



Information is Physical: Cross- Perspective Links in Relational Quantum Mechanics

EMILY ADLAM 

CARLO ROVELLI 

*Author affiliations can be found in the back matter of this article

Philosophy of Physics



RESEARCH



ABSTRACT

Relational quantum mechanics (RQM) is an interpretation of quantum mechanics based on the idea that quantum states do not describe an absolute property of a system but rather a relationship between systems. There have recently been some criticisms of RQM pertaining to issues around intersubjectivity. In this article, we show how RQM can address these criticisms by adding a new postulate which requires that all of the information possessed by a certain observer is stored in physical variables of that observer and thus is accessible by measurement to other observers. This makes it possible for observers to reach intersubjective agreement about quantum events that have occurred in the past. We suggest a possible ontology for a version of RQM employing this postulate; this ontology upholds the principle that quantum states are always relational, but it also postulates a set of quantum events that are not strictly relational. We show that the new postulate helps address the preferred basis problem in RQM.

**CORRESPONDING
AUTHOR:**

Emily Adlam

University of Western
Ontario, CA

eadlam2@uwo.ca

TO CITE THIS ARTICLE:

Adlam, Emily, and Carlo
Rovelli. 2023. "Information
is Physical: Cross-Perspective
Links in Relational Quantum
Mechanics." *Philosophy of
Physics* 1(1): 4, 1–19. DOI:
[https://doi.org/10.31389/
pop.8](https://doi.org/10.31389/pop.8)

1 INTRODUCTION

Relational quantum mechanics (RQM) is an interpretation of quantum mechanics based on the idea that quantum states do not describe an absolute property of a system but rather a relationship between systems. RQM has many very appealing features. It is a realist view that is compatible with relativity; it does not require us to add anything to the existing mathematical framework of quantum mechanics; it is a robustly naturalistic picture that does not attach any special significance to conscious minds or measurements; and it refrains from postulating unobservable, inaccessible levels of reality like hidden variables or other branches of an Everettian multiverse. Moreover, it seems likely that RQM will still be applicable in the context of relativistic quantum mechanics, quantum field theory, and quantum gravity, whereas many other proposed interpretations of quantum mechanics face significant difficulties when we try to extend them beyond nonrelativistic quantum mechanics.

However, some problems remain; in particular, there is a tension between RQM's naturalistic emphasis on the physicality of information and the inaccessibility of certain sorts of information in current formulations of RQM. Thus, in this article, we propose a new postulate for RQM which ensures that all of the information possessed by a certain observer is stored in physical variables of that observer and thus is accessible by measurement to other observers. The postulate of **cross-perspective links** makes it possible for observers to reach intersubjective agreement about quantum events that have occurred in the past, thus shoring up the status of RQM as a form of scientific realism and allowing that empirical confirmation is possible in RQM.

Adding this postulate requires us to update some features of the ontology of RQM, because it entails that not everything in RQM is relational. In this article, we suggest an ontology which upholds the principle that quantum states are always relational, but which also postulates a set of quantum events that are not relational. A quantum event arises in an interaction between two systems such that the values of some physical variables of one system become definite relative to another system, and these quantum events are observer-independent in the sense that any other observer can, in principle, obtain the same information about the values of the relevant variables by an appropriate measurement on either of the systems.

2 RQM

According to Rovelli's first paper on the topic (Rovelli 1996), the founding principle of RQM is the idea that "in quantum mechanics different observers may give different accounts of the same sequence of events." RQM has undergone significant development since this original proposal, but the basic idea remains the same: different observers may assign different quantum states to a given system, and moreover, in such cases, all of the different assignments are equally correct because the quantum state assigned to a system describes not only the system itself but also the relation between the system and the observer assigning the state. There exist other interpretations of quantum mechanics that take a similar view on the relational nature of quantum states (Brukner 2015; Bub 2012; Demopoulos 2012; Janas et al. 2021; Fuchs, Mermin, and Schack 2014), but these accounts typically regard (conscious) observers as playing some sort of privileged role. By contrast, RQM is built on strong naturalistic intuitions; therefore, in RQM, the term "observer" is understood in a broad sense, which allows that any physical system can be an "observer," so we do not have to accept that consciousness plays any fundamental role.

It is not straightforward to identify a single canonical formulation of RQM, as there is still vigorous ongoing discussion about it. However, in this article, we will take existing

formulations of RQM to be characterized by the following six postulates, which are endorsed in a recent presentation of RQM in (Di Biagio and Rovelli 2021a). Thus, it is this specific version of RQM to which our proposed modifications apply.

1. **Relative facts:** Events, or facts, can happen relative to any physical system.
2. **No hidden variables:** Unitary quantum mechanics is complete.
3. **Relations are intrinsic:** The relation between any two systems A and B is independent of anything that happens outside these systems' perspectives.
4. **Relativity of comparisons:** It is meaningless to compare the accounts relative to any two systems except by invoking a third system relative to which the comparison is made.
5. **Measurement:** An interaction between two systems results in a correlation within the interactions between these two systems and a third one; that is, with respect to a third system W , the interaction between the two systems S and F is described by a unitary evolution that potentially entangles the quantum states of S and F .
6. **Internally consistent descriptions:** In a scenario where F measures S , and W also measures S in the same basis, and W then interacts with F to “check the reading” of a pointer variable (i.e., by measuring F in the appropriate “pointer basis”), the two values found are in agreement.

Postulate four is particularly important to recent presentations of RQM. It is sometimes expressed by the slogan “there is no view from nowhere,” which is associated with a radically perspectival approach insisting there are no fact, observer-independent facts at all: every meaningful physical description must be relativized to something. This idea is the origin of some of the radical metaphysics that have been associated with RQM. The basic idea that some physical quantities previously thought to be absolute could in fact be relational is not conceptually novel (e.g., Rovelli (1996) emphasizes the analogy with the relativization of velocity in special relativity), but RQM as described in the postulates above takes this idea a step further—it is not just that quantum states are relational, but *all* facts are relational, even facts about relations must always be relativized to an observer.

Various different ontologies have been suggested to flesh out the postulates above; in this article, we will largely focus on the version espoused in Rovelli (2018) and Laudisa and Rovelli (2021), which suggests we should think of RQM in terms of a set of “quantum events” that occur relative to physical systems. The idea is that when two systems A , B interact, a “quantum event” typically occurs in which the value of some variable V of system A becomes definite relative to system B : if B is a conscious observer, this interaction may sometimes be understood as B measuring A and obtaining some definite value for the variable V , but values can also become definite relative to systems that are not conscious. Quantum states are therefore relational, because the observer relative to whom V has a definite value and observers relative to whom V does not have a different value will assign different states to A . But note that, in accordance with **no hidden variables** and **relative facts**, the value of V is not a hidden variable: it is simply true for some observers and not others. Thus, it is not the case that these relational states are purely epistemic descriptions based on observers having different information; each observer's quantum state is a correct and complete description of the facts relative to them.

This is perhaps best illustrated by an example (which we base on the famous “Wigner's Friend” case). Suppose Bob knows that his friend Alice is performing a measurement on a system S to determine the value of some variable V of S . When Alice performs this measurement, she witnesses some measurement outcome and thus learns the value of V , so

from her point of view, she and V are in a definite state corresponding to one measurement result. But if we assume that Bob should describe the whole interaction using unitary quantum mechanics, from his point of view, the interaction has only caused Alice and S to evolve into an entangled state $\sum_i c_i |A_i\rangle |S_i\rangle$ where $|S_i\rangle$ is the state of S corresponding to measurement outcome i and $|A_i\rangle$ is the state of S corresponding to Alice observing measurement outcome i ; so from Bob's point of view, the measurement has not selected any one measurement result. So we apparently have two contradictory descriptions of the situation. Which one is correct?

RQM answers these questions as follows. **Measurement** and **no hidden variables** tell us that Bob should indeed describe the measurement as an interaction that leads to an entangled state $\sum_i c_i |A_i\rangle |S_i\rangle$. Meanwhile, Alice and the system interact and undergo a quantum event in which the value of V becomes definite relative to her. However, according to **relative facts**, that event only occurs relative to her and thus is irrelevant to Bob's description of the scenario. Therefore there is no inconsistency in the two contradictory descriptions—each agent has a description of the scenario that is correct relative to them, and from **relativity of comparisons**, we cannot even meaningfully compare these descriptions. Moreover, due to **internally consistent descriptions**, we can be assured that if Alice and Bob subsequently discuss their experiences, they will always perceive the other as agreeing with them; so nobody will ever be aware of any inconsistencies.

We should also mention some other ideas that have played an important role in motivating RQM. First, one of the founding tenets of RQM is the idea that information is a part of the physical world and should be modeled appropriately within the theory. For example, Di Biagio and Rovelli (2021a) write,

In a naturalistic philosophy, what F “knows” regards physical variables in F . And this is accessible to W . If knowledge is physical, it is accessible by other systems via physical interactions. It is precisely for this reason that knowledge is also subjected to the constraints and the physical accidents due to quantum theory.

This idea is part of the motivation for the idea that any physical system can play the role of an “observer,” and for the postulate **measurement**, a measurement is not just an epistemic concept but a physical interaction that can be described as such from the point of view of an external observer.

In this spirit, we note that Rovelli's original presentation of RQM made use of several information-theoretic postulates (Rovelli 1996), such as “There is a maximum amount of relevant information that can be extracted from a system” and “It is always possible to acquire new information about a system.” We will refer to these postulates as “the 1996 postulates.” Rovelli showed that a significant amount of the structure of standard quantum mechanics can be extracted from the 1996 postulates, which he viewed as showing that “the notion of absolute observer-independent state of a system is replaced by the notion of information about a system that a physical system may possess.” This information-theoretic motivation is not always emphasized in later presentations of RQM, but it is an important part of the historical development of the theory.

3 INTERSUBJECTIVITY

A number of authors (Adlam 2022; Pienaar 2021; Brown 2007; van Fraassen 2010) have recently made criticisms of RQM related to intersubjectivity. To exhibit the problem,

Adlam (2022) considers the case where Bob knows that Alice is performing a measurement on a system S . When Alice performs the measurement, she witnesses some measurement outcome M_A and thus learns the value of some variable V of the system S . Now suppose Bob measures S in the same basis as Alice's measurement, and hence he obtains a measurement outcome M_B^S , which he will interpret as providing information about the result of Alice's measurement on S . Suppose that Bob also "measures" Alice herself and obtains a measurement outcome M_B^A for the value of some pointer variable that is supposed to be a record of her measurement result—for example, he could simply ask her what her measurement result was. So in this scenario, we have three different measurement outcomes M_A, M_B^S, M_B^A all supposedly providing information about the value of the same variable. What does RQM say about the relationships between these three measurement results?

Well, clearly **internally consistent descriptions** entails that M_B^S and M_B^A will agree. But this leaves a further question about whether M_B^S and M_B^A will match M_A . Unitary quantum mechanics does not provide any mechanism for a single measurement outcome to be selected and actualized for Alice in the first place, so it certainly cannot tell us anything about the relationship between her outcome and Bob's outcome—this is a question that lies entirely outside the unitary part of the theory. The purpose of interpreting quantum mechanics is precisely to tell us how the unitary part of the theory relates to the measurement outcomes witnessed by observers, but RQM as formulated in section 2 is also silent on this question. Indeed, **relativity of comparisons** implies that it is not even meaningful in this version of RQM to ask about the relationship between Alice's perspective and Bob's perspective, so we certainly cannot hope for any guarantee that Bob's measurement outcomes will match Alice's since relativity of comparisons denies that there is any fact at all about whether or not they match. Brown (2007) and Pienaar (2021) have analyzed similar cases and have similarly concluded that extant versions of RQM do not provide any reason to think that Bob's measurement outcomes will match Alice's.

Proponents of RQM may object that we can always compare these perspectives from the point of view of a third observer, Charlie: we can have Charlie measure both Alice and Bob to obtain outcomes $M_C^A, M_C^{B^A}$ for the value of some pointer variables that are supposed to be records of their measurement results M_A, M_B^A , respectively, and then compare these results. Again, **internally consistent descriptions** ensures that he will always find that these values match, thus confirming that from Charlie's point of view, Alice and Bob agree. But this just pushes the problem back to Charlie, and he still cannot find out the values or M_A or the value of M_B^S and M_B^A themselves, only the values of his own measurement outcomes on the relevant pointer variables; so he cannot do anything to determine whether Alice's and Bob's perspectives match from their own points of view. So it seems that there is no way for anybody but Alice to ever find out what Alice's measurement result was. Even when Alice tries to communicate to other observers what result she saw, **internally consistent descriptions** guarantees that everyone will always perceive her to be agreeing with them, and thus no form of communication will ever bridge the gap between Alice's perspective and the other observers around her. As Pienaar (2021) puts it, it seems as though RQM leads to a solipsistic ontology of "island universes."

Now, proponents of RQM might initially be unconcerned by this result; after all, RQM tells us that there is no view from nowhere, and therefore it is to be expected that there will not be any nonrelative facts about the relations between different perspectives. For example, Cuffaro and Hartmann (2021) write that proponents of RQM and other similar interpretations should hold that "the ideal of an observer-independent reality is not methodologically necessary for science and ... modern physics (especially, but not only,

quantum theory) has taught us ... that there is a limit to the usefulness of pursuing this ideal.” However, Adlam (2022) argues that such responses misunderstand the issues at stake. The problem with the lack of communication between perspectives in this version of RQM is not merely that it is unappealing or in tension with naive classical ideas; the problem is that intersubjectivity is a crucial part of our processes of empirical confirmation, because science is a collaborative endeavor that requires us to not only employ our own observations but also to rely on observations performed by other observers. And in RQM, as we have just described it, this is not really possible, because there is no way for Alice to communicate her observations to Bob, so he can never find out about the ways in which their perspectives differ or about the ways in which they are the same. It follows that if Bob is trying to carry out empirical confirmation, he will only be able to confirm a description of his *own* set of relative facts: he will never have any grounds for saying anything at all about the relative facts of other observers. Yet, RQM is clearly intended to be a theory describing the facts relative to all observers across all of space-time, not just the facts relative to a single observer—and it is hard to see how we could ever have adequate empirical evidence for such theory if our epistemic circumstances are really as implied by the six postulates above, so it would seem that this version of RQM itself implies that we should not believe RQM!

Furthermore, Adlam (2022) also argues that since there is not any objective and precise way of tracking the same observer over time, it follows that we must count different versions of the same person at different times as distinct “observers,” in the RQM sense. This means that observers in RQM, as formulated in section 2, cannot even trust that their memories correctly reflect the measurement outcomes obtained by their past selves, since consulting one’s memory is essentially a kind of measurement and RQM as defined in section 2 tells us that the results of this measurement cannot provide any information about the observations that were actually made by past versions of oneself. Thus, observers in this version of RQM cannot regard relative frequencies stored in their memories as meaningful because they have no idea how those relative frequencies are related to the actual measurement results that were observed by their past selves at the time of the measurements. So really, Bob cannot even confirm RQM as a description of his own relative facts, because he will only ever have access to a single observation at a single time, which is not enough information to empirically confirm anything. Thus, in this version of RQM, it appears that 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.

Therefore, in order for it be epistemically rational for us to believe RQM, it is necessary that there should be some mechanism that makes it possible to achieve intersubjective agreement between observers so we can make use of observations performed by other observers (including past versions of ourselves) in order to perform empirical confirmation. One suggestion for how to do this was made by van Fraassen, who proposes some new postulates to answer the question, “What relations are there between the descriptions that different observers give when they observe the same system?” (van Fraassen 2010). For example, he stipulates that for any systems S , O , P (witnessed by ROV), the state of S relative to O (if any) cannot at any time be orthogonal to the state of S relative to $O+P$ (if any). However, van Fraassen’s postulates do not fully solve the problems discussed here, because he uses an additional observer ROV, relative to whom these constraints hold. And as we saw with the example of Charlie above, if these postulates only constrain the relationships between S , O , and P relative to ROV, then they still fail to offer any grounds for a relationship between the perspectives of S , O , and P ; thus, they still do not give observers a means of getting outside their own perspective to learn about other perspectives, which is necessary if empirical confirmation is to be viable.

4 CROSS-PERSPECTIVE LINKS

We suggest that there is a natural way to resolve this issue, arising out of one of the original motivations for RQM. We noted earlier that one of the foundational principles of RQM was the idea that information is physical. And Di Biagio and Rovelli (2021a) explicitly tell us what that means: “If knowledge is physical, it is accessible by other systems via physical interactions.” The argument of the previous section demonstrates that the formulation of RQM we have just examined does not fully respect this guiding principle: Alice’s knowledge is not accessible to Bob by any possible physical interaction, and thus it fails to be physical in any meaningful sense because it does not satisfy the naturalistic criterion of Di Biagio and Rovelli.

This suggests that there may be a simple solution for the problem of intersubjectivity in RQM if we just take this principle seriously. Alice’s knowledge is physical; thus, there *must* be some measurement that Bob can perform on her that will reveal information about her measurement outcomes to him. This motivates the formulation of an alternative version of RQM. Specifically, we suggest removing postulate four and replacing it with the following postulate:

Cross-perspective links (CPL): In a scenario where some observer Alice measures a variable V of a system S , then provided that Alice does not undergo any interactions that 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.

Evidently, CPL solves the problem we introduced in section 3: it is now at least theoretically possible for information about physical values to be shared between different observers; so intersubjective agreement is possible under the right circumstances, although of course introducing CPL does not guarantee that intersubjective agreement will occur in any particular situation. Note also that CPL leads to the same empirical predictions as the original version of RQM and quantum mechanics itself. In particular, it follows from standard quantum mechanics that the expected statistics for Bob’s measurement are the same as the expected statistics for Alice’s measurement; so even though we have an extra constraint in the form of the matching requirement, over many repetitions of the experiment, both Alice and Bob will still see the correct Born rule statistics.

CPL can be understood physically as follows. When a system Alice has information about the variable V of system S , part of what it means for that information to be “physical” is that it should be accessible to other observers who have access to Alice and the ability to perform appropriate measurements. Thus, if Alice is not disturbed too much, it follows that when Bob measures the physical variable representing Alice’s information about S , then Bob should have information about the information that Alice has about S . That is to say, Bob can obtain information not only about the physical representation of Alice’s knowledge relative to him but also about the value of V relative to her, since that information is now understood to be encoded in physical variables that are accessible to B .

We note that the term “information stored in Alice’s physical variables” is being used loosely in this postulate: to say that information is “stored in Alice’s physical variables” simply means it is the case that other observers can access that information by means of subsequent physical interactions. One way of fleshing out this idea would be to say that Alice literally has an ontic state persisting over time that stores those physical variables, but this is not the only possible route. In older versions of RQM, it is usually said that values

have variables (relative to observers) only during quantum events, and in section 5.1, we will suggest this idea can be retained in RQM with CPL; so a measurement on Alice aiming to establish her information about the variable V should be understood not as probing her instantaneous state at the time of the measurement but as “looking back” in a nonlocal way at the value that becomes definite in the most recent interaction.

The postulate CPL, as stated above, is sufficient to achieve our stated purpose of making intersubjective agreement possible—as long as there are at least some cases in which information about V stored in Alice’s physical variables is not destroyed before subsequent measurements, there will be a possibility of achieving intersubjective agreement about measurement outcomes between different observers. However, if we want to apply to CPL to specific cases to determine which particular information may be shared, we must specify what kinds of interactions will destroy the information about V “stored in Alice’s physical variables,” or, more precisely, what kinds of disturbances have the result that subsequent measurements on Alice can no longer access the information. Here we suggest that if A_V is the “pointer” variable of Alice encoding the outcome of her measurement on V (i.e., it is the variable of Alice that Bob must measure if he wishes to know the outcome of her measurement on V), then the information about V stored in Alice’s physical variables is disturbed relative to Bob if and only if Alice subsequently undergoes an interaction with Bob in which a variable A_Q of Alice, which does not commute with A_V , becomes definite. This is in accordance with the standard picture in quantum mechanics where if we measure a system in a basis V and then a noncommuting basis Q , some of the information about the result of the measurement in the basis V is subsequently lost—that is, we cannot reliably find out about that result of that measurement by any subsequent measurements. We can quantify this disturbance by the Heisenberg disturbance relation: if A_Q has taken on a value to within precision δA_Q , then $\delta A_V \delta A_Q \propto \hbar$; that is, the disturbance to the information stored about A_V is inversely proportional to the precision with which A_Q has been measured.

Note that this has the following important consequence. If Alice is a macroscopic system, like a human being, standardly the variables that take definite values during her interactions will be position-basis variables, and these variables all commute with one another; so information stored in Alice’s physical variables will typically be very robust. The only way to erase that information would be to exert very fine microscopic control to measure Alice in a basis other than the position basis, which is not currently within the reach of experimental technique. However, if Alice is just a qubit that has interacted with some other qubit S , then there is a good chance that Alice will subsequently undergo an interaction in which the value of some noncommuting variable becomes definite, and thus information about past definite values stored in the physical variables of microscopic systems is not at all robust and frequently becomes inaccessible.

4.1 STABLE FACTS

Combining CPL with **internally consistent descriptions** implies that if Bob measures the variable V directly on S instead of measuring Alice, then provided that neither S nor Alice has been disturbed since the original interaction between S and Alice, it follows that Bob’s measurement on S will have the same result as Alice’s measurement on S . (Note that if S is subject to a nonzero Hamiltonian, then of course the variable V should be subject to appropriate time-evolution; e.g., if Alice measures the variable V , and Bob’s measurement of Alice takes place after a time t has passed according to some appropriate clock, then Bob will need to measure a variable corresponding to $U^{-1}VU$ rather than V , where U is the time-evolution operator $U^{-i\hbar t}$.) This is because **internally consistent descriptions** implies that if Bob measures both Alice and S , the results of the measurements must match

(in the notation of section 3, we must have $M_B^S = M_B^A$), and **cross-perspective links** tells us that if he measures Alice, his result will match hers (in the notation of section 3, we must have $M_A = M_B^A$) and thus by transitivity have $M_A = M_B^S$. So the combination of CPL and **internally consistent descriptions** entails that out of the substratum of relational facts, we will quickly arrive at a well-established set of intersubjective facts that command agreement across many different perspectives. This postulate entails that as an epistemic community of observers interact, they will build up a shared observable reality composed of a large number of variables whose values all the observers agree on. Healey (2021) offers a detailed account of the way in which such epistemic communities can arise within unified “decoherence environments.”

In particular, CPL plays an important role in the emergence of a stable and *shared* macroscopic reality in RQM. Di Biagio and Rovelli (2021b) demonstrate that within RQM decoherence, processes give rise to “stable facts ... whose relativity can effectively be ignored.” These stable facts arise in the situation when Alice measures a variable V of some system S . The variable V then takes on some definite value v relative to Alice but not relative to another observer Bob who has not interacted with Alice or S . However, the value of V can be considered stable for Bob if in computing the probability for some other variable Q to take the value q relative to Bob during a subsequent interaction involving Bob, and we can write the following:

$$P(q^B) = \sum_i P(q|v_i)P(v_i^A)$$

The point is that this expression looks like a classical mixture; there is no interference between branches of the superposition, and thus if this expression holds, Bob can reason as if V has some definite but unknown value relative to him. Moreover, if Alice and Bob are macroscopic observers, then decoherence will generally ensure that an expression of this kind does indeed hold (at least approximately), and therefore most facts about variables that Alice has observed will be stable relative to Bob in the sense that Bob can treat them as classical observables.

However, it must be stressed that this expression is a description of the situation for Bob, not for Alice. From Alice’s point of view, V already has a definite value at this time, so there can be no nontrivial classical mixture in her description of the situation. Thus, the variables v_i^A appearing in this equation must be understood as facts about the result of Alice’s measurement *relative to Bob*; that is, these variables do not denote the result that Alice has perceived herself as obtaining in the measurement that she has already performed on S . Indeed, since the value of A relative to Alice herself, as selected by the measurement she has already performed, plays no role in the expression above, it would seem that there is no connection between the “stable facts” about V relative to Bob and the value of V that Alice herself has observed. However, once we add the postulate of CPL, we are entitled to replace the facts about the outcome of Alice’s measurement *relative to Bob* with the facts about the outcome of Alice’s measurement *relative to Alice*, since the intersubjective agreement underwritten by CPL assures us that there cannot be any disagreement between these sets of facts. Thus, adding CPL to the theory of stable facts ensures that there not only exists a stable macroscopic reality for each individual observer, but the sets of stable facts making up macroscopic reality relative to different observers can also generally be expected to agree whenever they coincide.

4.2 THE 1996 POSTULATES

Because CPL is all about the circumstances in which agents can acquire information about values relative to other agents, it turns out that the 1996 postulates (Rovelli 1996),

which are not always emphasized in more recent formulations of RQM, are actually very relevant in the formulation of RQM with CPL. In particular, we note that the possibility of “destruction of information” follows immediately from the combination of the two 1996 postulates: together they imply that sometimes when an agent acquires new information about a system, some of their previous information becomes irrelevant. So Rovelli’s 1996 postulates help us understand why CPL holds only insofar as the relevant information is not disturbed.

The 1996 postulates also help make the point that the “destruction” of information as referenced in CPL must be relativized to an observer—for those postulates tell us that if Alice is in possession of the maximum amount of relevant information about *S*, which does not prevent Bob from obtaining some different information about *S*, provided that Bob does not currently have access to Alice’s information about *S*. However, because Bob cannot have more than the maximum amount of relevant information about *S*, it follows that if Bob obtains some different information about *S*, at least some part of the information that Alice has about *S* becomes irrelevant to Bob; that is, he is subsequently unable to access this information, and it will play no role in determining his future interactions with *S*. Thus, Alice’s information can be “destroyed” relative to Bob in the sense that it becomes irrelevant to Bob—although it could potentially still be relevant to some third observer who does not have access to the information that Bob has about *S*. So the question of whether or not information has been destroyed in RQM must be relativized to an observer, as one might expect from the fact that quantum states are relativized to an observer.

5 CHANGES TO RQM

Replacing postulate four with CPL may seem like a minor change, but it has some quite significant consequences for RQM—in particular, the resulting picture is no longer as radically relational as some extant versions of RQM appear to be. Thus, in this section, we will explore what a version of RQM with CPL might look like. Note that we do not necessarily claim that the suggestions we make here represent the only possible way of fleshing out a version of RQM with CPL, but they are useful to give a concrete picture of the ways in which adding CPL might change the theory. It will be evident that in some ways the resulting theory has quite a different character to older versions of RQM, but we will also see that this version of the theory shares significant continuity with older versions of RQM—in particular, the motivating idea of the relationality of *quantum states* is upheld.

Because adding CPL results in quite significant changes to RQM, one might perhaps wonder if the resulting theory is not just a new version of RQM but an entirely new interpretation. Ultimately, this comes down to a choice of naming convention, so we do not think it is a crucial question to resolve. However, in favor of the idea that this is still a form of RQM, we note that the initial formulation and development of RQM was motivated by a variety of different principles (we mentioned several of them in section 2). The standard version of RQM that has emerged in recent years—radical relationalism with no observer-independent facts at all—is one way of formalizing these principles. But RQM with CPL is also a valid way of formalizing these principles: as we have seen, it is particularly motivated by the principle “information is physical,” which was also part of the original motivations for RQM, although it arguably was not well-realized by the version of RQM that subsequently became orthodox. So although the two formulations are undoubtedly distinct, they have a common origin, and thus it is reasonable to refer to them both as versions of RQM.

5.1 ONTOLOGY

The ontology of RQM has been described in various different ways in the literature on the subject, but in at least some presentations, it appears that the whole ontology is supposed to be relational: for example, Smerlak and Rovelli (2007) write that “physical reality is taken to be formed by the individual quantum events (facts) through which interacting systems affect one another ... each quantum event is only relative to the system involved in the interaction.” Similarly, Wood (2014) describes RQM as asserting that “there is no such thing as an absolute, observer-independent physical value, but rather only values relative to observers.” But with the addition of CPL, it no longer seems possible to insist that everything is relational—or at least, it is no longer *necessary* to do so—because this postulate implies that the information stored in Alice’s physical variables about the variable V of the system S is accessible in principle to any observer who measures her in the right basis; so at least at an emergent level, this information about V is an observer-independent fact. This suggests that the set of “quantum events” should be regarded as absolute, observer-independent features of reality in RQM, although quantum states remain purely relational. Thus, we continue to endorse the sparse-flash ontology for RQM as advocated in Rovelli (2018) and Laudisa and Rovelli (2021); however, we now regard the point-like quantum events or “flashes” as absolute, observer-independent facts about reality rather than relativizing them to an observer.

We can continue to endorse the definition of “quantum event” used in previous versions of RQM: a quantum event arises in an interaction between two systems in which the variables of one system take on definite values relative to the other, and vice versa. For example, suppose there is an interaction between Alice and a system S which takes the form of a “measurement” of a variable V of the system S : then the corresponding quantum event can be loosely characterized as “variable V taking value v relative to Alice,” where the probability for the value v is given by the Born rule in the usual way. However, in RQM with CPL, we must be more careful about the phrase “relative to.” Since quantum events are no longer observer-dependent in this picture, we cannot say that the event itself is relativized to Alice. But although the *event* is an absolute, observer-independent fact, it is still correct to say that the *value* v is relativized to Alice. This is 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 . This means that at this stage the value v will not be reflected in any of the relational quantum states assigned by other systems; the value is true only relative to Alice, in the sense that it does not feature in the relational quantum states assigned by any other authors. Thus, RQM with CPL retains the idea that values of variables are relative.

We can also continue to endorse the idea from previous versions of RQM that values have variables only during quantum events. That is to say, although there are observer-independent facts about the quantum events, systems do not have observer-independent ontic states which persist through time storing their variables: variables are only ever defined instantaneously, and thus the only “states” in RQM are the relational quantum states, which of course are not observer-independent. Since CPL requires that the choice of which value becomes definite in a quantum event at one time will often depend on the values of variables in previous quantum events, this means there must be a nonlocal dependence of quantum events on one another. This is to be expected in a sparse-flash ontology, since other well-known flash ontologies, such as the GRW flash ontology (Tumulka 2021), are also nonlocal. However, as with other flash ontologies (Esfeld and Gisin 2013), this approach does not require superluminal signaling or some kind of collapse that takes place on a space-like hyperplane: because there are no states or persisting variables

in RQM with CPL, there is no need 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. Thus, although there is nonlocality in this approach to RQM with CPL, it is not of the objectionable kind that involves hidden influences or preferred reference frames.

Adding CPL to RQM helps us address concerns that have been raised around the definition of a quantum event in older versions of RQM. Specifically, the definition requires us to identify “systems,” but as noted in Wood (2014), it may seem hard to understand what a system really is if everything is supposed to be relative to a system! However, since RQM with CPL postulates a set of observer-independent absolute events, it has the resources to address this question: now a “system” can simply be identified with a set of quantum events that are related to one another in certain lawlike ways, as captured by the formalism of quantum mechanics. Each such system can be characterized by an algebra of physical variables, that is, the set of variables that can take on values in quantum events associated with the system. Recall that every quantum event involves two systems interacting; therefore, different systems do not have to be associated with disjoint sets of events, although no two systems will be associated with the *same* set of events.

Now, one might worry that there is some circularity here: events have previously been defined as interactions between systems, and now systems have been defined as sets of events. However, there is not actually any circularity if we start from the idea that RQM is a theory of a set of quantum events related to one another in lawlike ways: it then turns out that the lawlike relations work in such a way that it is possible to define “systems” such that every event can be regarded as an interaction between two systems. So, in RQM with CPL, the notion of a system is not necessarily fundamental but rather is used as an interpretative tool to help us make sense of the set of quantum events.

Given that RQM with CPL postulates a distribution of absolute, observer-independent events, one might naturally wonder whether it is possible to calculate that distribution as a whole. This cannot be done with pure unitary quantum mechanics because quantum mechanics only provides us with what might be described as a “patchwork” account of the distribution of quantum events, with each individual relational description characterizing the relation between some particular event and the most relevant events in the past, but nothing in the theory characterizing the distribution as a whole. Nor can we calculate the distribution from something like “the quantum state of the universe,” since, as noted in Rovelli (2018), in RQM, quantum states are by definition relational, and there is nothing for the quantum state of the whole universe to be relativized to.

So, perhaps we should conclude that there is actually no unified description of the full set of quantum events, and the “patchwork” description is truly fundamental. Or perhaps some way of giving a unified description will emerge from ongoing research on quantum gravity and quantum cosmology; for example, Höhn (2019) suggests an interpretation of the wave function of the universe in the framework of Dirac quantization, where it is understood as “a perspective-neutral global state, without immediate physical interpretation, that, however, encodes all the descriptions of the universe relative to all possible choices of reference system at once.” This approach might be seen as consistent with RQM, provided we are clear that the universal wave function obtained during Dirac quantization is not a quantum state in the ordinary sense, as it is not relativized to anything. But, in any case, RQM as it currently stands already provides us with a coherent understanding of standard quantum mechanics as a means of locally navigating the set of quantum events: we can continue to assert that “quantum mechanics provides a complete

and self-consistent scheme of description of the physical world, appropriate to our present level of experimental observations” (Rovelli 1996), because this unified description, if it exists, would certainly go beyond our present level of experimental observations.

5.2 RELATIONAL QUANTUM STATES

We have suggested that in our proposed ontology for RQM with CPL, each interaction involving S can be thought of as “looking back” at the most recent interactions involving S to determine the outcome of the new interaction. In this picture, systems do not have states; they just have histories. Thus, the appropriate way to think of the quantum state of S relative to Bob is as a characterization of the joint history of Bob and S ; that is, it describes some set of recent direct and indirect interactions between Bob and S (where “indirect” interactions mean cases where Bob interacts with other systems that are connected to S by some continuous chain of interactions). The possible future interactions for a pair of systems is determined by their joint history, and thus the relational quantum state determined by their joint history is the appropriate tool for predicting the outcomes of their future interactions.

Let us illustrate this idea by considering how it works in the context of quantum interference. If Bob does not know the result of Alice’s measurement of the variable V on the system S , he will describe Alice and S as being in a superposition of all the different possible values of the variable that she measured. We know that, in fact, in the interaction of Alice and S , a single value of V has become definite relative to Alice, and this value is an observer-independent physical fact in the sense that if Bob were to measure Alice or S in the same basis, he would obtain a result that agreed with Alice’s result. Nonetheless, if Bob does not perform this measurement and instead chooses to perform interference experiments on Alice (assuming he has access to sufficiently powerful technology to do so), he is able to see interference effects. This is because the outcome of the interaction is determined by the relational quantum state of Alice relative to Bob, which is a description of the joint history of Alice and Bob, that is, a specific set of past events involving both of the systems. Alice’s observation of the value V is a real observer-independent event, but that event is not a part of the joint history of Bob and Alice because it involves only Alice and not Bob; thus, it has no impact on the possibilities for future interactions of Bob and Alice, so it is not encoded in the quantum state of Alice relative to Bob. Therefore, Bob’s experiments can go ahead as if no definite value for V had ever been realized at all.

This account also helps us see why wave functions should be updated after measurements in the context of RQM, even though there is never any physical collapse or breakdown of unitarity. In this picture, the purpose of the quantum state assigned by Bob to S is to describe information about the joint history of Bob and S that is relevant to their future interactions. Furthermore, the original postulates of RQM set out in Rovelli (1996) together entail that “when new information is acquired, part of the old relevant-information becomes irrelevant.” In the language we have used here, this means that when a new interaction occurs, one or more earlier interactions cease to matter, in the sense that future interactions will no longer depend on them. So each interaction involving S only has to “look back” at a finite number of recent interactions involving S : most of S ’s history will be irrelevant to the outcome of the new interaction. Thus, when a new quantum event involving both Bob and S occurs, that event becomes relevant, and meanwhile some earlier event becomes irrelevant. Thus, the state of S relative to Bob must change to reflect a new set of possibilities for future interactions between Bob and S . Thus, there must be a state update, but that update is not a physical process located in space-time.

5.3 EPISTEMIC VS. ONTIC

The considerations addressed in section 5.2 make it clear that relational quantum states in RQM with CPL are objective physical facts and not merely a summary of an observer's knowledge. However, in the case where Bob is a *conscious* observer, the relational quantum state of system *S* relative to Bob may be linked to Bob's knowledge about the system: for if Bob has a piece of knowledge about *S*, then that knowledge must have its origins in an interaction between Bob and *S* in their joint history, and therefore any knowledge that Bob has about *S* will be in some way represented in the quantum state of *S* relative to Bob (though the converse implication does not necessarily hold—in general, Bob will not have conscious knowledge of all of the interactions in his joint history with *S*). Moreover, in order for RQM to be empirically adequate, clearly it must be assumed that when *S* is a quantum state in the laboratory that has been prepared in a controlled way, such that its past history is well known to the experimenters, then the state of *S* relative to the experimenters will be closely aligned with what they know about it; that is, it will be equivalent to the state that they would naturally assign if they correctly apply quantum mechanics based on their knowledge about how the state was prepared.

RQM with CPL gives us the resources to understand the connection between knowledge and relational quantum states: the reason that the quantum state of *S* relative to Bob is in some circumstances quite closely aligned with his knowledge about the recent history of *S* is because knowledge is not merely an abstract disembodied idea. What Bob does and does not know about the recent history of *S* is a result of the physical facts about his past interactions with *S*, and those are also the facts that determine the state of *S* relative to him. Thus, the idea that “information is physical” offers a new perspective on the traditional dichotomy between “epistemic” and “ontic” views of the quantum state (Spekkens 2007; Leifer 2014). For if we accept that information is always physical, then *knowledge* is physical; therefore, there can be no sharp distinction between epistemic and ontic approaches, since knowledge itself is ontic. Of course, traditionally ontic approaches insist that the quantum state is an ontic state *of the quantum system*, not of the observer who assigns the quantum state, but naturally proponents of a relational view will reject that distinction. RQM tells us that in order to understand the nature of the quantum state, we must consider the observer and quantum system together, and then we will appreciate that any knowledge that an observer has about a quantum system is necessarily included in the ontic facts about their joint history, thus playing a role in determining their possible future interactions.

5.4 METAPHYSICAL INDETERMINACY

The ontology for RQM with CPL that we have suggested here is quite different in character from some previous versions of RQM. In particular, in the past, RQM has been described as having a fairly radical form of metaphysical indeterminacy. For example, Calosi and Mariani (2020) distinguish between “gappy metaphysical indeterminacy” (where no determinate of a determinable is instantiated) and “glutty metaphysical indeterminacy” (where more than one determinate of a determinable is instantiated) and argues that both of these occur in RQM. By contrast, the version of RQM we have presented here exhibits gappy metaphysical indeterminacy, since variables have no definite values in regions between quantum events, but *not* glutty metaphysical indeterminacy, because CPL ensures that whenever there is some fact about the value taken by a variable in a given interaction relative to two different observers; those facts will always match, so we will never have a case where a physical variable takes two different values relative to different observers

(either the variable takes the same value relative to both observers or it takes no value at all relative to one of the observers). Of course, it is still the case in our version of RQM that a given system may be assigned two different quantum states by different observers, but not because of some kind of indeterminacy; quantum states differ between different observers simply because a quantum state describes the *relation* between the observer and the system rather than an absolute feature of the system. In a similar way, the person whom I refer to as “my mother” is probably not the same as the person whom you refer to as “my mother,” but this does not mean that there is any metaphysical indeterminacy about the term “my mother.” The point is that the phrase “my mother” does not describe an absolute feature of a person but a *relation* between the speaker and the person described, and my relation with my mother is not the same as your relation with that person.

6 THE PREFERRED BASIS PROBLEM

We close by showing that CPL also helps to solve a different problem that has sometimes been raised for previous versions of RQM. In particular, it has been objected that an “interaction” will generally not have the form of a measurement, and therefore it will not single out a unique variable of one system that should take a definite value relative to the other system during the interaction. In particular, Muciño, Okon, and Sudarsky (2021) and Brukner (2021) note that we can always rewrite an interaction Hamiltonian in a different basis, and a Hamiltonian that looks like it describes a measurement of variable V in one basis will typically look like it describes a measurement of some other variable V' when we write it in a different basis. So pure unitary quantum mechanics does not suffice to determine which variable in particular should take on definite values during an interaction that leads to a quantum event.

We see two options for RQM to respond to this objection. The first is to stipulate a preferred basis and insist that this is always the basis that takes on definite values during an interaction. For example, it has been noted in the context of the de Broglie-Bohm interpretation that all measurements are ultimately measurements of position, and therefore, for the purpose of explaining our definite macroscopic experiences, it is enough to ensure that some beables have definite values of position, at least during measurements. Thus, in principle, one could imagine a version of RQM in which systems always take in a definite value relative to one another in the *position basis* during an interaction, and no variables ever become definite in any other basis. However, we do not find this solution appealing because this preferred basis is not evident in the quantum formalism, so this approach seems somewhat ad hoc. In addition, we think it is clear that decoherence must play some role in the emergence of a definite macroscopic reality, so we do not consider it reasonable to expect that unique definite values will arise in fundamental interactions before any decoherence has taken place.

The alternative is to agree that quantum events do not typically have the simple form “variable V taking value v relative to Alice.” Rather, they must have a conjunctive form: “variable V_1 taking value v_1 relative to Alice, and variable V_2 taking value v_2 relative to Alice, ...” and so on, specifying definite values for each of the variables singled out by the interaction Hamiltonian in all of the different possible bases for it. The probability distribution over definite values in each disjunct would again be given by the Born rule, and the values in each conjunct would be probabilistically independent.

Now, this solution might seem to undermine the claim that RQM can explain why measurements have definite outcomes. However, RQM need not insist that an interaction singles out a unique value when the two systems involved are, for example, qubits. After

all, standard quantum mechanics does not say anything about what happens when one qubit “measures” another qubit, so there are no predictions here to reproduce. Thus, the information that a qubit has about another qubit need not be such that we can understand what predictions it would correspond to; therefore, RQM is under no obligation to solve the preferred basis problem in the case of interactions between individual fundamental particles. RQM need only show that in the limit, as one of the systems involved becomes macroscopic, then there is a unique choice of variable that takes definite values in the interaction in order that macroscopic conscious beings like ourselves can have definite experiences.

It is clear that decoherence should play some role in this story. And, in fact, decoherence provides exactly what it needed here: it picks out a basis that is dynamically favored and then disseminates information stored in that basis through the environment, for the values relative to a conscious observer in RQM must presumably arise in some way from the combination of the values relative to each of its constituent subsystem. Thus, the contents of the observer’s perspective are not defined by the information associated with a single fundamental particle but by the information that has been disseminated through their brain by decoherence processes. Typically, we would expect that the decoherence basis will favor at most one of the variables that took on a value during the original interaction, so decoherence effectively selects one variable out of the conjunction of variables that appeared in the original quantum event. It is that variable that then has a definite value in the perspective of the conscious observer. Of course, the decoherence process is not perfectly well-defined—there is no exact line between “decohered” and “nondecohered”—but that is not a problem because consciousness also does not seem to be perfectly well-defined: to our best current understanding, it appears to be some kind of emergent high-level feature of reality, so we are certainly entitled to suppose that consciousness can emerge only when enough decoherence has occurred to single out a well-defined preferred basis.

In more detail, consider a Stern-Gerlach experiment such that at the end of the experiment, the atom involved interacts with particles on the screen. The interaction Hamiltonian will not in general single out a unique variable of the atom; therefore, in this interaction, some set of variables of the atom take on definite values relative to the particles in the screen, with probabilities for each variable given separately by the Born rule. Now the particles on the screen interact with photons that in turn interact with receptors in my eyes, and thus decoherence propagates interactions through the particles of my brain. This is the point at which CPL becomes relevant to the story: CPL allows us to say that during the decoherence process, the interacting particles in my brain share information and thus become aligned on the values of certain variables in specific bases. In particular, in the case of nonrelativistic quantum mechanics, the dynamical processes involved in decoherence primarily favor the dissemination of information in a coarse-graining of the position basis (Wallace 2012), and therefore a significant number of particles in my brain will eventually share the same information about the definite value of the atom *in the coarse-grained position basis*—that is, the information about where the atom ended up on the screen. The information about the definite values in all the other bases that were realized in the original interaction does not get disseminated in the same way because these bases are not dynamically favored by the relevant decoherence processes. Indeed, because decoherence plays the role of a “measurement” of the definite values in the position basis, the information in the other bases necessarily becomes inaccessible, so no future interactions can obtain information about the definite values that were realized in all the other bases. Thus, assuming that my

REFERENCES

- Adlam, Emily.** 2022. “Does Science Need Intersubjectivity? The Problem of Confirmation in Orthodox Interpretations of Quantum Mechanics.” *Synthese* 200, article no. 522. DOI: <https://doi.org/10.1007/s11229-022-03989-0>
- Brown, Matthew.** 2007. “Relational Quantum Mechanics and the Determinacy Problem.” *British Journal for the Philosophy of Science* 60, no. 3. DOI: <https://doi.org/10.2139/ssrn.1006232>
- Brukner, Časlav.** 2015. “On the Quantum Measurement Problem.” In the Proceedings of the Conference “Quantum UnSpeakables II: 50 Years of Bell’s Theorem” (Vienna, 19–22 June 2014). DOI: https://doi.org/10.1007/978-3-319-38987-5_5
- Brukner, Časlav.** 2021. “Qubits Are Not Observers—A No-Go Theorem.” Unpublished manuscript. <https://arxiv.org/abs/2107.03513>
- Bub, Jeffrey.** 2012. “Bananaworld: Quantum Mechanics for Primates.” (Oxford, 2016; online edn, Oxford Academic, 24 Mar. 2016), <https://doi-org.proxy1.lib.uwo.ca/10.1093/acprof:oso/9780198718536.001.0001>, accessed 30 June 2023.
- Calosi, Claudio, and Cristian Mariani.** 2020. “Quantum Relational Indeterminacy.” *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 71: 158–169. DOI: <https://doi.org/10.1016/j.shpsb.2020.06.002>
- Cuffaro, Michael E., and Stephan Hartmann.** 2021. “The Open Systems View.” Unpublished manuscript. URL <https://arxiv.org/abs/2112.11095>
- Demopoulos, William.** 2012. “Generalized Probability Measures and the Framework of Effects.” In *Probability in Physics*, edited by Yemima Ben-Menahem, Meir Hemmo, 201–217. Berlin: Springer, Berlin, Heidelberg. DOI: https://doi.org/10.1007/978-3-642-21329-8_13
- Di Biagio, Andrea, and Carlo Rovelli.** 2021a. “Relational Quantum Mechanics Is about Facts, Not States: A Reply to Pienaar and Brukner.” *Foundations of Physics* 52, article no. 62. DOI: <https://doi.org/10.1007/s10701-022-00579-5>
- Di Biagio, Andrea, and Carlo Rovelli.** 2021b. “Stable Facts, Relative Facts.” *Foundations of Physics*, 51, no. 1. DOI: <https://doi.org/10.1007/s10701-021-00429-w>
- Esfeld, Michael, and Nicholas Gisin.** 2013. The GRW Flash Theory: A Relativistic Quantum Ontology of Matter in Space-Time? ArXiv e-prints. *Philosophy of Science* 81, no. 2 (April 2014): 248–264. DOI: <https://doi.org/10.1086/675730>
- Fuchs, C. A., N. D. Mermin, and R. Schack.** 2014. An Introduction to QBism with an Application to the Locality of Quantum Mechanics. *American Journal of Physics* 82, no. 8: 749–754. DOI: <https://doi.org/10.1119/1.4874855>
- Healey, Richard A.** 2021. “Scientific Objectivity and Its Limits.” To appear in the *British Journal for the Philosophy of Science* (August). DOI: <https://doi.org/10.1086/716169>
- Höhn, Philipp.** 2019. “Switching Internal Times and a New Perspective on the ‘Wave Function of the Universe.’” *Universe* 5, no. 5: 116. DOI: <https://doi.org/10.3390/universe5050116>
- Janas, Michael, Michael E. Cuffaro, and Michel Janssen.** 2021. *Understanding Quantum Raffles: Quantum Mechanics on an Informational Approach—Structure and Interpretation*. Cham, Switzerland: Springer. DOI: <https://doi.org/10.1007/978-3-030-85939-8>
- Laudisa, Federico, and Carlo Rovelli.** 2021. “Relational Quantum Mechanics.” In *The Stanford Encyclopedia of Philosophy*, Winter 2021 edition, edited by Edward N. Zalta. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/win2021/entries/qm-relational/>
- Leifer, M.** 2014. Is the Quantum State Real? An Extended Review of ψ -Ontology Theorems. *Quanta* 3, no. 1: 67–155. DOI: <https://doi.org/10.12743/quanta.v3i1.22>

