

A New Problem for Quantum Mechanics

Alexander Meehan

Abstract

In this article I raise a new problem for quantum mechanics, which I call the control problem. Like the measurement problem, the control problem places a fundamental constraint on quantum theories. The characteristic feature of the problem is its focus on state preparation. In particular, whereas the measurement problem turns on a premise about the completeness of the quantum state ('no hidden variables'), the control problem turns on a premise about our ability to prepare or control quantum states. After raising the problem, I discuss some applications. I suggest that it provides a useful new lens through which to view existing theories or interpretations, in part because it draws attention to aspects of those theories which the measurement problem does not (such as the role of conditional and relative states). I suggest that it also helps clarify the physical significance of the well-known no-go result—the no-cloning theorem—on which it is based.

1 Introduction

2 The Standard Measurement Problem

3 Quantum State Preparation

3.1 The preparation problem

4 The Control Problem

4.1 Set-up

4.2 Incompatibility argument

4.3 Other measurement problems

5 *Relation to Interpretations*

5.1 *Everettian mechanics*

5.2 *Bohmian mechanics*

5.3 *GRW theory*

5.4 *Other quantum theories*

6 *Relation to Foundations and the No-Cloning Theorem*

6.1 *Individual state determination*

7 *Conclusion*

Appendix

1 Introduction

There is a compelling argument that it is impossible to reconcile Schrödinger's quantum theory with the determinateness of measurement outcomes, unless we give up on the claim that the theory provides us with a complete description of physical systems. This is what I will call 'the standard measurement problem'. More precisely, the problem is that the following assumptions, though individually innocent, are jointly untenable (Maudlin [1995]):

- (A1) The quantum state of a system determines, directly or indirectly, all of its physical properties [**Completeness**].
- (A2) The quantum state of an isolated system always evolves in accord with a linear dynamical equation, such as the Schrödinger equation [**Linearity**].
- (A3) Given determinate inputs, our measuring devices always produce unique, determinate outcomes [**Determinate Outcome**].

One desideratum for a quantum theory (or interpretation) is that it solve the standard measurement problem. A theory must reject at least one of the assumptions listed, and explain how the resulting gap is to be filled in—preferably, without introducing 'measurement' as a

basic, unanalysed term in the dynamical laws.¹ Bohmian mechanics, for instance, rejects (A1), GRW theory rejects (A2), and Everettian mechanics rejects (A3).

I argue that there is another problem for quantum mechanics, distinct from the standard measurement problem. The characteristic feature of this new problem, which I call the control problem,² is that it introduces considerations of state preparation on top of considerations of measurement. In particular, the problem is that the following claims are jointly incompatible:

- (B1) We can successfully prepare quantum states: at least some of our preparation devices are such that, if determinately fed many inputs, they output a non-trivial fraction of those inputs in some specified range of quantum states [**Preparation**] (here the ‘inputs’ are subsystems, and we define ‘the quantum state of a subsystem’ in the standard way, as its reduced state).
- (B2) The quantum state of an isolated system always evolves in accord with a deterministic dynamical equation that preserves the inner product, such as the Schrödinger equation [**Unitarity**].
- (B3) It is always determinate whether or not a subsystem has been input into a given (measuring or preparation) device [**Determinate Input**].

(There is another background assumption needed to generate the incompatibility, which is also taken for granted in the standard measurement problem:

- (A4, B4) There is always, from a first-person view, an experienced ‘observed outcome’ of a measurement, even if—from a ‘God’s eye view’—there is no unique determinate outcome. Moreover, our measuring devices are at least somewhat informative and reliable. For example, if a spin- z measuring device is determinately fed an electron in the eigenstate \uparrow_z , it will output ‘UP’ (and not ‘DOWN’) [**Competent Measurement**].

¹ Authors often use the term ‘measurement problem’ to refer specifically to this deficiency of textbook collapse-on-measurement accounts. That is not how I am using the term here, but there is no real disagreement behind this difference; all agree that an account can reject one of (A1)–(A3), in this case (A2), and yet fail to be an acceptable fundamental physical theory for various reasons. For more on this terminological issue, see (Lewis [2007], pp. 750-1). For other presentations of the measurement problem that are similar to Maudlin’s above, see for example (Albert [1992]; Wallace [2012]; Healey [2017]; Myrvold [2018]).

²I thank an anonymous referee for suggesting this name.

This assumption would be disastrous to deny, which is why it is background.)

Another desideratum for a quantum theory is that it solve the control problem. A theory must reject at least one of the assumptions listed above and explain how the resulting gap is to be filled in—preferably, without introducing ‘measurement’ or ‘preparation’ as basic, unanalysed terms in the dynamical laws. GRW theory rejects (B2). Everettian mechanics rejects (B3). Bohmian mechanics, it turns out, rejects (B1) in favour of:

(B1’) Given determinate inputs, our preparation devices can prepare those inputs in desired conditional quantum states (relativized to an extra parameter) but not desired quantum states simpliciter [**Conditional Preparation**].

Why are (B1)–(B4) incompatible? Briefly, the idea is to consider a certain type of preparation protocol, in which one chooses which quantum states to prepare based on which outcome of some initial measurement one (first-personally) observes. It is straightforward to show that if the assumptions were all true, then these protocols would violate a variant of the ‘no-cloning theorem’ (Wootters and Zurek [1982]) and so would be impossible to implement. But this cannot be right; we implement such protocols all the time. Therefore, at least one of the assumptions must be false.

The main difference between the control problem and the standard measurement problem is the difference between (A1) and (B1). Whereas the standard measurement problem makes an assumption about what quantum states are (what physical properties they determine), the control problem makes an assumption about what we can do with quantum states—about the extent to which we can control the quantum states of the subsystems that we study in the laboratory. It is rare to find, for any interpretation or theory, detailed discussion of preparation and its limits. My hope is that the control problem will motivate such discussion.

Is the control problem just a ‘non-standard’ version of the measurement problem? On the one hand, the measurement problem is often glossed as a problem with interpreting quantum mechanics as a complete physical theory, and most discussions involve some sort of state completeness assumption, suggesting **Completeness** is a defining feature of the problem. If that is the case, then solving the measurement problem does not entail solving the control problem; consider, for example, a theory which denies that quantum states specify

macro-properties (like the readings on measuring devices), but maintains that we can prepare certain micro-systems in certain quantum states.

On the other hand, the control problem invokes **Competent Measurement**, and crucially involves a measurement-type interaction (combined with a preparation). So if the measurement problem is defined as the general puzzle that measurement interactions can be problematic under Schrödinger dynamics, then the control problem should be viewed as a specific aspect of it.

Of course, to some extent this question of how broadly to define ‘the measurement problem’ is a terminological one. What matters, for our purposes, is whether the control problem is of interest to the foundations of quantum mechanics. In this paper, I aim to show that it is.

The plan is as follows. In Section 2, I review the standard measurement problem. In Section 3, I clarify **Preparation** and discuss whether it is already in tension with **Unitarity**. It turns out that this question has been raised under the header of the ‘preparation problem’ (Ruetsche [1996]; Wessels [1997]), so I take the opportunity to review this problem as well. In Section 4, I present the main result of the paper, the incompatibility argument for the control problem. I also compare the control problem with some non-standard versions of the measurement problem that have been raised (Maudlin [1995]; Gao [2019]). In Section 5, I turn to consider how different interpretations may respond to the control problem. The treatment here is necessarily brief, but some basic direction is given. Finally, in Section 6, I relate the problem to another foundational issue. There is a folklore in the literature that the no-cloning theorem rules out the possibility of ‘individual state determination’ (D’Ariano and Yuen [1996]; Vaidman [2015]). In this section, I argue that the folklore is mistaken, and suggest that the right result in the vicinity is actually a variant of the control problem. The payoff is a revised understanding of the physical significance of the no-cloning theorem.

2 The Standard Measurement Problem

We begin by briefly recalling why (A1)–(A4) are jointly untenable. This will set the stage for my own incompatibility argument for (B1)–(B4), which has a similar set-up and structure.

First, the set-up. Zarna is a physicist who specializes in quantum mechanics. Her laboratory

is set up as shown in Figure 1. She has a collection of electrons, and a measuring device which takes those electrons as input and measures their spin in the z direction. We suppose, without loss of generality, that the lab is isolated, and so **Linearity** is applicable.

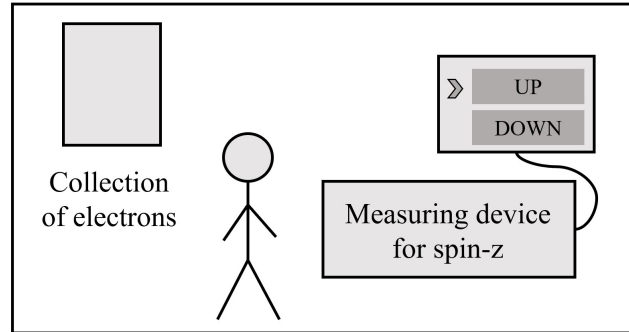


Figure 1: Set-up of Zarna's laboratory.

Suppose for *reductio* that (A1)–(A4) are all true. Suppose Zarna feeds an \uparrow_z electron from her collection into her measuring device. By **Competent Measurement** the device will read 'UP'. By **Completeness**, the final quantum state of the lab will specify this fact. So, letting \longrightarrow denote the evolution of the quantum state of the entire lab from the initial to final time, and letting 'initial' denote the initial quantum state of the rest of the lab, we have,

$$\uparrow_z \otimes \text{initial} \longrightarrow \text{device reads 'UP'}, \quad (1)$$

where the final quantum state of the lab is written 'device reads "UP"' to indicate any state which agrees on this specification of the device. (We stay neutral, for example, on the final state of the measured electron, since this was not specified by our assumptions.)

Next, we ask how the quantum state of the lab would have evolved if the electron were in the \downarrow_z spin state instead. By the same reasoning,

$$\downarrow_z \otimes \text{initial} \longrightarrow \text{device reads 'DOWN'}. \quad (2)$$

Finally, we ask how the quantum state of the lab would have evolved if the electron were \uparrow_x

instead. Since $\uparrow_x = \frac{1}{\sqrt{2}}\uparrow_z + \frac{1}{\sqrt{2}}\downarrow_z$, (1)–(2) and **Linearity** imply the following answer:

$$\uparrow_x \otimes \text{initial} \longrightarrow \frac{1}{\sqrt{2}} (\text{device reads 'UP'}) + \frac{1}{\sqrt{2}} (\text{device reads 'DOWN'}). \quad (3)$$

However, consider any possibility in which the right-hand side of (3) is the final quantum state of the lab. By **Completeness**, this quantum state must specify the reading on the measuring device, if there is one. And now we have a problem. As Maudlin ([1995], p. 8) puts it:

Since ‘UP’ and ‘DOWN’ enter symmetrically into the final state, by what argument could one attempt to show that this device is, in fact, in exactly one of the two indicator states?

So the only plausible option is to insist that, in all these possibilities, there is not a determinate fact of the matter as to what the measuring device reads (and so there is no failure of **Completeness**). But this contradicts **Determinate Outcome**.

As Maudlin acknowledges, ‘symmetry arguments [like this one] are not a matter of logic’ ([1995], p. 10) and so it is not a strict incompatibility. Nevertheless, the argument is highly illustrative. Decades of effort to convincingly remove the inconsistency by denying the symmetry, or by employing more detailed models of the measurement process, have not succeeded.

3 Quantum State Preparation

Like measurement, preparation plays a crucial role in our experimental practice. Countless quantum mechanical experiments, including many foundational Gedanken-experiments, involve steps of the form ‘prepare an ensemble of systems in such-and-such quantum state’.

The function of preparation, however, is rather distinct from that of measurement. In particular, it would be misleading to think of preparation as a special case of measurement. As Margenau ([1963], p. 6) writes:

In general, a preparation of state is any physical operation which assures that, if a system has been subjected to the operation, it will afterward be in a specifiable

quantum state (ψ , or perhaps a mixture of ψ 's). A measurement, on the other hand, [is] a physical operation which yields a number [an outcome] that refers to the state present before the operation.

Here ‘assures’ may be too strong, and we may not wish to require that the operation uniformly prepares the same ψ (or mixture of ψ 's) each time. But the basic contrast is clear, and can be restated in terms of devices. Roughly speaking, a preparation device is an object or physical set-up which, when fed the right sort of subsystem as input, subjects it to a preparation operation. Unlike measuring devices, there is no general requirement that a preparation device display an outcome, or that its final state be correlated with some aspect (initial quantum state, observable) of the input system. Indeed, at least for the purposes of this paper, all we will assume of a preparation device is that it satisfy the conditions of **Preparation**: it occasionally succeeds in outputting its inputs in some pre-specified range of quantum states. (We do require that this range is not trivial. In particular, in the next section we will take **Preparation** to imply that there are at least two ‘distinct’ preparation devices, in the sense that their ranges do not significantly overlap.)

While preparation is not measurement, the issue remains that the success of certain preparation procedures in the lab looks incompatible with the Schrödinger dynamics. This issue has become known as ‘the preparation problem’ (Ruetsche [1996]; Wessels [1997]). I discuss the issue in some detail now. An important upshot is that **Preparation** is nevertheless perfectly compatible with **Unitarity** and **Linearity**. Readers who therefore wish to move on to the incompatibility (B1)–(B4) may skip directly to Section 4. That said, the preparation problem is of independent interest, and provides valuable context for the result that follows.

3.1 The preparation problem

Consider the (apparently successful) preparation procedure depicted in Figure 2. Zarna feeds a beam of electrons through a z -oriented Stern-Gerlach apparatus that has its lower path blocked. She then goes on to perform experiments using the upper beam.

A quick analysis shows that, if **Linearity** is true, then almost none of Zarna’s electrons end up in a spin state close to \uparrow_z . Rather, almost all end up in some distinct (mixed) spin state ρ

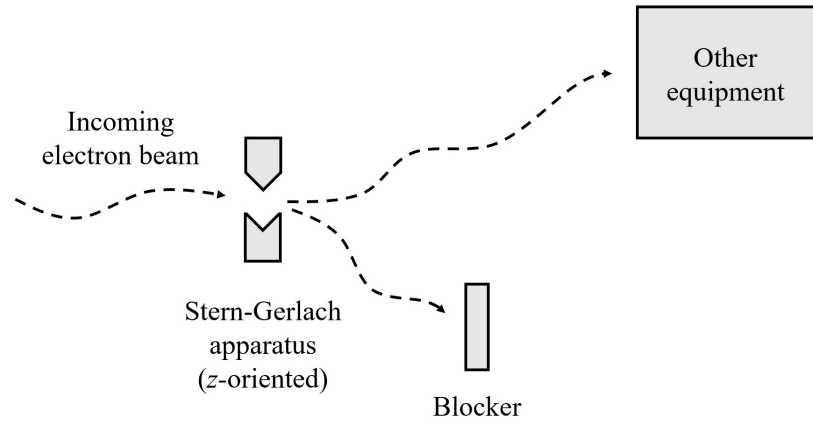


Figure 2: Purported preparation procedure for \uparrow_z .

which varies from electron to electron, depending on the electron's initial spin state $\omega_0 = \alpha \cdot \uparrow_z + \beta \cdot \downarrow_z$.³ Intuitively, this is because **Linearity** forbids simply 'throwing out' the \downarrow_z component (the β part of the superposition) just because it corresponds to hitting the blocker. This component of the superposition, suitably evolved, must also be considered when determining the final spin state of each electron.

At first glance, this result looks disastrous for **Linearity**. If this device does not really prepare \uparrow_z , then what explains why, when we measure the spin- z of its outputs, we obtain measurement statistics that correspond to \uparrow_z ? For example, what explains why, when Zarna places another z -oriented Stern-Gerlach apparatus as well as a scintillation screen along the upper path, she always obtains the outcome 'UP' at this screen?⁴

The result is not as disastrous as it appears. As Wessels argues ([1997], pp. 264-270), one can give a relatively interpretation-neutral story for why we should expect to observe the statistics of the filtered state even though (assuming **Linearity**) each subsystem is actually in an entangled, mixed state. The approach is piecemeal and involves taking into account how the measuring device works in each case. In the Stern-Gerlach case, note that the second Stern-Gerlach apparatus and scintillation screen are spatially localized to the upper beam, far from the blocker. So, while one should expect a position reading at the screen to see

³ Namely, $\rho \approx |\alpha|^2 \cdot |\uparrow_z\rangle\langle\uparrow_z| + |\beta|^2 \cdot \sigma$ where σ depends on how the blocker entangles with the electron's spin degree of freedom; if it doesn't, $\sigma = |\downarrow_z\rangle\langle\downarrow_z|$. See (Wessels [1997], 254) for details.

⁴The worry is not restricted to this Stern-Gerlach case, but extends to other uses of 'filters' such as the double slit experiment (Wessels [1997], p. 255). Unsurprisingly, it's common in the literature to model these filter procedures with a projection-type, non-unitary operation; see for example (Busch *et al.* [1996], §III.9) and (Ruetsche and Earman [2011], §2.2).

contributions from any components of the electron's spatial wave function that are localized to this upper region, one should not expect to see contributions from any components that are localized to the lower region next the blocker. (We assume, as seems reasonable, that the interaction between the electron and scintillation screen is highly local.) And precisely the effect of the first Stern-Gerlach apparatus was to correlate these upper and lower spatial components with the \uparrow_z and \downarrow_z spin components respectively.

This account at least temporarily rescues **Linearity** from making false empirical predictions.⁵ But it raises two more questions. First, if these particular preparation procedures do not really prepare their associated quantum states under the Schrödinger dynamics, are there any procedures which do? And second, if the above preparation procedures still ensure we obtain the measurement statistics corresponding to the desired quantum state \uparrow_z (for example), then why should we be interested in whether they 'really prepare' \uparrow_z ?

The answer to the first question, as foreshadowed, is yes. **Linearity** is compatible with **Preparation**, and so is **Unitarity**. This is thanks to the lack of constraints on the final state of a preparation device. For instance, take the ideal case of a 'uniform' preparation which, given an input in any initial quantum state ω_0 , will output it in the same desired quantum state ψ :

$$\nu \otimes \omega_0 \longrightarrow \nu'(\omega_0) \otimes \psi, \quad (4)$$

for all ω_0 , where ν is the fixed initial quantum state of the device and ν' is its final quantum state, written $\nu'(\omega_0)$ to indicate that it varies with ω_0 . Thanks to the lack of constraints on ν' , there do exist \longrightarrow which satisfy both (4) and **Unitarity** (Wessels [1997], pp. 249-250). Thus even this ideal kind of preparation is compatible, in principle, with **Unitarity**. And though they are not ideal, there are several laboratory procedures of this rough form. A simple example is the placement of a crystal in a cold vacuum, to cool it to its ground energy state.⁶ In fact this is only one of a wide class of 'environment-induced' preparations, a general technique in which one first unitarily tunes the dynamics of the system-plus-environment so

⁵It is possible that we will not be able to offer such a story for every case (see previous footnote), leaving some situations where the measurement statistics must be explained on an interpretation-by-interpretation basis. See also the problem of effect (Section 4.3).

⁶See (Park and Band [1992], §3). A similar procedure can be used to prepare, say, a collection of rubidium atoms in a Bose-Einstein condensate (Healey [2017], §5).

that the desired state-to-be-prepared is the ground state, and then lets the system relax for sufficient time (Fröhlich and Schubnel [2016]). It is also crucial to point out that even though it is inconsistent with **Unitarity** that the ‘filter’ procedures discussed above really prepare the quantum states we associate with them (like \uparrow_z), it is not inconsistent with **Unitarity** that they prepare a mixed state ρ , where ρ varies (non-uniformly) with ω_0 . And so, in light of these observations, I conclude that **Unitarity** and **Preparation** are perfectly compatible.

Schrödinger’s theory does not face the same general problem for preparations as it does for measurements. As Wessels ([1997]) writes:

No analogue of [the measurement problem] arises when quantum mechanics is applied to preparations [...] The preparation problem is not a foundational problem.

Of course, what I will argue next is that while **Unitarity** and **Preparation** are not incompatible taken on their own, they are incompatible when taken in conjunction with **Determinate Input** and **Competent Measurement**. This is the control problem.

Next we turn to the second question. One reason we might be interested (at least from a foundational perspective) in whether we ‘really prepare’ \uparrow_z is that we are interested in our control over quantum states themselves, not just our control over which patterns of outcomes will likely obtain when we make certain measurements. That being said, there may be a worry in the vicinity of this question: assuming we can reject **Preparation** without contradicting our empirical predictions (about the post-preparation statistics we observe), is the control problem, which pits **Preparation** against some other assumptions, really a ‘problem’?

In response to this worry, note that we can reject **Completeness** without contradicting our empirical predictions about measurement statistics. Yet, we do not say that the standard measurement problem is not really a problem. Indeed, just as rejecting **Completeness** requires a systematic account of what the quantum state is missing, and how to make sense of our usual talk of the quantum state as complete, rejecting **Preparation** requires a systematic account of our control over quantum states, and how to recover or translate our usual talk of preparing states. (We will see, for instance, that Bohmian mechanics accomplishes this using the notion of a conditional quantum state.) At least from the viewpoint of a theorist who is starting out

with the naïve Schrödinger account, these are not trivial tasks.

4 The Control Problem

4.1 Set-up

Zarna has installed an additional room in her laboratory, as shown in Figure 3. The room contains two preparation devices, D and D' , as well as a large collection of subsystems. The devices are designed to take these subsystems as input. We assume that the devices are distinct, that is, tend to prepare non-negligibly different quantum states, and can be used repeatedly, that is, can be fed many inputs if desired. Some examples of such devices are provided below. (Later on, it will be clear how ‘non-negligibly different’ and ‘many’ can be precisified in each example.)

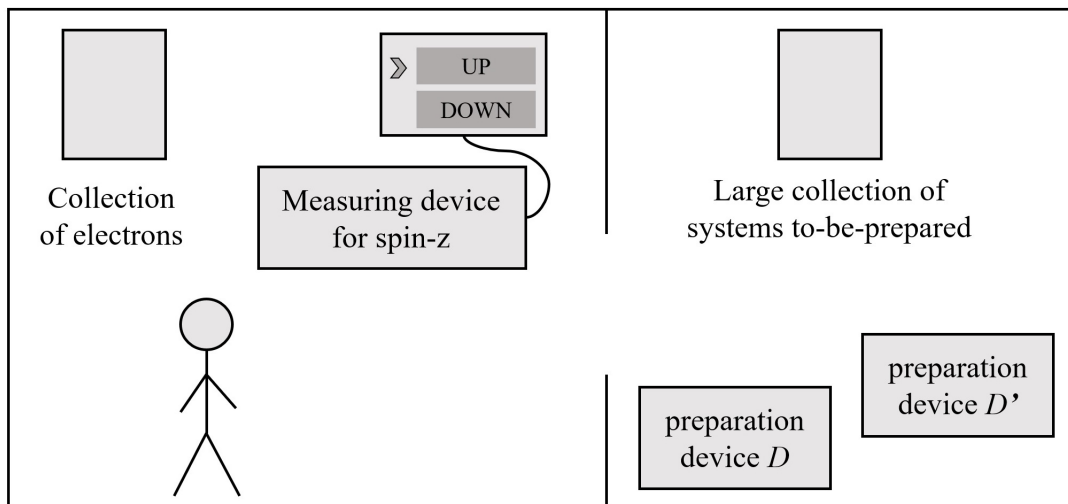


Figure 3: The new set-up of Zarna’s laboratory.

Example 1

The subsystems to-be-prepared are also electrons. D is a filtration which prepares about half of its inputs in the \uparrow_z spin state. D' is a distinct filtration which prepares about half of its inputs in \uparrow_x . As we discussed in Section 3.1, it is not obvious how to implement these filtrations under **Unitarity**. But we include the example, since D and D' do satisfy our criteria of being sufficiently distinct and reusable or repeatable.

Example 2

The subsystems to-be-prepared are crystals. D is an array of vacuums each enclosed in an absorbing wall and cooled as much as possible. They prepare their inputs in a ground energy state ψ or something near to ψ . D' is an array of containers which are heated to prepare their inputs in much more energetic states, so that many or most of their final states ψ' are different from ψ . Here ψ' may vary as long as it is usually distinct from ψ .

Example 3

The subsystems are electrons. D is a Stern-Gerlach apparatus oriented in the z direction and D' is a Stern-Gerlach apparatus oriented in the x direction. As discussed in Section 3.1, this may not be equivalent to Example 1 under **Unitarity**, even if Zarna blocks both the lower paths. Nevertheless, D and D' satisfy our criteria. If an electron is fed into D , it will usually end up in a spin state ρ very different from the ρ' that would have resulted if it were fed into D' .⁷ Also, the devices can be used repeatedly. (An important aspect of this repeatability assumption, which I have not made explicit, is that the repetition does not introduce unwanted entanglement between the prepared systems. This holds automatically if the prepared state is pure as in Examples 1 and 2, and is taken for granted as a property of any good ensemble preparation.⁸ However, given the complications of Section 3.1, it is worth flagging.⁹)

4.2 Incompatibility argument

Suppose for *reductio* that all four assumptions (B1)–(B4) are true. Note that, given the above set-up, it should be physically possible for Zarna to implement the following protocol:

⁷Continuing footnote 3, the z -oriented apparatus outputs $\rho \approx |\alpha|^2 |\uparrow_z\rangle\langle\uparrow_z| + |\beta|^2 |\downarrow_z\rangle\langle\downarrow_z|$ while the x -oriented apparatus outputs $\rho' \approx (1/2)(|\alpha + \beta|^2) |\uparrow_x\rangle\langle\uparrow_x| + (1/2)(|\alpha - \beta|^2) |\downarrow_x\rangle\langle\downarrow_x|$.

⁸Since, usually, the purpose of preparing many systems is so that they provide the basis for multiple independent trials or experiments. Quantum state tomography, for instance, would break down without this assumption (Paris and Rehacek [2004]).

⁹ Some remarks: First, we only require that the prepared subsystems are unentangled with each other; entanglement with the environment or the preparation device is allowed. Second, this unentanglement need only be in the degree of freedom being prepared, so for example if spin is being prepared, then entanglement in the spatial wave function is okay. And third, the prepared systems need only be approximately in a product state; perfect unentanglement is not required. (For the subsequent argument, we only require that having more members of the prepared ensemble makes its state more distinguishable. This is true even when the members are not perfectly unentangled.)

Zarna's protocol: I will feed an electron from my collection into my measuring device. If I observe 'UP',¹⁰ I will walk into the other room (the room containing the preparation devices D and D'), take all the subsystems, and feed them into D . Otherwise, if I do not observe 'UP', I will walk into the other room, take all the subsystems, and feed them into D' instead.

Suppose that, as a matter of fact, Zarna measures an electron which is in the \uparrow_z spin state. We ask how the quantum state of her lab will evolve from this initial time. By **Competent Measurement**, Zarna will observe 'UP'. So, by the protocol, she will end up feeding all the subsystems into D . By **Preparation**, many of these systems will end up, at the final time, in quantum states in the range associated with D . Let us call quantum states which fall within this specified range ' D -states'. Then the evolution of the quantum state of the lab will look like,

$$\uparrow_z \otimes \text{initial} \longrightarrow \text{many systems in } D\text{-states,} \quad (5)$$

where 'initial' is the initial quantum state of the rest of the lab in this particular instance.

Next, we ask how the quantum state of the lab would have evolved if the electron were \downarrow_z . By **Competent Measurement**, Zarna would not observe 'UP'. So, by the protocol, she would feed all her subsystems into D' instead of D . Again applying **Preparation**, we have:

$$\downarrow_z \otimes \text{initial} \longrightarrow \text{many systems in } D'\text{-states.} \quad (6)$$

(It is crucial that the right-hand sides of (5) and (6) are not misread as specifying some property of the preparation devices D or D' ; they are simply specifying the final quantum states of Zarna's subsystems. Importantly, there is no illicit reliance on **Completeness** here.)

Finally, we ask how the quantum state of the lab would have evolved if the electron were in the \uparrow_x spin state instead. By **Unitarity**, this question,

$$\uparrow_x \otimes \text{initial} \longrightarrow \boxed{\quad ? \quad}, \quad (7)$$

¹⁰There is (**Competent Measurement**) always, from a first-person view, an experienced observed outcome of a measurement, even if, from a third-person view, there is no unique determinate outcome.

is well-defined and must have a single definite answer (since \longrightarrow is deterministic, that is, a function from states to states, and ' $\uparrow_x \otimes \text{initial}$ ' is a possible state). Notice that this answer must take one of the following two forms:

$$\boxed{\text{many systems in } D\text{-states}}, \quad (\text{a})$$

or

$$\boxed{\text{many systems in } D'\text{-states}}. \quad (\text{b})$$

To see why, consider any possibility in which Zarna follows her protocol, and suppose toward contradiction that the final quantum state of the lab takes neither of these forms in that possibility. By **Determinate Input**, it cannot be indeterminate whether Zarna feeds her subsystems into a given preparation device at the end of her measurement. Now, if she did not feed the subsystems into either D or D' , then she did not follow the protocol, contrary to our assumption (the protocol always ends with her using one of D or D'). If she fed the subsystems into D , then by **Preparation** the final quantum state of the lab should indeed have form (a), contrary to our assumption. And if she fed the subsystems into D' , then by **Preparation** the final quantum state of the lab should indeed have form (b), contrary to our assumption. In all cases we reach a contradiction, as claimed.¹¹

Now suppose, on the one hand, that the answer is of form (a). Substituting (a) into (7), and putting it side by side with (6), we have:

$$\downarrow_z \otimes \text{initial} \longrightarrow \text{many systems in } D'\text{-states}, \quad (6)$$

$$\uparrow_x \otimes \text{initial} \longrightarrow \text{many systems in } D\text{-states}. \quad (7a)$$

Suppose, on the other hand, that the answer is of form (b). Substituting (b) into (7), and

¹¹Since we do not assume **Completeness**, we do not assume a general one-to-one correspondence between determinate states of affairs of the lab (this or that device being used) and its quantum state. However, we do assume that if two states of affairs involve the determinate use of distinct *preparation* devices, then they have different final quantum states; precisely what distinct preparations do is prepare subsystems in distinct quantum states.

putting it side by side with (5), we have:

$$\uparrow_z \otimes \text{initial} \longrightarrow \text{many systems in } D\text{-states,} \quad (5)$$

$$\uparrow_x \otimes \text{initial} \longrightarrow \text{many systems in } D'\text{-states.} \quad (7b)$$

The problem is that both answers lead to a violation of **Unitarity**. I give the argument for (5) and (7b); the reasoning for (6) and (7a) is the same.

According to **Unitarity**, the inner product of the left-hand sides of (5) and (7b) must equal the inner product of the right-hand sides. The basic observation is that since many D and D' -states were prepared, and many of those states are very different, the right-hand sides will actually be more orthogonal (i.e. more easily distinguishable, i.e. inner product closer to 0) than the left-hand sides. We illustrate this observation using our earlier examples.

Example 1 (cont.)

Recall that D prepares about half of the input electrons in \uparrow_z and D' prepares about half of the input electrons in \uparrow_x . Plausibly, there will be some subset of M electrons that were successfully prepared in both cases, allowing us to rewrite (5) and (7b) as follows:

$$\uparrow_z \otimes \text{initial} \longrightarrow \text{final} \otimes \overbrace{\uparrow_z \otimes \uparrow_z \otimes \dots \otimes \uparrow_z}^{M \text{ times}}, \quad (5-1)$$

$$\uparrow_x \otimes \text{initial} \longrightarrow \text{final}' \otimes \uparrow_x \otimes \uparrow_x \otimes \dots \otimes \uparrow_x, \quad (7b-1)$$

where ‘final’ and ‘final’ denote the final quantum state of everything in the lab except the (spin of the) successfully prepared electrons. **Unitarity** implies,

$$|(\uparrow_z, \uparrow_x)|^2 \cdot |(\text{initial}, \text{initial})|^2 = |(\text{final}, \text{final}')|^2 \cdot |(\uparrow_z, \uparrow_x)|^{2 \cdot M}, \quad (8)$$

which, since $|(\text{initial}, \text{initial})|^2 = 1$ and $|(\text{final}, \text{final}')|^2 \leq 1$, implies,

$$|(\uparrow_z, \uparrow_x)|^2 \leq |(\uparrow_z, \uparrow_x)|^{2 \cdot M}. \quad (9)$$

However, this clearly fails for any $M > 1$, since $|(\uparrow_z, \uparrow_x)|^2 = 1/2$ and $1/2 > (1/2)^M$ for all

$M > 1$. (If we suppose \uparrow_z and \uparrow_x are only approximately prepared, a modified argument still applies.¹²) Recall there was nothing to stop us from stipulating Zarna's collection of electrons to-be-prepared to be as large as we want, suggesting there are no principled limits on M .

Readers may recognize this derivation as a variant of the 'no-cloning theorem' (Wootters and Zurek [1982]). For a review of the theorem and a general statement of this variant, see Theorems 1 and 2(i)–(ii) of the Appendix.

Example 2 (cont.)

Similar reasoning applies. There will be some subset of M crystals which were successfully prepared in both cases:

$$\uparrow_z \otimes \text{initial} \longrightarrow \text{final} \otimes \psi_1 \otimes \psi_2 \otimes \dots \otimes \psi_M, \quad (5-2)$$

$$\uparrow_x \otimes \text{initial} \longrightarrow \text{final}' \otimes \psi'_1 \otimes \psi'_2 \otimes \dots \otimes \psi'_M, \quad (7b-2)$$

where we add subscripts to allow for variation from case to case, the requirement just being that ψ_i and ψ'_i are distinct so that $|(\psi_i, \psi'_i)|^2 < 1$ for all i . For instance if ψ_1 is the ground energy eigenstate and ψ'_1 is some other energy eigenstate then $|(\psi_1, \psi'_1)|^2 \approx 0$ which is already enough to violate **Unitarity**, since again it follows from **Unitarity** that we must have

$$|(\uparrow_z, \uparrow_x)|^2 \leq |(\psi_1, \psi'_1)|^2 \cdot |(\psi_2, \psi'_2)|^2 \cdot \dots \cdot |(\psi_M, \psi'_M)|^2. \quad (10)$$

Otherwise we can always stipulate that Zarna started out with more systems, so that M is plausibly high enough to violate (10). For the variant of the no-cloning theorem that abstracts this result, see Theorem 2(iii) of the Appendix.

Example 3 (cont.)

To model Example 3 we need the formal apparatus of density matrices. Since the upshot is the same, I relegate this more technical discussion to the Appendix, in particular to Theorem 3,

¹²As long as $|(\tilde{\uparrow}_z, \tilde{\uparrow}_x)|^2$ is less than 1, there will be some M above which the inequality is broken. For example, suppose we define an error term ϵ such that $|(\tilde{\uparrow}_z, \tilde{\uparrow}_x)|^2 = 0.5 \pm \epsilon$. Fix $\epsilon < 0.3$. Then for any $M > 25$, $|(\tilde{\uparrow}_z, \tilde{\uparrow}_x)|^{2 \cdot M}$ will be at least two degrees of magnitude smaller than $|(\uparrow_z, \uparrow_x)|^2$. (Note that if $\tilde{\uparrow}_z$ and $\tilde{\uparrow}_x$ are mixed states we rewrite these expressions in terms of the fidelity measure; see the Appendix. We assume the M systems are still approximately a product; see footnote 9.)

which generalizes the above variants of the no-cloning theorem to mixed states.

Does the incompatibility still hold if we replace **Unitarity** by **Linearity**? The natural strategy is to argue that the answer to (7),

$$\boxed{\frac{1}{\sqrt{2}}(\text{many systems in } D\text{-states}) + \frac{1}{\sqrt{2}}(\text{many systems in } D'\text{-states})}, \quad (c)$$

takes neither form (a) nor (b), contrary to **Determinate Input & Preparation**. However this reasoning may fail if the preparation devices are not assumed to work perfectly.¹³ Of course, this isn't to say that **Linearity** is not in tension with the other assumptions—just that the problem is more forceful when formulated in terms of the other major feature of the Schrödinger dynamics, namely its inner product-preserving property.

4.3 Other measurement problems

As I foreshadowed, the control problem crucially involves a measurement-type interaction, and so in this basic sense is a type of measurement problem. It is worth investigating, however, whether it shares any deeper similarities (or differences) with other existing versions of the measurement problem.

As one might expect, there are many versions of the measurement problem. In fact, immediately after presenting the standard measurement problem, Maudlin ([1995]) presents two others. The first, the problem of statistics, replaces (A2) and (A3) with:

(A2-S) the quantum state of an isolated system evolves deterministically;

(A3-S) two measurement situations which begin with the same initial quantum state can sometimes have different (unique) outcomes (where the probability of an outcome is given approximately by Born's rule).

This time (A4) is not needed; the **Completeness** assumption, (A1), does the rest of the work, since it rules out the strategy of distinguishing the two measurement situations using a 'hidden

¹³Consider Example 1. Suppose that whenever D fails to prepare \uparrow_z (about 50% of the time), it outputs \uparrow_x instead. Then (c) could actually be rewritten in form (b), since we could factor out many \uparrow_x states from both terms. While **Unitarity** would still be violated in this case (since there would still be many electrons which D prepares in \uparrow_z but D' prepares in \uparrow_x), **Linearity** would not.

variable' detector state. Here the contrast with the control problem is arguably more stark than before. (A3-S) and (B3) are less similar, and it's not clear there is any analogous reformulation (B3-S) which would bring the assumptions in parallel.¹⁴ And, most importantly, there remains the central difference between (A1) and (B1).

The other problem Maudlin presents, the problem of effect, has a deeper connection to the control problem ([1995], 13):

If the result of the first measurement [in spin- z of an \uparrow_x particle] is 'UP', a second measurement will certainly yield another 'UP', and similarly for 'DOWN'. The result of a measurement therefore has predictive power for the future [...] Any theory which seeks to replicate the empirical content of the traditional theory should have this feature. Let us call this the problem of effect [...].

To solve the problem of effect, a theory must explain why measurement results are repeatable. Strictly speaking, this problem is not about state preparation, since what is being discussed is the determinate record ('UP', 'UP', ...) associated with a measuring device (a record which, without **Completeness**, might not be encoded in any quantum state). However, it is tempting to infer from the record a claim about quantum state preparation, namely that the measuring device has genuinely prepared \uparrow_z . And then one quickly runs into trouble, since this claim contradicts **Unitarity**.¹⁵ In fact we have just circled back to the preparation problem, that certain 'filter' procedures do not really prepare their associated states under **Unitarity**, even though they seem (judging by the subsequent measurement statistics) to have succeeded. For discussion of the preparation problem, see Section 3.1.

Finally we consider a more recent formulation of the measurement problem, due to Gao ([2019]). Like some other presentations (Albert [1992], pp. 77-79; Barrett [1999], §9), this 'mentalistic' reformulation emphasizes the tension between **Linearity** and the claim that observers have unique, determinate experiences. It gets off the ground by reformulating (1)

¹⁴There is, though, a variant of the preparation problem which pits determinism (A2-S) directly against the following analogue of (AS-3): some of our 'filter' preparations for \uparrow_z begin with the same initial quantum state and yet, depending on the trial, do or do not prepare \uparrow_z .

¹⁵If this claim is true, then we have ' $\uparrow_x \otimes \text{initial} \rightarrow \uparrow_z \otimes \text{final}$,' where *initial* and *final* are the initial and final quantum states of everything in the lab except the measured electron. But then by the same reasoning ' $\downarrow_z \otimes \text{initial} \rightarrow \downarrow_z \otimes \text{final}$ '. The pair then contradicts **Unitarity**.

and (2) as follows,

$$\uparrow_z \otimes \text{initial} \longrightarrow \text{Zarna experiences 'UP'}, \quad (1\text{-M})$$

$$\downarrow_z \otimes \text{initial} \longrightarrow \text{Zarna experiences 'DOWN'}. \quad (2\text{-M})$$

Notice that the reformulation assumes that the final quantum state of the lab specifies Zarna's mental state. If we endorse the thesis that mental states supervene on physical states, then this claim follows automatically from (A1), since **Completeness** says the quantum state is the complete physical state. However if we reject physicalism then the two assumptions may come apart. Thus in a careful reformulation we must replace (A1) by:

(A1-M) Our mental states supervene on our quantum states.

The incompatibility is then with (Gao [2019], p. 302):

(A2-M) Same as (A2); the quantum state evolves linearly.

(A3-M) A measurement does not branch mental states; it yields a single mental state representing a single outcome.

(A4-M) Assuming no past mental state branching, implementing a spin- z measurement of an electron in the \uparrow_z quantum state will lead to one experiencing 'UP' as the outcome, and similarly for \downarrow_z and 'DOWN'.

This formulation is most illuminating when applied to non-physicalist interpretations like the Single Mind Interpretation (Albert and Loewer [1988]). This theory endorses (A1) but rejects (A1-M) by positing that an observer's mental state collapses into a particular branch of the wave function even while the physical state continues to evolve according to the Schrödinger equation. The theory thus manages to hold on to (A3-M) even while giving up on (A3).

What would a mentalistic formulation of the control problem look like? The answer reveals a deeper difference between the problems:¹⁶

¹⁶The incompatibility of (B1-M)–(B4-M) follows from an analogous argument, in which Zarna uses the same protocol. Note (B3-M) entails the antecedent to (B1-M) is satisfied in the scenario.

(B1-M) Assuming no past mental state branching, implementing a preparation leads to some of the systems being in the specified quantum states.

(B2-M) Same as (B2); the quantum state evolves unitarily.

(B3-M) Same as (A3-M); no mental state branching (even if there is physical branching).

(B4-M) Same as (A4-M); measurements are at least somewhat reliable and informative.

Note that while the assumption that the mental supervenes on the physical is needed for the mentalistic formulation of the traditional measurement problem, it is not needed for the mentalistic formulation of the control problem. What is needed instead is the assumption that our actions can affect the physical (specifically, the quantum state), and that different mental states (observing ‘UP’, observing ‘DOWN’) can be correlated with different actions (prepare with D , prepare with D'). This latter assumption is much less philosophically controversial, and shows that theories like the Single Mind Interpretation (since they endorse (B2-M)–(B4-M)) don’t get off so easy; in addition to rejecting physicalism, they must reject a certain degree of our control over the physical.

5 Relation to Interpretations

I now briefly classify the three ‘canonical’ quantum theories—Everettian mechanics, Bohmian mechanics, and GRW theory—according to which assumptions (B1)–(B4) they reject. I also discuss some other theories of interest. The main goal here is to begin looking at these theories through the lens of the control problem. The goal is not, however, to settle the question of whether they ultimately solve the problem in a satisfactory way; such a full treatment lies beyond the scope of this paper.

5.1 Everettian mechanics

Everettian mechanics rejects (B3), **Determinate Input**. According to this theory, when Zarna measures the spin of the electron, her lab branches into two ‘worlds’, one in which she observes ‘UP’ and then uses device D , and the other in which she observes ‘DOWN’ and then

uses D' . At the final time, the quantum state of the lab is of neither form (a) ‘many systems in D -states’ nor (b) ‘many systems in D' -states’. But crucially, there is no contradiction with **Preparation** here, since D was not determinately fed the input systems and neither was D' . (Intuitively, since there was not a determinate input, we should not expect the desired output.)

Taken in conjunction with the observation (Section 3.1) that a variety of preparations are not measurements (and so need not cause branching), this suggests that Everettian mechanics does not reject (B1), **Preparation**. There are devices which, as long as they are determinately used, will from a ‘God’s eye view’ prepare the desired quantum state. Other devices, like spin measuring devices, are such that even when determinately used, they only prepare the desired state relative to certain (decohered) branches, in other words only prepare desired ‘relative states’ as opposed to states simpliciter (Wallace [2012], §10).

Of course, if Everettian mechanics is true then realistically the universe has already branched such that most of our uses of preparation devices are not, from a God’s eye view, determinate; as a result, we are usually only preparing relative states in the lab, in any case. Since other than cases of branch re-combining, we only interact (from a first-person view) with ‘our’ world and its future sub-branches, the preparation of relative states is usually all that matters for practical purposes on this theory.

5.2 Bohmian mechanics

In Bohmian mechanics, by contrast, there is a determinate fact of the matter as to whether Zarna fed her systems into the device D or the device D' ; **Determinate Input**, (B3), is true. Nevertheless this device did not prepare her subsystems in the specified quantum state; **Preparation**, (B1), is false. To see why, consider the evolution of the wave function of Zarna’s lab, as predicted by **Unitarity** when the electron starts in \uparrow_x . There will be a component which leads to D -states and a component which leads to D' -states, and its final superposed form will be neither (a) ‘many systems in D -states’ nor (b) ‘many systems in D' -states’, even though Zarna determinately fed the subsystems into one of D or D' .

That being said, suppose one ‘conditions’ the quantum state of these subsystems (see the next paragraph) on the reality of which preparation device was used, that is, on a particular

configuration of Bohmian positions. Then one does get a ‘conditional’ or ‘effective’ state that takes one of these forms (a) or (b). In this sense, Zarna has indeed prepared.

Here the control problem highlights an aspect of Bohmian mechanics that the traditional measurement problems do not. Bohmian mechanics solves the standard measurement problem with its postulation of hidden variables; it solves the problem of effect with its account of the dynamics of those hidden variables; and it solves, or at least addresses, the problem of statistics with its statistical postulate (Albert [1992], 140). But in addition to these features, the theory offers up a new notion of ‘the quantum state of a subsystem’. This new notion is that of the conditional or effective state. Roughly speaking, it is the analogue of relative state from Everettian mechanics, obtained by conditioning the reduced state of the subsystem (what I have simply been calling its quantum state) on the hidden variable configuration of its environment (Dürr *et al.* [1992]; Norsen [2016]). In this theory, the conditional or effective state, not the reduced state, determines the measurement statistics we will (likely) observe.¹⁷ Note that while rejecting **Completeness** gives Bohmian mechanics the resources to define this new notion of a quantum state, it does not fall out automatically from the mere introduction of hidden variables; it takes further work to be motivated and developed. (See for instance (Dürr *et al.* [2005]) for an extension of the notion to density matrices.)

The development of this notion is not obviously needed to solve the above measurement problems. It is crucial, however, for addressing the control problem. Any theory which rejects **Preparation** must give a convincing story of how to recover talk of ‘preparation’, for instance by providing a sense in which our usual claims about preparation of quantum states are true. By motivating this revised notion, Bohmian mechanics makes the rejection of **Preparation** less a full-blown rejection of our ability to prepare, and more a rejection of the assumption that the ‘quantum state of a subsystem’ must be its reduced state or, more generally, must supervene only on the wave function or density matrix of the parent system. In this way, Bohmian mechanics successfully replaces **Preparation** with **Conditional Preparation**.

¹⁷In some situations, the conditional and reduced states align, namely when the subsystem is appropriately isolated from its environment (so that conditioning on the configuration of the environment does not make a difference). This happens for instance when the reduced state is pure, as in the scenario described by **Competent Measurement**. In this scenario, the initial reduced spin state of the electron is \uparrow_z , implying that the initial conditional spin state is also \uparrow_z . Thus Bohmian mechanics still predicts we always observe ‘UP’ in this scenario, as the **Competent Measurement** assumption requires.

5.3 GRW theory

GRW theory (Ghirardi *et al.* [1986]) rejects (B2), **Unitarity**. According to this theory, upon Zarna's initial measurement the quantum state of the lab very likely spontaneously collapses into either a state in which the device reads 'UP' or a state in which it reads 'DOWN'. (GRW endorses **Completeness**, so the quantum state must specify these facts.) Zarna then goes on to prepare using D or D' , respectively, and since the quantum state specifies exactly one of these situations, and the quantum state is complete, **Determinate Input** holds.¹⁸ The devices then genuinely prepare according to **Preparation**. Depending on how they operate, they may or may not induce more spontaneous collapses as part of their preparation process.

So far so good. But while GRW avoids the incompatibility by denying (B2), its approach raises another challenge. GRW permits non-unitary evolutions, yet we do not usually say that if GRW theory is true, then cloning is possible. So what is the sense in which GRW forbids the construction of cloning machines? The answer must be substantive, but it cannot rule out Zarna's measure-then-prepare schemes, since we know these are possible.

The needle is not trivial to thread. According to this theory (and other collapse theories), such schemes really do end with preparing many copies of the desired state, making them look like a kind of cloning process. To press the point, suppose we configure Zarna's set-up so that the lab itself has inputs and outputs like a machine, as shown in Figure 4. Then we can pick a scheme such that the machine behaves as follows:

The user feeds the machine an electron in either one of the non-orthogonal states \uparrow_z or \uparrow_x as input (top left arrow of Figure 4). With about 71% probability, the machine outputs nothing. But with about 29% probability, the machine outputs many (can be set to any desired finite number) electrons from one of its exits (bottom right arrows). Assuming GRW, these outputs are guaranteed to be genuine copies of the input, that is, to have the same spin state.

In some sense, this machine clones. But it is perfectly possible to build.¹⁹

¹⁸Here I ignore complications related to the tail problem (Albert and Loewer [1990]).

¹⁹Zarna carries out the unambiguous state discrimination in (Nielsen and Chuang [2002], p. 92); if she obtains the POVM result corresponding to \uparrow_z , she feeds the much larger collection of electrons (we can set this as arbitrarily large as we want) into the prepare- z -up device shown, and similarly for \uparrow_x . If she obtains the 'inconclusive' result (probability $1/\sqrt{2}$), she does nothing.

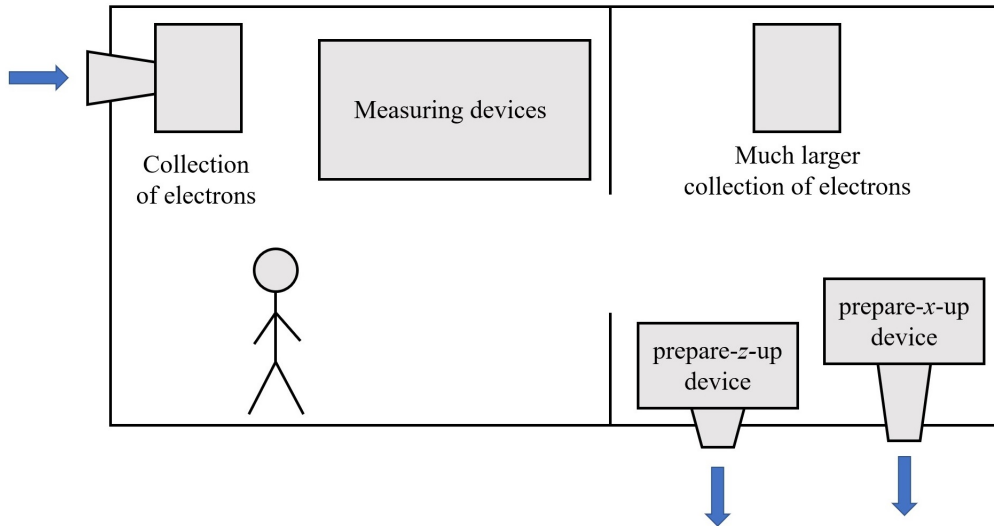


Figure 4: Cloning non-orthogonal spin states with Zarna's laboratory.

A promising response to these cases is to suggest that GRW only forbids the construction of machines which 'unitarily clone' quantum states. The basic idea is that we should allow for cloning by means of probabilistic collapse, but forbid cloning by means of a unitary process. After all, it is really the latter type of cloning which the no-cloning theorem deems problematic.²⁰

Although this response is on the right track, there is an initial complication regarding what to say about intermediate cases, like machines which clone spin states by means of an 'approximately unitary' process. Since the spatial wave function evolves indeterministically in GRW (regardless of whether a measurement is being made), small entanglements between the spatial and spin wave functions can make it so that the spin state dynamics is also indeterministic and hence non-unitary, but still approximately unitary, in the sense that the chance of a spin state collapse is still very small. Intuitively, although there is no unitary process here, the cloning of spin states should still be forbidden to some extent. To address this complication, a relevant starting place may be the literature on 'approximate cloning' and its limits (Werner [1998]; Scarani *et al.* [2005]), as well as the literature on 'probabilistic cloning' and its limits (Duan and Guo [1998a, 1998b]; Pati [2000]; Chen *et al.* [2011]).²¹

²⁰I thank an anonymous referee for suggesting a response along these lines, and for their helpful comments on this part of the discussion.

²¹The approximate cloning literature considers imperfections in the cloning process, but a drawback (for current purposes) is that it still assumes the process is deterministic. The probabilistic cloning literature argues that it is

5.4 Other quantum theories

One of the interesting differences between the control problem and the standard measurement problem is that the control problem makes no assumption about what quantum states specify or represent. Thus, we should expect the problem to be particularly revealing when applied to quantum theories of a more subjectivist, pragmatist or instrumentalist flavour, which deny that quantum states directly represent physical properties, but maintain that we measure and prepare them.

There are several factors which complicate the discussion of these theories. One is whether they accept **Unitarity** as formulated. Another is whether they accept a premise presupposed in the arguments above:

(A5, B5) All physical systems—including, for instance, Zarna’s lab—have a unique, determinate quantum state (wave function or density matrix) at a given time
[Determinate State].

Quantum Bayesianism (Caves *et al.* [2002]), for instance, rejects this assumption. Of course, in Quantum Bayesianism we can still attribute quantum states to systems as long as we also specify a third parameter: the epistemic perspective of a particular agent. Thus it is still possible to re-run the argument on this theory, by relativizing all the states to the perspective of a ‘friend’ outside the lab. (For hints of how a Quantum Bayesian solution to this revised problem might go, see the Quantum Bayesian analysis of how to translate talk of ‘not knowing which quantum state was prepared’ in tomography scenarios (Fuchs and Schack [2004]).) In any case, the point is that this is a complication which must be kept in mind.

Another theory of interest is Healey’s ([2017]) pragmatist interpretation, according to which quantum states prescribe rational degrees of belief but need not specify physical properties (**Completeness** is not assumed). Two agents may rationally assign different quantum states to the same system (**Determinate State** is false), but only if they are in different physical situations. In a Wigner’s friend scenario, the quantum state relative to the friend outside the

possible to clone non-orthogonal quantum states as long as the cloning is probabilistic in nature, and then derives several limitations on this kind of process, including maximum probabilities for success and maximum cloning efficiencies. None of these results are formulated in the GRW framework, and so they do not take into account the subtleties of the spontaneous collapse theory. But they provide a possible starting point. (It should be noted that, incidentally, Ghirardi himself was one of the early discoverers of the no-cloning theorem (Peres [2003]).)

lab satisfies **Unitarity** (Healey [2017], §11). Thus there is a question of how to cash out **Preparation** so that, when applied to Zarna's lab, there is no inconsistency with unitarity, yet we still recover a sense in which (even according to her friend) Zarna successfully prepares.

Of course, there are many other theories of interest, including Bub's information-theoretic interpretation (Bub [2005]), Rovelli's relational interpretation (Rovelli [1996]), and Peres's operationalist view (Peres [2006]). I do not purport to adequately treat these frameworks here. I leave it as an open question how they solve, or dissolve, the control problem.

6 Relation to Foundations and the No-Cloning Theorem

Let us now switch gears from interpretations to another topic in quantum foundations. It is customary, upon introducing the no-cloning theorem (Wootters and Zurek [1982]; Scarani *et al.* [2005]), to make the following remark: If the cloning of quantum states were possible, then it would be possible to reliably determine the initial, previously unknown quantum state of a single system. One would simply create many copies of the state, measure the copies along various bases, and then use the statistics to reliably reconstruct the state.

It is also common to encounter a remark in the converse direction (D'Ariano and Yuen [1996]; Vaidman [2015]): if it were possible to reliably determine the initial, previously unknown quantum state of a single system (with no other copies available), then it would be possible to clone in the sense forbidden by the original no-cloning theorem. Contrapositively: the no-cloning theorem rules out the possibility of individual state determination.

If both of these claims were true, the biconditional would suggest an attractively straightforward interpretation of the physical significance of the no-cloning theorem. According to this interpretation, the theorem gains physical significance, at least in part, because of its implications for the epistemology of 'individual' or 'single' quantum states.

In this section, I argue that the converse claim is mistaken; at the very least, the standard argument offered for this claim is unsound. Drawing on a variant of the control problem, I suggest a revised view of the connection between the no-cloning theorem and state determination. On my revision, the connection is not specific to the case of 'individual' or 'single' quantum states.

6.1 Individual state determination

Vaidman ([2015], p. 18) invokes the standard argument for this converse claim in his discussion of protective measurement:

If we are given a single quantum system in an unknown quantum state, there is no way to find out what this state is. It would contradict the no-cloning theorem: if we can find out what the state is, we can prepare many other systems in this state.

The argument traces back at least to D’Ariano and Yuen, who spell out the reasoning in more detail ([1996], pp. 2832-3):

In this Letter, we will show the impossibility of any measurement scheme for determining the wave function from a single copy of the system [...] A quantum cloning machine [...] has to be represented by a state transformation of the general form,

$$v \otimes \psi \otimes \underbrace{\omega_1 \otimes \dots \otimes \omega_{M-1}}_{M>1} \longrightarrow v'(\psi) \otimes \underbrace{\psi \otimes \dots \otimes \psi}_{M \text{ times}},$$

where both the transformation and the state preparation must be independent of ψ —which *a priori* is unknown [...] [By the no-cloning theorem,] one cannot clone the state of a system known to be in *any one* of two nonorthogonal states, say, ψ and φ , using a single state-independent unitary transformation. Let us now suppose that there is a device that is capable of determining (by a unitary process) in which of these two states the system is. Then, once the state is known, one can easily choose unitary transformations that generate an M clone of ψ or φ , depending on the two possible results of this device. Hence, such a ‘state detector’ would lead to a realization of the cloning machine.

Upon first inspection, the reasoning looks compelling. However, what the authors do not point out is that state transformations of the form,

$$v \otimes \underbrace{\psi \otimes \dots \otimes \psi}_{N \text{ times}} \otimes \underbrace{\omega_1 \otimes \dots \otimes \omega_{M-N}}_{M>N} \longrightarrow v'(\psi) \otimes \underbrace{\psi \otimes \psi \otimes \dots \otimes \psi}_{M>N \text{ times}}, \quad (11)$$

are also in violation of unitarity (see Theorem 2(ii) of the Appendix). This time the ensemble size N can be any finite number, so the result is not restricted to single systems. Indeed it was perhaps misleading to frame the no-cloning theorem as the result that ‘a single quantum cannot be cloned’ (Wootters and Zurek [1982]), given that a finite ensemble of identical quanta also cannot be cloned.

This is bad news for the argument. For if the argument were sound as stated, then it would also demonstrate the impossibility of ‘any measurement scheme for determining the wave function’, not just ‘from a single copy of the system’, but also from any finite number N copies of the system. (Recall there were no *a priori* restrictions on the number of prepared systems M .) But unless ‘determining’ is read in an overly stringent way, this result is absurd. When N is large, we can reliably determine the initial wave function. So something has gone wrong.

One line of response to this difficulty is to try to develop a different argument for the impossibility of individual state determination—in particular, one which is also based on cloning concepts, but does not overgeneralize to large N . (Alternative arguments for this conclusion in the literature, for instance those based on state disturbance (Fuchs [1998]) and superluminal signalling (Busch [1997]), do not overgeneralize.) However, while there exist promising strategies for pursuing this response, they involve going significantly beyond the scope of the no-cloning theorem, for instance by dropping the assumption that \longrightarrow is a deterministic map from states to states, or by dropping the assumption that the cloned states are a separable ensemble.²² And so this line of response leaves unanswered the question of how to understand the physical significance of the original theorem and its close variants, whose proof and formulation crucially rely on these features.

Another line of response is to insist that the basic argument is valid, but that rather than ruling out the possibility of individual state determination, it simply rules out the possibility of deterministic, perfectly reliable state determination for $N = 1$ and more generally $N < \infty$.

While this response appears to avoid the slippery slope worry I raised above, I suggest that

²²One strategy is to use the limits of ‘probabilistic cloning’ (Duan and Guo [1998a]) to derive limits on reliable (yet fallible) state determination as a function of N . Another strategy is to use the limits of ‘approximate’ or ‘imperfect’ cloning (Werner [1998]) to do the same thing. The latter has essentially been carried out, by Bae and Acín ([2006]). As they note, one must allow a high degree of entanglement between the outputs, which is not the case in preparation; on this issue, see also footnote 9.

this revised argument still misses the mark. The reason is that its preparation-related premises are still going to be in tension, not just with individual or perfect state determination, but with unproblematic, run-of-the-mill state determination. Indeed, the argument appears to take (B1)–(B3) for granted. There is the assumption that \longrightarrow is unitary, (B2). And there is the assumption that we can determinately prepare distinct non-orthogonal states as a function of what state determination result we obtain, which amounts to the conjunction of (B1) & (B3). But these assumptions, though individually innocent, are together incompatible with:²³

(B4') if we have a sufficiently large (finite) ensemble of subsystems which we know to be in a product of the same initial wave function, then we can come to know, with reasonable reliability and to a good approximation, what that initial wave function is [**Competent State Determination**].

(Note that (B4'), like (B4), would be disastrous to deny. All interpretations, including Bohmian and Everettian mechanics, endorse it.²⁴ What these unitary theories may deny, though, is that the antecedent of (B4') is ever realistically satisfied. On these theories, we are typically only ever going to know that an ensemble is in a product of the same effective wave function or relative quantum state, respectively.²⁵)

Drawing on this observation, I suggest an alternate view of the connection between the no-cloning theorem and state determination. Note that the source of the incompatibility with

²³The derivation proceeds much the same as in Section 4, except for the final step, which invokes the reliability condition in (B4'). Consider the set-up of Section 4.1, except Zarna is given state determination equipment, and informed that her collection of electrons is an N -product of some 'source state'; some wave function. Her protocol: 'If I end up confident that the source state is approximately \uparrow_z , I will feed all my many ($\gg N$) subsystems from the other room into D . Otherwise, I will feed them into D' .' Suppose the source state is \uparrow_z . We have (B1, B2, B4'): ' $\uparrow_z^{\otimes N} \otimes \text{initial} \longrightarrow \text{many systems in } D\text{-states}$ '. But now we ask how the quantum state of the lab would have evolved if the source state were ψ where ψ is non-orthogonal to \uparrow_z and non-negligibly different from it: ' $\psi^{\otimes N} \otimes \text{initial} \longrightarrow \boxed{\text{?}}$ '. The answer 'many systems in D' -states' is forbidden by (B2) for the same no-cloning reasons as before. But (B1, B3) entail that the answer is not of this form only if Zarna did not feed her subsystems into D' , which by the protocol means she is confident the source state was \uparrow_z . So if the source state were ψ , Zarna would still be confident it is \uparrow_z . And this violates the reliability clause in (B4').

²⁴It is important, for the claim to be plausible, that the systems do not start out unduly entangled with their environments, especially with the relevant measuring devices. I have bypassed these complications by restricting attention to pure states (wave functions). A more general formulation with density matrices should make sure to rule out cases where the systems have the same reduced state, but (thanks to environment entanglement) different Bohmian conditional states or Everettian relative states. For another place in the paper where this complication is relevant, see footnote 17.

²⁵See Sections 5.1 and 5.2. Regarding Bohmian mechanics, note that this claim—that we rarely know the quantum state simpliciter of a subsystem, but we do often know its effective wave function—is the counterpart to the claim that **Preparation** is false but **Conditional Preparation** is true.

(B4') is, as in our previous incompatibility result, the no-cloning theorem. In this way, the no-cloning theorem does have implications for state determination. On the one hand, if the kind of cloning described by the no-cloning theorem were possible, then in principle there would be no bounds on our state determination or preparation abilities. On the other hand, since the kind of cloning described by the no-cloning theorem is impossible, there is (assuming **Unitarity**) a tension between these abilities. This tension, put in more general terms below, is what I have been attempting to highlight throughout the paper, in the form of the control problem.

7 Conclusion

In this article I have argued that there is another problem for quantum mechanics, namely the control problem. Fortunately, the control problem does not straightforwardly rule out any major existing interpretations of quantum mechanics. However, it brings to the foreground a number of issues requiring attention, the most central being quantum state preparation. As we have seen, the control problem puts pressure on our current theories to spell out more carefully the limits of quantum state preparation. Thanks to the incompatibility (B1)–(B4), any interpretive proposal which neglects state preparation, or simply subsumes it under measurement, cannot be complete.

The problem also places a constraint on our future theorizing. Whereas **Completeness** is a claim about what quantum states are or what quantum states represent, **Preparation** is a claim about what we can do with quantum states. As I have remarked, these two issues can come apart. Thus a candidate future theory might successfully avoid the traditional measurement problem, but end up being ruled out by the control problem.

I conclude with one more gloss of the control problem, which I find illuminating. Suppose we hold fixed **Unitarity** and **Determinate Input** as background assumptions. From this point of view, what the control problem shows is that there is a tension between two of our alleged abilities as agents. First, our ability to gain any sort of reliable information about quantum states. Second, our ability to systematically control quantum states. This tension (precisified as the incompatibility between **Competent Measurement** or **Competent State**

Determination, on the one hand, and **Preparation**, on the other) exists regardless of whether we posit hidden variables in addition to quantum states (regardless of **Completeness**). Given our two background assumptions, it is a fundamental aspect of our epistemic relationship to quantum states; an aspect which is captured formally by the no-cloning theorem.

Appendix

We begin by recalling the original no-cloning theorem.

Definition 1

A linear map $U : \mathcal{H} \rightarrow \mathcal{H}$ is **unitary** if it has an inverse and $(U\psi, U\varphi) = (\psi, \varphi)$ for any vectors ψ and φ .

Theorem 1 (Original No-Cloning)

Let $\mathcal{H} = \mathcal{H}_1 \otimes \mathcal{H}_2$ and let U be a map from \mathcal{H} to itself. Fix any ψ, φ non-orthogonal and distinct. Suppose

$$U(\psi \otimes \nu) = \psi \otimes \psi$$

$$U(\varphi \otimes \nu) = \varphi \otimes \varphi$$

for some ν . Then U is not unitary (or linear).

Next we prove some close variants of the no-cloning theorem. We focus first on pure vector states, then generalize to mixed states.

Theorem 2 (Some Relevant Variants)

Let $\mathcal{H} = \otimes_{i \in I} \mathcal{H}_i$ and let U be a map from \mathcal{H} to itself. Fix any ψ, φ non-orthogonal and distinct. Then U is not unitary if:

(i) for some ν ,²⁶

$$U(\psi \otimes \nu) = f(\psi) \otimes \psi \otimes \psi \otimes \dots$$

$$U(\varphi \otimes \nu) = f(\varphi) \otimes \varphi \otimes \varphi \otimes \dots,$$

(ii) or more generally, for some ν and $M > N$,

$$U(\overbrace{\psi \otimes \dots \otimes \psi}^{N \text{ times}} \otimes \nu) = f(\psi) \otimes \overbrace{\psi \otimes \dots \otimes \psi}^{M \text{ times}} \otimes \dots$$

$$U(\varphi \otimes \dots \otimes \varphi \otimes \nu) = f(\varphi) \otimes \varphi \otimes \dots \otimes \varphi \otimes \dots,$$

(iii) or, more generally, there are pairs $\{\chi_j, \chi'_j : j = 1, 2, \dots, M\}$ such that $\chi_j \neq \chi'_j$ for each j and, for some ν ,

$$U(\overbrace{\psi \otimes \dots \otimes \psi}^{N \text{ times}} \otimes \nu) = f(\psi) \otimes \chi_1 \otimes \dots \otimes \chi_M \otimes \dots$$

$$U(\varphi \otimes \dots \otimes \varphi \otimes \nu) = f(\varphi) \otimes \chi'_1 \otimes \dots \otimes \chi'_M \otimes \dots,$$

where M exceeds the finite bound $N \cdot \log [|(\psi, \varphi)|^2] / \log [\max_j |(\chi_j, \chi'_j)|^2]$.

Proof

All of (i)–(iii) follow as corollaries to the next theorem, in particular the case $\rho = |\psi\rangle\langle\psi|$, $\sigma = |\varphi\rangle\langle\varphi|$. (Note in this case $F(\rho, \sigma) = |(\psi, \varphi)|^2$.) \square

To generalize to mixed states,²⁷ we need some more technical apparatus. Though generalizations to mixed states have been made (Barnum *et al.* [1996]), I am not aware of presentations of these particular variants. (Below we also relax from *unitary maps* to *quantum channels*. Physically this suggests we can relax the premise that Zarna's lab is isolated, since many open systems obey a reduced dynamics of this more general form (Rivas and Huelga [2012]).)

Definition 2

²⁶Throughout we just require $\nu \in \otimes_{i \geq 2} \mathcal{H}_i$ and $f(\psi), f(\varphi) \in \otimes_{i \in S} \mathcal{H}_i$ for some $S \subset I \subseteq \mathbb{N}$.

²⁷I restrict myself to an ordinary QM setting in which all mixed states are represented by density operators; for settings where this can break down, see (Ruetsche [2004]).

Let $\mathcal{D}(\mathcal{H})$ denote the convex set of density operators on \mathcal{H} . A map $\mathcal{E} : \mathcal{D}(\mathcal{H}) \rightarrow \mathcal{D}(\mathcal{H})$ is a **quantum channel** if it is *convex-linear* and *completely positive*; see (Nielsen and Chuang [2002], §8) for details. Two density operators $\rho, \sigma \in \mathcal{D}(\mathcal{H})$ are **non-orthogonal** if it's not the case that $\text{supp}(\rho) \perp \text{supp}(\sigma)$ (here $\text{supp}(\rho) \equiv (\ker(\rho))^\perp$). The **fidelity**, or closeness, between two density matrices is $F(\rho, \sigma) \equiv (\text{Tr} \sqrt{\rho^{1/2} \sigma \rho^{1/2}})^2$.

Definition 3

Fix $\mathcal{H} = \otimes_{i \in I} \mathcal{H}_i$ and $\rho \in \mathcal{D}(\mathcal{H})$. Given some subsystem $S \subset I$, let $\mathcal{R}_S(\rho)$ denote its **reduced state**, defined as the partial trace $\text{Tr}_{\bar{S}}(\rho)$ of the parent state ρ over all elements $i \in I$ not in S (Nielsen and Chuang [2002], §2.4).

Theorem 3 (Generalization)

Let $\mathcal{H} = \otimes_{i \in I} \mathcal{H}_i$ and let \mathcal{E} be a map from $\mathcal{D}(\mathcal{H})$ to itself. Fix any ρ, σ non-orthogonal. Then \mathcal{E} is not a quantum channel if:

(i) there is a subsystem $S \subset I$ such that, for some μ ,²⁸

$$\mathcal{R}_S[\mathcal{E}(\rho \otimes \mu)] = \rho \otimes \rho, \quad \mathcal{R}_S[\mathcal{E}(\sigma \otimes \mu)] = \sigma \otimes \sigma,$$

(ii) or, more generally, there is $S \subset I$ such that, for some μ and $M > N$,

$$\mathcal{R}_S[\mathcal{E}(\rho^{\otimes N} \otimes \mu)] = \rho^{\otimes M}, \quad \mathcal{R}_S[\mathcal{E}(\sigma^{\otimes N} \otimes \mu)] = \sigma^{\otimes M},$$

(iii) or, more generally, there is a set of pairs $\{\pi_j, \pi'_j : j = 1, 2, \dots, M\}$ with $\pi_j \neq \pi'_j$ for each j , and a subsystem $S \subset I$ such that, for some μ ,

$$\mathcal{R}_S[\mathcal{E}(\rho^{\otimes N} \otimes \mu)] = \pi_1 \otimes \dots \otimes \pi_M, \quad \mathcal{R}_S[\mathcal{E}(\sigma^{\otimes N} \otimes \mu)] = \pi'_1 \otimes \dots \otimes \pi'_M,$$

where M exceeds the finite bound $N \cdot \log [F(\rho, \sigma)] / \log[\max_j F(\pi_j, \pi'_j)]$.

Proof

To see the general strategy we prove (i) first. Then we prove (iii), which immediately entails

²⁸Here we have $\mu \in \mathcal{D}(\otimes_{i \geq 2} \mathcal{H}_i)$.

(ii).²⁹ For (i), suppose toward contradiction that \mathcal{E} is a quantum channel. First we invoke the fact that fidelity is *monotonic* under quantum channels (Nielsen and Chuang [2002], Theorem 9.6), that is

$$F(\rho \otimes \mu, \sigma \otimes \mu) \leq F(\mathcal{E}(\rho \otimes \mu), \mathcal{E}(\sigma \otimes \mu)). \quad (12)$$

Since fidelity is also monotonic under partial trace (indeed, partial trace can itself be viewed as a quantum channel or operation, and so this follows from the fact above (Nielsen and Chuang [2002], §8.3.1)), we also have

$$F(\mathcal{E}(\rho \otimes \mu), \mathcal{E}(\sigma \otimes \mu)) \leq F(\mathcal{R}_S[\mathcal{E}(\rho \otimes \mu)], \mathcal{R}_S[\mathcal{E}(\sigma \otimes \mu)]), \quad (13)$$

which under our assumptions imply

$$F(\rho \otimes \mu, \sigma \otimes \mu) \leq F(\rho \otimes \rho, \sigma \otimes \sigma). \quad (14)$$

But since fidelity distributes over products (Miszcak *et al.* [2009], §2), this implies

$$F(\rho, \sigma) \cdot F(\mu, \mu) \leq F(\rho, \sigma) \cdot F(\rho, \sigma). \quad (15)$$

And now we reach a contradiction. The fact that ρ and σ are non-orthogonal implies $0 < F(\rho, \sigma)$ (Miszcak *et al.* [2009], §2), and since ρ and σ are distinct, we have $F(\rho, \sigma) < 1$. But then it follows that $F(\rho, \sigma)^2 < F(\rho, \sigma)$ which, since $F(\mu, \mu) = 1$, contradicts the above.

For (iii), we use the same reasoning to obtain the inequality

$$F(\rho^{\otimes N}, \sigma^{\otimes N}) \cdot F(\mu, \mu) \leq F(\pi_1, \pi'_1) \cdot \dots \cdot F(\pi_M, \pi'_M). \quad (16)$$

Since the left-hand simplifies to $F(\rho, \sigma)^N$ and since each term in the right-hand is less than $\max_j F(\pi_j, \pi'_j)$, we obtain

$$F(\rho, \sigma)^N \leq (\max_j F(\pi_j, \pi'_j))^M \quad (17)$$

which implies $M \leq N \cdot \log(F(\rho, \sigma)) / \log(\max_j F(\pi_j, \pi'_j))$. Since π_j is distinct from π'_j for all i ,

²⁹Since then $N \cdot \log[F(\rho, \sigma)] / \log[\max_j F(\pi_j, \pi'_j)] = N \cdot \log[F(\rho, \sigma)] / \log[F(\rho, \sigma)] = N$.

we have $0 \leq \max_j F(\pi_j, \pi'_j) < 1$, and thus the bound is finite, where we use the fact that N is finite and ρ and σ are non-orthogonal, so $F(\rho, \sigma) < 1$ (Miszczak *et al.* [2009], §2). (Note that if $\max_j F(\pi_j, \pi'_j) = 0$ the denominator diverges; in this case we must have $M = 0$.) \square

Acknowledgements

I am very grateful to Hans Halvorson, Laura Ruetsche, and Adam Elga for their extensive guidance and feedback on this project. I would also like to thank Jeff Bub, Neil Dewar, Ben Feintzeig, Chip Sebens, Brad Skow, and two anonymous referees for their helpful comments on earlier drafts. Additionally I would like to thank David Albert, Thomas Barrett, Sheldon Goldstein, Tim Maudlin, Christopher Timpson, and David Wallace for their valuable conversation and/or correspondence surrounding this topic. Finally, I would like to thank audiences at the University of Washington philosophy of science reading group (spring 2018), University of Michigan foundations of modern physics reading group (fall 2018), Princeton Philosophical Society (spring 2019), Princeton dissertation seminar (spring 2019), and early career workshop in philosophy of physics at the City College of New York (spring 2019).

Alexander Meehan

Princeton University

Department of Philosophy

Princeton, NJ 08544

alexandermeehan@princeton.edu

References

- Albert, D. Z. [1992]: *Quantum Mechanics and Experience*, Harvard University Press.
- Albert, D. Z. and Loewer, B. [1988]: ‘Interpreting the Many Worlds Interpretation’, *Synthese*, 77(2), pp. 195–213.
- Albert, D. Z. and Loewer, B. [1990]: ‘Wanted Dead or Alive: Two Attempts to Solve Schrödinger’s Paradox’, in *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, vol. 1990, pp. 277–285. Philosophy of Science Association.

- Bae, J. and Acín, A. [2006]: ‘Asymptotic Quantum Cloning Is State Estimation’, *Physical Review Letters*, **97**(3), pp. 030402.
- Barnum, H., Caves, C. M., Fuchs, C. A., Jozsa, R. and Schumacher, B. [1996]: ‘Noncommuting Mixed States Cannot Be Broadcast’, *Physical Review Letters*, **76**(15), pp. 2818.
- Barrett, J. A. [1999]: *The Quantum Mechanics of Minds and Worlds*, OUP Oxford.
- Bub, J. [2005]: ‘Quantum Mechanics Is About Quantum Information’, *Foundations of Physics*, **35**(4), pp. 541–560.
- Busch, P. [1997]: ‘Is the Quantum State (an) Observable?’, in R. S. Cohen, M. Horne and J. Stachel (eds), *Potentiality, Entanglement and Passion-at-a-Distance*, Springer, pp. 61–70.
- Busch, P., Lahti, P. J. and Mittelstaedt, P. [1996]: *The Quantum Theory of Measurement*, Springer.
- Caves, C. M., Fuchs, C. A. and Schack, R. [2002]: ‘Quantum Probabilities as Bayesian Probabilities’, *Physical Review A*, **65**(2), pp. 022305.
- Chen, H., Lu, D., Chong, B., Qin, G., Zhou, X., Peng, X. and Du, J. [2011]: ‘Experimental Demonstration of Probabilistic Quantum Cloning’, *Physical Review Letters*, **106**(18), pp. 180404.
- D’Ariano, G. M. and Yuen, H. P. [1996]: ‘Impossibility of Measuring the Wave Function of a Single Quantum System’, *Physical Review Letters*, **76**, pp. 2832–2835.
- Duan, L. M. and Guo, G. C. [1998a]: ‘Probabilistic Cloning and Identification of Linearly Independent Quantum States’, *Physical Review Letters*, **80**(22), pp. 4999.
- Duan, L. M. and Guo, G. C. [1998b]: ‘A Probabilistic Cloning Machine for Replicating Two Non-Orthogonal States’, *Physics Letters A*, **243**(5-6), pp. 261–264.
- Dürr, D., Goldstein, S., Tumulka, R. and Zanghí, N. [2005]: ‘On the Role of Density Matrices in Bohmian Mechanics’, *Foundations of Physics*, **35**(3), pp. 449–467.

- Dürr, D., Goldstein, S. and Zanghi, N. [1992]: ‘Quantum Equilibrium and the Origin of Absolute Uncertainty’, *Journal of Statistical Physics*, **67**(5-6), pp. 843–907.
- Fröhlich, J. and Schubnel, B. [2016]: ‘The Preparation of States in Quantum Mechanics’, *Journal of Mathematical Physics*, **57**(4), pp. 042101.
- Fuchs, C. A. [1998]: ‘Information Gain vs. State Disturbance in Quantum Theory’, *Fortschritte der Physik: Progress of Physics*, **46**(4-5), pp. 535–565.
- Fuchs, C. A. and Schack, R. [2004]: ‘Unknown Quantum States and Operations, a Bayesian View’, in M. Paris and J. Rehacek (eds), *Quantum State Estimation*, Springer, pp. 147–187.
- Gao, S. [2019]: ‘The Measurement Problem Revisited’, *Synthese*, **196**(1), pp. 299–311.
- Ghirardi, G. C., Rimini, A. and Weber, T. [1986]: ‘Unified Dynamics for Microscopic and Macroscopic systems’, *Physical Review D*, **34**(2), pp. 470.
- Healey, R. [2017]: *The Quantum Revolution in Philosophy*, Oxford University Press.
- Lewis, P. J. [2007]: ‘How Bohm’s Theory Solves the Measurement Problem’, *Philosophy of Science*, **74**(5), pp. 749–760.
- Margenau, H. [1963]: ‘Measurements and Quantum States: Part I’, *Philosophy of Science*, **30**(1), pp. 1–16.
- Maudlin, T. [1995]: ‘Three Measurement Problems’, *Topoi*, **14**(1), pp. 7–15.
- Miszczak, J. A., Puchała, Z., Horodecki, P., Uhlmann, A. and Życzkowski, K. [2009]: ‘Sub- and Super-Fidelity as Bounds for Quantum Fidelity’, *Quantum Information & Computation*, **9**(1-2), pp. 103–130.
- Myrvold, W. [2018]: ‘Philosophical Issues in Quantum Theory’, in E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, Metaphysics Research Lab, Stanford University, fall 2018 edition.
- Nielsen, M. A. and Chuang, I. [2002]: *Quantum Computation and Quantum Information*, Cambridge University Press.

- Norsen, T. [2016]: ‘Bohmian Conditional Wave Functions (and the Status of the Quantum State)’, in *Journal of Physics: Conference Series*, vol. 701, p. 012003. IOP Publishing.
- Paris, M. and Rehacek, J. [2004]: *Quantum State Estimation*, Lecture Notes in Physics. Springer Berlin Heidelberg.
- Park, J. L. and Band, W. [1992]: ‘Preparation and Measurement in Quantum Physics’, *Foundations of Physics*, **22**(5), pp. 657–668.
- Pati, A. K. [2000]: ‘Probabilistic Exact Cloning and Probabilistic No-Signalling’, *Physics Letters A*, **270**(3-4), pp. 103–107.
- Peres, A. [2003]: ‘How the No-Cloning Theorem Got Its Name’, *Fortschritte der Physik: Progress of Physics*, **51**(4-5), pp. 458–461.
- Peres, A. [2002]: *Quantum Theory: Concepts and Methods*, Kluwer Academic Publishers.
- Rivas, A. and Huelga, S. F. [2012]: *Open Quantum Systems*, Springer.
- Rovelli, C. [1996]: ‘Relational Quantum Mechanics’, *International Journal of Theoretical Physics*, **35**(8), pp. 1637–1678.
- Ruetsche, L. [1996]: ‘Van Fraassen on Preparation and Measurement’, *Philosophy of Science*, **63**, pp. S338–S346.
- Ruetsche, L. [2004]: ‘Intrinsically Mixed States: An Appreciation’, *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, **35**(2), pp. 221–239.
- Ruetsche, L. and Earman, J. [2011]: ‘Interpreting Probabilities in Quantum Field Theory and Quantum Statistical Mechanics’, in S. Hartmann and C. Beisbart (eds), *Probabilities in Physics*, Oxford University Press.
- Scarani, V., Iblisdir, S., Gisin, N. and Acin, A. [2005]: ‘Quantum Cloning’, *Reviews of Modern Physics*, **77**(4), pp. 1225.

Vaidman, L. [2015]: ‘Protective Measurement of the Wave Function of a Single System’, in S. Gao (ed.), *Protective Measurement and Quantum Reality*, Cambridge University Press, pp. 15–26.

Wallace, D. [2012]: *The Emergent Multiverse: Quantum Theory According to the Everett Interpretation*, Oxford University Press.

Werner, R. F. [1998]: ‘Optimal Cloning of Pure States’, *Physical Review A*, **58**(3), pp. 1827.

Wessels, L. [1997]: ‘The Preparation Problem in Quantum Mechanics’, in J. Earman and J. Norton (eds), *The Cosmos of Science*, University of Pittsburgh Press, pp. 243–273.

Wootters, W. K. and Zurek, W. H. [1982]: ‘A Single Quantum Cannot Be Cloned’, *Nature*, **299**(5886), pp. 802–803.