories [240, 276], holographic dualities like the AdS/CFT correspondence between gauge theories and theories of quantum gravity in higher dimensions [10, 69, 160, 161, 228, 343], and various dualities that play key roles in string theory [151, 278, 291]. In particular, the AdS/CFT correspondence as well as string theory suggest that spacetime can undergo radical changes in structure [26, 341, 342] via intermediary non-geometric phases, and even perhaps that the “space” in “spacetime” is itself an emergent property of more primitive ingredients [108, 277].

7. Subjective Uncertainty and Proper Mixtures

As we explained in Section II B, we use nontrivial epistemic states $\{(p(q), q)\}_q$ in classical physics in order to account for subjective uncertainty about a classical system's actual ontic state at a given moment in time. Similarly, one way that epistemic states $\{(p_i, \Psi_i)\}_i$ can arise in quantum theory is when we have subjective uncertainty about a quantum system's actual ontic state at a particular moment in time, in which case we call the system's epistemic state a proper mixture.

Strictly speaking, the requirements of a logically rigorous probability distribution require that the possible ontic states Ψ_i that make up a proper mixture must be mutually exclusive and hence correspond to mutually orthogonal state vectors $|\Psi_i\rangle$ in accordance with (26). However, after deriving the Born rule later in this paper, we will describe in Section IV C how to accommodate contexts in which it is useful to relax this mutual exclusivity, just as we explained in Section II B 3 that we could consider *formal* probability distributions over non-exclusive possibilities in classical physics.

There is no real controversy or dispute over the metaphysical meaning of proper mixtures, at least in the sense that there is wide acceptance for regarding a proper mixture as a prosaic, subjective probability distribution over different possible underlying state vectors, one of which is the system's actual state vector. Nor are density matrices even strictly necessary when dealing with proper mixtures, as we could always work instead at the level of the individual state vectors that make up the proper mixture and then overlay our final results with the subjective probability distribution by hand. We will therefore put proper mixtures entirely aside for most of this paper, returning to them in Section IV C only after deriving the Born rule.

8. Objective Uncertainty and Improper Mixtures

Quantum theory also features what our minimal modal interpretation regards as being an *objective* kind of uncertainty that has no obvious classical counterpart, and in this case we will *always* insist upon logically rigorous epistemic states involving mutually exclusive possible ontic states. To see where this objective new form of uncertainty comes from, and to make clear the importance of the relationship between parent systems and subsystems in the context of our interpretation of quantum theory, we step back for a moment and compare the notions of parent systems and subsystems in classical physics and in quantum physics.

Classically, a system C with configuration space \mathcal{C}_C is said to be a parent or composite system consisting of two subsystems A and B with respective configuration spaces \mathcal{C}_A and \mathcal{C}_B if the configuration space of system C is expressible as the Cartesian product $\mathcal{C}_C = \mathcal{C}_A \times \mathcal{C}_B = \{(a, b) | a \in \mathcal{C}_A, b \in \mathcal{C}_B\}$, meaning that each element $c = (a, b)$ of \mathcal{C}_C is an ordered pair that identifies a specific element a of \mathcal{C}_A and a specific element b of \mathcal{C}_B . We then naturally denote the parent system by $C = A + B$, and we have the simple relation that $N_{A+B} = N_A N_B$, even in the possible case in which some of these configuration spaces have infinitely many elements. Note, of course, that either subsystem A or B (or both) could well be a parent system to subsystems that are even more elementary.

The quantum case is much more subtle because the Cartesian product of a pair of nontrivial Hilbert spaces is *not* another Hilbert space. We instead identify a given quantum system C having a Hilbert space \mathcal{H}_C as being the parent system of two quantum subsystems A and B with respective Hilbert spaces \mathcal{H}_A and \mathcal{H}_B if there exists an orthonormal basis for \mathcal{H}_C whose elements are of the tensor-product form $|a, b\rangle = |a\rangle \otimes |b\rangle$, where the sets of vectors $|a\rangle$ and $|b\rangle$ respectively constitute orthonormal bases for \mathcal{H}_A and \mathcal{H}_B . Then, by construction, every vector in \mathcal{H}_C consists of some linear combination of the vectors $|a, b\rangle$. We express this fact by writing the parent system's Hilbert space as the tensor product $\mathcal{H}_C = \mathcal{H}_A \otimes \mathcal{H}_B$, and we denote the parent system, as we did in the classical case, by $C = A + B$. Observe that the dimensions of these various Hilbert spaces (whether or not they are all of finite dimension) then satisfy $\dim \mathcal{H}_{A+B} = (\dim \mathcal{H}_A)(\dim \mathcal{H}_B)$.

Given this background, consider a parent quantum system $A + B$ consisting of two subsystems A and B and whose ontic state Ψ_{A+B} is described by a so-called entangled superposed state vector of the form

$$|\Psi_{A+B}\rangle = \alpha |\Psi_{A,1}\rangle |\Psi_{B,1}\rangle + \beta |\Psi_{A,2}\rangle |\Psi_{B,2}\rangle, \quad \alpha, \beta \in \mathbb{C}, \quad |\alpha|^2 + |\beta|^2 = 1. \quad (27)$$

What is the ontic state of the subsystem A ? What is the ontic state of the subsystem B ?

In the Copenhagen interpretation of quantum theory, these two questions do not possess well-defined answers, a problematic state of affairs because all realistic systems are at least slightly entangled with other systems. Instrumentalist interpretations do not regard even the questions themselves as being sensible or meaningful.

9. *The Everett-DeWitt Many-Worlds Interpretation of Quantum Theory*

By contrast, the Everett-DeWitt many-worlds interpretation of quantum theory axiomatically postulates that (27) implies the existence of two simultaneous “worlds” or “realities” or “branches,” at least if A and B are *sufficiently macroscopic* in some vaguely defined sense: In one world, associated in some sense with a probability $|\alpha|^2$, the ontic state of A is $\Psi_{A,1}$ and the ontic state of B is $\Psi_{B,1}$, whereas in the other world, associated in a similar sense with a probability $|\beta|^2$, the ontic state of A is $\Psi_{A,2}$ and the ontic state of B is $\Psi_{B,2}$. Note that attaching a many-worlds ontology to mathematical vectors in the first place is *itself* a nontrivial, affirmative axiom; indeed, the instrumentalist approach, which we describe in Appendix 1e, does not postulate *any* ontological status for state vectors, but regards them merely as abstract mathematical tools for predicting the statistical outcomes of experiments.

Of course, because we can always choose from a *continuously infinite* set of different orthonormal bases for the Hilbert space of the parent system $A + B$, as we discussed in Section II C 1, the two “worlds” associated with the entangled superposed state vector (27) are radically non-unique regardless of the physical sizes of A and B and even for a fixed parent-system state vector $|\Psi_{A+B}\rangle$; moreover, the different possible world-bases are not generally related to one another in a manner that can be conceptualized classically, and preserving manifest locality in the many-worlds interpretation generically requires switching from one world-basis to another, incompatible world-basis as a function of time. These issues, known collectively as the preferred-basis problem, imply a breakdown in the popular portrayal of the many-worlds interpretation as describing unfolding reality in terms of a well-defined forking structure in the manner of Borges’ *Garden of Forking Paths* [64] or Lewis’s modal realism [217, 218], and are unsolvable without introducing additional axiomatic ingredients into the interpretation. In particular, without additional postulates, one cannot evade the preferred-basis problem merely by appealing to environment-induced decoherence, because the many-worlds interpretation traditionally assumes as yet one more axiom that “the universe as a whole”—which determines the single branch-set shared by all systems in Nature—is described by an objectively always-pure state that never undergoes decoherence and therefore doesn’t possess a canonical preferred basis.¹⁸

Another conundrum of the many-worlds interpretation is the difficulty in making sense of the components α and β in (27) in terms of *probabilities* when both worlds are *deterministically* and *simultaneously* realized,¹⁹ especially given the fundamental logical obstruction to deriving probabilistic conclusions from deterministic assumptions and in light of the preferred-basis problem and the continuously infinite non-uniqueness of the choice of world-basis. (It’s also worth mentioning the possible conceptual difficulties that arise if, say, $|\Psi_{A,1}\rangle$ and $|\Psi_{A,2}\rangle$ are not orthogonal and thus not mutually exclusive, an issue that we’ll address in the context of our interpretation in Section III C 5.) Putting aside the preferred-basis problem, a common approach [9, 114] is to study a pure state defined in terms of a “limit” of many identical copies of a given measurement set-up, and then argue that “maverick branches”—meaning terms in the final-state superposition whose outcome frequency ratios deviate significantly from the Born rule—have arbitrarily small Hilbert-space amplitudes and thus can safely be ignored. However, infinite limits are always rigorously defined in terms of increasing but finite sequences, and for any finite number of copies of a given measurement set-up, maverick branches have nonzero amplitudes and outnumber branches with better-behaved frequency ratios. Asserting that the smallness of their amplitudes makes maverick branches “unlikely” therefore begs the question by implicitly assuming the very probability interpretation to be derived [72, 79, 230], and is akin to the sort of circular reasoning inherent in all attempts to use the law of large numbers to turn frequentism into a rigorous notion of probability.

Even if one could somehow add enough additional axioms to justify interpreting the absolute-value-squares of state-vector components like α and β appearing in (27) as *instantaneous, kinematical* probabilities, the traditional many-worlds interpretation of quantum theory lacks an explicit model describing how the experiential trajectory of an individual observer *dynamically* unfolds from moment to moment through the interpretation’s multitudinous (and ill-defined) branching worlds. That is, even if the many-worlds interpretation admits a sequence of static probability distributions at individual moments in time, it does not possess anything like the dynamical conditional probabilities (16) connecting one moment in time of an observer’s experiential trajectory to the next moment in time.

In particular, without adding on significantly more assumptions and metaphysical structure, the many-worlds interpretation is unable to ensure the *ontological stability* of such an observer’s experiential trajectory through time: Observers in the many-worlds interpretation of quantum theory are vulnerable to radical macroscopic instabilities in experience (and memory) that parallel the sorts of macroscopic instabilities that are possible in a hypothetical version of classical physics that lacks ontic-level dynamics, as we explained in Section II B 7 shortly after (17). We will detail a concrete quantum analogue of such dynamical ontological instabilities (eigenstate swaps) in Section III B 3, and we will eliminate them in the context of our own interpretation of quantum theory when we introduce a dynamical

¹⁸ For an explicit discussion of these points in the context of the EPR-Bohm thought experiment, including the issue of nonlocality, see Section VIF 4. There are additional reasons to be suspicious of attempts to give all the branches of the many-worlds interpretation an equal ontological meaning, as, for example, Aaronson describes in the context of quantum computing in [5].

¹⁹ Maudlin eloquently captures this key shortcoming of the many-worlds interpretation in “Problem 2: The problem of statistics” in [232], and [94, 324, 327] attempt to suppress the problem by burying it under the elaborate axiomatic apparatus of decision theory.

notion of quantum conditional probabilities in Section III D. Modal interpretations have been criticized in the past for lacking any such dynamical smoothing conditions on ontic-state trajectories to eliminate these sorts of metaphysical instabilities,²⁰ but to the extent that such criticisms of the traditional modal interpretations are justified, the same criticisms apply equally to the many-worlds interpretation as well. (It should also be noted that criticisms of modal interpretations on the grounds of not singling out a preferred tensor-product factorization of the Hilbert space of the universe into subsystem Hilbert spaces likewise applies equally to the many-worlds interpretation.)

Finally, because all realistic systems exhibit a nonzero degree of quantum entanglement with other systems, the traditional many-worlds interpretation depends crucially upon the axiomatic assertion of a maximal closed system—again, “the universe as a whole”—that admits a description in terms of an *always-exact* objectively pure state. Indeed, the many-worlds interpretation relies on the existence of this cosmic objectively pure state in order to identify the branches on which the superposed copies of various subsystems reside, and with what associated probabilities. However, as we discussed in Section ID 2, there are reasons to be skeptical that any such maximal closed system is physically guaranteed to exist or be mathematically well-defined—especially given that density matrices can exhibit an *objective* form of nontriviality—thus opening up the real possibility that the many-worlds interpretation isn’t fundamentally well-defined either.

10. Our Interpretation

As a first step toward answering our ontological questions about the state vector (27) in the context of our own interpretation of quantum theory, we begin by saying that our uncertainty over each subsystem’s actual ontic state is captured by what we call an objective epistemic state for that subsystem. A quantum epistemic state that includes at least some objective uncertainty of this kind—possibly together with some subjective uncertainty as well—is called an improper mixture. Turning things around, we can then identify an objective epistemic state as being an improper mixture in the extremal case of zero subjective uncertainty.

In contrast to proper mixtures (that is, wholly subjective epistemic states), improper mixtures do not have a widely accepted *a priori* meaning, and it is a central purpose of our minimal modal interpretation of quantum theory to provide one, starting with the requirement that an objective epistemic state must be a logically rigorous probability distribution and therefore must always involve mutually exclusive (26) possible ontic states.

In Section III B, we will define density matrices and relate them to objective epistemic states. Then, after developing a prescription for relating the density matrices of parent systems with those of their subsystems in Section III C, we will finally have in our hands a precise means of resolving the aforementioned questions about the ontic states of *A* and *B*.

B. The Correspondence Between Objective Epistemic States and Density Matrices

Even if we happen to know a system’s objective epistemic state $\{(p_i(t), \Psi_i(t))\}_i$ at a given time t , we have not yet presented a framework for determining the system’s objective epistemic state at any other time $t' \neq t$, nor a prescription for relating ontic states to each other over time or for relating objective epistemic states or ontic states between parent systems and their subsystems. Our first step in these directions will be to define a correspondence between the objective epistemic states of quantum systems and a class of mathematical objects—density matrices—that will serve as the fulcrum of this framework.

1. Density Matrices

Any objective epistemic state $\{(p_i(t), \Psi_i(t))\}_i$ of the kind defined in Section III A 2—meaning, in particular, that it involves *mutually exclusive* possible ontic states $\Psi_i(t)$ —can be identified uniquely²¹ with a unit-trace, positive

²⁰ See “Problem 3: The problem of effect” in [232].

²¹ We address subtleties regarding degeneracies $p_i(t) = p_j(t)$ ($i \neq j$) in Section III B 2.

semi-definite linear operator $\hat{\rho}(t)$ known as the system’s density matrix:

$$\left. \begin{aligned} \hat{\rho}(t) &= \sum_i p_i(t) |\Psi_i(t)\rangle \langle \Psi_i(t)|, & \text{with } \hat{\rho}(t)^\dagger &= \hat{\rho}(t), & \text{Tr}[\hat{\rho}(t)] &= 1, \\ 0 &\leq p_i(t) \leq 1, & \sum_i p_i(t) &= 1, \\ \langle \Psi_i(t) | \Psi_j(t) \rangle &= \delta_{ij}. \end{aligned} \right\} \quad (28)$$

By “identified” here, we mean that there exists a precise correspondence²² between the objective epistemic state $\{(p_i(t), \Psi_i(t))\}_i$ of the system and the eigenvalue-eigenvector pairs of the density matrix $\hat{\rho}(t)$ at each moment in time t :

$$\boxed{\{(p_i(t), \Psi_i(t))\}_i \leftrightarrow \hat{\rho}(t) = \sum_i p_i(t) |\Psi_i(t)\rangle \langle \Psi_i(t)|.} \quad (29)$$

Essentially, the correspondence (29) provides a means of encoding a list of non-negative real numbers—the system’s epistemic probabilities—and a list of state vectors—representing the system’s possible ontic states—in a basis-independent manner, namely, as the eigenvalue-eigenvector spectrum of the system’s density matrix.

Recalling the definition (2) of the Shannon entropy of a classical system’s epistemic state, the correspondence (29) immediately implies that we can express the Shannon entropy $-\sum_i p_i \log p_i$ of a quantum objective epistemic state $\{(p_i(t), \Psi_i(t))\}_i$ in terms of the corresponding density matrix $\hat{\rho}$ according to the basis-independent Von Neumann entropy formula

$$S \equiv -\text{Tr}[\hat{\rho} \log \hat{\rho}]. \quad (30)$$

Thinking about entropy gives yet one more reason for singling out the eigenbasis of a system’s density matrix as providing us with the best possible picture of the system’s underlying ontology: Jaynes proved in [193] that among all possible decompositions of a given system’s density matrix in orthogonal or non-orthogonal bases, it is precisely the decomposition of the density matrix in its orthogonal eigenbasis—that is, in its basis of mutually exclusive ontic states—that yields the smallest possible Shannon entropy $-\sum_i p_i \log p_i$ and thus the smallest possible epistemic uncertainty; this extremal value of the Shannon entropy precisely coincides with the Von Neumann formula (30).²³

In light of the association (29) between probabilities and matrices, together with the replacement of classical random variables (representing observables) with self-adjoint linear operators and the introduction of linear-algebraic expressions like the Von Neumann entropy formula (30), quantum theory has been called “a noncommutative generalization of classical probability theory” [215].

This noncommutativity presents key issues for interpreting quantum theory. The commutativity of *classical* projection functions P_E, P_F onto subsets E, F of a classical system’s sample space is deeply related to the commutativity of the intersection operation \cap of set theory ($P_E P_F = P_{E \cap F} = P_{F \cap E} = P_F P_E$), and so the generic *noncommutativity* of *quantum* projection operators represents an obstruction to naïve attempts at extending set-theoretic notions of sample spaces to quantum theory. Essentially, our minimal modal interpretation approaches this problem and resurrects sample spaces for quantum systems by singling out the subalgebra of a given quantum system’s projection operators that all commute with each other as well as commute with the system’s density matrix.

Note that we regard (29) as merely being a *correspondence* between an epistemic object and a mathematical object. Thus, just as we do not use the terms “ontic state” and “state vector” synonymously in this paper even though the associated objects are related by (25), we will not use the terms “(objective) epistemic state” and “density matrix” synonymously either, again in contrast to the practice of some authors.

The correspondence (29), which lies at the core of our minimal modal interpretation of quantum theory, provides a natural context for emphasizing the following tautological but crucial point:

$$\boxed{\text{If our interpretation does not explicitly identify a given quantity as being a literal epistemic probability, then that quantity is not—or at least not yet—a literal epistemic probability.}} \quad (31)$$

²² In the traditional language of the modal interpretations [316], this correspondence is termed a (core) property ascription or an ontic ascription, our ontic states Ψ_i are called (core) properties or value states, and our objective epistemic states $\{(p_i, \Psi_i)\}_i$ are called property sets, mathematical states, or dynamical states [317].

²³ For recent work attempting to derive important aspects of statistical mechanics from fundamentally quantum reasoning, including various fluctuation theorems and the evolution of a system toward thermal equilibrium, see, for example, [16, 110, 154, 199, 214, 222, 261, 262].

In particular, in keeping with our earlier comment in Section III A 4 that, insofar as our interpretation involves hidden variables, ontic states play the role of those hidden variables, and also our comment that we do not regard the state vectors representing them as being epistemic probability distributions for anything else, we do not assume the Born rule $\text{Prob}(\dots) = |\dots|^2$ *a priori*, nor do we correspondingly interpret the absolute-value-squared values of the complex components of state vectors as being *literal* probabilities for any of our hidden variables. According to our interpretation, and in accordance with the basic correspondence (29) above, we do not interpret the absolute-value-squared components of state vectors as describing literal probabilities for our hidden variables until some process—perhaps measurement- or environment-induced decoherence—turns them into the eigenvalues of *some* system’s density matrix. For these reasons, once we eventually do derive the Born rule in Section IV B, we will refer to the absolute-value-squared components of pre-measurement state vectors as empirical outcome probabilities, because it is only after a measurement device performs a specific measurement that they go from being mere “potentialities” to becoming *actualized* as true epistemic probabilities for some system.

2. Probability Crossings and Degeneracies

One frequently cited anomaly in the purportedly one-to-one relationship (29) between objective epistemic states and density matrices concerns the issue of eigenvalue degeneracies. However, although one can easily imagine a *classical epistemic state* evolving smoothly through a *probability crossing*

$$p_i(t_c) = p_j(t_c) \quad (32)$$

that occurs at some specific moment in time $t = t_c$, *eigenvalue crossings* in *density matrices* are arrangements that require infinitely sharp—that is, measure-zero—fine-tuning and thus never realistically occur: They are co-dimension-three events in the abstract four-dimensional space of 2×2 density-matrix blocks [30, 182] and would therefore require that the off-diagonal entries vanish *precisely* when the diagonal elements *exactly* agree.

Just for purposes of illustration, consider the following *static* example: Although the two-particle, entangled, spin-singlet state vector

$$|\Psi_{\text{EPR-Bohm}}\rangle = \frac{1}{\sqrt{2}} (|\uparrow\downarrow\rangle - |\downarrow\uparrow\rangle) \quad (33)$$

familiar from the EPR-Bohm thought experiment (to be discussed in detail in Section V B) would naïvely lead to degenerate 2×2 reduced density matrices for each individual particle, experimentally setting up the state vector (33) *exactly* would require unrealistically fine-tuning the total spin to $S_{\text{tot},z} = 0$ to infinite precision—a measure-zero state of affairs. Even putting aside the practical problem that no realistic two-particle system can ever be put in an *exactly* pure state having absolutely no external entanglement in the first place, there will unavoidably exist degeneracy-breaking deviations among the components of state vectors like $|\Psi_{\text{EPR-Bohm}}\rangle$, so that, at best, the state vector actually takes the form

$$|\Psi_{\text{EPR-Bohm}}\rangle = \frac{1}{\sqrt{2}} ((1 + \epsilon_1) |\uparrow\downarrow\rangle - (1 + \epsilon_2) |\downarrow\uparrow\rangle + \epsilon_3 |\uparrow\uparrow\rangle + \epsilon_4 |\downarrow\downarrow\rangle), \quad |\epsilon_1|, |\epsilon_2|, |\epsilon_3|, |\epsilon_4| \ll 1. \quad (34)$$

Some have suggested [11] that the problem with density matrices is how to interpret them when they exhibit *exact* degeneracies, which we now see are unphysical idealizations. However, as we will explain next, the *actual* danger for interpretations that are based on a correspondence like (29) (such as most modal interpretations) arises from the fact that density matrices can never realistically *have* exact degeneracies in the first place.

3. Near-Degeneracies and Eigenstate Swaps

The closest realistic density-matrix counterpart to a probability crossing (32) is a near-degeneracy in which two probability eigenvalues reach a point of closest approach

$$|p_i(t) - p_j(t)| \sim \rho_0 \xi \ll 1 \quad (35)$$

at some specific moment in time t_0 and then turn around again, while their associated orthogonal eigenstates $|\Psi_i(t)\rangle \perp |\Psi_j(t)\rangle$ undergo an ultra-fast eigenstate swap²⁴

$$\left. \begin{aligned} |\Psi_i(t)\rangle &\mapsto |\Psi_i(t + \delta t_{\text{swap}})\rangle \approx |\Psi_j(t)\rangle \perp |\Psi_i(t)\rangle, \\ |\Psi_j(t)\rangle &\mapsto |\Psi_j(t + \delta t_{\text{swap}})\rangle \approx |\Psi_i(t)\rangle \perp |\Psi_j(t)\rangle \end{aligned} \right\} \quad (36)$$

over an ultra-short time scale of order

$$\delta t_{\text{swap}} \sim \rho_0 \xi \tau. \quad (37)$$

The complex-valued dimensionless quantity

$$\xi \sim \exp(-\#\text{degrees of freedom}) \quad (38)$$

has exponentially small magnitude in the total number of degrees of freedom of the system itself and of all other systems that substantially interact and entangle with it, and characterizes the size of the off-diagonal elements in the relevant 2×2 density-matrix block in a basis in which the diagonal elements become momentarily equal at precisely $t = t_0$. The real-valued quantity ρ_0 is the would-be-degenerate eigenvalue in the idealized but unphysical measure-zero case $\xi = 0$. Hence, near $t = t_0$, the relevant 2×2 density-matrix block takes the approximate form

$$\begin{pmatrix} \rho_0 + (t - t_0) / \tau & \rho_0 \xi \\ \rho_0 \xi^* & \rho_0 - (t - t_0) / \tau \end{pmatrix} \subset \hat{\rho}. \quad (39)$$

Meanwhile, the real-valued quantity τ , which has units of time, is the characteristic time scale over which the probability eigenvalues $p_i(t)$ and $p_j(t)$ would be changing in the *absence* of eigenstate-swap effects. To the approximate extent that we can speak of well-defined energies E for a system that is open and thus whose dynamics is not strictly unitary, ordinary time evolution corresponds roughly to the $\exp(-iEt/\hbar)$ relative phase factors familiar from elementary quantum theory, and thereby implies that

$$\tau \sim \hbar/E. \quad (40)$$

Taken seriously, eigenstate swaps (36) could conceivably cause large systems to undergo frequent but sudden fluctuations between *macroscopically distinct* ontic states over arbitrarily short time scales, and thus eigenstate swaps have long been regarded as being serious ontological instabilities inherent in density-matrix-centered interpretations of quantum theory.²⁵ Fortunately, after we discuss the internal dynamics of quantum systems according to our minimal modal interpretation in Section III D, we will be able to show in detail that eigenstate swaps turn out to be a mirage.

4. The Fundamentally Unobservable Nature of Eigenstate Swaps

As a first hint that we should not take eigenstate swaps (36) seriously as physically real phenomena for macroscopic systems, notice that a macroscopic system's eigenstate-swap time scale δt_{swap} in (37) is *parametrically* small in the exponentially tiny quantity ξ described in (38), and is therefore always *exponentially* smaller than the system's ordinary characteristic time scale τ described in (40). Thus, the hierarchical discrepancy between a macroscopic system's characteristic time scale τ and its corresponding eigenstate-swap time scale δt_{swap} would actually *get worse* if we could somehow arrange for the system to approach the idealized limit of an exact degeneracy, because that would just make $\delta t_{\text{swap}} \sim \xi$ even smaller. Similarly, if we were to attempt to hook the system up to a macroscopic measuring device with the goal of trying to observe the eigenstate swap experimentally—perhaps introducing a lot more energy in order to achieve a very fine temporal measurement resolution—then we *would* decrease the overall system's ordinary characteristic time scale τ , but, simultaneously, we would *vastly* decrease ξ and thereby end up pushing the eigenstate-swap time scale even farther out of reach.

²⁴ Eigenstate swaps are called *core property instabilities* or *fluctuations* in [316] and *crossovers* in [182].

²⁵ As Vermaas writes on p. 133 of [316] in critiquing his own modal interpretation of quantum theory: “If one [takes eigenstate swaps seriously], it follows that the set of eigenprojections of a state that comes arbitrarily close to a degeneracy can change maximally in an arbitrarily small time interval $t_2 - t_1$. [...] [T]he set of core properties can change rapidly, resulting in an unstable property ascription during a finite time interval.” Later, on p. 260, he writes: “[I]f a state has a spectral resolution which is nearly degenerate, then an arbitrarily small change of that state (by an interaction with the environment or by internal dynamics) can maximally change the set of the possible core properties. Perhaps this instability is one of the more serious defects of modal interpretations because, firstly, it can have consequences for their ability to solve the measurement problem. For even when a modal interpretation ascribes readings to a pointer at a specific instant, a small fluctuation of the state may mean that at the next instant the pointer possesses properties which are radically different to readings.”

Interestingly, because a macroscopic quantum system's Hilbert space has dimension

$$\dim \mathcal{H} \sim \prod_{\substack{\text{range of each} \\ \text{degree of freedom} \\ \text{degrees} \\ \text{of freedom}}} \left(\frac{\text{range of each}}{\text{degree of freedom}} \right) \sim \exp(\#\text{degrees of freedom}) \quad (41)$$

and its maximum entropy (30) goes as

$$S \sim - \sum_{\text{basis}} \frac{1}{\dim \mathcal{H}} \log \frac{1}{\dim \mathcal{H}} \sim \log \dim \mathcal{H}, \quad (42)$$

we see that the parameter ξ from (38) loosely corresponds to our error-entropy bound (3):

$$\xi \sim \exp(-\#\text{degrees of freedom}) \sim e^{-S} \sim \text{minimum error}. \quad (43)$$

This result provides additional justification for regarding both the distance of closest approach (35) between the near-degenerate probability eigenvalues, as well as the eigenstate-swap time scale δt_{swap} in (37), as being unobservably small for macroscopic systems.

Eigenstate swaps for macroscopic systems therefore remain effectively decoupled from the observable predictions of quantum theory. This decoupling opens up an important window of opportunity for trying to smooth these kinds of instabilities out of existence altogether by a suitable choice of internal dynamics for quantum ontic states, with the added benefit of making the decoupling more manifest. We put forward just such a proposal in Section III D, where we show explicitly that our choice of internal dynamics ensures that a macroscopic system's time-evolving actual ontic state essentially never undergoes eigenstate swaps.

C. Parent Systems, Subsystems, the Partial-Trace Operation, and System-Centric Ontology

Up to now, our discussion has centered on the case in which we consider just one particular system of interest. Of course, any study of quantum theory must begin with some particular system, but there will generally be other systems that we can consider simultaneously, and each may naturally be interpreted as a parent system enclosing our original system, or as a subsystem, or as an adjacent system, or as something else entirely. In this section, we motivate and explain how our minimal modal interpretation of quantum theory defines the relationship between the objective epistemic states of parent systems and those of their subsystems, especially in situations like (27) that feature quantum entanglement.

1. The Partial-Trace Operation

Recalling our discussion of classical parent systems and subsystems in Section III A 8, and given an epistemic state p_{A+B} for a classical parent system $A+B$, we can naturally obtain reduced (or marginal) epistemic states p_A and p_B for the respective subsystems A and B by the familiar partial-sum (or marginalization) operation:

$$p_A(a) = \sum_b p_{A+B}(a,b), \quad p_B(b) = \sum_a p_{A+B}(a,b). \quad (44)$$

Our next goal will be to motivate an analogous *quantum* operation, called the partial-trace operation, that relates the objective epistemic states of parent systems with those of their subsystems, and we will see that density matrices play a crucial role.

It is important to note that we will make no appeals to the Born rule, Born-rule-based averages, or Von Neumann-Lüders projections in justifying the definition of the partial-trace operation. Indeed, we will ultimately find that we can *derive* the Born rule and its corollaries (as well as possible corrections described in Section IV B 4 that are invisible in the traditional Copenhagen interpretation) from the partial-trace operation's deeper principles of logical self-consistency, its relationship with classical partial sums, and the general axiomatic interpretation (29) that we attach directly to the reduced density matrices representing improper mixtures.

Central to our entire interpretation is the notion of system-centric ontology, by which we mean that every quantum system has an ontology and objective epistemology defined through a density matrix of its own. In particular, it is important to remember that in our interpretation of quantum theory, the density matrix or objective epistemic state of any *other* system, even that of a *parent* system, does not *directly* define the ontology or objective epistemology of the original system of interest.

Being mindful of the system-centric character of ontology in our interpretation, and given a quantum objective epistemic state for a parent system $W = A + B + C + D + \dots$ consisting of an arbitrary number of subsystems A, B, C, D, \dots and associated with some density matrix $\hat{\rho}_W$, we therefore require a *universal* prescription for *nontrivially* assigning unit-trace, positive semi-definite density matrices to all the various possible definable subsystems $A, B, C, D, \dots, A+B, A+C, A+D, \dots, A+B+C, A+B+D, \dots, A+B+C+D, \dots$ in such a way that there is no dependence on our arbitrary choice of orthonormal basis for each of the individual Hilbert spaces of the subsystems, nor any dependence on the arbitrary order in which we could imagine defining a descending sequence of density matrices for the subsystems. We must, for instance, end up with the *same* final density matrix $\hat{\rho}_A$ for the subsystem A whether we choose to define intermediate density matrices according to the sequence $W \mapsto A+B+C \mapsto A+B \mapsto A$ or according to the alternative sequence $W \mapsto A+B+D \mapsto A+D \mapsto A$. Equivalently, any arbitrarily complicated diagram typified by the following example must internally commute:

$$\left. \begin{array}{ccc} & \hat{\rho}_{A+B+C} & \\ \swarrow & & \searrow \\ \hat{\rho}_{A+B}, \hat{\rho}_C & & \hat{\rho}_A, \hat{\rho}_{B+C} \\ \swarrow & & \searrow \\ \hat{\rho}_A, \hat{\rho}_B, \hat{\rho}_C & & \end{array} \right\} \quad (45)$$

Furthermore, our final result for any one subsystem's density matrix—say, $\hat{\rho}_A$, for the subsystem A —cannot depend on our arbitrary choice among the *continuously infinite* different ways in which we could have defined the rest of the subsystems B, C, \dots . In particular, we must ultimately get the *same* answer for $\hat{\rho}_A$ whether we decompose \mathcal{H}_W as $\mathcal{H}_A \otimes \mathcal{H}_B \otimes \mathcal{H}_C \otimes \dots$ and then define $\hat{\rho}_A$ by the sequence $W \mapsto A+B \mapsto A$, or whether we instead decompose \mathcal{H}_W as $\mathcal{H}_A \otimes \mathcal{H}_{B'} \otimes \mathcal{H}_{C'} \otimes \dots$ for some different definitions B', C', \dots of the other subsystems and then define $\hat{\rho}_A$ by the sequence $W \mapsto A+B' \mapsto A$. That is, the continuous infinity of possible diagrams generalizing the following example must each likewise internally commute:

$$\left. \begin{array}{ccc} \hat{\rho}_{A+B+C} = \hat{\rho}_{A+B'+C'} & & \\ \swarrow & & \searrow \\ \hat{\rho}_{A+B} & & \hat{\rho}_{A+B'} \\ \swarrow & & \searrow \\ \hat{\rho}_A & & \end{array} \right\} \quad (46)$$

It is remarkable that a universal prescription meeting all these nontrivial requirements of logical self-consistency exists at all, much less that it turns out to be so straightforward. To define this prescription explicitly and to ensure that it is indeed nontrivial and has the correct behavior in the classical regime, we require furthermore that it should reduce to classical partial sums (44) if the density matrix $\hat{\rho}_{A+B}$ of a parent system $A+B$ happens to be diagonal in the tensor-product basis $|a, b\rangle = |a\rangle \otimes |b\rangle$ with probability eigenvalues $p_{A+B}(a, b)$ —that is, if the state vectors that make up the orthonormal eigenbasis of $\hat{\rho}_{A+B}$ exhibit no quantum entanglement between the subsystems A and B in the sense of (27):

$$\left. \begin{array}{l} \hat{\rho}_{A+B} = \sum_{a,b} p_{A+B}(a, b) |a\rangle \otimes |b\rangle \langle a| \otimes \langle b| \\ \implies \hat{\rho}_A = \sum_a \underbrace{\left(\sum_b p_{A+B}(a, b) \right)}_{p_A(a)} |a\rangle \langle a|, \quad \hat{\rho}_B = \sum_b \underbrace{\left(\sum_a p_{A+B}(a, b) \right)}_{p_B(b)} |b\rangle \langle b|. \end{array} \right\} \quad (47)$$

Additionally, we require that the prescription should be linear over the space of operators on the parent system's Hilbert space, in keeping with the linearity property that holds for classical partial sums over convex combinations $x p_{A+B} + y p'_{A+B}$ ($x, y > 0, x + y = 1$) of classical epistemic states p_{A+B} and p'_{A+B} .

The unique resulting partial-trace operation is well known and surprisingly simple to describe. We begin by considering a parent system C that we can regard as consisting of just two subsystems A and B , where either A or B could be parent systems encompassing subsystems of their own. Then we write the density matrix $\hat{\rho}_{A+B}$ for $C = A+B$ in the tensor-product basis $|a\rangle \otimes |b\rangle$, which, if A and B are entangled, *won't* be the diagonalizing basis for $\hat{\rho}_{A+B}$:

$$\hat{\rho}_{A+B} = \sum_{a,b,a',b'} \rho_{A+B}((a, b), (a', b')) |a\rangle \otimes |b\rangle \langle a'| \otimes \langle b'|. \quad (48)$$

Linearity and consistency with classical partial sums in the idealized zero-entanglement case

$$\rho_{A+B}((a, b), (a', b')) = p_{A+B}(a, b) \delta(a, a') \delta(b, b') \quad (49)$$

then immediately imply that the partial trace down to either subsystem A or B must generally be defined by formally “turning around” the bras and kets of the *other* subsystem and evaluating the resulting inner products, where the result then defines the reduced density matrix $\hat{\rho}_A$ or $\hat{\rho}_B$ of the subsystem A or B , respectively:

$$\left. \begin{aligned} \hat{\rho}_A &\equiv \text{Tr}_B [\hat{\rho}_{A+B}] \equiv \sum_{a,b,a',b'} \rho_{A+B}((a, b), (a', b')) |a\rangle \langle a'| \underbrace{\langle b| b'\rangle}_{\delta(b,b')} \\ &= \sum_{a,a'} \left(\sum_b \rho_{A+B}((a, b), (a', b)) \right) |a\rangle \langle a'|, \\ \hat{\rho}_B &\equiv \text{Tr}_A [\hat{\rho}_{A+B}] \equiv \sum_{a,b,a',b'} \rho_{A+B}((a, b), (a', b')) \underbrace{\langle a| a'\rangle}_{\delta(a,a')} |b\rangle \langle b'| \\ &= \sum_{b,b'} \left(\sum_a \rho_{A+B}((a, b), (a, b')) \right) |b\rangle \langle b'|. \end{aligned} \right\} \quad (50)$$

That is, the matrix elements of $\hat{\rho}_A$ and $\hat{\rho}_B$ are given in the respective orthonormal bases $|a\rangle$ for the Hilbert space \mathcal{H}_A of the subsystem A and $|b\rangle$ for the Hilbert space \mathcal{H}_B of the subsystem B by the natural matrix-generalizations of the classical partial sums (44):

$$\rho_A(a, a') = \sum_b \rho_{A+B}((a, b), (a', b)), \quad \rho_B(b, b') = \sum_a \rho_{A+B}((a, b), (a, b')). \quad (51)$$

2. An Example

With the formula (50) for partial traces in hand, we are finally ready to answer the questions that we posed in Section III A 8 regarding the respective ontic states of the two subsystems A and B belonging to the parent system $A + B$ described by the entangled state vector (27),

$$|\Psi_{A+B}\rangle = \alpha |\Psi_{A,1}\rangle |\Psi_{B,1}\rangle + \beta |\Psi_{A,2}\rangle |\Psi_{B,2}\rangle, \quad \alpha, \beta \in \mathbb{C}, \quad |\alpha|^2 + |\beta|^2 = 1,$$

where we assume for simplicity that $|\Psi_{A,1}\rangle \perp |\Psi_{A,2}\rangle$ and $|\Psi_{B,1}\rangle \perp |\Psi_{B,2}\rangle$ in (27).

Given that $A + B$ has the definite ontic state Ψ_{A+B} defined by (27), the density matrix of $A + B$ is

$$\left. \begin{aligned} \hat{\rho}_{A+B} &= |\Psi_{A+B}\rangle \langle \Psi_{A+B}| \\ &= (\alpha |\Psi_{A,1}\rangle |\Psi_{B,1}\rangle + \beta |\Psi_{A,2}\rangle |\Psi_{B,2}\rangle) (\alpha^* \langle \Psi_{A,1}| \langle \Psi_{B,1}| + \beta^* \langle \Psi_{A,2}| \langle \Psi_{B,2}|), \end{aligned} \right\} \quad (52)$$

thereby implying from the partial-trace prescription (50) that the respective reduced density matrices for A and B are

$$\left. \begin{aligned} \hat{\rho}_A &= |\alpha|^2 |\Psi_{A,1}\rangle \langle \Psi_{A,1}| + |\beta|^2 |\Psi_{A,2}\rangle \langle \Psi_{A,2}|, \\ \hat{\rho}_B &= |\alpha|^2 |\Psi_{B,1}\rangle \langle \Psi_{B,1}| + |\beta|^2 |\Psi_{B,2}\rangle \langle \Psi_{B,2}|. \end{aligned} \right\} \quad (53)$$

Hence, according to our interpretation of quantum theory, the possible ontic states of the subsystem A are $\Psi_{A,1}$ and $\Psi_{A,2}$, and the possible ontic states of the subsystem B are $\Psi_{B,1}$ and $\Psi_{B,2}$, with respective epistemic probabilities given by

$$\left. \begin{aligned} p_A(\Psi_{A,1}) &= |\alpha|^2, & p_A(\Psi_{A,2}) &= |\beta|^2, \\ p_B(\Psi_{B,1}) &= |\alpha|^2, & p_B(\Psi_{B,2}) &= |\beta|^2. \end{aligned} \right\} \quad (54)$$

3. Correlation and Entanglement

Given a parent quantum system $A + B$ consisting of two subsystems A and B , we say that A and B are correlated if the parent system's density matrix $\hat{\rho}_{A+B}$ is not a simple tensor product of the reduced density matrices $\hat{\rho}_A$ and $\hat{\rho}_B$:

$$A \text{ and } B \text{ correlated: } \hat{\rho}_{A+B} \neq \hat{\rho}_A \otimes \hat{\rho}_B. \quad (55)$$

Correlation is a property that exists even in classical physics and corresponds to the failure of the epistemic probabilities $p_{A+B}(a, b)$ for the composite system $A + B$ to factorize: $p_{A+B} \neq p_{APB}$.

But quantum systems are capable of an even stronger property known as entanglement, which we first encountered in (27) and which we define more generally within the context of our interpretation of quantum theory as describing the case in which the parent system's density matrix $\hat{\rho}_{A+B}$ is not diagonal in the tensor-product basis $|a, b\rangle = |a\rangle \otimes |b\rangle$ corresponding to the two subsystems:

$$A \text{ and } B \text{ entangled: } \text{eigenbasis of } \hat{\rho}_{A+B} \text{ is not } |a, b\rangle = |a\rangle \otimes |b\rangle. \quad (56)$$

It is precisely due to entanglement that the partial-trace operation (50) does not generically reduce to simple partial sums (44), and, indeed, is arguably the very reason why we need density matrices in quantum theory in the first place.²⁶

4. Proper Mixtures, Improper Mixtures, and the Measurement Problem

Recall from their definition in Section III A 7 that proper mixtures refer to subjective epistemic states, by which we mean epistemic states that merely encode *subjective* uncertainty about a system's actual ontic state and make perfect sense even in the context of classical physics. By contrast, improper mixtures, as we defined them in Section III A 10, refer to epistemic states that encode *objective* uncertainty arising from quantum entanglement with other systems, in addition to any subjective uncertainty that may also be present.

In the traditional language of quantum theory, there exists a sharp conceptual distinction between proper mixtures and improper mixtures. Indeed, the traditional measurement problem of quantum theory refers not to the question of how pure states evolve to become mixed states—decoherence adequately accounts for that behavior—but to the mystery of how improper mixtures resulting from measurement-induced decoherence become proper mixtures describing definite but classically unknown outcomes.

Our minimal modal interpretation of quantum theory sidesteps and thus resolves the measurement problem in part by blurring the distinction between proper and improper mixtures to a significant degree. Unlike in the Copenhagen interpretation, we do not assume the Born rule or a collapse postulate that converts improper mixture into proper mixtures, and we regard both kinds of mixtures as describing an epistemic state that hides the system's actual ontic state.

5. The Physical Significance of Objective Epistemic States

However, looking back at our basic correspondence (29), note that we have so far introduced density matrices *solely* for describing objective epistemic states—that is, for describing *fully objective* improper mixtures. (We will extend the use of density matrices to more general epistemic states in Section IV C only after deriving the Born rule.) Furthermore, it is important to keep in mind that *it really matters* whether or not a system's actual ontic state is hiding behind a density matrix describing a truly nontrivial objective epistemic state: There can exist physical differences between the case of a system whose density matrix happens to be pure $\hat{\rho} = |\Psi\rangle\langle\Psi|$ and a system whose density matrix is mixed but whose actual underlying ontic state nonetheless happens to be Ψ .

For example, consider a composite system $A + B$ whose density matrix has the *pure* form

$$\hat{\rho}_{A+B} = |1\rangle|\Psi\rangle\langle 1|\langle\Psi|, \quad (57)$$

where 1 is an allowed state of A and Ψ is an allowed state of B . We would naturally conclude not only that the parent system $A + B$ has actual ontic state $(1, \Psi)$ with unit epistemic probability, but, moreover, that the subsystems A and

²⁶ Note that the *traditional* formulation of quantum theory—unlike our own interpretation—does not single out the diagonalizing eigenbasis of a density matrix as being *fundamentally* preferred. Outside of the idealized case of a quantum system belonging to a parent system in an exactly pure state, entanglement (56) therefore ceases to be a more than a purely formal property of density matrices known as non-separability, and there exist various measures, such as quantum discord [177, 245], for characterizing the quantumness of the resulting correlations.

B themselves have respective actual ontic states 1 and Ψ each with unit epistemic probability, for the simple reason that the reduced density matrices for A and B are respectively of the pure forms $\hat{\rho}_A = |1\rangle\langle 1|$ and $\hat{\rho}_B = |\Psi\rangle\langle\Psi|$.

But now suppose instead that the parent system's density matrix has the *mixed* form

$$\hat{\rho}_{A+B} = p_1 |1\rangle|\Psi\rangle\langle 1|\langle\Psi| + p_2 |2\rangle|\Phi\rangle\langle 2|\langle\Phi|, \quad p_1 + p_2 = 1, \quad (58)$$

where we assume that $|1\rangle$ is orthogonal to $|2\rangle$ —that is, $\langle 1|2\rangle = 0$ —but *not necessarily* that $|\Psi\rangle$ is orthogonal to $|\Phi\rangle$. (Note, however, that the tensor-product state vector $|1\rangle|\Psi\rangle$ is nonetheless orthogonal to $|2\rangle|\Phi\rangle$.) Remember that in our interpretation of quantum theory, we determine a system's possible ontic states and associated epistemic probabilities from the orthonormal spectrum of that system's *own* density matrix. Hence, although it would be correct to say that the actual ontic state of the parent system could be $(1, \Psi)$ with epistemic probability p_1 or $(2, \Phi)$ with epistemic probability p_2 , it would be *incorrect* to conclude that the actual ontic state of the subsystem B is Ψ with epistemic probability p_1 or is Φ with epistemic probability p_2 , because $\langle\Psi|\Phi\rangle \neq 0$ means that Ψ and Φ are not mutually exclusive possibilities and therefore cannot correspond to the mutually orthogonal eigenstates of the density matrix of the subsystem B . As a further consequence, it would also be incorrect to say that if the actual ontic state of the parent system happens to be, say, $(1, \Psi)$, then the actual ontic state of the subsystem B must be Ψ .

The reason for all this trouble is, unsurprisingly, entanglement (56): The diagonalizing basis for $\hat{\rho}_{A+B}$ given in (58) is not the “correct” tensor-product orthonormal basis that corresponds to our original tensor-product factorization $\mathcal{H}_{A+B} = \mathcal{H}_A \otimes \mathcal{H}_B$ of the parent system's Hilbert space \mathcal{H}_{A+B} into the respective Hilbert spaces \mathcal{H}_A of the subsystem A and \mathcal{H}_B of the subsystem B . Of course, if the inner product $\langle\Psi|\Phi\rangle$ is very small in magnitude, then we are welcome to choose a *slightly different* tensor-product decomposition $\mathcal{H}_{A+B} = \mathcal{H}_{A'} \otimes \mathcal{H}_{B'}$ in order to make the diagonalizing basis for $\hat{\rho}_{A+B}$ the tensor-product basis $|a'\rangle|b'\rangle = |a'\rangle \otimes |b'\rangle$ corresponding to *suitably redefined* subsystems A' and B' . In that case, we could immediately read off from the eigenvalues of $\hat{\rho}_{A+B}$ the corresponding epistemic probabilities for the subsystems A' and B' , and declare that if the actual ontic state of the composite system happens to be, say, (a', b') , then the actual ontic states of A' and B' must respectively be a' and b' .

However, if we insist on working with our *original* subsystems A and B , then it might seem that all we can conclude about the subsystem B from the density matrix (58) of the parent system $A + B$ is that the reduced density matrix of B is given by the partial trace $\hat{\rho}_B = \text{Tr}_A[\hat{\rho}_{A+B}]$, as defined in (50). In particular, the fact that the parent system $A + B$ has a specific actual ontic state might not appear to imply anything about the actual ontic state of B alone.

There is a seemingly obvious connection between the actual ontic state of the parent system $A + B = A' + B'$ and the actual ontic state of B' , but does the actual ontic state of the parent system tell us anything about the actual ontic state of B , which is presumably “just a slightly different version” of the subsystem B' ? Without addressing this question, our interpretation of quantum theory would seem to be woefully inadequate, as our notion of an actual ontic state would be infinitely sensitive to arbitrarily small (and thus observationally meaningless) redefinitions of subsystems.²⁷ We will study and resolve this important issue in Section III D when we explicitly define the relationship between the ontic states of parent systems and the ontic states of their subsystems. In Section III E, we will introduce the relatively unexplored notion of “subsystem spaces” to describe the continuously infinite set of distinct ways of defining different versions of a particular subsystem.

6. Imperfect Tensor-Product Factorizations, Truncated Hilbert Spaces, Approximate Density Matrices, and Unstable Systems

Given a system's Hilbert space, there will generally exist many possible tensor-product factorizations defining valid subsystems. But, going the other way, there may be cases in which an *a priori* desired choice of subsystem cannot be realized as an *exact* tensor-product factorization of a given parent system's Hilbert space.

For example, the Hilbert space of Nature is presently unknown, and so we have no reason to believe that it admits a simple tensor-product factorization $\mathcal{H}_{\text{Nature}} = \mathcal{H}_{\text{system}} \otimes \mathcal{H}_{\text{other}}$ that includes the exact Hilbert space $\mathcal{H}_{\text{system}}$ of any of the kinds of systems known today. Nonetheless, we profitably employ such Hilbert spaces all the time. Indeed, even for the familiar example of a weakly interacting quantum field theory, we know that the full system's Hilbert-Fock space $\mathcal{H}_{\text{Fock}} = \mathcal{H}_0 \text{ particles} \oplus \mathcal{H}_1 \text{ particle} \oplus \mathcal{H}_2 \text{ particles} \oplus \dots$ does not neatly tensor-product-factorize as $\mathcal{H}_1 \text{ particle} \otimes \mathcal{H}_{\text{other}}$, where $\mathcal{H}_1 \text{ particle}$ is the Hilbert space of a single particle of a certain species, and yet we successfully make use of one-particle Hilbert spaces $\mathcal{H}_1 \text{ particle}$ whenever we wish to study nonrelativistic or semi-relativistic physics.

²⁷ Indeed, a potential instability somewhat analogous to eigenstate swaps (36) arises in the context of small changes in the definition of a subsystem, as explored in [30, 103]. As Vermaas writes on p. 134 of [316]: “A final remark concerns yet another source for incorrect property ascriptions. In this book I always assume that one can precisely identify the systems. Consequently, one can also precisely identify the composites of these systems. If, however, one proceeds the other way round and starts with a set of composites, one has to answer the question of how exactly to factor these composites into disjoint subsystems. In Bacciagaluppi, Donald and Vermaas (1995, Example 7.3) it is proved that the property ascription to a subsystem can depend with high sensitivity on the precise identification of that subsystem.”

We clearly need a prescription for obtaining *approximate* subsystem Hilbert spaces even when the full Hilbert space of the parent system is either unknown or does not exactly permit the desired tensor-product factorization. That prescription, which is implicit in all the foregoing examples, consists of first *truncating* the full Hilbert space—or, if the full Hilbert space is unknown, then regarding it as *already* being appropriately truncated—in order to make the desired tensor-product factorization possible, and only then taking appropriate partial traces to define reduced density matrices as needed.

An immediate corollary is that the restricted density matrix on the parent system's truncated Hilbert space, and thus also the associated reduced density matrix of the subsystem of interest, are only *approximate* objects and, despite still being positive semi-definite, will no longer have exactly unit trace, corresponding to the statement that their probability eigenvalues no longer add up all the way to unity: $\sum_i p_i < 1$. The discrepancy $(1 - \sum_i p_i) \in [0, 1]$ arises from the absence of the states in which our subsystem does not exist (or no longer exists)—that is, the discrepancy is directly related to our inability to capture the system's actual ontic state among the possible ontic states Ψ_i remaining in our description of the system—and so we naturally interpret the discrepancy as representing the probability that our system, which we now recognize as being unstable, has actually decayed.

D. Quantum Conditional Probabilities

We are now ready to say more about the time evolution and dynamics of ontic states and objective epistemic states, as well as address lingering questions regarding eigenstate swaps (36) that we first encountered in Section III B 3 and parent-subsystem discrepancies arising from nontrivial objective epistemic states that we first encountered in Section III C 5.

1. Classical Dynamics and Multilinear Dynamical Mappings

Recall from Section II B 7 that a classical system with well-defined dynamics is a system that, to an acceptable level of approximation, possesses an ontic-level dynamical mapping (16) that is independent of the system's epistemic state and naturally lifts to a multilinear dynamical mapping (17) relating initial and final epistemic states as well. In the special case of Markovian dynamics (19)—that is, for a system whose dynamical mapping is first order, meaning that it requires the input of initial data at only a *single* initial time—the ontic-level dynamics (16) takes the general form

$$p(\cdot; t' | \cdot; t) : \underbrace{(q; t)}_{\text{initial data}}, \underbrace{(q'; t')}_{\text{final data}} \mapsto \underbrace{p(q'; t' | q; t)}_{\text{conditional probabilities}} \quad (59)$$

and the corresponding epistemic-level dynamics (17) becomes the simple linear mapping

$$p(\cdot; t' | \cdot; t) : \underbrace{p(q'; t')}_{\text{epistemic state at } t'} = \sum_q \underbrace{p(q'; t' | q; t)}_{\substack{\text{conditional probability} \\ \text{(independent} \\ \text{of epistemic} \\ \text{states)}}} \underbrace{p(q; t)}_{\text{epistemic state at } t}. \quad (60)$$

We can regard this last equation as being a dynamical Bayesian propagation formula, or, equivalently, as being a self-consistent ensemble average over the ontic-level dynamics.

Consider now a classical system Q with a configuration space $\mathcal{C}_Q = \{q\}_q$ and suppose that Q is an open subsystem of some larger classical system $W = Q + E$. Here E is the environment of Q inside W and has configuration space $\mathcal{C}_E = \{e\}_e$ of its own, so that the configuration space of W is the Cartesian product $\mathcal{C}_W = \mathcal{C}_Q \times \mathcal{C}_E = \{w = (q, e) | q \in \mathcal{C}_Q \text{ and } e \in \mathcal{C}_E\}$. Then even if the parent system W as a whole has well-defined dynamics in the sense of a linear mapping $p_W(\cdot; t' | \cdot; t)$ that is independent of the epistemic state of W and that we assume for simplicity is of the Markovian form (60), the open subsystem Q of W will not necessarily have well-defined dynamics of its own.

Indeed, by expressing the time evolution of the epistemic state of Q over the time interval from t to t' as

$$\begin{aligned} p_Q(q'; t') &= \sum_{e'} p_W(w' = (q', e'); t') \\ &= \sum_{e', w} p_W(w' = (q', e'); t' | w = (q, e); t) p_W(w = (q, e); t) \end{aligned}$$

$$= \sum_q \left[\sum_{e',e} p_W(w' = (q', e'); t' | w = (q, e); t) \left(\frac{p_W(w = (q, e); t)}{p_Q(q; t)} \right) \right] p_Q(q; t), \quad (61)$$

we see immediately that the failure of Q to possess its own dynamics is characterized by the complicated *nonlinear* dependence of the ratio in parentheses (both in its numerator and in its denominator) on the epistemic state $p_Q(q; t)$ of Q . We readily identify this ratio as being the conditional probability that the environment E is in the ontic state e at the time t given that the subsystem Q is in the ontic state q at the same time t :

$$\frac{p_W(w = (q, e); t)}{p_Q(q; t)} = p_{E|Q}(e; t | q; t). \quad (62)$$

Notice that the open-subsystem evolution law (61) provides us with a *coarse-grained* or *effective* (albeit nonlinear) notion of dynamics $p_{Q \subset W}(\cdot; t' | \cdot; t)$ for Q ,

$$p_{Q \subset W}(\cdot; t' | \cdot; t) : \underbrace{p_Q(q'; t')}_{\text{epistemic state at } t'} = \sum_q \underbrace{p_{Q \subset W}(q'; t' | q; t)}_{\substack{\text{conditional} \\ \text{probability} \\ \text{(has dependence} \\ \text{on epistemic} \\ \text{states)}}} \underbrace{p_Q(q; t)}_{\text{epistemic state at } t}, \quad (63)$$

where we have defined the coarse-grained or effective conditional probabilities $p_{Q \subset W}(q'; t' | q; t)$ to be the factor appearing in brackets in (61)—that is, by multiplying each parent-system conditional probability $p_W(w'; t' | w; t)$ by the epistemic probability $p_W(w; t)$ for W at the time t , marginalizing over the environment E at both the initial time t and final time t' , and then conditioning on Q at the time t :

$$\begin{aligned} p_{Q \subset W}(q'; t' | q; t) &= \frac{p_{Q \subset W}(q'; t \text{ and } q; t)}{p_Q(q; t)} \\ &= \frac{1}{p_Q(q; t)} \sum_{e',e} p_W(w' = (q', e'); t' \text{ and } w = (q, e); t) \\ &= \sum_{e',e} p_W(w' = (q', e'); t' | w = (q, e); t) \underbrace{\left(\frac{p_W(w = (q, e); t)}{p_Q(q; t)} \right)}_{p_{E|Q}(e; t | q; t)}. \end{aligned} \quad (64)$$

In the case in which the correlations between Q and its environment E inside W wash out over some short characteristic time scale $\delta t_Q \ll t' - t$ —say, through irreversible thermal radiation into outer space—the epistemic probabilities for W approximately factorize,

$$p_W(w = (q, e)) \approx p_Q(q) p_E(e), \quad (65)$$

and so the ratio (62) appearing in parentheses in the evolution equation (61) for Q reduces to

$$p_{E|Q}(e; t | q; t) = \frac{p_W(w = (q, e); t)}{p_Q(q; t)} = \frac{p_Q(q; t) p_E(e; t)}{p_Q(q; t)} = p_E(e; t). \quad (66)$$

Hence, the dynamical mapping (63) for Q now defines properly linear dynamics $p_{Q \subset W}(\cdot; t' | \cdot; t) = p_Q(\cdot; t' | \cdot; t)$ for Q on its own,

$$p_Q(\cdot; t' | \cdot; t) : \underbrace{p_Q(q'; t')}_{\text{epistemic state at } t'} = \sum_q \underbrace{p_Q(q'; t' | q; t)}_{\substack{\text{conditional} \\ \text{probability} \\ \text{(independent} \\ \text{of epistemic} \\ \text{states)}}} \underbrace{p_Q(q; t)}_{\text{epistemic state at } t}, \quad (67)$$

where (64) has reduced to

$$p_Q(q'; t' | q; t) = \sum_{e',e} p_W(w' = (q', e'); t' | w = (q, e); t) p_E(e; t). \quad (68)$$

Nonetheless, even if the dynamics of the parent system W is deterministic, so that its own dynamical conditional probabilities $p_W(w'; t' | w; t)$ are trivial and relate each initial ontic state w to a unique final ontic state w' with unit probability, keep in mind that the presence of the environment's instantaneous epistemic probabilities $p_E(e; t)$ in the formula (68) for the conditional probabilities $p_Q(q'; t' | q; t)$ generally implies that the dynamics of Q is stochastic.

Notice that the characteristic time scale δt_Q over which correlations wash out between Q and E sets a natural short-time cutoff on the existence of dynamics for our open subsystem Q . Although Q is a sensible system at the level of *kinematics* over time scales shorter than δt_Q , it does not possess truly well-defined *dynamics* on time scales shorter than δt_Q .

2. Linear CPTP Dynamical Mappings

The classical linear dynamical mapping (60) has a much-studied quantum counterpart that generalizes the “deterministic” unitary Schrödinger dynamics of closed quantum systems to a form of linear *stochastic* dynamics governing the approximate time evolution of a large class of open quantum systems.²⁸ This quantum-dynamical mapping consists of a linear mapping $\mathcal{E}^{t' \leftarrow t}$ that relates an open system's initial and final density matrices at respectively initial and final times $t \leq t'$:

$$\hat{\rho}(t) \mapsto \hat{\rho}(t') = \mathcal{E}^{t' \leftarrow t}[\hat{\rho}(t)]. \quad (69)$$

Notice that $\mathcal{E}^{t' \leftarrow t}$ defines dynamics on *density matrices*, in contrast to classical dynamics (60) on epistemic states directly. Functions like $\mathcal{E}^{t' \leftarrow t}$ that map operators to operators are called superoperators.

Intuitively, if (69) describes a mapping that takes in an initial density matrix and produces a final density matrix, then the mapping should, in particular, preserve the *unit-trace* of density matrices:

$$\text{Tr}[\hat{\rho}(t)] = 1 \implies \text{Tr}[\mathcal{E}^{t' \leftarrow t}[\hat{\rho}(t)]] = 1. \quad (70)$$

The assumed linearity of $\mathcal{E}^{t' \leftarrow t}$, combined with its preservation of the trace of unit-trace operators in accordance with (70), means that it must in fact preserve the traces of *all* operators \hat{O} , and thus must in general be a trace-preserving (“TP”) mapping:

$$\mathcal{E}^{t' \leftarrow t} \text{ is TP: } \quad \text{Tr}[\mathcal{E}^{t' \leftarrow t}[\hat{O}]] = \text{Tr}[\hat{O}] \text{ for all } \hat{O}. \quad (71)$$

Furthermore, (69) should be a positive mapping, meaning that it preserves the positive semi-definiteness of density matrices:

$$\hat{\rho}(t) \geq 0 \implies \mathcal{E}^{t' \leftarrow t}[\hat{\rho}(t)] \geq 0. \quad (72)$$

If we were to imagine introducing an arbitrary, causally disconnected ancillary system with trivial dynamics determined by the identity mapping id on operators, then it would be reasonable to impose the requirement that the resulting composite dynamics $\mathcal{E}^{t' \leftarrow t} \otimes \text{id}$ governing the pair of systems should likewise be a positive mapping, a surprisingly nontrivial condition on our original dynamical mapping $\mathcal{E}^{t' \leftarrow t}$ called complete positivity (“CP”):

$$\mathcal{E}^{t' \leftarrow t} \text{ is CP: } \quad \mathcal{E}^{t' \leftarrow t} \otimes \text{id} \text{ is a positive mapping.} \quad (73)$$

We are therefore led to the study of linear completely-positive-trace-preserving (“CPTP”) dynamical mappings of density matrices. In generalizing unitary dynamics to linear CPTP dynamics in this manner, as is necessary in order to account for the crucial and non-reductive quantum relationships between parent systems and their subsystems, note that we are *not* proposing any fundamental modifications to the dynamics of quantum theory nor any new sources of violations of time-reversal symmetry, such as occur in GRW-type spontaneous-localization models [8, 39, 144, 249, 252, 334], but are simply accommodating the fact that generic mesoscopic and macroscopic quantum systems are typically open to their environments to some nonzero degree.²⁹ Indeed, linear CPTP dynamical mappings are widely used in quantum chemistry as well as in quantum information science, in which they are known as quantum operations; when specifically regarded as carriers of quantum information, they are usually called quantum channels.³⁰

²⁸ See [82, 198, 296] for early work in this direction, and see [275] for a modern pedagogical review.

²⁹ Given the generically local interactions found in realistic fundamental physical models like quantum field theories and the resulting tendency of decoherence to reduce open systems automatically to states of relatively well-defined spatial position [197], proponents of GRW-type spontaneous-localization models must justify why their modifications to quantum theory aren't redundant [195] and explain how we would experimentally distinguish those claimed modifications from the more prosaic effects of decoherence.

³⁰ Starting from a simple measure of distinguishability between density matrices that is non-increasing under linear CPTP dynamics [268], one can argue [66, 209] that linear CPTP dynamics implies the absence of any backward flow of information into the system from its environment. [73] strengthens this reasoning by proving that exact linear CPTP dynamics exists for a given quantum system if and only if the system's initial correlations with its environment satisfy a quantum data-processing inequality that prevents backward information flow.

3. Quantum Conditional Probabilities

To motivate our defining formula for the generalized quantum counterpart to the classical conditional probabilities appearing in (59), we begin by considering a collection of mutually disjoint quantum systems Q_1, \dots, Q_n that we can identify as being subsystems of some parent system $W = Q_1 + \dots + Q_n$, where we include the possibility that $n = 1$ (or, equivalently, that $Q_2 = \dots = Q_n = \emptyset$ are all trivial) so that $W = Q_1$. Suppose furthermore that the parent system W has well-defined dynamics over a given time interval $\Delta t \equiv t' - t \geq 0$, meaning that we can approximate the time evolution of the density matrix $\hat{\rho}_W$ of W over the time interval Δt by a linear CPTP dynamical mapping (69):

$$\hat{\rho}_W(t) \mapsto \hat{\rho}_W(t') = \mathcal{E}_W^{t'-t}[\hat{\rho}_W(t)]. \quad (74)$$

We can expand the density matrix $\hat{\rho}_W(t)$ of the parent system W at the initial time t in terms of its probability eigenvalues $p_W(w; t)$ and its eigenprojectors $\hat{P}_W(w; t) \equiv |\Psi_W(w; t)\rangle \langle \Psi_W(w; t)|$, meaning the projection operators onto the orthonormal eigenbasis of $\hat{\rho}_W(t)$:

$$\hat{\rho}_W(t) = \sum_w p_W(w; t) |\Psi_W(w; t)\rangle \langle \Psi_W(w; t)| = \sum_w p_W(w; t) \hat{P}_W(w; t). \quad (75)$$

Similarly, selecting the subsystem Q_1 without loss of generality, we can expand its (reduced) density matrix $\hat{\rho}_{Q_1}(t') = \text{Tr}_{Q_2+\dots+Q_n}[\hat{\rho}_W(t')]$ at the final time t' in terms of its own probability eigenvalues $p_{Q_1}(i_1; t')$ and its own eigenprojectors $\hat{P}_{Q_1}(i_1; t') \equiv |\Psi_{Q_1}(i_1; t')\rangle \langle \Psi_{Q_1}(i_1; t')|$:

$$\hat{\rho}_{Q_1}(t') = \text{Tr}_{Q_2+\dots+Q_n}[\hat{\rho}_W(t')] = \sum_{i_1} p_{Q_1}(i_1; t') |\Psi_{Q_1}(i_1; t')\rangle \langle \Psi_{Q_1}(i_1; t')| = \sum_{i_1} p_{Q_1}(i_1; t') \hat{P}_{Q_1}(i_1; t'). \quad (76)$$

We can also expand each of the identity operators $\hat{1}_{Q_2}, \dots, \hat{1}_{Q_n}$ on the respective Hilbert spaces of the other mutually disjoint subsystems Q_2, \dots, Q_n in terms of the respective eigenprojectors $\hat{P}_{Q_2}(i_2; t'), \dots, \hat{P}_{Q_n}(i_n; t')$ of their own respective (reduced) density matrices $\hat{\rho}_{Q_2}(t'), \dots, \hat{\rho}_{Q_n}(t')$,

$$\hat{1}_{Q_2} = \sum_{i_2} \hat{P}_{Q_2}(i_2; t'), \dots, \hat{1}_{Q_n} = \sum_{i_n} \hat{P}_{Q_n}(i_n; t'), \quad (77)$$

where we'll see that symmetry with Q_1 and consistency with the linear CPTP dynamical mapping (74) for the parent system W necessitates that these projection operators are all evaluated at the *same* final time t' .

From the spectral decompositions (75), (76), and (77), together with our formula (74) for the linear CPTP dynamics of the parent system W , we can trivially express the epistemic probability $p_{Q_1}(i_1; t')$ for the subsystem Q_1 to be in the ontic state $\Psi_{Q_1}(i_1; t')$ at the final time t' as

$$\begin{aligned} p_{Q_1}(i_1; t') &= \text{Tr}_{Q_1}[\hat{P}_{Q_1}(i_1; t') \hat{\rho}_{Q_1}(t')] \\ &= \text{Tr}_W \left[\left(\hat{P}_{Q_1}(i_1; t') \otimes \hat{1}_{Q_2} \otimes \dots \otimes \hat{1}_{Q_n} \right) \hat{\rho}_W(t') \right] \\ &= \text{Tr}_W \left[\left(\hat{P}_{Q_1}(i_1; t') \otimes \hat{1}_{Q_2} \otimes \dots \otimes \hat{1}_{Q_n} \right) \mathcal{E}_W^{t'-t}[\hat{\rho}_W(t)] \right] \\ &= \sum_w \text{Tr}_W \left[\left(\hat{P}_{Q_1}(i_1; t') \otimes \hat{1}_{Q_2} \otimes \dots \otimes \hat{1}_{Q_n} \right) \mathcal{E}_W^{t'-t}[\hat{P}_W(w; t)] \right] p_W(w; t), \\ &= \sum_{i_2, \dots, i_n, w} \text{Tr}_W \left[\left(\hat{P}_{Q_1}(i_1; t') \otimes \hat{P}_{Q_2}(i_2; t') \otimes \dots \otimes \hat{P}_{Q_n}(i_n; t') \right) \mathcal{E}_W^{t'-t}[\hat{P}_W(w; t)] \right] p_W(w; t). \end{aligned} \quad (78)$$

This last expression reduces to an intuitive Bayesian propagation formula with marginalization,

$$p_{Q_1}(i_1; t') = \sum_{i_2, \dots, i_n, w} p_{Q_1, \dots, Q_n|W}(i_1, \dots, i_n; t'|w; t) p_W(w; t), \quad (79)$$

provided that we interpret the trace over W appearing in the summation on w in (78) as the quantum conditional probability $p_{Q_1, \dots, Q_n|W}(i_1, \dots, i_n; t'|w; t)$ for each subsystem Q_α to be in the ontic state $\Psi_{Q_\alpha}(i_\alpha; t')$ at the time t' for $\alpha = 1, \dots, n$ given that the parent system W was in the ontic state $\Psi_W(w; t)$ at the time t :

$$\left. \begin{aligned} p_{Q_1, \dots, Q_n|W}(i_1, \dots, i_n; t'|w; t) &\equiv \text{Tr}_W \left[\left(\hat{P}_{Q_1}(i_1; t') \otimes \dots \otimes \hat{P}_{Q_n}(i_n; t') \right) \mathcal{E}_W^{t'-t}[\hat{P}_W(w; t)] \right] \\ &\sim \text{Tr}[\hat{P}_{i_1}(t') \dots \hat{P}_{i_n}(t') \mathcal{E}[\hat{P}_w(t)]] \end{aligned} \right\} \quad (80)$$

As we'll see, these quantum conditional probabilities serve essentially as smoothness conditions that sew together ontologies in a natural way.

Note that our quantum conditional probabilities (80) are *only* defined in terms of the projection operators onto the orthonormal eigenstates of density matrices—that is, the eigenstates representing each individual system's possible ontic states in accordance with our general correspondence (29)—and *not* in terms of projection operators onto *generic* state vectors, in contrast to the Born rule for computing empirical outcome probabilities. Furthermore, observe that the linear CPTP dynamical mapping $\mathcal{E}_W^{t' \leftarrow t}$ in our definition (80) plays the role of a “parallel-transport” superoperator that carries the parent-system projection operator $\hat{P}_W(w; t)$ from t to t' before we compare it with the subsystem projection operators $\hat{P}_{Q_1}(i_1; t'), \dots, \hat{P}_{Q_n}(i_n; t')$; the Born rule, on the other hand, involves the comparison of state vectors without any kind of parallel transport. Moreover, we emphasize that our definition (80), unlike the axiomatic wave-function collapse of the traditional Copenhagen interpretation of quantum theory, does not introduce any new *fundamental* violations of time-reversal symmetry; indeed, given a system that possesses well-defined *reversed* dynamics expressible in terms of a corresponding *reversed* linear CPTP dynamical mapping that evolves the system's density matrix *backward* in time, we are free to use that reversed linear CPTP dynamical mapping in our definition (80) of the quantum conditional probabilities.

Observe also that the tensor-product operator $\hat{P}_{Q_1}(i_1; t') \otimes \dots \otimes \hat{P}_{Q_n}(i_n; t')$ and the time-evolved projection operator $\mathcal{E}_W^{t' \leftarrow t}[\hat{P}_W(w; t)]$ are both (manifestly) positive semi-definite operators. Hence, each quantum conditional probability $p_{Q_1, \dots, Q_n|W}(i_1, \dots, i_n; t'|w; t)$ defined according to (80) consists of a trace over a product of two positive semi-definite operators, and is therefore guaranteed to be non-negative:³¹

$$p_{Q_1, \dots, Q_n|W}(i_1, \dots, i_n; t'|w; t) \geq 0. \quad (81)$$

From this starting point, we can show that the quantities $p_{Q_1, \dots, Q_n|W}(i_1, \dots, i_n; t'|w; t)$ miraculously satisfy several important and nontrivial requirements that support their interpretation as a quantum class of conditional probabilities.

1. From the completeness of the subsystem eigenprojectors $\hat{P}_{Q_1}(i_1; t'), \dots, \hat{P}_{Q_n}(i_n; t')$, the unit trace of the parent-system eigenprojectors $\hat{P}_W(w; t)$, and the trace-preserving property of the parent-system linear dynamical mapping $\mathcal{E}_W^{t' \leftarrow t}$, the set of quantities $p_{Q_1, \dots, Q_n|W}(i_1, \dots, i_n; t'|w; t)$ for any fixed value of w give unity when summed over the indices i_1, \dots, i_n labeling all the possible ontic states of the mutually disjoint partitioning subsystems Q_1, \dots, Q_n :

$$\sum_{i_1, \dots, i_n} p_{Q_1, \dots, Q_n|W}(i_1, \dots, i_n; t'|w; t) = 1. \quad (82)$$

2. Combining the properties (81) and (82), we see immediately that the quantities $p_{Q_1, \dots, Q_n|W}(i_1, \dots, i_n; t'|w; t)$ are each real numbers in the interval between 0 and 1:

$$p_{Q_1, \dots, Q_n|W}(i_1, \dots, i_n; t'|w; t) \in [0, 1]. \quad (83)$$

3. Recalling our derivation of (78), and from the linearity property of $\mathcal{E}_W^{t' \leftarrow t}$ together with the spectral decomposition (75) of the density matrix $\hat{\rho}_W(t)$ of W in terms of its eigenprojectors $\hat{P}_W(w; t)$ and the dynamical equation (74) for the time evolution of $\hat{\rho}_W(t)$, we see that multiplying the quantity $p_{Q_1, \dots, Q_n|W}(i_1, \dots, i_n; t'|w; t)$ by the epistemic probability $p_W(w; t)$ for W at the time t and summing over w and i_1, \dots, i_n *except for one subsystem index* i_α gives the epistemic probability $p_{Q_\alpha}(i_\alpha; t')$ for Q_α , in agreement with the rule (79) for Bayesian propagation and marginalization:

$$\sum_{i_1, \dots, (\text{no } i_\alpha), \dots, i_n, w} p_{Q_1, \dots, Q_n|W}(i_1, \dots, i_n; t'|w; t) p_W(w; t) = p_{Q_\alpha}(i_\alpha; t'). \quad (84)$$

4. In the idealized case in which $n = 1$ and $W = Q_1 \equiv Q$ is a closed system undergoing “deterministic” unitary dynamics, so that

$$\hat{\rho}_Q(t') = U_Q(t' \leftarrow t) \hat{\rho}_Q(t) U_Q^\dagger(t' \leftarrow t)$$

³¹ Proof: Let $A \geq 0$ and $B \geq 0$ be positive semi-definite matrices. Then \sqrt{A} and \sqrt{B} exist and are likewise positive semi-definite, and so, using the cyclic property of the trace, we have

$$\text{Tr}[AB] = \text{Tr}[\sqrt{A}\sqrt{A}\sqrt{B}\sqrt{B}] = \text{Tr}[\sqrt{B}\sqrt{A}\sqrt{A}\sqrt{B}] = \text{Tr}\left[\left(\sqrt{A}\sqrt{B}\right)^\dagger \sqrt{A}\sqrt{B}\right] \geq 0. \quad \text{QED}$$

Notice that this proof *does not* generalize to traces over products of more than two operators.

for some unitary time-evolution operator $U_Q(t' \leftarrow t)$, we have

$$\hat{P}_Q(j; t') = U_Q(t' \leftarrow t) \hat{P}_Q(j; t) U_Q^\dagger(t' \leftarrow t)$$

and thus, as expected, the quantities $p_{Q|Q}(i; t'|j; t)$ trivialize to the deterministic formula

$$p_{Q|Q}(i; t'|j; t) \equiv \text{Tr}_Q[\hat{P}_Q(i; t') \hat{P}_Q(j; t')] = \delta_{ij} = \begin{cases} 1 & \text{for } i = j, \\ 0 & \text{for } i \neq j. \end{cases} \quad (85)$$

Considering the inflexibility of quantum theory and its famed disregard for the sorts of familiar concepts favored by human beings, it's quite remarkable that the theory contains such a large class of quantities $p_{Q_1, \dots, Q_n|W}(i_1, \dots, i_n; t'|w; t)$ satisfying the properties (81)-(85) of conditional probabilities. Essentially, our interpretation of quantum theory takes this fact at face value by actually *calling* these quantities conditional probabilities.

4. Probabilistic and Non-Probabilistic Uncertainty

However, one should keep in mind that quantum theory does appear to limit what kinds of probabilities we can safely define. In particular, our derivation (78) does not extend to *non-disjoint* subsystems, nor to subsystems at *multiple distinct* final times, nor to *multiple parent* systems. Hence, certain kinds of hypothetical statements involving the ontic states of our interpretation of quantum theory turn out not to admit generally well-defined probabilities, despite lying behind a veil of uncertainty.

Physicists tend to use the terms “uncertain” and “probabilistic” essentially synonymously, but the two concepts are distinct. Indeed, there is no logically airtight *a priori* reason to believe as a general truth about Nature that the frequency ratio of *every* kind of repeatable event should “tend to” some specific value in the limit of many trials; the existence of such a limit is actually a highly nontrivial and non-obvious constraint because we could easily imagine outcomes instead occurring in a completely unpredictable way that never “settles down” in the long run. Science would certainly be a far less successful predictive enterprise if *observable* phenomena did not usually obey such a constraint, but it makes no difference to the predictive power of science if *hidden variables* do not always comply.

In fact, economists have known for many years that although certain types of uncertainty, called probabilistic uncertainty or risk, could safely be described in terms of probabilities, other kinds of uncertainty, called non-probabilistic uncertainty, could not be.³² This distinction arises in computer programming as the difference between the possibility of two *pseudo-random numbers* agreeing and the possibility of a *user* deciding to input two numbers that agree. As a more concrete example, consider a coin biased toward heads by 90% for its first 10 flips, then biased toward *tails* by 90% for its next 100 flips, then biased toward heads again by 90% for its next 1000 flips, and so forth; such a coin's long-running frequency ratios for heads would fluctuate between 90% and 10% and never settle down to a particular limit corresponding to a well-defined probability. Further examples closer to physics are examined in [298], which elegantly explains the nonexistence of a joint probability distribution $p(x, y, z, \dots)$ as the failure of its arguments x, y, z, \dots to possess a well-defined joint limiting relative frequency, perhaps due to obstructions arising from the need for proper subsets of the arguments x, y, z, \dots to satisfy their own mandated limiting relative frequencies.

In line with this last category of cases, we can present a simple example [179] demonstrating that joint probabilities may not exist for certain sets of random variables in classical probability theory. Consider a set of three classical random bits X, Y, Z with individual probabilities given by the table

| | | |
|----------|---------------|---------------|
| | + | - |
| $p_X(x)$ | $\frac{1}{2}$ | $\frac{1}{2}$ |
| $p_Y(y)$ | $\frac{1}{2}$ | $\frac{1}{2}$ |
| $p_Z(z)$ | $\frac{1}{2}$ | $\frac{1}{2}$ |

³² In his light-hearted paper [3], Aaronson refers to the latter as “Knightian” uncertainty, in honor of economist Frank Knight's seminal 1921 book [201] on the subject.

and pairwise-joint probabilities

| | (+, +) | (+, -) | (-, +) | (-, -) |
|-----------------|---|---|---|---|
| $p_{X,Y}(x, y)$ | $\frac{1}{4} \left(1 + \frac{1}{\sqrt{2}}\right)$ | $\frac{1}{4} \left(1 - \frac{1}{\sqrt{2}}\right)$ | $\frac{1}{4} \left(1 - \frac{1}{\sqrt{2}}\right)$ | $\frac{1}{4} \left(1 + \frac{1}{\sqrt{2}}\right)$ |
| $p_{X,Z}(x, z)$ | $\frac{1}{4} \left(1 + \frac{1}{\sqrt{2}}\right)$ | $\frac{1}{4} \left(1 - \frac{1}{\sqrt{2}}\right)$ | $\frac{1}{4} \left(1 - \frac{1}{\sqrt{2}}\right)$ | $\frac{1}{4} \left(1 + \frac{1}{\sqrt{2}}\right)$ |
| $p_{Y,Z}(y, z)$ | $\frac{1}{4}$ | $\frac{1}{4}$ | $\frac{1}{4}$ | $\frac{1}{4}$ |

satisfying the correct partial sums (44):

$$\begin{aligned} \sum_y p_{X,Y}(x, y) &= \sum_z p_{X,Z}(x, z) = \frac{1}{2} = p_X(x), \\ \sum_x p_{X,Y}(x, y) &= \sum_z p_{Y,Z}(y, z) = \frac{1}{2} = p_Y(y), \\ \sum_x p_{X,Z}(x, z) &= \sum_y p_{Y,Z}(y, z) = \frac{1}{2} = p_Z(z). \end{aligned}$$

Then it is impossible to define a joint probability distribution $p_{X,Y,Z}(x, y, z)$ for all three random bits X, Y, Z ,³³ meaning that statements about possible joint outcomes of the form (x, y, z) are uncertain but not probabilistic.

5. Hidden Ontic-Level Nonlocality

Note that the definition (80) of our quantum conditional probabilities is not manifestly local at the ontic level, as we'll make clear explicitly when we discuss the EPR-Bohm thought experiment in Section VB in the context of our minimal modal interpretation of quantum theory. However, the causal structure of special relativity only places constraints on *observable* signals, and our quantum conditional probabilities are, by construction, compatible with standard density-matrix dynamics and thus constrained by the no-communication theorem [166, 255] to disallow superluminal observable signals, as we will explain in greater detail in Section V.

The lack of manifest ontic-level locality in (80) is a feature, not a bug, because, as we'll see when we analyze the EPR-Bohm thought experiment in Section VB and the GHZ-Mermin thought experiment in Section VC, there is no way to accommodate ontic hidden variables without hidden ontic-level nonlocality.³⁴ Hence, if our formula (80) *didn't* allow for benign ontic-level nonlocality, then we would clearly be doing something wrong.

6. Kinematical Relationships Between Ontic States of Parent Systems and Subsystems

Setting $t' = t$ and suppressing time from our notation for clarity, the dynamical mapping $\mathcal{E}^{t' \leftarrow t}$ drops out and the quantum conditional probabilities (80) yield an explicit instantaneous kinematical (and generically probabilistic) relationship between the ontic states of the parent system $W = Q_1 + \dots + Q_n$ and the ontic states of the partitioning collection of mutually disjoint subsystems Q_1, \dots, Q_n :

$$\left. \begin{aligned} p_{Q_1, \dots, Q_n | W}(i_1, \dots, i_n | w) &\equiv \text{Tr}_W \left[\left(\hat{P}_{Q_1}(i_1) \otimes \dots \otimes \hat{P}_{Q_n}(i_n) \right) \hat{P}_W(w) \right] \\ &= \langle \Psi_{W,w} | (|\Psi_{Q_1, i_1}\rangle \langle \Psi_{Q_1, i_1}| \otimes \dots \otimes |\Psi_{Q_n, i_n}\rangle \langle \Psi_{Q_n, i_n}|) | \Psi_{W,w} \rangle. \end{aligned} \right\} \quad (86)$$

³³ We can prove this claim by contradiction: Supposing to the contrary that we could indeed assign joint probabilities to $(x, y, z) = (-, -, -)$ and $(x, y, z) = (-, -, +)$, we see that $p_{X,Y,Z}(-, -, -) \leq p_{Y,Z}(-, -) = 1/4$ and $p_{X,Y,Z}(-, -, +) \leq p_{X,Z}(-, +) = (1/4) \left(1 - 1/\sqrt{2}\right)$, but then the partial-sum formula (44) breaks down because $p_{X,Y,Z}(-, -, -) + p_{X,Y,Z}(-, -, +) \leq (1/4) \left(2 - 1/\sqrt{2}\right) < (1/4) \left(1 + 1/\sqrt{2}\right) = p_{X,Y}(-, -)$. QED

³⁴ See Section VIF 4 for a discussion of the status of locality in the Everett-DeWitt many-worlds interpretation.

In particular, this result leads to a simpler version of our formula (79) for Bayesian propagation and marginalization:

$$p_{Q_1}(i_1) = \sum_{i_2, \dots, i_n, w} p_{Q_1, \dots, Q_n | W}(i_1, \dots, i_n | w) p_W(w). \quad (87)$$

Notice how our quantum conditional probabilities play the role of sewing together the ontologies of W and Q_1, \dots, Q_n .

In the idealized case in which the actual ontic state of the parent system W is of a form $|\Psi_{W,w}\rangle \approx |\Psi_{Q_1, i_1}\rangle \otimes \dots \otimes |\Psi_{Q_n, i_n}\rangle$ that involves approximately no entanglement between its mutually disjoint subsystems Q_1, \dots, Q_n , we obtain the classical-looking result $p_{Q_1, \dots, Q_n | W}(j_1, \dots, j_n | w = (i_1, \dots, i_n)) = \delta_{j_1 i_1} \dots \delta_{j_n i_n}$, as expected. Because decoherence ensures that human-scale macroscopic systems, being in the highly classical regime, exhibit negligible quantum entanglement with one another, we see immediately that the ontologies of typical macroscopic parent systems and their macroscopic subsystems fit together in a classically intuitive, reductionist manner. However, naïve reductionism generically breaks down for microscopic systems like electrons, and thus demanding classically intuitive relationships³⁵ between *microscopic* parent systems and subsystems would mean committing a fallacy of composition or division.³⁶

With the formula (86) in hand, we can easily generalize our earlier entangled example (52) to the case in which the parent system $A + B$ itself has a nontrivial density matrix

$$\hat{\rho}_{A+B} = p_\Psi |\Psi_{A+B}\rangle \langle \Psi_{A+B}| + p_\Phi |\Phi_{A+B}\rangle \langle \Phi_{A+B}|, \quad p_\Psi, p_\Phi \in [0, 1], \quad p_\Psi + p_\Phi = 1,$$

where

$$\begin{aligned} |\Psi_{A+B}\rangle &= \alpha |\Psi_{A,1}\rangle |\Psi_{B,1}\rangle + \beta |\Psi_{A,2}\rangle |\Psi_{B,2}\rangle, & \alpha, \beta \in \mathbb{C}, \quad |\alpha|^2 + |\beta|^2 &= 1, \\ |\Phi_{A+B}\rangle &= \gamma |\Psi_{A,3}\rangle |\Psi_{B,3}\rangle + \delta |\Psi_{A,4}\rangle |\Psi_{B,4}\rangle, & \gamma, \delta \in \mathbb{C}, \quad |\gamma|^2 + |\delta|^2 &= 1 \end{aligned}$$

and where we assume for simplicity that the state vectors $|\Psi_{A,i}\rangle$ for the subsystem A are all mutually orthogonal and likewise that the state vectors $|\Psi_{B,i}\rangle$ for the subsystem B are all mutually orthogonal. The partial-trace prescription (50) then yields the reduced density matrix of A ,

$$\begin{aligned} \hat{\rho}_A &= p_\Psi |\alpha|^2 |\Psi_{A,1}\rangle \langle \Psi_{A,1}| + p_\Psi |\beta|^2 |\Psi_{A,1}\rangle \langle \Psi_{A,1}| \\ &\quad + p_\Phi |\gamma|^2 |\Psi_{A,3}\rangle \langle \Psi_{A,3}| + p_\Phi |\delta|^2 |\Psi_{A,4}\rangle \langle \Psi_{A,4}|, \end{aligned}$$

with a similar formula for B , and the instantaneous kinematical relationship (86) yields the conditional probabilities

$$\begin{aligned} p(\Psi_{A,1} | \Psi_{A+B}) &= p_\Psi |\alpha|^2, & p(\Psi_{A,2} | \Psi_{A+B}) &= p_\Psi |\beta|^2, \\ p(\Psi_{A,3} | \Psi_{A+B}) &= 0, & p(\Psi_{A,4} | \Psi_{A+B}) &= 0, \\ p(\Psi_{A,1} | \Phi_{A+B}) &= 0, & p(\Psi_{A,2} | \Phi_{A+B}) &= 0, \\ p(\Psi_{A,3} | \Phi_{A+B}) &= p_\Phi |\gamma|^2, & p(\Psi_{A,4} | \Phi_{A+B}) &= p_\Phi |\delta|^2, \end{aligned}$$

again with similar formulas for B .

The kinematical parent-subsystem formula (86) also allows us to resolve an issue that we brought up in our discussion surrounding (58) in Section III C 5, where we saw that the existence of entanglement between a pair of systems A and B seemed to prevent us from relating the ontic state of the subsystem B to a given ontic state of its parent system $A + B$ even if there existed a *slight redefinition* of our subsystem decomposition $A + B = A' + B'$ for which the parent system's ontic state had the simple non-entangled form (a', b') . In that case, the quantum conditional probability $p_{B|A+B}(b|w = (a', b')) \approx p_{B'|A'+B'}(b'|w = (a', b')) = 1$ obtained from (86) would be very close to unity, thereby smoothing out the supposed discrepancy and, furthermore, eliminating the need to define the subsystem B with measure-zero sharpness.

³⁵ Vermaas [316] refers to the two logical directions underlying these classically reductionist relationships as the property of composition and the property of division.

³⁶ Maudlin implicitly makes this kind of error in criterion 1.A of his three-part classification of interpretations of quantum theory in [232] when he assumes that the metaphysical completeness of state vectors implies that the state vector of a *parent system* completely determines the state vectors of all its *subsystems* even for the case of microscopic systems. Writing in [273] in 1935, Schrödinger was clear about this failure of naïve reductionism in quantum theory: “Maximal knowledge of a total system does not necessarily include total knowledge of all its parts, not even when these are fully separated from each other and at the moment are not influencing each other at all.” It's also worth mentioning that Maudlin's criterion 1.B that state vectors always evolve according to linear dynamical equations is somewhat misleading, because all realistic systems are always at least slightly open and thus must be described by density matrices that do not evolve according to exactly linear dynamical equations.

7. Quantum Dynamics of Open Subsystems

Our discussion of dynamics in quantum theory begins with the assumption of a system W having well-defined quantum dynamics to an acceptable level of approximation over a given time interval $\Delta t \equiv t' - t \geq 0$, an assumption that we take to mean that W approximately admits a well-defined linear CPTP dynamical mapping $\mathcal{E}_W^{t' \leftarrow t}$ in the sense of (74). Note that we include the possibility that $\mathcal{E}_W^{t' \leftarrow t}$ is a unitary mapping, as would be appropriate for the case in which W is a closed system and a good approximation for systems that are sufficiently microscopic and therefore relatively easy to isolate from their environments.

If we can regard this system as being a parent system $W = Q + E$ consisting of a subsystem Q and its larger environment E inside W , then there is no guarantee that Q itself has well-defined dynamics of its own. Indeed, the reduced density matrix of Q at the final time t' is given generally by

$$\hat{\rho}_Q(t') = \text{Tr}_E[\hat{\rho}_W(t')] = \text{Tr}_E[\mathcal{E}_W^{t' \leftarrow t}[\hat{\rho}_W(t)]], \quad (88)$$

which is not generically linear (or even analytic) in the reduced density matrix $\hat{\rho}_Q(t)$ at the initial time t because of the complicated way in which $\hat{\rho}_Q(t)$ is related to $\hat{\rho}_W(t)$ through the partial trace (50).³⁷

We encountered a similar issue in Section III D 1 when we discussed the general relationship between classical parent-system and subsystem dynamics, where we saw in (61) that a classical open subsystem's dynamical mapping is not generically linear. Mimicking our approach in that discussion, we begin by rewriting the subsystem evolution equation (88) as

$$\hat{\rho}_Q(t') = \text{Tr}_E \left[\mathcal{E}_W^{t' \leftarrow t} [\mathcal{A}_{Q \subset W}^t[\hat{\rho}_Q(t)]] \right], \quad (89)$$

where we have introduced the nonlinear, manifestly positive-definite (but not generally trace-preserving) mapping

$$\mathcal{A}_{Q \subset W}^t[\dots] \equiv \hat{\rho}_W^{1/2}(t) \left(\hat{\rho}_Q^{-1/2}(t) \otimes \hat{1}_E \right) \left((\dots) \otimes \hat{1}_E \right) \left(\hat{\rho}_Q^{-1/2}(t) \otimes \hat{1}_E \right) \hat{\rho}_W^{1/2}(t), \quad (90)$$

which is known as an assignment map [250]. The alternative version (89) of our original subsystem evolution equation (88) for Q is the natural noncommutative quantum counterpart to the classical formula (61) and motivates introducing a coarse-grained or effective version $p_{Q \subset W}(\cdot; t' | \cdot; t)$ of our original quantum conditional probabilities (80) for $Q_1 \equiv Q$ inside W in analogy with the classical formula (64):

$$\begin{aligned} p_{Q \subset W}(j; t' | i; t) &\equiv \text{Tr}_W \left[\left(\hat{P}_Q(j; t') \otimes \hat{1}_E \right) \mathcal{E}_W^{t' \leftarrow t} [\mathcal{A}_{Q \subset W}^t[\hat{P}_Q(i; t)]] \right] \\ &= \frac{1}{p_Q(i; t)} \text{Tr}_W \left[\left(\hat{P}_Q(j; t') \otimes \hat{1}_E \right) \mathcal{E}_W^{t' \leftarrow t} \left[\hat{\rho}_W^{1/2}(t) \left(\hat{P}_Q(i; t) \otimes \hat{1}_E \right) \hat{\rho}_W^{1/2}(t) \right] \right]. \end{aligned} \quad (91)$$

Again in parallel with the classical case (65), if initial correlations between Q and its environment E approximately wash out

$$\hat{\rho}_W \approx \hat{\rho}_Q \otimes \hat{\rho}_E \quad (92)$$

over some characteristic time scale δt_Q that is short when compared with the time scales relevant to our present discussion,

$$\delta t_Q \ll t' - t, \quad (93)$$

then

$$\hat{\rho}_W^{1/2}(t) \left(\hat{\rho}_Q^{-1/2}(t) \otimes \hat{1}_E \right) = \hat{1}_Q \otimes \hat{\rho}_E^{1/2}(t)$$

and thus the assignment map (90) trivializes:

$$\mathcal{A}_{Q \subset W}^t[\hat{\rho}_Q(t)] = \hat{\rho}_Q(t) \otimes \hat{\rho}_E(t).$$

³⁷ Arguments (such as in [275]) that the time evolution of generic subsystem density matrices must be linear typically make an implicit assumption that the density matrices in question describe proper mixtures. There is no general proof of linearity in the more general case in which the density matrices under investigation represent improper mixtures, and, indeed, we have already seen from (61) that the analogous classical case features manifest nonlinearity.

It follows as an immediate consequence that over time scales $t' - t \gg \delta t_Q$, the subsystem evolution equation (89) reduces to linear CPTP dynamics for the subsystem Q on its own,

$$\hat{\rho}_Q(t') = \text{Tr}_E \left[\mathcal{E}_W^{t' \leftarrow t} [\hat{\rho}_Q(t) \otimes \hat{\rho}_E] \right] = \sum_{\alpha} \hat{E}_{\alpha}^{t' \leftarrow t} \hat{\rho}_Q(t) \hat{E}_{\alpha}^{t' \leftarrow t \dagger}, \quad \sum_{\alpha} \hat{E}_{\alpha}^{t' \leftarrow t \dagger} \hat{E}_{\alpha}^{t' \leftarrow t} = \hat{1}_Q, \quad (94)$$

where $\{\hat{E}_{\alpha}^{t' \leftarrow t}\}_{\alpha}$ is the usual set of Kraus operators for the dynamics [204]. Indeed, this line of reasoning, together with the additional assumptions that the linear CPTP density-matrix dynamics for Q is Markovian and homogeneous in time, precisely lead to the well-known and highly effective Lindblad equation [219], which can be expressed in its diagonal form as [67, 196, 275]

$$\frac{\partial \hat{\rho}_Q}{\partial t} = -\frac{i}{\hbar} [\hat{H}, \hat{\rho}_Q] + \sum_{k=1}^{N^2-1} \gamma_k \left(\hat{A}_k \hat{\rho}_Q \hat{A}_k^{\dagger} - \frac{1}{2} \hat{A}_k^{\dagger} \hat{A}_k \hat{\rho}_Q - \frac{1}{2} \hat{\rho}_Q \hat{A}_k^{\dagger} \hat{A}_k \right) \quad (95)$$

for a suitable collection of parameters γ_k and operators \hat{H}, \hat{A}_k . Letting $\mathcal{E}_Q^{t' \leftarrow t}$ denote this reduced linear CPTP dynamics for Q alone,

$$\hat{\rho}_Q(t) \mapsto \hat{\rho}_Q(t') = \mathcal{E}_Q^{t' \leftarrow t} [\hat{\rho}_Q(t)], \quad (96)$$

the coarse-grained conditional probabilities (91) now coincide with our *exact* general definition (80) for $W = Q_1 \equiv Q$:

$$p_{Q \subset W}(j; t' | i; t) \rightarrow p_Q(j; t' | i; t) \equiv \text{Tr}_Q[\hat{P}_Q(j; t') \mathcal{E}_Q^{t' \leftarrow t}[\hat{P}_Q(i; t)]] \sim \text{Tr}[\hat{P}_j(t') \mathcal{E}[\hat{P}_i(t)]] \text{ for } t' - t \gg \delta t_Q. \quad (97)$$

In other words, by coarse-graining *in time* on the scale δt_Q over which correlations between the subsystem Q and its environment E wash out, we no longer need to coarse grain in the sense (91) of explicitly referring to the parent system $W = Q + E$.

As we remarked in the classical case, the characteristic time scale δt_Q determines a natural short-time cutoff on the existence of dynamics for our open subsystem Q . Notice also that our interpretation of quantum theory can fully accommodate the possibility that the parent system W likewise has a nonzero characteristic time scale $\delta t_W \neq 0$ that sets the cutoff on the existence of its own dynamics as well; for all we know, there is no maximal parent system in Nature for which this temporal cutoff scale exactly vanishes. That is, it may well be that *all* linear CPTP dynamics is ultimately only an approximate notion, although the same might well be true for density matrices themselves, as we explained in the context of discussing truncated Hilbert spaces in Section III C 6.

8. Dynamical Relationships Between Ontic States Over Time and Objective Epistemic States Over Time

Whether the dynamical quantum conditional probabilities $p_Q(j; t' | i; t)$ are exact

$$p_Q(j; t' | i; t) \equiv \text{Tr}_Q[\hat{P}_Q(j; t') \mathcal{E}_Q^{t' \leftarrow t}[\hat{P}_Q(i; t)]] \sim \text{Tr}[\hat{P}_j(t') \mathcal{E}[\hat{P}_i(t)]] \quad (98)$$

or coarse-grained either in the sense (91) of making explicit reference to a parent system or in the sense (97) of existing only over time scales exceeding some nonzero temporal cutoff δt_Q as introduced in (93), they manifestly define dynamics for the ontic states of Q in a manner that parallels the classical case (59).³⁸

$$p_Q(\cdot; t' | \cdot; t): \underbrace{(i; t)}_{\text{initial data}}, \underbrace{(j; t')}_{\text{final data}} \mapsto \underbrace{p_Q(j; t' | i; t)}_{\text{conditional probabilities}}. \quad (99)$$

The dynamical quantum conditional probabilities $p_Q(j; t' | i; t)$ also provide an explicit dictionary that translates between the manifestly quantum linear CPTP dynamics (96) for density matrices and the linear dynamical mapping of objective epistemic states analogous to the equation (60) for classical epistemic states:

$$p_Q(\cdot; t' | \cdot; t): \underbrace{p_Q(j; t')}_{\text{epistemic state at } t'} = \sum_i \underbrace{p_Q(j; t' | i; t)}_{\text{conditional probability (independent of epistemic states)}} \underbrace{p_Q(i; t)}_{\text{epistemic state at } t}. \quad (100)$$

³⁸ These dynamical quantum conditional probabilities also supply an important ingredient that is missing from the traditional modal interpretations and that is identified in “Problem 3: The problem of effect” in Maudlin’s paper [232], namely, a “detailed dynamics for the value [ontic] states.”

9. Entanglement and Imprecision in Coarse-Grained Quantum Conditional Probabilities

When working over time intervals so small that we *cannot* assume approximate factorization of the density matrix $\hat{\rho}_W \not\approx \hat{\rho}_Q \otimes \hat{\rho}_E$ of the parent system $W = Q + E$, we claim that quantum entanglement (56)—and not merely classical correlation (55)—between Q and its environment E inside W determines the size of our imprecision in using the coarse-grained quantum conditional probabilities $p_{Q \subset W}(j; t' | i; t)$ defined in (91).

As evidence in support of this claim, suppose that the dynamics for the parent system W over the time interval $t' - t$ exactly decouples into dynamics for the subsystem Q alone and dynamics for the environment E alone, with no interactions between the two subsystems, so that the linear CPTP dynamical mapping (74) for W factorizes:

$$\mathcal{E}_W^{t' \leftarrow t} = \mathcal{E}_Q^{t' \leftarrow t} \otimes \mathcal{E}_E^{t' \leftarrow t}. \quad (101)$$

(This circumstance includes the trivial case $t' = t$ in which $\mathcal{E}_W^{t' \leftarrow t}$ is just the identity mapping, as well as the case in which all these dynamical mappings describe unitary time evolution.) If Q and E are only classically correlated with each other at the initial time t but are not entangled in the sense of (56), then the density matrix $\hat{\rho}_W(t)$ of the parent system W takes the form

$$\hat{\rho}_W(t) = \sum_{i,e} p_W((i,e); t) \hat{P}_Q(i; t) \otimes \hat{P}_E(e; t), \quad (102)$$

and a simple calculation shows that the coarse-grained quantum conditional probabilities $p_{Q \subset W}(j; t' | i; t)$ defined in (91) reduce to the expected *exact* quantum conditional probabilities $p_Q(j; t' | i; t)$ defined in (98) for Q alone:

$$p_{Q \subset W}(j; t' | i; t) = \text{Tr}_Q \left[\hat{P}_Q(j; t') \mathcal{E}_Q^{t' \leftarrow t} [\hat{P}_Q(i; t)] \right] \equiv p_Q(j; t' | i; t). \quad (103)$$

Within the scope of our assumption of decoupled dynamics (101), deviations from (103) can therefore arise only if Q is entangled with its environment E . We interpret this result as implying that for general parent-system dynamics not necessarily factorizing (101) into independent dynamics for Q and E , we should only ever trust the validity of coarse-grained conditional probabilities (91) up to an intrinsic error—a new kind of “uncertainty principle”—determined by the amount of entanglement between Q and E .

10. Eigenstate Swaps

Recall that an eigenstate swap (36), as we have defined it in Section III B 3, describes a hypothetical exchange between two orthogonal density-matrix eigenstates over a time scale δt_{swap} that is exponentially small in the total number of degrees of freedom of both the system itself and of all other systems that substantially interact and entangle with it, in keeping with our formula (37) for δt_{swap} .

We can now explain why actual ontic states avoid eigenstate swaps: If an eigenstate swap in the system’s density matrix takes place during the time interval from t to $t + \delta t_{\text{swap}}$, where δt_{swap} is assumed to be larger than the minimal time scale δt_Q over which the dynamics of our system Q actually exists (an eigenstate swap has no meaning for $\delta t_{\text{swap}} < \delta t_Q$), then, in accordance with our formula (97) for the quantum conditional probabilities for Q and given that the system’s actual ontic state was $\Psi_{Q,i}(t)$ at the earlier time t , there is approximately zero conditional probability

$$p_Q(i; t + \delta t_{\text{swap}} | i; t) \sim |\langle \Psi_{Q,i}(t + \delta t_{\text{swap}}) | \Psi_{Q,i}(t) \rangle|^2 \approx 0$$

for the system’s ontic state to be

$$\Psi_{Q,i}(t + \delta t_{\text{swap}}) \approx \perp \Psi_{Q,i}(t)$$

at the later time $t + \delta t_{\text{swap}}$.³⁹ Indeed, the probability that the system’s actual ontic state will instead be

$$\Psi_{Q,j}(t + \delta t_{\text{swap}}) \approx \Psi_{Q,i}(t)$$

³⁹ Observe that the operation $\mathcal{E}_Q^{t' \leftarrow t}[\hat{P}_Q(i; t)]$ appearing in (97) evolves ontic states as though they were not part of density matrices, and so is blind to the eigenstate swap in the system’s density matrix. Note also that over long time scales $t' - t \gg \delta t_{\text{swap}} > \delta t_Q$, there can be appreciable conditional probabilities for the final ontic state of a system to end up being nearly orthogonal to its initial ontic state. That is, the present arguments eliminate only *ultra-fast* swaps to orthogonal ontic states, but allow for ordinary transitions that take place over reasonably long time intervals.

is approximately unity. Notice again the smoothing role played by our quantum conditional probabilities: They sew together the evolving ontology of Q in a manner that avoids eigenstate-swap instabilities.

These results may seem surprising given that $\Psi_{Q,i}(t)$ is *nearly certain* to become $\Psi_{Q,i}(t')$ over *even shorter* time scales $t' - t \ll \delta t_{\text{swap}}$, but it is important to keep in mind that our objective quantum conditional probabilities do not naïvely compose: We cannot generically express $p_Q(i; t + \delta t_{\text{swap}} | i; t)$, defined in accordance with (97), as a sum of products of quantum conditional probabilities of the form $p_Q(i_m; t_m | i_{m-1}; t_{m-1})$ over a sequence of many tiny intermediate time intervals $t_m - t_{m-1} \ll \delta t_{\text{swap}}$ going from t to $t + \delta t_{\text{swap}}$.

11. Ergodicity Breaking, Classical States, and the Emergence of Statistical Mechanics

On reasonably short time scales, our dynamical quantum conditional probabilities (98) generically allow a macroscopic system’s ontic-level trajectory to explore ergodically a large set of ontic states differing in the values of only a few of the system’s degrees of freedom, while super-exponentially suppressing transitions to ontic states that differ in the values of large numbers of degrees of freedom. Hence, the dynamical formula (98) leads emergently to a partitioning of the macroscopic system’s overall Hilbert space into distinct ergodic components that do not mix over short time intervals and that we can therefore identify as the system’s classical states, much in the way that a ferromagnet undergoes a phase transition from a single disordered ergodic component to a collection of distinct ordered ergodic components as we cool the system below its Curie temperature. For a treatment of ergodicity breaking and the related emergence of statistical mechanics from quantum considerations, see [183, 222, 261, 262].⁴⁰

12. Connections to Other Work

In a spirit reminiscent of our minimal modal interpretation of quantum theory, a series of approaches [74, 149] used in studying open quantum systems involve “unraveling” the dynamical equation for a given subsystem’s reduced density matrix as an ensemble average over a collection of stochastically evolving state vectors. Such techniques have been applied to quantum optics [74, 90], entanglement [76], decoherence [167], non-Markovian dynamics [100, 101, 133, 292], and geometrical phases [40].

Similarly, in [110], Esposito and Mukamel make use of a notion of “quantum trajectories” not unlike the evolving ontic-state trajectories in our interpretation of quantum theory. Esposito and Mukamel likewise note that defining these quantum trajectories for a given system Q requires working with the time-dependent eigenbasis that instantaneously diagonalizes the system’s density matrix $\hat{\rho}_Q(t)$, in contrast to the *fixed* sample space of a classical system. Moreover, defining a *differential* linear CPTP dynamical mapping \mathcal{K}_Q^t in terms of its finite-time counterpart $\mathcal{E}_Q^{t'-t}$ from (96) according to

$$\mathcal{K}_Q^t \equiv \frac{\mathcal{E}_Q^{t+\delta t_Q \leftarrow t} - \text{id}}{\delta t_Q}, \quad (104)$$

where δt_Q is the minimal time scale (93) over which our system’s dynamics (97) exists, we can naturally re-express our dynamical quantum conditional probabilities (98),

$$p_Q(j; t' | i; t) \equiv \text{Tr}_Q \left[\hat{P}_Q(j; t') \mathcal{E}_Q^{t'-t} [\hat{P}_Q(i; t)] \right],$$

in terms of a set of quantum transition rates

$$W_Q((j|i); t) \equiv \frac{p_Q(j; t + \delta t_Q | i; t) - p_Q(j; t | i; t)}{\delta t_Q} = \text{Tr}_Q \left[\hat{P}_Q(j; t') \mathcal{K}_Q^t [\hat{P}_Q(i; t)] \right] \quad (105)$$

that coincide in the formal limit $\delta t_Q \rightarrow 0$ with Esposito and Mukamel’s own definition of quantum transition rates. Esposito, Mukamel, and others [199, 214] use these quantum trajectories and quantum transition rates to study quantum definitions of work and heat, entropy production, fluctuation theorems, and other fundamental questions in statistical mechanics and thermodynamics.

⁴⁰ In particular, ergodicity breaking is crucial for explaining the complications that the authors of [222] encounter with their assumption that an open subsystem’s final equilibrium state should satisfy “subsystem state independence”—that is, that the open subsystem’s final equilibrium state should be independent of the subsystem’s initial state. This initial-state sensitivity also plays an important role in obtaining a robust resolution of the measurement problem for realistic measurement devices, an issue that arose in discussions with the author of [183] in the context of a similar modal interpretation of quantum theory developed concurrently with our own modal interpretation; we compare these two modal interpretations in Section VI G.

Our quantum conditional probabilities, in their dynamical manifestation (98) for a system Q having well-defined dynamics, are also closely related to the causal quantum conditional states of Leifer and Spekkens [215]. To introduce the Leifer-Spekkens construction and to explain this purported connection to our minimal modal interpretation of quantum theory, we begin by considering an arbitrary linear CPTP dynamical mapping $\mathcal{E}_Q^{t' \leftarrow t}$ of the form (96) for a system Q ,

$$\hat{\rho}_Q(t) \mapsto \hat{\rho}_Q(t') = \mathcal{E}_Q^{t' \leftarrow t} [\hat{\rho}_Q(t)].$$

Then, formally introducing a tensor-product Hilbert space $\mathcal{H}_{Q(t')} \otimes \mathcal{H}_{Q(t)}$ consisting of two time-separated copies of the Hilbert space of Q , the Choi-Jamiołkowski isomorphism [81, 82, 190], also known as the channel-state duality, implies the existence of a unique operator $\hat{\varrho}_{Q(t')|Q(t)}$ on $\mathcal{H}_{Q(t')} \otimes \mathcal{H}_{Q(t)}$ that implements precisely the same time evolution for Q through a quantum generalization of the classical Bayesian propagation rule $p_Q(j; t') = \sum_i p_Q(j; t' | i; t) p_Q(i; t)$, namely,

$$\hat{\rho}_Q(t') = \text{Tr}_{Q(t)} [\hat{\varrho}_{Q(t')|Q(t)} (\hat{1}_{Q(t')} \otimes \hat{\rho}_Q(t))], \quad (106)$$

where $\hat{1}_{Q(t')}$ is the identity operator on $\mathcal{H}_{Q(t')}$. Leifer and Spekkens regard the operator $\hat{\varrho}_{Q(t')|Q(t)}$, which they call a causal quantum conditional state, as being a noncommutative generalization of classical conditional probabilities, in much the same way that density matrices $\hat{\rho}_Q$ themselves serve as noncommutative generalizations of classical probability distributions p_Q .

Our dynamical quantum conditional probabilities (98) are then precisely the *diagonal* matrix elements of the Leifer-Spekkens quantum conditional state $\hat{\varrho}_{Q(t')|Q(t)}$ in the tensor-product basis $|\Psi_{Q,j}(t')\rangle \otimes |\Psi_{Q,i}(t)\rangle$ for $\mathcal{H}_{Q(t')} \otimes \mathcal{H}_{Q(t)}$ constructed out of the eigenstates $|\Psi_{Q,j}(t')\rangle$ of $\hat{\rho}_Q(t')$ and the eigenstates $|\Psi_{Q,i}(t)\rangle$ of $\hat{\rho}_Q(t)$:

$$\begin{aligned} p_Q(j; t' | i; t) &= \text{Tr}_{Q(t')} [\hat{P}_Q(j; t') \hat{\rho}_Q(t')] \\ &= \text{Tr}_{Q(t')} \left[\hat{P}_Q(j; t') \text{Tr}_{Q(t)} \left[\hat{\varrho}_{Q(t')|Q(t)} (\hat{1}_{Q(t')} \otimes \hat{P}_Q(i; t)) \right] \right] \\ &= (\langle \Psi_{Q,j}(t') | \otimes \langle \Psi_{Q,i}(t) |) \hat{\varrho}_{Q(t')|Q(t)} (| \Psi_{Q,j}(t') \rangle \otimes | \Psi_{Q,i}(t) \rangle). \end{aligned} \quad (107)$$

This result mirrors the way in which our interpretation of quantum theory identifies the diagonal matrix elements of a system's density matrix $\hat{\rho}_Q$ in its own eigenbasis as being the system's epistemic probabilities p_Q .

The connections between our interpretation of quantum theory and the approach of Leifer and Spekkens go deeper. For a system Q without well-defined dynamics of its own but belonging to a parent system $W = Q + E$ having its own linear CPTP dynamical mapping $\mathcal{E}_W^{t' \leftarrow t}$ as in (74),

$$\hat{\rho}_W(t) \mapsto \hat{\rho}_W(t') = \mathcal{E}_W^{t' \leftarrow t} [\hat{\rho}_W(t)],$$

we introduced coarse-grained quantum conditional probabilities $p_{Q \subset W}(j; t' | i; t)$ for Q according to (91),

$$p_{Q \subset W}(j; t' | i; t) \equiv \text{Tr}_W \left[\left(\hat{P}_Q(j; t') \otimes \hat{1}_E \right) \mathcal{E}_W^{t' \leftarrow t} [\mathcal{A}_{Q \subset W}^t [\hat{P}_Q(i; t)]] \right],$$

where we defined the nonlinear assignment map $\mathcal{A}_{Q \subset W}^t$ in (90),

$$\mathcal{A}_{Q \subset W}^t [\dots] \equiv \hat{\rho}_W^{1/2}(t) \left(\hat{\rho}_Q^{-1/2}(t) \otimes \hat{1}_E \right) ((\dots) \otimes \hat{1}_E) \left(\hat{\rho}_Q^{-1/2}(t) \otimes \hat{1}_E \right) \hat{\rho}_W^{1/2}(t).$$

It follows from a straightforward computation that we can equivalently express our coarse-grained quantum conditional probabilities $p_{Q \subset W}(j; t' | i; t)$ as the diagonal matrix elements of a corresponding coarse-grained quantum conditional state $\hat{\varrho}_{Q(t')|Q(t)}^{Q \subset W}$,

$$p_{Q \subset W}(j; t' | i; t) = \text{Tr}_{Q(t')} \left[\hat{P}_{Q(t'),j} \text{Tr}_{Q(t)} \left[\hat{\varrho}_{Q(t')|Q(t)}^{Q \subset W} (\hat{1}_{Q(t')} \otimes \hat{P}_{Q(t),i}) \right] \right]. \quad (108)$$

where $\hat{\varrho}_{Q(t')|Q(t)}^{Q \subset W}$ is an operator on the formal tensor-product Hilbert space $\mathcal{H}_{Q(t')} \otimes \mathcal{H}_{Q(t)}$ that we have defined by starting with the *exact* Leifer-Spekkens quantum conditional state $\hat{\varrho}_{W(t')|W(t)}$ of the parent system W , multiplying by the density matrix $\hat{\rho}_W(t)$ of W at the initial time t to obtain a natural causal quantum *joint* state $\hat{\varrho}_{W(t')+W(t)}$ for W , partial tracing over the environment E at both the initial time t and the final time t' to obtain a coarse-grained

causal quantum joint state $\hat{\rho}_{Q(t')+Q(t)}^{Q \subset W}$ for our original subsystem Q , and then “marginalizing” on Q at the initial time t by multiplying by its inverse density matrix $\hat{\rho}_Q^{-1}(t)$:

$$\begin{aligned} & \hat{\rho}_{Q(t')|Q(t)}^{Q \subset W} \\ &= \left(\hat{\mathbb{1}}_{Q(t')} \otimes \hat{\rho}_Q^{-1/2}(t) \right) \text{Tr}_{E(t')} \left[\text{Tr}_{E(t)} \left[\left(\hat{\mathbb{1}}_{W(t')} \otimes \hat{\rho}_W^{1/2}(t) \right) \hat{\rho}_{W(t')|W(t)} \left(\hat{\mathbb{1}}_{W(t')} \otimes \hat{\rho}_W^{1/2}(t) \right) \right] \right] \left(\hat{\mathbb{1}}_{Q(t')} \otimes \hat{\rho}_Q^{-1/2}(t) \right). \end{aligned} \quad (109)$$

As expected, this coarse-grained quantum conditional state exactly satisfies a propagation rule of the form (106):

$$\hat{\rho}_Q(t') = \text{Tr}_{Q(t)} \left[\hat{\rho}_{Q(t')|Q(t)}^{Q \subset W} \left(\hat{\mathbb{1}}_{Q(t')} \otimes \hat{\rho}_Q(t) \right) \right]. \quad (110)$$

E. Subsystem Spaces

An interesting geometrical structure emerges if we consider the continuously infinite set of different ways that we can define a particular subsystem Q with its own Hilbert space \mathcal{H}_Q by tensor-product-factorizing the Hilbert space $\mathcal{H}_W = \mathcal{H}_Q \otimes \mathcal{H}_E$ of a given parent system $W = Q + E$ that includes an environment E . The resulting geometrical structure, which we call the subsystem space for Q in W and which encompasses all the various possible versions of \mathcal{H}_Q , exists even for a parent system W with as few as four mutually exclusive states ($\dim \mathcal{H}_W = 4$) and has no obvious counterpart in classical physics.

1. Formal Construction

We parameterize the smooth family of choices of bipartite tensor-product factorization of the parent system’s Hilbert space \mathcal{H}_W using an n -tuple of complex numbers $\alpha = (\alpha_1, \dots, \alpha_n)$ for some integer n , writing $\mathcal{H}_W = \mathcal{H}_{Q(\alpha)} \otimes \mathcal{H}_{E(\alpha)}$. We thereby obtain a complex vector bundle—which we call the subsystem space for our subsystem Q of interest—that consists of an identical Hilbert space $\mathcal{H}_{Q(\alpha)} \cong \mathcal{H}_Q$ of some fixed dimension $\dim \mathcal{H}_Q$ attached to each point with coordinates α on a base manifold of complex dimension n . Each *specific* choice of tensor-product factorization $\mathcal{H}_W = \mathcal{H}_{Q(\alpha)} \otimes \mathcal{H}_{E(\alpha)}$ and partial trace down to $\mathcal{H}_{Q(\alpha)} \cong \mathcal{H}_Q$ then corresponds to just *one version* of our subsystem Q of interest, or, equivalently, to one fiber of this complex vector bundle—that is, to one “slice” of the subsystem space for our subsystem of interest. Because changes to our choice of tensor-product decomposition correspond to unitary transformations of the tensor-product basis for \mathcal{H}_W modulo unitary transformations of the individual tensor-product factors, we see that the resulting subsystem space is therefore essentially isomorphic to the quotient space $\mathcal{U}_W / (\mathcal{U}_Q \times \mathcal{U}_E)$, where \mathcal{U}_W , \mathcal{U}_Q , and \mathcal{U}_E are respectively the sets of unitary transformations acting on the Hilbert spaces \mathcal{H}_W , \mathcal{H}_Q , and \mathcal{H}_E .⁴¹

The parent system’s density matrix $\hat{\rho}_W$ on $\mathcal{H}_W = \mathcal{H}_{Q(\alpha)} \otimes \mathcal{H}_{E(\alpha)}$ then defines a natural Hermitian inner product between any pair of vectors $|\psi_{Q(\alpha)}\rangle \in \mathcal{H}_{Q(\alpha)}$ and $|\chi_{Q(\alpha')}\rangle \in \mathcal{H}_{Q(\alpha')}$ living in two fibers respectively located at α and α' :

$$\begin{aligned} h(\psi_{Q(\alpha)}, \chi_{Q(\alpha')}) &\equiv \text{Tr}_W \left[\hat{\rho}_W \left(|\psi_{Q(\alpha)}\rangle \langle \psi_{Q(\alpha)}| \otimes \hat{\mathbb{1}}_{E(\alpha)} \right) \left(|\chi_{Q(\alpha')}\rangle \langle \chi_{Q(\alpha')}| \otimes \hat{\mathbb{1}}_{E(\alpha')} \right) \right] \\ &= h(\chi_{Q(\alpha')}, \psi_{Q(\alpha)})^*. \end{aligned} \quad (111)$$

Here $\hat{\mathbb{1}}_{E(\alpha)}$ and $\hat{\mathbb{1}}_{E(\alpha')}$ are respectively the identity operators on the tensor-product factors $\mathcal{H}_{E(\alpha)}$ and $\mathcal{H}_{E(\alpha')}$.

2. Actual Ontic Layer of Reality

The smooth family of reduced density matrices $\hat{\rho}_{Q(\alpha)}$ that we can obtain by partial tracing $\hat{\rho}_W$ for different choices of α defines a set of $N = \dim \mathcal{H}_Q$ smooth sections $|\Psi_{Q(\alpha),i}\rangle$ ($i = 1, \dots, N = \dim \mathcal{H}_Q$) on this complex vector bundle. The notion of the subsystem’s actual ontic state therefore extends to a whole “actual ontic layer of reality” as we move through the different values of α , but we *cannot* generically identify this actual ontic layer of reality as a whole

⁴¹ After this work was substantially complete, we noticed that a similar idea appears in [186], where the authors parameterize their subsystem space using a coordinate label θ and employ it to study $\exp(-S)$ breakdowns in locality in the presence of a black hole of entropy S .




Figure 6. The schematic set-up for a Von Neumann experiment, consisting of a subject system Q together with a measurement apparatus A and a larger environment E .

with just a *single* section $|\Psi_{Q(\alpha),i}\rangle$ for fixed i , for the simple reason that the epistemic probabilities $p_{Q(\alpha),i}$ across this section may vary considerably from one value of α to another. However, our kinematical quantum conditional probabilities (86) ensure that if the actual ontic state of W is of the non-entangled form $(\Psi_{Q(\alpha),i}, \Psi_{E(\alpha),e})$ for a particular choice of α , then all versions of our subsystem Q near this value of α will be nearly certain to be in ontic states very similar to $\Psi_{Q(\alpha),i}$ in the Hilbert-space sense.

IV. THE MEASUREMENT PROCESS

A. Measurements and Decoherence

It is not our aim in this paper to investigate the measurement process and decoherence in any great depth, or to study realistic examples. For pedagogical treatments, see [58, 67, 196]. For additional detailed discussions, including a careful study of the extremely rapid rate of decoherence for typical systems in contact with macroscopic measurement devices or environments, see, for example, [197, 269, 270, 348, 349]. For an analysis of how realistic measuring devices with finite spatial and temporal resolution solve various delocalization [248] and instability [13–15] issues, and, in particular, for a treatment of the important role played both by environmental interactions and by ergodicity breaking in making sense of realistic measurement processes, see, for example, [31, 182, 183, 316].⁴²

B. Von Neumann Measurements

1. Motivation

For the purposes of establishing how our minimal modal interpretation makes sense of measurements, how decoherence turns the environment into a “many-dimensional chisel” that rapidly sculpts the ontic states of systems into their precise shapes,⁴³ and how the Born rule naturally emerges to an excellent approximation, we consider the idealized example of a so-called Von Neumann measurement. Along the way, we will address the status of both the measurement problem generally and the notion of wave-function collapse specifically in the context of our interpretation of quantum theory. Ultimately, we’ll find that our interpretation solves the measurement problem by replacing the textbook Copenhagen interpretation’s instantaneous axiomatic wave-function collapse with an interpolating ontic-level dynamics and thereby eliminates the need for an *ad hoc* Heisenberg cut.

⁴² We discuss the modal interpretation of [182, 183], which developed concurrently with our own interpretation of quantum theory, in Section VI G.

⁴³ The way that decoherence sculpts ontic states into shape is reminiscent of the way that the external pressure of air molecules above a basin of water maintains the water in its liquid phase.

2. Basic Set-Up

A Von Neumann measurement involves a subject system Q with a complete basis of mutually exclusive states $Q(i)$ (with $\langle Q(i)|Q(j)\rangle = \delta_{ij}$), a macroscopic measurement apparatus A in an initial actual ontic state $A("0")$ with its display dial set to “blank,” and an even more macroscopic environment E in an initial actual ontic state $E("0")$ in which it sees the dial on the apparatus A as being blank. (See Figure 6.) We suppose moreover that the time evolution is linear (an idealized assumption that could be strengthened further to require that the evolution be unitary) and that in the special case in which the actual ontic state of the subject system Q is precisely one of the definite states $Q(i)$, the time evolution proceeds according to the following idealized two-step sequence:

1. The apparatus A transitions to a new actual ontic state $A("i")$ in which its dial now displays the value “the subject system Q appears to be in the state $Q(i)$,” and then
2. the environment E subsequently changes to a new actual ontic state $E("i")$ in which its own configuration has been noticeably perturbed by the change in the apparatus A , such as through unavoidable and irreversible thermal radiation [165].

Written directly in terms of the appropriate state vectors, this two-step time evolution takes the form

$$\left. \begin{aligned} |Q(i)\rangle |A("0")\rangle |E("0")\rangle &\mapsto |Q(i)\rangle |A("i")\rangle |E("0")\rangle \\ &\mapsto |Q(i)\rangle |A("i")\rangle |E("i")\rangle. \end{aligned} \right\} \quad (112)$$

The different final ontic state vectors $|A("i")\rangle$ of the apparatus are generally not *exactly* mutually orthogonal for different values of i , and, similarly, the different final ontic state vectors $|E("i")\rangle$ of the environment E are generally not *exactly* mutually orthogonal. However, our assumption that the apparatus and the environment are macroscopic systems with large respective Hilbert spaces \mathcal{H}_A and \mathcal{H}_E ensures that inner products of the form $\langle A("i")|A("j")\rangle$ and $\langle E("i")|E("j")\rangle$ for $i \neq j$ are exponentially small in their respective numbers $\gg 10^{23}$ of degrees of freedom:

$$\langle A("i")|A("j")\rangle \sim \exp(-(\#\text{apparatus degrees of freedom}) \times \Delta t/\tau_A), \quad (113)$$

$$\langle E("i")|E("j")\rangle \sim \exp(-(\#\text{environment degrees of freedom}) \times \Delta t/\tau_E). \quad (114)$$

Here Δt is the time duration of the measurement process (112), τ_A is the characteristic time scale for the evolution of each individual degree of freedom of the apparatus A , and τ_E is similarly the characteristic time scale for the evolution of each individual degree of freedom of the environment E .

If we now instead arrange for the initial actual ontic state Ψ of the subject system Q to be a general superposition of the complete orthonormal basis of states $Q(i)$,

$$|\Psi\rangle = \sum_i \alpha_i |Q(i)\rangle = \alpha_1 |Q(1)\rangle + \alpha_2 |Q(2)\rangle + \dots, \quad \langle \Psi|\Psi\rangle = 1, \quad (115)$$

then the assumed linearity of the time evolution (112) implies the temporal sequence

$$\left. \begin{aligned} |\Psi\rangle |A("0")\rangle |E("0")\rangle &= \left(\sum_i \alpha_i |Q(i)\rangle \right) |A("0")\rangle |E("0")\rangle \mapsto \left(\sum_i \alpha_i |Q(i)\rangle |A("i")\rangle \right) |E("0")\rangle \\ &\mapsto \sum_i \alpha_i |Q(i)\rangle |A("i")\rangle |E("i")\rangle. \end{aligned} \right\} \quad (116)$$

The final density matrices of the subject system Q and of the measurement apparatus A , obtained by appropriately partial tracing (50) the composite final state appearing in (116), are respectively given in their diagonal form by

$$\hat{\rho}_Q = \sum_i |\alpha_i|^2 |Q(i)'\rangle \langle Q(i)'|, \quad (117)$$

$$\hat{\rho}_A = \sum_i |\alpha_i''|^2 |A("i")'\rangle \langle A("i")'| + (\text{small } \perp \text{ terms}), \quad (118)$$

where the primes indicate the *very tiny* changes in basis needed to diagonalize the two reduced density matrices and thereby eliminate exponentially small off-diagonal corrections of order the product of (113) with (114),

$$\text{off-diagonal corrections} \sim \exp\left(-\left(\frac{\#\text{apparatus} + \text{environment}}{\text{degrees of freedom}}\right) \times \Delta t/\tau\right), \quad (119)$$

where $\tau \sim \min(\tau_A, \tau_E)$. The *double* primes appearing on the eigenvalues of $\hat{\rho}_A$ in (118) account for the possibility that they might not precisely agree with the eigenvalues of $\hat{\rho}_Q$, again due to discrepancies of order (119). No realistic ontic basis is infinitely sharp anyway in the context of our interpretation of quantum theory, and certainly no sharper than these sorts of corrections.

3. Decoherence, Wave-Function Collapse, and the Effective Born Rule

The exponential speed with which the corrections (119) become negligibly small—a phenomenon called decoherence of the originally coherent superposition (115)—is perfectly in keeping with the notion that a large apparatus conducting a measurement while sitting in an even larger environment leads to the appearance of the “instantaneous” wave-function collapse of the textbook Copenhagen interpretation, and indeed *justifies* the use of wave-function collapse as a heuristic shortcut in elementary pedagogical treatments of the measurement process in quantum theory.⁴⁴

In keeping with the basic correspondence (29) between objective epistemic states and density matrices at the heart of our interpretation of quantum theory, we can then conclude that the final actual ontic state of the subject system Q is one of the possibilities $Q(i)' \approx Q(i)$ with epistemic probability given by the corresponding density-matrix eigenvalue $p_{Q,i} = |\alpha'_i|^2 \approx |\langle Q(i) | \Psi \rangle|^2$, up to exponentially small corrections (119). Similarly, the final actual ontic state of the measurement apparatus A is one of the possibilities $A(i)'' \approx A(i)$ with approximately the *same* epistemic probability $p_{A,i} = |\alpha''_i|^2 \approx |\langle Q(i) | \Psi \rangle|^2$, again up to exponentially small corrections (119). We therefore see that our interpretation of quantum theory solves the measurement problem by replacing axiomatic wave-function collapse with an underlying interpolating evolution for ontic states.

It is widely accepted that measurement-induced decoherence destroys quantum interference between distinct measurement outcomes and produces final objective density matrices for both subject systems and apparatuses that closely resemble the subjective mixtures that result from a so-called Von Neumann-Lüders projection [227, 323] without post-selection. (See Appendix 1b.) Because the latter operation constitutes a linear CPTP dynamical mapping, we see immediately that well-executed measurements by macroscopic apparatuses can be described to great accuracy by linear CPTP dynamical mappings as well. Hence, to the extent that the measurement process experienced by the subject system Q can be captured to an acceptable level of approximation by a linear CPTP dynamical mapping $\mathcal{E}_Q^{t+\Delta t \leftarrow t}$ for Q over the relevant time scale Δt of the measurement, we directly find that the stochastic dynamical conditional probability (98) for Q to end up in the final ontic state $Q(i)' \approx Q(i)$ at the time $t + \Delta t$ given that it was initially in the ontic state Ψ at the time t is $p_Q(i; t + \Delta t | \Psi; t) \approx |\langle Q(i) | \Psi \rangle|^2$, as expected. Similarly, if we can approximate the evolution of the apparatus A as a linear CPTP dynamical mapping, then $p_A(i; t + \Delta t | \emptyset; t) \approx |\langle Q(i) | \Psi \rangle|^2$.

By any of these approaches, the famous Born rule

$$\text{Prob}(\dots) = |\dots|^2 \quad (120)$$

for computing empirical outcome probabilities therefore emerges *automatically* to an excellent approximation without having to be assumed *a priori*.

Notice that in line with the central correspondence (29) between objective epistemic states and density matrices, and in accordance with the boxed statement (31) below it, the quantities $|\alpha_i|^2$ do not have an interpretation as *literal* epistemic probabilities for the underlying actual ontic states until they (or some approximate versions of them) have become *actualized* as eigenvalues of density matrices—that is, not until after the measurement is actually complete. Again, the *only* quantities that our minimal modal interpretation of quantum theory regards as being literal epistemic probabilities are those quantities that our interpretation’s axioms explicitly identify as such.

We can also go beyond the Born rule for Q and A as *individual* subsystems and, importantly, say something *jointly* about their actual ontic states. The post-measurement reduced density matrix $\hat{\rho}_{Q+A}$ of the composite system $Q + A$ consisting of the subject system together with the apparatus, computed according to the standard partial-trace operation (50), is not generically diagonal in exactly the “correct” tensor-product basis $|Q(j)', A(i)' \rangle \equiv$

⁴⁴ Typical decoherence time scales for many familiar systems are presented in [197, 269, 270, 349].

$|Q(j)'\rangle \otimes |A("i")'\rangle$ consistent with our original definitions of the subject system Q and the apparatus A as subsystems. Nonetheless, the diagonalizing basis for $\hat{\rho}_{Q+A}$ is *exponentially close* to this tensor-product basis,

$$\hat{\rho}_{Q+A} = \sum_{i,j} \alpha_i \alpha_j^* \underbrace{(\langle E("i") | E("j") \rangle)}_{\delta_{ij} + \left(\begin{smallmatrix} \text{exponentially} \\ \text{small} \\ \text{corrections} \end{smallmatrix} \right)} |Q(i)\rangle |A("i")\rangle \langle Q(j)| \langle A("j")|, \quad (121)$$

thereby implying that the actual ontic state for the composite system $Q + A$ is represented by a state vector *exponentially close* to a basis vector of the form $|Q(j)', A("i")'\rangle$ for $j = i$. Hence, given that $Q + A$ is in one of these actual ontic states, the kinematical parent-subsystem manifestation (86) of our quantum conditional probabilities smooths out the discrepancies and ensures with near-certainty that the final actual ontic states of the subsystems Q and A have correctly correlated values of the label i . We can therefore conclude that if the final actual ontic state of the apparatus is $A("i")' \approx A("i")$, then the final actual ontic state of the subject system must with near-certainty be $Q(i)' \approx Q(i)$ for the *same* value of i , and vice versa.

4. Corrections to the Born Rule

Looking back at our formula (43) relating the maximum entropy S of a system to its total number of degrees of freedom, we also see that the off-diagonal corrections (119) to the final density matrices resulting from the measurement process imply that the naïve Born rule (120) is itself only ever accurate up to corrections $\sim \exp(-S \times \Delta t/\tau)$, where $S = S_{A+E}$ is the combined entropy of the measurement apparatus A and the environment E . As a consequence, all statistical quantities derived from the naïve Born rule—including all expectation values, final-outcome state vectors, semiclassical observables, scattering cross sections, tunneling probabilities, and decay rates computed in quantum-mechanical theories and quantum field theories—likewise suffer from imprecisions $\sim \exp(-S \times \Delta t/\tau)$.

These results are nicely in keeping with the notion of our error-entropy bound (3), but are completely invisible in both the traditional Copenhagen interpretation of quantum theory as well as in the instrumentalist approach: In contrast to our own interpretation, both the Copenhagen interpretation and instrumentalism *derive* the partial-trace prescription (50) from the *a priori* assumption of the *exact* Born rule and feature a Von Neumann-Lüders projection postulate [227, 323] that *axiomatically* sets the off-diagonal entries of post-measurement density matrices precisely to zero and converts the corresponding improper mixtures into proper mixtures, as we explain in detail in Appendix 1.

On the other hand, for ordinary macroscopic measurement devices, our claimed corrections to the Born rule rapidly become far too small to notice, so our minimal modal interpretation of quantum theory is consistent with the predictions of the Copenhagen interpretation in most practical circumstances. Note that exponentially small corrections of this kind should, in principle, also generically arise in any other interpretations of quantum theory (such as the Everett-DeWitt many-worlds interpretation) that attempt to derive the Born rule from decoherence.⁴⁵

5. Repeatability and Robustness of Measurement Outcomes

Finally, it is easy to show within the scope of our analysis that if the dial on our apparatus A could display *sequences* of outcomes arising from *repeated* measurements conducted in rapid succession, then the final density matrix of the apparatus would have non-negligible probability eigenvalues *only* for possible ontic states of the form $A("1" \text{ and } "1" \text{ and } "1" \dots)$ and $A("2" \text{ and } "2" \text{ and } "2" \dots)$, but not for, say, $A("1" \text{ and } "2" \text{ and } "1" \dots)$. Thus, realistic measurements, at least of the simplified Von Neumann type, produce robust, persistent, repeatable outcomes, provided that we are working over sufficiently short time scales that that uncontrollable overall dynamics does not have time to alter the state of our subject system Q appreciably. This observation provides yet another reason why our interpretation of quantum theory, in contrast to the Copenhagen interpretation, does not need to assume wave-function collapse as a distinct axiomatic postulate.

C. Subjective Density Matrices and Proper Mixtures

Having derived the Born rule (120) at last, we can now safely introduce the conventional practice of employing density matrices to describe *subjective* epistemic states for quantum systems: Starting from the Born rule and given a

⁴⁵ These corrections to the Born rule have interesting consequences for gravitational physics: If every region of physical space has a bounded maximum entropy—given, say, in terms of the radius R and energy E of the region by the Bekenstein bound $S_{\max} = 2\pi k_B ER/\hbar c$ [42]—then observable notions of locality in such a region can never be more precise than $\sim \exp(-S \times \Delta t/\tau)$ for $S = S_{\max}$. See also footnote 11.

subjective probability distribution p_α over objective epistemic states described by a collection of density matrices $\hat{\rho}_\alpha$ for a given system, where we allow that p_α may be only a *formal* probability distribution in the spirit of Section II B 3 and thus we do not insist upon mutual exclusivity, standard textbook arguments show that we can compute empirical expectation values in terms of the subjective density matrix

$$\hat{\rho} = \sum_{\alpha} p_{\alpha} \hat{\rho}_{\alpha}. \quad (122)$$

Note, of course, that the mapping between formal subjective epistemic states and subjective density matrices is many-to-one [185] because the mapping does not invariantly encode the original choice of objective density matrices $\hat{\rho}_\alpha$: The same subjective density matrix $\hat{\rho}$ can be expressed in terms of infinitely many different choices of coefficients p_α and their associated objective density matrices $\hat{\rho}_\alpha$.

In the special case of a proper mixture, meaning an epistemic state that is *wholly* subjective in nature, each objective density matrix $\hat{\rho}_\alpha$ describes a single pure state $|\Psi_\alpha\rangle$ of the given system and thus the subjective density matrix (122) of the system takes the simpler form

$$\hat{\rho} = \sum_{\alpha} p_{\alpha} |\Psi_{\alpha}\rangle \langle \Psi_{\alpha}|. \quad (123)$$

Note, however, that the Von Neumann entropy formula $S = -\text{Tr}[\hat{\rho} \log \hat{\rho}]$ from (30) does not generically yield the Shannon formula $-\sum_{\alpha} p_{\alpha} \log p_{\alpha}$ for the proper mixture if p_{α} is merely a *formal* probability distribution over *non-exclusive* possible ontic states; the Von Neumann entropy formula gives $-\sum_{\alpha} p_{\alpha} \log p_{\alpha}$ only if p_{α} is a *logically rigorous* probability distribution involving *mutually exclusive* possibilities Ψ_{α} in the sense that the associated state vectors $|\Psi_{\alpha}\rangle$ are mutually orthogonal—see (26)—in which case the quantities p_{α} are the eigenvalues of $\hat{\rho}$.⁴⁶

It is important to note that a proper mixture (123) fundamentally describes classical uncertainty over the state vector $|\Psi_{\alpha}\rangle$ of a *single* physical system. For example, if a source generates a sequence of $N \gg 1$ physical copies of an elementary subsystem, with classical frequency ratios $p_{\alpha} \in [0, 1]$ for the emitted elementary subsystems to be described by definite corresponding state vectors $|\Psi_{\alpha}\rangle$ in the elementary-subsystem Hilbert space \mathcal{H} , then we can employ a subjective density matrix of the form (123) to provide an approximate description of the resulting proper mixture for a *single* randomly chosen copy of the elementary subsystem. However, in the idealized limit in which we can neglect any entanglement with the larger environment, the state of the parent system consisting of the *entire physical sequence* of copies is not described by a nontrivial density matrix at all, but instead by a tensor-product state vector of the form

$$|\Psi\rangle = |\Psi_{\alpha_1}\rangle \otimes |\Psi_{\alpha_2}\rangle \otimes \cdots \otimes |\Psi_{\alpha_N}\rangle \quad (124)$$

belonging to the tensor-product Hilbert space $\mathcal{H}_{\text{sequence}} = \mathcal{H} \otimes \cdots \otimes \mathcal{H}$ (N factors) of the parent system. In essence, formally replacing the state vector (124) of the full parent system with the subjective density matrix (123) of an arbitrarily chosen elementary subsystem of the parent system—as is often done merely for practical convenience in order to simplify calculations—implicitly assumes that we have erased all information about the original ordering of the physical sequence. By contrast, if we *do* wish to retain and make use of information about the original ordering of the sequence—such as if we wish to examine or select particular subsequences of interest—then we must be careful not to confuse the full sequence’s own state vector (124) with the subjective density matrix (123) of an arbitrarily chosen member of the sequence, lest we run into apparent contradictions of the sort uncovered in [237].⁴⁷

D. Paradoxes of Quantum Theory Revisited

1. Schrödinger’s Cat, Wigner’s Friend, and the System-Centric Nature of Ontology

In his famous 1935 thought experiment arguing against some of the interpretative approaches of his time [273],⁴⁸ Schrödinger considered the hypothetical case of an unfortunate cat trapped in a perfectly sealed box together with a

⁴⁶ As we noted immediately following our introduction (30) of the Von Neumann entropy, Jaynes proved in [193] that the *minimum* possible value of the Shannon entropy $-\sum_{\alpha} p_{\alpha} \log p_{\alpha}$ occurs precisely if the state vectors $|\Psi_{\alpha}\rangle$ are mutually exclusive, meaning that they constitute the diagonalizing orthonormal eigenbasis of $\hat{\rho}$ and thus the Shannon entropy $-\sum_{\alpha} p_{\alpha} \log p_{\alpha}$ exactly coincides with the Von Neumann entropy $S = -\text{Tr}[\hat{\rho} \log \hat{\rho}]$. A similar proof appears in [185].

⁴⁷ The authors of [237] elide these subtle distinctions in order to generate a paradox in which the proper mixture describing a carefully post-selected subsequence of non-entangled two-qubit systems can also seemingly be represented as an entangled pure state, with the ultimate goal of casting doubt on interpretations of quantum theory (such as our own) that attempt to ascribe an ontological meaning to state vectors.

⁴⁸ In describing his misgivings about taking macroscopic superpositions seriously as ontological facets of reality, Schrödinger wrote the following: “It is typical of these cases that an indeterminacy originally restricted to the atomic domain becomes transformed into macroscopic indeterminacy, which can then be resolved by direct observation. That prevents us from so naïvely accepting as valid a ‘blurred model’ for representing reality. In itself it would not embody anything unclear or contradictory. There is a difference between a shaky or out-of-focus photograph and a snapshot of clouds and fog banks.” Schrödinger himself coined the term “entanglement” (*Verschränkung*) in this 1935 paper.

killing mechanism triggered by the decay of an unstable atom. If the experiment begins at the time $t = 0$, and if the atom has a 50-50 empirical outcome probability of quantum-mechanically decaying over a subsequent time interval Δt , then the atom’s state vector at exactly the time $t = \Delta t$ is the quantum superposition

$$|\text{atom}\rangle = \frac{1}{\sqrt{2}} (|\text{not decayed}\rangle + |\text{decayed}\rangle). \quad (125)$$

Immediately thereafter, the presumably linear (if not unitary) time evolution of the closed system inside the box would seem to imply that the final state vector of the parent system

$$(\text{parent system}) \equiv (\text{atom}) + (\text{killing mechanism}) + (\text{cat}) + (\text{air in box}) + (\text{box}) \quad (126)$$

then has the highly entangled, quantum-superposed form

$$\begin{aligned} |\text{parent system}\rangle = \frac{1}{\sqrt{2}} & (|\text{not decayed}\rangle |\text{not triggered}\rangle |\text{alive}\rangle |\text{air in box}\rangle |\text{box}\rangle \\ & + |\text{decayed}\rangle |\text{triggered}\rangle |\text{dead}\rangle |(\text{air in box})'\rangle |(\text{box})'\rangle). \end{aligned} \quad (127)$$

Does this expression imply that the cat’s state of existence is smeared out into a quantum superposition of alive and dead, and that the cat remains in that superposed condition until a human experimenter opens the box to perform a measurement on the cat’s final state and thereby collapses the cat’s wave function?

The Everett-DeWitt many-worlds interpretation addresses these questions by dropping the assumption that the final measurement carried out by the human experimenter has a definite outcome in any familiar sense. Instead, according to the many-worlds interpretation, the human experimenter merely “joins” the superposition upon opening the box to look inside, thereby splitting into two clones that each observe a different outcome. As always, beyond whatever metaphysical discomfort accompanies the notion that the world “unzips” into multiple copies with each passing moment in time (and, in fairness, such discomfort may be nothing more than human philosophical prejudice), there remains the tricky issue of making sense of probability itself when all possible outcomes of an experiment simultaneously occur, as well as the trouble of understanding why the numerical coefficients $1/\sqrt{2}$ appearing in the superposed state vector (127) have the physical meaning of probabilities in the first place.⁴⁹

By contrast, our minimal modal interpretation of quantum theory resolves the problem much less extravagantly: We begin by noting that the moment the cat interacts with the killing mechanism (and inevitably also with the air in the box and with the walls of the box itself), the cat’s own density matrix rapidly decoheres to the approximate form

$$\hat{\rho}_{\text{cat}} = \frac{1}{2} |\text{alive}\rangle \langle \text{alive}| + \frac{1}{2} |\text{dead}\rangle \langle \text{dead}|,$$

up to small but unavoidable degeneracy-breaking corrections similar to those appearing in (34).⁵⁰ According to our interpretation, the cat’s actual ontic state is therefore one in which the cat is definitely alive or definitely dead almost immediately after the killing mechanism acts—each possibility with corresponding epistemic probability $1/2$ —and well before the human experimenter opens the box.

Of course, the parent system as a whole (126) has a strangely quantum-superposed pure state (127) until the human experimenter opens the box, but this parent system is simply *not* the system that we identify as the cat alone. The relationship between a parent system and any one of its subsystems must ultimately be postulated as part of the interpretation of *any* physical theory, like quantum theory, that describes the world in terms of systems. In our own interpretation, this relationship—and thus the cat’s ontology—is determined *locally*, in the sense of being system-centric in the language of Section III C 1, through the cat’s *own* reduced density matrix.

After the human experimenter has opened the box, the original parent system (126) undergoes rapid decoherence and no longer remains in a quantum-superposed pure state. However, the *larger* parent system

$$\begin{aligned} (\text{larger parent system}) \equiv & (\text{atom}) + (\text{killing mechanism}) + (\text{cat}) + (\text{air in box}) + (\text{box}) \\ & + (\text{human experimenter}) + (\text{environment}) \end{aligned} \quad (128)$$

does remain in a quantum-superposed pure state, provided that we allow our definition of the environment to grow rapidly in physical size with the unavoidable and irreversible outward emission of thermal radiation at the speed

⁴⁹ We will discuss possible connections between the many-worlds interpretation and our own interpretation in greater depth in Section V E.

⁵⁰ Contrary to some accounts [116, 158], at no time during the experiment is the cat alone ever represented by a pure-state superposition of the form $(1/\sqrt{2})(|\text{alive}\rangle + |\text{dead}\rangle)$.

of light. If a second human experimenter—“Wigner”—now enters the story, then even before Wigner makes causal contact with this larger parent system (128) to find out the result of the measurement by the first experimenter—whom we’ll call “Wigner’s friend” [338]—decoherence ensures that Wigner could still say that the cat and Wigner’s friend (WF) each have classical-looking reduced density matrices given respectively by

$$\hat{\rho}_{\text{cat}} = \frac{1}{2} |\text{alive}\rangle \langle \text{alive}| + \frac{1}{2} |\text{dead}\rangle \langle \text{dead}|$$

and

$$\hat{\rho}_{\text{WF}} = \frac{1}{2} |\text{WF (“alive”)}\rangle \langle \text{WF (“alive”)}| + \frac{1}{2} |\text{WF (“dead”)}\rangle \langle \text{WF (“dead”)}|,$$

all with classical-looking possible ontic states alive, dead, WF (“alive”), and WF (“dead”), and again up to tiny degeneracy-breaking corrections. In fact, thanks once more to decoherence, Wigner can even say that the composite system consisting of just (cat) + (Wigner’s friend) has a classical-looking reduced density matrix

$$\hat{\rho}_{\text{cat+WF}} = \frac{1}{2} |\text{alive}\rangle |\text{WF (“alive”)}\rangle \langle \text{alive}| \langle \text{WF (“alive”)}| + \frac{1}{2} |\text{dead}\rangle |\text{WF (“dead”)}\rangle \langle \text{dead}| \langle \text{WF (“dead”)}|.$$

Hence, our instantaneous kinematical parent-subsystem quantum conditional probabilities (86) ensure that if the composite system (cat) + (Wigner’s friend) has, say, the ontic state (alive, WF (“alive”)), then, with essentially unit conditional probability, the ontic state of the cat is alive and the ontic state of Wigner’s friend is WF (“alive”), as expected.

The preceding discussion makes clear that in our interpretation of quantum theory, ontology is a *local* phenomenon, where by “local” we again simply mean system-centric: In order to establish the ontology of a given system, one must examine that system’s *own* (reduced) density matrix and not the density matrices of any other systems—not even those of parent systems, except to the extent that a parent system’s density matrix determines the subsystem’s density matrix through the partial-trace operation (50).

This system-centric approach to ontology in our interpretation of quantum theory is motivated in part by the “Russian-doll problem” that we introduced in Section ID 2. As we explained in that section, there exists a very real possibility that Nature does not feature a well-defined maximal closed system in a pure state and enclosing all other systems. If indeed the “Russian doll” of open parent systems never terminates at a maximal closed system, then, for practical reasons, we eventually have to stop going up the unending hierarchy at some particular system’s reduced density matrix, and then we have no choice but to try to make sense of that reduced density matrix on its own terms. As a result, we are *forced* to approach notions of ontology in a system-centric way in any *realist* interpretation of quantum theory, including in the de Broglie-Bohm and Everett-DeWitt interpretations.

The way that our interpretation localizes notions of ontology is analogous to the way that general relativity localizes the notions of inertial reference frames and observables like relative velocities that are familiar from special relativity: Just as there’s no well-defined sense in which widely separated galaxies in an expanding universe have a sensible local rest frame or relative velocity (contrary to colloquial assertions that galaxies sufficiently far from our own are “receding from our own galaxy faster than the speed of light”), there’s no well-defined sense in which we can generically read off the ontologies of highly entangled subsystems directly from the density matrix of a larger parent system.

In particular, given a system whose ontic state involves a highly quantum superposition, naïvely assuming (as one does in the Everett-DeWitt many-worlds interpretation) that each of the system’s subsystems likewise has a highly quantum-superposed existence would mean committing the same kind of fallacy of division that we first described in the context of characterizing the precise kinematical relationship (86) between the ontic states of parent systems and the ontic states of their subsystems in Section IID 6.

Schrödinger intended his thought experiment to be a proof of principle against taking the notion of quantum superpositions too literally, but it is worth asking whether his experiment could ever practically be realized (barring obvious ethical questions). Microscopic versions of his experiment involving small assemblages of particles, known as “Schrödinger kittens,” are performed all the time [22, 23, 143, 247, 286], and these experiments have been scaled up beyond microscopic size in recent years [122, 246, 312, 339].

However, the systems involved in these experiments are far from what one would consider complex living organisms. The trouble is that maintaining a large system in a non-negligible quantum superposition of macroscopically distinct states requires completely decoupling the system of interest from the sort of warm, hospitable environment that life as we know it seems to require. When one recognizes moreover that all large living creatures constantly give off thermal radiation themselves, it becomes untenable to imagine that partial tracing down to just the degrees of freedom that we associate with a living cat could ever yield an actual ontic state for the cat that remains in such a distinctly non-classical superposition for more than an infinitesimal instant in time. One can certainly imagine

scientists eventually having the technology to construct a quantum superposition involving a frozen-solid cat at a temperature near absolute zero in an evacuated chamber, but such an experiment would hardly be in keeping with the original spirit of Schrödinger’s “paradox.”

2. The Quantum Zeno Paradox

The quantum Zeno paradox [35, 80, 92, 158, 239] demonstrates yet another limitation of the textbook Copenhagen interpretation, and so it is important to explain why the problem is avoided in decoherence-based interpretations of quantum theory such as the one that we present in this paper.

The paradox concerns the behavior of systems that experience rapid sequences of measurements but that otherwise undergo unitary time evolution according to some Hamiltonian \hat{H} . For time intervals T that are sufficiently large compared with the inverse energy differences (multiplied by \hbar) between the system’s energy eigenstates, but small enough that time-dependent perturbation theory is valid, the probability of observing a transition away from an unstable state Ψ_{initial} is approximately *linear* in T ,

$$p_{\text{initial} \rightarrow \text{anything} \neq \text{initial}}(T) \approx \alpha T, \quad (129)$$

where $\alpha \approx \text{const}$ is just the total transition rate away from Ψ_{initial} and where we assume for consistency with perturbation theory that α is much smaller than $1/T$. Hence, the probability that a measurement after a time T will find the system *still in its original* state Ψ_{initial} is

$$p_{\text{initial}}(T) \approx 1 - \alpha T. \quad (130)$$

It follows that if we perform $N \gg 1$ sequential measurements separated by time intervals $T = t/N$ over which (130) holds, then the probability of finding the system *still in the initial state* Ψ_{initial} at the final time $t = N \times (t/N)$ is given by the familiar exponential decay law

$$p_{\text{initial}}(t) \approx e^{-\alpha t}. \quad (131)$$

This argument helps explain why exponential laws are so ubiquitous among systems exhibiting quantum decay.

By contrast, for *extremely short* time intervals $\Delta t \rightarrow 0$, the Born rule (120) implies that the probability that a measurement after a time Δt will find the system still in its original state Ψ_{initial} depends *quadratically* on Δt ,

$$p_{\text{initial}}(\Delta t) = \left| \langle \Psi_{\text{initial}} | e^{-i\hat{H}\Delta t/\hbar} | \Psi_{\text{initial}} \rangle \right|^2 = 1 - \beta^2 \Delta t^2, \quad (132)$$

where

$$\beta = \frac{\sqrt{\langle (H - \langle H \rangle)^2 \rangle}}{\hbar} = \frac{\Delta E}{\hbar} \sim \frac{1}{\tau} \quad (133)$$

is just the variance of the system’s energy in the state Ψ_{initial} up to a factor of $1/\hbar$, and is related to the system’s characteristic time scale τ through the energy-time uncertainty principle $\Delta E \times \tau \sim \hbar$. (Brackets $\langle \dots \rangle$ denote expectation values in the state Ψ_{initial} .) If we carry out $N \gg 1$ sequential measurements separated by time intervals $\Delta t = t/N \ll \tau \sim 1/\beta$, then the probability that we will see the system still in its initial state Ψ_{initial} at the final time $t = N \times (t/N)$ seemingly asymptotes to unity,

$$p_{\text{initial}}(t) \rightarrow 1, \quad (134)$$

apparently implying that the quantum state never transitions. This result, known as the quantum Zeno paradox (or the watched-pot paradox), would appear to suggest that a *continuously observed* unstable system can never decay.

As with many of the supposed “paradoxes” of quantum physics, the quantum Zeno paradox can be resolved by examining the measurement process more carefully. To begin, any interpretation of quantum theory (unlike the Copenhagen interpretation) that regards measurements as being physical processes implies that they occur over a *nonzero* time interval, and therefore implies that Δt cannot, in fact, be taken to be arbitrarily small. The notion of a *continuously observed* system is therefore a fiction, and so we can safely rule out the quantum Zeno paradox in its strongest sense, namely, in its suggestion that one could prevent a system from *ever* decaying by observing it frequently enough.

But we can go farther and explain why even a weaker version of the quantum Zeno effect—that is, a noticeable delay in the system’s decay rate or a deviation in the exponential decay law (131)—is extremely difficult to produce and observe. By momentarily connecting our unstable system to a measurement apparatus A with its own temporal resolution scale $\Delta t_A \sim \hbar/\Delta E_A$ —meaning that we replace our initial state vector $|\Psi_{\text{initial}}\rangle$ with $|\Psi_{\text{initial}}\rangle|A\rangle$ —the resulting interactions risk enlarging the system’s total energy variance, with the effect of increasing the coefficient β appearing in the decay probability (132) and consequently causing a breakdown in the relevant quadratic approximation [145]. The quantum Zeno effect is therefore far from a generic process for systems exposed to environmental interactions, and successfully producing the effect requires very careful experiments [187, 345].

3. Particle-Field Duality

Disputes crop up occasionally over whether particles or fields represent the more fundamental feature of reality.⁵¹ Such questions are perfectly valid to ask in the context of certain interpretations of quantum theory, including the so-called “fixed” modal interpretations—the de Broglie-Bohm pilot-wave interpretation being an example—in which a particular Hilbert-space basis is singled out permanently as defining the system’s ontology.⁵²

From the standpoint of our own interpretation, however, these questions are much less meaningful. A quantum field is a quantum system having some specific dynamics and a Hilbert space that is usefully described at low energies and weak coupling as a Fock space. (Note that we use the term “quantum field” here to refer to the physical system itself and not to mathematical field operators.) Some of the state vectors in this Hilbert space look particle-like, in the sense of being eigenstates of a particle-number operator, whereas—at least in the case of a bosonic system—other state vectors in the Hilbert space look more like classical fields, in the sense of being coherent states [152, 272] with a well-defined classical field amplitude.⁵³ Most of the state vectors in the Hilbert space are highly quantum and don’t resemble either particle-like states or classical-field-like states.

According to our interpretation of quantum theory, the ontic basis of the system is contextual: Experiments that count quanta decohere the system to an ontic basis of particle-like states, whereas experiments that measure forces on test bodies decohere the system to an ontic basis of classical-field-like states. Hence, neither of these two possible classes of states represents the “fundamental” basis for the system in any permanent or universal sense. Indeed, to say one of these two classes of states is more fundamental than the other would be precisely like saying that the energy eigenstates of a harmonic oscillator are more fundamental than its coherent states or vice versa, or like saying that a harmonic oscillator “is really” made up of energy eigenstates or that it “is really” made up of coherent states.

V. LORENTZ INVARIANCE AND LOCALITY

A. Special Relativity and the No-Communication Theorem

Special relativity links locality with causality, but only to the extent of forbidding *observable signals* from propagating superluminally. Hidden variables that exhibit nonlocal dynamics are therefore not necessarily a problem: The no-communication theorem [166, 255] ensures that quantum systems with local density-matrix dynamics do not transmit superluminal observable signals, which could otherwise spell trouble for causality.

Indeed, consider the simplified example of a dynamically closed parent system $A + B$ with a Hamiltonian \hat{H} , in which case the parent system’s density matrix evolves according to the unitary Liouville-Von Neumann equation (24):

$$\frac{\partial \hat{\rho}(t)}{\partial t} = -\frac{i}{\hbar}[\hat{H}, \hat{\rho}(t)].$$

If the parent system’s Hamiltonian has a separable decomposition of the form $\hat{H} = \hat{H}_A \otimes \hat{1} + \hat{1} \otimes \hat{H}_B$ in terms of independent, commuting Hamiltonians \hat{H}_A and \hat{H}_B for the subsystems A and B , respectively—as is necessarily the case in practice if the subsystems A and B are well-separated from each other in physical space—then it is straightforward to show [38, 323] directly by taking appropriate partial traces (50) of the Liouville-Von Neumann

⁵¹ For a recent instance, see [180].

⁵² As we explained in Section III A 6, the reliance of the de Broglie-Bohm pilot-wave interpretation on a universal, permanent ontic basis leads to serious trouble in dealing with relativistic systems: States of particles with definite positions are unavailable in general, because relativistic systems do not admit a basis of sharp orthonormal position eigenstates, and neither coherent states nor field eigenstates exist for fermionic systems [293].

⁵³ Polchinski makes the same point in [259]: “[G]iven a QFT, one can take two different classical limits depending on what one holds fixed. One limit gives classical fields, the other classical particles. It is fruitless to argue whether the fundamental entities are particles or fields. The fundamental description (at least to the extent that we now understand) is a QFT.”




Figure 7. The schematic set-up for the EPR-Bohm thought experiment, consisting of particles 1 and 2 together with spin detectors A and B .

equation that the dynamics of the reduced density matrix $\hat{\rho}_A(t)$ of the subsystem A has no influence whatsoever on the reduced density matrix $\hat{\rho}_B(t)$ of the subsystem B , and vice versa.

However, a more subtle issue arises when we consider whether an interpretation’s hidden variables are well-defined under arbitrary choices of inertial reference frame, or, equivalently, under arbitrary choices of foliation of four-dimensional spacetime into the three-dimensional spacelike hyperplanes that define slices of constant time. We explore this and related questions in subsequent sections, including the well-known EPR-Bohm and GHZ-Mermin thought experiments, and ultimately conclude that our minimal modal interpretation of quantum theory is indeed consistent with Lorentz invariance and that its nonlocality is no more severe than is the case for classical gauge theories.

B. The EPR-Bohm Thought Experiment and Bell’s Theorem

In 1964, Bell [43] employed a thought experiment due to Bohm [58, 61] and based originally on a 1935 paper by Einstein, Podolsky, and Rosen [107] to prove that no realist interpretation of quantum theory in which experiments have definite outcomes could be based on *dynamically local* hidden variables.⁵⁴

Bell’s theorem has exerted a powerful influence on the foundations of quantum theory in all the years since, to the extent that all interpretations today drop at least one of Bell’s stated assumptions—either the existence of hidden variables, realism, dynamical locality, or the assertion that experiments have definite outcomes. For example, the de Broglie-Bohm pilot-wave interpretation [59, 60, 63, 91] involves a hidden level of *manifestly nonlocal* dynamics, as do all the modal interpretations [28, 29, 70, 75, 205–207, 313, 314, 316, 317], including our own. The Everett-DeWitt many-worlds interpretation [68, 93, 94, 96, 111–113, 325, 326, 337] does away with the assumption that experiments have definite outcomes, and instead asserts that all outcomes occur simultaneously and are thus equally real. (We argue in Section VIF 4 that the many-worlds interpretation, as traditionally formulated, still suffers from some residual dynamical nonlocality.)

It is interesting to recapitulate Bell’s arguments carefully in order to understand precisely how they interface with our interpretation of quantum theory. Along the way, we will spot a subtle *additional* loophole in his assumptions, although our interpretation does not ultimately exploit this loophole and thus does not void Bell’s conclusions about nonlocality.⁵⁵

1. The Basic Set-Up of the EPR-Bohm Thought Experiment

The EPR-Bohm set-up begins with a pair of distinguishable spin-1/2 particles—denoted particle 1 and particle 2—that are initially prepared in an entangled pure state of total spin zero, so that the ontic state Ψ of the composite two-particle system 1 + 2 is described by the rotation-invariant state vector (33):

$$|\Psi\rangle = \frac{1}{\sqrt{2}}(|\uparrow\downarrow\rangle - |\downarrow\uparrow\rangle), \quad \hat{S}_{1+2,z}|\Psi\rangle = 0. \quad (135)$$

The particles are then separated by a large distance and are subsequently measured by two spatially separated spin detectors A and B . Specifically, we suppose that A locally measures the component $S_{1,\vec{a}} = \vec{a} \cdot \vec{S}_1$ of the spin \vec{S}_1 of particle 1 along the direction of a three-dimensional spatial unit vector \vec{a} and that B locally measures the component $S_{2,\vec{b}} = \vec{b} \cdot \vec{S}_2$ of the spin \vec{S}_2 of particle 2 along the direction of a three-dimensional spatial unit vector \vec{b} . (See Figure 7.)

⁵⁴ It is important to note again that, despite Bell’s theorem, any interpretation of quantum theory consistent with the theory’s standard observable rules does not permit superluminal *observable signals*, as ensured by the no-communication theorem [166, 255].

⁵⁵ As we will discuss shortly, the same loophole also exists in the assumptions of similar no-go theorems, such as the CHSH theorem of [44, 83, 84, 213], although not in the GHZ-Mermin arguments to be described in Section VC.

2. Spin Correlation Predicted by Quantum Theory

After many repetitions of the entire experiment—each time preparing a composite two-particle system $1 + 2$ in the ontic state (135), separating the particles by a large distance, and then allowing the spin detectors A and B to perform local measurements on the spins of their respective individual particles—we can ask for the average value of the product $S_{1,\vec{a}}S_{2,\vec{b}}$, which is just the correlation function for the pair of measured spins. Working in units of $\hbar/2$ for clarity (that is, setting $\hbar/2 \equiv 1$), quantum theory predicts that this expectation value should be given by

$$\langle S_{1,\vec{a}}S_{2,\vec{b}} \rangle_{\text{QM}} = -\vec{a} \cdot \vec{b}, \quad (136)$$

a result that follows directly from linearity and rotation invariance.⁵⁶

3. The Bell Inequality

Bell argued that no *dynamically local* hidden variables could account for the result (136). To prove this claim, he derived a “triangle inequality”

$$\left| \langle S_{1,\vec{a}}S_{2,\vec{b}} \rangle_{\text{LHV}} - \langle S_{1,\vec{a}}S_{2,\vec{c}} \rangle_{\text{LHV}} \right| \leq 1 + \langle S_{1,\vec{b}}S_{2,\vec{c}} \rangle_{\text{LHV}} \quad (137)$$

that all interpretations based on local hidden variables (LHV) would necessarily have to satisfy, where \vec{c} is an alternative choice of unit vector for the alignment of the detector B . The correct quantum result (136) can easily violate this inequality as well as various generalizations, as has been confirmed repeatedly in experiments [18, 24, 25, 120, 150, 159, 267, 308, 309, 329].⁵⁷

To derive the inequality (137), Bell assumed that the hidden variables λ obey a probability distribution $p(\lambda)$ satisfying the standard Kolmogorov conditions [203]:

$$0 \leq p(\lambda) \leq 1, \quad \int d\lambda p(\lambda) = 1. \quad (138)$$

The proof of the Bell inequality (137), which we present in Appendix 2a, then consists of manipulations of averages weighted by $p(\lambda)$.

4. The Additional Loophole

Bell’s theorem and all the various generalizations that have arisen in subsequent papers by Bell and others, including the CHSH inequalities [44, 83, 84] and Legget’s inequality for “crypto-nonlocal” interpretations [213], implicitly assume that the hidden variables λ admit a sensible *joint* epistemic probability distribution $p(\lambda)$ as described in (138). As we made clear in Section IIID 4, this assertion is nontrivial, but none of these papers attempt to justify its validity. However, this assumption represents a significant loophole, one that appears to have been noticed first by Bene [49–54] in the context of his “perspectival” interpretation of quantum theory, but then also discovered independently and examined in a small number of additional papers [179, 311, 347].⁵⁸ Although this loophole is interesting to consider, we do not exploit it in our own interpretation.

⁵⁶ Proof: First, observe that $\langle S_{1,\vec{a}}S_{2,\vec{b}} \rangle = \sum_i \sum_j a_i b_j F_{ij}$, where $F_{ij} \equiv \langle S_{1,i}S_{2,j} \rangle$ is a real-valued 3×3 matrix that is manifestly independent of \vec{a} and \vec{b} . Due to the rotation invariance of the entangled state vector (135), our final result for $\langle S_{1,\vec{a}}S_{2,\vec{b}} \rangle$ must likewise be invariant under global rotations, and because it should obviously have the specific value -1 if $\vec{a} = \vec{b}$, we must have $F_{ij} = -\delta_{ij}$. Hence, $\langle S_{1,\vec{a}}S_{2,\vec{b}} \rangle = \sum_i \sum_j a_i b_j (-\delta_{ij}) = -\vec{a} \cdot \vec{b}$. QED

⁵⁷ For example, if we take \vec{a} and \vec{b} to be orthogonal and \vec{c} to be 45 degrees between both of them, then the quantum formula (136) for the spin correlation function yields $\langle S_{1,\vec{a}}S_{2,\vec{b}} \rangle = 0$, $\langle S_{1,\vec{a}}S_{2,\vec{c}} \rangle \simeq -0.707$, and $\langle S_{1,\vec{b}}S_{2,\vec{c}} \rangle \simeq -0.707$, in which case the Bell inequality (137) is expressly violated: $\left| \langle S_{1,\vec{a}}S_{2,\vec{b}} \rangle - \langle S_{1,\vec{a}}S_{2,\vec{c}} \rangle \right| \simeq 0.707 \not\leq 0.293 \simeq 1 + \langle S_{1,\vec{b}}S_{2,\vec{c}} \rangle$. For a discussion of potential shortcomings and loopholes in these Bell-test experiments, as well as a description of experimental work aimed at closing some of these loopholes, see, for example, [132, 134, 142, 178, 212].

⁵⁸ Fine [118, 119] appears to have worked on similar ideas in 1982. He studied whether one could assume the existence of the particular joint probability $p(A, B, 1, 2)$, but he did not regard the states of A , B , 1, and 2 as hidden variables themselves. Indeed, he assumed a deeper level of hidden variables λ with their own probability distribution $p(\lambda)$, and he never considered that $p(\lambda)$ might not exist. Some interpretations of quantum theory, such as Bene’s, essentially identify the states of A , B , 1, 2 as hidden variables and question the existence of $p(\lambda)$ itself.

5. Analysis of the EPR-Bohm Thought Experiment

To the extent that our interpretation of quantum theory involves hidden variables (see Section III A 4), they are just the actual ontic states $\lambda = \{\Psi, \Psi_1, \Psi_2\}$ of the two-particle parent system and of the individual one-particle subsystems, where Ψ is the entangled state defined in (135) and where $\Psi_i = \uparrow$ or \downarrow . Before any measurements take place, the epistemic probabilities for the composite two-particle parent system 1 + 2 are

$$\text{Before: } p_{1+2}(\Psi) = 1, \quad p_{1+2}(\text{states } \perp \Psi) = 0, \quad (139)$$

the epistemic probabilities for the individual one-particle subsystems 1 and 2 are⁵⁹

$$\text{Before: } p_1(\uparrow) = p_1(\downarrow) = \frac{1}{2}, \quad p_2(\uparrow) = p_2(\downarrow) = \frac{1}{2}, \quad (140)$$

and the epistemic probabilities for the individual spin detectors A and B are

$$\text{Before: } p_A(\text{"}\emptyset\text{"}) = 1, \quad p_B(\text{"}\emptyset\text{"}) = 1 \quad (\text{"}\emptyset\text{"} \equiv \text{"blank"}), \quad (141)$$

with all epistemic probabilities for orthogonal possible ontic states being zero. If we introduce a macroscopic measurement apparatus M (unavoidably coupled to an even larger environment E , at the very least through irreversible thermal radiation [165]) that will locally record and compare the two readings on the spin detectors after their pair of measurements, then before any measurements have taken place, the only nonzero epistemic probability for M is

$$\text{Before: } p_M(\text{"}\emptyset\text{"}) = 1. \quad (142)$$

Importantly, in accordance with our instantaneous kinematical conditional probabilities (86) (but contrary to the hypothetical loophole that we described in Section V B 4), the two particles 1 and 2 also have the initial *joint* epistemic probabilities

$$\text{Before: } \left. \begin{aligned} p_{1,2|1+2}(\uparrow, \downarrow | \Psi) &= \frac{1}{2}, & p_{1,2|1+2}(\downarrow, \uparrow | \Psi) &= \frac{1}{2}, \\ p_{1,2|1+2}(\uparrow, \uparrow | \Psi) &= 0, & p_{1,2|1+2}(\downarrow, \downarrow | \Psi) &= 0, \end{aligned} \right\} \quad (143)$$

which ensure that the individual ontic states of the two particles are initially anti-correlated.

It is easy to check that after a local spin measurement by either one of the spin detectors A or B aligned along the z direction, and the consequent rapid decoherence, the objective epistemic state of the composite two-particle system 1 + 2 becomes

$$\text{After } A \text{ or } B: \quad p_{1+2}((\uparrow\downarrow)') = p_{1+2}((\downarrow\uparrow)') = \frac{1}{2}, \quad p_{1+2}((\uparrow\uparrow)') = p_{1+2}((\downarrow\downarrow)') = 0, \quad (144)$$

where the primes denote corrections of order (119)—exponentially small in the number of (spin-detector) + (environment) degrees of freedom—to the two-particle basis states that are necessary to diagonalize the final density matrix $\hat{\rho}_{1+2}$ exactly. Given the assumption that the ontic state of the composite system 1 + 2 after the measurement is $(\uparrow\downarrow)'$, a simple application of the kinematical parent-subsystem formula (86) shows that there is nearly unit probability that particle 1 is in the ontic state \uparrow and particle 2 is in the ontic state \downarrow , whereas given the alternative assumption that the ontic state of 1 + 2 is $(\downarrow\uparrow)'$ after the measurement, there is nearly unit probability that particle 1 is in the ontic state \downarrow and particle 2 is in the ontic state \uparrow . These results are nicely consistent with our initial joint epistemic probabilities (143) for the two particles.

Taking appropriate partial traces (50) of the density matrix of the full system after both spin detectors A and B have performed their measurements, we find that the final epistemic probabilities for the individual one-particle subsystems 1 and 2 are given to very high numerical accuracy by

$$\text{After: } p_1(\uparrow') = p_1(\downarrow') = \frac{1}{2}, \quad p_2(\uparrow') = p_2(\downarrow') = \frac{1}{2}, \quad (145)$$

⁵⁹ As we explained in our discussion surrounding (34), there will unavoidably exist tiny degeneracy-breaking parameters in the entangled two-particle parent system's ontic state Ψ that—by an appropriate choice of our Cartesian coordinate system for three-dimensional physical space—pick out the spin- z basis for each one-particle subsystem.

in approximate agreement with (140) before the measurements took place, where, again, the primes denote exponentially tiny corrections (119) to the definitions of the ontic basis states. Likewise, the final *joint* epistemic probabilities for the two particles are given by

$$\text{After: } \left. \begin{aligned} p_{1,2|1+2}(\uparrow', \downarrow' | (\uparrow\downarrow)') &= 1, & p_{1,2|1+2}(\downarrow', \uparrow' | (\uparrow\downarrow)') &= 0, \\ p_{1,2|1+2}(\uparrow', \uparrow' | (\uparrow\downarrow)') &= 0, & p_{1,2|1+2}(\downarrow', \downarrow' | (\uparrow\downarrow)') &= 0, \\ p_{1,2|1+2}(\uparrow', \downarrow' | (\downarrow\uparrow)') &= 0, & p_{1,2|1+2}(\downarrow', \uparrow' | (\downarrow\uparrow)') &= 1, \\ p_{1,2|1+2}(\uparrow', \uparrow' | (\downarrow\uparrow)') &= 0, & p_{1,2|1+2}(\downarrow', \downarrow' | (\downarrow\uparrow)') &= 0. \end{aligned} \right\} \quad (146)$$

As for the individual spin detectors A and B , we find the final epistemic probabilities

$$\text{After: } \left. \begin{aligned} p_A(\text{“}\uparrow\text{”})' &= \frac{1}{2}, & p_A(\text{“}\downarrow\text{”})' &= \frac{1}{2}, \\ p_B(\text{“}\downarrow\text{”})' &= \frac{1}{2}, & p_B(\text{“}\uparrow\text{”})' &= \frac{1}{2}, \end{aligned} \right\} \quad (147)$$

and for the apparatus M that locally reads off the pair of results,

$$\text{After: } \left. \begin{aligned} p_M(\text{“found } A(\text{“}\uparrow\text{”}) \text{ and } B(\text{“}\downarrow\text{”})\text{”})' &= \frac{1}{2}, \\ p_M(\text{“found } A(\text{“}\downarrow\text{”}) \text{ and } B(\text{“}\uparrow\text{”})\text{”})' &= \frac{1}{2}. \end{aligned} \right\} \quad (148)$$

Notice that if our apparatus M could immediately check the spin detectors again and record their dials a second time, then its final epistemic probabilities would be

$$\text{After: } \left. \begin{aligned} p_M(\text{“found } A(\text{“}\uparrow\text{”}) \text{ and } B(\text{“}\downarrow\text{”}) \text{ twice”})' &= \frac{1}{2}, \\ p_M(\text{“found } A(\text{“}\downarrow\text{”}) \text{ and } B(\text{“}\uparrow\text{”}) \text{ twice”})' &= \frac{1}{2}, \end{aligned} \right\} \quad (149)$$

thereby implying that the measurement results are robust and persistent, as expected.

The spin detectors A and B play the role of macroscopic, highly classical intermediaries between the apparatus M and the individual particles 1 and 2. As a consequence, we would face troubling metaphysical issues if we were unable to assert a strong connection between the final actual ontic states of the spin detectors A and B and the final actual ontic state of the apparatus M that looks at them. Fortunately, after environmental decoherence, the composite system $W = 1 + 2 + A + B + M$ has the final epistemic state

$$\text{After: } \left. \begin{aligned} p_W(\text{“found } A(\text{“}\uparrow\text{”}) \text{ and } B(\text{“}\downarrow\text{”}) \text{ twice”, } A(\text{“}\uparrow\text{”}), B(\text{“}\downarrow\text{”}), \uparrow, \downarrow) &= \frac{1}{2}, \\ p_W(\text{“found } A(\text{“}\downarrow\text{”}) \text{ and } B(\text{“}\uparrow\text{”}) \text{ twice”, } A(\text{“}\downarrow\text{”}), B(\text{“}\uparrow\text{”}), \downarrow, \uparrow) &= \frac{1}{2}, \end{aligned} \right\} \quad (150)$$

where, as usual, primes denote exponentially suppressed corrections (119) in the definitions of the ontic states that are necessary to diagonalize the density matrix of W exactly. Hence, invoking again our kinematical parent-subsystem probabilistic smoothness condition (86), we can say with near-certainty at the level of *conditional* probabilities that if, say, the final actual ontic state of the apparatus M happens to be $(\text{“found } A(\text{“}\uparrow\text{”}) \text{ and } B(\text{“}\downarrow\text{”}) \text{ twice”})'$, then the final actual ontic states of the two spin detectors are really respectively $A(\text{“}\uparrow\text{”})$ and $B(\text{“}\downarrow\text{”})$ and, furthermore, the final actual ontic states of the individual particles are really \uparrow for 1 and \downarrow for 2. These results are obviously all consistent with our expectations for the post-measurement state of affairs.

6. More General Alignments and Nonlocality

Generating a violation of the Bell inequality (137) necessitates choosing *nonparallel* detector alignments, so it is no surprise that our investigation so far has not turned up any clear evidence of nonlocality. Hence, let us now suppose that the alignment \vec{a} of the spin detector A is along the x direction. For clarity, we will suppress the primes on ontic states, remembering always that decoherence is an imperfect process and that we should expect exponentially small corrections to all our results.

If the spin detector A performs its measurement on particle 1 first, then after the resulting measurement-induced decoherence, a simple calculation shows that particle 1 now has possible ontic spin states \leftarrow (spin- x left) and \rightarrow (spin- x right), each with epistemic probability $1/2$. But the local unitary dynamics of particle 2 implies that its own possible ontic states remain \uparrow and \downarrow , each with epistemic probability $1/2$. As expected, a subsequent local spin measurement by the spin detector B aligned along the z direction would yield the outcome “ \uparrow ” or “ \downarrow ,” each with corresponding empirical outcome probability $1/2$.

On the other hand, suppose that before performing its own local spin measurement, the spin detector B changes its alignment \vec{b} to point along the x direction just like the alignment \vec{a} of the spin detector A . In that case, if we use the local dynamics of particle 2 and the spin detector B in our single-system dynamical conditional probabilities (98), then we find that regardless of whether the initial actual ontic state of particle 2 was \uparrow or \downarrow , the conditional probability for B to obtain either “ \leftarrow ” or “ \rightarrow ” for particle 2 is $1/2$, independent of the result “ \leftarrow ” or “ \rightarrow ” obtained by the spin detector A for particle 1; this result is unsurprising because by partial tracing out the spin detector A and particle 1, we have chosen to ignore crucial information about the necessary final anti-correlation between the spins of the two particles. But if we instead condition on either the initial or final actual ontic state of the composite parent two-particle system $1 + 2$, which has evolved nonlocally during the experiment, then a simple calculation yields with unit probability that the final measurement outcome for the spin of particle 2 by the spin detector B will indeed show the correct anti-correlation. This nonlocal influence on the final measurement result obtained by B is in line with our claim in Section III D 5 that our quantum conditional probabilities safely allow for a hidden level of nonlocality.

C. The GHZ-Mermin Thought Experiment

In a 1989 paper [157], Greenberger, Horne, and Zeilinger (GHZ) showed that certain generalizations of the EPR-Bohm thought experiment involving three or more particles with spin could not be described in terms of dynamically local hidden variables. Unlike the arguments employed in Bell’s theorem, the GHZ result does not depend on the assumption of joint probability distributions for the hypothetical hidden variables, and can be captured by talking about just a single measurement outcome.⁶⁰

Following the simplest version of the GHZ argument, described in detail by Mermin in 1990 [235], we consider a system of three distinguishable spin- $1/2$ particles 1, 2, and 3 in an initial pure state represented by the entangled state vector

$$|\Psi_{\text{GHZ}}\rangle = \frac{1}{\sqrt{2}} (|\uparrow\uparrow\uparrow\rangle - |\downarrow\downarrow\downarrow\rangle), \quad (151)$$

where $|\uparrow\rangle$ and $|\downarrow\rangle$ are one-particle eigenstates of the spin- z operator \hat{S}_z and where again we work in units for which $\hbar/2 \equiv 1$:

$$\hat{S}_z |\uparrow\rangle = + |\uparrow\rangle, \quad \hat{S}_z |\downarrow\rangle = - |\downarrow\rangle. \quad (152)$$

It is then straightforward to show that the GHZ state vector described by (151) is a definite eigenstate of the three operators $\hat{S}_{1,x}\hat{S}_{2,y}\hat{S}_{3,y}$, $\hat{S}_{1,y}\hat{S}_{2,x}\hat{S}_{3,y}$, and $\hat{S}_{1,y}\hat{S}_{2,y}\hat{S}_{3,x}$ with eigenvalue $+1$,

$$\left. \begin{aligned} \hat{S}_{1,x}\hat{S}_{2,y}\hat{S}_{3,y} |\Psi_{\text{GHZ}}\rangle &= + |\Psi_{\text{GHZ}}\rangle, \\ \hat{S}_{1,y}\hat{S}_{2,x}\hat{S}_{3,y} |\Psi_{\text{GHZ}}\rangle &= + |\Psi_{\text{GHZ}}\rangle, \\ \hat{S}_{1,y}\hat{S}_{2,y}\hat{S}_{3,x} |\Psi_{\text{GHZ}}\rangle &= + |\Psi_{\text{GHZ}}\rangle, \end{aligned} \right\} \quad (153)$$

but is *also* a definite eigenstate of the operator

$$\hat{S}_{1,x}\hat{S}_{2,x}\hat{S}_{3,x} = - \left(\hat{S}_{1,x}\hat{S}_{2,y}\hat{S}_{3,y} \right) \left(\hat{S}_{1,y}\hat{S}_{2,x}\hat{S}_{3,y} \right) \left(\hat{S}_{1,y}\hat{S}_{2,y}\hat{S}_{3,x} \right)$$

with eigenvalue -1 ,

$$\hat{S}_{1,x}\hat{S}_{2,x}\hat{S}_{3,x} |\Psi_{\text{GHZ}}\rangle = - |\Psi_{\text{GHZ}}\rangle. \quad (154)$$

⁶⁰ Hardy [168, 169] has found certain two-particle variants of the GHZ construction that likewise do not depend on assumptions regarding joint probability distributions.




Figure 8. The schematic set-up for the GHZ-Mermin thought experiment, consisting of particles 1, 2, and 3 together with spin detectors A , B , and C .

Imagining that the particles are eventually separated and then measured by local spin detectors A , B , and C (see Figure 8), one can carefully list all the possible “initial local instructions” that we could have somehow packaged along with each of the individual particles in order to ensure agreement with the outcomes required by the first three eigenvalue equations (153), but not one of those sets of initial local instructions could then accommodate the fourth eigenvalue equation (154). The consequence is that any hidden variables accounting for all these possible measurement results must change nonlocally during the course of the experiment in order to ensure agreement with the necessary eigenvalue equations.

D. The Myrvold No-Go Theorem

Building on a paper by Dickson and Clifton [99] and employing arguments similar to Hardy [168], Myrvold [241] argued that modal interpretations are fundamentally inconsistent with Lorentz invariance at a deeper level than mere unobservable nonlocality, leading to much additional work [56, 106, 242] in subsequent years to determine the implications of his result. Specifically, Myrvold argued that assuming the existence of ontic states underlying density matrices could lead to paradoxes arising from Lorentz transformations. Seemingly the only way to avoid this conclusion would then be to break Lorentz symmetry in a fundamentally ontological way by asserting the existence of a universal “preferred” reference frame in which all ontic-state assignments must be made.

Myrvold’s claims, and those of Dickson and Clifton as well as Hardy, rest on assumptions that do not hold in our interpretation of quantum theory. In particular, Dickson and Clifton [99] assume that hidden ontic states admit certain joint epistemic probabilities that are conditioned on *multiple* disjoint systems at an initial time, and such probabilities are not a part of our interpretation, as we explained in Section III D 4.⁶¹ Similarly, all of Myrvold’s arguments hinge on the assumed existence of joint epistemic probabilities for two disjoint systems at *two or more separate times*—some of which are, moreover, *Lorentz-transformed* times—and again such probabilities are not present in our interpretation. As Dickson and Clifton point out in Appendix B of their paper, Hardy’s argument also depends on several inadmissible assumptions about ontic property assignments in modal interpretations.

E. Quantum Theory and Classical Gauge Theories

Suppose that we were to imagine reifying all the *possible* ontic states defined by every system’s density matrix as simultaneous *actual* ontic states in the sense of the many-worlds interpretation.⁶² Then because every density matrix *as a whole* evolves locally and there is no *single actual* ontic state that jumps between different *possible* ontic

⁶¹ The same implicit assumption occurs in Section 9.2 of [316].

⁶² Observe that the eigenbasis of each system’s density matrix therefore defines a preferred basis *for that system alone*—that is, in a system-centric manner. We *do not* assume the sort of universe-spanning preferred basis shared by all systems that is featured in the traditional many-worlds interpretation; such a universe-spanning preferred basis would lead to new forms of nonlocality, as we describe in Section VIF 4.

states, no nonlocal dynamics between the actual ontic states is necessary and our interpretation of quantum theory becomes manifestly dynamically local: For example, each spin detector in the EPR-Bohm or GHZ-Mermin thought experiments possesses all its possible results in actuality, and the larger measurement apparatus that later compares the results of the spin detectors locally “splits” into all the various possibilities when it visits each spin detector and looks at the detector’s final reading.

To help make sense of this step of adding unphysical actual ontic states into our interpretation of quantum theory, recall the story of classical gauge theories, and specifically the example of the Maxwell theory of electromagnetism. The physical states of the Maxwell theory involve only *two* possible field polarizations, corresponding fundamentally to the two physical spin states of the underlying species of massless spin-1 gauge boson that we call the photon.

In a unitary gauge, meaning a choice of gauge in which we describe the theory using just two polarization states for the gauge field A^μ , the theory appears to be nonlocal and not Lorentz covariant. For instance, in Weyl-Coulomb gauge, we impose the two manifestly non-covariant gauge conditions $A^0 = 0$ and $\vec{\nabla} \cdot \vec{A} = 0$, which together eliminate the unphysical timelike and longitudinal polarizations of the gauge field A^μ and thereby leave intact just the two physical transverse states orthogonal to the direction of wave propagation. In this “ontologically correct” choice of gauge, the vector potential \vec{A} becomes a nonlocal action-at-a-distance function of the distribution of electric source currents over all of three-dimensional space [188] and is no longer part of a true Lorentz four-vector because the gauge conditions are not preserved under arbitrary Lorentz transformations [332].

However, if we *formally* add an additional unphysical polarization state by replacing the two Weyl-Coulomb gauge conditions $A^0 = 0$ and $\vec{\nabla} \cdot \vec{A} = 0$ with the weaker but Lorentz-covariant single condition that defines Lorenz gauge, namely,

$$\partial_\mu A^\mu = \frac{1}{c} \frac{\partial A^0}{\partial t} + \vec{\nabla} \cdot \vec{A} = 0,$$

then we obtain a manifestly local, Lorentz-covariant description of the Maxwell theory in which the gauge potential becomes a true Lorentz four-vector given by *causal* integrals over appropriately *time-delayed* distributions of charges and currents. Because the Maxwell theory therefore has a mathematical description in which the superficial nonlocality of the unitary (that is, ontologically correct) Weyl-Coulomb gauge disappears, we see that the apparent nonlocality of the Maxwell theory is totally harmless.

Nonetheless, switching to Lorenz gauge does not imply that the extra unphysical polarization state that we have formally added to the Maxwell theory achieves a *true* ontological status, and we take precisely the same view toward the addition of extra unphysical “actual” ontic states that would provide a manifestly local mathematical description of our interpretation of quantum theory. Just as different choices of gauge for a given classical gauge theory make different calculations or properties of the theory more or less manifest—that is, each choice of gauge inevitably involves trade-offs—we see that switching from the “unitary gauge” corresponding to the original version of our interpretation of quantum theory to the “Lorenz gauge” in which it looks more like a density-matrix-centered version of the many-worlds interpretation makes the locality and Lorentz covariance of our interpretation more manifest at the cost of obscuring our interpretation’s underlying ontology and the meaning of probability.

In this analogy, gauge potentials A^μ correspond to ontic states Ψ_i , which can similarly undergo nonlocal changes. Gauge transformations $A^\mu \mapsto A^\mu + \partial^\mu \lambda$ describe an unobservable change in the gauge theory’s ontology, and are analogous in our interpretation of quantum theory to carrying out a simultaneous but fundamentally unobservable shift in the hidden actual ontic states (and thus also the memories) of all our systems, such as by reassigning each system’s hidden (“private”) actual ontic state to one of the other possible ontic states defined by the system’s density matrix. Just as all observable predictions of electromagnetism can be expressed in terms of gauge-invariant, dynamically local quantities like electric fields \vec{E} and magnetic fields \vec{B} , all outwardly observable statistical predictions of quantum theory ultimately derive from density matrices $\hat{\rho}$, which always evolve locally in accordance with the no-communication theorem [166, 255] and are insensitive to our identification of each system’s hidden actual ontic state from among the system’s possible ontic states.

Seen from this perspective, we can also better understand why it is so challenging [184] to make sense of a many-worlds-type interpretation as an ontologically and epistemologically reasonable interpretation of quantum theory: Attempting to do so leads to as much metaphysical difficulty as trying to make sense of the Lorenz gauge of Maxwell electromagnetism as an “ontologically correct interpretation” of the Maxwell theory.⁶³ Hence, taking a lesson from classical gauge theories, we propose that many-worlds-type interpretations should instead be regarded as being merely a convenient mathematical tool—a particular “gauge choice”—for establishing definitively that a given “unitary-gauge” interpretation of quantum theory like our own is ultimately consistent with locality and Lorentz invariance.

⁶³ Indeed, in large part for this reason, some textbooks [332] develop quantum electrodynamics fundamentally from the perspective of Weyl-Coulomb gauge.

VI. CONCLUSION

A. Summary

In this paper, we have introduced what we call the minimal modal interpretation of quantum theory. Our interpretation consists of several parsimonious ingredients:

1. **Ontic States:** In (25), we define quantum ontic states Ψ_i —meaning the states of a given system as it could actually exist in reality—in terms of arbitrary (unit-norm) state vectors $|\Psi_i\rangle$ in the system’s Hilbert space \mathcal{H} :

$$\Psi_i \leftrightarrow |\Psi_i\rangle \in \mathcal{H} \quad (\text{up to overall phase}).$$

This definition is the quantum counterpart to the classical notion of ontic states as elements in a configuration space.

2. **Epistemic States:** In Section III A 2, we define quantum epistemic states $\{(p_i, \Psi_i)\}_i$ as probability distributions over sets of possible quantum ontic states,

$$\{(p_i, \Psi_i)\}_i, \quad p_i \in [0, 1],$$

where, again, this definition parallels the corresponding notion from classical physics. We translate logical mutual exclusivity of ontic states Ψ_i as mutual orthogonality (26) of the associated state vectors $|\Psi_i\rangle$, and we make a distinction between subjective epistemic states (proper mixtures) and objective epistemic states (improper mixtures): The former arise from merely classical ignorance and are uncontroversial, whereas the latter arise from quantum entanglement with other systems and do not have a widely accepted *a priori* meaning outside of our interpretation of quantum theory. Indeed, the problem of interpreting objective epistemic states may well be unavoidable: All realistic systems are entangled with other systems to a nonzero degree and thus cannot be described exactly by pure states or by purely subjective epistemic states.⁶⁴ As part of our introduction of epistemic states into quantum theory, we posit a correspondence (29) between objective epistemic states $\{(p_i, \Psi_i)\}_i$ and density matrices $\hat{\rho}$:

$$\{(p_i, \Psi_i)\}_i \leftrightarrow \hat{\rho} = \sum_i p_i |\Psi_i\rangle \langle \Psi_i|.$$

The relationship between subjective epistemic states and density matrices is not as strict, as we explain in Section IV C.

3. **Partial Traces:** We invoke the partial-trace operation $\hat{\rho}_Q \equiv \text{Tr}_E [\hat{\rho}_{Q+E}]$, motivated and defined in (50) without appeals to the Born rule or Born-rule-based averages, to relate the density matrix (and thus the epistemic state) of any subsystem Q to that of any parent system $W = Q + E$.
4. **Quantum Conditional Probabilities:** We introduce a general class of quantum conditional probabilities (80),

$$\begin{aligned} p_{Q_1, \dots, Q_n | W}(i_1, \dots, i_n; t' | w; t) &\equiv \text{Tr}_W \left[\left(\hat{P}_{Q_1}(i_1; t') \otimes \dots \otimes \hat{P}_{Q_n}(i_n; t') \right) \mathcal{E}_W^{t' \leftarrow t} [\hat{P}_W(w; t)] \right] \\ &\sim \text{Tr} [\hat{P}_{i_1}(t') \dots \hat{P}_{i_n}(t') \mathcal{E}[\hat{P}_w(t)]], \end{aligned}$$

relating the possible ontic states of any partitioning collection of mutually disjoint subsystems Q_1, \dots, Q_n to the possible ontic states of a corresponding parent system $W = Q_1 + \dots + Q_n$ whose own dynamics is governed approximately by a linear completely-positive-trace-preserving (“CPTP”) dynamical mapping $\mathcal{E}_W^{t' \leftarrow t}$ over the given time interval $t' - t$. Here the operator $\hat{P}_W(w; t)$ denotes the eigenprojector onto the eigenstate $|\Psi_W(w; t)\rangle$ of the density matrix $\hat{\rho}_W(t)$ of the parent system W at the initial time t , and, similarly, for $\alpha = 1, \dots, n$, the operator $\hat{P}_{Q_\alpha}(i; t')$ denotes the eigenprojector onto the eigenstate $|\Psi_{Q_\alpha}(i; t')\rangle$ of the density matrix $\hat{\rho}_{Q_\alpha}(t')$ of the subsystem Q_α at the final time t' . In a rough sense, the dynamical mapping $\mathcal{E}_W^{t' \leftarrow t}$ acts as a parallel-transport superoperator that moves the parent-system eigenprojector $\hat{P}_W(w; t)$ from t to t' before we compare it with

⁶⁴ As we explained in Section ID 2, there are reasons to be skeptical of the common assumption that one can always assign an exactly pure state or purely subjective epistemic state to “the universe as a whole.”

the subsystem eigenprojectors $\hat{P}_{Q_\alpha}(i; t')$. As one special case, our quantum conditional probabilities provide a *kinematical* smoothing relationship (86),

$$\begin{aligned} p_{Q_1, \dots, Q_n | W}(i_1, \dots, i_n | w) &\equiv \text{Tr}_W \left[\left(\hat{P}_{Q_1}(i_1) \otimes \dots \otimes \hat{P}_{Q_n}(i_n) \right) \hat{P}_W(w) \right] \\ &= \langle \Psi_{W,w} | (|\Psi_{Q_1, i_1}\rangle \langle \Psi_{Q_1, i_1}| \otimes \dots \otimes |\Psi_{Q_n, i_n}\rangle \langle \Psi_{Q_n, i_n}|) | \Psi_{W,w} \rangle, \end{aligned}$$

between the possible ontic states of any partitioning collection of mutually disjoint subsystems Q_1, \dots, Q_n and the possible ontic states of the corresponding parent system $W = Q_1 + \dots + Q_n$ at any *single* moment in time. As another special case, if we take $Q \equiv Q_1 = W$, then our quantum conditional probabilities also provide a *dynamical* smoothing relationship (98),

$$p_Q(i'; t' | i; t) \equiv \text{Tr}_Q[\hat{P}_Q(i'; t') \mathcal{E}_Q^{t' \leftarrow t}[\hat{P}_Q(i; t)]] \sim \text{Tr}[\hat{P}_{i'}(t') \mathcal{E}[\hat{P}_i(t)]],$$

between the possible ontic states of the system Q over time and also between the objective epistemic states of the system Q over time.

Essentially, 1 establishes a linkage between ontic states and elements of Hilbert spaces, 2 establishes a linkage between (objective) epistemic states and density matrices, 3 establishes a linkage between parent-system density matrices and subsystem density matrices, and 4 establishes a linkage between parent-system ontic states and subsystem ontic states as well as between parent-system epistemic states and subsystem epistemic states, either at the same time or at different times.

After verifying in Section III D 3 that our quantum conditional probabilities satisfy a number of consistency requirements (81)-(85), we showed in Section III D 10 that they allow us to avoid ontological instabilities that have presented problems for other modal interpretations, analyzed the measurement process in Section IV, studied various familiar “paradoxes” and thought experiments in Section IV D, and examined the status of Lorentz invariance and locality in our interpretation of quantum theory in Section V. In particular, our interpretation accommodates the nonlocality implied by the EPR-Bohm and GHZ-Mermin thought experiments without leading to superluminal signaling, and evades claims purporting to show that interpretations like our own lead to unacceptable ontological contradictions with Lorentz invariance. As a consequence of its compatibility with Lorentz invariance, we claim that our interpretation is capable of encompassing all the familiar quantum models of physical systems in widespread use today, from nonrelativistic point particles to quantum field theories and even string theory. We also compared our interpretation with some of the other prominent interpretations of quantum theory, such as the de Broglie-Bohm pilot-wave interpretation in Section III A 6 and the Everett-DeWitt many-worlds interpretation in Section III A 9, and concluded in Section V E that we could view a density-matrix-centered version of the latter interpretation as being a more manifestly local “Lorenz gauge” of our own interpretation.

B. Classical Versus Quantum, Instrumentalism Versus Realism

In the following table, we compare the axiomatic content of our minimal modal interpretation of quantum theory with the corresponding (though often implicit) axiomatic content of the standard “realist interpretation” of classical physics:

| Realism | Classical | Quantum (Minimal Modal Interpretation) |
|---|---|--|
| 1. Ontic states | An ontic state q is represented by an element of a configuration space \mathcal{C} . (See Section II B 1.) | An ontic state Ψ is represented by an element $ \Psi\rangle$ of a Hilbert space \mathcal{H} (up to overall normalization). (See Section III A 1.) |
| 2. Epistemic states | An epistemic state $\{(p(q), q)\}_q$ consists of ordered pairs relating an epistemic probability $p(q) \in [0, 1]$ to a <i>possible</i> ontic state q , one of which is the system's <i>actual</i> ontic state. (See Section II B 2.) | An epistemic state $\{(p_i, \Psi_i)\}_i$ consists of ordered pairs relating an epistemic probability $p_i \in [0, 1]$ to a <i>possible</i> ontic state Ψ_i , one of which is the system's <i>actual</i> ontic state. (See Section III A 2.) There exists a correspondence (29) between objective epistemic states and density matrices: $\{(p_i, \Psi_i)\}_i \leftrightarrow \hat{\rho} = \sum_i p_i \Psi_i\rangle \langle \Psi_i .$ |
| 3. Relationship between parent-system epistemic states and subsystem epistemic states | The configuration space of a parent system $W = Q + E$ is the Cartesian product $\mathcal{C}_W = \mathcal{C}_Q \times \mathcal{C}_E$ of the configuration spaces of its subsystems Q and E . Given the epistemic state $\{(p_W(w), w)\}_w$ of a parent system $W = Q + E$, the epistemic state of the subsystem Q is obtained by partial summing (that is, marginalizing) the epistemic state of W over the other subsystem E as in (44): $p_Q(q) \equiv \sum_e p_W(w = (q, e)).$ | The Hilbert space of a parent system $W = Q + E$ is the tensor product $\mathcal{H}_W = \mathcal{H}_Q \otimes \mathcal{H}_E$ of the Hilbert spaces of its subsystems Q and E . Given the epistemic state $\{(p_W(w), \Psi_w)\}_w$ of a parent system $W = Q + E$, the density matrix of the subsystem Q is obtained by partial tracing the density matrix of W over the other subsystem E in accordance with our definition (50): $\hat{\rho}_Q \equiv \text{Tr}_E[\hat{\rho}_W] \iff \rho_Q(i, i') \equiv \sum_e \rho_W((i, e), (i', e)).$ |
| 4. (a) Kinematical relationship between parent-system ontic states and subsystem ontic states | (a) Given a parent system $W = Q + E$ in an ontic state $w = (q, e)$, the ontic state of the subsystem Q is q with unit conditional probability and the ontic state of the subsystem E is e with unit conditional probability: $p_{Q,E W}(q, e w = (q, e)) = 1.$ (See Section III A 8.) | Given a parent system $W = Q + E$ with well-defined dynamics in the sense of possessing a linear CPTP dynamical mapping $\mathcal{E}_W^{t' \leftarrow t}$ over the time interval $t' - t$ of interest, we introduce the general class of quantum conditional probabilities (80), $\begin{aligned} p_{Q,E W}(i, e; t' w; t) &\equiv \text{Tr}_W \left[\left(\hat{P}_Q(i; t') \otimes \hat{P}_E(e; t') \right) \mathcal{E}_W^{t' \leftarrow t} [\hat{P}_W(w; t)] \right] \\ &\sim \text{Tr}[\hat{P}_i(t') \hat{P}_e(t') \mathcal{E}[\hat{P}_w(t)]]. \end{aligned}$ (a) In the case $t' = t$, we obtain an instantaneous kinematical relationship (86) between the ontic state of the parent system $W = Q + E$ and the respective ontic states of its subsystems Q and E : $\begin{aligned} p_{Q,E W}(i, e w) &\equiv \text{Tr}_W \left[\left(\hat{P}_Q(i) \otimes \hat{P}_E(e) \right) \hat{P}_W(w) \right]. \end{aligned}$ |
| 4. (b) Dynamical relationship between ontic states at initial time(s) and ontic states at final time | (b) Given a system Q with well-defined dynamics and in ontic states q_1, q_2, \dots at appropriate respective initial times t_1, t_2, \dots (possibly infinitesimally close together), there exists a mapping (16), $\begin{aligned} p_Q(\cdot; t' (\cdot; t_1), (\cdot; t_2), \dots) : & \\ \underbrace{(q_1; t_1), (q_2; t_2), \dots}_{\text{initial data}}, & \underbrace{(q'; t')}_{\text{final data}} \\ \mapsto \underbrace{p_Q(q'; t' (q_1; t_1), (q_2; t_2), \dots)}_{\text{conditional probabilities}}, & \end{aligned}$ which yields the conditional probability for the system to end up in the ontic state q' at the final time t' and which lifts to a multilinear mapping (17) on the epistemic state of the system as a whole: $\begin{aligned} \underbrace{p_Q(q'; t')}_{\text{epistemic state at } t'} &= \sum_{q_1, q_2, \dots} \underbrace{p_Q(q'; t' (q_1; t_1), (q_2; t_2), \dots)}_{\text{conditional probabilities (independent of epistemic states)}} \\ &\quad \times \underbrace{p_Q(q_1; t_1) p_Q(q_2; t_2) \cdots}_{\text{epistemic states at } t_1, t_2, \dots}. \end{aligned}$ | (b) In the case $Q = W$, we obtain a dynamical relationship (98) between the ontic state of Q at the initial time t and its ontic state at the final time t' , $\begin{aligned} p_Q(i'; t' i; t) &\equiv \text{Tr}_Q[\hat{P}_Q(i'; t') \mathcal{E}_Q^{t' \leftarrow t}[\hat{P}_Q(i; t)]] \\ &\sim \text{Tr}[\hat{P}_{i'}(t') \mathcal{E}[\hat{P}_i(t)]], \end{aligned}$ which is just the conditional probability for the system to end up in the ontic state i' at the final time t' given that the system was in the ontic state i at the initial time t , and which lifts to a linear mapping (100) on the epistemic state of the system as a whole: $\underbrace{p_Q(i'; t')}_{\text{epistemic state at } t'} = \sum_i \underbrace{p_Q(i'; t' i; t)}_{\text{conditional probability (independent of epistemic states)}} \underbrace{p_Q(i; t)}_{\text{epistemic state at } t}.$ |

Note that classical physics admits an “instrumentalist interpretation” in which one downplays notions of ontology and works entirely at the level of classical probability theory. In contrast to our own interpretation, the *traditional* formulation of quantum theory (see Appendix 1) has a more natural correspondence to this instrumentalist approach toward classical physics, as we summarize in the following table:

| Instrumentalism | Classical | Quantum |
|-----------------------------------|---|---|
| 1. Epistemology | A system is described by a probability distribution p_a over the elementary outcomes a of a sample space Ω . | A system is described by a density matrix $\hat{\rho} \mapsto \rho_{aa'}$ on a Hilbert space \mathcal{H} having arbitrary orthonormal basis $\{ a\rangle\}_a$. |
| 2. Time evolution | The probability distribution of an approximately closed system evolves through time in an information-conserving manner. (Liouville evolution.) | The density matrix of an approximately closed system evolves through time in an information-conserving manner. (Unitary evolution.) |
| 3. Observables | Each observable Λ of the system corresponds to a random variable consisting of a collection of possible measurement outcomes λ_a that depend on the elementary outcomes a of the sample space Ω . | Each observable Λ of the system corresponds to a Hermitian operator $\hat{\Lambda} \mapsto \Lambda_{aa'}$ on the Hilbert space \mathcal{H} and whose eigenvalues λ_a (choosing $\{ a\rangle\}_a$ to be the orthonormal eigenbasis of $\hat{\Lambda}$) are the possible measurement outcomes. |
| 4. Empirical probabilities | <p>The empirical probability of obtaining a measurement outcome λ for an observable Λ is given by</p> $p(\lambda) = \sum_a p_a P_{\lambda,a},$ <p>where $P_{\lambda,a}$ is the projection function for λ on the sample space Ω and can be regarded as being the conditional probability $p_{\lambda a}$ of obtaining the measurement outcome λ for Λ given the elementary outcome a:</p> $P_{\lambda,a} = p_{\lambda a} = \begin{cases} 1 & \text{for } \lambda_a = \lambda, \\ 0 & \text{for } \lambda_a \neq \lambda. \end{cases}$ <p>These projection functions are idempotent $P_{\lambda,a}^2 = P_{\lambda,a}$, mutually orthogonal $P_{\lambda,a} P_{\lambda',a} = \delta_{\lambda\lambda'} P_{\lambda,a}$, and complete $\sum_\lambda P_{\lambda,a} = 1$. The observable Λ has expectation value</p> $\langle \Lambda \rangle = \sum_a p_a \lambda_a.$ | <p>The empirical probability of obtaining a measurement outcome λ for an observable Λ is given by</p> $p(\lambda) = \text{Tr}[\hat{\rho}\hat{P}_\lambda] = \sum_{a,a'} \rho_{aa'} P_{\lambda,a'a},$ <p>where \hat{P}_λ is the projection operator for λ on the Hilbert space \mathcal{H}:</p> $\hat{P}_\lambda \equiv \sum_{(\lambda_a=\lambda)} a\rangle \langle a $ <p>(choosing $\{ a\rangle\}_a$ to be the orthonormal eigenbasis of $\hat{\Lambda}$). These projection operators are idempotent $\hat{P}_\lambda^2 = \hat{P}_\lambda$, mutually orthogonal $\hat{P}_\lambda \hat{P}_{\lambda'} = \delta_{\lambda\lambda'} \hat{P}_\lambda$, and complete $\sum_\lambda \hat{P}_\lambda = \hat{1}$. The observable Λ has expectation value</p> $\langle \Lambda \rangle = \text{Tr}[\hat{\rho}\hat{\Lambda}] = \sum_{a,a'} \rho_{aa'} \Lambda_{a'a}.$ |
| 5. Epistemic update | <p>After obtaining a measurement outcome λ for an observable Λ, we update the probability distribution according to the Bayesian rule</p> $p_{\text{initial},a} \mapsto p_{\text{final},\lambda,a} \equiv p_{a \lambda} = \frac{p_{\lambda a} p_a}{p(\lambda)} = \frac{P_{\lambda,a} p_a}{\sum_{a'} P_{\lambda,a'} p_{a'}}.$ | <p>After obtaining a measurement outcome λ for an observable Λ, we update the density matrix according to the Von Neumann-Lüders projection postulate:</p> $\hat{\rho}_{\text{initial}} \mapsto \hat{\rho}_{\text{final},\lambda} = \frac{\hat{P}_\lambda \hat{\rho}_{\text{initial}} \hat{P}_\lambda}{\text{Tr}[\hat{P}_\lambda \hat{\rho}_{\text{initial}}]}.$ |

C. Falsifiability and the Role of Decoherence

As some other interpretations do, our own interpretation puts decoherence in a central role for transforming the Born rule from an axiomatic postulate into a derived consequence as part of our larger approach to resolving the measurement problem of quantum theory. We regard it as a positive feature of our interpretation that falsification of these decoherence-based claims would mean falsification of our interpretation. We therefore take great interest in the ongoing arms race between proponents and critics of decoherence, in which critics offer up examples of decoherence coming up short [11, 13–15, 248] and thereby push proponents to argue that increasingly realistic measurement set-ups involving non-negligible environmental interactions resolve the claimed inconsistencies [31, 182, 183, 316].

D. New Criteria for Foundational Work in Quantum Theory

In developing our new interpretation, we uncovered various implicit assumptions and other manifestations of conventional wisdom that have become entrenched in much of the current work on the foundations of quantum theory.

As a consequence, we have developed several new critiques and criteria that transcend our interpretation in particular and that we believe should be acknowledged and addressed explicitly in future research.

1. *Measure-Zero Idealizations and Robustness Against Arbitrarily Small Perturbations*

To start, we believe that many other realist interpretations of quantum theory rely too strongly in fundamental ways on infinitely idealized concepts, which make these interpretations susceptible to invalidation when they are forced to account for physically realistic perturbations of arbitrarily small magnitude. As an important example—and one that we will examine in greater detail below—some interpretative frameworks reserve physical meaning *only* for *exactly* pure state vectors evolving *exactly* according to the Schrödinger equation.⁶⁵ Although idealized assumptions such as these are perfectly acceptable for simplifying practical calculations, especially in instrumentalist approaches, they do not provide a safe foundation for a robust realist interpretation of quantum theory.

2. *Improper Mixtures and the Russian-Doll Problem*

Many realist interpretative frameworks for quantum theory do not directly assign any meaning to systems in improper mixtures, and so either assume that there exist microscopic systems in *exactly* pure states evolving *exactly* unitarily, or, alternatively, begin with the unexamined assertion that all systems are contained within some sort of maximal closed system that is *itself* in an exactly pure state (“the universal wave function”) evolving exactly unitarily. Absent such a state of affairs—in fact, even if such a state of affairs fails to hold only due to imperfections of order one part in a googol—these ostensibly realist interpretations of quantum theory have no mathematical objects on which to base a notion of physical quantum states, nor a way to specify the fundamental dynamics of the theory.⁶⁶ A particularly significant consequence is that the familiar textbook picture of particles evolving as well-defined wave functions breaks down in a severe manner.

Exactly pure states and *exactly* unitary time evolution are unphysical, measure-zero-sharp idealizations that do not hold for any realistic systems. Moreover, as we first explained in Section ID 2, there exists no physical evidence of a maximal closed system (“the universe as a whole”), and, indeed, there are several logical reasons to be suspicious that such a system is always rigorously guaranteed to be definable. Indeed, we know from experience that larger systems with a greater number of degrees of freedom are dynamically *less* closed and are exponentially *more* susceptible to the development of quantum entanglement with their larger environments. Especially given recent developments in cosmology, there is no logically airtight reason to be confident that as we imagine progressively enlarging our scope, the resulting “Russian doll” of increasingly open, outwardly entangled parent systems should suddenly and miraculously terminate at some dynamically closed system in an exactly pure state.

We were mindful of these basic issues in constructing our minimal modal interpretation, which directly addresses these problems: Our interpretation explicitly assigns a meaning to improper mixtures evolving non-unitarily and is fully consistent with the possible non-existence of a maximal closed system, but works just as well even if such system actually exists. All other interpretations of quantum theory should have to grapple seriously with these same concerns.

3. *Unitarity is Only a Useful Idealization*

A related issue is that some realist approaches to quantum foundations assume without justification that unitarity—that is, linear, unitary dynamics of state vectors or density matrices—is a fundamental axiom.⁶⁷ But, as explained above, unitarity is only an *approximation* that holds for idealized systems that are *perfectly* isolated and that therefore *precisely* conserve information.

Situating the case of unitarity in the broader historical development of physics, it’s important to note that many physical principles that were once thought to be fundamental axioms are now regarded more properly as being

⁶⁵ As an example of this point of view, Albert’s description [11] of the textbook algorithm of quantum theory begins in the following way: “Physical situations, physical states of affairs, are represented in this algorithm by vectors. They’re called state vectors. Here’s how that works: Every physical system (that is: every physical object, and every collection of such objects) to begin with is associated in the algorithm with some particular vector space; and the various possible physical states of any such system are stipulated by or correspond to vectors, and more particularly to vectors of length 1, in that system’s associated space; and every such vector is taken to pick out some particular such state; and the states picked out by all those vectors are taken to comprise all of the possible physical situations of that system.”

⁶⁶ Schrödinger addressed this very point in [273] in referring to the absence of state vectors for systems that are entangled with other systems: “The insufficiency of the ψ -function as model replacement rests solely on the fact that one doesn’t always have it.”

⁶⁷ For instance, in [11], Albert defines the fundamental dynamics of quantum theory as follows: “Given the state of any physical system at any ‘initial’ time (given, that is, the vector which represents the state of that system at that time), and given the forces and constraints to which that system is subject, there is a prescription whereby the state of that system at any later time (that is, the vector at any later time) can, in principle, be calculated. There is, in other words, a dynamics of the state vector; there are deterministic laws about how the state vector of any given system, subject to given forces and constraints, changes with time. Those laws are generally cast in the form of an equation of motion, and the name of that equation, for nonrelativistic systems, is the Schrödinger equation.”

derived approximations that follow from deeper, idealized assumptions. These derived principles are often still useful for theoretical and practical purposes, but are no longer assumed to be exact features of physical reality.

Energy conservation provides an archetypal example. Energy conservation is extremely useful in calculations and also for constraining candidate physical models. But although it is sometimes still depicted in elementary textbooks as being a perfectly obeyed aspect of the physical world, energy conservation, at least classically speaking, is now understood to follow from various idealizations: The physical system in question must be closed and have time-translation-invariant dynamics, and, from the standpoint of general relativity, the spacetime in which the system resides must also exhibit an asymptotically timelike Killing vector, a requirement that is not believed to be the case for our expanding universe as a whole. (In order for the total energy of the *gravitational field itself* to make well-defined sense, the spacetime must satisfy even stricter requirements, such as being asymptotically flat.) And, of course, quantum theory implies that even this idealized notion of energy conservation holds only at the level of expectation values.

Unitary dynamics in quantum theory follows a similar pattern. The idealization of unitarity is very useful at a practical level, especially in particle-physics calculations, where it leads to important constraints on scattering amplitudes and can even help identify regimes in which effective field theories break down. But unitarity is not a free-standing, foundational principle that can be taken for granted: It holds only under the assumptions that the physical system under consideration is isolated and conserves information, and there is no hard evidence that there exist any systems—let alone any sensibly defined notion of a “universe as a whole”—that exactly obey these conditions. Indeed, if there exists *just one* sufficiently pathological, nonlinearly evolving system residing in some hidden corner of outer space (a “quantum Russell’s teapot”) whose dynamics does not admit an embedding into a parent system that evolves unitarily, then the existence of such a system would present a clear obstruction to claims of a unitarily evolving universe containing all other systems.

Again, one feature that makes our interpretation distinct from some of the other realist interpretations of quantum theory is that it does not depend in any fundamental way on treating unitarity as a foundational axiom for any physical systems in Nature, but is fully consistent with the alternative proposition that unitarity is always just an approximation.

4. Rethinking the Traditional Starting Place for Quantum Theory

At a broader level, our overall approach to quantum foundations suggests that the *traditional* starting place—state vectors, wave functions, Hilbert spaces, and the Schrödinger equation—may be the *incorrect* starting place for defining and understanding quantum theory.

As we have emphasized in Section III B 1 and Appendix 1, and as we have laid out in our table comparing classical and quantum instrumentalism in Section VI B, there is a strong sense in which the instrumentalist approach to quantum theory is a straightforward noncommutative version of classical probability theory in which the probability distributions, commutative algebras of random variables, projection functions, probability formulas, Bayesian updates, Shannon entropy, and epistemic time evolution of classical probability theory are respectively replaced in quantum theory with density matrices, noncommutative algebras of self-adjoint operators, projection operators, the Born rule, the Von Neumann-Lüders projection formula, Von Neumann entropy, and density-matrix dynamics. Every one of these key ingredients of quantum theory therefore has a very natural classical counterpart. Moreover, as seen from this standpoint, the traditional language of state vectors and the Schrödinger equation consists of nothing more than employing convenient idealizations of purity and information conservation that make calculations more tractable by allowing us to work with N -dimensional vectors rather than with $N \times N$ noncommuting matrices.

In retrospect, it’s not altogether surprising that early formulations of quantum theory featured exotic entities like state vectors and wave functions that lack clear classical counterparts. Dirac’s ground-breaking text [102] comprehensively establishing quantum theory in essentially its traditional form was first published in 1930, years before the 1933 publication of Kolmogorov’s treatise [203] axiomatizing classical probability theory in the framework of measure theory and sample spaces, and both works predated the language of random variables and information theory as we know them today. Given the contingent nature of these historical developments, it might be worthwhile to reconsider the starting place both for teaching quantum theory as well as for studying the theory’s foundations.

As we noted in Section III B 1, taking an approach in which we regard quantum theory as being a noncommutative version of classical probability theory provides yet one more way to motivate our minimal modal interpretation, as it suggests that just as the domain of a classical probability distribution identifies the sample space of the corresponding classical system’s mutually exclusive, elementary ontic states, so too does the diagonalizing basis of a density matrix identify the sample space of a *quantum* system’s mutually exclusive, elementary ontic states.

5. The Real Measurement Problem

One finds in certain treatments of the foundations of quantum theory the measurement problem formulated essentially as the following question: “How can *smooth, deterministic* evolution for *unmeasured* systems coexist with *discontinuous, indeterministic* evolution for a system *being measured*?”⁶⁸ It would be fortunate if that question were indeed the actual measurement problem of quantum theory, because that question was answered a long time ago.

Whether in classical or quantum physics, we expect closed, isolated, information-conserving systems to evolve smoothly, deterministically, and, at the level of epistemic states or density matrices, linearly and unitarily. However, again in classical or quantum physics, we expect *open* systems to evolve discontinuously, indeterministically, and nonlinearly, as we explained in detail in Sections III D 1 and III D 7. As we’ve also explored in this paper, environmental influences in quantum theory typically lead to decoherence, which is a manifestly non-unitary, entropy-increasing transformation of a system’s density matrix that can take the system from a nearly pure state to an improper mixture reflecting the entanglement of the system’s degrees of freedom with those of its environment.⁶⁹

The *real* measurement problem of traditional quantum theory refers to the question about how we should reconcile the *improper* mixtures that result from physical measurement processes performed on a subject system—where the nontriviality of the subject system’s density matrix arises from external entanglement with the measurement device and the larger environment—with the seemingly *proper* mixtures required in the traditional formulation of quantum theory that would describe merely *classical* uncertainty about the subject system’s final, presumably definite, state vector.

Decoherence alone tells us how we get improper mixtures from pure states as the result of a measurement process, but *does not* resolve the question about how to make sense of those improper mixtures, let alone whether or how those *improper* mixtures become *proper* mixtures. In particular, decoherence does not tell us whether we should assume that systems, in fact, end up with definite state vectors at all, or whether, as in our approach to quantum theory, one can attach a direct ontological and epistemic interpretation to improper mixtures themselves.

6. Quantum Physics is No More Linear in Practice than Classical Physics

Quantum theory is sometimes described as a “linear theory.” Indeed, this purported linearity is often taken to be one of quantum theory’s defining features, in part as a way of distinguishing quantum physics from classical physics.⁷⁰

At the level of dynamics, however, the question of linearity is a murky one. It is true that in classical physics, the ontic-level dynamics of an isolated, information-conserving system can be nonlinear, as is well known from many examples phrased in Newtonian, Lagrangian, or Hamiltonian language. But quantum theory, at least as traditionally formulated, is not in a better position: Quantum theory generically *lacks* ontic-level dynamics altogether, such as if the system in question is described by an improper mixture due to external entanglement, a situation that is always true in practice for any realistic quantum system. And in instrumentalist approaches to quantum theory, there is no ontology in the first place to which ontic-level dynamics could even apply.

At best, one might argue that a perfectly isolated, information-conserving quantum system in an *exactly* pure state *must* evolve linearly—that is, according to the Schrödinger equation—whereas the same is not necessarily true in classical physics. Indeed, in the infinitely idealized case of a perfectly isolated quantum system in an exactly pure state, the Schrödinger equation features a notion of *kinematical* linearity that is totally foreign to classical mechanical systems, namely, the freedom to superpose state vectors. But, again, the notion of a perfectly isolated system in an exactly pure state is an unphysical, measure-zero possibility. Furthermore, the claim that the Schrödinger equation describes *linear, ontic-level* dynamics depends on whether we are to regard state vectors—and thus the Schrödinger equation as a whole—as being ontic ingredients of reality or, from a more instrumentalist point of view, as being just formal epistemic tools that result from splitting up pure density matrices into their bra-ket vector factors.

At the level of epistemic states, we have explained in Section III D 2 that quantum theory describes the time evolution of an appropriately idealized system’s density matrix or epistemic state in terms of a linear CPTP mapping. But at this level, quantum theory is no more or less linear than classical physics, in which the epistemic states of

⁶⁸ The dichotomy between these two forms of time evolution was clearly laid out by Von Neumann [323], and has a prominent role in various treatments [11, 38, 111] of quantum theory and its foundations.

⁶⁹ Writing presciently in [273], Schrödinger noted that the disappearance of a quantum system’s state vector as it becomes entangled with a measurement apparatus provided a way of sidestepping the problem of discontinuous measurement-induced collapse: “But it would not be quite right to say that the ψ -function of the object which changes *otherwise* according to a partial differential equation, independent of the observer, should *now* change leap-fashion because of a mental act. For it had disappeared; it was no more. Whatever is not, no more can it change. It is born anew, is reconstituted, is separated out from the entangled knowledge that one has, through an act of perception, which as a matter of fact is not a physical effect on the measured object. From the form in which the ψ -function was last known, to the new in which it reappears, runs no continuous road—it ran indeed through annihilation. Contrasting the two forms, the thing looks like a leap. In truth something of importance happens in between, namely the influence of the two bodies on each other, during which the object possessed no private expectation-catalog nor had any claim thereunto, because it was not independent.”

⁷⁰ For example, in [331], Weinberg disparages an English professor because he “calls quantum theory nonlinear, though it is the only known example of a precisely linear theory.”

analogous classical systems likewise evolve in a linear way, as explained in Sections [IIB 7](#) and [IIID 1](#). By contrast, *open* systems in both classical and quantum physics generically feature *nonlinear* epistemic dynamics, again as we have described in Sections [IIID 1](#) and [IIID 7](#).

Our approach to quantum theory addresses these issues directly. As we have emphasized, one of our key motivations in developing a new interpretation of quantum theory was not only to provide an ontology for generic quantum systems, but also to provide ontic-level dynamics for quantum systems whether or not they are in improper mixtures. As we have explained, this proposed ontic-level dynamics for an open subsystem is only approximate, and the accuracy of this approximation depends on the extent to which the dynamics of the enclosing parent system can itself be approximated as being linear.

7. *Insufficiency of Epistemic Dynamics, or “Hidden Variables Need Hidden Dynamics”*

Our minimal modal interpretation is not the only foundational approach to quantum theory that postulates the existence of a hidden ontic level of reality that lies below an epistemic layer of some kind. In contrast to our interpretation, in which state vectors represent the hidden ontic ingredients of reality and density matrices represent the epistemic layer, a number of other approaches, called “psi-epistemic” interpretations, propose that state vectors themselves constitute an epistemic layer on top of a deeper level of ontic variables. However, once again in contrast to our minimal modal interpretation, these other approaches often fail to specify dynamical laws for their ontic variables, and so run into a problem that we described in Sections [IB](#) and [IIB 7](#), namely, that dynamics at the level of the epistemic layer alone cannot protect the underlying ontic variables from experiencing drastic macroscopic fluctuations.

8. *Unstated Assumptions of Realism*

Classical realism is so intuitively familiar that certain hidden assumptions can be difficult to notice. One such assumption is that there exists a naïve reductionist relationship between the ontologies of parent systems and the ontologies of their subsystems, a supposition that our interpretation of quantum theory highlights and directly challenges, at least for sufficiently microscopic systems.

As summarized in Section [VIB](#), we have argued in this paper that when the hidden assumptions of classical realism are made explicit, then they form a clear correspondence with the axioms that underlie our interpretation of quantum theory. In particular, advocates of other realist interpretative frameworks for quantum theory—especially those who claim that their axiomatic foundations are exceptionally parsimonious—should be required to explain in careful detail how their own axioms line up against the axioms of classical realism.

9. *The Inadequacy of Interpretations of Nonrelativistic Quantum Theory*

It’s still very common to see talks or papers on the foundations of quantum theory that address only nonrelativistic physics, with the assumption that the relativistic regime can somehow be handled later. However, the landscape of quantum foundations is filled with a century’s worth of interpretations that work fairly well at the nonrelativistic level but fail spectacularly when attempting to accommodate special relativity, perhaps the most famous example being the de Broglie-Bohm pilot-wave formulation, which we have critiqued in detail in Section [IIIA 6](#).

In this paper, we have attempted to construct a realist interpretation of quantum theory that is capable of incorporating special relativity. Given the importance of this goal, we believe that other approaches should be expected to grapple with the same task going forward.

10. *Non-Probabilistic Uncertainty*

As we made clear in Section [IIID 4](#), and as examined in works such as [\[3\]](#), the words “uncertain” and “probabilistic” are not synonymous, and there is no logical way to prove that all uncertainty (let alone uncertainty characterizing hidden variables) can necessarily be captured in terms of probabilities.

Our minimal modal interpretation of quantum theory incorporates and deals with this dichotomy very explicitly. Specifically, our interpretation delineates which kinds of statements can be assigned probabilities and which cannot, and in a way that ensures that the empirical predictions of the theory—that is, predictions about what actually shows up in measurement statistics—essentially conform to the probabilistic structure of the Born rule, as we showed explicitly for the case of Von Neumann measurements in Section [IV B 3](#).

We are hardly the first to propose that not all statements in quantum theory admit meaningful conditional probabilities. As just one example, Friedman and Putnam presented a quantum-logic-based approach to eliminating quantum paradoxes in [123] that requires doing away with the assumption that conditional statements involving incompatible observables can be assigned conditional probabilities. It is entirely possible that further work on quantum foundations may require further allowances for non-probabilistic uncertainty.

11. Inexactness of the Naïve Born Rule

As explained in our table comparing classical and quantum instrumentalism in Section VIB, classical probability theory features a rule for extracting probabilities from probability distributions and random variables that has a clear quantum counterpart given by the traditional Born rule phrased generally in terms of density matrices and the linear operators representing observables. Our table also shows that the classical rule for updating probability distributions naturally corresponds to the Von Neumann-Lüders projection of density matrices.

However, the foregoing classical rules can be regarded as being *exact* axioms of classical probability theory. In quantum theory, by contrast, measurement-induced decoherence is able to furnish improperly mixed reduced density matrices that generically agree with the necessary post-measurement proper mixtures only up to exponentially small corrections, as described in Section IVB3. The Born rule and the Von Neumann-Lüders projection are therefore only ever approximate statements, and thus there are good reasons to be suspicious of interpretative frameworks for quantum theory in which they are treated as being *exact*. By contrast, in our interpretation, the Born rule is relegated to an approximate, derived corollary of other axioms.

12. Perspectivalism

Modal interpretations have been criticized in the past for perspectivalism (also known as relationalism), meaning the dependence of ontologies on seemingly arbitrary choices of observer or external reference system. Perspectivalism of this kind, as featured in alternative approaches such as [54, 56], would essentially bring back the “many worlds” of the Everett-DeWitt interpretation in a different guise, because a given system could then possess infinitely many distinct, equally real ontic states simultaneously.

Our interpretation features a much more benign variant of perspectivalism. In the sense that our interpretation assigns ontologies in a system-centric way, as explained in Sections IIIC1 and IVD1, we accept that an ontology cannot be specified until we have singled out a particular system of interest, and that changing our choice of system can change the corresponding ontology. However, this system-centric notion of ontology provides a unique ontology to a given system and has nothing to do with arbitrary choices of observer or external reference system.

13. Regarding Density Matrices as Elementary Ontic Ingredients

There exist alternative density-matrix-centered approaches to quantum foundations in which density matrices themselves are *directly* taken to be the elementary, irreducible ontic ingredients of the theory [335].

Of course, taking *all* density matrices to be elementary ontic ingredients leads to the trouble that epistemic uncertainty over a system’s density matrix produces a convex combination of density matrices that is then *itself* a density matrix, with the consequence that ontic density matrices are radically non-unique [20]. But one can avoid this problem by taking *only* density matrices representing *totally improper mixtures* to be ontic, in which case epistemic convex combinations represent the kind of benign *classical* uncertainty that we describe in Section IVC and therefore do not define new ontic density matrices.

However, improper mixtures generically become increasingly arcane under time evolution due to entanglement, and so an interpretation that regards density matrices as being the most fundamental ontic ingredients of reality suffers from some of the same metaphysical problems as the Everett-DeWitt many-worlds interpretation, albeit without the preferred-basis problem. Moreover, would-be ontic density matrices, which lack an epistemic interpretation, have no obvious classical limit, so it’s unclear how to recover a classical-looking reality from them.

E. Future Directions

1. Understanding and Generalizing the Hilbert-Space Structure of Quantum Theory

The Hilbert-space structure underlying quantum theory is closely related to the principle of linear superposition. Interesting recent work has examined whether this Hilbert-space structure can be motivated from more primitive ideas [170], perhaps by a “purification postulate” asserting that all mixed states should correspond to pure states in some formally defined larger system, together with the known linear structure of classical spaces of mixed epistemic states.

One can argue that requiring the vector space \mathcal{H} to be *complex* is necessary for the existence of energy as an observable, as well as for the existence of states having definite energy, at least for systems that are dynamically closed and thus possess a well-defined Hamiltonian in the first place; however, there exist subtle ways to get around these requirements [294, 295, 346]. And because the Born rule (120) involves *absolute-value-squares* of inner products, one could also, in principle, explore dropping the requirement that the Hilbert space’s inner product must be positive definite, although then avoiding the dynamical appearance of “null” state vectors having vanishing norm (and thus non-normalizable probability) requires a delicate choice of Hamiltonian.⁷¹ One could alternatively try to modify the definition of the inner product to involve a PT transformation [46–48]. These and other approaches [1] to altering the Hilbert-space structure of quantum theory may have interesting implications for our interpretation that would be worth exploring.

2. Coherent States

For a system with N continuously valued degrees of freedom q_α , where $\alpha = 1, \dots, N$, the similarity between the classical Liouville equation (22) and the quantum Liouville equation (24) becomes much closer if we re-express the $2N$ -dimensional classical phase space (q, p) and the Poisson brackets (20) in terms of the N dimensionless complex variables

$$z_\alpha \equiv \frac{1}{\sqrt{2}} \left(\frac{1}{\ell_\alpha} q_\alpha + i \frac{\ell_\alpha}{\hbar} p_\alpha \right) \quad (155)$$

and their complex conjugates z_α^* , where the quantities ℓ_α are characteristic length scales in the system [290] and the factor of \hbar is dictated by dimensional analysis. In that case, introducing the complexified Poisson brackets

$$\{f, g\}_{z, z^*} \equiv \sum_\alpha \left[\frac{\partial f}{\partial z_\alpha} \frac{\partial g}{\partial z_\alpha^*} - \frac{\partial g}{\partial z_\alpha} \frac{\partial f}{\partial z_\alpha^*} \right], \quad (156)$$

the classical Liouville equation takes a much more quantum-looking form, complete with the familiar prefactor of $-i/\hbar$ on the right-hand side:

$$\frac{\partial \rho}{\partial t} = -\frac{i}{\hbar} \{H, \rho\}_{z, z^*} \quad (157)$$

The complex variables z_α are natural for another reason: They label the corresponding quantum system’s coherent states $|z\rangle \equiv |\{z_\alpha\}_\alpha\rangle$ [152, 272], which are defined to be the solutions of the eigenvalue equations $\hat{z}_\alpha |z\rangle = z_\alpha |z\rangle$. For a system in a coherent state $|z\rangle$, the expectation values of the operators \hat{q}_α and \hat{p}_α are respectively given by the real and imaginary parts of z_α according to the formulas $q_\alpha \equiv \sqrt{2}\ell_\alpha \text{Re } z_\alpha$ and $p_\alpha \equiv \sqrt{2}\hbar \text{Im } z_\alpha / \ell_\alpha$.

Coherent states are the closest quantum analogues to classical states in a number important ways:

- They saturate the Heisenberg uncertainty bounds $\Delta q_\alpha \Delta p_\alpha \geq \hbar/2$;
- they have Gaussian wave functions in both coordinate space and momentum space that can both be made to approach delta functions in the limit $\hbar \rightarrow 0$;
- they become mutually orthogonal in the limit of large coordinate separation $\sum_\alpha |z_\alpha - z'_\alpha|^2 \gg 1$;

⁷¹ Unphysical null and negative-norm-squared (“ghost”) state vectors also arise when *formally* enlarging the Hilbert spaces of gauge theories in order to make their symmetries more manifest, but in the present context we are imagining that we treat state vectors of positive norm-squared and state vectors of negative norm-squared as both being physical.

- they each occupy an N -dimensional Gaussian disc of approximate volume $(2\pi\hbar)^N = h^N$ in N -dimensional phase space (q, p) and thereby neatly account for semiclassical phase-space quantization;
- they satisfy the overcompleteness relation $\int (1/\pi^N) d^2 z^N |z\rangle \langle z| = \hat{1}$, with an integration measure $d^2 z^N / \pi^N = dq^N dp^N / h^N$ that exactly replicates the familiar phase-space measure from semiclassical statistical mechanics;
- and, for coupled systems with small interactions between them, the rate at which coherent states become mutually entangled is very low, so that they remain approximately uncorrelated even in the macroscopic limit.⁷²

Moreover, we can regard delta-function coordinate-basis eigenstates $|q\rangle$ as being coherent states in the limit $\ell_\alpha \rightarrow 0$, and delta-function momentum-basis eigenstates $|p\rangle$ as being coherent states in the limit $\ell_\alpha \rightarrow \infty$. If we instead choose the length scales ℓ_α so that the terms in the Hamiltonian that are quadratic in coordinates and momenta are proportional to $\sum_\alpha \hat{z}_\alpha^\dagger \hat{z}_\alpha$, then the corresponding coherent states *remain* coherent states under unitary time evolution over short time intervals, in which case the expectation values $q_\alpha(t)$ and $p_\alpha(t)$ approximately follow the dynamical equations that would hold if instead the system were classical and governed by second-order dynamics; this last fact, in particular, helps explain the ubiquity of second-order dynamics among classical systems with continuous degrees of freedom.

Because there is no real observable whose Hermitian operator’s eigenstates are coherent states, systems don’t end up in coherent states as a consequence of a Von Neumann measurement. Instead, sufficiently large systems with continuously valued degrees of freedom end up approximately in coherent states due to messy environmental perturbations because coherent states are so physically robust [307, 350].

A set of coherent states farther apart in phase space than the Planck constant h forms an approximately orthogonal set, and we can imagine extending such a set to an orthonormal basis suitable for the spectrum of a macroscopic system’s density matrix by including additional state vectors as needed. For small h , coherent states are very, very close to being orthogonal, so this approximation becomes better and better in the formal classical limit $h \rightarrow 0$. It would be worthwhile to investigate this story in greater detail as it relates to our interpretation of quantum theory.

3. The First-Quantized Formalism for Quantum Theories, Quantum Gravity, and Cosmology

In describing first-quantized closed systems of particles or strings [41, 258, 260], as well as in canonical methods for describing quantum gravity [95, 238], it is often useful to work with a generalized Hilbert space consisting of infinitely many copies of the given system’s physical Hilbert space, each copy corresponding to one instant along a suitable temporal parameterization. Unitary dynamics is then expressed as a Hamiltonian constraint equation of the Wheeler-DeWitt form $\hat{\mathcal{H}}|\Psi\rangle = 0$. It would be an intriguing exercise to study how to re-express the formalism of our interpretation of quantum theory in this alternative framework, and, indeed, how to accommodate open systems exhibiting more general linear CPTP dynamics. More broadly, we look forward to exploring quantum features of black holes and cosmology, including the measure problem of eternal inflation [121, 136, 139, 163, 319], within the context of our interpretation.

4. Lorentz Non-Invariance of Reduced Density Matrices

A potential source of trouble for our interpretation of quantum theory arises from the fact that a quantum system’s ontology—and not merely its epistemology—is determined by the system’s *objective* density matrix (as opposed to the *subjective* density matrices considered in Section IV C). As a consequence, any fundamental non-uniqueness in assigning a system an objective density matrix can lead to perspectivalist ambiguities in the system’s ontology in the sense of Section VID 12.

The density matrix $\hat{\rho}$ of an isolated system transforms under the action of the unitary operator $\hat{U}(\Lambda)$ representing a Lorentz transformation Λ according to

$$\hat{\rho} \mapsto \hat{\rho}' = \hat{U}(\Lambda) \hat{\rho} \hat{U}^\dagger(\Lambda). \quad (158)$$

Equivalently, the members $|\Psi_i\rangle$ of the eigenbasis of the system’s density matrix transform as

$$|\Psi_i\rangle \mapsto |\Psi'_i\rangle = \hat{U}(\Lambda) |\Psi_i\rangle, \quad (159)$$

⁷² As explained in [197], “This argument may explain the dominance of the field aspect over the particle aspect for boson[ic] fields.”

meaning that each eigenvector $|\Psi_i\rangle$ corresponds naturally to an associated Lorentz-transformed eigenvector $|\Psi'_i\rangle$ and thus there's no essential change in the system's ontology.

However, any given system is generically a subsystem of a parent system that also includes an environment, and so we must actually perform the Lorentz transformation at the level of this parent system:

$$\hat{\rho}_{\text{parent}} \mapsto \hat{\rho}'_{\text{parent}} = \hat{U}_{\text{parent}}(\Lambda) \hat{\rho}_{\text{parent}} \hat{U}_{\text{parent}}^\dagger(\Lambda). \quad (160)$$

The trouble is that if the overall Hamiltonian of the parent system features nontrivial interactions between the original subsystem and its environment—meaning that the Hamiltonian does not decompose into a sum of mutually commuting terms for the subsystem and for its environment—then, due to well-known peculiarities of the Lorentz algebra, the parent system's Lorentz-transformation operator will not generically factorize as the tensor product of Lorentz-transformation operators for the original subsystem and for the environment separately. That is,

$$\hat{U}_{\text{parent}}(\Lambda) \neq \hat{U}(\Lambda) \otimes \hat{U}_{\text{env}}(\Lambda), \quad (161)$$

and so the reduced density matrix obtained by partial tracing the parent-system transformation rule (160) over the environment will not, in general, agree with the naïve transformation rule (158) that would hold if our original subsystem were isolated:

$$\text{Tr}_{\text{env}}[\hat{\rho}'_{\text{parent}}] = \text{Tr}_{\text{env}}\left[\hat{U}_{\text{parent}}(\Lambda) \hat{\rho}_{\text{parent}} \hat{U}_{\text{parent}}^\dagger(\Lambda)\right] \neq \hat{U}(\Lambda) \hat{\rho} \hat{U}^\dagger(\Lambda). \quad (162)$$

Consequently, there potentially exists an ambiguity in the definition of the subsystem's density matrix, and thus a corresponding ambiguity in identifying the subsystem's corresponding epistemic state and ontic basis, as implicitly noted in [241].

As a way out of this difficulty, one could appeal to the well-known corresponding problem in classical special relativity: Once we introduce the spacetime structure and Lorentz symmetries of classical special relativity, we simultaneously need to introduce new metaphysical criteria for defining what we mean by the fundamental ontological features of physical systems. Various classical ontic properties of objects—such as inertial masses, lengths, clock ticks, and perceived colors—are not Lorentz invariant and only have an intrinsic meaning when referring to the values that they take in a given object's rest frame. Following this example, one might therefore posit that the correct Lorentz frame for assigning a quantum system's reduced density matrix—and thus assigning the system's ontology—is the system's average rest frame, or perhaps the Lorentz frame in which the system's spatial size or Compton wavelength is maximal. These special, individualized notions of Lorentz frame become available as soon as one introduces the notion of Lorentz transformations in the first place.

Singling out a specific Lorentz frame for a given system for the purposes of identifying its ontology in our system-centric interpretation of quantum theory is no more out of step with the spirit of special relativity than is singling out a system's rest frame for identifying its rest mass, rest length, or rest color in the classical context. In particular, we do not have to assume the existence of a *globally* preferred, universe-spanning ontic reference frame that all systems must share.

5. Large Convex Combinations

A deeper source of trouble arises in the case of “large convex combinations” [171]. Suppose that the density matrix of the solar system has the form

$$\begin{aligned} \hat{\rho} = & \alpha \sum_{i \in I} \hat{\rho}_{\text{system},i} \otimes \hat{\rho}_{\text{atoms},i} \otimes \hat{\rho}_{\text{env},i} \\ & + \beta \sum_{i \in \bar{I}} \hat{\rho}_{\text{system},i} \otimes \hat{\rho}_{\text{atoms},i} \otimes \hat{\rho}_{\text{env},i}. \end{aligned}$$

Here α, β are real, positive constants; I and \bar{I} are non-overlapping index sets; and $\hat{\rho}_{\text{system},i}$, $\hat{\rho}_{\text{atoms},i}$, and $\hat{\rho}_{\text{env},i}$ denote positive semi-definite, unit-trace operators respectively for a subject system, a collection of atoms, and a larger environment. Suppose further that for $i \in I$, the operators $\hat{\rho}_{\text{atoms},i}$ each describe a configuration of atoms that constitute a sentient experimenter, whereas for $i \in \bar{I}$, they each describe a random, insentient assemblage of atoms. Then as far as the experimenter is concerned, the observed density matrix of the subject system might seem to involve just the operators $\hat{\rho}_{\text{system},i}$ for $i \in I$, whereas a birds-eye view of the entire solar system might seem to indicate that the correct density matrix of the subject system should involve also the operators $\hat{\rho}_{\text{system},i}$ for $i \in \bar{I}$; the inclusion

of these latter operators $\hat{\rho}_{\text{system},i}$ could conceivably have significant consequences for the underlying ontology of the subject system.

If we adopt a density matrix for the subject system conditioned on the existence of the experimenter, meaning that we include only the values of i in the first index set I , then we risk allowing the ontology of the subject system to depend on our arbitrary choice of external observer, thereby leading to perspectivalism as defined in Section VID 12. On the other hand, if we take the birds-eye view, then the density matrix of the subject system as actually used by the experimenter may differ radically from the true density matrix of the subject system, and thus experimenters can never trust that the density matrices that they use have anything to do with the actual ontologies of the systems that they study.

Taking the birds-eye view, one potential way out of this problem is to note that decoherence should make the diagonalizing bases of generic macroscopic systems relatively insensitive to our choice of $i \in I$ or $i \in I \cup \bar{I}$. In that case, as far as assigning ontologies is concerned, there is no practical trouble with restricting to $i \in I$, and we can feel free to work with a conditional density matrix for a macroscopic subject system to capture the practical epistemology as seen by the experimenter.

F. Relevant Metaphysical Speculations

1. The Status of Superdeterminism

Throughout this paper, we have repeatedly emphasized the importance and nontriviality of the existence of dynamics for a classical or quantum system. In particular, the existence of dynamics for a system is a much stronger property than the mere possession of a *particular* kinematical trajectory. Indeed, one could easily imagine a “superdeterministic” universe in which every system has some specific trajectory—written down at the beginning of time on some mystical “cosmic ledger,” say—but has no dynamics (not even deterministic dynamics) in the sense of the existence of a mapping (16) for *arbitrary* values of initial ontic states and that allows us to compute hypothetical *alternative* trajectories.⁷³

It is therefore a remarkable fact that so many systems in Nature—the Standard Model of particle physics being an example *par excellence*—are well described by dynamics simple enough that we can write them down on a sheet of paper. The time evolution of most systems apparently encodes *far* less information than their complicated trajectories might naïvely suggest—that is, the information encoded in their trajectories is highly compressible—and certainly far less information than could easily be possible for systems belonging to a superdeterministic universe.⁷⁴

2. Implications for Presentism and Block-Time Universes

A perennial debate in metaphysics concerns the fundamental nature of time itself: Presentism is the philosophical proposition that the present moment in time—the “moving now”—is an ontologically real concept, whereas an alternative proposition is that our notion of “the moving now” is merely an illusion experienced by beings like ourselves who actually live in a “block-time” universe whose past extent and future extent are equally real. In the classical deterministic case, we can always depict spacetime in block form, and we can even accommodate nontrivial stochastic dynamics by considering appropriate ensembles of block-time universes [3].

However, according to our interpretation of quantum theory, naïve notions of reductionism generically break down at the microscale, as we explained immediately following (86), and thus epistemic states of microscopic quantum systems each have their own time evolution and cannot generally be sewn together in a classically intuitive manner; indeed, the same may well be true on length scales that exceed the size of our cosmic horizon [65]. Hence, the notion of a block-time universe, which is, in a certain sense, a “many-times” interpretation of physics (“all times are equally real”), may turn out to be just as untenable as the many-worlds interpretation of quantum theory.

⁷³ It’s important to realize that the existence of dynamics—even deterministic dynamics—still leaves open a great deal of flexibility via initial conditions. Building on this idea, Aaronson suggests in [3] that if our universe is not superdeterministic but is instead governed by “merely” deterministic or probabilistically stochastic dynamical laws, then un-cloneable quantum details (“freebits”) of our observable universe’s initial conditions could allow for the kind of non-probabilistic uncertainty that we discussed in Section IID 4 and thus make room for notions of free will, much as the low entropy of our observable universe’s initial conditions makes room for a thermodynamic arrow of time.

⁷⁴ Following Aaronson’s classification of philosophical problems in [3], the Q question “Is the universe superdeterministic?” may be fundamentally unanswerable, but our present discussion suggests the more meaningful and well-defined Q' question “How compressible is trajectory information for the systems that make up the universe?”

3. Sentient Quantum Computers

In Section III A 4 and elsewhere in this paper, we have repeatedly emphasized that our interpretation of quantum theory regards ontic states as being irreducible objects rather than as being epistemic probability distributions over a more basic set of preferred basis states or hidden variables. Because human brains are warm systems in constant contact with a messy larger environment, their reduced density matrices and thus their possible ontic states are guaranteed by decoherence to look classical and not to involve macroscopic quantum superpositions [306]; putting a human brain into an overall quantum superposition would therefore seem to require completely isolating the brain from its environment (including its blood supply) and lowering its temperature to nearly absolute zero, in which case no human awareness would be conceivable. We can therefore sidestep metaphysical questions about the subjective first-person experiences of human observers who temporarily exist in quantum superpositions of mutually exclusive, classical-looking state vectors.

However, one might imagine someday building a sentient quantum computer capable of human-level intelligence. In principle, and in contrast to human beings, such a machine would be perfectly functional even at a temperature close to absolute zero and without any need for continual interactions with a larger environment. Could such a machine achieve something like subjective first-person experiences, and, if so, how would the machine experience existing in a quantum superposition of mutually exclusive state vectors? Or is there something about the existence of subjective first-person experiences that fundamentally requires continual decoherence and information exchange with a larger environment?⁷⁵

A recurring problem in the history of interpretations of quantum theory is the tendency for the subject to become mixed up with other thorny problems in philosophy. Just as our interpretation deliberately aims to be model independent and thereby attempts to avoid getting tangled up in the philosophical debate over the rigorous meaning of probability and its many schools of interpretation (from Laplacianism to frequentism to propensity to Bayesianism to decision theory), we also intentionally avoid making any definitive statements about a preferred philosophy of mind and the important metaphysical problem of understanding the connection between the physical third-person reality of atoms, planets, and galaxies and the subjective first-person phenomenology of colors, thoughts, and emotions. By design, our interpretation concerns itself solely with physical third-person reality, and does not favor any of the schools of thought on the reality of first-person experiences (from dualism to eliminativism to functionalism to panpsychism). That said, one of our motivations in developing a *realist* interpretation of quantum theory was in order to provide the physical world with a notion of physical systems (such as human brains) and physical states (such as brain configurations and patterns of synaptic activation) upon which the mental states of human minds could supervene.

4. Nonlocality in the Everett-DeWitt Many-Worlds Interpretation

Although often claimed to be manifestly local, the Everett-DeWitt many-worlds interpretation of quantum theory can only maintain this manifest locality by abandoning either any sharply defined probabilistic branching structure that spans all systems or else by abandoning assertions that it can solve the preferred-basis problem, as described in Section III A 9.

To see explicitly how this trouble arises, consider again the standard EPR-Bohm experiment that we originally examined in Section V B and which consists of a pair of spin-1/2 particles 1 and 2 together with a pair of spin detectors A (which locally measures the spin of particle 1) and B (which locally measures the spin of particle 2). By assumption, the initial state vector of the composite system $1 + 2 + A + B$ is given by

$$|\Psi_{1+2+A+B}\rangle = \frac{1}{\sqrt{2}} (|\uparrow\downarrow\rangle - |\downarrow\uparrow\rangle) |A(\text{“}\emptyset\text{”})\rangle |B(\text{“}\emptyset\text{”})\rangle.$$

How should one regard this state vector in terms of branches (“worlds”) and their associated probabilities according to the many-worlds interpretation?

If the claim is that there exists just one branch that possesses unit probability and in which the state vector of the two-particle system $1 + 2$ is really $|\Psi_{1+2}\rangle = (1/\sqrt{2})(|\uparrow\downarrow\rangle - |\downarrow\uparrow\rangle)$, then one runs into trouble with nonlocality, because after the spin detector A performs its local measurement on particle 1, two branches instantaneously emerge in which the widely separated particles 1 and 2 are suddenly classically correlated ($\uparrow\downarrow$ in one branch and $\downarrow\uparrow$ in the other branch) despite not having been classically correlated before.

⁷⁵ We thank Scott Aaronson (private communication [4]) for suggesting these points.

An alternative approach is to argue that there were actually two branches with respective probabilities of $1/2$ all along, meaning that one should have regarded the initial state vector of the composite system $1 + 2 + A + B$ as really being

$$|\Psi_{1+2+A+B}\rangle = \frac{1}{\sqrt{2}} |\uparrow\downarrow\rangle |A(\text{“}\emptyset\text{”})\rangle |B(\text{“}\emptyset\text{”})\rangle - \frac{1}{\sqrt{2}} |\downarrow\uparrow\rangle |A(\text{“}\emptyset\text{”})\rangle |B(\text{“}\emptyset\text{”})\rangle.$$

In that case, the two particles were really classically correlated in each branch even before the experiment and thus no “new” classical correlation suddenly appears nonlocally after the spin detector A carries out its local measurement on particle 1. But then one runs into a serious problem making the notion of the branches and their probabilities well-defined, because there is nothing that privileges splitting up the branches in the spin- z basis for the two-particle system $1 + 2$; indeed, one could just as well have split up the branches in the spin- x basis ($|\leftarrow\rangle$ and $|\rightarrow\rangle$) instead, in which case the initial state vector of the composite system $1 + 2 + A + B$ would have been better written as

$$|\Psi_{1+2+A+B}\rangle = \frac{1}{\sqrt{2}} |\leftarrow\rightarrow\rangle |A(\text{“}\emptyset\text{”})\rangle |B(\text{“}\emptyset\text{”})\rangle - \frac{1}{\sqrt{2}} |\rightarrow\leftarrow\rangle |A(\text{“}\emptyset\text{”})\rangle |B(\text{“}\emptyset\text{”})\rangle.$$

One is therefore confronted directly with the preferred-basis problem, and decoherence cannot help resolve the paradox because the spin detectors A and B haven’t performed their measurements yet and thus all the various possible bases are on an equal footing.⁷⁶

So what are the branches? Which basis is the “correct” one for deciding? As we have seen in this section, if we pick one preferred definition for the branches that span the systems under consideration—and also define their associated probabilities—then we immediately run into the nonlocality issue again; for example, if we pick, say, the spin- x branches, and then A performs a spin- z measurement, then again the branches need to change in a nonlocal way in order to ensure the correct final classical correlation. If we declare that we must be noncommittal about assigning entangled systems to branches, then we run into the trouble that entanglement is a ubiquitous phenomenon afflicting *all* systems to a nonzero degree, and thus we don’t obtain sharp definitions of what we mean by branches. Because there is no fixed choice of branch-set compatible with manifest locality, the many-worlds interpretation isn’t manifestly local unless we give up any notion of a preferred branch set that spans systems, but then we lose any hope of making sense of branch probabilities in the interpretation.

G. Comparison with the Hollowood Modal Interpretation

Our minimal modal interpretation of quantum theory differs in several key respects from the recently introduced “emergent Copenhagen interpretation” by Hollowood [181–183], with whom we have collaborated in the past. Specifically, we employ a different, manifestly non-negative, more general formula for our quantum conditional probabilities that doesn’t depend *fundamentally* on a temporal cut-off time scale; we use these quantum conditional probabilities not only for dynamical purposes but also to interpolate between the ontologies of parent systems and their sub-systems; our quantum conditional probabilities allow ontic states of disjoint systems to influence each other; joint ontic-state and epistemic-state assignments exist for mutually disjoint systems; and we accept the inevitability of hidden ontic-level nonlocality implied by the EPR-Bell and GHZ-Mermin thought experiments.

ACKNOWLEDGMENTS

J. A. B. has benefited tremendously from personal communications with Matthew Leifer, Timothy Hollowood, Francesco Buscemi, Alisa Bokulich, Elise Crull, and Gregg Jaeger. D. K. thanks Gaurav Khanna, Pontus Ahlqvist, Adam Brown, Darya Krym, and Paul Cadden-Zimansky for many useful discussions on related topics, and has been supported in part by FQXi minigrant #FQXi-MGB-1322. Both authors have greatly appreciated their interactions with the Harvard Philosophy of Science group and its organizers Andrew Friedman and Elizabeth Petrik, as well as with Scott Aaronson, Ned Hall, and Brian Greene. Both authors are indebted to Steven Weinberg, who has generously exchanged relevant ideas and whose own recent paper [335] also explores the idea of reformulating quantum theory in terms of density matrices. In addition, the work of Pieter Vermaas has been a continuing source of inspiration for both authors, as have many conversations with Allan Blaer. The authors would also like to thank Perimeter Institute,

⁷⁶ Note that these conclusions cannot be evaded by the assumption of small degeneracy-breaking effects of the form (34), which would merely have the effect of increasing the total number of branches.

and, in particular, Lucien Hardy and Matthew Pusey, for graciously inviting and hosting them to share their results. This research was supported in part by Perimeter Institute for Theoretical Physics. Research at Perimeter Institute is supported by the Government of Canada through Industry Canada and by the Province of Ontario through the Ministry of Economic Development & Innovation.

APPENDIX

In this appendix, we summarize the traditional Copenhagen interpretation, in part just to establish our notation and terminology. With this background established, we then define the measurement problem and systematically analyze attempts to solve it according to the various prominent interpretations of quantum theory, including the instrumentalist approach. Finally, we describe several important theorems that have been developed over the years to constrain candidate interpretations of quantum theory.

1. The Copenhagen Interpretation and the Measurement Problem

a. A Review of the Copenhagen Interpretation

The textbook Copenhagen interpretation asserts that every isolated quantum system is completely described by a particular unit-norm state vector $|\Psi\rangle$ in an associated Hilbert space \mathcal{H} . Furthermore, according to the Copenhagen interpretation, every observable property Λ of the system corresponds to a Hermitian operator $\hat{\Lambda} = \hat{\Lambda}^\dagger$ necessarily having a complete orthonormal basis of eigenstates $|a\rangle$ with corresponding real eigenvalues λ_a :

$$\hat{\Lambda}|a\rangle = \lambda_a|a\rangle, \quad \lambda_a \in \mathbb{R}, \quad \langle a|a'\rangle = \delta_{aa'}, \quad \sum_a |a\rangle\langle a| = \hat{1}. \quad (163)$$

The eigenvalues λ_a represent the possible measurement outcomes of the random variable Λ , and the empirical outcome probability $p(\lambda)$ with which a particular outcome $\lambda = \lambda_a$ is obtained is given by the Born rule,

$$p(\lambda) = \sum_{\substack{a \\ (\lambda_a=\lambda)}} |\langle a|\Psi\rangle|^2, \quad (164)$$

where, in order to accommodate the case of degeneracies in the eigenvalue spectrum of $\hat{\Lambda}$, the sum is over all values of the label a for which the eigenvalue λ_a of the eigenstate $|a\rangle$ is equal to the specified outcome λ . When degeneracy is absent, so that we can uniquely label the eigenstates of $\hat{\Lambda}$ by λ , the Born rule (164) reduces to the simpler expression

$$p(\lambda) = |\langle \lambda|\Psi\rangle|^2. \quad (165)$$

Either way, we can then express expectation values of observables Λ in terms of the system's state vector $|\Psi\rangle$ in the following way:

$$\langle \Lambda \rangle = \langle \Psi|\hat{\Lambda}|\Psi\rangle. \quad (166)$$

In particular, letting \hat{P}_λ denote the Hermitian projection operator onto eigenstates of $\hat{\Lambda}$ having the eigenvalue λ ,

$$\hat{P}_\lambda \equiv \sum_{\substack{a \\ (\lambda_a=\lambda)}} |a\rangle\langle a|, \quad (167)$$

we can use (166) to express the Born rule (164) in the alternative form

$$p(\lambda) = \langle P_\lambda \rangle = \langle \Psi|\hat{P}_\lambda|\Psi\rangle. \quad (168)$$

These statements all naturally extend to density matrices $\hat{\rho}$ according to the formulas

$$p(\lambda) = \text{Tr}[\hat{\rho}\hat{P}_\lambda], \quad \langle \Lambda \rangle = \text{Tr}[\hat{\rho}\hat{\Lambda}], \quad (169)$$

which we can identify as noncommutative generalizations of the respective classical formulas

$$p(\lambda) = \sum_a p_a P_{\lambda,a}, \quad \langle \Lambda \rangle = \sum_a p_a \lambda_a. \quad (170)$$

In these classical formulas, the quantities $p_a \in [0, 1]$ constitute a classical probability distribution over elementary outcomes a , and

$$P_{\lambda,a} \equiv \begin{cases} 1 & \text{for } \lambda_a = \lambda, \\ 0 & \text{for } \lambda_a \neq \lambda \end{cases} \quad (171)$$

is the projection (or indicator) function for the subset of elementary outcomes a corresponding to the generalized outcome λ . Moreover, just as the classical projection functions (171) are idempotent, mutually orthogonal, and complete,

$$P_{\lambda,a}^2 = P_{\lambda,a}, \quad P_{\lambda,a}P_{\lambda',a} = \delta_{\lambda\lambda'}P_{\lambda,a}, \quad \sum_{\lambda} P_{\lambda,a} = 1,$$

so too are the quantum projection operators (167):

$$\hat{P}_{\lambda}^2 = \hat{P}_{\lambda}, \quad \hat{P}_{\lambda}\hat{P}_{\lambda'} = \delta_{\lambda\lambda'}\hat{P}_{\lambda}, \quad \sum_{\lambda} \hat{P}_{\lambda} = \hat{1}.$$

Quantum systems that are dynamically closed undergo smooth, linear time evolution according to a unitary time-evolution operator $\hat{U}(t)$:

$$|\Psi(t)\rangle = \hat{U}(t)|\Psi(0)\rangle \quad \left[\hat{U}(t)^\dagger = \hat{U}(t)^{-1} = \hat{U}(-t), \quad \langle\Psi(t)|\Psi(t)\rangle = \langle\Psi(0)|\Psi(0)\rangle = 1 \right]. \quad (172)$$

If we can express the time-evolution operator in terms of a Hermitian, time-independent Hamiltonian operator \hat{H} according to

$$\hat{U}(t) = e^{-i\hat{H}t/\hbar}, \quad (173)$$

then the system's state vector obeys the famous Schrödinger equation:

$$i\hbar \frac{\partial}{\partial t} |\Psi(t)\rangle = \hat{H} |\Psi(t)\rangle. \quad (174)$$

At the level of the system's density matrix $\hat{\rho}(t)$, these integral and differential dynamical equations respectively become

$$\hat{\rho}(t) = \hat{U}(t)\hat{\rho}(0)\hat{U}(t)^\dagger, \quad \frac{\partial}{\partial t}\hat{\rho}(t) = -\frac{i}{\hbar}[\hat{H}, \hat{\rho}(t)]. \quad (175)$$

The Copenhagen interpretation axiomatically regards the Born rule (164) as being an exact postulate and defines the partial-trace prescription together with reduced density matrices for subsystems precisely to ensure that the Born-rule-based expectation value (169) obtained for any observable $\hat{\Lambda}$ of any subsystem agrees with the expectation value of the corresponding observable $\hat{\Lambda} \otimes \hat{1}$ of any parent system:

$$\text{Tr}_{\text{subsystem}}[\hat{\rho}_{\text{subsystem}}\hat{\Lambda}] = \underbrace{\text{Tr}_{\text{parent}}}_{\text{Tr}_{\text{subsystem}}\text{Tr}_{\text{other}}}[\hat{\rho}_{\text{parent}}\hat{\Lambda} \otimes \hat{1}] \implies \hat{\rho}_{\text{subsystem}} = \text{Tr}_{\text{other}}[\hat{\rho}_{\text{parent}}]. \quad (176)$$

Consequently, the dynamics governing the reduced density matrix $\hat{\rho}(t)$ of a general open subsystem of a dynamically closed parent system is determined by an equation that is generically non-unitary and, indeed, not even necessarily linear in $\hat{\rho}(t)$:

$$\frac{\partial}{\partial t}\hat{\rho}(t) = -\frac{i}{\hbar}\text{Tr}_{\text{parent}}[[\hat{H}_{\text{parent}}, \hat{\rho}_{\text{parent}}]]. \quad (177)$$

The final axiom of the Copenhagen interpretation stipulates that if an external observer measures an observable Λ of the system and obtains an outcome λ corresponding to the projection operator \hat{P}_{λ} defined in (167), then the system's density matrix instantaneously “collapses” according to a non-unitary rule called the Von Neumann-Lüders projection postulate [227, 323]: Regardless of the system's initial density matrix $\hat{\rho}_{\text{initial}}$, which may well be pure $|\Psi_{\text{initial}}\rangle\langle\Psi_{\text{initial}}|$ or may instead describe an improper mixture—meaning that the nontriviality of the density matrix

arises at least in part from quantum entanglement with other systems—the system’s *final* density matrix $\hat{\rho}_{\text{final},\lambda}$ is given by

$$\hat{\rho}_{\text{final},\lambda} = \frac{\hat{P}_\lambda \hat{\rho}_{\text{initial}} \hat{P}_\lambda}{\text{Tr} [\hat{P}_\lambda \hat{\rho}_{\text{initial}}]}. \quad (178)$$

The denominator ensures that the final density matrix $\hat{\rho}_{\text{final},\lambda}$ has unit trace, and, from (169), is just the initial probability $p(\lambda)$ of obtaining the outcome λ . In the special case in which the eigenvalue spectrum of the operator $\hat{\Lambda}$ representing the observable Λ has no degeneracies, so that the projection operator \hat{P}_λ reduces to the simple form $|\lambda\rangle\langle\lambda|$ for a single eigenstate $|\lambda\rangle$ of $\hat{\Lambda}$, the Von Neumann-Lüders projection postulate (178) reduces to

$$\hat{\rho}_{\text{final},\lambda} = |\lambda\rangle\langle\lambda|, \quad (179)$$

so that the system ends up in a pure state represented by $|\lambda\rangle$.

b. Wave-Function Collapse and the Measurement Problem

The Von Neumann-Lüders projection postulate (178)—known informally as wave-function collapse—both accounts for the statistical features of the post-measurement state of affairs and also ensures that definite measurement outcomes *persist* under identical repeated experiments performed over sufficiently short time intervals. However, (178), which holds for systems that are open enough to be measured by an external apparatus, represents a discontinuous departure from the smooth time evolution experienced by closed quantum systems as determined by the Schrödinger equation (174), a discrepancy that forms part of the historical measurement problem of quantum theory.

As a first step toward better characterizing the measurement problem, notice that we can gather together the different possible post-measurement density matrices $\hat{\rho}_{\text{final},\lambda}$ appearing in (178) into a subjective probability distribution that we can represent by

$$\{p(\lambda) \implies \hat{\rho}_{\text{final},\lambda}\}_\lambda. \quad (180)$$

This subjective probability distribution over the different possible final density matrices $\hat{\rho}_{\text{final},\lambda}$ yields the same statistical predictions as the block-diagonal “subjective” density matrix

$$\hat{\rho}_{\text{final}} = \sum_\lambda p(\lambda) \hat{\rho}_{\text{final},\lambda} = \sum_\lambda \hat{P}_\lambda \hat{\rho}_{\text{initial}} \hat{P}_\lambda, \quad (181)$$

which describes the post-measurement system in the absence of post-selecting or conditioning on the actually observed value λ . Note that the specific decomposition appearing on the right-hand side of (181) is *preferred* among all possible decompositions of $\hat{\rho}_{\text{final}}$ because we know that the system’s *true* density matrix is *really* one of the possibilities $\hat{\rho}_{\text{final},\lambda}$ obtained from the Von Neumann-Lüders projection postulate (178). That is, in the present situation, our system is fundamentally described by the subjective probability distribution (180), and we have introduced the subjective density matrix (182) merely for mathematical convenience.

In the special case (179) in which the set of possible measurement outcomes exhibits no degeneracies, the subjective post-measurement density matrix (181) reduces to a proper mixture, meaning that the nontriviality of the density matrix arises *solely* from subjective uncertainty over the true underlying state vector $|\lambda\rangle$ of the system:

$$\hat{\rho}_{\text{final}} = \sum_\lambda p(\lambda) |\lambda\rangle\langle\lambda|. \quad (182)$$

c. Decoherence

Remarkably, provided that the measurement in question is performed by a sufficiently macroscopic observer or measuring device, the quantum phenomenon of decoherence [58, 67, 196, 197, 269, 270, 349], described more fully in Section IV B 3, naturally produces *reduced* (“objective”) density matrices describing *improper* mixtures that look *formally* just like the subjective density matrix (182) up to tiny corrections, with eigenstates that exhibit negligible quantum interference with one another under further time evolution. However, the subjective density matrix (182), unlike a decoherence-generated density matrix, is not a *reduced* density matrix arising from external quantum entanglement; instead, the subjective density matrix (182) is ultimately just a mathematically convenient stand-in for

the subjective probability distribution (180) over *persistent* measurement outcomes, and, in the simplest case (182), is a proper mixture. Moreover, the Copenhagen interpretation provides no canonical recipe for assigning preferred decompositions to the *reduced* density matrices that describe improper mixtures.

Hence, as emphasized in Section VID 5, the measurement problem essentially reduces to the question not of how pure states evolve to become mixed states, a phenomenon adequately explained by decoherence, but of how *reduced* density matrices representing *improper* mixtures arising from measurement-induced decoherence can turn into *subjective* density matrices (which, in the simplest case, represent *proper* mixtures) describing definite outcomes with corresponding probabilities. Solving this manifestation of the measurement problem requires either additional ingredients or a new interpretation altogether.

d. A Variety of Approaches

Within the framework of the textbook Copenhagen interpretation, one simply declares that observers or measurement devices that are “sufficiently classical”—meaning that they are on the classical side of the so-called Heisenberg cut—cause decoherence-generated density matrices to *cease* being *reduced* density matrices by *breaking* their quantum entanglement with any external systems, and, furthermore, cause them to develop the necessary preferred decomposition appearing on the right-hand side of (181) or (182); the overall effect is therefore to convert decoherence-generated reduced density matrices into subjective probability distributions of the form (180). Equivalently, one can phrase the Heisenberg cut as a threshold on the amount of quantum interference between the eigenstates of a decoherence-generated density matrix: If the amount of quantum interference falls below that threshold, then we can treat those eigenstates as describing classical possibilities in a subjective probability distribution (180). Unfortunately, according to either formulation, the Heisenberg cut remains ill-defined and has not been identified in any experiment so far.

An alternative approach is to postulate that *all* reduced density matrices gradually evolve into corresponding subjective probability distributions (180), with macroscopic systems evolving in this way more rapidly than microscopic systems. Achieving these effects requires altering the basic dynamical structure of quantum theory along the lines of GRW dynamical-collapse or spontaneous-localization constructions [8, 39, 144, 249, 252, 334].

The Everett-DeWitt many-worlds interpretation [68, 93, 94, 96, 111–113, 325, 326, 337] attempts to solve the measurement problem by reifying all the members of a suitably chosen basis as simultaneous “worlds” or “branches” while somehow also regarding them as the elements of an appropriate subjective probability distribution. Apart from the difficulty of trying to make sense of probabilities when all outcomes are simultaneously realized, one key trouble with this interpretation is deciding which basis to choose, a quandary known as the preferred-basis problem and that we describe in Sections III A 9 and VIF 4.

The modal interpretations instead reify *just one member* of a suitable basis. In “fixed” modal interpretations—of which the de Broglie-Bohm interpretation [59, 60, 63, 91] is the most well-known example—this preferred basis is fixed for all systems, leading to problems that we detail in Section III A 6. By contrast, in density-matrix-centered modal interpretations such as the one that we introduce in this paper, one chooses the basis to be the diagonalizing eigenbasis of the density matrix of whatever system is currently under consideration. Essentially, the central idea of our own minimal modal interpretation is a conservative one: We identify improper density matrices as closely as possible with subjective probability distributions, and add the minimal axiomatic ingredients that are necessary to make this identification viable.

e. The Instrumentalist Approach

A final prominent option is the instrumentalist approach, in which one formally accepts the basic axioms of the Copenhagen interpretation—regarding them as describing a kind of noncommutative version of classical Bayesian probability theory—but without taking a definitive stand on the ontological meaning of state vectors or density matrices, or on any reality that underlies the formalism of quantum theory more generally.

As motivation, we begin by noticing that the classical projection function $P_{\lambda,a}$ defined in (171) for the generalized outcome λ of an observable Λ agrees numerically with the conditional probability $p_{\lambda|a}$ of the generalized outcome λ given the elementary outcome a , and so we can recast the formula (170) for the probability $p(\lambda)$ of the generalized outcome λ as the Bayesian propagation rule

$$p(\lambda) = \sum_a p_{\lambda|a} p_a.$$

We can therefore analogously regard the Born-rule formula for $p(\lambda)$ in (169) as describing a quantum version of a Bayesian propagation rule.

Similarly, we can regard the Von Neumann-Lüders projection postulate (178) as being the natural noncommutative generalization of the classical Bayesian post-measurement probability-update formula for the conditional probability $p_{a|\lambda}$ of obtaining the elementary outcome a given that a prior measurement has yielded the generalized outcome λ :

$$p_{\text{initial},a} \mapsto p_{\text{final},\lambda,a} = p_{a|\lambda} = \frac{p_{\lambda|a} p_a}{p(\lambda)} = \frac{P_{\lambda,a} p_a}{\sum_{a'} P_{\lambda,a'} p_{a'}}. \quad (183)$$

Merely replacing the ill-defined notions of an observer and the Heisenberg cut with the notion of an abstract, unphysical agent carrying out idealized measurements and updating probabilities using Bayes' theorem is not, in itself, progress. Moreover, attempting to regard the Von Neumann-Lüders projection as being an instantaneous quantum version of a classical post-measurement probability update runs into trouble with the quantum Zeno paradox described in Section IV D 2.

However, by borrowing from our discussion of decoherence earlier in this section, we can eliminate the need for a Heisenberg cut and define more concretely what we mean by an agent by introducing an appropriate self-consistency condition. Specifically, in order for a system such as a creature or a measurement device to count as an agent, we can require that whenever it performs generic measurements and thereby causes measurement-induced decoherence, the observable Q representing the true-or-false question “Did the agent obtain a set of frequency ratios in agreement with the Born formula?” and the observable Q' representing the true-or-false question “Did the agent find a persistent measurement outcome that appears to be in keeping with the Von Neumann-Lüders projection postulate?” each have respective Born-rule probabilities $p(Q = \text{“true”})$ and $p(Q' = \text{“true”})$ computed from (169) that are acceptably close to unity—say, 99.9%. Provided that the agent is sufficiently macroscopic, decoherence will generally guarantee that these conditions hold.

It is important to keep in mind, however, that the Born rule (169) and the Von Neumann-Lüders projection postulate (178) are still nontrivial axioms and cannot be dropped within the instrumentalist approach: Without these axioms, we cannot derive the statements $p(Q = \text{“true”}) \approx 1$ and $p(Q' = \text{“true”}) \approx 1$ in the first place, nor can we conclude *solely* from those statements that sufficiently macroscopic measurement devices (such as human beings) will experience the Born rule and the Von Neumann-Lüders projection postulate, the reason being that neither of the associated Hermitian operators \hat{Q} and \hat{Q}' actually correspond to *unique* questions; indeed, each of these operators generically has two *highly* degenerate eigenvalues 1 (“true”) and 0 (“false”), meaning that neither \hat{Q} nor \hat{Q}' can uniquely pick out an orthonormal basis of states describing classically sensible realities. Thus, at best, the conditions $p(Q = \text{“true”}) \approx 1$ and $p(Q' = \text{“true”}) \approx 1$ merely supply us with a self-consistency check—that is, a necessary but not sufficient condition—on our definition of agents and a quantitative criterion for determining how macroscopic a valid agent must be, rather than making possible an *ab initio* derivation of the Born rule or the Von Neumann-Lüders projection postulate.

A somewhat more serious issue is that the decoherence process, while exponentially rapid, never completely finishes, and thus the reduced density matrix of a system following measurement-induced decoherence never precisely agrees with the subjective density matrix (181) that describes definite measurement outcomes and their associated probabilities. Even if we decided to increase our self-consistency threshold from 99.9% to a value closer to unity, instrumentalism requires that we replace reduced density matrices representing improper mixtures with subjective density matrices differing by exponentially small corrections. While unimportant in most practical cases, these corrections are presumably inherent in all quantities that ultimately depend on the Born rule—including all probabilities, expectation values, tunneling rates, and semiclassical observables—and may have a deep significance for some of quantum theory's fundamental theoretical problems, such as making sense of black-hole evaporation. We discuss these corrections and their possible implications in Section IV B 3.

2. Foundational Theorems

We discuss the Bell theorem in Section VB of the main text, and the Myrvold theorem in Section VD. Here we present a brief proof of the Bell theorem and describe several other important theorems that put strong constraints on candidate interpretations of quantum theory.

a. Proof of the Bell Theorem

As we explained in Section VB, the Bell theorem involves a pair of spin detectors A and B aligned respectively along two unit vectors \vec{a} and \vec{b} and that make measurements on pairs of spin-1/2 particles governed by local hidden variables λ . Granting the observed fact that spin is quantized in units of $\hbar/2$ (one can account for this condition

in classical language by insisting that the particles always automatically line themselves up along the local detector alignments as they are being measured), Bell assumed that the respective results $A(\vec{a}, \lambda) = \pm 1$ and $B(\vec{b}, \lambda) = \pm 1$ (in units of $\hbar/2$) of the two spin detectors depend only on data local to each detector. In keeping with the given anti-correlated state (135) of each pair of spin-1/2 particles, Bell also required that if the two detectors are aligned, then they must always measure opposite spins:

$$A(\vec{a}, \lambda) = -B(\vec{b} = \vec{a}, \lambda).$$

Finally, Bell posited the existence of a probability distribution (138) for the hidden variables λ themselves:

$$0 \leq p(\lambda) \leq 1, \quad \int d\lambda p(\lambda) = 1.$$

Introducing an alternative unit vector \vec{c} for the orientation of the spin detector B , the Bell theorem then asserts that the three average spin correlations

$$\begin{aligned} \langle S_{1,\vec{a}} S_{2,\vec{b}} \rangle_{\text{LHV}} &= \int d\lambda p(\lambda) A(\vec{a}, \lambda) B(\vec{b}, \lambda), \\ \langle S_{1,\vec{a}} S_{2,\vec{c}} \rangle_{\text{LHV}} &= \int d\lambda p(\lambda) A(\vec{a}, \lambda) B(\vec{c}, \lambda), \\ \langle S_{1,\vec{b}} S_{2,\vec{c}} \rangle_{\text{LHV}} &= \int d\lambda p(\lambda) A(\vec{b}, \lambda) B(\vec{c}, \lambda) \end{aligned}$$

must satisfy the inequality (137),

$$\left| \langle S_{1,\vec{a}} S_{2,\vec{b}} \rangle_{\text{LHV}} - \langle S_{1,\vec{a}} S_{2,\vec{c}} \rangle_{\text{LHV}} \right| \leq 1 + \langle S_{1,\vec{b}} S_{2,\vec{c}} \rangle_{\text{LHV}},$$

as follows from a straightforward computation:

$$\begin{aligned} \left| \langle S_{1,\vec{a}} S_{2,\vec{b}} \rangle_{\text{LHV}} - \langle S_{1,\vec{a}} S_{2,\vec{c}} \rangle_{\text{LHV}} \right| &= \left| \int d\lambda p(\lambda) \left(A(\vec{a}, \lambda) \underbrace{B(\vec{b}, \lambda)}_{-A(\vec{b}, \lambda)} - \underbrace{1}_{(A(\vec{b}, \lambda))^2} A(\vec{a}, \lambda) B(\vec{c}, \lambda) \right) \right| \\ &\leq \int d\lambda p(\lambda) \underbrace{|A(\vec{a}, \lambda) A(\vec{b}, \lambda)|}_{\pm 1} \underbrace{(1 + A(\vec{b}, \lambda) B(\vec{c}, \lambda))}_{>0} \\ &= 1 + \langle S_{1,\vec{b}} S_{2,\vec{c}} \rangle_{\text{LHV}}. \quad \text{QED} \end{aligned}$$

b. Gleason's Theorem

Gleason's theorem [153] asserts that the only consistent probability measures for closed subspaces of a Hilbert space of dimension ≥ 3 must be given by $p_i = \text{Tr}[\hat{\rho} \hat{P}_i]$, which generalizes the state-vector version of the Born rule. Here $\hat{\rho}$ is a unit-trace, positive semi-definite operator that we interpret as the system's density matrix and $\hat{P}_i = |\Psi_i\rangle \langle \Psi_i|$ is a projection operator onto some state vector $|\Psi_i\rangle$.

All the well-known interpretations of quantum theory satisfy Gleason's theorem. In particular, the basic correspondence (29) at the heart of our own minimal modal interpretation is consistent with the theorem.

c. The Kochen-Specker Theorem

It might seem plausible for an interpretation of quantum theory to claim that when we measure observables $\Lambda_1, \Lambda_2, \dots$ belonging to a quantum system and obtain some set of outcome values $\lambda_1, \lambda_2, \dots$, we are merely revealing that those observables secretly possessed all those values $\lambda_1, \lambda_2, \dots$ simultaneously even before the measurement took place. The Kochen-Specker theorem [202] presents an obstruction to such claims. More precisely, the theorem

rules out hidden-variables interpretations that assert that all observables that *could* be measured have simultaneous, sharply defined values *before* they are measured—values that depend *only* on the observables themselves and on the system to which they belong—and that measurements merely reveal those supposedly preexisting values. For viable hidden-variables interpretations of quantum theory, an immediate consequence of the theorem is that the properties of a quantum system must generically be contextual, in the sense that they can depend on the kinds of measurements performed on the system. Here we present one of the simplest versions of the theorem, due to Peres [254].

We begin by considering a four-state quantum system, which we can regard as consisting of a pair of two-state spin-1/2 subsystems and therefore having a Hilbert space \mathcal{H} of the form

$$\mathcal{H} = \mathcal{H}_{1/2} \otimes \mathcal{H}_{1/2}, \quad \dim \mathcal{H} = 2 \times 2 = 4. \quad (184)$$

We consider also the following nine Hermitian matrices on this four-dimensional Hilbert space:

$$\left. \begin{aligned} P_{2z} &\equiv (1 - \sigma_{2z})/2, & P_{1z} &\equiv (1 - \sigma_{1z})/2, & P_{1z,2z} &\equiv (1 - \sigma_{1z}\sigma_{2z})/2, \\ P_{1x} &\equiv (1 - \sigma_{1x})/2, & P_{2x} &\equiv (1 - \sigma_{2x})/2, & P_{1x,2x} &\equiv (1 - \sigma_{1x}\sigma_{2x})/2, \\ P_{1x,2z} &\equiv (1 - \sigma_{1x}\sigma_{2z})/2, & P_{1z,2x} &\equiv (1 - \sigma_{1z}\sigma_{2x})/2, & P_{1y,2y} &\equiv (1 - \sigma_{1y}\sigma_{2y})/2. \end{aligned} \right\} \quad (185)$$

In our notation here, $1 \equiv 1_{2 \times 2} \otimes 1_{2 \times 2}$ is the 4×4 identity matrix, and $\sigma_{1i} \equiv \sigma_i \otimes 1_{2 \times 2}$ and $\sigma_{2i} \equiv 1_{2 \times 2} \otimes \sigma_i$ are respectively the Pauli sigma matrices for the first and second spin-1/2 subsystems.

Each 4×4 matrix appearing in (185) has two +1 eigenvalues and two -1 eigenvalues; in the usual language of quantum theory, the associated observables therefore have possible measured values +1 and -1. Furthermore, all three matrices in each column of (185) commute with each other, and, likewise, all three matrices in each row commute with each other; hence, for each column or row, the three corresponding observables are mutually compatible and thus can be made simultaneously sharply defined by an appropriate preparation of the state of our system.

We next define three new 4×4 matrices A_1, A_2, A_3 by respectively summing each column of (185),

$$\left. \begin{aligned} A_1 &\equiv P_{2z} + P_{1x} + P_{1x,2z} && (\text{eigenvalues } +2, +2, 0, 0), \\ A_2 &\equiv P_{1z} + P_{2x} + P_{1z,2x} && (\text{eigenvalues } +2, +2, 0, 0), \\ A_3 &\equiv P_{1z,2z} + P_{1x,2x} + P_{1y,2y} && (\text{eigenvalues } +3, +3, +1, +1), \end{aligned} \right\} \quad (186)$$

and three 4×4 matrices B_1, B_2, B_3 by respectively summing each row of (185),

$$\left. \begin{aligned} B_1 &\equiv P_{2z} + P_{1z} + P_{1z,2z} && (\text{eigenvalues } +2, +2, 0, 0), \\ B_2 &\equiv P_{1x} + P_{2x} + P_{1x,2x} && (\text{eigenvalues } +2, +2, 0, 0), \\ B_3 &\equiv P_{1x,2z} + P_{1z,2x} + P_{1y,2y} && (\text{eigenvalues } +2, +2, 0, 0). \end{aligned} \right\} \quad (187)$$

Finally, we introduce an observable Σ defined to be the sum (times two) of the nine original observables defined in (185), or, equivalently, defined to be the sum of the six observables $A_1, A_2, A_3, B_1, B_2, B_3$:

$$\begin{aligned} \Sigma &\equiv 2P_{2z} + 2P_{1x} + 2P_{1x,2z} \\ &\quad + 2P_{1z} + 2P_{2x} + 2P_{1z,2x} \\ &\quad + 2P_{1z,2z} + 2P_{1x,2x} + 2P_{1y,2y} \\ &= A_1 + A_2 + A_3 + B_1 + B_2 + B_3. \end{aligned} \quad (188)$$

Each of the nine observables defined in (185) has permissible values +1 or -1, and so, from the first expression for Σ in (188), we see that the existence of simultaneous pre-measurement values for these nine observables would imply that Σ has an *even* pre-measurement value:

$$\left. \begin{aligned} \Sigma &= 2 \times (+1 \text{ or } -1) + 2 \times (+1 \text{ or } -1) + 2 \times (+1 \text{ or } -1) \\ &\quad + 2 \times (+1 \text{ or } -1) + 2 \times (+1 \text{ or } -1) + 2 \times (+1 \text{ or } -1) \\ &= \text{even}. \end{aligned} \right\} \quad (189)$$

However, if we assume that the six observables $A_1, A_2, A_3, B_1, B_2, B_3$ likewise have simultaneous pre-measurement values, then we find instead that Σ has an *odd* pre-measurement value, a contradiction:

$$\left. \begin{aligned} \Sigma &= (0 \text{ or } 2) + (0 \text{ or } 2) + (1 \text{ or } 3) + (0 \text{ or } 2) + (0 \text{ or } 2) + (0 \text{ or } 2) \\ &= \text{odd}. \end{aligned} \right\} \quad (190)$$

We are therefore forced to give up our assumption that a quantum system can always have simultaneous sharply defined pre-measurement values for all its observables. QED

Neither the Copenhagen interpretation nor the minimal modal interpretation that we introduce in this paper asserts that all observables have well-defined values before measurements. Indeed, our own interpretation makes manifest that a system's set of possible ontic states can change contextually from one orthonormal basis of the system's Hilbert space to another orthonormal basis in the course of interactions with other systems. Thus, both interpretations are consistent with the Kochen-Specker theorem. The de Broglie-Bohm pilot-wave interpretation evades this theorem as well, because although the interpretation assumes that systems have hidden well-defined values of both canonical coordinates and canonical momenta at all times, the canonical momenta are not identified with the *observable* momenta that actually show up in measurements.

Note that regardless of one's interpretation of quantum theory, decoherence ensures that typical observables for *macroscopic* systems simultaneously develop what are perceived to be *approximate* pre-measurement values, in agreement with classical expectations.

d. The Pusey-Barrett-Rudolph (PBR) Theorem

The Pusey-Barrett-Rudolph (PBR) theorem [37, 87, 264] rules out so-called psi-epistemic interpretations [172] of quantum theory that directly regard *state vectors* (as opposed to density matrices) as merely being epistemic probability distributions for hidden variables whose configurations determine the outcomes of measurements. Essentially, the theorem derives a one-to-one correspondence between each specific configuration of hidden variables and each state vector (up to trivial overall phase factor), thereby implying that each configuration singles out a unique state vector. It is therefore impossible to regard each state vector as being an epistemic probability distribution over a nontrivial collection of *different* configurations of hidden variables.

To illustrate the theorem in a simple case, the authors consider a two-state system with an orthonormal basis $|0\rangle, |1\rangle$ and a second orthonormal basis defined by

$$|+\rangle \equiv \frac{1}{\sqrt{2}} (|0\rangle + |1\rangle), \quad |-\rangle \equiv \frac{1}{\sqrt{2}} (|0\rangle - |1\rangle). \quad (191)$$

If the two state vectors $|0\rangle$ and $|+\rangle$ merely describe probability distributions over configurations of hidden variables, and if those probability distributions are allowed to overlap on the sample space of configurations of hidden variables, then there exists some nonzero probability $p > 0$ that a particular configuration λ of hidden variables will reside in both probability distributions. Hence, if the system's hidden variables have the configuration λ , then the corresponding probability distribution could be either $|0\rangle$ or $|+\rangle$.

If we now set up a pair of independent such systems as a composite system, then, with probability p^2 , the configurations λ_1 and λ_2 of the two subsystems would permit being jointly described by the probability distributions arising from any of the possible tensor-product state vectors $|0\rangle|0\rangle, |0\rangle|+\rangle, |+\rangle|0\rangle, |+\rangle|+\rangle$. But then a measurement of an observable whose corresponding basis of orthonormal eigenstates is

$$\left. \begin{aligned} |\xi_1\rangle &= \frac{1}{\sqrt{2}} (|0\rangle|1\rangle + |1\rangle|0\rangle), \\ |\xi_2\rangle &= \frac{1}{\sqrt{2}} (|0\rangle|-\rangle + |1\rangle|+\rangle), \\ |\xi_3\rangle &= \frac{1}{\sqrt{2}} (|+\rangle|1\rangle + |-\rangle|0\rangle), \\ |\xi_4\rangle &= \frac{1}{\sqrt{2}} (|+\rangle|-\rangle + |-\rangle|+\rangle) \end{aligned} \right\} \quad (192)$$

would have zero probability of yielding $|\xi_1\rangle$ if the composite system's state vector happened to be $|0\rangle|0\rangle$, zero probability of yielding $|\xi_2\rangle$ if the composite system's state vector happened to be $|0\rangle|+\rangle$, zero probability of yielding $|\xi_3\rangle$ if the composite system's state vector happened to be $|+\rangle|0\rangle$, and zero probability of yielding $|\xi_4\rangle$ if the composite system's state vector happened to be $|+\rangle|+\rangle$. Hence, whichever result is found by the measurement, there exists a contradiction with the notion that the original simultaneous configurations λ_1 and λ_2 of the conjoined subsystems were really compatible with *all four* possible state vectors $|0\rangle|0\rangle, |0\rangle|+\rangle, |+\rangle|0\rangle, |+\rangle|+\rangle$. The authors then extend this general argument to a much larger class of possibilities beyond the particular pair $|0\rangle$ and $|+\rangle$ to argue that no two state vectors could ever describe overlapping probability distributions over hidden variables.

The Copenhagen interpretation and our own interpretation satisfy the PBR theorem, because they both regard state vectors as irreducible features of reality rather than as mere epistemic probability distributions over a deeper

layer of hidden variables. The de Broglie-Bohm pilot-wave interpretation also evades the theorem because, in that interpretation, the state vector plays both the role of an epistemic probability distribution as well as a physical pilot wave that (nonlocally if necessary) guides the hidden variables during measurements to values that are always consistent with final measurement outcomes.

e. The Fine and Vermaas No-Go Theorems

Finally, there exist no-go theorems due to Fine [118, 119] and Vermaas [315] suggesting that axiomatically imposing joint probability distributions for properties of *non-disjoint* subsystems of a larger parent system leads to contradictions with axiomatic impositions of joint probability distributions for properties of *disjoint* subsystems. In our minimal modal interpretation of quantum theory, we do not postulate joint probability distributions for non-disjoint subsystems of a given parent system, and so our interpretation evades both theorems.

-
- [1] Scott Aaronson. “Is Quantum Mechanics An Island In Theoryspace?”. 2004. [arXiv:quant-ph/0401062](https://arxiv.org/abs/quant-ph/0401062).
 - [2] Scott Aaronson. “Quantum Computing and Hidden Variables I: Mapping Unitary to Stochastic Matrices”. 2004. [arXiv:quant-ph/0408035](https://arxiv.org/abs/quant-ph/0408035).
 - [3] Scott Aaronson. “The Ghost in the Quantum Turing Machine”. 2013. [arXiv:1306.0159](https://arxiv.org/abs/1306.0159).
 - [4] Scott Aaronson. Private communication, 2014.
 - [5] Scott Aaronson and Luke (interviewer) Muehlhauser. Scott Aaronson on Philosophical Progress. December 2013. In particular, in discussing quantum computing: “You can’t just ‘fan out’ and have one branch try each possible solution—twenty years of popular articles notwithstanding, that’s not how it works! We also know today that you can’t encode more than about n classical bits into n quantum bits (qubits), in such a way that you can reliably retrieve any one of the bits afterward. And we have all [*sic*] lots of other results that make quantum-mechanical amplitudes feel more like ‘just souped-up versions of classical probabilities,’ and quantum superposition feel more like just a souped-up kind of potentiality. I love how the mathematician Boris Tsirelson summarized the situation: he said that ‘a quantum possibility is more real than a classical possibility, but less real than a classical reality.’ It’s an ontological category that our pre-mathematical, pre-quantum intuitions just don’t have a good name for”. URL: <http://intelligence.org/2013/12/13/aaronson/>.
 - [6] Stephen L. Adler. “Statistical Dynamics of Global Unitary Invariant Matrix Models as Pre-Quantum Mechanics”. 2002. [arXiv:hep-th/0206120](https://arxiv.org/abs/hep-th/0206120).
 - [7] Stephen L. Adler. “Incorporating gravity into trace dynamics: the induced gravitational action”. *Classical and Quantum Gravity*, 30(19):195015, 2013. URL: <http://stacks.iop.org/0264-9381/30/i=19/a=195015>, [arXiv:1306.0482](https://arxiv.org/abs/1306.0482), doi: [10.1088/0264-9381/30/19/195015](https://doi.org/10.1088/0264-9381/30/19/195015).
 - [8] Stephen L. Adler and Angelo Bassi. “Is Quantum Theory Exact?”. *Science*, 325:275–276, 2009. URL: <http://www.sciencemag.org/content/325/5938/275.abstract>, [arXiv:0912.2211](https://arxiv.org/abs/0912.2211), doi: [10.1126/science.1176858](https://doi.org/10.1126/science.1176858).
 - [9] Anthony Aguirre and Max Tegmark. “Born in an infinite universe: A cosmological interpretation of quantum mechanics”. *Physical Review D*, 84(10):105002, 2011. URL: <http://link.aps.org/doi/10.1103/PhysRevD.84.105002>, [arXiv:1008.1066](https://arxiv.org/abs/1008.1066), doi: [10.1103/PhysRevD.84.105002](https://doi.org/10.1103/PhysRevD.84.105002).
 - [10] Ofer Aharony, Steven S. Gubser, Juan Maldacena, Hiroshi Ooguri, and Yaron Oz. “Large N Field Theories, String Theory and Gravity”. *Physics Reports*, 323(3–4):183–386, 2000. URL: <http://www.sciencedirect.com/science/article/pii/S0370157399000836>, [arXiv:hep-th/9905111](https://arxiv.org/abs/hep-th/9905111), doi: [10.1016/S0370-1573\(99\)00083-6](https://doi.org/10.1016/S0370-1573(99)00083-6).
 - [11] David Z. Albert. *Quantum Mechanics and Experience*. Harvard University Press, 1st edition, 1994. For the question of degeneracies, see Appendix: The Kochen-Healy-Dieks Interpretations.
 - [12] David Z. Albert and Barry Loewer. “Interpreting the many worlds interpretation”. *Synthese*, 77(2):195–213, 1988. URL: [article/10.1007/2FBF00869434](https://doi.org/10.1007/2FBF00869434), doi: [10.1007/BF00869434](https://doi.org/10.1007/BF00869434).
 - [13] David Z. Albert and Barry Loewer. “Wanted Dead or Alive: Two Attempts to Solve Schrödinger’s Paradox”. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, 1:277–285, 1990. URL: <http://www.jstor.org/stable/192710>, doi: [10.1007/BF00485415](https://doi.org/10.1007/BF00485415).
 - [14] David Z. Albert and Barry Loewer. “The measurement problem: Some “solutions””. *Synthese*, 86(1):87–98, 1991. URL: <http://link.springer.com/article/10.1007/2FBF00485415>, doi: [10.1007/BF00485415](https://doi.org/10.1007/BF00485415).
 - [15] David Z. Albert and Barry Loewer. “Non-ideal measurements”. *Foundations of Physics Letters*, 6(4):297–305, 1993. URL: [http://dx.doi.org/10.1007/BF00665649](https://doi.org/10.1007/BF00665649), doi: [10.1007/BF00665649](https://doi.org/10.1007/BF00665649).
 - [16] Janet Anders and Vittorio Giovannetti. “Thermodynamics of discrete quantum processes”. *New Journal of Physics*, 15(3):033022, March 2013. URL: <http://stacks.iop.org/1367-2630/15/i=3/a=033022>, [arXiv:1211.0183](https://arxiv.org/abs/1211.0183), doi: [10.1088/1367-2630/15/3/033022](https://doi.org/10.1088/1367-2630/15/3/033022).
 - [17] Dionysios Anninos. “De Sitter Musings”. *International Journal of Modern Physics A*, 27(13), 2012. URL: <http://www.worldscientific.com/doi/abs/10.1142/S0217751X1230013X>, [arXiv:1205.3855](https://arxiv.org/abs/1205.3855), doi: [10.1142/S0217751X1230013X](https://doi.org/10.1142/S0217751X1230013X).
 - [18] Markus Ansmann, H. Wang, Radoslaw C. Bialczak, Max Hofheinz, Erik Lucero, M. Neeley, A. D. O’Connell, D. Sank, M. Weides, J. Wenner, A. N. Cleland, and John M. Martinis. “Violation of Bell’s inequality in Josephson phase qubits”.

- Nature*, 461:504–506, September 2009. URL: <http://www.nature.com/nature/journal/v461/n7263/full/nature08363.html>, doi:10.1038/nature08363.
- [19] D. M. Appleby, Åsa Ericsson, and Christopher A. Fuchs. “Properties of QBist State Spaces”. *Foundations of Physics*, 41(3):564–579, 2011. URL: <http://dx.doi.org/10.1007/s10701-010-9458-7>, doi:10.1007/s10701-010-9458-7.
- [20] Mateus Araújo. “Quantum realism and quantum surrealism”. 2012. Section 1.2.1 (“The naïve ontology”). [arXiv:1208.6283](https://arxiv.org/abs/1208.6283).
- [21] Juan Sebastián Ardenghi, Mario Castagnino, and Olimpia Lombardi. “Modal-Hamiltonian interpretation of quantum mechanics and Casimir operators: the road toward quantum field theory”. *International Journal of Theoretical Physics*, 50(3):774–791, 2011. URL: <http://dx.doi.org/10.1007/s10773-010-0614-9>, [arXiv:1012.0970](https://arxiv.org/abs/1012.0970), doi:10.1007/s10773-010-0614-9.
- [22] Markus Arndt, Angelo Bassi, Domenico Giulini, Antoine Heidmann, and Jean-Michel Raimond. “Fundamental Frontiers of Quantum Science and Technology”. *Procedia Computer Science*, 7(0):77–80, 2011. Proceedings of the 2nd European Future Technologies Conference and Exhibition 2011 (FET 11). URL: <http://www.sciencedirect.com/science/article/pii/S1877050911006892>, doi:10.1016/j.procs.2011.12.024.
- [23] Markus Arndt, Olaf Nairz, Julian Vos-Andreae, Claudia Keller, Gerbrand van der Zouw, and Anton Zeilinger. “Wave-particle duality of C₆₀ molecules”. *Nature*, 401:680–682, 1999. URL: <http://www.nature.com/nature/journal/v401/n6754/abs/401680a0.html>, doi:10.1038/44348.
- [24] Alain Aspect, Jean Dalibard, and Gérard Roger. “Experimental Test of Bell’s Inequalities Using Time-Varying Analyzers”. *Physical Review Letters*, 49:1804–1807, December 1982. URL: <http://link.aps.org/doi/10.1103/PhysRevLetters49.1804>, doi:10.1103/PhysRevLetters49.1804.
- [25] Alain Aspect, Philippe Grangier, and Gérard Roger. “Experimental Tests of Realistic Local Theories via Bell’s Theorem”. *Physical Review Letters*, 47:460–463, August 1981. URL: <http://link.aps.org/doi/10.1103/PhysRevLetters47.460>, doi:10.1103/PhysRevLetters47.460.
- [26] Paul S. Aspinwall, Brian R. Greene, and David R. Morrison. “Calabi-Yau moduli space, mirror manifolds and spacetime topology change in string theory”. *Nuclear Physics B*, 416(2):414–480, 1994. URL: <http://www.sciencedirect.com/science/article/pii/0550321394903212>, [arXiv:hep-th/9309097](https://arxiv.org/abs/hep-th/9309097), doi:10.1016/0550-3213(94)90321-2.
- [27] Guido Bacciagaluppi. “Kochen-Specker theorem in the modal interpretation of quantum mechanics”. *International Journal of Theoretical Physics*, 34(8):1205–1216, 1995.
- [28] Guido Bacciagaluppi and W. Michael Dickson. “Dynamics for Density Operator Interpretations of Quantum Theory”. 1997. [arXiv:quant-ph/9711048](https://arxiv.org/abs/quant-ph/9711048).
- [29] Guido Bacciagaluppi and W. Michael Dickson. “Dynamics for Modal Interpretations”. *Foundations of Physics*, 29(8):1165–1201, 1999. URL: <http://link.springer.com/article/10.1023/A:1018803613886>, doi:10.1023/A:1018803613886.
- [30] Guido Bacciagaluppi, Matthew J. Donald, and Pieter E. Vermaas. “Continuity and discontinuity of definite properties in the modal interpretation”. *Helvetica Physica Acta*, 68(7):679–704, 1995.
- [31] Guido Bacciagaluppi and Meir Hemmo. “Modal Interpretations, Decoherence and Measurements”. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 27(3):239–277, 1996. URL: <http://www.sciencedirect.com/science/article/pii/S1355219896000020>, doi:10.1016/S1355-2198(96)00002-0.
- [32] Vijay Balasubramanian and Bartłomiej Czech. “Quantitative approaches to information recovery from black holes”. *Classical and Quantum Gravity*, 28(16):163001, 2011. URL: <http://stacks.iop.org/0264-9381/28/i=16/a=163001>, [arXiv:1102.3566](https://arxiv.org/abs/1102.3566), doi:10.1088/0264-9381/28/16/163001.
- [33] Vijay Balasubramanian, Michael B. McDermott, and Mark Van Raamsdonk. “Momentum-space entanglement and renormalization in quantum field theory”. *Physical Review D*, 86(4):045014, 2012. URL: <http://link.aps.org/doi/10.1103/PhysRevD.86.045014>, [arXiv:1108.3568](https://arxiv.org/abs/1108.3568), doi:10.1103/PhysRevD.86.045014.
- [34] Leslie E. Ballentine. “The Statistical Interpretation of Quantum Mechanics”. *Reviews of Modern Physics*, 42(4):358–381, 1970. URL: <http://link.aps.org/doi/10.1103/RevModPhys.42.358>, doi:10.1103/RevModPhys.42.358.
- [35] Leslie E. Ballentine. *Quantum Mechanics: A Modern Development*. World Scientific, 1998. With a treatment of the watched-pot paradox in Section 12.2 (“Exponential and Nonexponential Decay”).
- [36] Jacob A. Barandes and David Kagan. “A Synopsis of the Minimal Modal Interpretation of Quantum Theory”. 2014. To be published. [arXiv:1405.6754](https://arxiv.org/abs/1405.6754).
- [37] Jonathan Barrett, Eric G. Cavalcanti, Raymond Lal, and Owen J. E. Maroney. “No ψ -epistemic model can fully explain the indistinguishability of quantum states”. 2013. [arXiv:1310.8302](https://arxiv.org/abs/1310.8302).
- [38] Jean-Louis Basdevant and Jean Dalibard. *Quantum Mechanics*, pages 444–445. Springer, 1st edition, 2002. Appendix D, Section 4.2 (“Evolution of a Reduced Density Operator”).
- [39] Angelo Bassi and Gian Carlo Ghirardi. “Dynamical reduction models”. *Physical Reports*, 379(5):257–426, 2003. URL: <http://www.sciencedirect.com/science/article/pii/S0370157303001030>, [arXiv:quant-ph/0302164](https://arxiv.org/abs/quant-ph/0302164), doi:10.1016/S0370-1573(03)00103-0.
- [40] Angelo Bassi and Emiliano Ippoliti. “Geometric phase for open quantum systems and stochastic unravellings”. 2005. [arXiv:quant-ph/0510184](https://arxiv.org/abs/quant-ph/0510184).
- [41] Katrin Becker, Melanie Becker, and John H. Schwarz. *String Theory and M-Theory: A Modern Introduction*. Cambridge University Press, 2007. With the relevant discussion for point particles in Section 2.1 (“ p -brane actions”).
- [42] Jacob D. Bekenstein. “Universal upper bound on the entropy-to-energy ratio for bounded systems”. *Physical Review D*, 23:287–298, January 1981. URL: <http://link.aps.org/doi/10.1103/PhysRevD.23.287>, doi:10.1103/PhysRevD.23.287.
- [43] John S. Bell. “On the Einstein-Podolsky-Rosen Paradox”. *Physics*, 1(3):195–200, 1964.

- [44] John S. Bell. In *Foundations of Quantum Mechanics, Proceedings of the International School of Physics “Enrico Fermi,” Course XLIX*, page 171. Academic, New York, 1971.
- [45] John S. Bell. *Speakable and Unspeakable in Quantum Mechanics: Collected papers on quantum philosophy*. Cambridge University Press, 2004. With Bell’s discussion of beables in Chapter 19 (“Beables for Quantum Field Theory”), which also includes his discussion of fermionic fields.
- [46] Carl M. Bender. “Introduction to PT-Symmetric Quantum Theory”. 2005. [arXiv:quant-ph/0501052](https://arxiv.org/abs/quant-ph/0501052).
- [47] Carl M. Bender and Stefan Boettcher. “Real Spectra in Non-Hermitian Hamiltonians Having PT Symmetry”. *Physical Review Letters*, 80(24):5243–5246, June 1998. URL: <http://link.aps.org/doi/10.1103/PhysRevLett.80.5243>, [arXiv:physics/9712001](https://arxiv.org/abs/physics/9712001), doi:10.1103/PhysRevLett.80.5243.
- [48] Carl M. Bender, Stefan Boettcher, and Peter N. Meisinger. “PT-Symmetric Quantum Mechanics”. *Journal of Mathematical Physics*, 40(5):2201–2229, 1998. URL: http://jmp.aip.org/resource/1/jmapaq/v40/i5/p2201_s1, [arXiv:quant-ph/9809072](https://arxiv.org/abs/quant-ph/9809072), doi:10.1063/1.532860.
- [49] Gyula Bene. “On the nature of the quantum states of macroscopic systems”. 1997. [arXiv:quant-ph/9711072](https://arxiv.org/abs/quant-ph/9711072).
- [50] Gyula Bene. “Quantum phenomena do not violate the principle of locality—a new interpretation with physical consequences”. 1997. [arXiv:quant-ph/9706043](https://arxiv.org/abs/quant-ph/9706043).
- [51] Gyula Bene. “Quantum reference systems: a new framework for quantum mechanics”. *Physica A: Statistical Mechanics and its Applications*, 242(3):529–565, 1997. URL: <http://www.sciencedirect.com/science/article/pii/S0378437197002549>, [arXiv:quant-ph/9703021](https://arxiv.org/abs/quant-ph/9703021), doi:10.1016/S0378-4371(97)00254-9.
- [52] Gyula Bene. “Quantum reference systems: reconciling locality with quantum mechanics”. *Decoherence and its Implications in Quantum Computation and Information Transfer*, 182:98, 2001. [arXiv:quant-ph/0008128](https://arxiv.org/abs/quant-ph/0008128).
- [53] Gyula Bene. “Relational modal interpretation for relativistic quantum field theories”. 2001. [arXiv:quant-ph/0104111](https://arxiv.org/abs/quant-ph/0104111).
- [54] Gyula Bene and Dennis Dieks. “A perspectival version of the modal interpretation of quantum mechanics and the origin of macroscopic behavior”. *Foundations of Physics*, 32(5):645–671, 2002. URL: <http://dx.doi.org/10.1023/A%3A1016014008418>, [arXiv:quant-ph/0112134](https://arxiv.org/abs/quant-ph/0112134), doi:10.1023/A:1016014008418.
- [55] Felix Alexandrovich Berezin. *The Method of Second Quantization*. Academic Press, 1966.
- [56] Joseph Berkovitz and Meir Hemmo. “Modal Interpretations of Quantum Mechanics and Relativity: A Reconsideration”. *Foundations of Physics*, 35(3):373–397, 2005. URL: <http://link.springer.com/article/10.1007/s10701-004-1980-z>, doi:10.1007/s10701-004-1980-z.
- [57] Garrett Birkhoff and John Von Neumann. “The Logic of Quantum Mechanics”. *Annals of Mathematics*, 37(4):823–843, 1936. URL: <http://www.jstor.org/stable/1968621>.
- [58] David Bohm. *Quantum Theory*. Courier Dover Publications, 1951. With Bohm’s prescient discovery of decoherence in Chapters 22–23.
- [59] David Bohm. “A Suggested Interpretation of the Quantum Theory in Terms of “Hidden” Variables. I”. *Physical Review*, 85:166–179, January 1952. URL: <http://link.aps.org/doi/10.1103/PhysRev.85.166>, doi:10.1103/PhysRev.85.166.
- [60] David Bohm. “A Suggested Interpretation of the Quantum Theory in Terms of “Hidden” Variables. II”. *Physical Review*, 85:180–193, January 1952. URL: <http://link.aps.org/doi/10.1103/PhysRev.85.180>, doi:10.1103/PhysRev.85.180.
- [61] David Bohm and Yakir Aharonov. “Discussion of Experimental Proof for the Paradox of Einstein, Rosen, and Podolsky”. *Physical Review*, 108(4):1070, 1957. URL: <http://link.aps.org/doi/10.1103/PhysRev.108.1070>, doi:10.1103/PhysRev.108.1070.
- [62] David Bohm and Basil J. Hiley. “Measurement Understood Through the Quantum Potential Approach”. *Foundations of Physics*, 14(3):255–274, 1984. URL: <http://dx.doi.org/10.1007/BF00730211>, doi:10.1007/BF00730211.
- [63] David Bohm and Basil J. Hiley. *The Undivided Universe: An Ontological Interpretation of Quantum Theory*. Routledge, 1993.
- [64] Jorge Luis Borges. *El jardín de senderos que se bifurcan (The Garden of Forking Paths)*. Editorial Sur, 1941. “I paused, as you may well imagine, at the sentence ‘I leave to several futures (not to all) my garden of forking paths.’ Almost instantly, I saw it—the garden of forking paths was the chaotic novel; the phrase ‘several futures (not all)’ suggested to me the image of a forking in *time*, rather than in space. A full rereading of the book confirmed my theory. In all fictions, each time a man meets diverse alternatives, he chooses one and eliminates the others; in the work of the virtually impossible-to-disentangle Ts’ui Pen, the character chooses—simultaneously—all of them. *He creates*, thereby, ‘several futures,’ several *times*, which themselves proliferate and fork. [...] In Ts’ui Pen’s novel, *all* the outcomes in fact occur; each is the starting point for further bifurcations. [...] Unlike Newton and Schopenhauer, [Ts’ui Pen] did not believe in a uniform and absolute time; he believed in an infinite series of times, a growing, dizzying web of divergent, convergent, and parallel times. The fabric of times that approach one another, fork, are snipped off, or are simply unknown for centuries, contains *all* possibilities. In most of those, we do not exist; in some, you exist but I do not; in others, I do and you do not; in others still, we both do”.
- [65] Raphael Bousso and Leonard Susskind. “Multiverse interpretation of quantum mechanics”. *Physical Review D*, 85(4):045007, February 2012. URL: <http://link.aps.org/doi/10.1103/PhysRevD.85.045007>, [arXiv:1105.3796](https://arxiv.org/abs/1105.3796), doi:10.1103/PhysRevD.85.045007.
- [66] Heinz-Peter Breuer, Elsi-Mari Laine, and Jyrki Piilo. “Measure for the Degree of Non-Markovian Behavior of Quantum Processes in Open Systems”. *Physical Review Letters*, 103:210401, November 2009. URL: <http://link.aps.org/doi/10.1103/PhysRevLett.103.210401>, [arXiv:0908.0238](https://arxiv.org/abs/0908.0238), doi:10.1103/PhysRevLett.103.210401.
- [67] Heinz Peter Breuer and Francesco Petruccione. *The Theory of Open Quantum Systems*. Oxford university press, 2002.
- [68] Harvey R. Brown and David Wallace. “Solving the measurement problem: de Broglie-Bohm loses out to Everett”. *Foundations of Physics*, 35, 2005. URL: <http://dx.doi.org/10.1007/s10701-004-2009-3>, [arXiv:quant-ph/0403094](https://arxiv.org/abs/quant-ph/0403094),

[doi:10.1007/s10701-004-2009-3](https://doi.org/10.1007/s10701-004-2009-3).

- [69] J. D. Brown and Marc Henneaux. “Central Charges in the Canonical Realization of Asymptotic Symmetries: An Example from Three Dimensional Gravity”. *Communications in Mathematical Physics*, 104(2):207–226, June 1986. URL: <http://dx.doi.org/10.1007/BF01211590>, doi:10.1007/BF01211590.
- [70] Jeffrey Bub. “Quantum mechanics without the projection postulate”. *Foundations of Physics*, 22(5):737–754, 1992. URL: <http://link.springer.com/article/10.1007/BF01889676>.
- [71] Roman V. Buniy, Stephen D. H. Hsu, and Anthony Zee. “Is Hilbert space discrete?”. *Physics Letters B*, 630(1):68–72, 2005. URL: <http://www.sciencedirect.com/science/article/pii/S0370269305013286>, arXiv:hep-th/0508039, doi:10.1016/j.physletb.2005.09.084.
- [72] Roman V. Buniy, Stephen D. H. Hsu, and Anthony Zee. “Discreteness and the origin of probability in quantum mechanics”. *Physics Letters B*, 640(4):219–223, 2006. URL: <http://www.sciencedirect.com/science/article/pii/S0370269306009373>, arXiv:hep-th/0606062, doi:10.1016/j.physletb.2006.07.050.
- [73] Francesco Buscemi. “Complete positivity in the presence of initial correlations: necessary and sufficient conditions given in terms of the quantum data-processing inequality”. 2013. arXiv:1307.0363.
- [74] Howard J. Carmichael. *An Open Systems Approach to Quantum Optics: Lectures Presented at the Universite Libre De Bruxelles, October 28 to November 4, 1991*. Springer, 1st edition, 1993. Section 7.4 (“Unravelling the master equation for the source”) and Section 7.5 (“Stochastic wavefunctions”).
- [75] Nancy D. Cartwright. “Van Fraassen’s Modal Model of Quantum Mechanics”. *Philosophy of Science*, 41(2):199–202, 1974. URL: <http://www.jstor.org/stable/186870>.
- [76] André R. R. Carvalho, Marc Busse, Olivier Brodier, Carlos Viviescas, and Andreas Buchleitner. “Unravelling Entanglement”. 2005. arXiv:quant-ph/0510006.
- [77] Carlton M. Caves, Christopher A. Fuchs, and Rüdiger Schack. “Quantum probabilities as Bayesian probabilities”. *Physical Review A*, 65(2):022305, 2002. URL: <http://link.aps.org/doi/10.1103/PhysRevA.65.022305>, arXiv:quant-ph/0106133, doi:10.1103/PhysRevA.65.022305.
- [78] Carlton M. Caves, Christopher A. Fuchs, and Rüdiger Schack. “Subjective probability and quantum certainty”. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 38(2):255–274, 2007. URL: <http://www.sciencedirect.com/science/article/pii/S1355219807000354>, arXiv:quant-ph/0608190, doi:10.1016/j.shpsb.2006.10.007.
- [79] Carlton M. Caves and Rüdiger Schack. “Properties of the frequency operator do not imply the quantum probability postulate”. *Annals of Physics*, 315(1):123–146, January 2005. Special issue. URL: <http://www.sciencedirect.com/science/article/pii/S0003491604001794>, doi:10.1016/j.aop.2004.09.009.
- [80] C. B. Chiu, E. C. G. Sudarshan, and B. Misra. “Time evolution of unstable quantum states and a resolution of Zeno’s paradox”. *Physical Review D*, 16:520–529, July 1977. URL: <http://link.aps.org/doi/10.1103/PhysRevD.16.520>, doi:10.1103/PhysRevD.16.520.
- [81] Man-Duen Choi. “Positive linear maps on C*-algebras”. *Canadian Journal of Mathematics*, 24(3):520–529, 1972. URL: <http://cms.math.ca/10.4153/CJM-1972-044-5>, doi:10.4153/CJM-1972-044-5.
- [82] Man-Duen Choi. “Completely Positive Linear Maps on Complex Matrices”. *Linear Algebra and its Applications*, 10(3):285–290, 1975. URL: <http://www.sciencedirect.com/science/article/pii/0024379575900750>, doi:10.1016/0024-3795(75)90075-0.
- [83] John F. Clauser and Michael A. Horne. “Experimental consequences of objective local theories”. *Phys. Rev. D*, 10:526–535, July 1974. URL: <http://link.aps.org/doi/10.1103/PhysRevD.10.526>, doi:10.1103/PhysRevD.10.526.
- [84] John F. Clauser, Michael A. Horne, Abner Shimony, and Richard A. Holt. “Proposed Experiment to Test Local Hidden-Variable Theories”. *Physical Review Letters*, 23:880–884, October 1969. URL: <http://link.aps.org/doi/10.1103/PhysRevLett.23.880>, doi:10.1103/PhysRevLett.23.880.
- [85] Rob Clifton. “Why Modal Interpretations of Quantum Mechanics Must Abandon Classical Reasoning About Physical Properties”. *International Journal of Theoretical Physics*, 34(8):1303–1312, 1995. URL: <http://link.springer.com/article/10.1007/BF00676242>.
- [86] Nino B. Cocchiarella and Max A. Freund. *Modal Logic: An Introduction to its Syntax and Semantics*. Oxford University Press, 2008.
- [87] Roger Colbeck and Renato Renner. “A system’s wave function is uniquely determined by its underlying physical state”. 2013. arXiv:1312.7353.
- [88] Sidney Coleman. “Quantum sine-Gordon equation as the massive Thirring model”. *Physical Review D*, 11(8):2088–2097, April 1975. URL: <http://link.aps.org/doi/10.1103/PhysRevD.11.2088>, doi:10.1103/PhysRevD.11.2088.
- [89] Sidney Coleman. “Quantum Mechanics in Your Face”, 1995. Lecture video.
- [90] Jean Dalibard, Yvan Castin, and Klaus Mølmer. “Wave-function approach to dissipative processes in quantum optics”. *Physical Review Letters*, 68:580–583, February 1992. URL: <http://link.aps.org/doi/10.1103/PhysRevLett.68.580>, doi:10.1103/PhysRevLett.68.580.
- [91] Louis de Broglie. *An introduction to the Study of Wave Mechanics*. E. P. Dutton and Company, Inc., 1930.
- [92] A. Degasperis, L. Fonda, and G.C. Ghirardi. “Does the Lifetime of an Unstable System Depend on the Measuring Apparatus?”. *Il Nuovo Cimento A Series 11*, 21(3):471–484, 1974. URL: <http://link.springer.com/article/10.1007/BF02731351>, doi:10.1007/BF02731351.
- [93] David Deutsch. “Quantum theory as a universal physical theory”. *International Journal of Theoretical Physics*, 24(1):1–41, 1985.
- [94] David Deutsch. “Quantum Theory of Probability and Decisions”. 1999. arXiv:quant-ph/9906015.

- [95] Bryce S. DeWitt. “Quantum Theory of Gravity. I. The Canonical Theory”. *Physical Review*, 160:1113–1148, August 1967. URL: <http://link.aps.org/doi/10.1103/PhysRev.160.1113>, doi:10.1103/PhysRev.160.1113.
- [96] Bryce S. DeWitt. “Quantum mechanics and reality”. *Physics Today*, 23(9):30–35, September 1970. URL: <http://scitation.aip.org/content/aip/magazine/physicstoday/article/23/9/10.1063/1.3022331>, doi:10.1063/1.3022331.
- [97] W. Michael Dickson. “Faux-Boolean algebras and classical models”. *Foundations of Physics Letters*, 8(5):401–415, 1995. URL: <http://link.springer.com/article/10.1007/BF02186577>.
- [98] W. Michael Dickson. “Faux-boolean algebras, classical probability, and determinism”. *Foundations of Physics Letters*, 8(3):231–242, 1995. URL: <http://link.springer.com/article/10.1007/BF02187347>.
- [99] W. Michael Dickson and Rob Clifton. “Lorentz-invariance in modal interpretations”. In *“The Modal Interpretation of Quantum Mechanics”*, pages 9–47. Springer, 1998. URL: http://link.springer.com/chapter/10.1007/978-94-011-5084-2_2.
- [100] Lajos Diósi, Nicolas Gisin, and Walter T. Strunz. “Non-Markovian quantum state diffusion”. *Physical Review A*, 58(3):1699–1712, September 1998. URL: <http://link.aps.org/doi/10.1103/PhysRevA.58.1699>, arXiv:quant-ph/9803062, doi:10.1103/PhysRevA.58.1699.
- [101] Lajos Diósi and Walter T. Strunz. “The non-Markovian stochastic Schrödinger equation for open systems”. *Physics Letters A*, 235:569–573, 1997. URL: <http://www.sciencedirect.com/science/article/pii/S0375960197007172>, arXiv:quant-ph/9706050, doi:10.1016/S0375-9601(97)00717-2.
- [102] Paul Adrien Maurice Dirac. *The Principles of Quantum Mechanics*. Oxford University Press, 1st edition, 1930.
- [103] Matthew J. Donald. “Discontinuity and continuity of definite properties in the modal interpretation”. In *“The Modal Interpretation of Quantum Mechanics”*, pages 213–222. Springer, 1998. URL: http://link.springer.com/chapter/10.1007/978-94-011-5084-2_8.
- [104] Detlef Dürr, Sheldon Goldstein, Roderich Tumulka, and Nino Zanghì. “Bell-type quantum field theories”. *Journal of Physics A: Mathematical and General*, 38(4):R1, 2005. URL: <http://stacks.iop.org/0305-4470/38/i=4/a=R01>, arXiv:quant-ph/0407116, doi:10.1088/0305-4470/38/4/R01.
- [105] Detlef Dürr, Sheldon Goldstein, Roderich Tumulka, and Nino Zanghì. “Quantum Hamiltonians and Stochastic Jumps”. *Communications in Mathematical Physics*, 254(1):129–166, February 2005. URL: <http://dx.doi.org/10.1007/s00220-004-1242-0>, arXiv:quant-ph/0303056, doi:10.1007/s00220-004-1242-0.
- [106] John Earman and Laura Ruetsche. “Relativistic Invariance and Modal Interpretations”. *“Philosophy of Science”*, 72(4):557–584, 2005. URL: <http://www.jstor.org/stable/10.1086/505448>.
- [107] Albert Einstein, Boris Podolsky, and Nathan Rosen. “Can Quantum-Mechanical Description of Physical Reality Be Considered Complete?”. *Physical Review*, 47:777–780, May 1935. URL: <http://link.aps.org/doi/10.1103/PhysRev.47.777>, doi:10.1103/PhysRev.47.777.
- [108] Sheer El-Showk and Kyriakos Papadodimas. “Emergent spacetime and holographic CFTs”. *Journal of High Energy Physics*, 2012(10):1–72, 2012. URL: <http://dx.doi.org/10.1007/JHEP10%282012%29106>, arXiv:1101.4163, doi:10.1007/JHEP10(2012)106.
- [109] H.-Th. Elze. “Relativistic quantum transport theory”. 2002. arXiv:hep-ph/0204309.
- [110] Massimiliano Esposito and Shaul Mukamel. “Fluctuation theorems for quantum master equations”. *Physical Review E*, 73(4):046129, April 2006. URL: <http://link.aps.org/doi/10.1103/PhysRevE.73.046129>, arXiv:cond-mat/0602679, doi:10.1103/PhysRevE.73.046129.
- [111] Hugh Everett III. ““Relative state” formulation of quantum mechanics”. *Reviews of Modern Physics*, 29(3):454–462, July 1957. URL: <http://link.aps.org/doi/10.1103/RevModPhys.29.454>, doi:10.1103/RevModPhys.29.454.
- [112] Hugh Everett III. “The theory of the universal wave function”. In *“The many-worlds interpretation of quantum mechanics”*, volume 1, page 3, 1973.
- [113] Hugh Everett III, Bryce Seligman DeWitt, and Neill Graham. *The many-worlds interpretation of quantum mechanics*. Princeton University Press, 1973.
- [114] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. “How Probability Arises in Quantum Mechanics”. *Annals of Physics*, 192(2):368–382, June 1989. URL: <http://www.sciencedirect.com/science/article/pii/0003491689901413>, doi:10.1016/0003-4916(89)90141-3.
- [115] Richard P. Feynman. *The Character of Physical Law*. British Broadcasting Corporation, 1965. From Feynman’s 1964 Messenger Lectures delivered at Cornell University.
- [116] Richard P. Feynman, Fernando Bernardino Morinigo, William G. Wagner, and Brian Hatfield. *Feynman Lectures on Gravitation*. Westview Press, 2002. Section 1.4 (“On the Philosophical Problems in Quantizing Macroscopic Objects”).
- [117] Richard P. Feynman and Steven Weinberg. *Elementary Particles and the Laws of Physics*. Cambridge University Press, 1987. In particular, the lecture “The Reason for Antiparticles” by Richard Feynman.
- [118] Arthur Fine. “Hidden Variables, Joint Probability, and the Bell Inequalities”. *Physical Review Letters*, 48:291–295, February 1982. URL: <http://link.aps.org/doi/10.1103/PhysRevLett.48.291>, doi:10.1103/PhysRevLett.48.291.
- [119] Arthur Fine. “Joint distributions, quantum correlations, and commuting observables”. *Journal of Mathematical Physics*, 23(7):1306, 1982. URL: http://jmp.aip.org/resource/1/jmapaq/v23/i7/p1306_s1, doi:10.1063/1.525514.
- [120] Stuart J. Freedman and John F. Clauser. “Experimental Test of Local Hidden-Variable Theories”. *Physical Review Letters*, 28:938–941, April 1972. URL: <http://link.aps.org/doi/10.1103/PhysRevLetters28.938>, doi:10.1103/PhysRevLetters28.938.
- [121] Ben Freivogel. “Making predictions in the multiverse”. *Classical and Quantum Gravity*, 28(20):204007, 2011. URL: <http://stacks.iop.org/0264-9381/28/i=20/a=204007>, arXiv:1105.0244, doi:10.1088/0264-9381/28/20/204007.

- [122] Jonathan R. Friedman, Vijay Patel, Wei Chen, S. K. Tolpygo, and James E. Lukens. “Quantum superposition of distinct macroscopic states”. *Nature*, 406(6791):43–46, 2000. URL: <http://www.nature.com/nature/journal/v406/n6791/abs/406043a0.html>, doi:10.1038/35017505.
- [123] Michael Friedman and Hilary Putnam. “Quantum Logic, Conditional Probability, and Interference”. *Dialectica*, 32(3/4):305–315, 1978. URL: <http://www.jstor.org/stable/42970322>, doi:10.1038/35017505.
- [124] Christopher A. Fuchs. “Quantum Mechanics as Quantum Information (and only a little more)”. 2002. [arXiv:quant-ph/0205039](https://arxiv.org/abs/quant-ph/0205039).
- [125] Christopher A. Fuchs. “QBism, the Perimeter of Quantum Bayesianism”. 2010. [arXiv:1003.5209](https://arxiv.org/abs/1003.5209).
- [126] Christopher A. Fuchs. “Quantum Bayesianism at the Perimeter”. 2010. [arXiv:1003.5182](https://arxiv.org/abs/1003.5182).
- [127] Christopher A. Fuchs. “Interview with a Quantum Bayesian”. 2012. [arXiv:1207.2141](https://arxiv.org/abs/1207.2141).
- [128] Christopher A. Fuchs and Rüdiger Schack. “Quantum-Bayesian Coherence”. 2009. [arXiv:0906.2187](https://arxiv.org/abs/0906.2187).
- [129] Christopher A. Fuchs and Rüdiger Schack. “A Quantum-Bayesian Route to Quantum-State Space”. *Foundations of Physics*, 41(3):345–356, 2011. URL: <http://dx.doi.org/10.1007/s10701-009-9404-8>, [arXiv:0912.4252](https://arxiv.org/abs/0912.4252), doi:10.1007/s10701-009-9404-8.
- [130] Christopher A. Fuchs and Rüdiger Schack. “Bayesian Conditioning, the Reflection Principle, and Quantum Decoherence”. In *Probability in Physics*, pages 233–247. Springer, 2012. URL: http://dx.doi.org/10.1007/978-3-642-21329-8_15, [arXiv:1103.5950](https://arxiv.org/abs/1103.5950), doi:10.1007/978-3-642-21329-8_15.
- [131] Christopher A. Fuchs and Rüdiger Schack. “Quantum-Bayesian Coherence: The No-Nonsense Version”. 2013. [arXiv:1301.3274](https://arxiv.org/abs/1301.3274).
- [132] Jason Gallicchio, Andrew S. Friedman, and David I. Kaiser. “Testing Bell’s Inequality with Cosmic Photons: Closing the Settings-Independence Loophole”. 2013. [arXiv:1310.3288](https://arxiv.org/abs/1310.3288).
- [133] Jay Gambetta, Tomas Gustav Askerud, and Howard M. Wiseman. “Jump-like unravelings for non-Markovian open quantum systems”. 2004. [arXiv:quant-ph/0401171](https://arxiv.org/abs/quant-ph/0401171).
- [134] Anupam Garg and N. D. Mermin. “Detector inefficiencies in the Einstein-Podolsky-Rosen experiment”. *Phys. Rev. D*, 35:3831–3835, Jun 1987. URL: <http://link.aps.org/doi/10.1103/PhysRevD.35.3831>, doi:10.1103/PhysRevD.35.3831.
- [135] Jaume Garriga, Alan H. Guth, and Alexander Vilenkin. “Eternal inflation, bubble collisions, and the persistence of memory”. *Physical Review D*, 76(12):123512, 2007. URL: <http://link.aps.org/doi/10.1103/PhysRevD.76.123512>, [arXiv:hep-th/0612242](https://arxiv.org/abs/hep-th/0612242), doi:10.1103/PhysRevD.76.123512.
- [136] Jaume Garriga, Delia Schwartz-Perlov, Alexander Vilenkin, and Sergei Winitzki. “Probabilities in the inflationary multiverse”. *Journal of Cosmology and Astroparticle Physics*, 2006(01):017, 2006. URL: <http://stacks.iop.org/1475-7516/2006/i=01/a=017>, [arXiv:hep-th/0509184](https://arxiv.org/abs/hep-th/0509184), doi:10.1088/1475-7516/2006/01/017.
- [137] Jaume Garriga and Alexander Vilenkin. “Recycling universe”. *Physical Review D*, 57:2230–2244, Feb 1998. URL: <http://link.aps.org/doi/10.1103/PhysRevD.57.2230>, [arXiv:astro-ph/9707292](https://arxiv.org/abs/astro-ph/9707292), doi:10.1103/PhysRevD.57.2230.
- [138] Jaume Garriga and Alexander Vilenkin. “Many worlds in one”. *Physical Review D*, 64:043511, July 2001. URL: <http://link.aps.org/doi/10.1103/PhysRevD.64.043511>, [arXiv:gr-qc/0102010](https://arxiv.org/abs/gr-qc/0102010), doi:10.1103/PhysRevD.64.043511.
- [139] Jaume Garriga and Alexander Vilenkin. “Watchers of the multiverse”. *Journal of Cosmology and Astroparticle Physics*, 2013(05):037, 2013. URL: <http://stacks.iop.org/1475-7516/2013/i=05/a=037>, [arXiv:1210.7540](https://arxiv.org/abs/1210.7540), doi:10.1088/1475-7516/2013/05/037.
- [140] Murray Gell-Mann and James B. Hartle. “Classical equations for quantum systems”. *Physical Review D*, 47(8):3345, 1993. URL: <http://link.aps.org/doi/10.1103/PhysRevD.47.3345>, [arXiv:gr-qc/9210010](https://arxiv.org/abs/gr-qc/9210010), doi:10.1103/PhysRevD.47.3345.
- [141] Murray Gell-Mann and James B. Hartle. “Adaptive Coarse Graining, Environment, Strong Decoherence, and Quasiclassical Realms”. 2013. [arXiv:1312.7454](https://arxiv.org/abs/1312.7454).
- [142] Ilja Gerhardt, Qin Liu, Antía Lamas-Linares, Johannes Skaar, Valerio Scarani, Vadim Makarov, and Christian Kurtsiefer. “Experimentally Faking the Violation of Bell’s Inequalities”. 2011. [arXiv:1106.3224gr-qc/9210010](https://arxiv.org/abs/1106.3224gr-qc/9210010).
- [143] Stefan Gerlich, Sandra Eibenberger, Mathias Tomandl, Stefan Nimmrichter, Klaus Hornberger, Paul J. Fagan, Jens Tüxen, Marcel Mayor, and Markus Arndt. “Quantum interference of large organic molecules”. *Nature Communications*, 2:263, 2011. URL: http://www.nature.com/ncomms/journal/v2/n4/supinfo/ncomms1263_S1.html, doi:10.1038/ncomms1263.
- [144] Gian Carlo Ghirardi, Alberto Rimini, and Tullio Weber. “Unified dynamics for microscopic and macroscopic systems”. *Physical Review D*, 34(2):470–491, 1986. URL: <http://link.aps.org/doi/10.1103/PhysRevD.34.470>, doi:10.1103/PhysRevD.34.470.
- [145] Gian Carlo Ghirardi, C. Omero, Tullio Weber, and Alberto Rimini. “Small-Time Behaviour of Quantum Nondecay Probability and Zeno’s Paradox in Quantum Mechanics”. *Il Nuovo Cimento A*, 52(4):421–442, 1979. URL: <http://link.springer.com/article/10.1007/BF02770851>, doi:10.1007/BF02770851.
- [146] Peter Gibbins. *Particles and Paradoxes: The Limits of Quantum Logic*. Cambridge University Press, 1987.
- [147] Steven B. Giddings. “The Black Hole Information Paradox”. 1995. [arXiv:hep-th/9508151](https://arxiv.org/abs/hep-th/9508151).
- [148] Steven B. Giddings. “Black hole information, unitarity, and nonlocality”. *Physical Review D*, 74:106005, November 2006. URL: <http://link.aps.org/doi/10.1103/PhysRevD.74.106005>, [arXiv:hep-th/0605196](https://arxiv.org/abs/hep-th/0605196), doi:10.1103/PhysRevD.74.106005.
- [149] Nicolas Gisin and Ian C. Percival. “The quantum-state diffusion model applied to open systems”. *Journal of Physics A: Mathematical and Theoretical*, 25(21):5677, 1992. URL: <http://stacks.iop.org/0305-4470/25/i=21/a=023>, doi:10.1088/0305-4470/25/21/023.

- [150] Marissa Giustina, Alexandra Mech, Sven Ramelow, Bernhard Wittmann, Johannes Kofler, Jorn Beyer, Adriana Lita, Brice Calkins, Thomas Gerrits, Sae Woo Nam, Rupert Ursin, and Anton Zeilinger. “Bell violation using entangled photons without the fair-sampling assumption”. *Nature*, 497:227–230, May 2013. URL: <http://www.nature.com/nature/journal/v497/n7448/full/nature12012.html>, doi:10.1038/nature12012.
- [151] Amit Giveon, Massimo Porrati, and Eliezer Rabinovici. “Target Space Duality in String Theory”. *Physics Reports*, 244(2–3):77–202, 1994. URL: <http://www.sciencedirect.com/science/article/pii/0370157394900701>, arXiv:hep-th/9401139, doi:10.1016/0370-1573(94)90070-1.
- [152] Roy J. Glauber. “Coherent and Incoherent States of the Radiation Field”. *Physical Review*, 131:2766–2788, September 1963. URL: <http://link.aps.org/doi/10.1103/PhysRev.131.2766>, doi:10.1103/PhysRev.131.2766.
- [153] Andrew M. Gleason. “Measures on the Closed Subspaces of a Hilbert Space”. *Journal of Mathematics and Mechanics*, 6(6):885–893, 1957.
- [154] Sheldon Goldstein, Joel L. Lebowitz, Roderich Tumulka, and Nino Zanghì. “Canonical Typicality”. *Physical Review Letters*, 96(5):050403, 2006. URL: <http://link.aps.org/doi/10.1103/PhysRevLett.96.050403>, arXiv:cond-mat/0511091, doi:10.1103/PhysRevLett.96.050403.
- [155] Jeffrey Goldstone. “Field theories with Superconductor solutions”. *Il Nuovo Cimento*, 19(1):154–164, 1961. URL: <http://link.springer.com/article/10.1007/BF02812722>, doi:10.1007/BF02812722.
- [156] Jeffrey Goldstone, Abdus Salam, and Steven Weinberg. “Broken Symmetries”. *Physical Review*, 127:965–970, August 1962. URL: <http://link.aps.org/doi/10.1103/PhysRev.127.965>, doi:10.1103/PhysRev.127.965.
- [157] Daniel M. Greenberger, Michael A. Horne, and Anton Zeilinger. “Going Beyond Bell’s Theorem”. In *Bell’s Theorem, Quantum Theory and Conceptions of the Universe*, Fundamental Theories of Physics, pages 69–72. Springer, 1989. URL: http://dx.doi.org/10.1007/978-94-017-0849-4_10, arXiv:0712.0921, doi:10.1007/978-94-017-0849-4_10.
- [158] David J. Griffiths. *Introduction to Quantum Mechanics*. Pearson, 2nd edition, 2005. With treatments of Schrödinger’s cat in Section 12.4 (“Schrödinger’s Cat”) and the quantum Zeno paradox in Section 12.5 (“The Quantum Zeno Paradox”).
- [159] Simon Groblacher, Tomasz Paterek, Rainer Kaltenbaek, Caslav Brukner, Marek Zukowski, Markus Aspelmeyer, and Anton Zeilinger. “An experimental test of non-local realism”. *Nature*, 446:871–875, April 2007. URL: <http://www.nature.com/nature/journal/v446/n7138/full/nature05677.html>, arXiv:0704.2529, doi:10.1038/nature05677.
- [160] Steven S. Gubser, Igor R. Klebanov, and Alexander M. Polyakov. “Gauge theory correlators from non-critical string theory”. *Physics Letters B*, 428(1–2):105–114, 1998. URL: <http://www.sciencedirect.com/science/article/pii/S0370269398003773>, arXiv:hep-th/9802109, doi:10.1016/S0370-2693(98)00377-3.
- [161] Monica Guica, Thomas Hartman, Wei Song, and Andrew Strominger. “The Kerr/CFT correspondence”. *Physical Review D*, 80(12):124008, 2009. URL: <http://link.aps.org/doi/10.1103/PhysRevD.80.124008>, arXiv:0809.4266, doi:10.1103/PhysRevD.80.124008.
- [162] Alan H. Guth. “Inflation and eternal inflation”. *Physics Reports*, 333–334(0):555–574, 2000. URL: <http://www.sciencedirect.com/science/article/pii/S0370157300000375>, doi:10.1016/S0370-1573(00)00037-5.
- [163] Alan H. Guth. “Eternal inflation and its implications”. *Journal of Physics A: Mathematical and Theoretical*, 40(25):6811, 2007. URL: <http://stacks.iop.org/1751-8121/40/i=25/a=S25>, arXiv:hep-th/0702178, doi:10.1088/1751-8113/40/25/S25.
- [164] Alan H. Guth and David I. Kaiser. “Inflationary Cosmology: Exploring the Universe from the Smallest to the Largest Scales”. *Science*, 307:884–890, 2005. URL: <http://www.sciencemag.org/content/307/5711/884.abstract>, arXiv:astro-ph/0502328, doi:10.1126/science.1107483.
- [165] Lucia Hackermüller, Klaus Hornberger, Björn Brezger, Anton Zeilinger, and Markus Arndt. “Decoherence of matter waves by thermal emission of radiation”. *Nature*, 427(6976):711–714, 2004. URL: <http://www.nature.com/nature/journal/v427/n6976/abs/nature02276.html>, doi:10.1038/nature02276.
- [166] Michael J. W. Hall. “Imprecise measurements and non-locality in quantum mechanics”. *Physics Letters A*, 125:89–91, 1987. URL: <http://www.sciencedirect.com/science/article/pii/0375960187901277>, doi:10.1016/0375-9601(87)90127-7.
- [167] Jonathan Halliwell and Andreas Zoupas. “Quantum state diffusion, density matrix diagonalization, and decoherent histories: A model”. *Physical Review D*, 52(12):7294–7307, December 1995. URL: <http://link.aps.org/doi/10.1103/PhysRevD.52.7294>, arXiv:quant-ph/9503008, doi:10.1103/PhysRevD.52.7294.
- [168] Lucien Hardy. “Quantum Mechanics, Local Realistic Theories, and Lorentz-Invariant Realistic Theories”. *Physical Review Letters*, 68:2981–2984, May 1992. URL: <http://link.aps.org/doi/10.1103/PhysRevLett.68.2981>, doi:10.1103/PhysRevLett.68.2981.
- [169] Lucien Hardy. “Nonlocality for Two Particles without Inequalities for Almost All Entangled States”. *Physical Review Letters*, 71:1665–1668, September 1993. URL: <http://link.aps.org/doi/10.1103/PhysRevLett.71.1665>, doi:10.1103/PhysRevLett.71.1665.
- [170] Lucien Hardy. “Quantum theory from five reasonable axioms”. 2001. arXiv:quant-ph/0101012.
- [171] Lucien Hardy. Private communication, 2014.
- [172] Nicholas Harrigan and Robert W. Spekkens. “Einstein, Incompleteness, and the Epistemic View of Quantum States”. *Foundations of Physics*, 40(2):125–157, 2010. URL: <http://dx.doi.org/10.1007/s10701-009-9347-0>, arXiv:0706.2661, doi:10.1007/s10701-009-9347-0.
- [173] James B. Hartle. “Spacetime coarse grainings in nonrelativistic quantum mechanics”. *Physical Review D*, 44(10):3173, 1991. URL: <http://link.aps.org/doi/10.1103/PhysRevD.44.3173>, doi:10.1103/PhysRevD.44.3173.
- [174] James B. Hartle. “Spacetime quantum mechanics and the quantum mechanics of spacetime”. 1993. arXiv:gr-qc/9304006.

- [175] Stephen W. Hawking. “Particle Creation by Black Holes”. *Communications in Mathematical Physics*, 43(3):199–220, 1975. URL: <http://dx.doi.org/10.1007/BF02345020>, doi:10.1007/BF02345020.
- [176] Stephen W. Hawking. “Breakdown of predictability in gravitational collapse”. *Physical Review D*, 14:2460–2473, November 1976. URL: <http://link.aps.org/doi/10.1103/PhysRevD.14.2460>, doi:10.1103/PhysRevD.14.2460.
- [177] Leah Henderson and Vlatko Vedral. “Classical, quantum and total correlations”. *Journal of Physics A: Mathematical and Theoretical*, 34(35):6899, 2001. URL: <http://stacks.iop.org/0305-4470/34/i=35/a=315>, arXiv:quant-ph/0105028, doi:10.1088/0305-4470/34/35/315.
- [178] Bas Hensen et al. “Experimental loophole-free violation of a Bell inequality using entangled electron spins separated by 1.3 km”. 2015. arXiv:1508.05949.
- [179] Karl Hess and Walter Philipp. “Bell’s theorem: Critique of proofs with and without inequalities”. 2004. arXiv:quant-ph/0410015.
- [180] Art Hobson. “There are no particles, only fields”. *American Journal of Physics*, 81:211–223, March 2013. With the resulting debate in the September 2013 issue of American Journal of Physics. arXiv:1204.4616.
- [181] Timothy J. Hollowood. “New Modal Quantum Mechanics”. 2013. arXiv:arXiv:1312.4751.
- [182] Timothy J. Hollowood. “The Copenhagen Interpretation as an Emergent Phenomenon”. 2013. arXiv:1302.4228.
- [183] Timothy J. Hollowood. “The Emergent Copenhagen Interpretation of Quantum Mechanics”. 2013. arXiv:1312.3427.
- [184] Stephen D. H. Hsu. “On the origin of probability in quantum mechanics”. *Modern Physics Letters A*, 27(12), 2012. arXiv:1110.0549.
- [185] Lane P. Hughston, Richard Jozsa, and William K. Wootters. “A complete classification of quantum ensembles having a given density matrix”. *Physics Letters A*, 183:14–18, 1993. URL: <http://www.sciencedirect.com/science/article/pii/0375960193908809>, doi:10.1016/0375-9601(93)90880-9.
- [186] Norihiro Iizuka and Daniel Kabat. “On the Mutual Information in Hawking Radiation”. 2013. arXiv:1308.2386.
- [187] Wayne M. Itano, D. J. Heinzen, J. J. Bollinger, and D. J. Wineland. “Quantum Zeno effect”. *Physical Review A*, 41:2295–2300, March 1990. URL: <http://link.aps.org/doi/10.1103/PhysRevA.41.2295>, doi:10.1103/PhysRevA.41.2295.
- [188] John D. Jackson. *Classical Electrodynamics*. John Wiley and Sons, 3rd edition, 1998. Section 6.3 (“Gauge Transformations, Lorenz Gauge, Coulomb Gauge”).
- [189] Gregg Jaeger. *Entanglement, Information, and the Interpretation of Quantum Mechanics*. Springer-Verlag Berlin Heidelberg, 1st edition, 2009. Section 3.1 (“Interpretation and Metaphysics”), p. 116.
- [190] A. Jamiolkowski. “Linear transformations which preserve trace and positive semidefiniteness of operators”. *Reports on Mathematical Physics*, 3(4):275–278, 1972. URL: <http://www.sciencedirect.com/science/article/pii/0034487772900110>, doi:10.1016/0034-4877(72)90011-0.
- [191] Josef M. Jauch. “The quantum probability calculus”. *Synthese*, 29(1):131–154, 1974. URL: <http://link.springer.com/article/10.1007/2FBF00484955>, doi:10.1007/BF00484955.
- [192] Edwin Thompson Jaynes. “Information Theory and Statistical Mechanics”. *Physical Review*, 106:620–630, May 1957. URL: <http://link.aps.org/doi/10.1103/PhysRev.106.620>, doi:10.1103/PhysRev.106.620.
- [193] Edwin Thompson Jaynes. “Information Theory and Statistical Mechanics. II”. *Physical Review*, 108:171–190, October 1957. URL: <http://link.aps.org/doi/10.1103/PhysRev.108.171>, doi:10.1103/PhysRev.108.171.
- [194] Edwin Thompson Jaynes. “Gibbs vs Boltzmann Entropies”. *American Journal of Physics*, 33:391–398, 1965. URL: <http://scitation.aip.org/content/aapt/journal/ajp/33/5/10.1119/1.1971557>, doi:10.1119/1.1971557.
- [195] Erich Joos. “Comment on ‘Unified dynamics for microscopic and macroscopic systems’”. *Physical Review D*, 36(10):3285–3286, November 1987. URL: <http://link.aps.org/doi/10.1103/PhysRevD.36.3285>, doi:10.1103/PhysRevD.36.3285.
- [196] Erich Joos, editor. *Decoherence and the appearance of a classical world in quantum theory*. Springer, 2003.
- [197] Erich Joos and H. Dieter Zeh. “The emergence of classical properties through interaction with the environment”. *Zeitschrift für Physik B Condensed Matter*, 59(2):223–243, 1985. URL: <http://link.springer.com/article/10.1007/BF01725541>.
- [198] Thomas F. Jordan and E. C. G. Sudarshan. “Dynamical Mappings of Density Operators in Quantum Mechanics”. *Journal of Mathematical Physics*, 2(6):772–775, 1961. URL: <http://scitation.aip.org/content/aip/journal/jmp/2/6/10.1063/1.1724221>, doi:10.1063/1.1724221.
- [199] Tatsuhiro Kawamoto and Naomichi Hatano. “Test of fluctuation theorems in non-Markovian open quantum systems”. *Physical Review E*, 84(3):031116, September 2011. URL: <http://link.aps.org/doi/10.1103/PhysRevE.84.031116>, arXiv:1105.3579, doi:10.1103/PhysRevE.84.031116.
- [200] Andrei Khrennikov. “Violation of Bell’s inequality and postulate on simultaneous measurement of compatible observables”. *Journal of Computational and Theoretical Nanoscience*, 8(6):1006–1010, 2011. URL: <http://www.ingentaconnect.com/content/asp/jctn/2011/00000008/00000006/art00012>, arXiv:1102.4743, doi:10.1166/jctn.2011.1780.
- [201] Frank Knight. *Risk, Uncertainty, and Profit*. Hart, Schaffner and Marx. Houghton Mifflin Company, 1921.
- [202] Simon Kochen and E. Specker. “The Problem of Hidden Variables in Quantum Mechanics”. *Indiana University Mathematics Journal*, 17:59–87, 1968.
- [203] Andrey N. Kolmogorov. *Grundbegriffe der Wahrscheinlichkeitsrechnung, Ergebnisse Der Mathematik (Foundations of the Theory of Probability)*. 1993.
- [204] Karl Kraus. “General State Changes in Quantum Theory”. *Annals of Physics*, 64(2):311–335, June 1971. URL: <http://www.sciencedirect.com/science/article/pii/0003491671901084>, doi:10.1016/0003-4916(71)90108-4.
- [205] Henry P. Krips. “Two Paradoxes in Quantum Mechanics”. *Philosophy of Science*, 36(2):145–152, June 1969. Where we note that although van Fraassen is usually credited with the idea of the modal interpretations, and Vermaas and Dieks for

- its spectral version, Krips appears to have arrived at these ideas first. URL: <http://www.jstor.org/stable/186167>.
- [206] Henry P. Krips. “Statistical interpretation of quantum theory”. *American Journal of Physics*, 43(5):420–422, 1975. URL: <http://dx.doi.org/10.1119/1.9804>, doi:10.1119/1.9804.
- [207] Henry P. Krips. *The Metaphysics of Quantum Theory*. Clarendon Press Oxford, 1987.
- [208] Henry P. Krips. “A Propensity Interpretation for Quantum Probabilities”. *The Philosophical Quarterly*, 39(156):308–333, July 1989. URL: <http://www.jstor.org/stable/2220174>.
- [209] Elsi-Mari Laine, Jyrki Piilo, and Heinz-Peter Breuer. “Measurement for the non-Markovianity of quantum processes”. *Physical Review A*, 81(6):062115, June 2010. URL: <http://link.aps.org/doi/10.1103/PhysRevA.81.062115>, arXiv:1002.2583, doi:10.1103/PhysRevA.81.062115.
- [210] Lev D. Landau. “The damping problem in wave mechanics”. *Z. Phys*, 45:430, 1927.
- [211] Nicolaas P. Landsman. “Between classical and quantum”. In *“Handbook of the Philosophy of Science”*, pages 417–553. Elsevier, 2007. With a review of the Heisenberg cut and its history in Section 3.2 (“Object and apparatus: the Heisenberg cut”). arXiv:quant-ph/0506082.
- [212] Jan-Åke Larsson. “Bell’s inequality and detector inefficiency”. *Phys. Rev. A*, 57:3304–3308, May 1998. URL: <http://link.aps.org/doi/10.1103/PhysRevA.57.3304>, doi:10.1103/PhysRevA.57.3304.
- [213] Anthony J. Leggett. “Nonlocal Hidden-Variable Theories and Quantum Mechanics: An Incompatibility Theorem”. *Foundations of Physics*, 33(10):1469–1493, 2003. URL: <http://link.springer.com/article/10.1023/A:1026096313729>, doi:10.1023/A:1026096313729.
- [214] Bruno Leggio, Anna Napoli, Heinz-Peter Breuer, and Antonino Messina. “Fluctuation theorems for non-Markovian quantum processes”. *Physical Review E*, 87(3):032113, March 2013. URL: <http://link.aps.org/doi/10.1103/PhysRevE.87.032113>, arXiv:1301.1553, doi:10.1103/PhysRevE.87.032113.
- [215] Matthew S. Leifer and Robert W. Spekkens. “Towards a formulation of quantum theory as a causally neutral theory of Bayesian inference”. *Physical Review A*, 88(5):052130, November 2013. URL: <http://link.aps.org/doi/10.1103/PhysRevA.88.052130>, arXiv:1107.5849, doi:10.1103/PhysRevA.88.052130.
- [216] Clarence I. Lewis and Cooper H. Langford. *Symbolic Logic*. New York: Century Company, 1932.
- [217] David K. Lewis. *Counterfactuals*. Oxford: Blackwell Publishers. Cambridge: Harvard University Press., 1973.
- [218] David K. Lewis. *On the Plurality of Worlds*, volume 322. Cambridge University Press, 1986.
- [219] G. Lindblad. “On the Generators of Quantum Dynamical Semigroups”. *Communications in Mathematical Physics*, 48(2):119–130, 1976. URL: <http://dx.doi.org/10.1007/BF01608499>, doi:10.1007/BF01608499.
- [220] Andrei D. Linde. “Eternal chaotic inflation”. *Modern Physics Letters A*, 1(02):81–85, 1986.
- [221] Andrei D. Linde. “Eternally existing self-reproducing chaotic inflationary universe”. *Physics Letters B*, 175(4):395–400, 1986.
- [222] Noah Linden, Sandu Popescu, Anthony J. Short, and Andreas Winter. “Quantum mechanical evolution towards thermal equilibrium”. *Physical Review E*, 79(6):061103, June 2009. URL: <http://link.aps.org/doi/10.1103/PhysRevE.79.061103>, arXiv:0812.2385, doi:10.1103/PhysRevE.79.061103.
- [223] A. M. Lisevski. “On the classical hydrodynamic limit of quantum field theories”. 1999. arXiv:quant-ph/9905014.
- [224] Olimpia Lombardi, Juan Sebastián Ardenghi, Sebastian Fortin, and Martin Narvaja. “Foundations of quantum mechanics: decoherence and interpretation”. *International Journal of Modern Physics D*, 20(05):861–875, 2011. URL: <http://www.worldscientific.com/doi/abs/10.1142/S0218271811019074>, arXiv:1009.0322, doi:10.1142/S0218271811019074.
- [225] Olimpia Lombardi and Mario Castagnino. “A modal-Hamiltonian interpretation of quantum mechanics”. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 39(2):380–443, 2008. URL: <http://www.sciencedirect.com/science/article/pii/S1355219808000051>, arXiv:quant-ph/0610121, doi:10.1016/j.shpsb.2008.01.003.
- [226] Fernando Lombardo and Francisco D. Mazzitelli. “Coarse graining and decoherence in quantum field theory”. *Physical Review D*, 53(4):2001, 1996. URL: <http://link.aps.org/doi/10.1103/PhysRevD.53.2001>, arXiv:hep-th/9508052, doi:10.1103/PhysRevD.53.2001.
- [227] Gerhart Lüders. “Über die Zustandsänderung durch den Meßprozeß”. *Annalen der Physik*, 443(5-8):322–328, 1950. URL: <http://onlinelibrary.wiley.com/doi/10.1002/andp.19504430510/abstract>, doi:10.1002/andp.19504430510.
- [228] Juan Maldacena. “The Large-N Limit of Superconformal Field Theories and Supergravity”. *International Journal of Theoretical Physics*, 38(4):1113–1133, June 1999. URL: <http://dx.doi.org/10.1023/A:1026654312961>, arXiv:hep-th/9711200, doi:10.1023/A:1026654312961.
- [229] Stanley Mandelstam. “Soliton operators for the quantized sine-Gordon equation”. *Physical Review D*, 11(10):3026–3030, May 1975. URL: <http://link.aps.org/doi/10.1103/PhysRevD.11.3026>, doi:10.1103/PhysRevD.11.3026.
- [230] André L.G. Mandolesi. “Analysis of Wallace’s Decision Theoretic Proof of the Born Rule in Everettian Quantum Mechanics”. 2015. The author refers to maverick branches with small but nonzero amplitudes as macrostate discontinuities. arXiv:1504.05259.
- [231] Andrey A. Markov. The theory of algorithms. *Trudy Matematicheskogo Instituta im. VA Steklova*, 42:3–375, 1954.
- [232] Tim Maudlin. “Three Measurement Problems”. *Topoi*, 14(1):7–15, 1995.
- [233] Tim Maudlin. “Why Bohm’s theory solves the measurement problem”. *Philosophy of Science*, 62:479–483, 1995.
- [234] N. David Mermin. “What’s Wrong with this Pillow?”. *Physics Today*, 42(4):9–11, April 1989. URL: <http://scitation.aip.org/content/aip/magazine/physicstoday/article/42/4/10.1063/1.2810963>, doi:10.1063/1.2810963.
- [235] N. David Mermin. “Quantum mysteries revisited”. *American Journal of Physics*, 58(8):731–734, 1990. URL: <http://dx.doi.org/10.1119/1.16503>.

- [236] N. David Mermin. “Could Feynman have said this?”. *Physics Today*, 57(5):10–11, May 2004. URL: <http://scitation.aip.org/content/aip/magazine/physicstoday/article/57/5/10.1063/1.1768652>, doi:10.1063/1.1768652.
- [237] David J. Miller and Matthew Farr. “On the Possibility of Ontological Models of Quantum Mechanics”. 2014. [arXiv:1405.2757](https://arxiv.org/abs/1405.2757).
- [238] Charles W. Misner. “Feynman Quantization of General Relativity”. *Reviews of Modern Physics*, 29(3):497–509, July 1957. URL: <http://link.aps.org/doi/10.1103/RevModPhys.29.497>, doi:10.1103/RevModPhys.29.497.
- [239] Baidyanath Misra and Ennackal Chandy George Sudarshan. “The Zeno’s paradox in quantum theory”. *Journal of Mathematical Physics*, 18(4):756, 1977. URL: http://jmp.aip.org/resource/1/jmapaq/v18/i4/p756_s1, doi:10.1063/1.523304.
- [240] Claus Montonen and David Olive. “Magnetic monopoles as gauge particles?”. *Physics Letters B*, 72(1):117–120, 1997. URL: <http://www.sciencedirect.com/science/article/pii/0370269377900764>, doi:10.1016/0370-2693(77)90076-4.
- [241] Wayne C. Myrvold. “Modal Interpretations and Relativity”. *Foundations of Physics*, 32(11):1773–1784, 2002. URL: <http://dx.doi.org/10.1023/A:1021406924313>, [arXiv:quant-ph/0209109](https://arxiv.org/abs/quant-ph/0209109), doi:10.1023/A:1021406924313.
- [242] Wayne C. Myrvold. “Chasing Chimeras”. *The British Journal for the Philosophy of Science*, 60(3):635–646, 2009. URL: <http://bjps.oxfordjournals.org/content/60/3/635.short>.
- [243] Yoichiro Nambu. “Quasi-Particles and Gauge Invariance in the Theory of Superconductivity”. *Physical Review*, 117:648–663, February 1960. URL: <http://link.aps.org/doi/10.1103/PhysRev.117.648>, doi:10.1103/PhysRev.117.648.
- [244] Stefan Nimmrichter, Klaus Hornberger, Philipp Haslinger, and Markus Arndt. “Testing spontaneous localization theories with matter-wave interferometry”. *Physical Review A*, 83:043621, April 2011. URL: <http://link.aps.org/doi/10.1103/PhysRevA.83.043621>, [arXiv:1103.1236](https://arxiv.org/abs/1103.1236), doi:10.1103/PhysRevA.83.043621.
- [245] Harold Ollivier and Wojciech H. Zurek. “Quantum Discord: A Measure of the Quantumness of Correlations”. *Physical Review Letters*, 88(1):017901, December 2001. URL: <http://link.aps.org/doi/10.1103/PhysRevLett.88.017901>, [arXiv:quant-ph/0105072](https://arxiv.org/abs/quant-ph/0105072), doi:10.1103/PhysRevLett.88.017901.
- [246] Alexei Ourjoumtsev, Hyunseok Jeong, Rosa Tualle-Brouri, and Philippe Grangier. “Generation of optical ‘Schrödinger cats’ from photon number states”. *Nature*, 448:784–786, 2007. URL: <http://www.nature.com/nature/journal/v448/n7155/abs/nature06054.html>, doi:10.1038/nature06054.
- [247] Alexei Ourjoumtsev, Rosa Tualle-Brouri, Julien Laurat, and Philippe Grangier. “Generating Optical Schrödinger Kittens for Quantum Information Processing”. *Science*, 312(5770):83–86, 2006. URL: <http://www.sciencemag.org/content/312/5770/83.short>, doi:10.1126/science.1122858.
- [248] Don N. Page. “Quantum Uncertainties in the Schmidt Basis Given by Decoherence”. 2011. [arXiv:1108.2709](https://arxiv.org/abs/1108.2709).
- [249] Philip Pearle. “Combining stochastic dynamical state-vector reduction with spontaneous localization”. *Physical Review A*, 39:2277–2289, March 1989. URL: <http://link.aps.org/doi/10.1103/PhysRevA.39.2277>, doi:10.1103/PhysRevA.39.2277.
- [250] Philip Pechukas. “Reduced Dynamics Need Not Be Completely Positive”. *Physical Review Letters*, 73:1060–1062, August 1994. URL: <http://link.aps.org/doi/10.1103/PhysRevLett.73.1060>, doi:10.1103/PhysRevLett.73.1060.
- [251] Philip Pechukas. “Pechukas Replies”. *Physical Review Letters*, 75:3021–3021, October 1995. “In general, for arbitrary strength of [the coupling with the environment], this is a realization of Alicki’s suggestion that we abandon [the assumption that the assignment map preserves mixtures] and accept nonlinearity as a feature of reduced dynamics outside the weak coupling regime.”. URL: <http://link.aps.org/doi/10.1103/PhysRevLett.75.3021>, doi:10.1103/PhysRevLett.75.3021.
- [252] Roger Penrose. *The Emperor’s New Mind: Concerning Computers, Minds, And The Laws Of Physics*. Oxford University Press, 1989.
- [253] Asher Peres. “Two simple proofs of the Kochen-Specker theorem”. *Journal of Physics A: Mathematical and General*, 24(4):L175–L178, 1991. URL: <http://iopscience.iop.org/0305-4470/24/4/003>, doi:10.1088/0305-4470/24/4/003.
- [254] Asher Peres. “Generalized Kochen-Specker theorem”. *Foundations of Physics*, 26(6):807–812, 1996. URL: <http://link.springer.com/article/10.1007/BF02058634>, [arXiv:quant-ph/9510018](https://arxiv.org/abs/quant-ph/9510018), doi:10.1007/BF02058634.
- [255] Asher Peres and Daniel R. Terno. “Quantum Information and Relativity Theory”. *Reviews of Modern Physics*, 76:93–123, January 2004. URL: <http://link.aps.org/doi/10.1103/RevModPhys.76.93>, [arXiv:quant-ph/0212023](https://arxiv.org/abs/quant-ph/0212023), doi:10.1103/RevModPhys.76.93.
- [256] J.K. Perring and Tony H.R. Skyrme. A model unified field equation. *Nuclear Physics*, 31(0):550–555, 1962. URL: <http://www.sciencedirect.com/science/article/pii/0029558262907745>, doi:10.1016/0029-5582(62)90774-5.
- [257] Michael E. Peskin and Daniel V. Schroeder. *An Introduction to Quantum Field Theory*. Westview Press, 1995. Section 2.1 (“The Necessity of the Field Viewpoint”) discusses the failure of the position eigenbasis to be orthonormal, and Section 2.4 (“The Klein-Gordon Field in Space-Time”) explains the resulting necessity of antiparticles.
- [258] Joseph Polchinski. *String Theory, Volume 1: An Introduction to the Bosonic String*. CRC Press, 2005. With a discussion of point particles in Section 1.2 (“Action Principles”).
- [259] Joseph Polchinski. “Dualities of Fields and Strings”. 2014. [arXiv:1412.5704](https://arxiv.org/abs/1412.5704).
- [260] Alexander M. Polyakov. *Gauge Fields and Strings*. CRC Press, 1987. With a discussion of point particles in Section 9.11 (“Fermi Particles”).
- [261] Sandu Popescu, Anthony J. Short, and Andreas Winter. “The foundations of statistical mechanics from entanglement: Individual states vs. averages”. 2005. [arXiv:quant-ph/0511225](https://arxiv.org/abs/quant-ph/0511225).
- [262] Sandu Popescu, Anthony J. Short, and Andreas Winter. “Entanglement and the foundations of statistical mechanics”. *Nature Physics*, 2(11):754–758, 2006. URL: <http://www.nature.com/nphys/journal/v2/n11/full/nphys444.html>, doi:

- 10.1038/nphys444.
- [263] John Preskill. “Do Black Holes Destroy Information?”. 1992. [arXiv:hep-th/9209058](https://arxiv.org/abs/hep-th/9209058).
- [264] Matthew F. Pusey, Jonathan Barrett, and Terry Rudolph. “On the reality of the quantum state”. *Nature Physics*, 8(6):475–478, 2012. URL: <http://www.nature.com/nphys/journal/v8/n6/full/nphys2309.html>, [arXiv:1111.3328](https://arxiv.org/abs/1111.3328), [doi:10.1038/nphys2309](https://doi.org/10.1038/nphys2309).
- [265] J. Rau and Berndt Müller. “From Reversible Quantum Microdynamics to Irreversible Quantum Transport”. *Physics Reports*, 272(1):1–59, 1996. URL: <http://www.sciencedirect.com/science/article/pii/0370157395000771>, [arXiv:nuc1-th/9505009](https://arxiv.org/abs/nuc1-th/9505009), [doi:10.1016/0370-1573\(95\)00077-1](https://doi.org/10.1016/0370-1573(95)00077-1).
- [266] Carlo Rovelli. “Relational Quantum Mechanics”. *International Journal of Theoretical Physics*, 35(8):1637–1678, 1996. URL: <http://dx.doi.org/10.1007/BF02302261>, [arXiv:quant-ph/9609002](https://arxiv.org/abs/quant-ph/9609002), [doi:10.1007/BF02302261](https://doi.org/10.1007/BF02302261).
- [267] M. A. Rowe, D. Kielpinski, V. Meyer, C. A. Sackett, W. M. Itano, C. Monroe, and D. J. Wineland. “Experimental violation of a Bell’s inequality with efficient detection”. *Nature*, 409:791–794, February 2001. URL: <http://www.nature.com/nature/journal/v409/n6822/full/409791a0.html>, [doi:10.1038/35057215](https://doi.org/10.1038/35057215).
- [268] Mary Beth Ruskai. “Beyond Strong Subadditivity? Improved Bounds on the Contraction of Generalized Relative Entropy”. *Reviews in Mathematical Physics*, 6(05a):1147–1161, 1994. With a proof of the inequality in Section 1.6 (“Trace Norm Relative Entropy”). URL: <http://www.worldscientific.com/doi/abs/10.1142/S0129055X94000407>, [doi:10.1142/S0129055X94000407](https://doi.org/10.1142/S0129055X94000407).
- [269] Maximilian Schlosshauer. “Decoherence, the measurement problem, and interpretations of quantum mechanics”. *Reviews of Modern Physics*, 76(4):1267, 2005. URL: <http://link.aps.org/doi/10.1103/RevModPhys.76.1267>, [arXiv:quant-ph/0312059](https://arxiv.org/abs/quant-ph/0312059), [doi:10.1103/RevModPhys.76.1267](https://doi.org/10.1103/RevModPhys.76.1267).
- [270] Maximilian Schlosshauer and Kristian Camilleri. “The quantum-to-classical transition: Bohr’s doctrine of classical concepts, emergent classicality, and decoherence”. 2008. [arXiv:0804.1609](https://arxiv.org/abs/0804.1609).
- [271] Maximilian Schlosshauer, Johannes Kofler, and Anton Zeilinger. “A snapshot of foundational attitudes toward quantum mechanics”. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 44(3):222–230, 2013. URL: <http://www.sciencedirect.com/science/article/pii/S1355219813000397>, [arXiv:1301.1069](https://arxiv.org/abs/1301.1069), [doi:10.1016/j.shpsb.2013.04.004](https://doi.org/10.1016/j.shpsb.2013.04.004).
- [272] Erwin Schrödinger. “Der stetige Übergang von der Mikro- zur Makromechanik”. *Naturwissenschaften*, 14(28):664–666, 1926. URL: <http://dx.doi.org/10.1007/BF01507634>, [doi:10.1007/BF01507634](https://doi.org/10.1007/BF01507634).
- [273] Erwin Schrödinger. “Die gegenwärtige Situation in der Quantenmechanik”. *Naturwissenschaften*, 23(49):823–828, 1935. With the English translation by John D. Trimmer in the Proceeding of the American Philosophical Society, 1980.
- [274] Erwin Schrödinger. “Discussion of Probability Relations between Separated Systems”. *Mathematical Proceedings of the Cambridge Philosophical Society*, 31(04):555–563, October 1935. URL: http://journals.cambridge.org/article_S0305004100013554, [doi:10.1017/S0305004100013554](https://doi.org/10.1017/S0305004100013554).
- [275] Benjamin Schumacher and Michael Westmoreland. *Quantum Processes[,] Systems, and Information*. Cambridge University Press, 1st edition, 2010. Appendix D (“Generalized Evolution”). Note that the authors’ general ensemble argument for linearity implicitly assumes that the density matrix describes a proper mixture.
- [276] Nathan Seiberg. “Electric-magnetic duality in supersymmetric non-Abelian gauge theories”. *Physics Letters B*, 435(1–2):129–146, 1995. URL: <http://www.sciencedirect.com/science/article/pii/0550321394000238>, [doi:10.1016/0550-3213\(94\)00023-8](https://doi.org/10.1016/0550-3213(94)00023-8).
- [277] Nathan Seiberg. “Emergent Spacetime”. 2006. [arXiv:hep-th/0601234](https://arxiv.org/abs/hep-th/0601234).
- [278] Ashoke Sen. “Strong-Weak Coupling Duality in Four Dimensional String Theory”. *International Journal of Modern Physics A*, 09(21):3707–3750, 1994. URL: <http://www.worldscientific.com/doi/abs/10.1142/S0217751X94001497>, [arXiv:hep-th/9402002](https://arxiv.org/abs/hep-th/9402002), [doi:10.1142/S0217751X94001497](https://doi.org/10.1142/S0217751X94001497).
- [279] Ramamurti Shankar. *Principles of Quantum Mechanics*. Plenum Press, 2nd edition, 1994.
- [280] Claude E. Shannon. “A Mathematical Theory of Communication”. *Bell System Technical Journal*, 27:379–423, 1948. URL: <http://www.alcatel-lucent.com/bstj/vol27-1948/articles/bstj27-3-379.pdf>.
- [281] Claude E. Shannon and Warren Weaver. *The Mathematical Theory of Communication*. University of Illinois Press, 1949.
- [282] Tony H. R. Skyrme. “A Non-Linear Field Theory”. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 260(1300):127–138, 1961. URL: <http://rspa.royalsocietypublishing.org/content/260/1300/127.abstract>, [doi:10.1098/rspa.1961.0018](https://doi.org/10.1098/rspa.1961.0018).
- [283] Tony H. R. Skyrme. “Particle States of a Quantized Meson Field”. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 262(1309):237–245, 1961. URL: <http://rspa.royalsocietypublishing.org/content/262/1309/237.abstract>, [doi:10.1098/rspa.1961.0115](https://doi.org/10.1098/rspa.1961.0115).
- [284] Tony H.R. Skyrme. A unified field theory of mesons and baryons. *Nuclear Physics*, 31(0):556–569, 1962. URL: <http://www.sciencedirect.com/science/article/pii/0029558262907757>, [doi:10.1016/0029-5582\(62\)90775-7](https://doi.org/10.1016/0029-5582(62)90775-7).
- [285] Matteo Smerlak and Carlo Rovelli. “Relational EPR”. *Foundations of Physics*, 37(3):427–445, 2007. URL: <http://dx.doi.org/10.1007/s10701-007-9105-0>, [arXiv:quant-ph/0604064](https://arxiv.org/abs/quant-ph/0604064), [doi:10.1007/s10701-007-9105-0](https://doi.org/10.1007/s10701-007-9105-0).
- [286] Shang Song, Carlton M. Caves, and Bernard Yurke. “Schrödinger Kittens from Optical Back-Action Evasion”. In Joseph H. Eberly, Leonard Mandel, and Emil Wolf, editors, *Coherence and Quantum Optics VI*, pages 1107–1111. Springer, 1989. URL: http://link.springer.com/chapter/10.1007/978-1-4613-0847-8_200, [doi:10.1007/978-1-4613-0847-8_200](https://doi.org/10.1007/978-1-4613-0847-8_200).
- [287] Robert W. Spekkens. “Contextuality for preparations, transformations, and unsharp measurements”. *Physical Review A*, 71(5):052108, May 2005. URL: <http://link.aps.org/doi/10.1103/PhysRevA.71.052108>, [arXiv:quant-ph/0406166](https://arxiv.org/abs/quant-ph/0406166), [doi:10.1103/PhysRevA.71.052108](https://doi.org/10.1103/PhysRevA.71.052108).

- [288] Robert W. Spekkens. “Negativity and Contextuality are Equivalent Notions of Nonclassicality”. *Physical Review Letters*, 101(2):020401, July 2008. URL: <http://link.aps.org/doi/10.1103/PhysRevLett.101.020401>, arXiv:0710.5549, doi: 10.1103/PhysRevLett.101.020401.
- [289] Robert W. Spekkens. “What is the appropriate notion of noncontextuality for unsharp measurements in quantum theory?”. 2013. arXiv:1312.3667.
- [290] Franco Strocchi. “Complex Coordinates and Quantum Mechanics”. *Reviews of Modern Physics*, 38(1):36–40, 1966. URL: <http://link.aps.org/doi/10.1103/RevModPhys.38.36>, doi:10.1103/RevModPhys.38.36.
- [291] Andrew Strominger, Shing-Tung Yau, and Eric Zaslow. “Mirror Symmetry is T-Duality”. *Nuclear Physics B*, 479(1–2):243–259, 1996. URL: <http://www.sciencedirect.com/science/article/pii/0550321396004348>, arXiv:hep-th/9606040, doi:10.1016/0550-3213(96)00434-8.
- [292] Walter T. Strunz, Lajos Diósi, and Nicolas Gisin. “Open System Dynamics with Non-Markovian Quantum Trajectories”. *Physical Review Letters*, 82:1801–1805, March 1999. URL: <http://link.aps.org/doi/10.1103/PhysRevLett.82.1801>, arXiv:quant-ph/9803079, doi:10.1103/PhysRevLett.82.1801.
- [293] Ward Struyve. “On the zig-zag pilot-wave approach for fermions”. *Journal of Physics A: Mathematical and Theoretical*, 45(19):195307, 2012. URL: <http://stacks.iop.org/1751-8121/45/i=19/a=195307>, arXiv:1201.4169, doi:10.1088/1751-8113/45/19/195307.
- [294] Ernst C. G. Stueckelberg. “Field quantization and time reversal in real Hilbert space”. *Helvetica Physica Acta*, 33(254), 1960.
- [295] Ernst C. G. Stueckelberg. “Quantum theory in real Hilbert space”. *Helvetica Physica Acta*, 33(727), 1960.
- [296] E. C. G. Sudarshan, P. M. Mathews, and Jayaseetha Rau. “Stochastic Dynamics of Quantum-Mechanical Systems”. *Physical Review*, 121:920–924, February 1961. URL: <http://link.aps.org/doi/10.1103/PhysRev.121.920>, doi:10.1103/PhysRev.121.920.
- [297] Leonard Susskind. *The Cosmic Landscape: String Theory and the Illusion of Intelligent Design*. Little, Brown, 2005.
- [298] George Svetlichny, Michael Redhead, Harvey Brown, and Jeremy Butterfield. “Do the Bell Inequalities Require the Existence of Joint Probability Distributions?”. *Philosophy of Science*, 55(3):387–401, 1988. URL: <http://www.jstor.org/stable/187655>.
- [299] Gerard ’t Hooft. “Determinism Beneath Quantum Mechanics”. In *Quo Vadis Quantum Mechanics?*, The Frontiers Collection, pages 99–111. Springer, 2005. URL: http://dx.doi.org/10.1007/3-540-26669-0_8, arXiv:quant-ph/0212095, doi:10.1007/3-540-26669-0_8.
- [300] Gerard ’t Hooft. “*The Mathematical Basis for Deterministic Quantum Mechanics*”, pages 3–19. World Scientific, 2007. Chapter 1. URL: http://www.worldscientific.com/doi/abs/10.1142/9789812771186_0001, arXiv:quant-ph/0604008, doi:10.1142/9789812771186_0001.
- [301] Gerard ’t Hooft. “How a wave function can collapse without violating Schrödinger’s equation, and how to understand Born’s rule”. 2011. arXiv:1112.1811.
- [302] Gerard ’t Hooft. “Discreteness and Determinism in Superstrings”. 2012. arXiv:1207.3612.
- [303] Gerard ’t Hooft. “Superstrings and the Foundations of Quantum Mechanics”. *Foundations of Physics*, pages 1–9, 2014. URL: <http://link.springer.com/article/10.1007/s10701-014-9790-4>, doi:10.1007/s10701-014-9790-4.
- [304] Gerard ’t Hooft. “The Cellular Automaton Interpretation of Quantum Mechanics: A View on the Quantum Nature of our Universe, Compulsory or Impossible?”. 2014. arXiv:1405.1548.
- [305] Max Tegmark. “The Interpretation of Quantum Mechanics: Many Worlds or Many Words?”. *Fortschritte der Physik*, 46(6-8):855–862, 1998. URL: [http://dx.doi.org/10.1002/\(SICI\)1521-3978\(199811\)46:6/8<855::AID-PROP855>3.0.CO;2-Q](http://dx.doi.org/10.1002/(SICI)1521-3978(199811)46:6/8<855::AID-PROP855>3.0.CO;2-Q), arXiv:quant-ph/9709032, doi:10.1002/(SICI)1521-3978(199811)46:6/8<855::AID-PROP855>3.0.CO;2-Q.
- [306] Max Tegmark. “Importance of quantum decoherence in brain processes”. *Physical Review E*, 61(4):4194–4206, April 2000. URL: <http://link.aps.org/doi/10.1103/PhysRevE.61.4194>, arXiv:quant-ph/9907009, doi:10.1103/PhysRevE.61.4194.
- [307] Max Tegmark and Harold S. Shapiro. “Decoherence produces coherent states: An explicit proof for harmonic chains”. *Physical Review E*, 50(4):2538–2547, October 1994. URL: <http://link.aps.org/doi/10.1103/PhysRevE.50.2538>, doi:10.1103/PhysRevE.50.2538.
- [308] W. Tittel, J. Brendel, B. Gisin, T. Herzog, H. Zbinden, and N. Gisin. “Experimental demonstration of quantum correlations over more than 10 km”. *Physical Review A*, 57:3229–3232, May 1998. URL: <http://link.aps.org/doi/10.1103/PhysRevA.57.3229>, arXiv:quant-ph/9707042, doi:10.1103/PhysRevA.57.3229.
- [309] W. Tittel, J. Brendel, H. Zbinden, and N. Gisin. “Violation of Bell Inequalities by Photons More Than 10 km Apart”. *Physical Review Letters*, 81:3563–3566, October 1998. URL: <http://link.aps.org/doi/10.1103/PhysRevLetters81.3563>, arXiv:quant-ph/9806043, doi:10.1103/PhysRevLetters81.3563.
- [310] Myron Tribus. *Thermostatistics and Thermodynamics*. Center for Advanced Engineering Study, Massachusetts Institute of Technology, 1961. Where Tribus defines the “surprisal” as the “unpredictability of a single digit or letter” in a word.
- [311] Angel G. Valdenebro. “Assumptions Underlying Bell’s Inequalities”. *European Journal of Physics*, 23(5):569–577, 2002. URL: <http://stacks.iop.org/0143-0807/23/i=5/a=313>, arXiv:quant-ph/0208161, doi:10.1088/0143-0807/23/5/313.
- [312] Caspar H. Van Der Wal, A. C. J. Ter Haar, F. K. Wilhelm, R. N. Schouten, C. J. P. M. Harmans, T. P. Orlando, Seth Lloyd, and J.E. Mooij. “Quantum Superposition of Macroscopic Persistent-Current States”. *Science*, 290(5492):773–777, 2000. URL: <http://www.sciencemag.org/content/290/5492/773.short>, doi:10.1126/science.290.5492.773.
- [313] Bas C. Van Fraassen. “A Formal Approach to the Philosophy of Science”. In R. Colodny, editor, *Paradigms and Paradoxes: The Philosophical Challenge of the Quantum Domain*, pages 303–366. University of Pittsburgh Press, 1972.

- [314] Bas C. Van Fraassen. *Quantum Mechanics: An Empiricist View*. Oxford University Press, 1991.
- [315] Pieter E. Vermaas. “A No-Go Theorem for Joint Property Ascriptions in Modal Interpretations of Quantum Mechanics”. *Physical Review Letters*, 78:2033–2037, March 1997. URL: <http://link.aps.org/doi/10.1103/PhysRevLett.78.2033>, doi:10.1103/PhysRevLett.78.2033.
- [316] Pieter E. Vermaas. *A Philosopher’s Understanding of Quantum Mechanics: Possibilities and Impossibilities of a Modal Interpretation*. Cambridge University Press, 1999. With a discussion of realistic measurement devices and how they evade criticisms of the modal interpretations in Chapter 10 (“The measurement problem”). URL: <http://books.google.com/books?id=X2V769s0QIYC>.
- [317] Pieter E. Vermaas and Dennis Dieks. “The Modal Interpretation of Quantum Mechanics and Its Generalization to Density Operators”. *Foundations of Physics*, 25(1):145–158, 1995. URL: <http://link.springer.com/article/10.1007/BF02054662>.
- [318] Alexander Vilenkin. *Many Worlds in One: The Search for Other Universes*. Macmillan, 2006.
- [319] Alexander Vilenkin. “Global Structure of the Multiverse and the Measure Problem”. *AIP Conference Proceedings*, 1514(1):7–13, 2013. URL: <http://scitation.aip.org/content/aip/proceeding/aipcp/10.1063/1.4791716>, arXiv:1301.0121, doi:10.1063/1.4791716.
- [320] Hans Christian von Baeyer. “Quantum Weirdness? It’s All in Your Mind”. *Scientific American*, pages 46–51, June 2013.
- [321] John von Neumann. “Wahrscheinlichkeitstheoretischer Aufbau der Quantenmechanik”. *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse*, 1927:245–272, 1927. URL: <https://eudml.org/doc/59230>.
- [322] John von Neumann. *Mathematische Grundlagen der Quantenmechanik*. Berlin: Springer, 1932.
- [323] John von Neumann. *Mathematical Foundations of Quantum Mechanics*. Princeton University Press, 1955.
- [324] David Wallace. “Quantum Probability and Decision Theory, Revisited”. 2002. arXiv:quant-ph/0211104.
- [325] David Wallace. “Worlds in the Everett Interpretation”. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 33(4):637–661, 2002. URL: <http://www.sciencedirect.com/science/article/pii/S1355219802000321>, arXiv:quant-ph/0103092, doi:10.1016/S1355-2198(02)00032-1.
- [326] David Wallace. “Everettian Rationality: defending Deutsch’s approach to probability in the Everett interpretation”. 2003. arXiv:quant-ph/0303050.
- [327] David Wallace. “A formal proof of the Born rule from decision-theoretic assumptions”. 2009. arXiv:0906.2718.
- [328] David Wallace. “Inferential vs. Dynamical Conceptions of Physics”. 2013. arXiv:1306.4907.
- [329] Gregor Weihs, Thomas Jennewein, Christoph Simon, Harald Weinfurter, and Anton Zeilinger. “Violation of Bell’s Inequality under Strict Einstein Locality Conditions”. *Physical Review Letters*, 81:5039–5043, December 1998. URL: <http://link.aps.org/doi/10.1103/PhysRevLetters81.5039>, arXiv:quant-ph/9810080, doi:10.1103/PhysRevLetters81.5039.
- [330] Steven Weinberg. *Gravitation and Cosmology: Principles and Applications of the General Theory of Relativity*. John Wiley and Sons, 1972. With a discussion of the necessity of antiparticles in Section 2.13 (“Temporal Order and Antiparticles”).
- [331] Steven Weinberg. “Sokal’s Hoax”. *New York Review of Books*, XLIII(13):11–15, August 1996.
- [332] Steven Weinberg. *The Quantum Theory of Fields*, volume 1. Cambridge University Press, 1996. For a development of classical gauge theories from the perspective of Weyl-Coulomb gauge, in which the gauge potential is not a proper Lorentz four-vector, see Section 5.9.
- [333] Steven Weinberg. *Lectures on Quantum Mechanics*. Cambridge University Press, 1st edition, 2012. In particular, the end of Section 3.7 (“The Interpretations of Quantum Mechanics”): “... [I]t would be disappointing if we had to give up the ‘realist’ goal of finding complete descriptions of physical systems, and of using this description to derive the Born rule, rather than just assuming it. We can live with the idea that the state of a physical system is described by a vector in Hilbert space rather than by numerical values of the positions and momenta of all the particles in the system, but it is hard to live with no description of physical states at all, only an algorithm for calculating probabilities. My own conclusion (not universally shared) is that today there is no interpretation of quantum mechanics that does not have serious flaws, and that we ought to take seriously the possibility of finding some more satisfactory other theory, to which quantum mechanics is merely a good approximation”.
- [334] Steven Weinberg. “Collapse of the State Vector”. 2013. arXiv:1109.6462.
- [335] Steven Weinberg. “Quantum Mechanics Without State Vectors”. 2014. arXiv:1405.3483.
- [336] Steven Weinberg and Jermy N. A. (interviewer) Matthews. “Questions and answers with Steven Weinberg”. *Physics Today*, July 2013. In particular, “... Some very good theorists seem to be happy with an interpretation of quantum mechanics in which the wavefunction only serves to allow us to calculate the results of measurements. But the measuring apparatus and the physicist are presumably also governed by quantum mechanics, so ultimately we need interpretive postulates that do not distinguish apparatus or physicists from the rest of the world, and from which the usual postulates like the Born rule can be deduced. This effort seems to lead to something like a ‘many worlds’ interpretation, which I find repellent. Alternatively, one can try to modify quantum mechanics so that the wavefunction does describe reality, and collapses stochastically and nonlinearly, but this seems to open up the possibility of instantaneous communication”. URL: <http://scitation.aip.org/content/aip/magazine/physicstoday/news/10.1063/PT.4.2502>, doi:10.1063/PT.4.2502.
- [337] John A. Wheeler. “Assessment of Everett’s “Relative State” Formulation of Quantum Theory”. *Reviews of Modern Physics*, 29(3):463–465, July 1957. URL: <http://link.aps.org/doi/10.1103/RevModPhys.29.463>, doi:10.1103/RevModPhys.29.463.
- [338] Eugene Paul Wigner. *Symmetries and Reflections*. Indiana University Press, Bloomington, Ind., 1967.

- [339] F. K. Wilhelm, C. H. Van Der Wal, A. C. J. Ter Haar, R. N. Schouten, C. J. P. M. Harmans, J. E. Mooij, T. P. Orlando, and Seth Lloyd. “Macroscopic quantum superposition of current states in a Josephson-junction loop”. *Physics-Uspekhi*, 44(10S):117, 2001. URL: <http://iopscience.iop.org/1063-7869/44/10S/S26>, doi:10.1070/1063-7869/44/10S/S26.
- [340] Timothy Williamson. *Modal Logic as Metaphysics*. Oxford University Press, 2013.
- [341] Edward Witten. “Phases of $N = 2$ theories in two dimensions”. *Nuclear Physics B*, 403(1–2):159–222, 1993. URL: <http://www.sciencedirect.com/science/article/pii/055032139390033L>, arXiv:hep-th/9301042, doi:10.1016/0550-3213(93)90033-L.
- [342] Edward Witten. “Phase transitions in M-theory and F-theory”. *Nuclear Physics B*, 471(1–2):195–216, 1996. URL: <http://www.sciencedirect.com/science/article/pii/055032139600212X>, arXiv:hep-th/9603150, doi:10.1016/0550-3213(96)00212-X.
- [343] Edward Witten. “Anti-de Sitter Space and Holography”. *Advances in Theoretical and Mathematical Physics*, 2:253–291, 1998. arXiv:hep-th/9802150.
- [344] Ludwig Wittgenstein. *Tractatus Logico-Philosophicus*. Routledge & Kegan Paul, 1921. URL: <http://www.gutenberg.org/ebooks/5740>.
- [345] Janik Wolters, Max Strauß, Rolf Simon Schoenfeld, and Oliver Benson. “Observation of the Quantum Zeno Effect on a Single Solid State Spin”. 2013. arXiv:1301.4544.
- [346] William K. Wootters. “Optimal Information Transfer and Real-Vector-Space Quantum Theory”. 2013. arXiv:1301.2018.
- [347] Hai-Long Zhao. “On the Implication of Bell’s Probability Distribution and Proposed Experiments of Quantum Measurement”. 2009. arXiv:0906.0279.
- [348] Wojciech H. Zurek. “Pointer basis of quantum apparatus: Into what mixture does the wave packet collapse?”. *Physical Review D*, 24:1516–1525, September 1981. URL: <http://link.aps.org/doi/10.1103/PhysRevD.24.1516>, doi:10.1103/PhysRevD.24.1516.
- [349] Wojciech H. Zurek. “Decoherence and the transition from quantum to classical”. 2003. arXiv:quant-ph/0306072.
- [350] Wojciech H. Zurek, Salman Habib, and Juan Pablo Paz. “Coherent States via Decoherence”. *Physical Review Letters*, 70:1187–1190, March 1993. URL: <http://link.aps.org/doi/10.1103/PhysRevLett.70.1187>, doi:10.1103/PhysRevLett.70.1187.