

One World Is (Probably) Just as Good as Many

Jeremy Steeger
jsteeger@uw.edu

*Department of Philosophy
University of Washington*

May 24, 2021

Abstract

One of our most sophisticated accounts of objective chance in quantum mechanics involves the Deutsch-Wallace theorem, which uses state-space symmetries to justify agents' use of the Born rule when the quantum state is known. But Wallace (2003, 2012) argues that this theorem requires an Everettian approach to measurement. I find that this argument is unsound. I demonstrate a counter-example by applying the Deutsch-Wallace theorem to the de Broglie-Bohm pilot wave theory.

Contents

1	Introduction	2
2	Gleason's theorem has a problem of coordination	5
2.1	Setting the stage: what is an explanation of chance?	5
2.2	Frequentism is far from fatal	8
2.3	Gleason's theorem and the coordination problem	11
3	A symmetry theorem for pilot wave and many-worlds theories	12
3.1	An interpretation-neutral approach to decoherence	12
3.2	The Hydra, Lewisian, and Bohm-style views	16
3.3	How each view justifies state supervenience	20
3.4	The symmetry theorem	22
4	Discussion	25
A	Proof of the symmetry theorem	28

1 Introduction

Both defenders and critics of Everettian quantum mechanics (EQM) know the so-called “problem of probability” well. Roughly put: Everettians claim that every quantum measurement outcome occurs, so what could it mean for events to be more or less probable? Surely, every event occurs with certainty, and so with probability one! Brown puts it strikingly, imagining that you awake from a dream in which Prof. X managed to toss a quantum coin:

The next day, relieved in the knowledge that there is at most one Prof. X, you recall the moment in the dream when he claimed that the probability of heads for the biased coin was around 0.7; it was before you were aware of the bizarre consequences of tossing the coin. You now find yourself idly wondering *what Prof. X could have meant*. [...] From the God’s-eye perspective, everything that could happen was happening, and there was no uncertainty about the outcome of the tosses. Was Prof. X not talking then about genuine probabilities at all? (2011, p. 9)

On the one hand, this story raises a genuine philosophical puzzle. But on the other, it can easily be used as a sort of incredulous-stare response to EQM. In light of physicist-philosophers making truly wild claims, it is natural to cast about for surface-level signs of contradiction. Probability-talk is an obvious place to start.

Which, unfortunately, puts defenders of EQM on the defensive. On some level, the situation is a bit odd. Everyone encounters similar skeptical hypotheses in childhood: “the sun rises and sets,” but the earth actually does the turning; “the planets wander through the stars,” but the stars are actually quite a bit further away.¹ We all have gone through the (not unchallenging!) process of resolving surface-level conflicts between theory and the immediate grammar of our experience. It seems distinctly uncharitable not to lend EQMers the same effort. But many do not. And so before EQMers can even get around to extolling the virtues of the view, they have to shore it up against charges of incoherence. No one ought to envy the dreams of EQMers, plagued as they are with bad-faith skeptics.

Upon reflection, the EQMer finds that this situation allows for a tempting rhetorical gestalt: perhaps they could show that the many-worlds interpretation makes sense of probability and makes *more* sense of it than non-branching theories. After all, most single-world theorists tacitly adopt frequentism. And any trained philosopher worth their salt knows that the traditional analyses of chance in terms of frequency are doomed to failure. So perhaps a cutting critique of frequentism, coupled with a many-worlds-exclusive derivation of chance, will be enough for the EQMer to pull the rug out from under the skeptic. Once put off-balance, they think, the skeptic will come around.

Enter the Deutsch-Wallace theorem, a derivation of chance values (i.e., the Born rule) from a thin, operational definition of objective probability and basic facts about the quantum state space within the EQM framework. This theorem has been much-celebrated within the philosophical literature—and rightly so! It is a lovely result that deepens our understanding of quantum probability. But couched in the above context, one can see the temptation to argue that this theorem *only* holds within the EQM framework.

Perhaps in this spirit, Wallace, after giving a thorough critique of frequentism, puts the following words into the mouth of his anti-skeptic:

Anti-sceptic: We’re totally used to probability in Everettian contexts. Okay, it might be *philosophically* a bit puzzling, but those puzzles don’t really matter from the point of

¹I borrow the latter analogy from Wallace (2012, pp. 427–428).

view of physics: practically speaking, we’ve got a sufficiently solid grip on probability to do science. In the single-world interpretation, though, we’re worried that the whole idea of probability makes no sense at all. Failure to understand probability in a satisfactory way in the single-world interpretation isn’t just problematic or puzzling: it’s fatal. (2012, p. 246)

Wallace, in his own voice, validates the anti-skeptic’s claim: “as long as probability talk is understood operationally,” he writes, “the Everett interpretation is actually *better off* than non-branching theories in making sense of that talk” (2012, p. 275). In sum, Wallace seems to claim that the Deutsch-Wallace justification of the Born rule holds in a many-worlds approach and not in single-world approaches, and the failure of the latter to achieve anything similar is fatal.

But what if a single-worlorder could make *just as good* use of the Deutsch-Wallace theorem? Then, by Wallace’s lights, this single-worlorder and the EQMer would be on even footing—at least regarding probability.²

At this point, it is probably worth stressing that the anti-skeptic’s claim is about as dramatic as it gets. Other proponents of EQM defend a weakened version of it. Saunders asserts that “nothing comparable has been achieved for any other physical theory of chance” (2010, p. 184). Read affirms that Deutsch and Wallace’s derivation of the Born rule “renders the notion of objective probability less mysterious in EQM than in antecedent physical theories” (2018, p. 5). But I would wager that the weakened form of the claim makes the underlying philosophical question all the more pressing: is it *exclusively* many-worlds that can avail itself of the derivation, or can a single-world approach do so, too?

This paper demonstrates how a proponent of the de Broglie-Bohm pilot wave theory can make excellent use of Wallace’s proof of the Born rule. There are two motives behind this demonstration. First, it shows that (by Wallace’s lights) one world is just as good as many as far as explaining probability goes—so if we wish to find unique advantages for the Everettian, we ought to look for them elsewhere. Second, it affords a template for how various single-world theorists might find value in the Deutsch-Wallace theorem.

For the first purpose, it is useful to assess Wallace’s discussion of a simple, non-decision-theoretic version of his Born-rule proof, which I will call the symmetry theorem (2012, pp. 148–156).³ A key premise in this theorem is *state supervenience*, the assumption that chances supervene on the wave function Ψ . Roughly, Wallace suggests that EQMers ought to justify state supervenience by noting that Ψ captures all that exists. He then asserts that particles’ configurations in pilot wave theories violate the premise by breaking symmetries in Ψ . But the precise values of configurations q in Bohm’s theory are, indeed, “hidden” from agents—at least as a practical limitation on their ability to measure a system (Bohm, 1952a,b). Bohm’s theory does not take an agent to observe precise particle configurations q directly, but rather indirectly and approximately via the particle’s selection of a particular branch of the overall wavefunction Ψ of the system and the measuring device (Barrett, 2019). Thus, an agent only ever has *approximate* knowledge of q before measurement—an epistemic fact that I call *q-ignorance*. Given that chance is “understood operationally” (as detailed in §2.1), *q-ignorance* implies state supervenience. So a Bohm-style theorist can use the symmetry theorem to derive Born-rule chances. As a bonus, I show that the theorem admits a plausible frequentist

²I am not making any claims about how well these theories deal with locality. In fact, I hope this paper emphasizes locality as a deciding factor between interpretations by demonstrating the relative flexibility of the probability problem.

³Something like the symmetry theorem lies at the heart of several other recent derivations of the Born rule, including the self-locating uncertainty approach of Sebens and Carroll (2018) and the “envariance” approach of Zurek (2005, 2009). Both seem to require an Everettian approach. I focus exclusively on Wallace’s derivation to provide one clear account of how to disentangle many-worlds assumptions from the symmetry theorem.

justification of the rules of chance (i.e., Kolmogorov’s rules).⁴ So Wallace’s critique of frequentism fails to cast doubt on our ability to justify probability in a single world, and pilot wave theory provides an example of how to do it.

But is this justification *just as good* as the Everettian’s? I think so, at least on a charitable approach to many-worlds. On the “Hydra” reading, where Ψ is really all that exists, metaphysics alone justifies state supervenience. But the Hydra view requires users of quantum theory to accept speaking falsehoods the majority of the time. Thus, Wallace encourages us to adopt the more charitable “Lewisian” view, which posits individual space-time worms corresponding to different branches (2012, Ch. 7). But just like particle configurations, these worms break symmetry: after a branching event, one branch will contain worms that the other does not. So we need a new justification of supervenience. Wallace appeals to what I call *self-ignorance*, the assumption that an agent cannot reliably know the worm-identity fixing their future path in a branching event (2012, p. 150). I argue that environmental decoherence motivates q -ignorance and self-ignorance equally well, yielding a precise sense in which one world is (probably) just as good as many.

For the second purpose, it helps to contrast my derivation of the symmetry theorem for pilot wave theory with a very different argument for the Born rule’s validity in Bohmian mechanics—namely, that of Dürr, Goldstein, and Zanghì (henceforth DGZ). I do not seek to criticize DGZ’s derivation, but rather to use their setup as a foil to clarify my own approach. Diverging from the approach I take here, DGZ (1992) do not use decoherence to justify q -ignorance. Instead, they seek to explain q -ignorance by deriving it as an additional consequence of a different sort of proof of the Born rule. Very roughly, they argue as follows. They first suppose that an agent’s knowledge of a system’s q must be mediated by their knowledge of the system’s environment. Then they derive that, in a “typical” Bohmian universe, preparing subsystems with the same effective wavefunction (using some fixed states of the environment) must yield precise values of q that obey Born-rule statistics. Thus, they conclude that an agent could do no better than assign q the distribution given by the system’s wave function—even if (per impossibile) they had precise knowledge of the environment’s configurations! Note, however, that this argument asserts that the measure of the “typicality” of a Bohmian universe is, itself, given by the Born rule. If one worries about the coherence of distinguishing “typicality” from “chance,” then one might worry that this argument is viciously circular—as, e.g., Valentini (2020) recently has.

I do not wish to take a side in this debate: I will neither vindicate nor contest DGZ’s argument here. But I think it is instructive to note that the symmetry theorem provides an alternative approach that seems to recover much of what both Valentini and DGZ want. This approach proceeds by reversing DGZ’s order of explanation. Rather than attempting to use “typicality” to prove that some environmental mechanism must restrict an agent’s knowledge, it simply takes decoherence to give a plausible physical description of such a mechanism. Positing this mechanism entails q -ignorance (in its domain of applicability). At this point, one can use the symmetry theorem to derive the agreement of actual, precise frequencies of q with Born-rule statistics without distinguishing “typicality” from “chance.” I conclude with a brief discussion of how this argument would go and the degree to which it might help adjudicate the dispute between these latter two stripes of Bohmians.

The paper is structured as follows. Section 2 sets the stage by describing the minimal “operational” definition of probability needed for the subsequent results (§2.1), and it illustrates one sort of frequentist justification of this definition (§2.2). It then motivates the symmetry theorem by noting that it addresses a question that Gleason’s theorem fails to answer (§2.3). Section 3 shows how the Hydra, Lewisian, and Bohm-style views can each use the symmetry theorem. §3.1 introduces the

⁴Note well, however, that only the operational definition is strictly necessary for the single-world derivation of the symmetry theorem!

decoherent histories formalism as a framework to house these views. §3.2 gives a precise statement of each view. §3.3 describes how each view justifies state supervenience, argues that self-ignorance and q -ignorance are epistemically on par, and assesses what goes wrong in Wallace’s argument. §3.4 gives a precise statement of the symmetry theorem that applies to all three views. Section 4 concludes with a brief consideration of how the symmetry theorem might help adjudicate the dispute between DGZ and Valentini.

2 Gleason’s theorem has a problem of coordination

2.1 Setting the stage: what is an explanation of chance?

It may be useful to start by reviewing just what Deutsch and Wallace’s symmetry theorem is doing and why it is so compelling. The symmetry theorem affords one sort of explanation of chance values, i.e., the objective probability values associated with some physical system. This explanation comes in three components: a theory of physical states, a theory of chances, and some link between the two. We intend this link to provide a telling answer to the following sort of why-question: given that the physical state of a system is χ , why should we assign measurement outcomes the chances ch , rather than others? We take a relevant answer to show how ch depends on what χ represents about the system in question.⁵

Our physical theory of interest is non-relativistic quantum mechanics (NRQM). NRQM specifies kinematics, dynamics, and observables for microscopic systems. We begin in the usual way, by associating a unit vector in a Hilbert space \mathcal{H}_S with our system of interest, S . S could include an electron’s spin, location, and any other observables that we intend to measure. We represent the preparation and measurement of such systems with projections onto linear subspaces of \mathcal{H}_S ; let $\mathcal{P}(\mathcal{H}_S)$ denote the set of these projections. Recall that the spectra of the familiar self-adjoint observables are associated with projection-valued measures (PVMs)—and thereby with σ -algebras of projections $\{P_i\} \subseteq \mathcal{P}(\mathcal{H}_S)$ —via the spectral theorem.⁶

We allow the system to be open, i.e., possibly coupled with its environment, E . E might include air, dust, photons, measurement devices, and so on; associate all these with a Hilbert space \mathcal{H}_E . We suppose that the S and E together describe a closed system, i.e., a system that is not subject to any other influences from *its* environment. This total system is described by a unit vector $\Psi \in \mathcal{H}_{SE} = \mathcal{H}_S \otimes \mathcal{H}_E$.

I ought to offer a brief caveat regarding the scope of NRQM before proceeding. I agree with Wallace (2020) that it is most perspicuous to view NRQM as a non-cosmological theory. In other words, I do not wish to demand that some wave function Ψ describes the entire universe, regardless of whether I adopt a Bohmian or an Everettian attitude towards measurement. My approach to the de Broglie-Bohm theory differs on this point from the approaches of both DGZ (1992) and Valentini (2019). In §4, I will discuss how the symmetry theorem might still interest these cosmological Bohmians. Until then, however, the reader should take Ψ to denote the wavefunction of a closed system that need not be the universe.

Possible trajectories are given by the action of unitary maps U_t on unit vectors of \mathcal{H}_{SE} , while

⁵I use van Fraassen’s (1980, pp. 134–157) pragmatic model as a framework, but I do not require this model to give the final word on explanation.

⁶Recall that a *projection-valued measure*, for some measurable space (X, Σ) and some subset $\{P_i\}$ of $\mathcal{P}(\mathcal{H}_S)$, is a map $E : \Sigma \rightarrow \{P_i\}$ that is non-negative, normalized, and countably additive. For ease of exposition, I will not be treating positive-operator-valued measurements (POVMs) in this paper. But note that we can recover all POVMs as PVMs on closed systems via Neumark’s theorem. See Busch et al. (1995) for a comprehensive discussion of these concepts.

the Schrödinger equation,

$$i\hbar \frac{d}{dt} \Psi(t) = H\Psi(t), \quad (1)$$

for some Hamiltonian H acting on the system, picks out which of these trajectories are dynamical (once appropriate initial and boundary conditions are specified).⁷ At any time t , the state ρ of the open system S is given by tracing out the environmental degrees of freedom; that is,

$$\rho(t) = \text{Tr}_E |\Psi(t)\rangle\langle\Psi(t)|, \quad (2)$$

where $\rho(t)$ is a density operator on \mathcal{H}_S . In the special case that S is not coupled with E , then $\rho(t)$ will equal $|\psi\rangle\langle\psi|$ for some wave function ψ in \mathcal{H}_S . Before §4, I always use the upper-case Ψ to refer to a closed, total system that need not be the universe and the lower-case ψ to refer to an open system decoupled from its environment.

Note that the partial trace specifies the relationship between various physical degrees of freedom at a given time. NRQM also specifies a temporal system-subsystem relationship: the projection postulate. According to this postulate, when a total system Ψ yields an outcome P_i at time t (via a “projective measurement”), the system at t is effectively described by $P_i\Psi$ (appropriately normalized). This postulate is crucial for ensuring that we can reliably prepare states within a certain range. However, it does not entail anything so strong as “the collapse of the wave function.” We will add to this skeleton to recover many-worlds and pilot wave theories that all deny this collapse (see §3.1). As such, we say for now that the system’s dynamical state includes *at least* Ψ , which evolves according to equation (1).

What about the chances of measurement outcomes for that system? It is widely agreed that states in a formal theory of chance ought to be functions that satisfy some (usually set-theoretic, logical, or algebraic) formalization of Kolmogorov’s three axioms—namely, non-negativity, normality, and countable additivity. Various formalizations add surprising complications. The usual set-theoretic approach, for instance, turns out to be too strict for our purposes. So I will stick to a simple and general algebraic approach. Let Σ be a (possibly partial) σ -algebra with top and bottom elements \top and \perp .⁸ Elements of Σ represent events or utterances that given events occur. A chance function $ch : \Sigma \rightarrow \mathbb{R}$ from the algebra of events to the real numbers must satisfy

$$ch(e) \geq 0, \quad (3)$$

$$ch(\top) = 1, \text{ and} \quad (4)$$

$$ch\left(\bigvee_i e_i\right) = \sum_i ch(e_i) \text{ when } e_i \leq \neg e_j \text{ for } i \neq j, \quad (5)$$

where the later condition holds only when the argument of the function is defined. Probability theory, like quantum theory, also has a system-subsystem relation—namely, conditional probability, the definition of which allows for modifications to the algebra Σ of events. For instance, suppose we wish to restrict our attention to only those elements of Σ smaller than or equal to e (informally, events that occur given that e occurs); these elements form a subalgebra, $\downarrow e$. Via the usual definition of conditional probability, the state ch on Σ yields the following state on the subsystem $\downarrow e$:

$$ch(f|e) := \frac{ch(f \wedge e)}{ch(e)}. \quad (6)$$

⁷Nothing hinges on the choice of the Schrödinger picture, here; the aim is merely to get one well-defined notion of kinematics and dynamics on the table.

⁸Recall that a partial σ -algebra is a partial complemented lattice, i.e., a lattice with partial operations \vee and \wedge denote the least upper bound and greatest lower bound of a set of elements, respectively, and the operation \neg , which denotes the complement of an element.

This way of thinking about conditional probability will prove instructive: to answer our why-question, we will strive to link the system-subsystem relations in both NRQM and probability theory.

Now we can begin to fill in the variables in our why-question: given that the physical state of a closed system is χ , where $\Psi \in \chi$, why should we assign measurement outcomes $P_i \in \mathcal{P}(\mathcal{H}_{SE})$ the chances given by the Born rule, i.e.,

$$ch_{\Psi}(P_i) = \langle \Psi, P_i \Psi \rangle, \tag{7}$$

rather than others? As flagged above, a relevant answer shows how ch_{Ψ} depends on what χ represents. To motivate this dependence—and assess the goodness of an answer—it helps to say a bit more about what ch represents.

Many agree that whatever ch represents, it must account for the functional roles that chance-talk plays in our day-to-day and scientific reasoning. EQMers, in particular, identify two roles that chances must recover:

1. **the inferential link**, i.e., the chance of an event is measured (roughly) by (actual) relative frequencies of that event; and
2. **the credential link**, i.e., all else being equal, one’s subjective degree of belief or credence in an event ought to equal the chance of that event,

where we suppose that repeatable processes yield chancy events (making good sense of relative frequencies). Note well that the “ought” in the credential link roughly implies “can.” We assume that we can roughly measure chance values and use the results of such measurements to make predictions. One can concoct “chances” that are less accessible to agents—but these would fail to capture the function of “chances” in *scientific* reasoning.

Papineau (1996) introduces these roles under slightly different names, and Saunders (2010), Brown (2011), Wallace (2012), and Read (2018) all endorse them. Wallace assumes that the credential link benefits an agent’s pragmatic aims, e.g., their desire to avoid losing money in bets. Thus, he calls it the “decision-theoretic link.” However, as Brown (2011) notes, the link with credence need not be spelled out in terms of pragmatic decision-making—it could, instead, be a matter of epistemic (truth-seeking) aims. Additionally, as Saunders (2010) notes, there is another notable role that chance-talk plays in our discourse—namely,

3. **the link with uncertainty**, i.e., chance events, prior to their occurrence, are uncertain.

However, it is not clear that this last link is *essential* to chance-talk. So, assuming that the inferential and credential links capture the essential bits of this talk, we may try to define chance as the thing satisfying them.

Wallace follows this strategy. He formalizes an agent’s credences with a function cr satisfying the probability axioms (3)–(6). Then he defines chance in terms of credence using Lewis’s (1980) principal principle (PP).

Principal principle (PP). ch is a chance function iff for any event e , if the theory t together with admissible background information b entails that $ch(e) = x$, then an agent ought to set their credence as $cr(e | b \wedge t) = x$.⁹

⁹I borrow this statement of Lewis’s PP from Wallace (2012, p. 141). PP is a specific formalization of the intuitions described by the credential and inferential links. The links themselves are ambiguous between Lewis’s formalization and, e.g., those of Hall (2004) and Ismael (2008). The differences among various formalizations are relevant for chance functions that are self-undermining, i.e., that do not assign $ch(t) = 1$ (Pettigrew, 2012). But on the operational approach, we assume that chances are not self-undermining.

On the operational approach, it is crucial that b only includes information that agents can reliably access. As long as it does, agents can increase their credence in the right theory t by updating. Suppose our agent notes the number of times that e occurs times a large number of trials N all satisfying b . Let e_m represent them seeing e a total of m times. By PP and assuming each e is independent, their prior credence $cr(e_m | b \wedge t)$ is given by the binomial distribution for $ch(e)$ —which is well-approximated by a sharply-peaked Gaussian centered on the chance value. After seeing M occurrences of e , our agent updates using Bayes’s theorem

$$cr(t | e_M \wedge b) = \frac{cr(e_M | b \wedge t)cr(t | b)}{cr(e_M | b)}, \quad (8)$$

which, recall, is a consequence of (6). Thus, their new credence in t will be very low unless t and b imply that $ch(e) \approx M/N$. Relative frequency thereby (roughly) measures chance. And in the idealized case where an agent has full credence in t , they will update to set their credence in e equal to the appropriate chance value given b . So if something plays the role of chance in PP (so defined!), it satisfies both the inferential and credential links. PP (once combined with Bayesian updating) thereby elucidates the standard scientific use of “chance.”

This much—equations (3)–(6) and PP—suffices to establish a thin, operational definition of chance. Note well that the non-decision-theoretic version of Deutsch-Wallace that I aim to generalize assumes only this operational definition. I will argue that single-world theorists can use this version of Deutsch-Wallace, too, provided they assume at least this definition (although they may assume more).

The bare operational definition has a rather glaring weakness: it posits Kolmogorov’s rules (3)–(6) by fiat, and it does little to elucidate why these rules have anything to do with chance. Here, decision theory or relative frequencies may enter the story. But nothing in a single-world approach requires a commitment to any sort of frequentism. It just happens to be the most popular single-world theory of chance on hand. So while Wallace (2012, Ch. 4) seeks to motivate the anti-skeptic’s view by offering a lengthy critique of frequentism, it is not clear that the single-world theorist needs anything so contentious to motivate Kolmogorov’s rules.

But it is worth taking a moment to talk about the hypothetical frequentist’s explanation of these rules for at least two reasons. First, it is useful to expand on the operational definition. With a justification of Kolmogorov’s rules in hand, we can better assess how well a single-world symmetry theorem answers our why-question. And I think that frequentist tools afford one expedient justification of those rules. Second, we ought to assess Wallace’s harsh words for the frequentist. Do frequentism’s flaws motivate us to look away from single-world approaches to probability? If not, then the argument for EQM-exclusivity already stands on shaky dialectical ground.

2.2 Frequentism is far from fatal

It turns out that the tools of hypothetical frequentism work wonders, so long as we put them to their proper use. Suppose that, following Wallace, we adopt the operational definition of chance in terms of PP as our starting point. I claim that we can still use the infinite limit of relative frequencies to explain why PP’s chance function follows Kolmogorov’s rules. To do so, we just recast the hypothetical frequentist’s infinite limit as an approximation.

My explanation stems from a different why-question you may have posed after reading the previous section: why *ought* an agent set their credence equal to chance? Let us follow Brown’s lead and take a relevant answer to invoke an agent’s epistemic aim to have true beliefs—specifically, to have beliefs for the values of relative frequencies that are close to their actual values.¹⁰ Now

¹⁰I follow van Fraassen’s lead in not requiring such an answer to provide *the* reason for the “ought” claim in PP.

note that we may infer from experience that relative frequencies are well-modeled by infinite limits. Specifically, letting $k_n(e)$ denote the number of times that e occurs in n trials, define the approximate description $ch(e)$ of the relative frequency of e to be

$$ch(e) := \lim_{n \rightarrow \infty} \frac{k_n(e)}{n}, \quad (9)$$

where we stipulate that the limits exist for all elements e of Σ . We must also suppose that we can minimally read off from $e_i \leq \neg e_j$ that when e_i occurs, e_j does not, and from $e_i \wedge e_j$ that both e_i and e_j occur. This second assumption is not quite as innocuous as it may seem (as we will see in §3.1), but single-world theorists happily endorse it. It then follows from (9) that ch must obey Kolmogorov’s three rules and the definition of conditional probability. Because we know that the approximation is a good one, setting $cr(e) = ch(e)$ satisfies our epistemic aim to have true beliefs. Barring knowledge of the future, it is not clear how we could do better. Thus, PP’s claim that agents *ought* to set credence equal to chance follows.

Note well that limiting frequency, on this view, is *not* an idealization. Here, I make use of Norton’s (2012) distinction between approximation and idealization, where the former is a partial description of some target system, and the latter is some entirely novel (and possibly fictional) system that bears a crucial analogy with the target. Norton helpfully illustrates the difference by elongating a unit sphere into an ellipsoid with semimajor axis a . In the infinite limit, this ellipsoid becomes an infinitely-long cylinder—a new system that may or may not serve as a useful idealization. The ellipsoid’s volume is $4\pi a/3$, and as it elongates, its surface area gets arbitrarily close to $\pi^2 a$. Thus, in the infinite limit, the ratio of surface area to volume is given by $3\pi/4$ —and this limiting value is an excellent approximation of the ratio for ellipsoids with large a . But can we reason the other way, from the infinite cylinder back to an approximate description of the ellipsoid? Certainly not! The surface-area-to-volume ratio of the infinite cylinder is ambiguous. We could obtain the infinite cylinder by elongating a finite cylinder—and if we take this process to define the surface-area-to-volume ratio, we get a value of 2 instead! The limit system, in this case, is simply too impoverished to serve as an idealization.

As Norton notes, the same thing happens when we take the Boltzmann-Grad limit. This limit generates the Boltzmann equation, which approximately governs the time evolution of the distribution of particles in an ideal gas. Supposing that the gas consists of n hard spheres of diameter d , the approximation is a good one. But it only holds if we assume (a) that the density of particles is low enough to treat them independently and (b) that a typical particle undergoes on the order of one collision per unit of time. Another way of putting (b) is that the mean free path of a particle, given approximately by $\lambda = 1/(2\pi n d^2)$, should be of order one. By this reasoning, the Boltzmann equation holds in the Boltzmann-Grad limit, where we take n to infinity and d to zero in such a way that $n d^2$ remains a constant of order one (Lanford III, 1981, p. 72). But it would be a fatal mistake to conclude that this limit yields an idealization. In the derivation of Boltzmann’s equation, we assume that d is non-zero to determine the state of particles after collisions. If d were equal to 0, then there would be no preferred plane of collision—and so we could not determine the post-collision state (Norton, 2012, p. 219). A literal infinity of point particles does not support Boltzmann’s equation as an approximate description.

Similarly, it is far from clear that an infinite set of trials supports the limit in equation (9) as an approximate description of some finite sequence. The question of whether or not it does is, charitably, one for future research.

I take Wallace’s decision-theoretic explanation to offer an account that complements, rather than challenges, my frequentist one.

At this stage, the unsympathetic reader will have many reasonable concerns. The literature arguing against hypothetical frequentism is vast and nuanced, and I cannot do it full justice within the present paper’s scope. I will instead focus on how my approach defuses Wallace’s three main criticisms.

First, Wallace argues that the frequentist should not invoke arbitrarily long sequences. “The die is not going to be thrown arbitrarily many times,” he writes; “even if it were, it will have abraded away to nothing long before the quintillionth throw” (2012, p. 123). But an infinite approximation can usefully hold for large n below a threshold. We do not demand the Boltzmann equation to describe a gas of, say, a centillion particles—a gas with far more particles than exist in the known universe, by many orders of magnitude!

Second, Wallace cites the oft-repeated argument that the Law of Large Numbers (LLN) shows that hypothetical frequentism is circular: “if probability *is* limiting relative frequency, what can it possibly mean to say that the long-run relative frequency approaches the probabilities with high probability?” (2012, pp. 123–124) I stipulate that the approximation in equation (9) is a good one; I do not argue for this fact with LLN or with any other theorem. I would nonetheless like for chance functions defined by equation (9) to be self-consistent. I would like them to assign values close to one for relative frequencies of e in n trials close to the chance value for e . And LLN guarantees precisely this, for large n !

Third, Wallace argues that using an actual infinity of trials is inadequate, at least because “it is a well-known result in analysis that this limit depends on the order in which the infinite set is arranged,” and probabilities should not depend on the ordering of events (2012, p. 124). I think that this criticism is useful insofar as it provides a reason to be skeptical of infinite idealizations like von Mises’s collectives. Collectives are infinite sets of trials where all limits of the form (9) exist and which satisfy an axiom of randomness, which roughly enforces the invariance of these limits under any recursively-specifiable re-ordering. Just as in the case of the infinite cylinder and the point-particle gas, a collective is an infinite system with a property that strictly lies about the target—for the limits can’t remain invariant under *all* re-orderings. But that is fine for my view! I have no stake in the success of collectives. I only wish to use (9) as an approximation to provide one explanation for Kolmogorov’s rules.

The preceding ought to provide at least preliminary reasons to be skeptical of the usual arguments against frequentism. More pressingly, it shows a deep tension in Wallace’s framing of the dialectic. My treatment of the frequentist’s limits as approximations is no different from how physicists approach other infinite limits. In particular, my treatment mirrors Norton (2012) and Lanford’s (1981) approach to the Boltzmann-Grad limit. Given that our overall project aims to provide a naturalistic explanation of chance values, it would be odd to reject this approach to limits outright. Von Mises makes a similar point by drawing an analogy with continuum mechanics. “Nobody will deny the utility and theoretical importance of the abstraction underlying the concept of a material point,” he writes, “and this despite the fact that we now have theories of mechanics which are not based on the consideration of discrete points” (von Mises, 1981, p. 82). So what justifies changing our epistemic standards just for theories of chance?

The reader already sympathetic to frequentism will likely have a different concern. Many self-avowed frequentists cite Gleason’s theorem as providing all the explanation that quantum probabilities need. So with this much scaffolding in place, let us assess whether Gleason’s theorem answers our why-question.

2.3 Gleason’s theorem and the coordination problem

Does Gleason’s theorem tell us why we ought to assign chances ch when the physical state of a system includes Ψ ? I argue that it does not because it suffers from a problem of coordination: it identifies the right set of chance states, but it does not favor any particular link between Ψ and ch over any other. Thus, the casual frequentist would be mistaken to think that Gleason’s theorem renders the Deutsch-Wallace argument otiose.

This thought originates in Barnum, Caves, Finkelstein, Fuchs, and Schack’s (2000) infamous response to Deutsch’s (1999) original decision-theoretic argument. “*By assuming that measurements are described by probabilities that are consistent with the Hilbert-space structure of the observables, Gleason’s theorem derives in one shot the state-space structure of quantum mechanics and the probability rule,*” Barnum et al. forcefully claim (2000, p. 1182; emphasis theirs). They are right about the state-space structure but wrong about the probability rule.

Specifically, they are right that Gleason’s theorem pins down chance states’ structure for open systems, i.e., for systems $\rho = \text{Tr}_E|\Psi\rangle\langle\Psi|$. It is easy to see that the Born rule yields functions that must follow Kolmogorov’s rules. In more detail: for a set $\{\Pi_i\}$ of projections that pairwise commute, define $\bigvee_i \Pi_i$ as the projection onto the closed linear subspace spanned by the ranges of the Π_i and define $\bigwedge_i \Pi_i$ as the projection onto the intersection of those ranges, and for each Π , define $\neg\Pi := 1 - \Pi$. Then $\mathcal{P}(\mathcal{H}_S)$ is a (partial) σ -algebra, and it is straightforward to check that the Born rule maps states in NRQM to functions that satisfy (3)–(5).¹¹ But it is a far subtler matter to verify that *only* Born-rule functions satisfy these rules. Gleason (1957) shows us one way to do it.

Gleason’s theorem. For $\dim(\mathcal{H}_S) \geq 3$, A function $ch : \mathcal{P}(\mathcal{H}_S) \rightarrow \mathbb{R}$ satisfies Kolmogorov’s rules if and only if there exists some density operator ρ such that

$$ch(P_i) = \text{Tr}(\rho P_i). \tag{10}$$

Here, equation (10) states the Born rule for open systems. By this theorem, all and only Born-rule functions yield probability functions. So Gleason’s theorem proves that there is a bijection between chance-states and QM-states for (nearly all) open quantum systems.¹²

But note that this theorem does not favor any particular bijection between QM states and chance states over any other. The necessary and sufficient condition that it identifies for ch to be Kolmogorovian is that there is *some* density operator that yields it via the Born rule. Therefore, Gleason’s theorem does not decide between the Born rule and, say, the Shmorn Rule—a rule which instructs the user of QM to first rotate a unit vector ninety degrees about some one-dimensional subspace before applying the Born rule. In other words, Gleason’s theorem straightforwardly yields the following:

Shmleason’s theorem. For $\dim(\mathcal{H}_S) \geq 3$, a function $ch : \mathcal{P}(\mathcal{H}_S) \rightarrow \mathbb{R}$ satisfies Kolmogorov’s rules if and only if there exists a density operator ρ' such that

$$ch(P_i) := \text{Tr}(U^\dagger \rho' U P_i). \tag{11}$$

for some fixed unitary operator $U \neq I$.

¹¹Note that while $\mathcal{P}(\mathcal{H}_S)$ is a partial complemented lattice, it is *not* a Hilbert lattice; these latter lattices are complete, but non-intuitive from the standpoint of probability. For more technical details on Hilbert lattices, see Rédei (1998); for a critical perspective, see Kochen (2015).

¹²As Busch (2003) has shown, we can remove the parenthetical by attending to POVMs, which generalize PVMs. I will be sticking with PVMs for ease of exposition, but the preceding arguments all generalize naturally to allow for POVMs.

For the proof, let ρ be the state from Gleason’s theorem and define $\rho' := U\rho U^\dagger$. Now note that Shmleason’s theorem “endorses” a different probability rule. Shmleason’s theorem is just as sound as Gleason’s theorem; as of yet, we have no reason to privilege (10) over (11).¹³

Thankfully, Gleason’s theorem invokes very little of the structure of NRQM. In particular, it does not invoke the projection postulate. This postulate is the key ingredient of the symmetry theorem. However, its physical significance has long been a matter of controversy—so we had better justify it if we wish for the symmetry theorem to answer our why-question properly! Following Wallace, I begin my justifications of the projection postulate by positing environmental decoherence.

3 A symmetry theorem for pilot wave and many-worlds theories

In the previous section, I asserted that the symmetry theorem derives the Born rule by invoking the projection postulate. Thus, to answer our why-question, we need to ensure that this postulate has physical meaning. This section demonstrates how both many-worlds and pilot wave theorists can use decoherence to do so. In §3.1, I briefly review how to use decoherence to recover quasi-classical histories in an interpretation-neutral way, closely following the accounts of Wallace (2012) and Rosaler (2016). In §3.4, I use this account to state an interpretation-neutral version of the symmetry theorem.

3.1 An interpretation-neutral approach to decoherence

Recall the decoherence program’s core idea: when a subsystem of interest couples with its environment, coherence among its pointer states leaks into the total system, leaving the subsystem in a mixture of these states. As Schlosshauer (2007) cogently argues, positing just this much on top of the bare quantum theory described above does not solve the measurement problem. Instead, it addresses two closely related issues: why some pointer bases (like Gaussian wave packets or spin eigenstates) seem to be preferred by given observations and why it is so difficult to observe the effects of quantum coherence at macroscopic levels. It resolves the former by specifying a physical mechanism—namely, the system-environment interaction—that picks out the preferred basis. It addresses the latter by positing that macroscopic systems undergo much quicker decoherence than their microscopic cousins. But the program does not try to explain the appearance of *specific* outcomes. This final question is one that the traditional interpretations of quantum theory (many-worlds and pilot wave theories among them) are poised to answer. All three of these elements combined provide one or another physical justification of the projection postulate as a description of repeatable preparations and measurements.

If the preceding is right, then decoherence does not favor any given traditional interpretation. Likewise, any tool that tracks it—including the decoherent histories formalism—cannot carry interpretive commitments. Nonetheless, it is worth reviewing the formalism, as its relationship to interpretations is often unclear. Gell-Mann and Hartle’s (1990) original proposal does not spell out the role of measurement, and their subsequent works take interestingly divergent approaches. For example, while one may read Hartle (2010) as recommending a many-worlds approach, Gell-Mann and Hartle (2012) explicitly endorse just one history as “real.” So, closely following Wallace (2012) and Rosaler (2016), I will briefly detail how I take their histories formalism to accommodate both many-worlds and pilot wave theories.

¹³In (2003, pp. 433–434), Wallace points to this underlying problem when he asserts that were we to try to use Gleason’s theorem to derive the Born rule, we would still need to appeal to the “games” developed in Deutsch’s original proof of the symmetry theorem.

We start with our total, closed system described by some state $\Psi \in \mathcal{H}_{SE}$ at time t_0 . Now suppose that on a very brief timescale τ_D , the interaction between the system and the environment H_{SE} dominates the Hamiltonian in equation (1). Pick an \mathcal{H}_S -spanning set of states $\{\psi_j\}$, the pointer states, that are robust under the action of H_{SE} . More precisely, for some $\Delta t \ll \tau_D$, a system prepared in a pointer state (with some environmental “ready state” E) couples with its environment and evolves as

$$|\psi_j\rangle \otimes |E\rangle \xrightarrow{\Delta t} |\psi_j\rangle \otimes |E_j\rangle, \quad (12)$$

where E_j is some final environmental state. If, instead, the system begins in a superposition of pointer states, then due to the linearity of (1), the total system evolves as

$$(a|\psi_1\rangle + b|\psi_2\rangle) \otimes |E\rangle \xrightarrow{\Delta t} a|\psi_1\rangle \otimes |E_1\rangle + b|\psi_2\rangle \otimes |E_2\rangle. \quad (13)$$

This completes our first step of decoherence, which we suppose ends at t_1 . Applying the partial trace of equation (2), we see that the final state of the system is given by

$$\rho_S(t_1) = |a|^2 |\psi_1\rangle\langle\psi_1| + |b|^2 |\psi_2\rangle\langle\psi_2| + a^* b \langle E_1|E_2\rangle |\psi_2\rangle\langle\psi_1| + b^* a \langle E_2|E_1\rangle |\psi_1\rangle\langle\psi_2|. \quad (14)$$

If the environment interacts strongly enough with the system, then we can suppose that $\langle E_1|E_2\rangle \simeq 0$. For example, take ψ_1 and ψ_2 to be well-separated coherent states of a heavy dust particle S , and take E to be a well-localized air particle scattering strongly off the dust. Then E_1 and E_2 would be roughly orthogonal, well-localized states of the air particle. So we have

$$\rho_S(t_1) \simeq |a|^2 |\psi_1\rangle\langle\psi_1| + |b|^2 |\psi_2\rangle\langle\psi_2|, \quad (15)$$

an approximate mixture of pointer states. In this way, decoherence has (approximately) moved the initial superposition of the system into its environment.

Now let the system continue to evolve to a later time. Once the system is in a mixture of pointer states, the environment has very little effect on its further evolution—so, to a good approximation and for some intermediate time $\Delta t' \gg \tau_D$, the system and the environment evolve independently under their own self-Hamiltonians. Since this dynamical evolution is linear, each of the terms in equation (13) evolves into a new superposition of pointer states tensored with a new ready state of the environment. For example, the evolution might look like

$$a|\psi_1\rangle \otimes |E_1\rangle + b|\psi_2\rangle \otimes |E_2\rangle \xrightarrow{\Delta t'} a(a_1|\psi_1\rangle + a_2|\psi_2\rangle) \otimes |E'_1\rangle + b(b_1|\psi_1\rangle + b_2|\psi_2\rangle) \otimes |E'_2\rangle. \quad (16)$$

Then decoherence will occur again, but for *each* of the terms in the superposition on the right-hand side of equation (16). This process yields the total evolution

$$(a|\psi_1\rangle + b|\psi_2\rangle) \otimes |E\rangle \xrightarrow{\Delta t + \Delta t' + \Delta t} |\psi_1\rangle \otimes (aa_1|E_{1,1}\rangle + bb_1|E_{2,1}\rangle) + |\psi_2\rangle \otimes (aa_2|E_{1,2}\rangle + bb_2|E_{2,2}\rangle), \quad (17)$$

completing our second instance of decoherence (which we suppose ends at t_2). Note that decoherence ensures that either of $E_{1,1}$ or $E_{2,1}$ is approximately orthogonal with either of $E_{1,2}$ or $E_{2,2}$. But we also expect the information recorded by the first decoherence event to be distributed widely throughout the environmental degrees of freedom (so new bits of the environment, e.g., other particles in the air scattering off of our system, are doing the second bit of decoherence). Thus we expect the pairs $E_{1,1}, E_{2,1}$ and $E_{1,2}, E_{2,2}$ to be approximately orthogonal, as well. In this way, decoherence yields a natural branching structure wherein the environment records four different sequences of pointer states (ψ_1 then ψ_1 , ψ_2 then ψ_1 , ψ_1 then ψ_2 , and ψ_2 then ψ_2), and none of these sequences interfere with each other. Naturally, this process may be iterated an arbitrary number of times.




Figure 1: Schematic illustration of a branching history algebra with $n = 3$ steps of decoherence. Events have been labelled for the two histories α and β (note that $\alpha_1 = \beta_1$).

Note well that the set of pointer states might be uncountable and overcomplete. Both Schlosshauer (2007, §2.8, §5.2) and Wallace (2012, §3) argue that a particularly natural choice of pointer states (for a single spinless particle) is the set of coherent states

$$\psi_{(q,p)}(x) = \langle x|q,p\rangle = ae^{-\lambda(x-q)^2} e^{ipx} \quad (18)$$

(for q and p reflecting “position” and “momentum” values ranging over the reals, and where λ denotes the width of the Gaussian wave-packet). For these pointers, an arbitrary initial state $\Psi \in \mathcal{H}_{SE}$ may be written as

$$\Psi(t_0) = \int dq_0 dp_0 c(q_0, p_0) |q_0, p_0\rangle \otimes |E(q_0, p_0)\rangle, \quad (19)$$

where $c(q_0, p_0)$ is the coefficient for the appropriate pointer state and $|E(q_0, p_0)\rangle \in \mathcal{H}_E$ is its associated environment state. To clean up our notation a bit, let us introduce the sequence $A = (A_1, \dots, A_n)$ to represent a sequence of pointer states like $|q_0, p_0\rangle, \dots, |q_n, p_n\rangle$ for n steps of decoherence. Let $c(A)$ be the product of transition amplitudes for the appropriate pointer states—so, for example, in equation (17), we have $c(\psi_1, \psi_2) = aa_2$ and $c(\psi_1, \psi_1) = aa_1$. At t_n , the state in equation (19) has evolved to

$$\Psi(t_n) = \int dq_0 \dots dq_n dp_0 \dots dp_n c(A) |q_n, p_n\rangle \otimes |E(A)\rangle, \quad (20)$$

where $|q_n, p_n\rangle$ is the final state in A and

$$\langle E(A')|E(A)\rangle \simeq 0 \text{ if } A' \neq A. \quad (21)$$

Equation (21) points to the idea that n steps of decoherence result in a number of incompatible environment states, each recording a different history of n pointer states.

Capitalizing on this idea, let us idealize equation (21) such that at each of the n steps, the environmental states are *exactly* orthogonal. Given this idealization, consider all possible pointer states A_i at time t_i and note that their associated projections

$$\alpha_i = |A_i\rangle\langle A_i| \otimes |E(A_1, \dots, A_i)\rangle\langle E(A_1, \dots, A_i)| \quad (22)$$

form a mutually orthogonal set.

To derive the Born rule, it would help to organize these projections within a familiar algebraic structure. Luckily, every set of histories that satisfy the branching criterion may be associated with

a natural σ -algebra of events. First, note that each time-indexed projection α_i is associated with a PVM, which in turn specifies a σ -algebra \mathcal{S}^i of projections on \mathcal{H}_{SE} . So, following Wallace (2012, pp. 95–96), define the *history algebra* $\{\mathcal{S}^i\}$ as the n -fold direct product of such σ -algebras of projections,

$$\{\mathcal{S}^i\} := \mathcal{S}^1 \times \dots \times \mathcal{S}^n, \quad (23)$$

the elements of which are given by complete specifications of histories $\alpha = (\alpha_1, \dots, \alpha_n)$ (for the appropriate initial state and dynamics, and where the algebraic operations are performed pointwise). Our idea now is to characterize when a history algebra witnesses the branching structure described above.

One natural way to do so is to consider the transition weights for projections in atomic coarse-grainings of the algebra. We say that a *coarse-graining* of $\{\mathcal{S}^i\}$ is a history algebra $\{\mathcal{C}^i\}$ where every projection in \mathcal{C}^i is a sum of projections in \mathcal{S}^i (for every i). A coarse-graining is *atomic* if each \mathcal{C}^i is the free σ -algebra generated by some countable set of projections. Now define the transition weight between any two projections, for $t_i < t_j$ as follows:

$$\mathcal{T}(\alpha_i, \beta_j) := \frac{|\beta_j U(t_i, t_j) \alpha_i U(t_0, t_i) \Psi|^2}{|\alpha_i U(t_0, t_i) \Psi|^2}, \quad (24)$$

where $U(t_0, t_i) \Psi = \Psi(t_i)$ and $U(t_i, t_j) \Psi(t_i) = \Psi(t_j)$. We say that a history algebra is *branching* when each projector “receives weight” from just one of its predecessors, which amounts to satisfying the follow condition.

Branching criterion. A history algebra $\{\mathcal{S}^i\}$ is *branching* for Ψ (or Ψ -*branching*) when it admits an atomic coarse-graining $\{\mathcal{C}^i\}$ such that, for any $\alpha, \alpha', \beta \in \mathcal{C}$ and $t_i < t_j$,

$$\text{if } \mathcal{T}(\alpha_i, \alpha'_j) \neq 0 \text{ and } \mathcal{T}(\beta_i, \alpha'_j) \neq 0, \text{ then } \alpha_i = \beta_i. \quad (25)$$

If the above is satisfied for $\{\mathcal{C}^i\} = \{\mathcal{S}^i\}$, then we say that $\{\mathcal{S}^i\}$ is *strictly branching* (or *strictly Ψ -branching*).

Figure 1 gives a heuristic illustration of the branching that results. We need to invoke one final idealization before proceeding: we assume that every kinematic trajectory is dynamically possible. As such, we can avail ourselves of environmental decoherence along any pointer basis for any (coarse-grained) dynamics:

Decoherence availability. Let $\{\mathcal{S}^i\}$ be an atomic history algebra with n steps of decoherence. For any $\Psi(t_0)$ and any set of unitary maps $\{U_1, \dots, U_n\}$, there is an evolution of $\Psi(t_0)$ for which $\{\mathcal{S}^i\}$ is strictly Ψ -branching and such that $U_i \Psi = \Psi(t_i)$ for each i .

The coarse-grained dynamics $\{U_i\}$ idealize the evolution generated by the system’s self-Hamiltonian and the environment’s. In other words, *decoherence availability* captures the idea that the system can evolve in any unitary manner during the time intervals $\Delta t'$ between decoherence events along an arbitrary pointer basis (as described in the account of §3.1).

We can now see how decoherence conditionally justifies the projection postulate. In short, (25) ensures that every branch is dynamically isolated from all the others. So as long as we can associate projective measurements with branches, the projection postulate holds. Thus, each branching event α_i creates a new closed system with the wave function $\Psi' = \alpha_i \Psi(t_i) / |\alpha_i \Psi(t_i)|$. We might loosely call this recursive specification of closed systems “subsystem-recursivity in time”; Figure 2 illustrates this recursivity. We also have what we might roughly call “subsystem-recursivity in space.” At every




Figure 2: Subsystem-recursivity in time shared by both many-worlds and the Bohm-style pilot wave theory.

time step t_i , the projected wave function $\Psi'(t_i) = |A_i\rangle \otimes |E(A_1, \dots, A_i)\rangle$ is separable. Thus, the system’s wave function $|A_i\rangle \in \mathcal{H}_S$ is decoupled from the environment, and it may be promoted to a closed system.¹⁴ Assuming that agents reliably know pointer states, they can reliably know which new closed systems a branching event yields:

Decoherence reliability. Agents gain reliable information about quantum systems through decoherence along known pointer states.

Decoherence reliability justifies our repeated “discovery” of closed systems with specific, explicit wave functions Ψ —so long as we can explain why we see only one of the options in the post-branching superposition (or otherwise state that this question is ill-posed). Thus, to complete the justification of the projection postulate, we turn to how the many-worlds and pilot wave theorists respond to the question of specific outcomes.

3.2 The Hydra, Lewisian, and Bohm-style views

To respond to the question of specific outcomes, many-worlds approaches begin with a simple premise: given a branching history algebra, merely suppose that all histories (with non-zero weight) occur. Each of the histories represents a different version of the initial observer’s experience. However, this premise is ambiguous. It points to several distinct ways of dealing with the identity of branches.

I focus on two of these: the Hydra view and the Lewisian view.¹⁵ On the Hydra view, we suppose that no single history α corresponds to an individual object. The entire branching structure of Ψ is all there is—so the entire physical state of the system is $\chi = \Psi$. The consequence of this view for the semantics of our history algebra is dramatic. Each α_t with a non-zero weight ought to be true, because every event occurs. This semantics, illustrated in Figure 3a, instructs us to ignore the question of specific outcomes: strictly speaking, such outcomes are illusory. But this response is distinctly uncharitable to users of the theory. Suppose you and a friend toss a fair coin to decide who leaves a tip for your waiter: if you claim that the coin did not land heads, clearing

¹⁴Wallace (2019) provides a rigorous, formal theory of subsystem-recursivity that informs my brief and qualitative discussion here.

¹⁵I borrow this terminology from Wallace (2012, p. 281). These views have natural analogs that replace Lewisian worms with Siderian stages—Wallace refers to these as the Disconnected view and the Stage view, respectively (2012, p. 282). Tappenden (2011) has extensively developed the Stage view, and he saliently notes that the view yields post-measurement, pre-observation uncertainty that is quite similar to single-world uncertainty about chancy events. But since the Hydra and Lewisian views suffice to make my point about the symmetry theorem, I will not cover the Disconnected and Stage views here.

you of responsibility, you are, strictly speaking, wrong. By the lights of the Hydra view, users of quantum theory speak falsehoods the vast majority of the time. So the view is not terribly kind to its adherents (although it may, incidentally, turn out to be quite a bit kinder to waiters).

Notably, the Hydra view eliminates uncertainty about future events. Every proposition in $\{\mathcal{S}^i\}$ with non-zero weight is true. As flagged by the story of Prof. X, this view must sever the link between chance and uncertainty flagged in §2.1. On the one hand, the insight that we can sever this link is one of the modern Everettian program’s crowning achievements. On the other, it shows another way in which the Hydra view is uncharitable: agents are, strictly speaking, incorrect to claim that chancy events are uncertain. Given these difficulties, it would be nice to have a more charitable alternative on hand.

One such alternative is the Lewisian view developed by Saunders and Wallace (2008). This view admits branches as bona fide individuals. Every complete history α corresponds to a particular spacetime worm, an individual object to which any observer can refer. Figure 3b gives a schematic illustration, singling out the spacetime worm for α (a complete history with a non-zero weight). So a complete specification of the observer’s physical state would be $\chi = (\Psi, \alpha)$, where α is the identity of their spacetime worm. Crucially, on the Lewisian view, an observer can be uncertain about the identity of their branch. Uncertainty about the future—i.e., about whether α_t or β_t will occur—is recovered as a form of self-locating uncertainty (namely, uncertainty about branch-identity).

There remains a question as to whether worms share their temporal parts before branching. If they do so “overlap,” then there is numerically one worm before branching, and numerically two after. Nevertheless, it may be the case that there are numerically two worms at all times, worms that are qualitatively the same before branching but that “diverge” afterward. Saunders (2010) and Wilson (2012) argue for divergence, but Wallace suggests that the difference ultimately does not matter (2012, pp. 286–287). What is crucial, he asserts, is that “each agent does have a unique future [on the Lewisian view], but it is in principle impossible for him to possess reliable knowledge of that future” (2012, p. 150, fn. 25). In other words, what is crucial is that agents cannot know the identity of their time-extended, four-dimensional worm α with a grain fine enough to fix future facts (regardless of how we divvy up their worm’s temporal parts). I will call this assertion *self-ignorance*.

Self-ignorance. An agent cannot reliably know the identity of the worm α that determines which branch they will take before the branching.

Principles that set limits on reliable knowledge are crucial for our discussion of chance. Recall that (operational) chance can only depend on the information that agents can reliably access. Thus, self-ignorance directly entails that chances cannot depend on α .

One might try to argue that self-ignorance follows from the dynamics. After all, on both the divergent and overlap views, worms are spatially and temporally coincident—and so dynamically identical—before branching. But this observation alone does not imply self-ignorance on either view. For the divergent view, suppose, for example, that my chair and the corresponding chair-wise arrangement of atoms are numerically two objects. These objects are spatially and temporally coincident, and they evolve the same way in time. But I see both, and I am clear on which is which. For the overlap view, suppose one version of me decides to remove my chair’s armrests. Another decides otherwise. There is numerically one chair before the time my second self removes the armrests and two afterward. At the moment I make my decision, I am quite clear on which of the two future chairs is mine—even though there is numerically one chair at that moment! So in either case: why should a lack of dynamical difference before branching limit agents’ knowledge?

The Lewisian might respond by claiming that an agent *only* divines reliable information about a quantum system with an unknown state via branching along known pointer states. It immediately




Figure 3: Schematic illustrations of the many-worlds completions of a branching history algebra on the Hydra and Lewisian views, as well as a sketch of the Bohm-style pilot wave completion.

follows that the projection postulate sets a strict limit on reliable knowledge, i.e., Ψ is the *most* our agent can reliably know:

Decoherence exclusivity. Agents gain reliable information about quantum systems *exclusively* through decoherence along known pointer states, and so Ψ is the most that they can reliably know.

This principle (which implies decoherence reliability, but not decoherence availability) suggests that decoherence accounts for all preparations and measurements an agent could conceivably make. By securing the physics of these processes, it makes a compelling case for self-ignorance. But it is not exclusive to the Lewisian. The Hydra view is also free to endorse decoherence exclusivity. As we will see, a pilot wave approach can use it, too.

The pilot wave theory that I consider in depth adapts Barrett’s (2019) presentation of Bohm’s (1952a, 1952b) theory; henceforth, I will refer to this theory as the Bohm-style pilot wave view (or just “the Bohm-style view”). The Bohm-style view assigns precisely the same semantics to the branching history algebra as the Lewisian view. However, it differs dramatically in its metaphysics. We replace the four-dimensional worms with (point-like) particles whose motion is governed by the wave function. We might view the wave function as a real, physical field or merely as a governing-law. In either case, its physical significance is chiefly dynamical. In turn, only one history is “real” in the sense that only one can (approximately!) describe the particles’ actual trajectories. Still, the wave function contains information regarding all the other possible paths (represented by the dotted lines in Figure 3c).¹⁶

We may describe pilot wave metaphysics with a bit more formalism. Specifically, we add states and dynamics directly describing the particles themselves. Describe a total system of N particles evolving in \mathbb{R}^3 with a point $q = (q_1, \dots, q_N)$ in the configuration space \mathbb{R}^{3N} . For a given time t , Ψ is a spinor-valued function on this space. (Recall that “spinors” are vectors of complex numbers that allow for the description of systems that “have spin”; see Norsen (2014) for a more detailed discussion.) While equation (1) still governs Ψ , it leaves the evolution of the point-particles under-determined.

¹⁶Note well that I am not making any specific ontological claims here about Ψ ! The dotted lines in Figure 3c are meant to represent possible particle paths in (four-dimensional) spacetime, much like the solid lines in Figure 3b are meant to represent worms in (four-dimensional) spacetime. This view of Bohmian particles is compatible with taking Ψ to be an object that lives in an ontic $3N$ -dimensional configuration space. It is compatible with taking Ψ to specify a multi-valued field in spacetime in the style of, e.g., Romano (2020). It is also compatible with taking Ψ to be purely nomic, nothing more than a law governing the motions of particles in spacetime. The reader should feel free to choose whichever view of Ψ seems most natural! Nothing in the subsequent argument will hinge on this choice.

We adopt one natural option, namely, the guiding equation

$$\frac{dq_i}{dt} = \frac{\hbar}{m_i} \operatorname{Im} \frac{\Psi^* \nabla_i \Psi}{\Psi^* \Psi}(q), \quad (26)$$

where $q_i = (x_i, y_i, z_i)$ is the position of the i th particle and $\nabla_i = (\partial/\partial x_i, \partial/\partial y_i, \partial/\partial z_i)$ is the gradient with respect to that position (and the products of spinors in the numerator and the denominator are scalar products). Just as it does for the Schrödinger dynamics, decoherence guarantees that these particle dynamics are always well-defined: the motion of the particles is the same regardless of whether or not we apply the projection postulate to Ψ after decoherence. If we do apply the projection postulate, we get one instance of what DGZ call an *effective wave function*—very roughly, a wave packet in Ψ whose support contains the actual particle configuration and is macroscopically distinct from the supports of the other wave packets that comprise Ψ .¹⁷ The complete state of a closed system, then, is $\chi = (\Psi, q)$, where (1) and (26) specify the dynamics (and where Ψ may be an effective wave function).

Note well, however, that an agent cannot reliably access the particle configurations q in the typical domains where NRQM applies! Barrett (2019) nicely illustrates this fact with a toy Bohmian model of a Stern-Gerlach experiment. In this toy model, an electron e in a superposition of spins travels a path B if its configuration lies in the support of the spin-down wave packet (and it otherwise travels a path A). He then introduces a particle p that acts as an idealized measuring device (or, in the parlance of the decoherent histories formalism, an environmental record): the particle’s wavefunction entangles with the electron’s such that p moves to a region b if and only if e took path B . Of this model, Barrett writes:

Suppose that the recording particle p in the two-path experiment moves to region b . This does not tell us precisely where the electron is. It might be anywhere in the wave packet that traveled path B . Rather, it tells us *which wave packet* the configuration (and hence the electron) is associated with. That is, the empirical content of the record, what one can deduce from the value of the record, is given by *the effective wave function selected by the current particle configuration*. This is what an observer has epistemic access to given her measurement record. In this precise sense, this is what she sees. (2019, pp. 213–214, emphasis Barrett’s)

Thus, the Bohmian agent only ever knows q approximately, at least in domains where this model of measurement applies. And, thankfully, we have already encountered a separate dynamical condition that will *guarantee* that this model applies: decoherence exclusivity. (As noted in the above parenthetical, the environmental record states of the decoherent histories formalism simply generalize Barrett’s recording particle p .) So decoherence exclusivity implies the weaker claim that agents cannot reliably know the precise value of q before measurement. This claim is the one that I call *q-ignorance* in the introduction.

q-ignorance. An agent cannot reliably know a system’s configuration q before measurement.

On my non-cosmological approach, *q-ignorance* has the same scope as self-ignorance. Each claim is epistemic: each has to do with what sort of reliable information an agent can possess. Moreover, they seem to enjoy equally good dynamical justifications. True, there is a dynamical difference between different configurations q before branching, and there is no such dynamical difference

¹⁷See §5 and fn. 20 of DGZ (1992) for precise definitions and a brief discussion of how decoherence yields effective wave functions. For a longer discussion, see Romano (2016).

between different worm-identities α . But the natural story that the Lewisian can tell about why dynamical coincidence limits agents' knowledge—decoherence exclusivity—works just as well to justify why q is inaccessible to the pilot wave theorist (thereby implying q -ignorance). So, by my tally, the two ignorance principles are on par.

Note well that this last step—taking decoherence exclusivity to justify q -ignorance—is where my approach diverges from that of DGZ. Recall that DGZ seek to justify q -ignorance via a different route that considers the actual relative frequencies of precise values of q . Explicitly, they consider linear combinations of delta functions on \mathbb{R}^{3N} , functions which do not figure in my proof. I will discuss how the symmetry theorem interacts with these functions and DGZ's overall approach in §4.

For now, note that each of the two ignorance principles can do the key work that the symmetry theorem needs: each entails that chances must supervene on Ψ . As such, Wallace's argument that no single-world theory can justify this supervenience principle must be unsound. In the next section, I diagnose what goes wrong in Wallace's argument by detailing how the Hydra, Lewisian, and Bohm-style views each obtain state supervenience.

3.3 How each view justifies state supervenience

Here is Wallace's argument that only a many-worlds theory can justify state supervenience, in brief. Suppose, contra the many-worlds hypothesis, that measurement is a single-world event. Then the measurement process is either stochastic or deterministic. If it is stochastic, we could either specify the process or leave it unstated. The latter is question-begging. The former is inadequate because the process must break the symmetries of the deterministic evolution of Ψ . If the process is deterministic, then we ought to justify a distribution over microstates. But then the actual microstate (e.g., for the pilot wave theory, the particle positions q) will break Ψ 's symmetries. In either case, Wallace argues that symmetry-breaking is enough to scuttle the argument that chances ought to supervene on Ψ . The conclusion that only a many-worlds theory can justify state supervenience follows.

Wallace summarizes the moral:

Whether we are considering a stochastic or a deterministic process, the problem is ultimately the same. We are attempting to use a dynamical symmetry between two possible outcomes to argue that the outcomes are equally likely. But since only one outcome actually occurs, something must break the symmetry—be it the actual microconditions of the system, or the actual process that occurs in a stochastic situation. Either way, we have to build probabilistic assumptions into that symmetry-breaking process, and in doing so we effectively abandon the goal of explicating probability. (Wallace, 2012, pp. 147–148)

This objection originates in (Wallace, 2003, §7), where it arises as part of a response to Barnum et al. (2000). In addition to their claim that Gleason's theorem renders Deutsch's argument otiose, Barnum et al. level two other criticisms. First, they note that while Deutsch claims to derive probability using the non-probabilistic part of NRQM and decision theory, one can derive the probability axioms from decision-theoretic principles on their own. Second, they claim that Deutsch's proof contains a technical non-sequitur. Wallace (2003) handily dispenses with the first of these criticisms. In so many words: Deutsch is deriving the NRQM-chance link, not probability *tout court*. The supposed non-sequitur to which the second criticism points turns out to be a consequence of measurement neutrality, which in turn follows from the assumption of state supervenience (Wallace 2003, p. 432; Wallace 2012, p. 197). If we can defend state supervenience, then the criticisms of Barnum et al. (2000) dissolve. With the above argument, Wallace (2003, 2012) claims that EQM is unique in its ability to give such a defense.

One thing to note about this argument is that if it were sound, it would be quite a bit stronger than Wallace lets on. The Hydra view’s ontology contains only Ψ , so nothing could break its symmetries. Different branches are nothing more than different parts of the same object. But the Lewisian’s ontology *does* break these symmetries! Individual branches (qua individual Lewisian worms) break the symmetries that Ψ establishes across all the branches. For example, let α be Alice’s spin-up worm and let α' be Alice’s spin-down worm. Before branching, these worms are spatiotemporally coincident (and so the total state is symmetric). After branching, one branch contains α , and the other does not. Symmetry-wise, this case is no different from that in which one pilot wave branch “contains” the particle while the other does not.¹⁸ So Wallace’s Everettian can have probability (on the Hydra view) or charity (on the Lewisian view), but not both. Given that Wallace argues that we ought to adopt the Lewisian view because of its charity, the above argument’s soundness would put the aspiring EQMer in a rather uncomfortable position.

Fortunately, it is not sound. Specifically, the premise that symmetry-breaking scuttles state supervenience is false. As noted in the last section, supervenience follows both from the Lewisian’s self-ignorance and the pilot waver’s q -ignorance. Supposing otherwise leads to a contradiction with our definition of chance. Note first that by decoherence reliability, chances can depend on Ψ . But self-ignorance (q -ignorance) implies that any knowledge of α (q) is unreliable. So α (q) is not admissible background information b in PP, and chances cannot vary with it. Since Ψ and α (q) give a complete description of the system, Ψ is the *only* reliable information an agent can access. (Note, too, that we can bypass the ignorance principles by inferring this conclusion directly from decoherence exclusivity.) So it had better be the case that chances vary only when Ψ varies!

There is no denying that the Hydra view’s proponents have a particularly elegant strategy for defending state supervenience: they simply note that Ψ is all there is. But the strategy of Lewisians and pilot wavers, while not quite as simple, seems no less elegant: they simply note that Ψ is all we can reliably know. Supervenience then follows from Wallace’s operational approach to chance. Moreover, given that Wallace appears to adopt this strategy already for the Lewisian (2012, p. 150), it is hard to see how he could deny it to the pilot waver.¹⁹ So, for charitable theorists, the game is zero-sum. Nothing probabilistic is gained or lost when switching from many worlds to one.

Given the above, I take the Lewisian and Bohm-style views to motivate state supervenience equally well. Nevertheless, a skeptical pilot waver would be right to anticipate at least one looming challenge. First, Bohm (1952a, 1952b) motivates q -ignorance as a “practical limitation” rather than as a consequence of decoherence exclusivity, as I do here. But decoherence exclusivity is not a strict dynamical law: it describes how agents secure information. On these grounds, some might wish to deny that it can play a role in a proper physical theory.

While I sympathize with the desire to take dynamical explanations as far as possible, I do not wish to restrict physical theories to dynamical laws (and thereby demand an exclusively dynamical explanation of probability). As Callender (2007) saliently notes, there are plenty of non-dynamical laws in physics already (such as the Pauli exclusion principle), and some views of laws (especially Lewis’s best-system account) make a restriction to dynamics appear unnatural.

This response might raise another worry: if I have no compunction with non-dynamical laws, why not simply adopt the Born rule as an axiom? I would agree that this move is not problematic,

¹⁸How seriously the reader wants to think of a wavefunction “containing” a particle will vary with the ontological choice flagged in fn. 16. But, e.g., a Bohmian taking a nomic view of Ψ should feel free to replace “contains” with “approximately describes.”

¹⁹In an early footnote, Wallace may concede that this argumentative strategy is possible. He notes that the relevant symmetries may only appear at the level of distributions on phase space (2003, p. 435, fn. 8). But using the symmetries in *distributions* raises the specter of circularity. My approach assigns only *dynamical* significance to the pilot wave Ψ (in the manner discussed in §3.1).

per se. Nonetheless, we might wish to make decoherence exclusivity axiom independently of any concerns about the Born rule. In this case, we would like to know whether the Born rule remains independent of the other axioms, and the symmetry theorem would show that it does not. So: do we have good, independent reasons to promote decoherence exclusivity to axiom-status?

I think that we do. At the very least, I wish to echo the familiar reason that it provides thorough dynamical detail about the process of measurement that forgoes any physical collapse of the state. But I also think that a non-cosmological approach to NRQM goes a long way towards motivating axiom-status. In this case, neither the Schrödinger nor the guiding equation is an exceptionless, universal law. Each is explicitly domain-specific. Therefore, decoherence exclusivity is an idealization that is no more “practical” or “contingent” than our initial choice of dynamics.

As an aside, it is worth noting that Bohm paves the way for a non-cosmological approach in his initial presentation of the pilot wave theory. He notes that his interpretation of NRQM might merely serve as a stepping-stone for further physical theories. For example, he asserts that “wherever the present form of the quantum theory is correct, our interpretation cannot lead to inconsistencies with relativity” and concedes that his search for relativistic extensions might just provide an “important heuristic principle in the search for *new* physical laws” (1952b, p. 187, emphasis mine). Keeping these caveats in mind, a careful Bohmian is always free to claim that their theory only applies to non-relativistic subsystems of the universe rather than the whole of it. I think that this attitude towards the de Broglie-Bohm theory naturally motivates the promotion of decoherence exclusivity to an axiom, thereby securing (a fortiori) q -ignorance and state supervenience. To derive the Born rule, all that remains is to draw structural links between our chance states and our NRQM states.

3.4 The symmetry theorem

The symmetry theorem flows from a simple observation: if a system starts in a superposition of two pointer states with the same weights, nothing changes upon re-labeling its final states. More explicitly, suppose that one unitary map U acts on Ψ as

$$(a|\uparrow\rangle + b|\downarrow\rangle) \otimes |E\rangle \xrightarrow{U} a|\text{up}\rangle + b|\text{down}\rangle \quad (27)$$

and consider the symmetric map U' obtained by swapping the final states,

$$(a|\uparrow\rangle + b|\downarrow\rangle) \otimes |E\rangle \xrightarrow{U'} a|\text{down}\rangle + b|\text{up}\rangle. \quad (28)$$

Now if $a = b$, $U\Psi = U'\Psi$. Thus, given state supervenience (and the rules of probability), the chances assigned to up and down both ought to equal one-half. In brief, the symmetry theorem uses decoherence availability and minimal chance-NRQM links to extend the above argument to arbitrary Ψ and arbitrary observables. Specifically, we need two premises that connect the structure of chance states to that of NRQM states. But if these premises hold for any one of the Hydra, Lewisian, or Bohm-stye views, then they hold for all three.

First, we assign chance functions to our branching history algebras that satisfy the rules of probability:

Probability. For every Ψ -branching history algebra $\{\mathcal{S}^i\}$, there is a chance function $\mathbf{ch}_\Psi : \{\mathcal{S}^i\} \rightarrow \mathbb{R}$ that satisfies Kolmogorov’s rules (3)–(6). This assignment defines a chance function $ch_\Psi : \mathcal{S}^i \rightarrow \mathbb{R}$ on each \mathcal{S}^i as follows:

$$ch_\Psi(\alpha_i) := \mathbf{ch}_\Psi\left(\bigvee\{\gamma \mid \gamma_i = \alpha_i\}\right). \quad (29)$$

In other words, the chance that α_i is the i th event is given by the chance that any one of the histories γ with that event occurs. It is easy to check that ch_Ψ also follows all of Kolmogorov’s rules (a property it inherits from \mathbf{ch}_Ψ). The Hydra, Lewisian, and Bohm-style views may all adopt Wallace’s operational definition of chance with PP. This definition stipulates Kolmogorov’s rules. If we are not worried about motivating them, the story ends here.

The Lewisian and Bohm-style views can easily do more: they can make quick and ready use of the frequentist justification of Kolmogorov’s rules given in §2.2. Each imparts semantics to \mathcal{S}^i on which $\alpha_i \leq \neg\beta_i$ means that when α_i occurs, β_i does not. Each takes repeated preparations of Ψ to yield finite sequences of independent events: identifications of worms for the former and approximate determinations of a particle’s position for the latter. In the former case, the agent notes at some t_{i+1} that, of the worms coincident with them at t_i , they were really in the set of such-and-such worms all along (e.g., those leading to a spin-up event), and similarly at time t_{i+2} , and so on. In the latter case, the agent notes at t_{i+i} that, of the possible particle positions at t_i , the particle was really in such-and-such a range, and similarly down the line. We suppose that these sequences are well-modeled by their infinite limits, and then the justification I give in §2.2 holds. This frequentist approach cannot work for the Hydra view, of course, since all events occur (i.e., any pair of α_i and β_i with non-zero weights both occur, even if $\alpha_i \leq \neg\beta_i$).

Second, we assume that normalization and conditionalization of chance and NRQM states agree with each another. Explicitly, we assume:

Structural links. Two of them:

1. *Normalization link.* $ch_\Psi(\alpha_i) = 1$ if and only if $\Psi(t_i)$ lies in the range of α_i .
2. *Temporal link.* The conditional quantum state $(\Psi, \alpha_i) := \alpha_i\Psi(t_i)/|\alpha_i\Psi(t_i)|^2$ agrees with the definition of conditional probability, i.e.,

$$\mathbf{ch}_{(\Psi, \alpha_i)}(\cdot) = \mathbf{ch}_\Psi\left(\cdot \mid \bigvee\{\gamma \mid \gamma_i = \alpha_i\}\right). \quad (30)$$

The normalization link connects NRQM states that lie in the range of projections with probability-one chance states. It is a weaker version of the eigenstate-eigenvalue link: a state lies in the range of a projection only if it possesses the relevant observable property (as opposed to “if and only if”). The temporal link simply coordinates the conditional chance and quantum states that are defined when some event α_i occurs. Each of the three views of measurement adds strictly more structure to NRQM-plus-decoherence. So if the above coordination of structure is justified for one view, it is justified for all three.

Before stating the symmetry theorem, we formalize state supervenience with the chance functions defined by *probability*.

State supervenience. If $\Psi(t_i) = \Phi(t_i)$, then $ch_\Psi(\alpha_i) = ch_\Phi(\alpha_i)$.

As flagged at the jump, state supervenience is the core assumption of the symmetry theorem. Accordingly, the theorem lives or dies on our ability to justify it. But the premise follows from decoherence exclusivity, and (as we have seen) all three views can adopt this axiom. The Hydra theorist might wish to forgo this axiom, as their metaphysics suffice to justify state supervenience. But decoherence exclusivity seems to be the best that either the Lewisian or the pilot wave theorist can do.

The preceding allows for a simple statement of the symmetry theorem for all three views:

The symmetry theorem. For a quantum system satisfying *decoherence availability*, *decoherence exclusivity*, *probability*, and *structural links*, the chance functions are given by the Born rule,

$$ch_{\Psi}(\alpha_i) = \langle \Psi(t_i), \alpha_i \Psi(t_i) \rangle. \quad (31)$$

And so the symmetry theorem explains why systems with the physical state χ including Ψ ought to be assigned the chance state ch on multiple accounts of what χ represents. I provide a proof in Appendix A.

While Saunders (2004) provides an explicitly operational derivation of the symmetry theorem, the role of decoherence in his proof is not immediately apparent. I borrow Wallace’s (2012) more transparent strategy for my proof. Saunders also suggests that his operational derivation is only applicable to one pointer basis and serves as an alternative to Gleason’s theorem (2004, pp. 1786–1787). But the above shows that the symmetry theorem applies to all pointer bases by attending to all the dynamical possibilities.

Moreover, while Wallace asserts that Gleason’s theorem does not add anything to the symmetry argument (2003, p. 434), the former complements the latter. The symmetry theorem proves that the NRQM-chance link ought to be given by the Born rule, while Gleason’s theorem shows that this rule does not miss any of the possible chance states for open quantum systems. Explicitly, so long as the projections α have the form (22), the Born rule for open quantum systems, equation (10), follows from the Born rule for closed quantum systems, (31), for every density operator ρ . (This fact is an immediate consequence of Stinespring’s (1955) dilation theorem, which guarantees that every density operator ρ can be expressed as $\text{Tr}_E |\Psi\rangle\langle\Psi|$ for some \mathcal{H}_E and some $\Psi \in \mathcal{H}_S \otimes \mathcal{H}_E$.)

As a bonus, the symmetry theorem derives, rather than assumes, measurement neutrality, the independence of chance on measurement context (i.e., the particular PVM that we use to measure $P_i \in \mathcal{P}(\mathcal{H}_S)$). We do not assume that anything is “noncontextual”: the context-dependence of Bohmian properties is consistent with each of our four premises.²⁰ Indeed, the theorem provides an explanation of why the chances in the pilot wave theory do not depend on the context of measurement even while the values of properties do.

As an aside, one might say that Gleason tacitly assumes measurement neutrality by assigning probability functions directly to the partial σ -algebra, $\mathcal{P}(\mathcal{H}_S)$. In a branching history algebra, we associate each full σ -subalgebra of $\mathcal{P}(\mathcal{H}_S)$ with a given PVM for one \mathcal{S}^i —so the same event occurring in different measurement contexts need not have the same chance by default. This way of framing the matter adds an asterisk to the relationship between Gleason’s theorem and the symmetry theorem: the former verifies that the Born rule for open systems (which the latter endorses) recovers all of the *measurement-neutral* chance states.

I do think that there is a sense in which the symmetry theorem is easier to visualize when we assume the Hydra view rather than the Lewisian or the Bohm-style view. On the Hydra view, all branches are on equal ontic footing, and so we are invited to view the dynamical symmetries of Ψ in much the same way that we envision the spatial symmetries of a die or a coin. But as shown above, Ψ ’s symmetries can have the same import for chance regardless of whether Ψ is the totality of ontology or merely a law governing particle motions—whether it is an “is” or a “tends to.” And I am more than happy to grant that many-worlds has been productive as a research program. This claim, after all, is far weaker than the anti-skeptic’s.

²⁰Wallace (2012, p. 197) refers to measurement neutrality as “noncontextuality.” But since this usage is at odds with the more prevalent notion of Kochen-Specker contextuality (1975), I will stick with the former term. The pilot wave theory is both measurement-neutral and (Kochen-Specker) contextual (because the values of self-adjoint observables typically depend on the context of measurement).

4 Discussion

On the one hand, this paper has been a work of criticism. I have shown that Wallace’s (2012) claim that many-worlds theories are better off than single-world theories in making sense of probability fails by his own lights. While Wallace argues that only a many-worlds theorist can justify state supervenience, I have shown that a user of the Bohm-style pilot wave theory can do so by noting that agents cannot know the precise configurations of particles. This single-world theorist can therefore make equally good use of the Deutsch-Wallace theorem.

On the other hand, the main takeaway of this paper ought to be positive. I hope to have shown that the Deutsch-Wallace theorem should not be an item of niche interest, a feather in the cap of the most philosophically adventurous physicists—on the contrary, a whole host of approaches to quantum theory might appeal to it to explain the ascription of particular chance values. Note that this explanation presupposes not just Hilbert space structure but also decoherence-governed preparations and measurements. But I do not wish to claim that the theorem gives the final word on quantum probability. I merely wish to stress that it is one important step forward in our understanding, one that might hold interest for theorists of many different philosophical stripes.

In this spirit, I would like to conclude with a brief sketch of how the symmetry theorem might help settle a dispute between Dürr, Goldstein, and Zanghì (1992) and Valentini (2020). As flagged in §2.1, DGZ’s approach to pilot wave theory differs from mine in at least one crucial way: they take a cosmological approach and assign a state (Ψ, \mathbf{q}) to the entire universe. Agents, of course, cannot access the universe’s wave function Ψ directly. They only gain knowledge of the effective wave functions ψ of isolated subsystems. Nonetheless, the universal wavefunction Ψ crucially affords DGZ an extra conceptual resource when it comes the question of q -ignorance.

DGZ recognize the importance of q -ignorance (which they call “absolute uncertainty”): they are well aware that an agent who knows both the precise q and effective wave function ψ of a subsystem could predict future outcomes with certainty, thereby violating NRQM’s empirical adequacy (1992, pp. 855, 883). Ambitiously, however, they want to avoid adding q -ignorance as axiom to their theory. Moreover, they do not wish to appeal to decoherence to explain it. Instead, they seek to use facts about the universal wave function to derive *both* q -ignorance and the Born rule in one go.

Here is a brief and qualitative sketch of their argument. They start by splitting the universal configuration \mathbf{q} into a microscopic subsystem’s configuration, q , and that of its environment, e , so that $\mathbf{q} = (q, e)$. The environment couples to the subsystem at most through the latter’s conditional wave function,

$$\psi(q) := \Psi(q, E), \tag{32}$$

for some specification of the environment’s configuration $e = E$. Then they introduce a measure of “typicality” over initial configurations of the universe given by the *quantum equilibrium distribution*

$$\mathbf{P}(d\mathbf{q}) = |\Psi(\mathbf{q})|^2 d\mathbf{q}, \tag{33}$$

It follows immediately from (32) that

$$\mathbf{P}(dq | E) = |\psi(q)|^2 dq. \tag{34}$$

In other words, the system and the environment must be \mathbf{P} -independent given ψ . But as of yet, \mathbf{P} only has meaning as a measure of typicality; we have not (strictly speaking) introduced anything about chance or frequency into the picture. Thus, DGZ apply the (weak) Law of Large Numbers to an ensemble of empirical subsystems that have the same effective wave function given the appropriate states of the environment. They show that the distribution of precise values of q across these subsystems must agree with the subsystem equilibrium (34) in the long run, at least for “typical”

initial configurations \mathbf{q} of the universe. Thus, we have derived an agreement of actual frequencies of actual values of q with the Born rule. But how does this fact imply a restriction on agents' knowledge?

To answer this question, DGZ note that “whatever [...] could conceivably be regarded as knowledge of, or information concerning, the systems under investigation, must be a part of or grounded in the environment of these systems” (1992, p. 883). That is, an agent learns about a subsystem (ψ, q) by studying its environment. It follows that, for each subsystem in an empirical ensemble, the best an agent could do (in principle!) would be to know the precise configuration $e = E$ of q 's environment. This information would straightforwardly provide the agent with precise knowledge of each subsystem's effective wave function ψ via equation (32). But the foregoing LLN argument proves that it can provide no more detailed knowledge of these subsystem's qs than that given by equilibrium. So, they have derived q -ignorance, as desired.

For clarity, it is worth spelling out a simple, equal-time version of this typicality argument. Start with a universal wave function Ψ yielding systems q_1, \dots, q_M with the environment $e_t = E_M$ at time t , each with the same conditional wave function ψ . Letting z range across the possible configurations of each subsystem, we represent a measurement of z with the distribution

$$\rho_{\text{emp}}(z) = \frac{1}{M} \sum_{i=1}^M \delta(q_i - z). \quad (35)$$

Now pick any bounded function f and $\epsilon > 0$, and say that the ensemble *agrees* with equilibrium when

$$\|\rho_{\text{emp}}(z) - |\psi(z)|^2\|_f = \left| \int dz f(z) (\rho_{\text{emp}}(z) - |\psi(z)|^2) \right| < \epsilon. \quad (36)$$

Let $\mathbf{A}(M, f, \epsilon, t)$ be the set of all values q_0 for the initial configuration of the universe that yield such agreement. Dürr et al. use (weak) LLN to show that, by the lights of the quantum equilibrium distribution \mathbf{P} defined by (33),

$$\mathbf{P}(q_0 \in \mathbf{A}(M, f, \epsilon, t) | e_t = E_M) = 1 - \delta(M, f, \epsilon) \quad (37)$$

where $\delta(M, f, \epsilon)$ goes to zero as M goes to infinity. (37) gives the precise sense in which agreement holds for “typical” initial q : namely, such agreement is typical by the lights of the Born-rule measure on the universe. So, even if an agent knew the environmental configuration E_M precisely, they could not derive information about M subsystems beyond what is encoded in their effective wavefunctions. Dürr et al. then generalize this result for random multi-time ensembles with random environmental conditions.

The success of this argument hinges, in part, on whether we can justify the choice of quantum equilibrium as the measure of typicality. Dürr et al. aim to do so by invoking the fact that this choice is equivariant: so long as the measure is equilibrium at *some* time, then it must be equilibrium at *all* times. But there are good reasons to worry about their appeal. As they note themselves, *any* initial \mathbf{P} will agree with (single-time) ensembles once it is time-evolved to the appropriate \mathbf{P}_t (1992, p. 874). Equivariance simplifies the analysis, but it is not clear that simplicity is enough to motivate ignoring alternative initial measures—measures that may approximate the Born rule well in a suitable domain.

Valentini (2020) argues for a different cosmological approach to pilot wave theory that admits such non-equilibrium measures. He identifies the same worry about the sufficiency of equivariance in Dürr et al.'s argument (2020, pp. 24–25). However, he also challenges the coherence of distinguishing typicality from probability; in his view, they are synonymous. Thus, he paints their situation

as rather more dire: “[There] is no scientific basis for the claim that quantum non-equilibrium is intrinsically unlikely. This claim stems, as we have seen, from a circular argument in which the Born-rule measure is taken to define ‘typicality’ for the initial conditions of the universe” (2020, p. 455). I prefer to be a bit more cautious: like it or not, we can make sense of these concepts coming apart. But I think Valentini’s worry is one well worth taking seriously. For even if we do take typicality and chance to diverge, it is far from clear to me that a measure of one can provide a telling explanation of a measure of the other.

Thus, while I do not wish to take a side in this particular dispute, I think that it is instructive to note how the symmetry theorem might help adjudicate it. To start, note that DGZ cannot fully eliminate epistemic premises from their derivation: they still need to claim that agents’ knowledge of a subsystem is “grounded in [its] environment” in order to derive q -ignorance. But this claim, to me, seems just as unsatisfactory as Wallace’s claim that the Lewisian can help themselves to self-ignorance. Absent an explicit, physical model of measurement, I am not sure why I ought to accept it. But the explicit, physical models of measurement that spring most readily to mind are those provided by decoherence. And as long as we are clear on which subsystems of the universe fit this model, we can help ourselves to decoherence exclusivity.

So consider a subsystem (Ψ, q) of DGZ’s universal state (Ψ, \mathbf{q}) where decoherence exclusivity applies, and suppose (Ψ, q) contains an empirical ensemble of subsystems with effective wave functions ψ . The symmetry theorem yields a Born-rule chance function for Ψ . The distribution $\mathbf{P}(dq) = |\Psi(q)|^2$ is a straightforward consequence of that function.²¹ Dürr et al.’s LLN demonstration still applies to ψ , more-or-less as stated. It just does different work. It verifies the consistency of long-run frequencies with our derived chances, in the sense sketched for the frequentist approach to chance in §2.2. As long as typicality aligns with chance, that agreement with the Born rule holds in a *typical* Bohmian system (Ψ, q) is just an analytic consequence of chance’s self-consistency.

There are at least three virtues to this approach. First, it does not stipulate any measure of typicality or chance: it derives all measures from state-space symmetries. Second, it does not need any robust distinction between typicality and chance: it only needs the thin, operational definition of chance from §2.1. Third, it makes the environmental mechanism that restricts agent’s knowledge explicit: rather than attempting to show that some such mechanism *must* exist, we just directly describe it via the decoherence program.

This approach charts a middle ground between the philosophies of DGZ and Valentini, and in so doing it might invoke commitments that either (or both) find objectionable. In particular, DGZ may want a frequentist-style explanation that reduces all facts about agents’ ignorance to facts about (actual) frequencies. They might provide such an explanation with their “typicality” argument (if it works), but clearly they would have to give up on this goal if they were to adopt my operational definition of chance (which invokes agents’ credences directly). Nonetheless, my approach allows for both non-standard initial distributions on a universal configuration \mathbf{q} and a comparatively flexible approach to interpretations of chance, all while securing q -ignorance and the empirical adequacy of Bohmian NRQM in a suitable domain of applicability. This approach strikes me as promising, but it assuredly requires more work. In particular, I will leave for future inquiry the question of how a Bohmian might wish to view the symmetry theorem in light of Valentini’s (2019) quantum H -theorem.

For now, though, I want to stress that the aim of the present work is to highlight the flexibility and fecundity of the symmetry theorem as a conceptual resource, not to rock any particular Bohmian’s boat. I have sketched an alternative to the typicality approach, but its proponents are free to ignore

²¹For example, consider a probability space with an atomless σ -algebra generated by the projections $\chi(dq)$ onto position intervals and a measure given by the Born rule. Then push that measure forward to \mathbb{R}^{3N} . See Stroock (2010) for more details on this construction.

my sketch if they wish. My point is just that Wallace’s attempts to reserve a symmetry-based strategy for Everettians do not succeed.

Of course, applications of the symmetry theorem to various *other* single-world interpretations of NRQM are bound to involve idiosyncrasies that I have not treated here. To that, I can only say that I hope the above discussion provides a few important heuristics for future attempts to use the theorem more broadly.

Acknowledgements

The author would like to thank Harvey Brown, Benjamin H. Feintzeig, James Read, and David Wallace for their invaluable discussions and suggestions. Many thanks as well to the thorough and rigorous criticisms of three anonymous referees, without which this paper would be in a much sorrier state (though any remaining errors are, of course, the author’s sole responsibility!). Finally, much thanks to the participants of the 2020 Michaelmas Oxford Philosophy of Physics Seminar for a fantastic discussion of this work. The author was supported during the completion of this work by the National Science Foundation under Grant No. 1846560.

A Proof of the symmetry theorem

Before proving the symmetry theorem, we prove as a lemma an important condition that follows from the *structural links* assumption.

Branching link. When $\{\mathcal{S}^i\}$ is strictly Ψ -branching and $\mathcal{T}(\alpha_i, \beta_j) \neq 0$ for $t_i < t_j$,

$$ch_{(\Psi, \alpha_i)}(\beta_j) = \frac{ch_{\Psi}(\beta_j)}{ch_{\Psi}(\alpha_i)}. \quad (38)$$

Proof of the branching link. First, we show that chances respect branching, i.e., $\mathbf{ch}_{\Psi}(\gamma) = 0$ whenever γ contains α_i, β_j such that $\mathcal{T}(\alpha_i, \beta_j) = 0$. Explicitly,

$$\mathbf{ch}_{\Psi}(\gamma) \leq \mathbf{ch}_{\Psi}\left(\bigvee\{\gamma \mid \gamma_j = \beta_j \wedge \gamma_i = \alpha_i\}\right) = \mathbf{ch}_{\Psi}\left(\bigvee\{\gamma \mid \gamma_j = \beta_j\} \mid \bigvee\{\gamma \mid \gamma_i = \alpha_i\}\right) \quad (39)$$

$$= \mathbf{ch}_{(\Psi, \alpha_i)}\left(\bigvee\{\gamma \mid \gamma_j = \beta_j\}\right) \quad (40)$$

$$= ch_{(\Psi, \alpha_i)}(\beta_j) \quad (41)$$

$$= 0, \quad (42)$$

where (39) follows from additivity and the definition of conditional probability, (40) follows from *temporal link*, i.e., (30), (41) follows from the definition (29) in the *probability* assumption, and (42) follows from *normalization link*, the rules of probability, and the fact that $\mathcal{T}(\alpha_i, \beta_j) = 0$. The rules of probability then imply that $\mathbf{ch}_{\Psi}(\gamma) = 0$.

Now suppose that $\{\mathcal{S}^i\}$ is strictly Ψ -branching and $\mathcal{T}(\alpha_i, \beta_j) \neq 0$ for $t_i < t_j$. Note that

$$ch_{(\Psi, \alpha_i)}(\beta_j) = \mathbf{ch}_{(\Psi, \alpha_i)}\left(\bigvee\{\gamma \mid \gamma_j = \beta_j\}\right) \quad (43)$$

$$= \mathbf{ch}_{\Psi}\left(\bigvee\{\gamma \mid \gamma_j = \beta_j\} \mid \bigvee\{\gamma \mid \gamma_i = \alpha_i\}\right) \quad (44)$$

$$= \frac{\mathbf{ch}_{\Psi}(\bigvee\{\gamma \mid \gamma_j = \beta_j \wedge \gamma_i = \alpha_i\})}{\mathbf{ch}_{\Psi}(\bigvee\{\gamma \mid \gamma_i = \alpha_i\})} \quad (45)$$

$$= \frac{\mathbf{ch}_{\Psi}(\bigvee\{\gamma \mid \gamma_j = \beta_j\})}{\mathbf{ch}_{\Psi}(\bigvee\{\gamma \mid \gamma_i = \alpha_i\})} \quad (46)$$

$$= \frac{ch_{\Psi}(\beta_j)}{ch_{\Psi}(\alpha_i)}, \quad (47)$$

where (43) follows from *probability*, (44) follows from *temporal link*, (45) follows from the definition of conditional probability (and a bit of algebra), (46) follows from \mathbf{ch}_{Ψ} respecting branching and additivity (and our supposition), and (47) follows from (29). \square

With *branching link* in hand, we can quickly prove the symmetry theorem.

Proof of the symmetry theorem. Following (Wallace, 2012, Ch. 4), we aim to prove that $ch_{\Psi}(\alpha_i) = \langle \Psi(t_i), \alpha_i \Psi(t_i) \rangle$ in four steps. First (i), we generalize the intuitive re-labeling argument to show that projections with equal Born weights must have equal chances—i.e., we show that ch_{Ψ} must be a function of Born weights. Second (ii), we invoke some of our environmental degrees of freedom to show that this function is increasing. Third (iii), we invoke N environmental degrees of freedom to show that function must equal the Born weight when it is rational. Fourth (iv), we use a simple limiting argument (and the second and third steps) to obtain agreement for arbitrary Born weights.

(i) Suppose $\langle \Psi(t_i), \alpha_i \Psi(t_i) \rangle = \langle \Psi(t_i), \beta_i \Psi(t_i) \rangle$.

First, suppose that both sides equal zero. Then $\Psi(t_i)$ lies in the range of both $\neg\alpha_i$ and $\neg\beta_i$, and so by *normalization link*, $ch_{\Psi}(\neg\alpha_i) = ch_{\Psi}(\neg\beta_i) = 1$. Thus, by *probability*, $ch_{\Psi}(\alpha_i) = ch_{\Psi}(\beta_i) = 0$.

Now suppose otherwise. By the above reasoning, we get that $ch_{\Psi}(\alpha_i) \neq 0$ and $ch_{\Psi}(\beta_i) \neq 0$. Next, we run the analog of the intuitive re-labelling argument. By *decoherence availability*, we can consider two different strictly branching history algebras for the next step of decoherence: one that gives a projection α_{i+1} weight from α_i and one that gives it weight from β_i . To distinguish these, let Ψ evolve with the first dynamics, and let Φ evolve with the second dynamics. Now define $\Phi(t_j) = \Psi(t_j)$ for $j \leq i$ and let $\Psi(t_{i+1}) = X\Psi(t_i)$ and $\Phi(t_{i+1}) = Y\Psi(t_i)$ for the unitary operators

$$X := V^{\alpha}\alpha_i + W^{\alpha}(1 - \alpha_i) \quad Y := V^{\beta}\beta_i + W^{\beta}(1 - \beta_i) \quad (48)$$

where V^{α} and V^{β} share the range of α_{i+1} and W^{α} and W^{β} share the range of some mutually orthogonal projection β_{i+1} . By construction, $\Psi(t_{i+1}) = \Phi(t_{i+1})$.

Note, too, that $X\alpha_i\Psi(t_i)$ and $Y\beta_i\Psi(t_i)$ both lie in the range of α_{i+1} . So by *branching link*, we have

$$ch_{(\Psi, \alpha_i)}(\alpha_{i+1}) = \frac{ch_{\Psi}(\alpha_{i+1})}{ch_{\Psi}(\alpha_i)} = 1, \quad (49)$$

and similarly

$$ch_{(\Phi, \beta_i)}(\alpha_{i+1}) = \frac{ch_{\Phi}(\alpha_{i+1})}{ch_{\Phi}(\beta_i)} = 1. \quad (50)$$

By *state supervenience*, $ch_{\Phi}(\alpha_{i+1}) = ch_{\Psi}(\alpha_{i+1})$, and so $ch_{\Psi}(\alpha_i) = ch_{\Phi}(\beta_i)$. Applying *state supervenience* once more, we get

$$ch_{\Psi}(\alpha_i) = ch_{\Psi}(\beta_i). \quad (51)$$

(Note that, again by *state supervenience*, this last equality holds even if our original history algebra was not strictly branching.)

(ii) Suppose $\langle \Psi(t_i), \alpha_i \Psi(t_i) \rangle > \langle \Psi(t_i), \beta_i \Psi(t_i) \rangle$.

By *decoherence availability*, we may assume that the $i + 1$ step of decoherence is given by a unitary Y such that, for γ_{i+1} and ω_{i+1} two mutually orthogonal projections, we have

$$\begin{aligned} \langle Y^\dagger \Psi(t_i), (\gamma_{i+1} \vee \omega_{i+1}) Y \Psi(t_i) \rangle &= \langle \Psi(t_i), \alpha_i \Psi(t_i) \rangle \\ \langle Y^\dagger \Psi(t_i), \gamma_{i+1} Y \Psi(t_i) \rangle &= \langle \Psi(t_i), \beta_i \Psi(t_i) \rangle \end{aligned} \quad (52)$$

where Y maps vectors in the range of α_i to the range of $\gamma_{i+1} \vee \omega_{i+1}$. Then, by step (i),

$$ch_{\Psi}(\gamma_{i+1} \vee \omega_{i+1}) = ch_{\Psi}(\alpha_i), \quad ch_{\Psi}(\gamma_{i+1}) = ch_{\Psi}(\beta_i). \quad (53)$$

By *probability*,

$$ch_{\Psi}(\gamma_{i+1} \vee \omega_{i+1}) \geq ch_{\Psi}(\gamma_{i+1}) \quad (54)$$

and so

$$ch_{\Psi}(\alpha_i) \geq ch_{\Psi}(\beta_i). \quad (55)$$

(iii) Suppose $\langle \Psi(t_i), \alpha_i \Psi(t_i) \rangle$ is rational, i.e. equal to $\frac{M}{N}$ for some positive integers M, N .

By *decoherence availability*, we may pick some \mathcal{H}_{SE} -spanning sequence of orthogonal projections $\gamma^1, \dots, \gamma^m, \dots, \gamma^N$ that generate a σ -algebra containing α_i such that, for some Φ such that $\Phi(t_i) = \Psi(t_i)$,

$$\langle \Phi(t_i), \gamma^m \Phi(t_i) \rangle = \frac{1}{N} \quad (56)$$

for all m . By (i), $ch_{\Phi}(\gamma^m)$ must be independent of m . Thus, by *probability*, $ch_{\Phi}(\gamma^m) = \frac{1}{N}$.

Now let $\omega := \bigvee_{i \leq M} \gamma^i$. We have that $\langle \Psi(t_i), \alpha_i \Psi(t_i) \rangle = \langle \Phi(t_i), \omega \Phi(t_i) \rangle = \langle \Psi(t_i), \omega \Psi(t_i) \rangle$, so by step (i), $ch_{\Psi}(\alpha_i) = ch_{\Psi}(\omega)$. By *probability*, $ch_{\Psi}(\omega) = \sum_{m=1}^M ch_{\Psi}(\gamma^m) = \frac{M}{N}$, and so we get that

$$ch_{\Psi}(\alpha_i) = \frac{M}{N}. \quad (57)$$

(iv) Suppose $\langle \Psi(t_i), \alpha_i \Psi(t_i) \rangle = r \in [0, 1]$, where r may not be rational.

By (i), ch is a function of Born-rule weights, i.e.

$$ch_{\Psi}(\alpha_i) = f(\langle \Psi(t_i), \alpha_i \Psi(t_i) \rangle) \quad (58)$$

for some $f : [0, 1] \rightarrow [0, 1]$. By step (ii), f is increasing; by step (iii), $f(M/N) = M/N$.

So let $\{a_i\}$ and $\{b_i\}$ be, respectively, increasing and decreasing sequences of rational numbers in $[0, 1]$ converging to r . Since $f(b_i) = b_i$ for all i , $f(r) \leq r$. And since $f(a_i) = a_i$ for all i , $f(r) \geq r$ —thus, $f(r) = r$, and so

$$ch_{\Psi}(\alpha_i) = \langle \Psi(t_i), \alpha_i \Psi(t_i) \rangle. \quad (59)$$

□

References

- Barnum, H., Caves, C. M., Finkelstein, J., Fuchs, C. A., and Schack, R. (2000). Quantum probability from decision theory? *Proceedings of the Royal Society of London*, A456:1175–1182.
- Barrett, J. A. (2019). *The Conceptual Foundations of Quantum Mechanics*. Oxford University Press, Oxford.
- Bohm, D. (1952a). A suggested interpretation of the quantum theory in terms of “hidden” variables. I. *Physical Review*, 85(2):166–179.
- Bohm, D. (1952b). A suggested interpretation of the quantum theory in terms of “hidden” variables. II. *Physical Review*, 85(2):180.
- Brown, H. R. (2011). Curious and sublime: the connection between uncertainty and probability in physics. *Philosophical Transactions of the Royal Society of London A*, 369(1956):4690–4704.
- Busch, P. (2003). Quantum states and generalized observables: a simple proof of Gleason’s theorem. *Physical Review Letters*, 91(12):120403.
- Busch, P., Grabowski, M., and Lahti, P. J. (1995). *Operational Quantum Physics*. Springer-Verlag, Berlin.
- Callender, C. (2007). The emergence and interpretation of probability in Bohmian mechanics. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 38(2):351–370.
- Deutsch, D. (1999). Quantum theory of probability and decisions. *Proceedings: Mathematical, Physical and Engineering Sciences*, 455(1988):3129–3137.
- 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.
- Gell-Mann, M. and Hartle, J. B. (1990). Quantum mechanics in the light of quantum cosmology. In Zurek, W. H., editor, *Complexity, Entropy and the Physics of Information*, page 425–459. Addison-Wesley, Redwood City.
- Gell-Mann, M. and Hartle, J. B. (2012). Decoherent histories quantum mechanics with one real fine-grained history. *Physical Review A*, 85(6):062120.
- Gleason, A. M. (1957). Measures on the closed subspaces of a Hilbert space. *Journal of Mathematics and Mechanics*, 6(6):885–893.
- Hall, N. (2004). Two mistakes about credence and chance. *Australasian Journal of Philosophy*, 82(1):93–111.
- Hartle, J. B. (2010). Quasiclassical realms. In Saunders, S., Barrett, J., Kent, A., and Wallace, D., editors, *Many Worlds? Everett, Quantum Theory, & Reality*, page 73–98. Oxford University Press, Oxford.
- Ismael, J. (2008). Raid! Dissolving the big, bad bug. *Noûs*, 42(2):292–307.
- Kochen, S. (2015). A reconstruction of quantum mechanics. *Foundations of Physics*, 45(5):557–590.

- Kochen, S. and Specker, E. P. (1975). The problem of hidden variables in quantum mechanics. In *The Logico-Algebraic Approach to Quantum Mechanics: Volume I: Historical Evolution*, pages 293–328. Springer.
- Lanford III, O. E. (1981). The hard sphere gas in the Boltzmann-Grad limit. *Physica A*, 106:70–76.
- Lewis, D. (1980). A subjectivist’s guide to objective chance. In Jeffrey, R. C., editor, *Studies in Inductive Logic and Probability*, volume 2, pages 263–294. University of California Press, Berkeley.
- Norsen, T. (2014). The pilot-wave perspective on spin. *American Journal of Physics*, 82(4):337–348.
- Norton, J. D. (2012). Approximation and idealization: Why the difference matters. *Philosophy of Science*, 79(2):207–232.
- Papineau, D. (1996). Many minds are no worse than one. *The British Journal for the Philosophy of Science*, 47(2):233–241.
- Pettigrew, R. (2012). Accuracy, chance, and the principal principle. *Philosophical Review*, 121(2):241–275.
- Read, J. (2018). In defence of Everettian decision theory. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*.
- Rédei, M. (1998). *Quantum Logic in Algebraic Approach*. Kluwer Academic Publishers, Dordrecht.
- Romano, D. (2016). Bohmian classical limit in bounded regions. In Feline, L., Ledda, A., Paoli, F., and Rossanese, E., editors, *New Directions in Logic and Philosophy of Science*. Lightning Source, Milton Keynes.
- Romano, D. (2020). Multi-field and Bohm’s theory. *Synthese*. <https://doi.org/10.1007/s11229-020-02737-6>.
- Rosaler, J. (2016). Interpretation neutrality in the classical domain of quantum theory. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 53:54–72.
- Saunders, S. (2004). Derivation of the Born rule from operational assumptions. *Proceedings: Mathematical, Physical and Engineering Sciences*, 460(2046):1771–1788.
- Saunders, S. (2010). Chance in the Everett interpretation. In Saunders, S., Barrett, J., Kent, A., and Wallace, D., editors, *Many Worlds? Everett, Quantum Theory, & Reality*, book section 6, page 181–205. Oxford University Press, Oxford.
- Saunders, S. and Wallace, D. (2008). Branching and uncertainty. *The British Journal for the Philosophy of Science*, 59(3):293–305.
- Schlosshauer, M. (2007). *Decoherence and the Quantum-to-Classical Transition*. Springer Science & Business Media.
- Sebens, C. T. and Carroll, S. M. (2018). Self-locating uncertainty and the origin of probability in Everettian quantum mechanics. *The British Journal for the Philosophy of Science*, 69(1):25–74.
- Stinespring, W. F. (1955). Positive functions on C*-algebras. *Proceedings of the American Mathematical Society*, 6(2):211–216.

- Stroock, D. W. (2010). *Probability Theory: An Analytic View*. Cambridge University Press, Cambridge, 2 edition.
- Tappenden, P. (2011). Evidence and uncertainty in Everett’s multiverse. *British Journal for the Philosophy of Science*, 62(1):99–123.
- Valentini, A. (2019). Foundations of statistical mechanics and the status of the born rule in de Broglie-Bohm pilot-wave theory.
- Valentini, A. (2020). Foundations of statistical mechanics and the status of the born rule in de Broglie-Bohm pilot-wave theory. In Allori, V., editor, *Statistical Mechanics and Scientific Explanation: Determinism, Indeterminism and Laws of Nature*, page 423–478. World Scientific, Singapore.
- van Fraassen, B. C. (1980). *The Scientific Image*. Clarendon Press.
- von Mises, R. (1981). *Probability, Statistics, and Truth*. Dover, New York, 2 edition.
- Wallace, D. (2003). Everettian rationality: defending Deutsch’s approach to probability in the Everett interpretation. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 34(3):415–439.
- Wallace, D. (2012). *The Emergent Multiverse: Quantum Theory According to the Everett Interpretation*. Oxford University Press.
- Wallace, D. (2019). Isolated systems and their symmetries, part II: local and global symmetries of field theories. <http://philsci-archive.pitt.edu/16624/>.
- Wallace, D. (2020). On the plurality of quantum theories: Quantum theory as a framework, and its implications for the quantum measurement problem. In French, S. and Saatsi, J., editors, *Scientific Realism and the Quantum*, book section 5, page 78–102. Oxford University Press, Oxford.
- Wilson, A. (2012). Everettian quantum mechanics without branching time. *Synthese*, 188(1):67–84.
- Zurek, W. H. (2005). Probabilities from entanglement, Born’s rule $p_k = |\psi_k|^2$ from enviance. *Physical Review A*, 71(5):052105.
- Zurek, W. H. (2009). Quantum Darwinism. *Nature Physics*, 5(3):181.