A dilemma for relational quantum mechanics

Peter J. Lewis
Dartmouth College, USA
peter.j.lewis@dartmouth.edu

August 13, 2024

Abstract

Relational quantum mechanics (RQM) is an interesting alternative to the standard responses to the measurement problem in quantum mechanics. But it suffers from a distinctive kind of epistemic solipsism: an observer can't in principle know anything beyond their immediate present experience. This makes RQM self-undermining: it takes away the evidence we have for believing in quantum mechanics in the first place. Recently, Adlam and Rovelli have proposed a solution to this problem in the form of a new postulate they call *cross-perspective links*. Here I argue that this postulate does indeed solve the skeptical problem, but it also removes those aspects of RQM that make it distinctively *relational*. Nevertheless, the result of equipping RQM with *cross-perspective links* is an interesting interpretation in its own right.

Keywords: Relational quantum mechanics, solipsism, locality, explanation

1 Some background

Consider a spin-1/2 particle p in a symmetric superposition state

$$\frac{1}{\sqrt{2}}(|{\uparrow_z}\rangle_p + |{\downarrow_z}\rangle_p). \tag{1}$$

Suppose Alice measures the spin of the particle in the z-direction. Then assuming that the unitary Schrödinger dynamics are universally applicable, the final state is

$$\frac{1}{\sqrt{2}}(|\uparrow\rangle_A\,|{\uparrow_z}\rangle_p + |\downarrow\rangle_A\,|{\downarrow_z}\rangle_p), \tag{2}$$

where $|\uparrow\rangle_A$ and $|\downarrow\rangle_A$ represent combined states of Alice and her measuring device in which she gets the result "spin-up" and "spin-down" respectively. But (2) is not a state in which Alice gets the result "spin-up," and it is not a state in which she gets the result "spin-down." So it looks like quantum mechanics is incomplete, as it does not represent the unique outcome of Alice's measurement. This is one version of the quantum measurement problem. One could propose that Alice's measurement "collapses" state (2) to one term or another, but measurements do not constitute a physical natural kind that could precipitate such a collapse. This is another version of the measurement problem.

A number of responses to (both versions of) the measurement problem are available. First, one can add some further representational machinery to quantum mechanics that picks out one or the other of the terms in (2) as the actual result of the measurement. This is the hidden variable approach, epitomized by Bohm's theory. Second, one can reject the assumption that the unitary Schrödinger dynamics are universally applicable, and modify the dynamics so that one or other term in (2) disappears or becomes very small. This is the spontaneous collapse approach, epitomized by the GRW theory. Third, one can deny our tacit assumption that Alice's measurement has a unique outcome, and embrace the idea that there are two outcomes, one relative to each term in (2). This is the many-worlds approach.

None of these approaches is fully satisfactory.[1] The hidden variable approach and the spontaneous collapse approach both have trouble with locality. Suppose we have a pair of spin-1/2 particles in an entangled state

$$\frac{1}{\sqrt{2}}(|{\uparrow_z}\rangle_1\,|{\downarrow_z}\rangle_2 - |{\downarrow_z}\rangle_1\,|{\uparrow_z}\rangle_2). \tag{3}$$

Suppose the two particles are separated: Alice measures the z-spin of particle 1 at location A, and Bob measures the z-spin of particle 2 at space-like separated location B. The unitary Schrödinger dynamics produces a final state

$$\frac{1}{\sqrt{2}}(|\uparrow\rangle_A\,|\downarrow\rangle_B\,|{\uparrow_z}\rangle_1\,|{\downarrow_z}\rangle_2 - |\downarrow\rangle_A\,|\uparrow\rangle_B\,|{\downarrow_z}\rangle_1\,|{\uparrow_z}\rangle_2), \tag{4}$$

---

[1] What follows is a *very* quick gloss. Each approach has been defended against the problems I raise here: see Lewis (2016) for a summary. My point here is just to *motivate* looking beyond the "big three," not to show that it is required.

where $|\uparrow\rangle_B$ and $|\downarrow\rangle_B$ are defined for Bob just as $|\uparrow\rangle_A$ and $|\downarrow\rangle_A$ were defined for Alice. According to the hidden variable approach, the hidden variables determine whether Alice gets "spin-up" and Bob gets "spin-down" or vice versa. In Bohm's theory the dynamics governing the motion of the hidden variables is explicitly causally non-local, in the sense that Alice's measurement on particle 1 instantly affects the hidden variable that determines the spin of particle 2. According to the spontaneous collapse approach, Alice's measurement on particle 1 instantly affects the amplitudes of both terms in the superposition—again, non-locally influencing the spin of particle 2. Furthermore, Bell's theorem (Bell 1964; Bell 1975) shows that no *local* causal mechanism, hidden variable or collapse, can recover all the predictions of quantum mechanics for entangled states.

The many-worlds approach (arguably) avoids this non-locality. But it faces a problem with probability: since both outcomes of each measurement actually occur, and Alice and Bob know that they will both occur, it is hard to see what sense it makes to ascribe each measurement outcome a probability of 1/2. In light of these difficulties, it would be good to have another approach to the measurement problem. That is precisely why relational quantum mechanics is so welcome.

2 Relational quantum mechanics

Relational quantum mechanics (RQM) is a distinctive approach to the measurement problem that is neither a hidden variable approach, nor a spontaneous collapse approach, nor a many-worlds approach (Rovelli 1996; Laudisa and Rovelli 2021). It achieves this via a non-standard treatment of the properties of quantum systems, which in turn leads to a non-standard interpretation of the quantum state. This enables RQM to explain the correlations between entangled systems without non-local causation and without denying that measurements have unique outcomes. Here I will lay out the ingredients that allow RQM to achieve this.

2.1 Relational properties

Typically, properties such as position or spin are taken to be intrinsic properties of quantum systems.[2] RQM, as its name implies, takes such properties to be inherently *relational*. That is, the possession of a property by a system is always relative to some other system. So, for example, when Alice measures the z-spin of particle 1, what she measures is a property *relative to Alice*. Similarly, when Bob measures the z-spin of particle 2, what he measures is a property *relative to Bob*. This is the first step in an account of entangled states that avoids non-local causation: if Alice gets the result spin-up, she learns the property of particle 1 relative to Alice, but not the property relative to Bob. But the devil is in the details, as we shall see.

2.2 Discrete properties

The first detail: what, exactly are the relational properties relative *to*? Alice is an observer making a measurement—but we don't want to say that systems have properties relative to *observers* or to *measurements* for fear of reintroducing the measurement problem. Neither "measurement" nor "observation" is a well-defined physical attribute, so neither has a place in the statement of a physical theory. Instead, RQM makes properties relative to an *interaction* between physical systems.[3] Alice is a physical system, and particle 1 is a physical system, and when they interact, particle 1 acquires a spin property relative to Alice. (Strictly, of course, Alice interacts with a measuring device which in turn interacts with a particle: this will be important later.)

In general, then, whenever two physical systems interact in such a way that some observable P of the first becomes correlated with some observable Q of the second, the first system acquires a value of P relative to the second, and the second acquires a value of Q relative to the first. So when Alice measures the spin of particle 1, particle 1 acquires a spin property, relative to Alice, and Alice acquires a value of the observable corresponding to seeing a particular measurement outcome, relative to the particle. This story makes no ineliminable reference to "measurement"—any interaction between

---

[2] There are nuances, of course: for example, in Bohm's theory, arguably the only intrinsic property that particles have is position.

[3] Of course, "interaction" and "system" are also tricky to define (Adlam and Rovelli 2022, 14), but at least a case can be made that it is *appropriate* that these terms (and not "observation" and "measurement") play a prominent role in our physical theories.

physical systems is analyzed in the same way—and so does not reintroduce the measurement problem.

A consequence of the relativization of properties to interactions is that physical properties become temporally *discrete*. Normally, we think of a property like position as something that a system has over an extended period of time. Even a quantized property like spin is typically thought of as continuous over time: a particle that acquires a z-spin property *retains* that property until it is changed. But according to RQM, systems have properties when, and only when, they interact with other systems. Hence the ontology of RQM should be thought of as a set of spatio-temporally discrete property instantiations.

2.3 States as bookkeeping devices

An ontology of discrete, relational properties is not enough by itself to circumvent the need for non-local causation in explaining the correlations exhibited by entangled systems. To see this, we can appeal to the EPR criterion of reality (Einstein, Podolsky, and Rosen 1935, 777). When Alice measures the z-spin of particle 1 and gets the result spin-up, she can conclude that the proper state to ascribe to the two-particle system is not (3), but $|{\uparrow_z}\rangle_1 |{\downarrow_z}\rangle_2$. In that case, she knows what Bob's outcome (relative to Bob) will be: he will get spin-down. But then it looks like particle 2 must already have a z-spin property (relative or no): Alice's measurement seems to non-locally make it the case that particle 2 is spin-down.

To avoid this consequence, RQM reverses the priority of properties and states. In the major approaches to the measurement problem, the properties of a system are determined, directly or indirectly, by the quantum state (Lewis 2016, 87). In RQM, the properties are fundamental, and the state is simply a convenient bookkeeping device for keeping track of them. And since the properties are inherently relational, the appropriate state for one observer to use may differ from that for a different observer.

So when Alice measures the z-spin of particle 1 and gets the result spin-up, the proper state for her to use henceforth is indeed $|{\uparrow_z}\rangle_1 |{\downarrow_z}\rangle_2$. It tells her what to expect if she were to travel over to particle 2 and measure its z-spin.[4] But it does not entail anything about what state Bob should use.

---

[4] Although cross-temporal comparisons of this kind are in fact problematic in RQM, as will be explored in section 3.

Prior to making his measurement, Bob should still use state (3): particle 2 has no spin property relative to Bob, so the best he can do is to use (3) to generate probabilities for the two possible outcomes. After he makes his measurement, the appropriate state for Bob to use is one of the terms in (3), depending on which outcome he gets. But there is nothing to guarantee that he should use the same state as Alice: it is entirely possible for particle 1 to be spin-up relative to Alice *and* particle 2 to be spin-up relative to Bob, so that Alice should use the first term in (3) and Bob should use the second.

Hence, according to RQM, there is no such thing as *the* quantum state of a system. A quantum state is a tool, to be used by a particular observer to make predictions about their future observations on the basis of their past observations—where all such observations are properties relative to that observer. As a consequence, the appropriate state for Alice to use, based on her observation of particle 1, tells her nothing about the properties of particle 2 relative to Bob.

2.4 Correlations as local

At this point it is reasonable to wonder what has become of the distinctive quantum correlations exhibited by entangled systems. If Alice's outcome tells her nothing about Bob's outcome, in what sense are the spins correlated? The answer, according to Martin-Dussaud, Rovelli and Zalamea (2019), is that any such correlation is a *property* of the pair of measurement results, and hence, according to RQM, only takes a value relative to some *interaction* with the pair itself. Neither Alice nor Bob interact with the *pair* of particles, so no correlation property is instantiated relative to Alice or to Bob. But suppose there is a third observer, Olive, who can observe the outcomes at A and B and see whether or not they are correlated. That is, Olive is located at a point O within the future light-cones of both A and B. When Olive makes her observation, the pair of spin measurements acquires a value for the correlation property—but that value is acquired *relative to Olive.* That is, relative to Olive, Alice and Bob get opposite spin results.

Hence the correlation is a property instantiated relative to an interaction at a specific space-time point, and is in that sense a *local* property. Furthermore, the spin measurements at A and B—and indeed the source of the entangled particles itself—are in the past light-cone of O, so there is no need for any non-local causation to explain this correlation property.

Putting it all together, RQM's combination of discrete, relational properties and observer-relative quantum states offers a distinctive solution to the measurement problem. It is not a many-worlds account: Alice gets a unique result to her spin measurement. It is not a collapse theory: there is no objective quantum state of a system to collapse, and any reduction of a state to one of its terms is a broadly epistemic process, reflecting what information is available to a particular system. It is not a hidden variable theory, at least, not in the sense of an interpretation in which an objective global quantum state is *supplemented* by additional intrinsic properties. But nevertheless, it accounts for the (unique) outcomes of observers, and it does so without the need for non-local causation to account for quantum correlations.

3 Island universes

This all looks very promising. However, the standard RQM account comes with a cost. It is important that Alice's measurement outcome doesn't immediately allow her to infer Bob's measurement outcome, or else we could appeal to the EPR criterion of reality to conclude that particle 2 already has a spin value. The relativity of properties in RQM achieves precisely that. But of course Alice and Bob could later meet up, and Alice could *ask* Bob about his measurement outcome. Note that this puts Alice in the role of Olive: they meet at a point in the future light-cones of A and B, and Alice's question amounts to a measurement of the correlation property. Alice asks Bob "What measurement outcome did you get?" and she hears his reply. But that reply is a physical property of Bob *relative to Alice*: it tells her nothing about the earlier spin property of particle 2 relative to Bob. Indeed, it *can't* tell her about this earlier property, or we could run a slightly more indirect version of the EPR reality criterion: at A, Alice could reason that, since she knows that when she later meets up with Bob she will find that he got the opposite z-spin result, particle 2 must already have the corresponding spin property.

So asking Bob what result he got will not enable Alice to find out the spin of particle 2 relative to Bob. And in fact, it is easy to see that RQM property relativity prevents Alice from finding out Bob's measurement result by *any* means: whatever observation she makes will be an interaction between Alice and some other system, resulting in a property relative to Alice, and telling her nothing about properties relative to Bob. And there is nothing special about Bob in this regard: Alice can't find out about the measurement results obtained by any other observer. Further, recall that RQM regards a measurement by a human observer as a special case of an interaction between physical systems: it doesn't matter whether Bob is a human being or a simple physical recording device. But then, by the same token, Alice can't even know the spin result obtained by her *own* measuring device at A: the outcome relative to Alice, produced by the interaction between Alice and the device, tells her nothing about the outcome relative to the device, produced by the interaction between the device and particle 1. Finally, note that, physically speaking, there is no real distinction between checking the result on a measuring device and checking a *memory.* The outcome relative to the interaction between Alice's cognitive center and her memory center tells her nothing about the outcome relative to the earlier interaction between her memory center and the world.

The result is a kind of epistemic solipsism that Pienaar (2021) aptly describes in terms of "island universes". All Alice has access to are her momentary experiences. Those experiences tell her nothing about her earlier experiences, or about events in the world, including those events having to do with the experiences of other people. The quantum state, regarded as a bookkeeping device, makes sure that her momentary experiences are consistent with the physical laws—for example, that if she remembers getting z-spin up for her spin measurement, she also remembers Bob telling her that he got z-spin down. But these memories tell her nothing about her measurement, Bob's measurement, or the particles themselves.

Di Biagio and Rovelli (2022) point out that RQM is not solipsistic in any *metaphysical* sense: the theory presupposes that there are many physical systems in mutual interaction, including many conscious systems. But a theory that blocks an agent, in principle, from knowing anything beyond their momentary experiences is solipsistic in a broader, epistemic, sense. It is also rather self-undermining: the theory says that a given agent could never know whether the other physical systems presupposed by the theory behave as the theory predicts (Adlam 2022b). However we categorize it, RQM is clearly deeply and problematically skeptical.

Rovelli has, in the past, been resistant to this line of criticism of RQM. Laudisa and Rovelli (2021) argue against the accusation of solipsism as follows: "imagine experimenter S' measures the spin of the electron S, and writes the value of this spin on a piece of paper... if S'' measures the spin and reads the piece of paper, she will find that experimenter S' has seen the same spin as herself." But there is an important ambiguity in the claim that S' and S'' see the same spin. As Laudisa and Rovelli point out, the form of the quantum state ensures that what S'' sees written on the paper, relative to S'', matches the spin result, relative to S''. But it doesn't follow that what S'' sees on the paper, relative to S'', matches the result of the spin measurement, relative to S'—which is what would be required to avoid the "island universe" objection.

4 Cross-perspective links

More recently, in joint work with Adlam, Rovelli has expressed more sympathy with the solipsism objection, noting that within RQM, "each observer is trapped inside their own instantaneous perspective, unable to get information about what the world is like for other observers or at other times" (Adlam and Rovelli 2022, 4). Adlam and Rovelli (2022, 5) propose a solution, namely to add a new postulate to RQM they call *cross-perspective links*:

> In a scenario where some observer Alice measures a variable V of a system S, then provided that Alice does not undergo any interactions which destroy the information about V stored in Alice's physical variables, if Bob subsequently measures the physical variable representing Alice's information about the variable V, then Bob's measurement result will match Alice's measurement result.

So, in particular, if Bob asks Alice about her measurement result, since asking this question is tantamount to measuring the physical variables representing Alice's information about result, the answer he receives genuinely tells him about Alice's result, and hence gets him outside his solipsistic prison.

This postulate has to be understood quite broadly if it is to completely solve the problem. Recall that the problem with the original formulation of RQM is not just that Bob can't find out about Alice's measurement result, but that he can't find out about his own measurement result as recorded on his measuring device, or about his earlier measurement result as recorded in his memory. But since "observers" in RQM are just physical systems, and measurements are just a particular kind of interaction between physical systems, we can regard "Alice" and "Bob" in the statement of *cross-perspective links* as names of any suitable physical system. On that reading, when Bob looks at his measuring device, then assuming that the device hasn't undergone any unfortunate interactions, the result he sees matches the result the device recorded. And similarly, when Bob consults his memory, the result of his act of remembering matches the original observation that led to the formation of the memory.

One might wonder how such a postulate might be *motivated* except as an *ad hoc* fix for RQM. But for the moment, let us accept RQM as amended by *cross-perspective links*: we can call it RQM+. Adlam and Rovelli (2022, 1) admit that the move to RQM+ "requires us to clarify some features of the ontology of RQM," but contend that it does not depart from the spirit of RQM entirely. In particular, they argue that RQM+ requires "an ontology which upholds the principle that quantum states are always relational, but which also postulates a set of quantum events which are not relational."

Let us see why they say that. Suppose that Alice measures the z-spin of particle p, prepared in state (1) and gets the result spin-up. According to *cross-perspective links*, the information about p stored in Alice's physical variables is now available to any observer—indeed, any physical system—that interacts with Alice in a suitable way. Hence, while RQM asserts that quantum events of this kind are inherently relational, RQM+ "suggests that the set of 'quantum events' should be regarded as absolute, observer-independent features of reality" (2022, 8). That is, Adlam and Rovelli "continue to endorse the sparse-flash ontology for RQM... however we now regard the pointlike quantum events or 'flashes' as absolute, observer-independent facts about reality, rather than relativizing them to an observer" (2022, 8).

Nevertheless, even though the *events* are non-relative in RQM+, Adlam and Rovelli insist that *value* spin-up is relativized to Alice, since it is only Alice who has obtained this information (2022, 8). They make a similar point about quantum *states*. After the measurement, Alice should ascribe the eigenstate $|{\uparrow_z}\rangle_p$ to p. But Bob should ascribe a different state to the system: he is in no position to predict with certainty what value he would obtain were he to measure the z-spin of p. The spin of p is now correlated with the state of Alice, and since Alice is macroscopic, she interacts strongly with her environment, suppressing interference terms via decoherence. Hence the appropriate state for Bob to ascribe to p is an equally-weighted classical mixture of spin-up and spin-down states, reflecting his ignorance of the result he would obtain on measurement (2022, 7).

As a result, then, Adlam and Rovelli argue that adding *cross-perspective links* to RQM requires only a small change to the interpretation. There is now an objective, observer-independent set of quantum events at which properties are instantiated. However, those properties are still relativized to an observer, and are still temporally discrete. And quantum states are still to be regarded as observer-relative bookkeeping devices rather than objective descriptions of the world. However, I think Adlam and Rovelli significantly undersell the differences between RQM and RQM+, as spelled out in the following section.[5]

5 Comparing the accounts

5.1 Properties

Consider again Alice's z-spin measurement on state (1). If she gets spin-up, and Bob asks her what result she got, then *cross-perspective links* entails that she will reply "spin-up". Similarly, if Bob measures the z-spin of p himself, he will get spin-up. And the same is true of any system that interacts with p in the relevant way. So then in what sense is the value observer-relative?

Adlam and Rovelli claim that "the value v is relativized to Alice because at this stage Alice is the only observer who has this information about S, although other observers could later come to have the same information by interacting appropriately with either Alice or S" (2022, 8). But *information* needs to be distinguished from what that information is *about*. This is so even if (as Adlam and Rovelli insist) "information is physical" (2022, 2). Only Alice has the information about p, in the sense that the information is stored in Alice's physical variables. But it does not follow that the *property* that this information represents is in any sense private to Alice. Indeed, part of the point of *cross-perspective links* is to make certain properties *public*: the spin property of p is publicly accessible by any suitable measuring device. This is just what it means for p to have the corresponding property (non-relationally).

One could cash this out dispositionally: what it means for a system to have a given property is for it to be disposed to result in the corresponding value on measurement. But I don't think a dispositional account of properties is required here. All that is required is a minimal form of realism, such that properties are not *identified* with observations of them, but rather are *revealed* by observations. Adlam and Rovelli regard RQM and RQM+ as realist interpretations of quantum mechanics (2022, 1). And all suitable interactions with p are bound to reveal the value spin-up. So it seems hard

---

[5] A similar sentiment is expressed by Stacey (2022).

for them to resist the claim that p has the corresponding spin property—not just relative to Alice, but absolutely and non-relatively.

Of course, both the property of p and Alice's information about it could be disturbed by physical interactions. As far as Alice's information goes, this is explicitly reflected in the proviso in *cross-perspective links*: the guarantee that a measurement of Alice's information will match Alice's result holds only so long as Alice's relevant physical variables are not significantly disturbed. And we can say the same about the spin property itself: the guarantee that a z-spin measurement of p will yield spin-up holds only so long as p is not significantly disturbed, for example by an x-spin measurement. But of course, this is exactly what you would expect for *any* physical property: it persists only so long as it isn't destroyed.

Similar comments apply to Adlam and Rovelli's insistence that properties in RQM+ inherit the temporal discreteness of those in RQM. After Alice has obtained spin-up for her measurement, a z-spin measurement *at any subsequent time* will yield spin-up. If, as just argued, system S has the corresponding property, then it has that property through time, not at a single time. One might insist that p has the property only at the multiple discrete times at which it is actually measured, but this too seems to flirt with anti-realism. While the *observations* may be discrete in time, the property so revealed is not discrete in time, since the same spin value would have resulted at any time after the initial measurement.[6]

Perhaps, though, I am being too permissive in my criteria for property possession. Adlam and Rovelli (2022, 9) concede that "although strictly speaking the variables of a system do not have values in between interactions, nonetheless one may sometimes wish to speak colloquially of variables 'having values' at other times." Even if, colloquially speaking, we might count p as spin-up when a measurement would yield that value, we should use a stricter standard when positing fundamental ontology. Certainly it is coherent to think that p has a spin property exactly when its spin is measured, where the quantum state ascribed to p by Alice is simply a guide to future point-like property instantiations, should there be future spin measurements. But

---

[6] One might worry about the temporal directedness of counterfactuals like this in the context of the atemporal "all at once" mode of explanation endorsed by Adlam and Rovelli (2022, 13). But even if atemporal explanation turns out to show that there are cases in which forward-looking counterfactual analyses are inappropriate, it looks like this one is innocuous. This is not a context in which inserting an additional measurement would change the whole distribution of events.

I think it is reasonable to demand a motivation for withholding property ascriptions except during such interactions. RQM has such a motivation: a subsequent measurement might not reveal the same spin value. But this motivation is absent in RQM+.[7]

A further argument for the temporal discreteness of properties in RQM+ offered by Adlam and Rovelli (2022, 8) is that it is required to satisfy the Kochen-Specker theorem (Kochen and Specker 1967). The theorem shows that no set of possessed properties at a time can reproduce the quantum statistics for the complete set of measurements that could be performed on a (suitably complex) system. The temporally discrete ontology of RQM certainly doesn't try to ascribe properties to a system at a single time corresponding to every measurement that might be performed on it: at most one property is ascribed to a system at a time, and even then only for an instant. But such a sparse ontology is not *required.* As just argued, in RQM+ a system acquires a property when a measurement is performed on it, and retains that property until a disturbing measurement is performed on it. Since measurements of non-commuting observables are mutually disturbing, a system will never have a set of properties for all potential measurements, such as would violate the Kochen-Specker theorem.

All this is strikingly different from the original version of RQM. In the original version, properties really are relative: just because Alice got spin-up for her z-spin measurement on p, it doesn't follow that Bob will also get spin-up for his. So p has one property relative to Alice, and another relative to Bob. It doesn't follow that outcomes relative to later interactions with p will agree with outcomes relative to earlier interactions, so property possession is discrete in time. These features of RQM really do lead to properties being relative and temporally discrete—but they are removed by the assumption of *cross-perspective links.* The observer-relativity and temporal discreteness of properties is a natural consequence of RQM, but not of RQM+.

---

[7] Bell (1987) first introduces the ontology of point-like flashes to find a place for three-dimensional ontology in a wavefunction-only interpretation of the GRW theory. But this motivation is absent in RQM+, where the wavefunction is not treated ontologically. Tumulka (2006) adopts this ontology to construct a relativistic account of the GRW collapse events But again, this motivation is not present in RQM+: there are no *collapse events* in RQM+, and a relativistic theory doesn't preclude properties that persist over time.

5.2 States

Adlam and Rovelli might seem to be on firmer ground regarding their claim that *states* remain relative to observers in RQM+. After all, states are intimately connected to *prediction*. Consider again the situation in which Alice has measured the z-spin of p, prepared in state (1), and obtained the result spin-up. Suppose Bob then performs his own z-spin measurement on p. Bob's lack of information concerning Alice's result means that he can't predict his own outcome with certainty: instead, his predictive ability is accurately reflected by a suitable *mixture* of z-spin eigenstates, as noted by Adlam and Rovelli (2022, 7). Alice, on the other hand, can predict that if she were to repeat the measurement, she would get spin-up again: for Alice, the appropriate state is the spin-up eigenstate $|{\uparrow_z}\rangle_p$.

However, there is some subtlety regarding proper state ascription here. RQM (and RQM+) posit that "unitary quantum mechanics is complete," in the sense that there are no hidden variables (Adlam and Rovelli 2022, 2). That is, the quantum state is taken to be the *best* available predictor of the properties of a system: there are no additional descriptions of the system—hidden variables—that would enable better predictions. But now consider Bob's predictive abilities. Although the mixture accurately represents Bob's predictive abilities (given his information), there is a *better* predictor available, namely the eigenstate $|{\uparrow_z}\rangle_p$. After all, Bob *will* get spin-up when he makes his z-spin measurement: *cross-perspective links* guarantees it. Again, a moderate realism seems to require the use of the z-spin eigenstate here: since (as just argued) p has the z-spin-up property (absolutely and non-relatively), the appropriate state to use—by anyone—for making predictions about future measurement of p is the corresponding eigenstate. Of course, the same caveats apply concerning future disturbances: if p is disturbed, then the properties, and the state, might change.

So far, then, I have argued that after Alice gets spin-up for her z-spin measurement, the appropriate state for *anyone* to use for making predictions about p is $|{\uparrow_z}\rangle_p$. But a case can also be made that the appropriate state to use is the entangled superposition state (2). Note that RQM (and RQM+) assign no special status to an observer: an observer is simply a physical system that becomes entangled with the target system. So Alice needn't be a *person*: "Alice" could be the name of a single particle whose position becomes correlated with the spin of p. Suppose that after "Alice" performs her "measurement," Bob performs a measurement on Alice and p that has state (2) as an eigenstate, with eigenvalue +1. Bob can predict with certainty that the measurement will result in outcome +1. But to make this prediction, he needs to use the *whole* of (2), not just the first term.

Even if Alice *is* a person, there is still an argument that Bob should use (2) to describe Alice and p. Adlam and Rovelli (2022, 18) discuss extended Wigner's friend scenarios, which raise the possibility of Bob performing the measurement just described *even if* Alice is a macroscopic object like a person. Of course, such measurements are not practically possible, but a universal interpretation of quantum mechanics should be able to make sense of them nonetheless. In that case, even if Alice is a person, Bob should use state (2) to make predictions: he can predict with certainty that his outcome will be +1.

One might take my argument so far to support Adlam and Rovelli's contention that states remain observer-relative in RQM+: Alice should use $|{\uparrow_z}\rangle_p$ to make predictions, Bob should use state (2). But this is not my point: the state to use depends on the *measurement*, not the observer. For future measurements of the z-spin of p, both Alice and Bob should use $|{\uparrow_z}\rangle_p$. But for future Wigner's-friend measurements, both Alice and Bob should use (2). This is obvious for Bob, but even for Alice, if she contemplates such a measurement being performed on her, she should predict the outcome using (2)—she should predict that the outcome will definitely be +1. Of course, she should also predict that such a measurement will destroy her information about the z-spin of p. In fact, Adlam and Rovelli (2022,20) note precisely this consequence of RQM+: since it is an observer-independent fact that Bob will get outcome +1, Alice should use (2) to predict the result of Bob's measurement. So perhaps the observer-relativity claim is that *anyone* should use the spin-up eigenstate to make predictions about future measurements *by Alice*, but (2) to make predictions about future measurement *by Bob*. But this, too, doesn't seem to put the distinction in the right place: Bob might perform a z-spin measurement, in which case Alice should make predictions about *Bob's measurement* using the spin-up eigenstate.

Elsewhere, Adlam and Rovelli describe the relativity of state-ascriptions in terms of the *history* of the interacting systems: the state of S relative to Bob "characterizes the *joint* history of Bob and S" (2022, 11). Alice has interacted with the spin of p and Bob has not, and this is reflected in the fact that the state of p relative to Alice is $|{\uparrow_z}\rangle_p$, but the state of Alice plus p relative to Bob is (2). But note that the joint history of Alice and p is relevant to *both* state-ascriptions: the entangled state (2) also reflects a measurement interaction between Alice and p. Hence the history of interaction between Alice and p doesn't distinguish between $|{\uparrow_z}\rangle_p$ and (2) as the state relative to Alice. Similarly, the lack of interaction between Bob and p doesn't distinguish whether Bob should ascribe (2) to Alice plus p, or just one of its terms.

In sum, then, there is nothing particularly observer-relative about states in RQM+. Perhaps Bob should use (2) to make his prediction, or perhaps he should use one of its terms: it depends on the measurement he is contemplating. Perhaps Alice should self-ascribe (2), or perhaps one of its terms: again, it depends on the measurement in question. And the move from RQM to RQM+ also challenges the posit that unitary quantum mechanics is *complete*. If this posit holds, then there should be a quantum state that encodes all the available information about a system. But neither $|{\uparrow_z}\rangle_p$ nor (2) is adequate to this task. $|{\uparrow_z}\rangle_p$ is incomplete when it comes to predicting the results of Wigner's-friend measurements: we need the full superposition. And (2) is incomplete when it comes to predicting the result of future z-spin measurements, given that Alice got "spin-up": we need to single out the first term in the superposition somehow. A *complete* representation of the situation after Alice has performed her measurement requires using state (2), plus a "pointer" picking out the first term in (2). That is, it seems to require a *hidden variable*.[8] But hidden variables are one of the things that Rovelli's program explicitly set out to avoid (Adlam and Rovelli 2022, 2).

5.3 Non-locality

Recall that one of the main motivations for RQM was to develop an interpretation of quantum mechanics that avoids non-local causation when applied to entangled systems. The way RQM achieves this is via its strict relativization of properties and states to observers: when Alice measures the z-spin of particle 1 in state (3), she cannot on that basis predict the outcome of Bob's measurement on particle 2. In RQM, there is no correlation between Alice's result, relative to Alice, and Bob's result, relative to Bob: a correlation only exists relative to some future observer who can see both measurements. But this kind of resolution is blocked in RQM+: if Alice gets spin-up for her measurement on particle 1, she can predict with certainty that Bob will get

---

[8] Perhaps the response might be that the hidden variable *just is* the ascription of $|{\uparrow_z}\rangle_p$ to p relative to Alice. I have no objection to this way of putting it, as long as it is clear that the hidden variable is not the *state*, relative to Alice or to anyone: it is not by itself a complete bookkeeping device.

spin-down for his measurement on particle 2. That is, she can appeal to the EPR criterion of reality to conclude that particle 2 already has a spin property. *Cross-perspective links* ensures that there *is* an objective, non-relative correlation between Alice's outcome and Bob's space-like separated outcome.

Adlam and Rovelli recognize this consequence of the move from RQM to RQM+: "the version of RQM that we have set out here is nonlocal in a straightforward way" (2022, 13). How can this non-locality be reconciled with relativity? The strategy they suggest starts from the observation that "we do not need to think of the set of events as being generated in some particular temporal order" (2022, 13). That is, they propose that RQM+ is best understood via "a metaphysical picture in which the laws of nature apply atemporally to the whole of history, fixing the entire distribution of quantum events all at once" (2022, 13). This is contrasted with the usual, temporal mode of explanation, in which the laws of nature fix *later* events on the basis of *earlier* events.

Recall that Bell's theorem entails that any past-to-future explanation, deterministic or probabilistic, of the correlations exhibited by entangled particles would have to involve non-local action. What Adlam and Rovelli suggest is that we give up the attempt to explain the correlated events in terms of their local *pasts*: "the outcome of a given quantum event is not determined by the information locally available at that event, but rather by the probability distribution assigned from the outside, which necessarily contains information about all the other quantum events happening elsewhere and at other times" (2022, 13). Nothing in Bell's theorem precludes sprinkling space-time with quantum events, so that the properties instantiated at those events satisfy the probabilistic predictions of quantum mechanics. Adlam and Rovelli stress that "this kind of nonlocality does not involve superluminal signalling" because RQM+ "is not required to tell any story about the spatiotemporal unfolding of beables in between quantum events, or to say anything about the path along which an influence travels from one quantum event to another" (2022, 14).

This is an interesting solution to the problem of locality in quantum mechanics—and a distinct one from the one offered by the original version of RQM.[9] But note that it leans heavily on the idea that RQM+ retains the

---

[9] For examples of atemporal accounts of quantum correlations, see Adlam (2022a), Price (2012), Price and Wharton (2015), and Silberstein, Stuckey, and McDevitt (2018). For an overview of the approach, see Friederich and Evans (2019) and Wharton and Argaman (2020).

temporally discrete ontology of RQM: there is no “spatiotemporal unfolding of beables in between quantum events” because there are only the discrete events themselves. However, in section 5.1 I argued that RQM+ does away with the discrete ontology of RQM: systems generally *retain* their properties until those properties are disturbed. Hence we cannot take the ontology to consist entirely of spatiotemporally discrete property instantiations: RQM+ does, after all, admit something like a “spatiotemporal unfolding of events”.

However, I think Adlam and Rovelli’s atemporal approach can still mitigate the problem of non-locality. First, even though quantum systems in RQM+ retain their properties over time, RQM+ plausibly still lacks the chains of spatio-temporally continuous properties required for past-to-future explanation. Before Alice and Bob make their measurements, there is no reason to think that their respective particles have spin properties, or position properties, or any other properties that might constitute a spatio-temporal *trajectory.* The right way to understand explanation in RQM+ might still be in terms of an atemporal distribution of properties over spacetime, even if those properties are somewhat extended in time.

Second, even if an ontology of continuous space-time trajectories *is* compatible with RQM+, atemporal explanation can accommodate this as well. What blocks a local, past-to-future explanation of entanglement correlations is that the properties ascribed to the particles prior to measurement must be able to account for the results of *any* measurements that might be made on the particles in the future. But an atemporal account allows the properties of the particles to depend on the *actual* measurements that *will* be performed on them (Price 1994). This is the key insight that allows atemporal explanations to remain local, the insight that Adlam and Rovelli rely on, and it is still available if one moves away from a temporally discrete ontology.

6 The dilemma

Where does all this leave us? I think it presents relational quantum mechanics with a dilemma. On the one hand, we could stick with the original version of RQM, but this results in an extreme, solipsistic skepticism. On the other hand, we could modify RQM by the addition of *cross-perspective links*, but this undoes most of the distinctive features of RQM, turning it into a hidden variable theory—albeit of an interesting, non-standard kind.

I seems clear to me which horn to take. Something like *cross-perspective* *links* looks like a required addition to RQM: without it, the resulting skepticism is epistemically devastating, taking away the reasons one might have for believing in RQM in the first place (Adlam 2022b). Adlam and Rovelli argue that this addition requires only minor changes to the ontology of relational quantum mechanics, making quantum events non-relational, but leaving intact relational, temporally discrete properties and observer-relative states. However, I think this underestimates the changes wrought by *cross-perspective links*: I have argued that it makes quantum properties non-relational and temporally continuous, and reintroduces objective quantum states. Furthermore, if relational quantum mechanics is to retain the requirement of descriptive completeness, quantum states need to be supplemented by hidden variables representing the outcomes of measurements.

A further cost of the move from RQM to RQM+, recognized by Adlam and Rovelli, is that RQM+ needs a new account of entanglement correlations: property relativization can no longer do the trick. They suggest an account of explanation in terms of laws conceived as atemporal, global constraints rather than in terms of dynamical, past-to-future constraints. I think that this is very promising, and can apply even if we take on board the sweeping ontological changes required by the move to RQM+. The result—an atemporal hidden variable theory—is worth further investigation. It is not really a *relational* form of quantum mechanics but it nevertheless provides an interesting alternative to the existing interpretive options.

Acknowledgments

I would like to thank Emily Adlam for helpful discussions about these issues, and the audience at the 13th Principia International Symposium in Florianópolis, Brazil for useful feedback on an earlier version of this paper.

References

Adlam, Emily (2022a), “Two roads to retrocausality,” *Synthese* 200: 422.

Adlam, Emily (2022b), “Does science need intersubjectivity? The problem of confirmation in orthodox interpretations of quantum mechanics,” *Synthese* 200: 522.

Adlam, Emily, and Carlo Rovelli (2022), “Information is physical: Cross-perspective links in relational quantum mechanics,” *Philosophy of Physics* 1: 4.

Bell, John S. (1964), “On the Einstein-Podolsky-Rosen paradox,” *Physics* 1: 195–200. Reprinted in Bell (2004), 14–21.

Bell, John S. (1975), “The theory of local beables,” CERN preprint, https://cds.cern.ch/record/980036/files/197508125.pdf. Reprinted in Bell (2004), 52–61.

Bell, John S. (1987), “Are there quantum jumps?” in C. W. Kilmister (ed.), *Schrödinger: Centenary Celebration of a Polymath*, Cambridge University Press, 41–52. Reprinted in Bell (2004), 201–212.

Bell, John S. (2004), *Speakable and Unspeakable in Quantum Mechanics*, Cambridge University Press.

Di Biagio, Andrea, and Carlo Rovelli (2022), “Relational quantum mechanics is about facts, not states: A reply to Pienaar and Brukner,” *Foundations of Physics* 52: 62.

Einstein, Albert, Boris Podolsky, and Nathan Rosen (1935), “Can quantum-mechanical description of physical reality be considered complete?,” *Physical Review* 47: 777–780.

Friederich, Simon and Peter W. Evans (2019), “Retrocausality in quantum mechanics,” in Edward N. Zalta (ed.), *Stanford Encyclopedia of Philosophy* (Summer 2019 Edition), https://plato.stanford.edu/archives/sum2019/entries/qm-retrocausality/.

Kochen, Simon, and E. P. Specker (1967), “The problem of hidden variables in quantum mechanics,” *Journal of Mathematics and Mechanics* 17: 59–87.

Laudisa, Federico and Carlo Rovelli (2021), “Relational quantum mechanics,” in Edward N. Zalta (ed.), *Stanford Encyclopedia of Philosophy* (Winter 2021 Edition), https://plato.stanford.edu/archives/win2021/entries/qm-relational/.

Lewis, Peter J. (2016), *Quantum Ontology: A Guide to the Metaphysics of Quantum Mechanics*, Oxford University Press.

Martin-Dussaud, Pierre, Carlo Rovelli, and Federico Zalamea (2019), "The notion of locality in relational quantum mechanics," *Foundations of Physics* 49: 96–106.

Pienaar, Jacques (2021), "A quintet of quandaries: Five no-go theorems for relational quantum mechanics," *Foundations of Physics* 51: 97.

Price, Huw (1994), "A neglected route to realism about quantum mechanics," *Mind* 103: 303–336.

Price, Huw (2012), "Does time-symmetry imply retrocausality? How the quantum world says "Maybe"?" *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 43: 75–83.

Price, Huw, and Ken Wharton (2015), "Disentangling the quantum world," *Entropy* 17: 7752–7767.

Rovelli, Carlo (1996), "Relational quantum mechanics," *International Journal of Theoretical Physics* 35: 1637–1678.

Silberstein, Michael, William Mark Stuckey, and Timothy McDevitt (2018), *Beyond the Dynamical Universe: Unifying Block Universe Physics and Time as Experienced*, Oxford University Press.

Stacey, Blake C. (2022), "The de-relationalizing of relational quantum mechanics," arXiv preprint, https://arxiv.org/pdf/2211.03230

Tumulka, Roderich (2006), "A relativistic version of the Ghirardi–Rimini–Weber model," *Journal of Statistical Physics* 125: 821–840.

Wharton, K. B., and N. Argaman (2020), "Colloquium: Bell's theorem and locally mediated reformulations of quantum mechanics," *Reviews of Modern Physics* 92: 021002.