Studies in History and Philosophy of Science 98 (2023) 40–48

Contents lists available at ScienceDirect

Studies in History and Philosophy of Science

journal homepage: www.elsevier.com/locate/shpsa

Consistent histories through pragmatist lenses

Dr Quentin Ruyant

MSCA Postdoctoral Fellow at Departamento de Lógica y Filosofía Teórica, Facultad de Filosofía, Universidad Complutense de Madrid Ciudad Universitaria, Madrid, Spain

ARTICLE INFO

Keywords:
Quantum mechanics
Consistent histories
Pragmatism
Causation
Modalities
Situations

ABSTRACT

This article adopts a bottom-up approach to theory interpretation, following the slogan “meaning is use”, and applies it to quantum mechanics. I argue that it fits very well with the Consistent Histories formulation of quantum mechanics, interpreted in a particular way that is not the interpretation favoured by original proponents of the formulation. I examine the difficulties and advantages of this interpretation.

1. A bottom-up approach to theory interpretation

The debate on how to interpret quantum mechanics is ongoing perhaps since the first formulations of the theory. What is paradoxical with this situation is that there is a sense in which quantum theory is well understood and not particularly problematic: scientists know perfectly well how to handle its formalism and how to apply the theory in various situations. So, perhaps the problem lies no so much with how to interpret the theory than with what to expect from an interpretation. Perhaps if we follow the slogan “meaning is use”, and if we are careful enough, while climbing up the ladder of abstraction, to maintain the ties to potential applicative contexts, then the infamous interpretive problems of quantum mechanics would simply vanish.

Such a bottom-up approach to theory interpretation in general has been developed in my recent work (Ruyant, 2021a ch. 3) in the form of a theory of scientific representation. This theory first proposes an account of the application of theoretical models in concrete experimental situations, and then considers the way more abstract models and theories relate to their potential applications. The main originality of this account is that it considers that models are only representational in concrete contexts. Scientific models outside of a concrete context of application play a normative role with regards to their potential applications. Although the notion of a context of application can be understood in a rather permissive way, the experimental situations implemented in order to confront theoretical models with experience constitute our main exemplars, and so, they should play a central role in theory intepretation. This is, in essence, how the slogan “meaning is use” is understood in this account.

This account did not originally aim at solving the interpretive problems of quantum mechanics. The purpose of this article is to examine to what extent it could be informative in this respect.

I first present the theory of representation itself, and then I apply it to quantum mechanics. The best formulation of quantum mechanics for this purpose is the Consistent Histories formulation (Griffiths, 2003, 2013; Omnes, 1999; Gell-Mann & Hartle, 1993; Hartle, 2011), where the content of models is relative to a *framework*. I examine the problematic interpretations of the framework that are traditionally proposed, and argue that a pragmatic interpretation along the line of the theory of representation presented here does not encounter the same difficulties. In conclusion, I consider what the metaphysical lessons of quantum mechanics are if we adopt this perspective.

2. A theory of representation

2.1. Experimental contexts

According to the theory of representation that I will present here, the interpretation of a theory should be primarily based on the way the theory is applied in concrete experimental contexts, so we first need an account of these experimental contexts.

When applying a theory (in order to confirm it or to use it), epistemic agents are typically interested in particular properties of a concrete target system. They have empirical access to these properties by means of measuring instruments with limited precision. These properties are *determinables* that can take several possible values, and the agents typically

---

E-mail address: qruyant@ucm.es.

https://doi.org/10.1016/j.shpsa.2023.01.001
Received 22 June 2022; Received in revised form 3 January 2023;

want to know which values or which complex combinations of values can actually be realised and which cannot, so as to guide their interaction with the target system. They use a theoretical model to make their inferences. The notion of a *context* aims at capturing the conditions of application of a theoretical model, which correspond to the interests of such agents with their limited abilities.

A physicist, for example, could be interested in the distribution of the position, along a particular physical direction, of a population of electron in a magnetic field. She will set up her apparatus accordingly (the measuring instrument, the generator of the magnetic field, etc.), and measure positions with a certain precision. In this context, the property "position" can take a finite set of conceivable values corresponding to the intervals that the measuring instrument can discriminate. From this basic example alone we can observe the following features that our notion of context should capture:

Theory ladenness: Properties of interest are generally described in a theoretical vocabulary (needles and computer displays are of interest only insofar as they inform us about these theoretical properties).

Coarse-graining: Properties of interest are generally coarse-grained, because the precision of measurements is limited, and they can be aggregate (or statistical) when measurements are performed on a population of individuals.

Finiteness: Properties of interest and their possible values are described by finite means: experimenters do not have infinite resources.

Taking into account these three features, we shall formalise a context of application as a finite set of possible experimental outcomes described in a theoretical vocabulary. An experimental outcome corresponds to a finite set of co-instantiated coarse-grained properties. In the case of a dynamical system, this would be a *history* for the system, that is, a finite sequence of *events*, where an event is the instantiation of a coarse-grained property in a particular time interval, corresponding to a measurement outcome. The context is therefore a set of possible histories in this case.

Since the agent does not know what to expect from the experiment beforehand, these possible outcomes must cover the entire space of a priori possibilities (we know a priori that *at least one* outcome must be realised), and since the experimenter has empirical access to the properties and expects the experiment to be fully informative, each outcome must be distinguishable from any other (we know a priori that *at most one* outcome can be realised). So, in the end, we know that one and exactly one of the outcomes specified by the context will be obtained, but we do not know which. In other words, the context is a *partition* of the space of a priori possibilities for the target system. There are infinitely many ways of partitioning this space of possibilities, but the selected partition corresponds exactly to the interests of the agent: what she wishes to, and can discriminate during the experiment.

This formalisation is close to the concept of a *model of experiment*, which Suppes (1969) describes as a set of conceivable data models. It is also related to the notion of a *question* asked to the world used by Rovelli (1996), and, as we will see later, to what Griffiths (2003) calls a *framework* in the Consistent Histories formulation of quantum mechanics.

Framed in this way, the notion of a context could seem purely conceptual, corresponding to something that is entirely "in the mind" of the agent, a focus on a particular way of coarse-graining the state of the entire universe maybe. However, there are two more features of the context that one should consider:

Indexicality: Properties of interest are generally specified relative to the position of the experimenter rather than in absolute terms. For example, the direction of position measurements, or the target system itself, can be identified by ostentation (*this* direction of the laboratory, the electrons generated by *this* apparatus).

Performativity: The context usually prescribes certain manipulations in order to be realised. Very often, the experimenter must actively create the situation she is interested in, and not simply observe the world.

These features have to do with how the context relates to experience, and more precisely, with the way the theoretical vocabulary is translated into concrete observations and manipulations. We can assume that this translation follows experimental norms that tell the agent how particular physical quantities should be operationalized in a given context.[1] Indexicality implies that experimental norms are indexical: they are a function of the broad context in which the agent is situated. Performativity implies that experimental norms can prescribe actions, and not only observations. These actions are not always on the target system directly (observing a distant star only requires manipulating instruments here on Earth), but sometimes they are. This entails that, as a matter of generality, the context is not merely "in the head" of the agent: it is actually instantiated by the agent. Moreover, this instantiation is intentional, so experimental contexts are at least apparently teleological.

2.2. Models and theories

It is commonplace, in metaphysics, to interpret the relation between abstract theoretical models and reality in terms of a direct correspondence between the two: the model literally describes some part of reality. The theory of representation presented here claims that a finer analysis is required. In particular, we should distinguish two levels of representation: (1) the way models are used to represent and make inferences in applicative contexts, and (2) the way standard theoretical models *prescribe* appropriate representational uses within a scientific community.

The first level is considered primary, in the sense that the second level can only be understood with reference to the first level.[2] Among the motivations for this approach is the fact that modelling practices are sensitive to the intentions of model users and to contextual aspects (Giere, 2010; Morgan and Morrison 1999) while still following communal norms (Boesch, 2017).

Without going into too much details, we can understand an interpreted model (our first level) as a vehicle (for example, equations on paper) endowed with demonstration and interpretation rules (Hughes, 1997). These rules are followed by the user of the model in order to make inferences about the target system. The demonstration rules, together with the structure of the vehicle, constrain the inferences that can be made, and the interpretation rules map elements of the vehicle and context, for example, a symbol on paper and a property that would be measured, so that the inferences can be about the target system. The main idea is that the interpreted model, assuming that it is reliable, tells its user which, among the conceivable outcomes specified by the context, are *really* possible and which are not (perhaps with a probability weight). For example, when using a mechanical model in order to anticipate the position of a pendulum that will be measured, an agent can infer which among the outcomes that are a priori conceivable given the instrument that is used are in fact possible given the nature of the pendulum.

So, an experimental context defines a set of a priori possibilities for a system of interest, and the inferential content of an interpreted model can ultimately be reduced to a weighting of these possibilities (which can be all-or-nothing in the case of deterministic models). We can interpret this weighting as corresponding to objective (e.g. natural or physical) possibilities, because they are independent of the contextual knoweldge state of the user. Assuming a counterfactual theory of causation, such as

---

[1] Norms of experimentation are presumably geared towards the stability and reproductibility of experimental results and the convergence of various operationalisations. They can be relatively independent from the theory and survive theory changes. They do not necessarily imply that the theoretical terms of the context refer to real properties in the world, although they do not imply the contrary. For the sake of exposition, it will be convenient to use a (entity) realist way of talking, and to state that a context directly refers to real properties (determinables) of a concrete target system.

[2] This approach is actually a transposition of Grice's distinction, in philosophy of language, between speaker-meaning (the beliefs that a speaker overtly intends to transmit to the audience by uttering a sentence in context) and expression-meaning (the appropriate use of words and sentences in order to transmit beliefs within a community). It roughly follows the same slogan: "meaning is use".

Woodward (2004)'s, such modal content can often be interpreted as a causal or dispositional structure assigned to the target system.[3]

As for the second level, a theoretical model such as the quantum model of the hydrogen atom is not a concrete vehicle, but an abstract entity, a mathematical structure. It is not destined to be applied to only one type of context, because there are many ways of measuring a hydrogen atom or a population of such atoms. However, it is defined on a range of potential contexts, its domain of application, and given any particular context in this range, it should *prescribe* appropriate interpretation rules, for example, that some element of its structure should be interpreted as the energy of the electron in a context where energy would be measured, and not as anything else. According to this theory of representation, prescribing contextual interpretations (including for mererly hypothetical contexts) is the main function of a theoretical model. A theoretical model can therefore be analysed as a function from context to interpreted model. If we restrict ourselves to the inferential content of interpreted models, this is a function from context to weighting of the context.

Does this interpretation match the structure of actual scientific models? In order to be such a function, the model should incorporate an abstract representation of the range of contexts to which it potentially applies. The state-space on which a model is typically defined in physics should be considered such a representation. A region of state-space corresponds to a coarse-grained property for the target system of a context, and a finite partition of state-space (or a sequence of partitions at various times for dynamical systems) therefore corresponds to a context of application, assuming reference to a concrete system. The complete state-space encodes all possible types of contexts of application, which are its possible partitions.

We have seen that contexts are finite and discrete structures. On the contrary, state-spaces are often defined on a continuum. The reason is that the range of potential contexts of application is a priori infinite. A theoretical model in physics will typically specify a continuous trajectory on state-space, or a set of trajectories if initial conditions are left unspecified, or a probabilistic distribution on state-space. This structure can be projected onto a context (a finite partition of state-space) in order to retrieve a weighting of the elements of this context. If, for example, the model specifies a trajectory and the context a set of possible histories defined on a set of discrete instants, we have an all-or-nothing weighting: only the history constituted of all the regions through which the trajectory passes at these instants is selected as possible. In the case of a probabilistic distribution over state-space, a probability weight will be applied to each contextual property or history by integrating the distribution over the corresponding regions. So, in essence, a mechanical model of this kind is indeed a function from context to weighting of context.[4]

In sum, according to the theory of representation presented here, a theoretical model defined on a continuum is not directly representational, in so far as it abstracts away from contexts (or generalises over contexts). It has a normative and indexical status, which can be formalised as a function from context to content. Only its interpretations in finite, bounded contexts, actual or hypothetical, are directly representational, and they usually represent modal relations between the properties of the system that are accessible in this context. But it is a mistake to reify the content of an abstract theoretical model.

We can make one final step up the ladder of abstraction by understanding theories as families of abstract models related by a common vocabulary and common rules of construction. The theory typically tells us how a model should be constructed, how models should be combined or separated into sub-models, etc. This guides the extension of the theoretical framework to new domains of application. In physics, we could see the theory and its symmetry principles as encoding interesting relations between various possible contexts of application.

2.3. Realism and explanations

This theory of representation is deflationary in spirit. It has instrumentalist flavours. This could raise suspicion. A scientific realist, in particular, could have hoped for a more "pictorial" account of representation, where models do not merely account for possible measurement results in a range of possible contexts, but somehow describe "what there is" independently of these contexts, and *explain* contextual measurement results. Nothing prevents the realist from completing this theory of representation with such aspects, and maybe justify this move by inference to the best explanation. However, I believe that the widespread use of idealisations in science, as well as problems of theory change, plead in favour of attributing a mere psychological role to the pictorial content of models. At most, they facilitate their manipulation, but we should not take this content too seriously when it has no practical implication. Other reasons to be sceptical are the difficulties that plague realist interpretations of quantum mechanics, and the mutiplicity of practically equivalent formulations of the theory: it is far from clear that any theory provides us with a unique, straightforward picture of reality. Having said that, it should be noted that this theory of representation is a priori compatible with the kind of entity realism defended by Cartwright (1983) and Hacking (1983). Theoretical models do not represent directly, but their contextual interpretations do represent modal relations between the instantiated coarse-grained properties of particular physical systems with which experimenters interact.

A worry could be that this deflationary stance makes us unable to account for the explanatory power of models and theories. But remember that the representational content of interpreted models is *modal* (it tells us about objective possibilities in context), which fits with counterfactual and causal theories of explanation (Woodward, 2004), and that theoretical models are *unifying* (they synthetise an infinite range of possible applications), which fits with unificationist theories of explanation (Kitcher, 1989). It is also noteworthy that the relevance of explanations often depends on a certain focus on properties of interest (van Fraassen, 1980 ch. 5), which comes out naturally with this theory of representation. So, even if deflationary, this theory of representation can account for the fact that models and theories offer us a certain grasp or understanding of the phenomena they represent, because they inform us about what to expect in a large range of possible contexts.

Another worry could be that this theory of representation implies reinterpreting scientific discourse rather than taking it "at face value". However, the notion of "face-value interpretation" is far from clear, and nothing indicates that it implies a direct, metaphysical correspondance between the content of abstract models and reality. On the contrary, as I argued in (Ruyant, 2021a ch. 8), the way scientists talk about their models and theories in abstract contexts (and notably when they consider

---

[3] Arguments to that effect can be found in the literature on causal modelling. In particular, Pearl (1988, ch. 3) has shown that a causal graph can be uniquely identified from a probability distribution over variables assuming that (i) the variables are time-indexed, (ii) a probability is assigned to every possible values of the variables, (iii) there are no latent variables in the graph and (iv) the probability measure satisfies a Markov and minimality condition with respect to the graph. Conditions (i) and (ii) concern the probabilistic structure itself, and arguably, they are fullfiled in many contexts of application. Conditions (iii) and (iv) roughly amount to assuming that the variables of the context are causally relevant, with appropriate fine-graining, and that there is no latent common cause (so as to eliminate spurious correlations when identifying the relevant causal graph). We could assume that this is ensured by the stability sought by experimental practice.

[4] This is an example from physics, and the reader could wonder if the approach is applicable in other fields, for example in biology. I believe that it is. The recipe is roughly the following: (1) extract the propositional content of the model, (2) apply a standard possible world (intensional) analysis to these propositions and (3) interpret a context as a partition of the sub-set of possible worlds where some relevant experimental conditions are present.

different theories to be equivalent) favours a more pragmatic interpretation of their discourse, identifying the truth of theoretical statements with their ability to withstand any possible situation, given some general rules of contextual interpretation. This is exactly the approach adopted here.

The problem could be that many models in science do not have any direct interpretation in terms of experiment or observations, and yet, they do not seem devoid of content. What are the potential experimental contexts associated with a model of the core of a neutron star, or with a model of the early universe, for instance? Such models have indirect implications for direct observations that we could make here and now, but they also speculate far beyond them. If they are interpreted exclusively in terms of their indirect observable consequences *for us now*, they are not really interpreted "at face value".

Whether this is a serious problem for the present theory of representation depends on how permissive we are with respect to our understanding of potential contexts of application. If we accept as legitimate hypothetical contexts corresponding to direct observations of some characteristics of the early universe or of the core of a neutron star, this is less a problem: the model still says something about what would be the case if these experimental contexts were instantiated. Now, the counterfactual "if we were to measure such characteristics of the early universe directly" might seem problematic in several respects, since it would be impossible for us to do so. Furthermore, we might think that scientific models of neutron stars purport to be informative about what is *actually* the case even when no experiment is carried out, and not only about counterfactual states of affairs.

My suggestion is that we replace the formulation in terms of counterfactual experiments with ampliative reasoning, so as to incorporate more, in our range of hypothetical contexts, than what could effectively be experimented upon given our situation here on Earth, and more, in our range of actually instantiated contexts, than what we, human, experiment on. This idea will be made more precise towards the end of this article. The important point is that all potential contexts are still conceptualised and hypothesised in reference to the actual ones that we do implement, so that experimental practice remains our ultimate basis for interpreting models and theories.

Even if we are permissive in our identification of potential contexts of application, this way of thinking is still markedly distinct from the approach towards theory interpretation that is usually adopted in contemporary metaphysics. This will appear clearly when exploring its implications for quantum mechanics, which is the main objective of this paper.

3. Consistent histories

3.1. *The formulation*

I mentioned above that the notion of a context is formally similar to the notion of a framework introduced in the Consistent Histories formulation of quantum mechanics (Griffiths, 2003, 2013; Omnes, 1999; Gell-Mann & Hartle, 1993; Hartle, 2011). Indeed, this formulation (but not its usual interpretation) is particularly adapted for the approach advocated here.

As explained in the previous section, a coarse-grained property for a system can be associated with a region of state-space in classical mechanics. The analog, in quantum mechanics, is a sub-space of Hilbert space, or, equivalently, a projection operator on a sub-space. If the projector is indexed to a particular time, it corresponds to an *event*, that is, an instantiation of the property at this time. A *history* for a dynamical system is a sequence of events. A *framework* is a family of mutually exclusive histories that sum to identity, which means that they constitute a partition of the space of possibilities. The framework is *consistent* if it respects a consistency condition (related to the notion of decoherence) which guarantees that the events and histories of the framework consitute an event space. Roughly speaking, this means that we will be able to apply standard probability calculus without running into contradictions. This consistency condition depends on the Hamiltonian of the system.[5]

Different frameworks can be related by relations of refinement or coarse-graining. A refinement is just a finer partition of the Hilbert space that discriminates more possible states. Two frameworks are *compatible* if they have a common refinement. In the classical analog, this is always the case, since it is always possible to find a common refinement for any two partitions of state-space by intersecting them. So, one can dispense with the notion of a framework and work directly on state-space, which is the common refinement of all possible frameworks. The main peculiarity of quantum mechanics is that some frameworks are incompatible (when their respective sub-spaces are not orthogonal). There is no common refinement to all possible frameworks, and so, a choice has to be made.

Once a consistent framework is selected, it is possible, given an Hamiltonian, to use a generalisation of the Born rule in order to assign probabilities to the events and histories of this framework (or conditional probabilities if inital conditions are left unspecified). One can then make inferences about the properties of the system that are included in the framework and their possible values, including counterfactual inferences. So, the standard quantum theoretical model is precisely used as a function from context (framework) to weighting of the context, in accord with the theory of representation presented in the previous section.

According to Griffiths, we should only make inferences, including counterfactual inferences, within a framework, and refrain from changing the framework in the middle of our reasonning, so as to avoid the paradoxes normally associated with quantum mechanics. This is what he calls the *single framework rule*. However, if two frameworks are compatible, they have a common refinement, and then it is possible to make inferences using this common refinement.

Note that the single framework rule does not necessarily bring strong limitations on the kind of counterfactual inferences that it is possible to make and on the possibility to recover ordinary counterfactual talk. So, for instance, the single framework rule prevents us from considering alternative possibilities where a given photon would be polarised either in the x direction or in the y direction at the same time, because these possibilities belong to incompatible frameworks. However, it is possible to consider as our object a larger system composed of the photon and an apparatus measuring its polarisation either in the x or in the y direction, depending on the orientation of the apparatus. With an appropriate framework, one can consider in the same piece of reasoning alternative possibilities corresponding to "what would happen if the polarisation were measured in the x/y direction", yielding definite polarisation values for the photon in each case: these possibilities no more belong to incompatible frameworks, because of their association with different possible orientations for the apparatus (these orientations can be considered orthogonal projectors, since they are the eigenstate of a macroscopic observable). This is how Griffiths (2003, ch. 24) represents Bell-type experiments (connecting the two apparatus to "quantum dice" that affect their x/y orientation randomly – see also ch. 19 for his analysis of counterfactuals).

The upshot is that which properties (determinable) of an object are compatible or not depends on the boundaries of the system considered. It is generally possible to broaden our scope in order to access the alternative states of affairs that we are interested in within a single consistent framework. In particular, when considering ordinary macroscopic systems with many components, we are almost guaranteed to find a coarsed-grained framework that is both consistent and arbitrarily close to a representation of the classical determinables that we are interested in:

---

[5] Alternatively, a *histories Hilbert space* can be defined for the dynamical system as the tensor product of copies of the Hilbert space of the system indexed to different times, and a history then corresponds to a complex projector on this space, which is a tensor product of instantaneous projectors. The consistency condition is then an orthogonality condition on the histories Hilbert space.

what Griffiths (2003, ch. 26) calls a *quasi-classical framework*. Therefore, the single framework rule does not rule out ordinary counterfactual talk about classical properties.

The interpreted models that are obtained once the framework has been selected are very well behaved: they do not feature any. of the problematic aspects of quantum mechanics. In particular,

The measurement problem: Wave-function models do not predict the unicity of measurement outcomes without additional structure. The interpreted models of Consistent Histories predict this unicity by assigning well-behaved probabilities to possible outcomes.

The ontological interpretation: There is no straightforward ontological interpretation of a wave-function model, but several options that require additional structure (Belot, 2012). The ontological interpretation of an interpreted model of Consistent Histories is straightforward: it is a modal structure of events.[6]

The relation to classicality: Because of the two previous points, the relation between quantum mechanics and classical mechanics remains unclear with wave-function models. With Consistent Histories, choosing a coarse-grained quasi-classical framework allows one to recover the predictions of classical mechanics in approximation. Assuming a counterfactual theory of causation, it might be possible to interpret the probabilistic structure of events described by Consistent Histories as a causal structure (although this would require further analysis that would not fit here[7]). This would provide an ontological continuity with higher-level disciplines such as chemistry and biology, where causation is central.

Locality: An interpreted model of Consistent Histories is dynamically local. According to Griffiths (2003, ch. 24), although the framework can incorporate non-local *properties*, there are no non-local *influences*: any correlation at a distance can be explained by a common cause within a choosen framework. Incidentally, the interpretation can be extended to relativistic spacetime (Griffiths, 2002) (it actually comes close to Feynman's path integral formulation).

There is of course a price to pay for all these advantages, which is the relativity to the framework. The Consistent Histories formulation has received various criticisms in this respect. Let us examine them.

3.2. The problem of the framework

We expect from a good theory that it correctly account for the phenomena that we observed in the past, and that it correctly predict the future. Quantum mechanics can be considered very successful in this respect. But according to some critics, the Consistent Histories interpretation proposed by Griffiths (2003), and elaborated by Omnès (1995) and Gell-Mann and Hartle (1993), cannot fulfill this aim (Dowker & Kent, 1996; Okon & Sudarsky, 2014a; 2014b).

In order to understand the criticisms, it is important to first understand how Griffiths interprets frameworks (the interpretations from other authors are similar and encounter similar problems). We must select a consistent framework, that is, an event space, in order to represent and make inferences on a system, but according to him, the framework need not correspond to what is measured on the represented system: this would mean introducing an unexplicated notion of measurement, which is illegitimate for a realist interpretation of the theory. So, he takes the choice of framework to be arbitrary: it is merely a perspective adopted by the modeler, which depends on the properties she is interested in. Strictly speaking, all consistent frameworks are on a par.

The problem with this approach is that it leaves us unable to account for past measurement outcomes. We know, from experience, that outcomes corresponding to one specific framework (the one associated with measurements) have occurred. So, our experience seems to tell us that one particular framework was actually instantiated, but the Consistent Histories approach, in Griffiths's interpretation, is unable to account for this. Unfortunately, incorporating the full experimental setup in the model, including measuring instruments, does not help, because the theory still does not tell us that we should adopt a framework that corresponds to the pointers of these instruments. Okon and Sudarsky (2014b) argue that something extra-theoretical is required to identify measurement situations in order to connect the theory and our experience, so that the aim of providing a self-sufficient realist theory ultimately fails.

This problem concerns the account of our past experiences, but an analogous problem occurs with respect to our future experiences. We seem to know that the world behaves in a classical manner at macroscopic scales, at least in very good approximation. Objects generally have well-defined positions and velocities. This classicality does not depend on an arbitrary choice from us, and we expect it to persist in the future. However, Griffiths's Consistent Histories interpretation fails to explain or to predict the quasi-classicality of the universe. As explained before, quasi-classical frameworks can be used to recover the predictions of classical mechanics. However, these frameworks are not the only consistent ones, far from it. What is missing is an account of why quasi-classical frameworks in particular are selected by nature. Furthermore, most consistent frameworks that are quasi-classical in the past cease abruptly to be quasi-classical in the future, so adding constraints from the past does not resolve the issue.

In the end, all these problems originate in the interpretation of the framework as an arbitrary perspective choosen by model users. This interpretation entails that the content of representation is relative to an arbitrary choice, which is somehow puzzling. This leaves unclear what, in our representation, does correspond to a reality that is independent from the user of the model. On the one hand, the single framework rule precludes the content of all framework-relative representations to correspond to the same reality, because this would lead to contradictions. On the other hand, no framework-relative representation should be considered more veridical than any other. The remaining options are not very satisfying: either none of these representations is veridical, but then the theory is not a realist picture after all, or what these representations describe cohabit in “parallel universes”, but then the interpretation is much more metaphysically revisionary than it looks at first sight.[8]

Now if, as I propose, we associate the framework not with an arbitrary representational choice, but with an experimental context that is intentionally instantiated by model users, that is, if we *reify* the framework, we lose one of the main motivations of Griffiths's interpretation, which is to get rid of the anthropocentric notion of measurement. However, we solve at least some of these problems. In particular, there should not be any problems in accounting for past experimental results, assuming that we know which framework was instantiated by the agents during these experiments. Frameworks correspond to the properties (determinables) of a physical system that are made accessible to its environment, including us: what I have called the context. The context that is instantiated can be

---

[6] Note that there are no collapses in this view. Since the wave-function is not representational, nothing collapses. However, the dynamics is stochastic.

[7] See footnote 3: the probabilistic structure of a Consistent Histories model respects conditions (i) and (ii) by construction. Remember that conditions (iii) and (iv) concern the causal relevance or appropriate fine-graining of the variables that figure in the framework. Arguably, this could depend on the selected framework, and the question is whether and why this should be assumed of quasi-classical frameworks.

[8] This last option is briefly discussed in Dowker and Kent (1996). It is similar in spirit to the Many-Worlds Interpretation, but quite different, since it is not different histories, but different frameworks or perspectives that constitute alternative worlds. However, since the same measurements are compatible with many different frameworks, it is still true that incompatible outcomes of a given measurement are sometimes instantiated in parallel universes, as in the standard Many-Worlds approach (the contrary assumption would violate the probabilistic predictions of quantum mechanics (Kent, 1997)).

inferred from an experimental configuration, presumably by means of connections between theoretical properties and their standard operationalisations that are given by norms of experimentation. The veridical representations are simply the ones that are obtained by choosing frameworks that correspond to contexts that are actually instantiated in the world.

In one sense, it should not be a surprise that we need some extra-theoretical inputs (an actual experimental configuration) and some rules connecting them to theoretical properties in order to interpret the theory and to know how to apply it. The point of a theory is to explain the phenomena we experience, and so, the theory must assume a pre-theoretical grasp of these phenomena, otherwise it is of no import. The fact that these extra-theoretical inputs can be local, associated with particular contexts, is not surprising either. In order to make predictions, we must translate our experience into theoretical descriptions of an initial configuration, and even in classical or relativistic physics, a local coordination is required in order to interpret a model: one must associate, for example, the X axis of a mathematical space with a direction in physical space that is referred to from one's perspective, for example by pointing at a measuring instrument. The framework is similar to a coordinate system in this respect: providing a framework is just describing some aspect of an experimental configuration from our own perspective, in ways that are relevant for making predictions.

An important difference between frameworks and coordinate systems is the single framework rule. In classical mechanics, we are free to switch from one coordinate system to another in our reasonning. Here, as in Griffiths's interpretation, counterfactual reasoning should be restricted to a context associated with a framework that specifies the determinables involved. Then the counterfactuals concern the possible values of these determinables only.[9] So, if our experimental context concerns the measurement of the spin of an electron in direction x, we are only allowed to reason about the possible spin values in this precise direction (the measuring environment is held fixed "in all possible worlds", so to speak). This allows us to avoid the paradoxes usually associated with quantum mechanics. As explained before, this limitation is not very strong: it is still possible to consider what would have happened if something else had been measured on the electron by stepping back and considering a more comprehensive context that includes the measuring apparatus and its orientation. Then the orientation of the apparatus is no more held fixed, and spin values for the electron in different directions can appear in a single consistent framework if they are combined with different orientations for the apparatus.[10] But the single framework rule could still seem puzzling. What shall we make of it?

It has been argued that this rule is ad-hoc in Griffiths's interpretation (Kent, 1997). However, assuming, as we did in the previous section, that only interpreted models are truly representational, I think that it makes perfect sense: changing the framework means switching to a model that applies to a different concrete context, and we cannot assume a priori that we are still talking about the same objects and properties or that the two contexts can co-exist (remember that contexts can be performative). Unless we are able to provide a more comprehensive context (a common refinement), that is. For example, we cannot assume a priori that it is possible to measure the spin of an electron in two different directions at the same time. What we took for granted concerning coordinate transformations in classical mechanics appears not be an innocuous assumption after all (see (Ruyant 2021b) for a pragmatist analysis of coordinate transformations). But again, this should not bring strong limitations to ordinary discourse, since a consistent quasi-classical framework will normally be available to represent the classical determinables we could be interested in and their possible values.

However, not all problems are resolved by this approach, so, more should be said. Firstly, it does not seem that the theory can really account for or predict the classicality of the world under this interpretation. The question is why quasi-classical frameworks are apparently selected by nature, but if frameworks are actually selected by experimenters, the classicality of the world remains mysterious. Secondly, the idea that experimenters instantiate contexts, and thereby choose which properties objects have or not, is problematic. We are able to represent galaxies that are remote both in space and in the past, and the idea that they would only have their properties in virtue of us instantiating an experimental context to observe them seems far too unreasonable. More generally, the approach seems too instrumentalist. Common sense tells us that objects continue to exist even when nobody is looking, and that the universe existed before cognitive agents were able to experience it. If what exists only exists relative to an experimental context, these common-sense intuitions are lost, and that seems too high a price to pay to make sense of quantum mechanics (this echoes the objections against our theory of representation discussed in the previous section).

3.3. *The problem of classicality*

Let us address the last concern first, and remark that the claim that framework-relative representations are veridical when the corresponding context is instantiated does not imply that objects and properties cease to exist when no context is implemented by any cognitive agent. The point is epistemological rather than ontological: the theory tells us something about the world only if we provide as input a viewpoint associated with a real context. This does not entail that there is nothing outside of a viewpoint, only that what there is lies beyond the reach of the theory.

Now we could follow a radical neo-Kantian line and assume that what lies outside of a viewpoint is simply not representable. Or we could follow a more common-sensical approach, and accept that there are objects with well-defined properties out there, even when no one experiments on them. If these objects are describable, this must be relative to a framework choice, so we should accept that nature does select frameworks outside of our activities, that is, that contexts can exist in the world without any cognitive agent instantiating them. If this is so, such contexts should be associated not with experimenters, but rather with natural objects and the way they relate to their environment. We could call these "natural contexts" and associated objects *situations* (Ruyant, 2021a ch. 4). For any identified situation (where the identification includes an event space), the theory can tell us what properties it is disposed to instantiate. What the theory does not say is which contexts and situations are instantiated: for this, we need direct experience and intentionality.

In this view, an experimental context is just the special case of a context that is intentionally created by cognitive agents. Agents place a particular object in a controled environment, which allow them to know the context that is instantiated for this object. If intentional contexts are necessary for representation, this is simply because experimenters must take part in the relevant environment of an object in order to know anything about it, and the experimental context then represents their

---

[9] In a first approach, we could consider a simple analysis of "If A had been the case, then B would have been the case" as equivalent to "for all the histories of the framework, if A then B". The fact that counterfactuals are relative to a context by stipulation could be enough to solve some of the traditional issues with counterfactuals (see for example Williamson, 2020, pt. II).

[10] Note that in this more comprehensive context, the norms of interpretation that are used in order to determine the framework from our experience are different. It is our direct observation of the apparatus that determines the relevant framework for the electron + apparatus system, instead of the apparatus itself determining the framework for the electron alone (but the existence of this broader context does not imply that the narrow context, restricted to one particular measurement being performed on the electron, is nonexistent).

partial perspective.[11]

A consequence of this proposal is that it is not entirely up to agents what context to implement: there are constraints on what can and cannot be done, because agents are themselves part of and related to natural situations beyond their control. We know by experience that some contexts are easily implemented and that others are not. For example, measuring microscopic properties with very high precision requires very big and finely tuned experimental setups, such as particle colliders, whereas measuring macroscopic distances is much more straightforward from our situation. And we know for sure that measuring a property that would formally correspond to a macroscopic superposition of positions for a distant galaxy is practically impossible. "Practical" here means: as a consequence of our situation in the universe. This responds to one of the worries mentioned previously: if it is not entirely up to experimenters to choose which context to implement, then is it not necessarily up to them to decide which kind of properties remote galaxies must have and which they cannot have.

These remarks can help us solve the problem of classicality as well. True, quantum theory, thus interpreted, does not tell us why the world is quasi-classical, nor does it predict that the world will remain quasi-classical. But we know by direct experience that quasi-classical situations are accessible to us with no particular effort, while purely quantum contexts are hard to implement. It has always been this way, and we can be confident, by induction on situations, that it will remain so. This knowledge is knowledge of our own contingent situation in the universe, and of the state of the universe, and this factual knowledge can be used as input for model building, but it is not something that quantum theory tells us about or explains. The demand that quantum theory, suitably interpreted, predict the classicality of the universe is a realist demand that we have no reason to fulfill (but of course, after having noticed that quasi-classical frameworks are more commonly instantiated in the universe, we are free to use them when building cosmological models).

If quantum mechanics does not predict the quasi-classicality of the universe, then classical mechanics does not strictly reduce to quantum mechanics. It gives us information that quantum mechanics does not, namely that some situations, the quasi-classical ones, are much more common and easy to access from our epistemic vantage point than others. We can live with this non-reducibility. Another way to say this is that quantum mechanics (just like relativity theory) is applicable to a wider range of potential contexts than classical mechanics, and so, in retrospective, classical mechanics appears to be informative about which kinds of contexts are more relevant or accessible to us: the ones involving classical properties in flat space-time regions.

4. What can be known about "worldly" contexts?

In sum, the metaphysical interpretation I propose is the following one: the world is populated by *situations*, each characterised by a coarse-grained event space (its context) and a modal structure of events. Situations presumably have mereological as well as coarse-graining relations. Experimental situations are only a special case of situations: the ones to which we have direct empirical access. Our understanding of what situations are and what kinds of properties they instantiate ultimately rests on our knowledge of experimental situations, so they are epistemically privileged, but ontologically on a par with other situations. Once a situation is identified, hypothetised or intentionally created, quantum theory can be used to know its modal structure. But the theory does not tell us what situations there are in the world. Only direct experience does.

Quantum theory is not entirely silent on this, though. The consistency condition and the single framework rule imply that some situations are impossible, and that if one context is instantiated, incompatible contexts on the same object cannot be instantiated as well, otherwise we would run into contradictions in probability calculus (this is where quantum theories departs from classical theories, where we could simply assume an ideally refined context). This is informative.

The single framework rule and the consistency condition have a counterpart in experience, which is the impossibility of implementing an experiment that would measure non-commuting observables on the same object at the same time, for example the position and velocity of the same particle with unlimited precision.[12] It is quite standard to analyse impossibilities of this kind by providing a model that incorporates the measuring apparatus: then we can explain, from the theory, why these measurements are incompatible.

Following this approach, one could hope to provide a reductionist analysis of our notion of context from standard quantum-theoretical models. Contexts would no more be extra-theoretical inputs required for representation, but rather represented objects, or an aspect of object–environment or object–instrument relations when considering more encompassing models. Typically, the bi-orthonormal decomposition of a wave-function can be used to tell us "what is measured" on each part of a system by the other parts (Bene & Dieks, 2002), and the theory of decoherence can tell us what context is induced by an environment on an object. The problem with this approach is that decoherence or biorthonormal decomposition assume an object–environment cut or a particular decomposition of the represented system into sub-parts, which only a larger context can provide.[13] So, our more encompassing model must also have a framework, or something equivalent, and if we want this larger context to be determined by an even larger target system, we run into an infinite regress (see Wallace, 2010 for an exposition of this problem in the case of decoherence). Therefore, it does not seem that a reductionist approach towards situations, based on standard abstract (context-free) quantum mechanical models, is viable. The alternative adopted here, which follows from a "meaning is use" approach, is simply to assume that situations are primitive objects that are directly known by experience.

A more subtle proposal is to refuse to reduce situations and contexts, but to consider it possible, assuming a situation to which we have direct access, to make ampliative inferences about which contexts and situations are instantiated outside of it. The kind of analysis that we get by considering composite systems constituted of a measuring apparatus and a measured object seems informative after all, and there is no reason that it could not be carried out using the interpreted models of Consistent Histories instead of the standard wave-function models, in order to infer that a particular situation is instantiated for an object by the instrument that measures it.

Assume, for example, an instrument that can measure the spin of an electron in two direction, X and Y, depending on its state, S_X or S_Y. The outcome of the measurement can be either x^+ or x^- if the spin is measured in direction X, and it can be either y^+ or y^- if the spin is measured in direction Y. Our context is given by the following consistent histories: $[S_X \odot x^+, S_X \odot x^-, S_Y \odot y^+, S_Y \odot y^-]$. I presume that we could infer from this representation that (1) all coarse-grainings of this context are

---

[11] If this is correct, the more general representation of the perspective of an agent should involve weak measurements, because experimenters do not constitute the full environment of their objects, but only part of it. We can postulate the existence of situations associated with strong measurements on our objects, but in general, we do not have full epistemic access to these situations.

[12] This impossibility is also there when considering more than one agent, each wanting to measure a different property: they are in conflict and at least one agent must fail (assuming they experiment on the object only, which excludes the idea that one agent would experiment on the composite system composed of the object and of the other agent, as in Wigner's friend thought experiment). This is one reason why frameworks should ultimately be associated with objects or situations, and not with particular cognitive agents. Another reason is that experimental contexts are usually implemented by *teams* of agents.

[13] This unless we fix a preferred context or object decomposition at the theory level, which is what Bohmian mechanics does by considering particle positions to be the only fundamental properties.

also instantiated, for example, $[S_X \odot I, S_Y \odot I]$, with I the identity projector, (2) a context associated with the instrument only is instantiated, namely $[S_X, S_Y]$ and (3) a context associated with the electron only is instantiated, it can be either $[x^+, x^-]$ or $[y^+, y^-]$, but *we do not know which* (or only probabilistically), since it depends on the state of the instrument.[14] These inferences follow merely intuitive rules as they stand, but the point is that they are carried out from our contextual representation alone, without considering wave-functions, biorthonormal decomposition and the like (although a form of decoherence and object decomposition is implicit in the consistency condition of the original framework).

I suspect that such inferences from one's own situation to external contexts is implicit in scientific practice. Strictly speaking, what an experimenter observes is the pointer of a measuring instrument, and yet, the inference that is made is that the measured object has a particular microscopic property. So, inferences from a broad, macroscopic situation that a certain microscopic context is instantiated seem to be routine, even if not always strictly formalised in the theory.

As said before, such inferences seem to rest primarily on experimental practice and its norms. However, the compatibility of these norms with theoretical reconstructions of measurement situations is, arguably, a desirable feature. The theory of decoherence often plays a central role in these reconstructions in quantum mechanics. In the present interpretation, this is because decoherence, which takes the form of a consistency condition built in the definition of a framework, tells us about which contexts are possible in the world and which are not.[15] This puts constraints on our inferences regarding which microscopic situation can be instantiated given a macroscopic situation that we experience more directly. This can help us make sense of the informal inferences found in experimental practice, but this is not a complete reduction, because a macroscopic situation, including an event space, must first be given and translated in theoretical terms for the theoretical reconstruction to have any import.[16]

These inferences from context to context concern concrete experimental situations. What can we say about the cases where scientists model systems on which we cannot experiment directly, such as the core of neutron stars or the early universe?

If these models are formulated in quantum theory without any framework specification, they cannot inform us about real situations. At most they support conditional claims, such as: if such context is instantiated in the neutron star, then such modal structure is present (which is still somehow informative). However, if these models implicitly incorporate object–environment cuts or decompositions into parts, and if they make use of the theory of decoherence, then they might be interpreted as incorporating claims about the kind of contexts and situations that are instantiated in these remote parts of the universe. These claims could be supported by inferences from our observations of these objects that follow the schema just described in the case of experimental situations. We could also wonder if an induction on the worldly situations encountered so far is implicitly involved here or in model building in general, in cosmology or elsewhere. This is the kind of inference that supports our previous reasoning about the classicality of the universe, and we could wonder to what extent, and in which conditions, it is warranted.[17]

All this to say that this approach is not purely instrumentalist: experimental contexts have a privileged epistemological and interpretive role, but quantum mechanics is potentially informative with respect to inferences that go beyond the experimental contexts to which we are directly acquainted.

5. Conclusion

In this article, I have proposed to interpret quantum mechanics in terms of situations, which are modal structures of events relativised to a context. This interpretation comes with an epistemic limitation: we can only identify the relevant situations by direct experience, and then the theory informs us about their modal structures. However, this limitation can be partly overcome if we assume the validity of probabilistic inferences from context to context under the constraints of quantum theory. I am convinced that analysing the imports of quantum theory with regards to this kind of ampliative inference, taking into account scientific practice (both experimental inferences and model building practices) rather than pure theoretical reasoning, is a research program worse pursuing.

The situations and contexts that we are postulating here are indispensable in order to apply quantum mechanics. They are required in order to make sense of our experience and of the common-sense intuition that objects have well-defined properties outside of our experience. So, in a sense, they should be acceptable by almost anyone and constitute a neutral ground from which interpretive questions can be asked. In this respect, standard realist interpretations of quantum mechanics, such as Bohmian Mechanics or the Many-Worlds Interpretation, merely assume that these contexts and situations are grounded in something more fundamental and acontextual. The proposal of this article, which is motivated by a "meaning is use" approach to theory interpretation, is to take them to be primitive entities instead. This allows to take all the benefits of the Consistent Histories formulation without facing the same limitations.

Although there is no space for a full analysis, it should be noted that there are similarities between the present approach and the relational interpretation (Rovelli, 1996; Smerlak & Rovelli, 2007), modal interpretations (van Fraassen, 1991), including perspectival ones (Bene & Dieks, 2002; Dieks, 2022) and QBism (Fuchs et al., 2014). There are also important differences. For example, although similar, QBism does not interpret quantum probabilities as objective, and it does not consider cross-contextual ampliative inferences. Modal interpretations postulate continuous states instead of discrete events. The relational interpretation takes events to be relative to other systems in interaction. The most closely related program might be Healey's pragmatist interpretation (Healey, 2012; 2022). As far as I am aware, this program does not make an explicit link to representation, direct experience, epistemology and ampliative cross-contextual inferences, although this proves fruitful.

The proposal of this article is originally motivated by considerations that have to do with scientific representation in general (notably with the

---

[14] The rules involved in (2) and (3) are different: in (2), we merely factorise and shrink the context obtained in (1), which would give us $[I]$ if applied to the electron, whereas in (3), we infer possible fine-grainings of $[I]$ given the original situation. This seems warranted if we want to say that the instrument induces a local context on the electron (and we certainly want to say this if quantum mechanics is to be applicable). Also note that neither the context $[I \odot x^+, I \odot x^-]$ nor the context $[I \odot y^+, I \odot y^-]$ can be instantiated, because they are incompatible with the original context, so they should not be conflated with the contexts restricted to the electron: a context only makes sense once a bounded system has been specified, considering the environment to be fixed (to say it differently, the parthood relation between situations is not a mere coarse-graining).

[15] This view is similar to Healey (2012)'s, according to whom decoherence warrants the ascription of magnitudes to physical quantities.

[16] More generally, according to the present theory of representation, any theoretical reconstruction of an experimental practice must also be interpreted in terms of potential contexts of application (since it is a function from context to concrete representation), so attempting to reduce all experimental practice to theory is a non-starter. This is a direct consequence of the "meaning is use" approach.

[17] We could assume the existence of a cosmic situation from which all local situations and the classicality of the world could be inferred. However, this cosmic situation is radically underdetermined by our experience, and local contexts can only be determined probabilistically from it (in the best case), so it is not a very useful concept. But perhaps we can determine coarse-grainings of this cosmological situation from our observations, which seems to correspond more closely to what cosmological models do.

idea that acontextual models are not representational), and not by the interpretive problems of quantum mechanics. But applying this theory of representation to quantum mechanics apparently resolves many of the difficulties associated with its interpretation, which constitutes a further motivation for it.

Acknowledgment

I would like to thank Raoni Arroyo for inviting me to present these views for the first time at the VIII International Workshop on Quantum Mechanics and Quantum Information.

This work was supported by the Marie Skłodowska-Curie Actions of the European Commission Individual Fellowships, Grant No. 101022338.

References

Griffiths, R. (2003). *Consistent quantum theory*. Cambridge University Press.

Belot, G. (2012). Quantum states for primitive ontologists. *European Journal for Philosophy of Science, 2*(1), 67–83.

Bene, G., & Dieks, D. (2002). A perspectival version of the modal interpretation of quantum mechanics and the origin of macroscopic behavior. *Foundations of Physics, 32*(5), 645–671.

Boesch, B. (2017). There *is* a special problem of scientific representation. *Philosophy of Science, 84*(5), 970–981. https://doi.org/10.1086/693989

Cartwright, N. (1983). *How the laws of physics lie, 34*. Oxford University Press.

Dieks, D. (2022). Perspectival quantum realism. *Foundations of Physics, 52*(4), 95. https://doi.org/10.1007/s10701-022-00611-8

Dowker, F., & Kent, A. (1996). On the consistent histories approach to quantum mechanics. *Journal of Statistical Physics, 82*(5–6), 1575–1646. https://doi.org/10.1007/BF02183396

Fuchs, C. A., David Mermin, N., & Schack, R. (2014). An introduction to QBism with an application to the locality of quantum mechanics. *American Journal of Physics, 82*(8), 749–754. https://doi.org/10.1119/1.4874855

Gell-Mann, M., & Hartle, J. B. (1993). Classical equations for quantum systems. *Physical Review D, 47*(8), 3345–3382. https://doi.org/10.1103/PhysRevD.47.3345

Giere, R. (2010). An agent-based conception of models and scientific representation. *Synthese, 172*(2), 269–281. https://doi.org/10.1007/s11229-009-9506-z

Griffiths, R. (2002). Consistent resolution of some relativistic quantum paradoxes. *Physical Review A, 66*(6), Article 062101. https://doi.org/10.1103/PhysRevA.66.062101

Griffiths, R. (2013). A consistent quantum ontology. *Studies in History and Philosophy of Science Part B, 44*(2), 93–114.

Hacking, I. (1983). *Representing and intervening: Introductory topics in the philosophy of natural science, 22*. Cambridge University Press.

Hartle, J. B. (2011). The quasiclassical realms of this quantum universe. *Foundations of Physics, 41*(6), 982–1006. https://doi.org/10.1007/s10701-010-9460-0

Healey, R. (2012). Quantum theory: A pragmatist approach. *The British Journal for the Philosophy of Science, 63*(4), 729–771.

Healey, R. (2022). Securing the objectivity of relative facts in the quantum world. *Foundations of Physics, 52*(4), 88. https://doi.org/10.1007/s10701-022-00603-8

Hughes, R. (1997). Models and representation. *Philosophy of Science, 64*, S325–S336. https://doi.org/10.1086/392611

Kent, A. (1997). Consistent sets yield contrary inferences in quantum theory. *Physical Review Letters, 78*(15), 2874–2877. https://doi.org/10.1103/PhysRevLett.78.2874

Kitcher, P. (1989). Explanatory unification and the causal structure of the world. In Philip Kitcher, & Wesley Salmon (Eds.)*, 8*. *Scientific explanation* (pp. 410–505). Minneapolis: University of Minnesota Press.

Morgan, M., & Morrison, M. (Eds.). (1999). *Models as mediators: Perspectives on natural and social science*. Cambridge: Cambridge University Press. https://doi.org/10.1017/CBO9780511660108.

Okon, E., & Sudarsky, D. (2014a). On the consistency of the consistent histories approach to quantum mechanics. *Foundations of Physics, 44*(1), 19–33.

Omnès, R. (1995). A new interpretation of quantum mechanics and its consequences in epistemology. *Foundations of Physics, 25*(4), 605–629.

Omnes, R. (1999). *Understanding quantum mechanics*. Princeton, N.J: Princeton University Press.

Pearl, J. (1988). *Probabilistic Reasoning in intelligent systems: Networks of plausible inference. The morgan Kaufmann Series in representation and reasoning*. San Mateo, Calif: Morgan Kaufmann Publishers.

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

Ruyant, Q. (2021a). *Modal empiricism: Interpreting science without scientific realism, 440*. Cham: Synthese Library, Springer Nature. https://doi.org/10.1007/978-3-030-72349-1

Smerlak, M., & Rovelli, C. (2007). Relational EPR. *Foundations of Physics, 37*(3), 427–445.

Suppes, P. (1969). Models of data. In Patrick Suppes (Ed.)*, Vols. 24–35*. *Studies in the methodology and foundations of science*. Dordrecht: Springer Netherlands. https://doi.org/10.1007/978-94-017-3173-7_2.

van Fraassen, B. (1980). *The scientific image*. Oxford: Oxford University Press. https://doi.org/10.1093/0198244274.001.0001

van Fraassen, B. (1991). *Quantum mechanics: An empiricist view*. Oxford University Press.

Wallace, D. (2010). Decoherence and ontology, or: How I learned to stop worrying and love FAPP. In Simon Saunders, Jonathan Barrett, Adrian Kent, and David Wallace (Ed.), *Many worlds? Everett, quantum theory, and reality*. Oxford University Press.

Williamson, T. (2020). *Suppose and tell: The Semantics and Heuristics of conditionals* (1st ed.). New York, NY: Oxford University Press.

Woodward, J. (2004). *Making Things happen: A theory of causal explanation*. Oxford: Oxford University Press. https://doi.org/10.1093/0195155270.001.0001

Okon, E., & Sudarsky, D. (2014b). Measurements according to consistent histories. *Studies In History and Philosophy of Science Part B: Studies In History and Philosophy of Modern Physics, 48*(November), 7–12. https://doi.org/10.1016/j.shpsb.2014.08.011

Okon, E., & Sudarsky, D. (2021b). Symmetries, indexicality and the perspectivist stance. *International Studies in the Philosophy of Science*. https://doi.org/10.1080/02698595.2021.2001723. November, 1–19.