Typical Quantum States of the Universe are Observationally Indistinguishable

Eddy Keming Chen* and Roderich Tumulka†

October 22, 2024

Abstract: This paper is about the epistemology of quantum theory. We establish a new result about a limitation to knowledge of its central object—the quantum state of the universe. We show that, if the universal quantum state can be assumed to be a typical unit vector from a high-dimensional subspace of Hilbert space (such as the subspace defined by a low-entropy macro-state as prescribed by the Past Hypothesis), then no observation can determine (or even just narrow down significantly) which vector it is. Typical state vectors, in other words, are observationally indistinguishable from each other. Our argument is based on a typicality theorem from quantum statistical mechanics. We also discuss how theoretical considerations that go beyond the empirical evidence might bear on this fact and on our knowledge of the universal quantum state.

Key words: limitation to knowledge; empirical equivalence; past hypothesis; quantum statistical mechanics.

Contents

*Department of Philosophy, University of California San Diego, 9500 Gilman Dr, La Jolla, CA 92093, USA. Email: eddykemingchen@ucsd.edu

†Fachbereich Mathematik, Eberhard-Karls-Universität Tübingen, Auf der Morgenstelle 10, 72076 Tübingen, Germany. Email: roderich.tumulka@uni-tuebingen.de

1 Introduction

This paper is about the epistemology of quantum theory. We establish a new result about a limitation to knowledge of the quantum state of the universe. We show that typical universal quantum states are observationally indistinguishable, in the sense that no observation can distinguish typical universal quantum states in a high-dimensional subspace of the Hilbert space, such as the subspace characterized by a low-entropy macro-state as prescribed by the Past Hypothesis (see §2.2 and (Albert 2000, Goldstein et al. 2020)). We call this fact *observation typicality*.

Observation typicality follows from a result in quantum statistical mechanics called *distribution typicality* (Reimann 2007, Teufel et al. 2023): for any observable, most quantum states in a high-dimensional Hilbert subspace lead to very nearly the same probability distribution for that observable. We show that this implies observation typicality.

Observation typicality is a surprisingly strong result. There are other known examples of epistemic limitations in physics. Dürr et al. (2004) already proved that it is impossible to measure a wave function precisely. Observation typicality is a further step in this direction, and it delivers a stronger conclusion—even an imprecise measurement is impossible. No observation, as long as it does not disconfirm the physical theory, will give us any substantive information about the quantum state. In this sense, limitation to knowledge is more pervasive and persistent in a quantum universe than has been recognized.

A statement about *most* quantum states in the full Hilbert space $\mathscr{H}$ of the universe has limited applicability because it may fail for the quantum states of interest. For example, most quantum states in $\mathscr{H}$ are in thermal equilibrium, but the actual quantum state of our universe is not. To use typicality theorems in our situation, we apply them to a subspace $\mathscr{H}_0$ of $\mathscr{H}$ such as the subspace comprising the quantum states Ψ compatible with a certain macro-state of the universe that may be far from thermal equilibrium. We assume that $\mathscr{H}_0$ has finite but huge dimension d_0 (while $\mathscr{H}$ may have finite or infinite dimension). Observation typicality then holds for most quantum states Ψ in $\mathscr{H}_0$. We can use the result while accepting certain inductive hypotheses about the physical laws of the universe, including those pertaining to initial conditions, according to which our universe is a typical member of the possible universes compatible with such laws.

We assume here that there is a wave function Ψ_t of the universe and that at the initial time t_0 of the universe (say, at the big bang), Ψ_{t_0} had to lie in a particular subspace $\mathscr{H}_0$ of the Hilbert space $\mathscr{H}$ of the universe (say, corresponding to a low-entropy macro-state, as prescribed by the Past Hypothesis). Observation typicality asserts that our empirical data at any time t will reveal very little about Ψ_{t_0} (beyond its membership in $\mathscr{H}_0$, which we take as known) because most Ψ_{t_0} from $\mathscr{H}_0$ lead to very nearly the same probability distribution for these empirical data. In other words, typical initial quantum states of the universe are observationally indistinguishable from each other. If the time evolution is unitary, then the same arguments yield that also Ψ_t cannot be distinguished empirically from typical alternative vectors in the appropriate subspace, i.e., in $U_t\mathscr{H}_0$ with $U_t = \exp(-iH(t-t_0))$ and H the Hamiltonian of the universe.

The analogous statement fails in classical mechanics. Observation typicality holds in quantum mechanics because, roughly speaking, quantum states can be in superpositions, and the superposition weights of any particular observation given by typical quantum states are nearly the same. The latter is related to the mathematical fact, known as the concentration of measure phenomenon (the situation where random quantities become nearly deterministic), which applies in any high-dimensional Hilbert space.

We present two main results. First (§2 and §3), that a typical unit vector $\Psi_{t_0} \in \mathscr{H}_0$ cannot be reliably distinguished from the density matrix ρ_0 associated with a uniform distribution in $\mathscr{H}_0$, i.e., $\rho_0 = P_0/d_0$ with P_0 the projection operator to $\mathscr{H}_0$ and $d_0 = \dim \mathscr{H}_0$; as a consequence, typical unit vectors cannot be reliably distinguished from each other. The second result (§3.4) is even stronger and asserts that even *unreliably*, they cannot be distinguished; that is, our empirical data do not even yield (any significant amount of) partial or probabilistic information about Ψ_{t_0}. This result can be expressed in terms of Bayesian credences as follows: if we start from the uniform probability distribution u_0 over the unit sphere $\mathbb{S}(\mathscr{H}_0)$ (containing all normalized wave functions) in $\mathscr{H}_0$ as the prior distribution, and if we update our credences in a Bayesian way on the basis of our empirical observation (that is not too improbable), then the updated distribution will still be very nearly uniform over $\mathbb{S}(\mathscr{H}_0)$.

Some empirical observations might, of course, change our mind about which subspace $\mathscr{H}_0$ is the correct one. Since we regard here the specification of $\mathscr{H}_0$ as part of the given theory, such an observation would amount to a disconfirmation of the theory. This leads to the interesting question of which kind of observations should be regarded as disconfirming the theory. We will look more into this question elsewhere (Chen and Tumulka 2024), take here for granted that the theory is correct and not getting disconfirmed, and remark only briefly that since the specification of $\mathscr{H}_0$ has the status of a law of nature, we may expect it to be simple and natural in order to be convincing; as a consequence, we would not be inclined to adjust the choice of $\mathscr{H}_0$ in complicated ways just to adapt it to some random-looking features of our world (such as, e.g., the shape of Ireland).

Our goal is to further the discussion about the epistemology of physical theories and highlight the limitations of (direct) empirical knowledge in quantum universes. Insofar as we have detailed knowledge about the universal quantum state, it is even more theoretical than has been recognized. Our result should be of interest to anyone interested in the epistemological implications of physics. We highlight several distinctive features of the present approach:

- Observation typicality applies to all interpretations of quantum mechanics. It is also compatible with generous inductive assumptions about what we can infer based on available evidence. For example, it is compatible with our knowledge of the dynamical laws of quantum mechanics, the Past Hypothesis, and the aggregation of individual observations into a collection.

- Our argument relies on the concentration of measure phenomenon instead of special constructive techniques, such as the cut-and-paste method used in Manchak (2009)'s argument for observational indistinguishability in general relativity. The standard "physical unreasonableness" objections against arguments based on special constructions do not apply to our argument, which establishes observational indistinguishability for *generic* pairs of models.
- Our claims are based on quantitative bounds (recently proven mathematically) on how large the deviations of certain probabilities can at most become for which fraction of all wave functions in a high-dimensional Hilbert subspace.

In §2, we formulate distribution typicality and explain how it applies to the special case of the Past Hypothesis. In §3, we clarify our criterion for observational indistinguishability, derive observation typicality from distribution typicality, apply the fact to the universal quantum state, and provide a Bayesian analysis. In §4, we compare observation typicality to known results about epistemic limitations in physics. In §5, we discuss potential philosophical implications. For more detail, see the longer version, Chen and Tumulka (2024).

2 Distribution Typicality

In the 21st century, many exciting new results have been proven in quantum statistical mechanics and quantum thermodynamics, leading to improved understanding in the thermodynamic behavior of a closed quantum system in a pure state (Gemmer et al. 2009, Tasaki 2016, Mori et al. 2018). Such results often assume an "individualist" attitude, according to which an individual quantum system in a pure state, as opposed to an ensemble, is studied with respect to its thermodynamic properties. These results often concern *typical properties* of those states.[1]

We shall focus on *distribution typicality* in this paper. In §3, we show that distribution typicality implies observation typicality. In this section, we formulate distribution typicality and specialize it to the Past Hypothesis.

2.1 The Key Theorem

Before stating the key theorem, we highlight its implications for orthodox quantum mechanics (OQM), Bohmian mechanics (BM), and the Ghirardi-Rimini-Weber spontaneous collapse theory (GRW) (Tumulka 2022b, Dürr and Teufel 2009, Ghirardi and Bassi 2020). Distribution typicality entails that, in OQM, if an observer at time t conducts an

[1]There is a growing philosophical literature about typicality and its role in physical theories. See, for example, Goldstein (2012), Reichert (2018), Wilhelm (2022), Lazarovici (2023), Hubert (2024), (Frigg and Werndl 2024, sec. 4.5), and the references therein.

experiment, the probability distribution of the outcome is nearly independent of Ψ_{t_0}; in BM, at any time t, the probability distribution of the Bohmian configurations Q_t of the universe is nearly independent of Ψ_{t_0}; in GRW with a flash ontology, for any time $t > t_0$, the probability distribution of the pattern of flashes up to t is nearly independent of Ψ_{t_0}. We will elucidate in §3 why this near-independence entails the impossibility for observers according to OQM, BM, or GRW to distinguish among different Ψ_{t_0}.

Terminology. Recall that $\mathbb{S}(\mathscr{H}_0)$ is the unit sphere in the subspace $\mathscr{H}_0$, $d_0 = \dim \mathscr{H}_0$, and u_0 the uniform probability distribution over $\mathbb{S}(\mathscr{H}_0)$; that is, for any subset S of $\mathbb{S}(\mathscr{H}_0)$, $u_0(S)$ is the surface area of S, normalized by dividing through the surface area of $\mathbb{S}(\mathscr{H}_0)$. We say that a statement $s(\Psi)$ is true for $(1-\varepsilon)$-most $\Psi \in \mathbb{S}(\mathscr{H}_0)$ if and only if the set S of $\Psi \in \mathbb{S}(\mathscr{H}_0)$ satisfying $s(\Psi)$ has $u_0(S) \geq 1-\varepsilon$ (equivalently, if S has at least the fraction $1-\varepsilon$ of the surface area, or if a purely random point on the sphere has probability $\geq 1-\varepsilon$ to lie in S). Averages on $\mathbb{S}(\mathscr{H}_0)$ will also be taken with respect to u_0.[2] We also need the concept of POVM.

POVM. POVM stands for positive-operator-valued measure, which is a generalization of the concept of an observable as given by a self-adjoint operator. Technically, a POVM on a Hilbert space $\mathscr{H}$ is a family of positive operators E_z on $\mathscr{H}$ that add up to the identity operator. POVMs have an important representational role in quantum mechanics, given by the Main Theorem about POVMs: *for every quantum physical experiment $\mathscr{E}$ on a quantum system $\mathscr{S}$ whose possible outcomes lie in a space $\mathscr{Z}$, there exists a POVM E on $\mathscr{Z}$ such that, whenever $\mathscr{S}$ has wave function Ψ at the beginning of $\mathscr{E}$, the random outcome Z has probability distribution given by*

$$\mathbb{P}(Z = z) = \langle\Psi| E_z |\Psi\rangle. \tag{1}$$

See Dürr et al. (2004) and Tumulka (2022b) for proofs.

The Main Theorem about POVMs provides a generalization of the Born rule to arbitrary experiments (instead of ideal quantum measurements). The operator E_z is associated with the possible outcome z; in the case of an ideal quantum measurement, E_z would be the projection operator to the eigenspace of the eigenvalue z.

We can now state the key theorem as follows:

Theorem 1 (Distribution Typicality)**.** *Let $\mathscr{H}_0$ be an arbitrary subspace with finite dimension d_0 in $\mathscr{H}$, let the density matrix $\rho_0 = P_0/d_0$ be the normalized projection to $\mathscr{H}_0$, and let the operator E_z be an element of an arbitrary POVM on $\mathscr{H}$. Then for every $\varepsilon > 0$, for $(1-\varepsilon)$-most $\Psi \in \mathbb{S}(\mathscr{H}_0)$,*

$$\Big|\langle\Psi| E_z |\Psi\rangle - \mathrm{tr}(\rho_0 E_z)\Big| \leq \frac{1}{\sqrt{\varepsilon d_0}}\,. \tag{2}$$

Theorem 1 is a particular case of Theorem 3 proven by Teufel et al. (2023). It can also be derived from a theorem proven by Reimann (2007).

It can be shown that $\mathrm{tr}(\rho_0 E_z)$ is exactly the average of $\langle\Psi|E_z|\Psi\rangle$ over $\mathbb{S}(\mathscr{H}_0)$ (using the uniform distribution u_0); thus (2) expresses that $\langle\Psi|E_z|\Psi\rangle$ is close to its average

[2] The results in this paper can be extended to the Gaussian adjusted projected (GAP) measures introduced by Goldstein et al. (2006b), as is done by Teufel et al. (2024). But here, we consider only the uniform measure u_0.

value, provided that the right-hand side is small, which will be the case whenever

$$d_0 \gg \frac{1}{\varepsilon}. \tag{3}$$

Inequality (2) is an example of the "concentration of measure phenomenon," which can be expressed as saying that some random quantity is nearly deterministic (i.e., very probably close to a deterministic value). Here, we have that *for any observation represented by an element in a POVM, typical quantum states in a high-dimensional Hilbert space have very nearly the same Born probability distribution for that observation*.

Let us clarify and elaborate on this result.

Example. A macroscopic system with $N = 10^{20}$ particles is in macro-state μ. The dimension of the Hilbert space $\mathscr{H}_0$ of the system scales exponentially with N. Let us suppose $d_0 = 10^{10^{20}}$. Consider a very small error $\varepsilon = 10^{-200}$. Inequality (2) tells us that, for any quantum experiment with observational outcome z, for $(1-\varepsilon)$-most pure states in $\mathscr{H}_0$, the Born probability of outcome z given by the pure state differs from the ensemble average by at most $10^{100-0.5\times 10^{20}} \approx 3^{-10^{20}}$, which is close to zero.

Minimal Assumptions. Theorem 1 holds in great generality. First, it applies to every POVM. By the Main Theorem about POVMs, the theorem covers every observational outcome arising from an arbitrary quantum physical experiment. Second, it does not make assumptions about the interaction between the subsystem and the environment. Third, it does not invoke chaos, ergodicity, or mixing. The typicality in distribution typicality comes from the typicality of quantum states in $\mathbb{S}(\mathscr{H}_0)$, and the large number comes from d_0, the large dimension of $\mathscr{H}_0$. Finally, distribution typicality applies to thermal equilibrium and non-equilibrium. For concreteness, in §2.2 we provide a physical interpretation for a universe in thermal non-equilibrium, taking $\mathscr{H}_0$ to be the low-entropy initial macro-state of the universe prescribed by the Past Hypothesis.

Difference from the Classical Case. We do not have distribution typicality in classical mechanics. The prediction of any observation from a particular *micro-state* in the classical phase space has a trivial probability of 0 or 1, while the expectation values for many observables with respect to a uniform probability distribution (the ensemble averages) are often non-trivial. It is false that for every experimental outcome, typical micro-states will assign very nearly the same probability distribution. For example, we can design a coin-flip experiment where half the microstates assign probability 0 and the other half assign probability 1 to the "Tails" outcome. The relevant way in which quantum mechanics differs from classical mechanics is that if a wave function $\Psi = \Psi_{t_0}$ gets chosen randomly, then the outcome Z of an observation is "doubly random": Given Ψ, it is random with the Born distribution $\langle\Psi|E_z|\Psi\rangle$, and in addition Ψ is itself random. The crucial point is that in our situation, different Ψs have very similar superposition weights, with the consequence that they have nearly the same $\langle\Psi|E_z|\Psi\rangle$.

2.2 Main Motivation: The Past Hypothesis

A typicality statement about most quantum states in some space has limited applicability if the relevant quantum states are atypical in that space. For example, most quantum states in the energy shell are in thermal equilibrium, but many quantum states we observe are not. That is not a problem in our case, as Theorem 1 is quite general. To apply Theorem 1 to the relevant class of quantum states, we can specialize the arbitrary subspace $\mathscr{H}_0$ to any macro-state that may be far from thermal equilibrium. We are particularly motivated by the thought that the physical laws may require the initial wave function of the universe to lie in a particular subspace $\mathscr{H}_0$.

In fact, it has been suggested for the explanation of the thermodynamic arrow of time (Feynman 1963, p. 115), (Penrose 1979, Lebowitz 1993, Albert 2000, Goldstein et al. 2020) that the initial state of the universe must be restricted to a set of low-entropy states; in a quantum theory, such a set would be given by a suitable subspace $\mathscr{H}_0$ of the Hilbert space $\mathscr{H}$ of the universe, and the presumed additional law might then be formulated as follows:

Past Hypothesis (PH) Ψ_{t_0} is a typical element of $\mathscr{H}_0$.

Here, $\mathscr{H}_0$ is assumed to contain all wave functions compatible with a certain macro-state. For example, Penrose's Weyl curvature hypothesis might amount to taking as $\mathscr{H}_0$ something like the joint eigenspace with all eigenvalues 0 of all Weyl curvature operators at t_0. To explain the arrow of time, the initial macro-state should have very low entropy. Taking as the definition of entropy the quantum Boltzmann entropy (Lebowitz 1993, Goldstein et al. 2010, 2020)

$$S_{qB}(\Psi_{t_0}) = k_B \log(d_0)\,, \tag{4}$$

the condition of low entropy corresponds the condition that the dimension d_0 of $\mathscr{H}_0$ is much smaller than the dimension of the full Hilbert space $\mathscr{H}$, in fact much smaller than the dimension of the subspace $\mathscr{H}_{eq}$ corresponding to thermal equilibrium in the same energy shell, $d_0 \ll d_{eq}$. Then the PH requires Ψ_{t_0} to have low entropy and to be far from thermal equilibrium.

Even though $\mathscr{H}_0$ has much lower dimension than $\mathscr{H}_{eq}$, $\mathscr{H}_0$ is still, like all subspaces representing macro-states, a high-dimensional subspace. Realistically, we can expect $d_0 >> 10^{10^{20}}$. With respect to the normalized uniform measure u_0, we expect that typical wave functions in $\mathbb{S}(\mathscr{H}_0)$ will evolve in a way such that S_{qB} increases in the medium and the long run, satisfying the Developmental Conjecture formulated by Goldstein et al. (2020), a version of the Second Law of Thermodynamics for a quantum universe.

Hence, if we choose (say) $\varepsilon = 10^{-200}$, then for any E_z, typical individual wave functions in $\mathscr{H}_0$ will all lead to nearly the same probability distribution as ρ_0. In the philosophical literature, ρ_0 has been called the *Wentaculus* density matrix, corresponding to the initial state of the Wentaculus theory (Chen 2021, 2024a).

3 Observation Typicality

In this section, we show that distribution typicality implies observation typicality.

3.1 General Distinctions

There are different kinds of observational indistinguishability (OI) discussed in the philosophical literature. We draw some general distinctions regarding (1) OI and empirical equivalence, and (2) in-principle and in-practice OI.

Empirical equivalence is a relation holding among certain physical theories. Two theories may be regarded as empirically equivalent when they make the same predictions about all possible observations. For example, one can explore whether BM is empirically equivalent to OQM, GRW, or Everettian quantum mechanics by checking whether BM makes the same empirical predictions as the latter.

In contrast, OI is primarily a relation holding among models in the same theory. In the philosophical literature, there is a large literature about the OI of certain models of general relativity. Manchak (2009) showed that for almost every general relativistic spacetime, we can construct another one in such a way that no observation will distinguish between the two spacetimes.

In this paper, we focus on OI as it is applied to models of quantum theory. Here, we shall take an interpretation-neutral approach, and discuss how to apply the result in different interpretations in §3.2. Nevertheless, what we show may also be relevant to the issue of empirical equivalence among quantum theories.

The second distinction is between in-principle OI and in-practice OI. People often focus on exact equivalence of predictions (or retrodictions) about outcomes of observations. This is done in the general relativistic case in terms of exact agreement (up to an isometry transformation) of past lightcones of observers, and in quantum foundations in terms of exact equality of theoretical distributions. In many cases, exact equivalence is sufficient but not necessary for OI. We never did and never will probe the exact microscopic detail of the past lightcone. In practice we will quickly exhaust our resources long before reaching the in-principle limits. Nevertheless, to prove sharp mathematical results, it is sometimes easier to focus on in-principle limits. Of course, in-principle OI implies in-practice OI.

Observation typicality concerns the stronger sense of OI—in principle and not just in practice. We show that in-principle OI follows from *approximate* equality of probabilistic predictions. The approximate equality of predictions is to be distinguished from the approximation of frequency to chance. For example in Bohmian mechanics, we justify the quantum equilibrium hypothesis using a law-of-large-number argument (Dürr et al. 1992), with frequency converging to chance: typically (with respect to the quantum equilibrium distribution), the empirical distribution converges to the theoretical distribution. The result establishes exact equality of the chance functions (theoretical distributions). While exact equality of probabilities is sufficient for in-principle OI, our argument below shows that it is not necessary.

According to a standard criterion (Nielsen and Chuang 2010, p.86), two quantum states are perfectly distinguishable by some experiment if and only if they are orthogonal. If we do not insist on perfect distinguishability, we may say that two quantum states are reliably distinguishable by some experiment if and only if they are (at least) approximately orthogonal.

We can specialize the criterion for OI to a particular observation E_z. For two quantum states Ψ_1 and Ψ_2, we say that they are perfectly distinguishable by observing the outcome $Z = z$ if and only if $\langle\Psi_1| E_z |\Psi_1\rangle = 0$ and $\langle\Psi_2| E_z |\Psi_2\rangle = 1$ or vice versa; let us assume the former is the case. If we know that the actual quantum state is either Ψ_1 or Ψ_2, the outcome $Z = z$ is a perfect indicator of Ψ_2. The indicator is perfect in the sense that, in 100% of the worlds compatible with the outcome $Z = z$, it correctly indicates the actual quantum state. If we do not insist on perfect distinguishability, we can allow the probabilities to be approximately zero or one, so that they are reliably distinguishable by E_z just in case $\langle\Psi_1| E_z |\Psi_1\rangle \approx 0$ and $\langle\Psi_2| E_z |\Psi_2\rangle \approx 1$, or $\langle\Psi_1| E_z |\Psi_1\rangle \approx 1$ and $\langle\Psi_2| E_z |\Psi_2\rangle \approx 0$, i.e., the absolute difference in their probabilities is close to 1. The outcome $Z = z$ is a reliable indicator of the quantum state in the sense that, in nearly 100% of the worlds compatible with the outcome $Z = z$, it correctly indicates the actual quantum state. This suggests the following sufficient condition:

OI with respect to E_z If the Born probability distributions of E_z assigned by a set of quantum states are all within (say) 10^{-200} of each other, then observing the outcome $Z = z$ does not reliably distinguish among those states.

In what follows, unless specified otherwise, we shall understand OI with this criterion.

3.2 Observational Indistinguishability of Typical Quantum States

In this section, we show that distribution typicality implies observation typicality. We do this separately for OQM, BM, and GRW. The same arguments as for OQM and BM also provide OI for Everettian quantum mechanics.

(1) OI in OQM. Suppose we carry out any experiment $\mathscr{E}$; suppose $\mathscr{E}$ is associated, according to the Main Theorem about POVMs, with the POVM E; suppose we obtained the outcome $Z = z$ associated with the operator E_z. Then typical quantum states in a high-dimensional Hilbert space $\mathscr{H}_0$ (with $d_0 > 10^{1000}$) are not distinguished by this outcome from ρ_0.

Indeed, consider $\mathscr{H}_0$ with $d_0 > 10^{1000}$ and $\varepsilon = 10^{-200}$. By Theorem 1 for any POVM element E_z, for $(1-\varepsilon)$-most Ψ in $\mathscr{H}_0$, $|\langle\Psi| E_z |\Psi\rangle - \mathrm{tr}(\rho_0 E_z)| \leq \frac{1}{\sqrt{\varepsilon d_0}} < 10^{-400}$. Hence, the Born probability distributions of E_z assigned by $(1-\varepsilon)$-most Ψ are all within $\frac{1}{2}\varepsilon$ of that of ρ_0, and thus within ε of each other. As discussed above, a reliable distinction would require the observation of an event that has probability near 1 in one case and near 0 in the other, which does not happen here. Therefore, observing E_z does not reliably distinguish those states from each other or from the normalized projection. So, we arrive at observation typicality in $\mathscr{H}_0$.

Again, we may think of $\mathscr{H}_0$ as the subspace demanded by the Past Hypothesis. Thus, for any observation E_z, typical quantum states in $\mathscr{H}_0$ are not distinguished by

E_z from the Wentaculus density matrix. Neither does it distinguish typical quantum states in $\mathscr{H}_0$ from each other.

In quantum theory with unitary evolution, this extends to later states too: Since typical Ψ_t from the appropriately evolved subspace $U_t\mathscr{H}_0$ lead to nearly the same probability for outcome $Z = z$, the observation of this outcome does not distinguish among different Ψ_t.

(2) OI in BM. In BM, a key fact about OI is this:

$$\text{\textit{the probability distribution of the configuration } } Q_t \text{\textit{ of the universe at time t, i.e., the }} |\Psi_t|^2 \text{\textit{ distribution in configuration space, is nearly independent of }} \Psi_{t_0} \text{\textit{ for typical }} \Psi_{t_0} \in \mathbb{S}(\mathscr{H}_0). \tag{5}$$

This follows from Theorem 1 by taking the POVM E to be the Heisenberg-evolved configuration observable, i.e., the projection-valued measure jointly diagonalizing all position operators at time t, or

$$E(B) = U_t^\dagger 1_B U_t, \tag{6}$$

where B is any subset of configuration space and 1_B the multiplication operator multiplying by the characteristic function of B. Note that the probability distribution associated with this POVM is exactly the $|\Psi_t|^2$ distribution:

$$\langle \Psi_{t_0}|E(B)|\Psi_{t_0}\rangle = \int_B dq\, |\Psi_t(q)|^2 \tag{7}$$

for any subset B of configuration space.

Now Q_t comprises the exact positions of all particles in the universe, so observers inside a Bohmian universe cannot possibly measure Q_t; at best, they can measure a very coarse-grained version of Q_t. But the point is that even complete information about Q_t would not help with distinguishing between typical Ψ_{t_0} and ρ_0 (or between two typical wave functions); and *a fortiori*, a coarse-grained version does not help.

The essence of the difficulty with distinguishing is an instance of the general problem of deciding between two probability distributions f_1 and f_2 (say, on $\mathbb{R}^n$) after we are given just one point X that was randomly chosen with distribution either f_1 or f_2. If f_1 and f_2 are disjoint, then this is possible; if they are nearly disjoint, then it is still possible quite reliably (for example, at the confidence level of 95% often used in statistics). But if f_1 and f_2 overlap significantly, and if the observed X lies in the overlap region, then this value of X could have arisen from either one, so we cannot decide which f was used. A reliable decision is possible only if $f_1(X) \approx 0$ and $f_2(X)$ is significantly non-zero or vice versa. Now in our case, $f_i = |\Psi_i|^2$ and $f_1 \approx f_2$ by (5), so they overlap almost exactly and everywhere; values of $X = Q_t$ with $f_1(X) \approx 0$ will also have $f_2(X) \approx 0$ and therefore, with overwhelming probability, are not going to occur for either f_1 or f_2. So, the two cannot be distinguished at any time t (not reliably and in fact, as we will discuss in §3.4, not even unreliably).

This conclusion covers also the possibility that observers might make observations (say, with telescopes) or experiments, as the outcomes of these observations or experiments will be recorded in the configuration of some particles (e.g., when scientists publish their findings). It also covers the possibility of making observations or experiments at several times $t_1, t_2, \ldots, t_k$, since the outcomes will be recorded and thus can be read off from Q_t at $t \geq t_k$.

On the other hand, the precise microscopic trajectories of all Bohmian particles may contain more information about Ψ_{t_0}, but it is known (Dürr et al. 1992) that observers inside a Bohmian universe cannot measure positions at several times without introducing decoherence and thereby changing the trajectory. Therefore, they do not have access to the full trajectories, and the observational indistinguishability just derived remains valid. In other words, there is a known limitation to knowledge about full trajectories—the absolute uncertainty in Bohmian mechanics.

(3) OI in GRWf. We consider GRWf, the GRW theory with a flash ontology in spacetime (e.g., Tumulka 2022b, sec. 3.3.4). A key fact about OI in GRWf is this:

$$\text{\textit{the joint probability distribution of all flashes in the universe up to time } t \text{ \textit{is nearly independent of} } \Psi_{t_0} \text{ \textit{for typical} } \Psi_{t_0} \in \mathbb{S}(\mathscr{H}_0). \tag{8}$$

This follows from Theorem 1 by taking the POVM E to be the POVM governing the joint distribution of all flashes (Tumulka 2022b, sec. 5.1.1 and 7.8). As in BM, observers in a GRWf universe do not have access to the full pattern of flashes, only to a coarse-grained version of it. But even the exact pattern would not provide the information needed for distinguishing a typical Ψ_{t_0} from ρ_0 or from another typical initial wave function, for the same reasons based on Theorem 1 as discussed for BM above.

A difference to the Bohmian case is that the exact *history* of the universal configuration in BM may provide sufficient information to determine Ψ_{t_0}, whereas in GRWf even the exact history of all flashes is of no help. Thus, GRWf provides, in a sense, a stronger kind of limitation to knowledge than BM.

Another difference to BM is that inhabitants of a GRWf universe living at time t may very well find out a lot about Ψ_t. After all, Ψ_t is not unitarily evolved from Ψ_{t_0} but has collapsed; for example, if we find Schrödinger's cat alive then its wave function will have collapsed to (approximately) the wave function of a live cat. (Note that this still does not give us new information about Ψ_{t_0}.)

3.3 Remarks

Let us return to the general result of observation typicality and discuss the order of quantifiers. It has the form "for any E and for most Ψ, $s(\Psi, E)$." This is weaker than the statement "for most Ψ and for any E, $s(\Psi, E)$," which is *false* in this context. It is easy to confuse the two. The false but stronger statement corresponds to the following:

Super-Strong Observation Typicality For most quantum states in a high-dimensional Hilbert subspace $\mathscr{H}_0$, for any observation E_z, they cannot be distinguished by E_z.

Super-strong observation typicality implies observation typicality; it is false because if $E_z = |\psi\rangle\langle\psi|$ is a 1d projection in the direction of a particular $\psi \in \mathbb{S}(\mathscr{H}_0)$, then $\mathbb{P}_{\Psi=\psi}(Z = z) = 1$, whereas $\mathbb{P}_{\rho_0}(Z = z) = 1/d_0 \approx 0$. The super-strong property requires a uniformity in the class of typical quantum states, while observation typicality is compatible with the possibility that the typicality class depends on the choice of E_z. However, the weaker statement—observation typicality—is sufficient for our purposes. OI of typical quantum states is true for any chosen experiment, so it is true for the actual experiment that we happen to perform at an arbitrary time.

Another remark concerns alternative arguments for OI using other typicality theorems, specifically canonical typicality (Gemmer et al. 2009, Popescu et al. 2006, Goldstein et al. 2006a) and dynamical typicality as stated in (Teufel et al. 2023). We will discuss their connection to distribution typicality and observation typicality elsewhere (Chen and Tumulka 2024). Here we make two brief comments:

1. Suppose we have access to only local observables pertaining to a small subsystem $\mathscr{S}$ of the universe (less than half the number of degrees of freedom in the universe). Then our observations cannot distinguish typical quantum states Ψ_{t_0} of the universe from each other. After all, canonical typicality, applied to a subspace $\mathscr{H}_R$ now not given by the micro-canonical energy shell (as in most discussions of canonical typicality) but by $\mathscr{H}_R = U_t\mathscr{H}_0$, asserts that most $\Psi_t \in \mathscr{H}_R$ have nearly the same reduced density matrix $\rho_{\mathscr{S}} = \mathrm{tr}_{\mathscr{S}^c} |\Psi_t\rangle\langle\Psi_t|$ in $\mathscr{S}$ (where $\mathscr{S}^c$ denotes the complement of $\mathscr{S}$). But the Born distribution of any observable in $\mathscr{S}$ depends on Ψ_t only through $\rho_{\mathscr{S}}$; thus, the Born distribution accessible to observers like us is nearly independent of Ψ_{t_0}.

   In this context, the super-strong form of OI is actually true: for a typical quantum state Ψ_{t_0}, there is no local observation *at all* that can distinguish it from ρ_0. After all, once $\mathrm{tr}_{\mathscr{S}^c} |\Psi_t\rangle\langle\Psi_t| \approx \mathrm{tr}_{\mathscr{S}^c}(U_t\rho_0U_t^\dagger)$, all observables in $\mathscr{S}$ have nearly the same Born distribution in Ψ_t as in $U_t\rho_0U_t^\dagger$.

2. Dynamical typicality means here the statement that most $\Psi_{t_0} \in \mathbb{S}(\mathscr{H}_0)$ have for every t nearly the same Born distribution for all macroscopic observables (and nearly the same as ρ_0). If any scientist could distinguish Ψ_{t_0} from ρ_0, then she or he could publish the result, say at time t, so there would be a macroscopic difference between Ψ_{t_0} and ρ_0, which is impossible by dynamical typicality.

3.4 A Bayesian Analysis

So far, we have argued that no observations can reliably distinguish typical universal quantum states from each other. What about unreliable tests that may nonetheless yield probabilistic information about the quantum state? We can use the previous results to give a Bayesian analysis. Here we provide two quantitative bounds of the probabilistic information obtainable from each observation. It turns out that insofar as the observation is not too improbable, it will not give us substantial probabilistic information about the universal quantum state. Hence, there is a sense that even unreliably, universal quantum states cannot be distinguished.

Suppose the prior probability distribution over quantum states in $\mathscr{H}_0$, $\mathbb{P}(\Psi_{t_0})$, is given by the Past Hypothesis, i.e., by u_0. We can show that, if the observation is not too improbable, the posterior probability distribution is very close to the prior.

Corollary 1. *Suppose $d_0 > 1/\varepsilon^5$ and $\operatorname{tr}(\rho_0 E_z) > \varepsilon$. Consider a prior distribution given by u_0, and let $f(\psi)$ be the density relative to u_0 of the posterior distribution obtained by Bayesian updating, given that $Z = z$. Then*

$$1 - \varepsilon < f(\psi) < 1 + \varepsilon \tag{9}$$

for $(1-\varepsilon)$-most $\psi \in \mathbb{S}(\mathscr{H}_0)$.

This follows from Theorem 1 via a short calculation.[3] We can also reformulate the result in terms of the probability of any set $S \subseteq \mathbb{S}(\mathscr{H}_0)$ instead of the density function f:

Corollary 2. *Suppose $d_0 > 1/\varepsilon^5$ and $\operatorname{tr}(\rho_0 E_z) > \varepsilon$. For any subset $S \subseteq \mathbb{S}(\mathscr{H}_0)$:*

$$\mathbb{P}(\Psi_{t_0} \in S \,|\, Z = z) \in \Big[\mathbb{P}(\Psi_{t_0} \in S) - 2\varepsilon, \mathbb{P}(\Psi_{t_0} \in S) + 3\varepsilon\Big] \tag{10}$$

This follows from Corollary 1 via standard arguments of probability and integration theory.[4]

The assumption $d_0 > \varepsilon^{-5}$ is reasonable; even when $\varepsilon = 10^{-200}$, the condition is still satisfied as $\varepsilon^{-5} = 10^{1000} \leq 10^{10^{20}} \leq d_0$. In this case, the posterior probability distribution differs minimally from the prior and is almost independent of any observation. This conclusion requires the assumption that the observation is not too improbable.

Hence, distribution typicality entails not only OI, but also that, in general, observations reveal very little about the initial quantum state. Notice how this fact fails in classical mechanics. Observations, even when they are probable, can reveal a great deal about the classical initial state. For example, observing the Tails outcome of a classical coin-flip experiment can rule out half the initial conditions, namely those that predict a Heads outcome. Observations cannot reveal so much in quantum mechanics because of distribution typicality and observation typicality, two facts whose analogues fail in classical statistical mechanics.

4 Comparisons with Known Results

In this section, we briefly compare our results for OI with some known results in philosophy of science and foundations of physics.

In quantum mechanics, there are several results about limitation to knowledge about the quantum state. For a survey, see (Tumulka 2022b, ch.5). Here we highlight

[3]Proof: From $\mathbb{P}(\Psi \in d\psi) = u_0(d\psi)$ for every infinitesimal set $d\psi \subset \mathbb{S}(\mathscr{H}_0)$ and $\mathbb{P}(Z = z|\Psi) = \langle\Psi|E_z|\Psi\rangle$, we find that $\mathbb{P}(\Psi \in d\psi, Z = z) = u_0(d\psi)\langle\psi|E_z|\psi\rangle$ and $\mathbb{P}(Z = z) = \int_{\mathbb{S}(\mathscr{H}_0)} u_0(d\phi)\langle\phi|E_z|\phi\rangle = \operatorname{tr}(\rho_0 E_z)$. Therefore, $u_0(d\psi)\, f(\psi) = \mathbb{P}(\Psi \in d\psi|Z = z) = u_0(d\psi)\langle\psi|E_z|\psi\rangle/\operatorname{tr}(\rho_0 E_z)$. By Theorem 1 for $(1-\varepsilon)$-most ψ, $\big|\langle\psi|E_z|\psi\rangle - \operatorname{tr}(\rho_0 E_z)\big| < \varepsilon^2$, and thus $\big|\langle\psi|E_z|\psi\rangle/\operatorname{tr}(\rho_0 E_z) - 1\big| < \varepsilon^2/\operatorname{tr}(\rho_0 E_z) < \varepsilon$.

[4]Proof: Let $M \subseteq \mathbb{S}(\mathscr{H}_0)$ be the set where (9) holds. By Corollary 1, $u_0(M) \geq 1 - \varepsilon$. Thus, writing $M^c = \mathbb{S}(\mathscr{H}_0) \setminus M$ for the complement of M and using (9), $\mathbb{P}(\Psi \in S|Z = z) = \int_S u_0(d\psi)\, f(\psi) \geq \int_{S\cap M} u_0(d\psi)\, f(\psi) \geq \int_{S\cap M} u_0(d\psi)\,(1-\varepsilon) = (1-\varepsilon)\, u_0(S \cap M) \geq (1-\varepsilon)\,(u_0(S) - u_0(M^c)) \geq (1-\varepsilon)\,(u_0(S) - \varepsilon) \geq u_0(S) - 2\varepsilon$. On the other hand, since f is normalized, $\int_{\mathbb{S}(\mathscr{H}_0)} u_0(d\psi)\, f(\psi) = 1$, we have that $\int_{M^c} u_0(d\psi)\, f(\psi) = 1 - \int_M u_0(d\psi)\, f(\psi) \leq 1 - \int_M u_0(d\psi)\,(1-\varepsilon) = 1 - (1-\varepsilon)\, u_0(M) \leq 1 - (1-\varepsilon)^2 < 2\varepsilon$. Thus, using (9) again, $\mathbb{P}(\Psi \in S|Z = z) = \int_S u_0(d\psi)\, f(\psi) = \int_{S\cap M} u_0(d\psi)\, f(\psi) + \int_{S\cap M^c} u_0(d\psi)\, f(\psi) \leq \int_{S\cap M} u_0(d\psi)\,(1+\varepsilon) + \int_{M^c} u_0(d\psi)\, f(\psi) < (1+\varepsilon)\, u_0(S \cap M) + 2\varepsilon \leq (1+\varepsilon) u_0(S) + 2\varepsilon \leq u_0(S) + 3\varepsilon$.

a few:

- Because of the Main Theorem about POVMs, we cannot measure the wave function of a given system. Relatedly, because of the no-cloning theorem, we cannot reliably copy the wave function of a given system.
- It can be shown (Tumulka 2022a) that in any ontological theory of quantum mechanics, it is impossible to measure the ontic state. Thus, whatever the true theory of quantum mechanics may be, there must be facts that cannot be empirically determined.
- Because of the Pusey-Barrett-Rudolph (PBR) theorem (Pusey et al. 2012), two different ensembles of wave functions with the same density matrix are physically distinct. However, they are observationally indistinguishable by all possible observations directly on the ensembles.

The last item also asserts a kind of OI, but differs from our result here in that it provides a condition (equal density matrix) under which two situations (ensembles of wave functions) are indistinguishable, whereas we show here that most wave functions from $\mathscr{H}_0$ are indistinguishable. Moreover, all three items are compatible with substantive learning about the actual quantum state. One may still hope to rule out a significant fraction of quantum states based on some observation. Moreover, a Bayesian can, in general, update their probability distribution based on observations after which their posterior may be highly peaked. In contrast, observation typicality tells us that observations in a typical quantum universe do not rule out more than a tiny fraction of universal quantum states, and (if the observation is not too improbable) do not have a substantial influence on a uniform prior probability distribution about initial quantum states. It stays nearly uniform.

We now compare our results to a much discussed underdetermination result in the philosophy of general relativity (GR). Manchak (2009) showed that for almost every spacetime model of GR, there exists another spacetime that is physically distinct (in the sense of having a different global structure) but observationally indistinguishable from it.[5]

Both Manchak's result and ours are compatible with a number of inductive hypotheses. In Manchak's case, one can assume that the same law of nature (Einstein's equation of GR) applies to all spacetime points. Moreover, Manchak's result is compatible with any set of local conditions (conserved under local isometries). In our case, one can also assume that the same laws of nature (within quantum theory) apply to the entire model, so that we are considering the same Hilbert space $\mathscr{H}$, the same Hamiltonian H, and the same subspace $\mathscr{H}_0$.

There are important differences too. First, Manchak's argument shows the existence of an OI counterpart for almost every model of GR, while our argument shows that almost every model is an OI counterpart to almost every other model of QM. Hence, his result is compatible with the fact that we can observationally rule out a significant

[5]For some philosophical discussions about this result, see Beisbart (2009), Norton (2011), Butterfield (2012), and Cinti and Fano (2021).

fraction of possible universes while ours is not. Second, Manchak's proof uses a "cut-and-paste" method to construct special models of GR that are OI from the original model, while ours uses the probabilistic method to show that OI must hold among typical models of QM with generic features.

Related to the second point, Manchak's argument has been challenged on philosophical grounds that the constructed spacetimes are not "physically reasonable," because some people regard the "cut-and-paste" models as implausible candidates for the actual universe. (For example, see Cinti and Fano (2021).) One may respond, as Manchak does, that "the theorem can have physical relevance even if the particular model constructed in the proof does not" (Manchak 2011, p.7). We do not take a stance on that issue, but merely point out that the objection of "physical unreasonableness" does not even work as a *prima facie* objection against our argument for OI in QM. We do not show that almost any model has at least one specially constructed counterpart. Rather, we show that typical (generic) models of QM in a high-dimensional Hilbert space are OI from each other.

It is interesting to note that GR and QM, two pillars of modern physics, both lead to OI.

5 Philosophical Implications

We highlight some potential implications of observation typicality for philosophy of science. First, observation typicality places severe limitations to knowledge in a quantum universe. In a large universe, no experiment can distinguish the overwhelming majority of universal quantum states. This is an in-principle limitation that cannot be overcome by new technologies.

Second, the flip side is that, concerning the observable properties of the universal quantum state, we may know a lot about them because of the typicality results. For any observation, almost all universal quantum states look the same as the normalized projection ρ_0. Assuming our universe is typical among those that satisfy the PH, even though we do not know which it is exactly, we know there is uniformity in their predictions—they are very nearly the same as the probabilistic prediction of the Wentaculus density matrix. Fixing the PH nearly fixes the probabilistic predictions of typical quantum states.

It provides a new perspective on the predictive power of initial condition laws such as the PH. When the initial condition law delivers a high-dimensional subspace in the full Hilbert space, the precise choice of a typical quantum state makes almost no difference to the probabilistic prediction. For any observation, typical individual quantum states predict with nearly the same probability distribution. This raises an interesting question, which we leave to future work, whether typical choices of the initial condition, in some sense to be defined (e.g. typical subspace), lead to nearly the same probability distributions.

Third, observation typicality may conflict with positivism, the idea that a statement is unscientific or meaningless if it cannot be tested experimentally, and a variable is not well-defined if it cannot be measured. Given the PBR theorem, and the various solutions to the quantum measurement problem, we have good reasons to accept that the universal quantum state is an objective feature of the universe. Due to observation typicality, however, it is undetertermined by observational evidence. This is ironic, to echo a point made by Cowan and Tumulka (2016): positivism is sometimes supported by examples from quantum mechanics, and yet quantum mechanics tells us that the central object in a quantum universe is well-defined, objective, but non-measurable, which undermines positivism.

Finally, our result has potential implications for realism about the quantum state. On realism, it is an open question what the universal quantum state represents and how best to understand its objectivity and reality. (For a survey, see Chen (2019).) Within realism, there are two main approaches. The first is to regard it as a physical thing, like a physical field, that exists as part of the basic building blocks of the universe (Albert 1996, Ney 2021). The second is to regard it as a physical law, like the Hamiltonian function in classical mechanics, that tells physical things (such as particles and other fields) how to move (Goldstein and Zanghì 2013).

If the quantum state is a physical thing, OI suggests an in-principle limit to how precisely it can be measured: we cannot gain substantive information about which quantum state is the actual one. Physical objects in classical mechanics do not have this character. They can be measured to arbitrary precision. Almost all quantum states of the universe will appear the same to us. This should be challenging to the empiricists who hold out hope that the basic building blocks of the universe should be in-principle measurable.

Perhaps the result is less surprising if *the universal quantum state has the status of a physical law*. Since it is underdetermined by observational evidence, we have to use theoretical virtues, such as simplicity and elegance, to pin it down. The theoretical flavor of how we choose the universal quantum state is similar to how we choose other physical laws, whose precise forms are constrained but still vastly underdetermined by observational evidence.

Realism about the universal quantum state also leaves open whether it is necessarily a pure state represented by a wave function, or possibly an impure ("mixed") one represented by a density matrix. Following Chen (2021, 2024b), we call the first option *wave function realism* (WFR) and the second *density matrix realism* (DMR). It is already known that WFR and DMR are empirically equivalent, given appropriate choices of the universal quantum states. Observation typicality suggests that it is even harder to distinguish the two theories. Not only are typical individual wave functions OI from each other, but they are also OI from the normalized projection, an impure density matrix. One may regard it as a potential reason to prefer DMR over WFR. If typical pure states Ψ_{t_0} compatible with the Past Hypothesis do not lead to different predictions, why not just use the Wentaculus density matrix ρ_0 instead? After all, ρ_0 may be regarded as much simpler than any typical Ψ_{t_0} even though they are nearly predictively equivalent.

Here we do not take a firm stance on these implications, merely exploring how they may be suggested by our results.

6 Conclusion

We have shown that the central object in quantum theory—the quantum state of the universe—is most likely hidden from observation, because typical quantum states in a high-dimensional Hilbert subspace $\mathscr{H}_0$ are observationally indistinguishable from the density matrix $\rho_0 = P_0/d_0$ and thus from each other. In fact, no observation will yield any substantial information at all, even partial or probabilistic, about Ψ_{t_0}. That is in a sense the strongest known result about a limitation to knowledge in a quantum universe. Assuming the universe we inhabit is typical, our observational data alone will tell us very little about exactly which one it is that we inhabit. Nature is more secretive than we have realized.

Acknowledgment. We thank Sheldon Goldstein for extensive feedback and for suggesting the term "observation typicality." For helpful discussions, we also thank Emily Adlam, David Albert, Fabio Anza, Patricio Avila Cardenas, Jeff Barrett, Jeremy Butterfield, Craig Callender, Branden Fitelson, Hans Halvorson, Alan Hájek, Tyler Hildebrand, Christopher Jarzynski, Joel Lebowitz, JB Manchak, Kerry McKenzie, Marcos Rigol, Daniel Rubio, Tony Short, Stefan Teufel, Isaac Wilhelm, Nicole Yunger Halpern, the UC San Diego Philosophy of Physics Group, and the participants at the 2024 Quantum Thermodynamics Workshop at the University of Maryland, College Park.

Funding. EKC is supported by Grant 63209 from the John Templeton Foundation. The opinions expressed in this publication are those of the authors and do not necessarily reflect the views of the John Templeton Foundation.

Conflict of interests. The authors declare no conflict of interest.

References

Albert, D. Z. (1996). Elementary quantum metaphysics. Pages 277–284 in Cushing, J. T., Fine, A., and Goldstein, S. (ed.s), *Bohmian Mechanics and Quantum Theory: An Appraisal*. Dordrecht: Kluwer Academic Publishers.

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

Beisbart, C. (2012). Can we justifiably assume the cosmological principle in order to break model underdetermination in cosmology? *Journal for General Philosophy of Science*, 40:175–205.

Butterfield, J. (2012). Underdetermination in cosmology: An invitation. *Aristotelian Society Supplementary Volume*, 86(1):1–18.

Chen, E. K. (2021). Quantum mechanics in a time-asymmetric universe: On the nature of the initial quantum state. *The British Journal for the Philosophy of Science*, doi.org/10.1093/bjps/axy068.

Chen, E. K. (2019). Realism about the wave function. *Philosophy Compass*, 14(7).

Chen, E. K. (2024). The Wentaculus: Density Matrix Realism Meets the Arrow of Time

In Bassi, A., Goldstein, S., Tumulka, R., and Zanghì, N. (ed.s), *Physics and the Nature of Reality: Essays in Memory of Detlef Dürr*. Springer.

Chen, E. K. (2024). Density Matrix Realism. In Cuffaro, M. E., and Hartmann, S. (ed.s), *The Open Systems View: Physics, Metaphysics and Methodology*, Oxford University Press, forthcoming, `http://arxiv.org/abs/2405.01025`

Chen, E. K. and Tumulka, R. (2024). Observation Typicality. In preparation.

Cinti, E. and Fano, V. (2021). Careful with those scissors, Eugene! Against the observational indistinguishability of spacetimes. *Studies in History and Philosophy of Science Part A*, 89:103–113.

Cowan, C. W. and Tumulka, R. (2016). Epistemology of wave function collapse in quantum physics. *The British Journal for the Philosophy of Science*, 67(2):405–434.

Dürr, D., Goldstein, S., and Zanghì, N. (1992). Quantum equilibrium and the origin of absolute uncertainty. *Journal of Statistical Physics*, 67(5-6):843–907.

Dürr, D., Goldstein, S., and Zanghì, N. (2004). Quantum equilibrium and the role of operators as observables in quantum theory. *Journal of Statistical Physics*, 116:959–1055.

Dürr, D. and Teufel, S. (2009). *Bohmian Mechanics.* Heidelberg: Springer.

Feynman, R. P. (1963, published 1995). *Lectures on Gravitation.* Edited by F. B. Morínigo, W. G. Wagner, and B. Hatfield. Westview Press.

Frigg, R. and Werndl, C. (2024). Philosophy of Statistical Mechanics. In Zalta, E. N. and Nodelman, U. (ed.s), *Stanford Encyclopedia of Philosophy (Spring 2024 Edition)*, published online by Stanford University `http://plato.stanford.edu/archives/spr2024/entries/statphys-statmech/`

Gemmer, J., Michel, M., and Mahler, G. (2009). *Quantum Thermodynamics: Emergence of Thermodynamic Behavior within Composite Quantum Systems (Second Edition)*. Springer.

Ghirardi, G.C. and Bassi, A. (2020). Collapse Theories. In Zalta, E. N. and Nodelman, U. (ed.s), *Stanford Encyclopedia of Philosophy*, published online by Stanford University `http://plato.stanford.edu/entries/qm-collapse`

Goldstein, S., Lebowitz, J. L., Tumulka, R., and Zanghì, N. (2006a). Canonical Typicality. *Physical Review Letters*, 96:050403.

Goldstein, S., Lebowitz, J. L., Tumulka, R., and Zanghì, N. (2006b). On the distribution of the wave function for systems in thermal equilibrium. *Journal of Statistical Physics*, 125:1193–1221.

Goldstein, S., Lebowitz, J. L., Tumulka, R., and Zanghì, N. (2010). Long-Time Behavior of Macroscopic Quantum Systems. *European Physical Journal H*, 35:173–200.

Goldstein, S. (2012). Typicality and notions of probability in physics. Pages 59-71 in Ben-Menahem, Y. and Hemmo, M. (ed.s), *Probability in Physics*. Springer.

Goldstein, S., Lebowitz, J. L., Tumulka, R., and Zanghì, N. (2020). Gibbs and Boltzmann entropy in classical and quantum mechanics. In Allori, V. (ed.), *Statistical Mechanics and Scientific Explanation: Determinism, Indeterminism and Laws of Nature*. Singapore: World Scientific.

Goldstein, S. and Zanghì, N. (2013). Reality and the role of the wave function in quantum theory. Pages 91–109 in Albert, D. and Ney, A. (ed.s), *The Wave Function: Essays on the Metaphysics of Quantum Mechanics*. Oxford University Press.

Hubert, M. (2024). The History of Moral Certainty as the Pre-History of Typicality. Pages 431–450 in Bassi, A., Goldstein, S., Tumulka, R., and Zanghì, N. (ed.s), *Physics and the Nature of Reality: Essays in Memory of Detlef Dürr*. Cham: Springer.

Lazarovici, D. (2023). *Typicality Reasoning in Probability, Physics, and Metaphysics.* Palgrave Macmillan.

Lebowitz, J. L. (1993). Macroscopic laws, microscopic dynamics, time's arrow and Boltzmann's entropy. *Physica A*, 194:1–27.

Manchak, J. B. (2009). Can we know the global structure of spacetime? *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 40(1):53–56.

Manchak, J. B. (2011). What is a physically reasonable space-time? *Philosophy of Science*, 78(3):410–420.

Mori, T., Ikeda, T.N., Kaminishi, E., and Ueda, M. (2018). Thermalization and prethermalization in isolated quantum systems: a theoretical overview. *Journal of Physics B*, 51:112001.

Ney, A. (2021). *The World in the Wave Function: A Metaphysics for Quantum Physics*. Oxford University Press.

Nielsen, M. A. and Chuang, I. L. (2010). *Quantum Computation and Quantum Information, 10th Anniversary Edition*. Cambridge University Press.

Norton, J. D. (2011). Observationally indistinguishable spacetimes: A challenge for any inductivist. Pages 164-176 in Morgan, G. J. (ed), *Philosophy of Science Matters: The Philosophy of Peter Achinstein*. Oxford University Press.

Penrose, R. (1979). Singularities and time-asymmetry. Pages 581–638 in Hawking, S. and Israel, W. (ed.s), *General Relativity: An Einstein Centenary Survey*. Cambridge University Press.

Popescu, S., Short, A. J., and Winter, A. (2006). Entanglement and the foundation of statistical mechanics. *Nature Physics*, 2:754–758.

Pusey, M. F., Barrett, J., and Rudolph, T. (2012). On the reality of the quantum state. *Nature Physics*, 8(6):475–478.

Reichert, P. (2018). *On the Mathematical Foundations of the Statistical Analysis in Physics*. Doctoral dissertation, Ludwig Maximilian University of Munich.

Reimann, P. (2007). Typicality for generalized microcanonical ensembles. *Physical Review Letters*, 99(16):160404.

Tasaki, H. (2016). Typicality of thermal equilibrium and thermalization in isolated macroscopic quantum systems. *Journal of Statistical Physics*, 163:937–997.

Teufel, S., Tumulka, R., and Vogel, C. (2023). Time evolution of typical pure states from a macroscopic Hilbert subspace. *Journal of Statistical Physics*, 190(3):69.

Teufel, S., Tumulka, R., and Vogel, C. (2024). Canonical typicality for other ensembles than micro-canonical. *Annales Henri Poincaré*, online first.

Tumulka, R. (2022a). Limitations to Genuine Measurements in Ontological Models of Quantum Mechanics. *Foundations of Physics*, 52:110.

Tumulka, R. (2022b). *Foundations of Quantum Mechanics*. Springer.

Wilhelm, I. (2022). Typical: A theory of typicality and typicality explanation. *The British Journal for the Philosophy of Science*, 73:2.