

Knowledge of the Quantum Domain: An Overlap Strategy

James D. Fraser and Peter Vickers

Abstract

The existence of multiple interpretations of quantum mechanics appears to pose a serious challenge for knowledge claims about the quantum domain. Hoefer (2020) argues that a scientific realist epistemology must be abandoned in this context, while Callender (2020) argues that the realist's only option is to break the underdetermination between rival interpretations by appealing to extra-empirical virtues. We develop a different response to the quantum underdetermination problem based on identifying statements about the unobservable which all the major ontic interpretations of quantum mechanics agree on. It is commonly believed that Everettian, Bohmian and GRW quantum mechanics share nothing but empirical content. We argue that, while they say very different things about the fundamental nature of quantum systems, they can be understood as agreeing on a plethora of more abstract theoretical claims. In our view, focusing on this descriptive overlap represents the most promising strategy for defending knowledge claims in the quantum domain. We close by considering how this overlap strategy relates to working posits formulations of scientific realism.

Table of Contents

1 Introduction

2 Three Responses to the Quantum Underdetermination Problem

2.1 Response 1: Abandon Knowledge Claims About the Quantum Domain

2.2 Response 2: Break (or Dissolve) the Underdetermination

2.3 Response 3: Identify Overlap Between the Underdetermined Interpretations

3 The Overlap Strategy

3.1 Opening Gambit: A Hydrogen Atom Test Case

3.2 General Strategy: Against Textbook Quantum Mechanics

3.3 Next Steps: The Wavefunction, Decoherence and an Interpretation-Neutral Notion of Possibility

3.4 Options for Ontology and Semantics

4 The Overlap Strategy and Working Posits Realism

4.1 Unconceived Alternatives?

4.2 Explanatory Deficit?

1 Introduction

Quantum mechanics appears to have all the qualities which scientific realists typically take to warrant knowledge claims about unobservable aspects of the world: novel predictive success, theoretical maturity, explanatory power, and so on. Saying what we can really claim to know about the quantum domain has long been a project fraught with difficulty, however.

Historically, quantum mechanics was often taken to be incompatible with a realist semantics of theories. The Copenhagen interpretation, at least on some popular readings, denied that quantum mechanics can be understood as a description of the way the sub-atomic world is. In the contemporary debate, a variety of epistemic and pragmatist interpretations continue to push this kind of view of the theory (see Healey (2016) for a review of these approaches).¹ Over the past 50 years, however, a different challenge has arisen. There are now too many ways of understanding quantum mechanics as a descriptive theory. The most prominent so-called ontic interpretations of quantum mechanics – the Everett or many-worlds interpretation, Bohmian hidden variable interpretations and GRW spontaneous collapse interpretations – give radically different accounts of what the world is fundamentally like. In some the wavefunction is everything, in others, it is only part of the physical description of the quantum realm; in some, the basic laws are deterministic, in others indeterministic; and so on.

We seem, therefore, to be faced with a severe case of underdetermination of theory by evidence in the foundations of quantum theory. Different paths have been taken here in the recent literature. Hoefer (2020) argues that we must abandon knowledge claims in the quantum domain; the scientific realist should fall back to sciences like chemistry and biology where claims to bona fide scientific knowledge are more secure. Conversely, Callender (2020) argues that the scientific realist must commit herself to a particular interpretation, breaking the underdetermination by appealing to extra-empirical virtues.

A key premise in both Hoefer's and Callender's arguments is that the scientific realist cannot remain neutral between the various descriptive interpretations of quantum theory. This paper argues that their dismissal of this response is too quick and that this is in fact the most promising option for the realist. We develop a strategy for articulating modest knowledge claims in the

¹ We will be putting epistemic interpretations of quantum theory to one side from this point on. The nominal justification for this is that we are interested in whether a scientific realist epistemology can be maintained in the quantum domain. Since proponents of QBism (Fuchs 2017) and pragmatism (Healey 2020) have claimed that these approaches can be understood as compatible with forms of scientific realism, this may be more accurately characterised as a limitation of our discussion.

quantum domain based on identifying descriptive overlap between the various ontic interpretations. Contrary to what has sometimes been claimed in the literature, we argue that the overlap between Everett, Bohm and GRW quantum mechanics is not purely empirical and that it is possible to find non-trivial theoretical claims endorsed by all three interpretations.

The paper is structured as follows. Section 2 discusses Hofer's and Callender's responses to the underdetermination problem in the foundations of quantum theory and motivates the need for another option. Section 3 develops a programme for identifying descriptive overlap between the various ontic interpretations and defends its viability. Section 4 concludes by considering an apparent conflict between this overlap strategy and the more familiar selective realist notion of working posits. Two more general concerns arise from this discussion: that the overlap strategy falls prey to an unconceived alternatives problem and that it leads to a problematic weakening of the explanatory component of scientific realism. We leave these issues as fodder for future investigation and debate.

2 Three Responses to the Quantum Underdetermination Problem

In understanding the epistemic problem posed by the existence of multiple ontic interpretations of quantum mechanics as a case of underdetermination of theory by evidence we are following a growing body of philosophical literature (Saatsi 2019; Callender 2020; Hofer 2020; Acuña 2021; Egg 2021). Still, some clarifications are in order.

By ontic interpretations we mean, roughly, proposals that purport to, (i) solve the measurement problem, and (ii) provide an account of what the quantum domain is like. For the sake of concreteness, we shall use generic versions of the Everett,² Bohmian and GRW interpretations as reference points throughout this paper, in the hope that most of what we say can be generalised without drastic revision when the full space of viable ontic interpretations is taken into account.³

² By default, we shall have in mind the modern decoherence-based formulation of the Everett interpretation of the kind defended in Wallace (2012).

³ Everett, Bohm and GRW are really families of interpretations, since there are variants of each approach. Indeed, when one includes accounts of the fundamental ontologies posited by these interpretations considered in the recent metaphysics of science literature there appears to be a large amount of underdetermination within each family. Section 3.4 brings these more fine-grained metaphysical readings into the discussion.

It is sometimes suggested that, since these approaches differ at the level of mathematical formalism, the term “interpretation” is misleading here, and it would be more accurate to classify them as rival theories (Myrvold 2016). Nothing in our discussion will hang on this distinction and we will continue to use the term “interpretation” broadly. A more significant point, which marks a departure from the underdetermination scenarios traditionally focused on in the philosophy of science literature, is that they are not all empirically equivalent. GRW spontaneous collapse models, in particular, introduce new parameters characterising the rate and width of collapse events that can, in principle, be measured. As Stanford (2006) persuasively argues, however, if we are interested in the question of what we know right now, strict empirical equivalence is something of a red herring. Underdetermination with respect to available evidence is sufficient to challenge the present justification of scientific claims about the unobservable. The problem, in a nutshell, is this: if the Everettian, Bohmian and GRW interpretations are all consistent with extant empirical data how can we justifiably believe what one says about the quantum domain over the others?

There is a second respect in which our case differs from traditional underdetermination scenarios: it is only in the rather special case of non-relativistic quantum mechanics that the underdetermination can be made fully explicit. The philosophical and foundational literature has focused on this particular quantum theory and it is in this context that Everettian, Bohmian and GRW approaches have been worked out in detail. When it comes to other quantum theories, especially relativistic quantum field theories, the situation is rather different. Whether hidden variables or spontaneous collapse versions of empirically supported quantum field theories exist is currently an open question. We will discuss this issue further below, but our basic position again follows Stanford (2006). Contrary to what has sometimes been argued, underdetermination does not have to be made fully concrete to have epistemic weight. The existence of research programmes developing relativistic Bohmian and GRW models is enough to pose a significant challenge for knowledge claims in quantum field theory. Methodologically, we think it makes sense to start with the special case of underdetermination between Everettian, Bohmian and GRW versions of quantum mechanics and consider how possible responses can be generalised from there.

Granting that the existence of multiple ontic interpretations poses an underdetermination problem – which we will call the quantum underdetermination problem for short – the question

is, how should we respond? We consider – and criticise – two recently advanced answers in order to motivate a third response, which will form the basis of our positive proposal in section 3.

2.1 Response 1: Abandon Knowledge Claims About the Quantum Domain

The conclusion that the scientific anti-realist wants to draw from underdetermination is that a scientific theory's claims about the unobservable are not epistemically justified. Hofer (2020) argues that this is the right response to the quantum underdetermination problem but claims that this does not mean that we must adopt an anti-realist epistemology in other scientific contexts where the problem of underdetermination does not arise with the same force. To use Callender (2020)'s terminology, Hofer suggests that the underdetermination problem facing quantum mechanics can be quarantined – we can continue to hold a positive epistemic attitude towards claims about the unobservable made in chemistry and virology, say, while giving up on knowledge claims about the quantum domain.

We have two things to say about this quarantine response. Firstly, we have doubts about whether it can be implemented in a principled way. If a no-miracles style predictive-success-to-truth inference fails in the case of quantum physics, how can it be trusted elsewhere? Hofer is well-aware of this issue and responds by rejecting the traditional no-miracles argument for scientific realism. In its place, he proposes a new way of motivating scientific realist commitments based on the existence of diverse and interconnected “epistemic handles” on unobservable entities (see also Hofer and Martí 2020). This way of motivating realism requires further development before its promise can be fully assessed, but we suspect that the problem of giving a principled reason for weakening one's commitments in the context of quantum physics alone will simply recur. Many of the features Hofer associates with epistemic handles – “direct or indirect observation; theoretical reasons for believing that an entity of such a type should exist; knowing how to produce or manipulate the entity; taking the entity to be the cause of a certain observable event or phenomenon” (Hofer, 2020, 32) - appear to be amply present, and highly diverse and interconnected, in quantum physics.⁴

⁴ Another way to pose this worry about Hofer's response is to ask how we know that troubling underdetermination is not present with respect to unobservable chemical and biological entities when we know that it is for quantum entities. “Epistemic handles” cannot be the answer if they are amply present for both quantum and non-quantum scientific entities.

Secondly, Hoefer explicitly styles his approach as a solution of last resort – he endorses it because he does not think that any of the other possible responses to the quantum underdetermination problem are credible. In particular, he argues that one cannot maintain knowledge claims in the quantum domain while remaining neutral between the various ontic interpretations. Our main goal in this paper is to chip away at this antecedent claim, calling the motivation for Hoefer’s response into question. As the previous paragraph suggests, Hoefer’s preferred route should not be taken lightly as it requires one to carefully calibrate one’s general epistemology of science so as to exclude the quantum domain from the scope of scientific knowledge in a non-ad-hoc way. It seems prudent to thoroughly explore the possibility of defending substantive knowledge claims in the quantum domain before resorting to the nuclear option of abandoning them entirely.⁵

2.2 Response 2: Break (or Dissolve) the Underdetermination

A common general response to the problem of underdetermination of theory by evidence in the philosophy of science is to suggest that extra-empirical virtues like simplicity and explanatory power can be used to break evidential ties. Appeals to extra-empirical virtues are sometimes used to justify committing oneself to a particular interpretation, especially in the more metaphysical wing of the literature (for instance, Allori 2020; Esfeld 2014). Callender (2020) argues that this is the only option for the scientific realist: the realist must commit herself to a particular interpretation of quantum mechanics and justify her choice via extra-empirical virtues.

We are sceptical about this response to the quantum underdetermination problem. One difficulty with evaluating this proposal is the lack of agreement amongst philosophers about the justificatory basis of non-empirical theory assessment – which virtues have evidential power, how they should be weighed against each other, and so on. Even granting that extra-empirical virtues do sometimes give us rational reasons to increase our credence in a theory, and suppressing questions about the details, we think it is implausible that they carry enough weight in this particular case to support scientific knowledge claims. Note that merely breaking the evidential tie is not enough here. If, hypothetically, extra-empirical virtues were to tip the scales in favour of Bohmian mechanics, so that we ought rationally to assign a credence of 0.5 to this

⁵ Hoefer ultimately pulls back from striking all claims about the quantum domain from the canon of established scientific knowledge and admits that some minimal claims about quantum entities, such as “electrons exist and have rest mass m_e ” (Hoefer and Martí 2020) are good candidates for realist commitment. We suspect that Hoefer’s intuition that statements of this sort can be trusted is underwritten by the assumption that they will come out as true on any reasonable interpretation of quantum mechanics and is therefore implicitly reliant on something like the overlap strategy.

interpretation, and 0.25 to Everett and GRW, this would not be enough to give us knowledge.⁶ We find it implausible that extra-empirical virtues justify assigning an extremely high credence to a particular interpretation, as would be required to rescue knowledge claims about the quantum domain in this way.

One way to support this judgement is our subjective assessment of the extra-empirical virtues of the various interpretations. Arguments can be levied in support of the theoretical virtuousness of all three major ontic interpretations, none of which seem to us to be anywhere near compelling enough to underwrite knowledge claims. Another way of supporting our pessimistic take on the justificatory power of extra-empirical virtues in this case is a community level argument. The Everett, Bohm, and GRW interpretations of quantum mechanics have been rivals for decades. During this time, a huge number of scientists and philosophers have given conflicting evaluations of which is most theoretically virtuous, and no interpretation has won overwhelming support in the scientific community – not even close. In this context, it seems cavalier to fully commit oneself to a particular interpretation on the basis of one’s personal assessment of its theoretical virtues. This is not the sort of safe, modest, theoretical claim the contemporary scientific realist wants to endorse.

There is another style of response to putative instances of underdetermination: argue that the underdetermination is merely apparent. It is not uncommon for advocates of a particular interpretation of quantum mechanics to claim, not just that rival approaches are theoretically unvirtuous, but that they do not solve the measurement problem at all. In the case of the Everett interpretation, for instance, opponents have argued that it is conceptually incoherent or fails to adequately recover the empirical predictions of the Born rule (Albert 2010; Dawid and Thébault 2015). Arguing the other way, Everettians sometimes suggest that the underdetermination problem can be dissolved in their favour. As we mentioned above, hidden variable and spontaneous collapse interpretations modify the formalism of conventional non-relativistic quantum mechanics making the project of extending them into the relativistic context highly non-trivial. Crucially, it is currently unclear whether Bohmian and GRW versions of the standard model of particle physics exist. Since the Everett interpretation does not modify the conventional formalism, this approach can be straightforwardly extended beyond the confines of non-relativistic quantum mechanics, or so its advocates claim. Thus it might be argued that the huge

⁶ On this point, note that the overlap strategy developed in section 3 is not based on the idea that one must be strictly neutral between all ontic interpretations, in the sense of assigning equal credence to each. A fairly strong advocate of a particular interpretation who nevertheless admits that they do not know that it is true can make use of our proposal.

amount of empirical evidence for the standard model of particle physics supports the Everett interpretation, as it is the only interpretative framework that can (currently) recover these predictions.

We have more sympathy with these lines of argument than with appeals to extra-empirical virtues. It is certainly true that there is some uncertainty about how much underdetermination there really is in the foundations of quantum theory. There is scope to whittle down the space of viable ontic interpretations, either through new experiments (which could, for instance, rule out GRW models) or through new theoretical results. If no-go theorems could be proven which showed that Bohmian and GRW formulations of the standard model do not exist, for instance, this would be a major advance with important ramifications for the present discussion. In any case, it is more plausible that the underdetermination issue might someday be resolved in this way than that a scientific consensus will form around the superior non-empirical theoretical virtues of any currently available interpretation. Crucially though, we are interested in whether we have knowledge of the quantum domain now, not in some hypothetical future scenario that might never materialise.

It is sometimes suggested that the difficulty with generalising Bohmian and GRW interpretations to empirically successful quantum field theories stems from an inevitable clash between these approaches and relativity theory. These arguments are heuristic, however, and are called into question by the existence of toy Bohmian and GRW type field theories. Even if it is reasonable to doubt that the Bohmian and GRW approaches will ever be able to recover all of the empirical predictions made by the full range of quantum theories, this will again not be sufficient to underwrite a strong enough commitment to the Everett interpretation unless the credence we should assign to this possibility is extremely small. Remember that good reasons to lean towards a particular interpretation are not sufficient to underwrite knowledge claims. We can rerun our community level argument here: given the widespread disagreement amongst the relevant experts about the prospects of recovering the predictions of relativistic quantum theories within the Bohmian and GRW frameworks, it would be cavalier to dismiss this possibility entirely.

In sum, we consider plumping for a particular ontic interpretation, on whatever grounds, to be a poor option for someone wishing to defend knowledge claims about the quantum domain. Doing so is simply too epistemically risky given the theoretical and empirical information that is currently available. The scientific realist wants to make safe commitments that can be trusted to stand the

test of time. There is thus a strong motivation for looking at a third response to the quantum underdetermination problem.

2.3 Response 3: Identify Overlap Between the Underdetermined Interpretations

A final response to underdetermination, which has also been discussed in the general philosophy of science literature to some extent (for instance, French 2011), is to look for common claims made by all of the equally supported theories. The justification for these common claims is not called into question by the underdetermination, it is only the points on which the theories disagree that are threatened by it. Adopting this sort of response to the quantum underdetermination problem would mean putting one's trust in the claims about the unobservable which the various ontic interpretations agree on – their descriptive overlap – while remaining uncommitted on the points on which they disagree. Variations of this response have, in fact, been advanced by a number of authors (Belousek 2005, Cordero 2001, Ney 2012, Egg 2021).⁷

It certainly isn't obvious that any meaningful overlap is there to be found, however. Callender (2020) gives a clear statement of what may represent the orthodox take on this question amongst philosophers of quantum physics:

Most of what we say about the quantum realm is 'interpretation' dependent. The research programs described here portray radically different worlds from top to bottom, agreeing on little more than what is observable. (p. 72)

The claim that searching for descriptive overlap between rival ontic interpretations is a dead-end for the scientific realist drives both Callender's and Hoefer's responses to the quantum underdetermination problem. Given the objections to their proposed solutions raised above, there is a strong motivation to carefully reconsider the viability of this third response, however. In the following section, we combat the worry that the descriptive overlap between Everettian, Bohmian, and GRW quantum mechanics is empty by developing an approach to identifying non-trivial, interpretation-neutral theoretical claims. Rather than a dead-end, this turns out to be the most promising route to defending knowledge claims about the quantum domain.

⁷ Saatsi (2019) also suggests that a form of scientific realism can be maintained while remaining neutral between competing ontic interpretations. However, he gives up on the knowledge claim component of traditional scientific realism, which is our primary interest in this paper.

3 The Overlap Strategy

This section lays groundwork for this third response to the quantum underdetermination problem. It is structured as follows. We start, in 3.1, by examining a simple test case of a theoretical claim that all three interpretations seem to endorse. Our dialectical strategy here is to establish proof of principle by showing that there is at least some descriptive overlap between our three interpretations. 3.2 discusses how to go about delimiting the overlap in general terms. We endorse a case-by-case approach and criticise the idea, suggested in the recent literature, that quantum mechanics textbooks might be used to articulate the overlap. 3.3 sketches a more ambitious deployment of the overlap strategy to claims about the wavefunction (or quantum state), drawing on work on interpretation-neutrality in decoherence theory. 3.4 concludes by providing greater detail on the semantic and meta-ontological positions an overlap theorist might adopt, including Egg's "effective ontology" approach.

3.1 Opening Gambit: A Hydrogen Atom Test Case

Let us start our discussion by considering a simple test case:

The electron in a hydrogen atom transitions between discrete energy states. ()*

The thought is that, while the three interpretations do disagree in various ways about what hydrogen atoms are, and how they behave, they do at least agree on this much. Furthermore, this is not (see Section 3.4) simply a statement about the observed emission spectrum of Hydrogen; it is a substantive claim about the dynamical behaviour of an electron-proton bound state.

A worry here, however, is that, in their detailed accounts of what is going on in a hydrogen atom, our three interpretations of quantum mechanics unpack the key terms "electron", "transition" and "state" very differently. It might be thought that, unless one commits to one of the precise meanings of these terms supplied by a particular interpretation, it is not clear what one is being realist about when one asserts (*). Our goal in this section will be to make an initial case that these terms can be used in a broad sense, which straddles all three interpretations, while still saying something non-trivial about unobservable aspects of the physical world, so (*) really is an example of descriptive overlap between the Everett, Bohm and GRW interpretations. (Note that the *full* case for this conclusion includes the more detailed discussion of ontology and semantics, in Section 3.4.)

We want to understand this as an instance of a general linguistic tactic in which one abstracts away from details one is unsure about to arrive at a statement that is less specific but more likely to be true. In everyday life as well as the sciences we are constantly using terms in a broad sense – compatible with many possible precisifications – in order to make epistemically safer claims. If one does not know whether Alice drove or took the train to Durham, for instance, one might simply say “Alice travelled to Durham”, using “travelled” as an umbrella term that admits driving and getting the train as specific instances, in order to ensure we say something true. Similarly, scientists with conflicting commitments and hunches are happy to assent to the statement “A bolide impact triggered the extinction of non-avian dinosaurs”. Here the term “bolide impact” allows us to hedge our bets about the nature of the object which impacted Earth and the term “triggered” allows one to remain neutral between various competing theories about the detailed causal story leading to the extinction event.

As well as replacing a more specific term like “drove” with a more permissive one like “travelled” we can also use a *single* word in a broader and narrower sense to achieve a similar result. An interesting case to consider here is the term “gene”. This term features in many statements which are considered by the relevant scientific communities to be clear cases of secure knowledge; for instance, “*the SRY gene is required for the development of mammalian testes*” (Berta et al. 1990; see also Vickers 2022, Chapter 1). However, debate has raged for many decades about how the term “gene” should be defined, with no consensus being reached (Falk 2010). Scientists were still able to advance substantive claims in the domain of genetics without committing themselves to one of the competing precise definitions of the concept by using the term “gene” in a broad sense.

The question then is can we employ the same kind of abstraction manoeuvre with the terms “electron” “transition” and “state” in (*)? Is there any principled reason why we cannot use these terms in a broad sense to express something non-trivial about the quantum regime while remaining neutral on the questions which our three interpretations disagree on?

One might think that an important disanalogy between (*) and the examples of secure but non-specific knowledge mentioned above is that quantum physics, at least putatively, concerns the fundamental ontology of the physical world. There is a natural thought that when it comes to fundamental ontology one has to state with maximal precision what kind of entities and properties one is talking about. Here we need to make explicit that the overlap strategy as we understand it gives up on the idea that we can read the fundamental ontology of our world off from quantum mechanics. The Everett, Bohm and GRW interpretations disagree quite radically

on fundamental metaphysics; they disagree about whether there is one quasi-classical world or many, whether the fundamental laws are deterministic or indeterministic, and so on. All of this is clearly not going to be in the descriptive overlap and is therefore threatened by the quantum underdetermination problem.⁸ Admitting that we do not know what the quantum domain is fundamentally like does not mean that we don't know anything about it at all, however. The promise of the overlap strategy is based on the idea that there may be non-fundamental truths about quantum systems that are better candidate knowledge claims.

Arguably, necessary concessions to anti-realist arguments pull the scientific realist away from making claims about fundamental metaphysics in any case. Furthermore, as Egg (2021) points out non-relativistic quantum mechanics (and, in fact, quantum field theory) are not fundamental theories, so we also have context-specific reasons to distrust what these say about the fundamental. Egg develops an "effective ontology" approach to quantum mechanics, according to which we can understand interpretation-neutral claims as referring to non-fundamental entities and properties. As we discuss further in section 3.4, we consider this meta-ontological approach to be a non-compulsory extension of the basic overlap strategy. For present purposes, the key point is that (*) should not be read as a claim about the fundamental, and intuitions about fundamentality should therefore not be used to disqualify it as a candidate non-specific but safe knowledge claim.

Separately from these concerns about fundamentality, however, a critic might hold that the specific notions of "electron", "transition" and "state" which one finds in the three interpretations of the hydrogen atom are so radically different, perhaps even incommensurable, that we cannot enact the relevant kind of semantic abstraction (or if we try to we are left with no descriptive content at all). In our view, examining the details of the three rival accounts of the hydrogen atom goes a long way towards dispelling this worry, however.

All three of our interpretations employ the discrete spectrum of energy eigenstates obtained from an exact solution of the standard hydrogen atom Hamiltonian to account for its observed emission spectrum. There are two central points that they disagree about, however, which might, *prima facie*, be thought to problematize (*): i) they disagree about whether the wavefunction provides a complete or incomplete description of a quantum system, leading to differences in their

⁸ Note that this marks a departure from Ney (2012)'s otherwise closely related ideas about interpretation neutrality since she appears to be looking for fundamental metaphysical principles which are common to all ontic interpretations of quantum mechanics.

characterisation of the electron and its states, ii) they disagree about whether the Schrödinger equation is exceptionless or not potentially leading to differences in their description of atomic transitions.

Let us start with the disagreement over dynamics, as, in this special case, it does not appear to pose as much of a challenge for the overlap theorist. In general terms, by an atomic “transition” we mean a temporal evolution from one eigenstate of the hydrogen Hamiltonian to another. A perfectly isolated hydrogen atom will not, in fact, display this behaviour. However, if we couple it to a large number of environmental quantum degrees of freedom, for instance by allowing photon excitations of the electromagnetic field, the excited states of the hydrogen subsystem are no longer stationary. Rather, a hydrogen atom which is prepared in an excited state will evolve, via the Schrödinger equation, into a quantum state which has most of its mass in the hydrogen ground state. Since Everett and Bohmian quantum mechanics both take the Schrödinger equation to completely describe the dynamics of quantum systems they essentially agree about this story (though, as we discuss shortly, they give different physical interpretations of the end state of this evolution). Furthermore, GRW’s modifications of the dynamics have a relatively minor impact on the description of atomic transitions.⁹ In any case, the dynamical differences which do obtain between the GRW, Bohm and Everett hydrogen atoms do not seem to problematically stretch the standard broad meaning of “transition” as a time-evolution from one state into another. We can plausibly remain neutral about whether the evolution is a continuous or discontinuous processes, for instance, and still use the term “transition” to make a non-trivial descriptive claim about atomic systems.

What may appear more challenging for the overlap theorist is the difference between the way that the electron, and its states, are characterised in Bohmian mechanics as compared to the other interpretations. Alongside the standard wavefunction energy eigenstates of hydrogen the Bohmian account posits a set of periodic trajectories for the “hidden” position variable of the electron. Famously, the Bohmian trajectory associated with the hydrogen ground state is at rest at the Bohr radius from the proton (Holland 1995). This means that the Bohmian electron has additional properties which the Everett and GRW electron does not have. This, in itself, does not seem too problematic. The idea that we can maintain a realist attitude towards the atomic

⁹ One difference is that GRW “hit” events lead to spontaneous excitations from the ground state of a completely isolated hydrogen atom, which does not occur on the other interpretations. This has been held up as a potentially measurable signature of the presence of a spontaneous collapse mechanism (Squires 1990).

transitions posited by the old quantum theory without committing ourselves to the orbital trajectories posited by Bohr and his followers has been developed in detail in the literature (Norton 2000, Vickers 2020), and maintaining (*) while remaining agnostic about the existence of stable hidden variable trajectories seems to be a natural extension of these ideas. Perhaps more troubling, however, is the fact that the Bohmian electron also, in some sense, lacks core properties which are associated with electrons, such as charge and spin. These are not intrinsic properties of the Bohmian corpuscle but instead arise from its coupling to the wavefunction via the guiding equation. If the three interpretations disagree about the most basic properties of the electron it might seem doubtful that it makes sense to use the term “electron” in a broad sense which straddles all three accounts.

Consider how causal and descriptivist theories of reference treat this case, however. The causal account requires only that the Bohmian and Everettian electron concepts are appropriately causally related to an initial baptism event (presumably Thomson’s discovery of the electron) for them to be talking about the same thing. Meanwhile, sophisticated forms of descriptivism associate theoretical terms with a cluster of descriptive claims, not fundamental properties. Crucially, we do not think that a proponent of the Bohmian interpretation will want to deny that statements like “electrons have charge $-e$ ” and “electrons have spin $\frac{1}{2}$ ” are true (cf. Egg 2021). While they give a different account of the deep metaphysics which makes these statements true they still want to associate these properties with electrons and use them to individuate them from other particles. Arguably, therefore, the fact that these descriptive claims are true in a Bohmian possible world is sufficient to dispel the worry that the Bohmian electron concept is incommensurable with that posited by the other two interpretations. Consequently, we think that a standard textbook notion of the electron as “a sub-atomic particle with mass m_e , charge $-e$ and spin $\frac{1}{2}$ ” is already capable of supplying a non-trivial meaning to (*) which is common to all three interpretations. Contra the incommensurability worry raised above, then, we hold that it is possible to use the term “electron” in a broad sense while remaining neutral on the question of whether charge and spin are intrinsic properties, or extrinsically associated with the coupled system of a Bohmian corpuscle and its wavefunction.

If the reader is still not convinced by the above discussion we will return to the question of how we might understand the semantics of interpretation-neutral theoretical statements in 3.4. For the moment, we take (*) to indicate that the descriptive overlap between our three

interpretations is not in fact empty. The question now is where do we go from here. How do we go about characterising the full descriptive overlap of Everett, Bohm, and GRW quantum mechanics?

3.2 General Strategy: Against Textbook Quantum Mechanics

One answer which has been considered in the recent literature is to point to “textbook quantum mechanics”, the presentation of non-relativistic quantum mechanics found in physics textbooks, as an encapsulation of the putative descriptive overlap of the rival interpretations. Callender (2020) reads Cordero (2001) this way, and Egg (2021) explicitly takes up and defends this approach. We think this is a mistake for a number of reasons. Firstly, what counts as textbook quantum mechanics is vague owing to the diversity that exists within the textbook tradition. Secondly, the measurement problem arguably shows that attempts to precisify what is meant by textbook quantum mechanics turn out to be either inconsistent or incoherent, hence the need for an interpretation of quantum theory in the first place.¹⁰ Finally, textbooks are pedagogical tools not compendiums of reliable scientific knowledge; they often contain material that is instructive for grasping important concepts but does not (and is not intended to) provide an accurate representation of real-world phenomena.

Confusion about this last point, in particular, has sidetracked the recent discussions of an overlap response to the quantum underdetermination problem, in our view. Callender’s critique of this option focuses explicitly on the fact that textbook accounts of important quantum phenomena are directly contradicted by the corresponding Bohmian accounts. A standard problem in introductory quantum mechanics texts, for instance, is the treatment of quantum tunnelling through a potential barrier using stationary plane wave states. Callender points out that the transmission and reflection coefficients that are calculated in this way do not correspond to anything physically meaningful in the corresponding Bohmian system – indeed, there are no reflected Bohmian trajectories associated with these plane wave states. He sees this as one of a number of examples that tell against the viability of an overlap response to the quantum

¹⁰ We do not take the overlap strategy to be allied to attempts to dissolve the measurement problem on behalf of the scientific realist (Ladyman and Ross 2007). While the overlap theorist does not commit herself to one of the solutions of the measurement problem provided by particular ontic interpretations, the existence of these accounts seems to us to be necessary in order to assign meaningful physical content to quantum theories. This marks a difference with Egg (2021) who apparently takes textbook quantum mechanics to provide a functional characterization of the properties and behaviour of quantum systems without solving the measurement problem. This contentious claim strikes us as unnecessary in order to implement the overlap strategy.

underdetermination problem. But, while his discussion suggests that textbooks are not the right place to look for putative interpretation-neutral knowledge claims, it does not show that such claims cannot be found elsewhere.

The use of stationary unnormalized plane waves in the usual textbook account is physically problematic for reasons which have nothing to do with the subtleties of individual interpretations (as Callender actually remarks in a footnote). Prospects for finding interpretation-neutral claims about tunnelling phenomena look much better when we turn to more sophisticated approaches found in the physics literature, such as Norsen's (2013) treatment of tunnelling in terms of normalized wavepackets. Roughly speaking, what happens in this more realistic treatment is that a wavepacket incident on a finite potential barrier splits into two components, one of which is transmitted through the barrier and the other of which is reflected. As Norsen discusses in detail, in the Bohmian picture this corresponds to one subset of the initial conditions for the hidden variables leading to transmitted trajectories, and another subset leading to reflected trajectories. In the Everett and GRW approaches to quantum theory one simply has a transition to a superposition of a transmitted and reflected wavepacket: these interpretations then have their own stories to tell about how determinate particle positions are recovered from this superposition. Callender's suggestion that there is an irreconcilable conflict between the physical accounts of quantum tunnelling provided by the various ontic interpretations now evaporates and the stage appears to be set for identifying interpretation-neutral theoretical claims.

We take examples like these to problematise the textbook approach to characterising the descriptive overlap between ontic interpretations. Egg (2021) tries to patch up the textbook approach, suggesting that one can take a "selective" attitude towards the content of textbooks. But we are not really being selective when we *replace* a textbook account of a quantum phenomenon with a superior account found elsewhere. Furthermore, as we have already suggested, the motivation for starting from textbooks was flimsy to begin with. There may be an intuition underlying this move that if one is not a scientific realist about Everettian, Bohmian, or GRW quantum mechanics then one must point to some fourth theory which one is going to be realist about, at which point "textbook quantum mechanics" enters the stage. This line of thought should simply be rejected. Rather than characterising the overlap all at once, as it were, we advise taking a case-by-case approach, looking at the best-supported models of a given quantum phenomenon within each interpretative framework and identifying claims about the unobservable which are true in all accounts.

This will usually involve leaving textbooks behind and looking at specialist physics and foundations of physics literature. In some cases it may require novel theoretical work – identifying interpretation-neutral claims that can be advanced in quantum chemistry, for instance, is a substantive project which may turn on new theoretical results and conceptual moves. Unlike the textbook approach then, our implementation of the overlap strategy is more of a research programme than a concrete answer to the question of the scope of scientific knowledge in the quantum domain. To us, this is not a problem but a reflection of the fact that it is hard work to delimit what we can really claim to know about the world based on the predictive success of mathematised physical theories.

3.3 Next Steps: The Wavefunction, Decoherence and an Interpretation-Neutral Notion of Possibility

Further reflection on the tunnelling example just discussed raises some doubts about the promise of this programme, however. It might be suspected that a special feature of our initial exemplar statement (*) which made interpretation-neutrality possible was that it does not make any explicit reference to the wavefunction (or Hilbert space vectors, or other means of representing the quantum state).¹¹ Since the interpretations say such different things about the wavefunction one might think that it will be impossible to recover any interpretation-neutral claims once it is being referred to directly. If this were correct, the scope of the overlap strategy would be severely limited as the wavefunction plays a central role in the characterisation of almost all quantum phenomena. In the example of quantum tunnelling discussed above, for instance, all three interpretations agree that one component of the wavefunction is transmitted through the potential barrier while another is reflected, but if the physical contents assigned to this superposition are completely disjoint this will apparently not amount to any substantive agreement on descriptive claims about the world.

In this subsection, we sketch a more ambitious implementation of the overlap strategy which attempts to combat this worry. Generalising the discussion of the hydrogen atom in 3.1, we can identify two dimensions which a proponent of the overlap strategy needs to remain neutral on: whether the wavefunction provides a complete or incomplete characterisation of the state of a quantum system, and whether there is a dynamical collapse mechanism that violates the

¹¹ In the following discussion we focus on the wavefunction representation of the quantum state for the sake of rhetorical simplicity. We are not assuming, as wavefunction realists do, that configuration space provides an ontically preferred mathematical representation of the quantum state.

Schrödinger dynamics or not. How our three interpretations answer these questions is indicated in Figure 1. We will now argue that one can plausibly remain neutral on both of these dimensions while still assigning physical significance to statements about the wavefunction.

	No-Collapse	Collapse
Ψ -complete	Everett	GRW
Ψ -incomplete	Bohm	GRWm

Figure 1. Table categorising the interpretations of quantum mechanics we are considering on Ψ -complete/ Ψ -incomplete and collapse/no-collapse dimensions. GRWm is a variant of GRW quantum mechanics in which a mass density field is added to the theory (see for instance Esfeld 2014) which is included for illustrative purposes.

Let us start with neutrality on the Ψ -complete/ Ψ -incomplete dimension. At least at the formal level, there appears to be substantial overlap in the kinematic structures posited by the three interpretations. Everettian, Bohmian, and GRW quantum mechanics all posit the wavefunction and use the same Hilbert space formalism to describe it. The key point they differ on is whether something else needs to be added to the wavefunction in order to fully characterise the kinematically possible states of a quantum system: Bohmian mechanics adds particle positions, and some GRW variants (labelled GRWm in Figure 1) add a mass density field. The overlap theorist will remain agnostic about whether these additional structures are needed to represent all of the properties of quantum systems. But, *prima facie*, the door appears to be open to interpret statements about the state of the wavefunction as true descriptive claims on all three interpretations.

It is instructive at this point to clarify how the overlap strategy, as we are conceiving of it, differs from the so-called “redundancy argument” for the Everett interpretation (Brown and Wallace 2005), which has a superficially similar structure. The redundancy argument is also based on the observation that there is substantial overlap in the kinematic structures posited by Everett and (typically) Bohmian quantum mechanics. Since one can already locate definite outcomes for classical observable quantities within the structure of the wavefunction, according to the many-worlds theorist, the hidden variables posited by the Bohmian are redundant. The Bohmian is thus sometimes accused of being an Everettian in denial. Will this accusation not also apply to the overlap theorist? If they believe claims about the wavefunction will they not, thereby, end up believing in many worlds?

As has been noted in the literature, the redundancy argument only has force if it is admitted that one can solve the measurement problem without positing additional kinematic structure (Callender 2011). Proponents of Bohmian mechanics, of course, typically will not concede this point, holding instead that the hidden variables they posit are required in order to explain how definite outcomes are obtained in agreement with the Born rule. Since this is one of the central points of contention in the debate surrounding the interpretation of quantum theory the overlap theorist wants to remain neutral on this point, not join the Everettian camp and commit themselves to the existence of many worlds.¹² They are able to do this because, unlike the many-worlds theorist, they do not commit themselves to the descriptive completeness of the wavefunction.¹³ They are open to the possibility that hidden variables (or a spontaneous collapse mechanism) selects one of the Everettian worlds as ontologically preferred.

This brings us to the collapse/no-collapse dimension, which, at first blush, seems to be more problematic for the overlap theorist. The non-linear, indeterministic dynamics for the wavefunction posited by GRW quantum mechanics appears to be fundamentally different from the linear, deterministic Schrödinger evolution posited by the Everettian and Bohmian approaches. Nevertheless, this disagreement on the dynamical behaviour of the wavefunction is arguably less drastic than it initially seems. Decoherence theory provides a framework for making this point manifest. As Rosaler (2016) discusses in detail, decoherence theory allows us to make considerable progress in recovering classical trajectories from quantum theory in an interpretation-neutral way; it describes how the coupling of a quantum system to a large number of environmental degrees of freedom leads the components of its wavefunction to separate into weakly interfering “branches” which can be approximately identified with classical trajectories. Spontaneous collapse events posited by empirically viable GRW models will arguably take place on a long enough time scale that the branching dynamics described by decoherence theory will not be spoiled (Bacciagaluppi, 2003).¹⁴ Thus, decoherence theory captures a substantial and important overlap in the dynamical content of all three of our interpretations.

¹² In fact, if we accept the premise that the wavefunction alone is sufficient to solve the measurement problem it still does not follow that we should believe in many worlds. Hidden variables could still exist even if they are not required to solve the measurement problem and act to select a particular branch of the wavefunction as ontologically preferred. If we remain agnostic about this possibility we will thus remain agnostic about the existence of many worlds.

¹³ Note that this verdict is also reinforced by the standard scientific realist line that we can only claim *approximate* truth for successful theories. The realist should, therefore, be open to adjustments to the structure of their theories and should not commit themselves to the completeness of their representations; cf. Vickers (2020) for a pertinent case study.

¹⁴ If decoherence-incompatible GRW models do turn out to be empirically viable this will be highly problematic for our attempt to identify dynamical overlap between the three interpretations, however.

From this point of view, in fact, the main point on which the Everett, Bohm, and GRW interpretations disagree concerns how one of the decoherent branches of the wavefunction is selected as physically salient for a particular observer (see Figure 2). Following a suggestion from Rosaler (2016), we can understand the story each interpretation provides about the selection of a particular branch as enacting a kind of “effective collapse”. In the case of GRW, of course, there are real stochastic collapse events, which move most of the mass of the wavefunction into a single branch. In the other two approaches, however, one still has an apparent collapse relative to an observer. On the Everett interpretation, measuring a quantum observable allows the observer to locate themselves within the structure of the wavefunction; on the Bohm interpretation measurements allow an observer to identify which branch is “occupied” by the hidden variables. On both accounts, the other branches of the wavefunction become essentially irrelevant to the future dynamical evolution of the selected branch. Just as we proposed using “electron” and “transitioned” in an abstract way, we can arguably use the term “effective collapse” to abstractly refer to the mechanism by which a particular branch is realised for an observer, while remaining non-committal about which of the specific mechanisms posited by each interpretation is operative.¹⁵




Figure 2. Depicting the branching structure described by decoherence theory and the status of the different branches in our three interpretations. Taken from Rosaler (2016) (with permission).

Another insight from Rosaler (2016) points to a powerful sense in which we can associate non-trivial physical content with the wavefunction without committing ourselves to a particular interpretation. He points out that we can take the decoherent branches of the wavefunction to represent “possible” outcomes for quantum observables on all three interpretations:

¹⁵ There is a legitimate worry that the notion of “effective collapse” just sketched is too thin to underwrite genuine knowledge claims about unobservable aspects of quantum systems. We acknowledge that the positive proposals discussed in this section are speculative and that more discussion of this notion of “effective collapse” in particular is needed. Remember, however, that scepticism about the notion of “effective collapse” need not imply scepticism about the overlap strategy as a whole since our previous examples of common theoretical claims arguably do not hang on the more ambitious proposals put forward in this section.

At the coarse level of description where interpretation-specific details of collapse are omitted, there is a correspondingly coarse sense of “possibility” on which future descendants [sic] of one's current branch all represent distinct possible future evolutions, and on which branches other than the one that happens to have been realized could have possibly been realized instead. (p. 63, fn. 22)

On the Bohmian interpretation the different branches of the wavefunction correspond to outcomes of quantum observables which could have occurred if the initial conditions of the hidden variables had been different. In addition, if we lack complete knowledge of the initial conditions, we can see them as representing states of affairs which are “possible” in an epistemic sense. On the GRW interpretation, the different branches of the wavefunction can be read as corresponding to genuine ontic possibilities given the indeterministic nature of the dynamics. On the Everett interpretation, all branches of the quantum state are realised in the structure of the multiverse, however, there is arguably a sense in which they nevertheless correspond to distinct possibilities for particular observers. Alastair Wilson has developed an account of modality in the Everett interpretation which makes this explicit (Wilson, 2020), but even an Everrettian who pulls back from a thick modal reading of decoherent branches has to admit that they play the functional role of alternative possibilities in decision-theoretic derivations of the Born rule (see Greaves, 2007, for relevant discussion).¹⁶ Our suggestion then is that we can understand statements about the wavefunction as encoding claims about the physically possible states and evolutions of a quantum system while remaining open to different, more precise, analyses of the nature of these “possibilities” provided by particular ontic interpretations.

In light of the foregoing discussion, let us now return to our example of quantum tunnelling to illustrate the kind of positive claims the overlap theorist can make with reference to the wavefunction. After the incident wavepacket has split into a transmitted and reflected component, interactions with the environment will give rise to two decoherent branches – all of the interpretations agree on this part of the story. We can then say that it is “possible” that the particle was transmitted through the potential barrier, or was reflected, in the abstract sense indicated above, and that one of these possibilities will be actualised via an “effective collapse”. The different interpretations disagree about the precise nature of these possibilities, and about the details of the mechanism which actualises one of them, but we can arguably still take the

¹⁶ A more careful treatment of the characterisations of quantum modalities associated with the rival interpretations is clearly needed here, but we hope our impressionistic discussion will form a useful starting point. Thanks to Alastair Wilson for useful correspondence on these matters.

wavefunction to encode descriptive claims that go well beyond anything a constructive empiricist is willing to countenance by using the terms “possibility” and “effective collapse” abstractly.

3.4 Options for Ontology and Semantics

We conclude this discussion of the overlap strategy by considering its relationships to more general positions in the metaphysics of science and the semantics of scientific theories. This will also allow us to address some final lines of objection to our programme.

Consider the following worry about the approach to the wavefunction put forward in the previous subsection. Some contemporary metaphysical readings of Bohmian mechanics do not take the wavefunction to represent part of the kinematic state of a quantum system, as was implicitly suggested in our discussion of neutrality on the Ψ -complete/ Ψ -incomplete dimension, but instead want to read it as representing a *law* (Esfeld et al, 2014). Thus, it would seem, the overlap theorist cannot be a realist about the wavefunction when the full family of rival ontological packages is taken into consideration.

It is important to emphasise, however, that we were not in section 3.3 endorsing wavefunction realism in the sense of Albert (1996). Wavefunction realism is a thesis about the fundamental ontology of the quantum domain, but, as we flagged at the outset, we doubt that the overlap will contain claims about fundamental metaphysics. Furthermore, nothing in the discussion of section 3.3 hangs on the idea, endorsed by wavefunction realists, that configuration space provides an ontologically preferred mathematical representation of the quantum state - the story could have been told equally well in the language of matrix mechanics. What we were concerned with in 3.3 was whether claims about the quantum state can be taken to express non-trivial physical content endorsed by all three interpretations, not whether we should admit the wavefunction into our fundamental ontology.

A follow-up worry arises here, however, which may have been lingering in the minds of some readers since the discussion of our initial test statement (*) in section 3.1. Surely, the thought goes, if we are to assign meaning to statements about the electron and the quantum state we need to provide a metaphysical account of what those entities are, and this will inevitably require

commitment to *one* of the ontological packages associated with a *particular* solution to the measurement problem.¹⁷

There are, we think, a number of lines of response an aspiring overlap theorist might adopt here; in this subsection, we highlight two options.

First, the overlap theorist might push back against the presumption that engaging in the metaphysics of quantum mechanics necessarily means doing *fundamental* metaphysics. Egg (2021) advances an “effective ontology” approach to quantum metaphysics which takes quantum particles and the wavefunction to be non-fundamental physical entities. Drawing on scientific realist readings of the effective field theory framework in high energy physics put forward by Fraser (2018) and Williams (2019), Egg argues that we need to adopt a meta-ontological framework which recognises a plethora of non-fundamental scientific entities in addition to the fundamental ontology posited by a theory, and that doing so helps us to articulate substantive ontological overlap between the rival interpretations of quantum mechanics. According to Egg, the wavefunction, and properties like charge and spin, do genuinely exist in a sparse Bohmian possible world, it is just that they are not fundamental. The idea is that these effective entities are common to all three interpretations so ontologically committing to them is not problematised by underdetermination. If one accepts the premise that commitment to an ontological picture of the quantum domain is required for extra-empirical claims about quantum systems to be assigned non-trivial meanings then the effective ontology approach provides resources for a direct response to the challenge on behalf of the overlap theorist.

Admittedly, the effective ontology framework stands in need of further development – see Saatsi (2021) for a recent critique. A second possible line of response is to reject the claim that one must engage in metaphysical analysis in order to assign a meaning to theoretical statements about electrons, spin and the wavefunction. It is notable that, while metaphysical theorising has become more prominent in the literature on the philosophy of quantum mechanics, contemporary forms of scientific realism have become increasingly quietist about metaphysics. An overlap theorist wanting to avoid the more metaphysically inflationary route represented by the effective ontology approach might instead respond to the challenge by sharply distinguishing semantics from ontology.

¹⁷ Thanks to an anonymous reviewer for pushing us on this point.

Why think that we need to commit ourselves to an account of the fundamental ontology of quantum systems to assign meaning to claims like “The electron in a hydrogen atom transitions between discrete energy states”? This thought might be motivated by a form of semantic reductionism: the view that in order to specify the meaning of a statement we need to unpack each of the terms contained within it until we reach a base level of primitive concepts. If we take the reduction base to be statements about fundamental metaphysics then the desired conclusion will follow.¹⁸ However, this extreme form of semantic reductionism is certainly not a popular position in the philosophy of language literature and seems highly implausible when applied to everyday language and the special sciences.¹⁹ We do not typically take perennial debates about the metaphysics of properties and mereological composition to problematise everyday knowledge claims; we need not contemplate “simples arranged catwise” to know that the cat is on the mat, for instance. Furthermore, we do not expect chemists or biologists to explicate their scientific concepts in terms of fundamental metaphysical posits. This suggests that an opponent can, at most, demand a more moderate, *partial* unpacking of terms like “wavefunction” and “electron”. However, we would claim that we have partially unpacked these terms in an interpretation-neutral fashion in the above discussion. Characterising the electron as “a subatomic particle with mass m_e , charge $-e$ and spin $\frac{1}{2}$ ” already assigns meaningful content to that concept, as we urged in 3.1, and it is not clear that associating some descriptive statements with a scientific concept has to be accompanied by explicit ontological commitments.

One might cash this second options out in terms of the distinction between the truth conditions and its truthmakers. While the truth conditions of a statement are often connected to its semantic content, its truthmakers are ontological items out in the world which, according to many contemporary truthmaker theories, need not mirror its semantic structure.²⁰ The overlap theorist might admit (*pace* Egg 2021) that the truthmakers for a statement like (*) are very different in

¹⁸ An alternative end-point for semantic reductionism is ‘observation statements’. On a semantic theory like this, (*) might come out as true on all three interpretations but ends up being a statement about observed spectral lines (so not a candidate for realist commitment). Of course, this approach to theory semantics has been thoroughly rejected by both sides of the modern scientific realism debate. Cf. Van Fraassen (1980, 1989, 1991) who develops an empiricist epistemology of science whilst rejecting the reductionist semantics of the logical empiricists (see Douven 1998, p. 127).

¹⁹ The other end of the spectrum to semantic reductionism is Fodor’s conceptual atomism which accepts no (or very little) reduction of natural language concepts (Fodor 1978, 1981). Between these extreme positions, which are largely unpopular today, lie a number of more moderate positions regarding the relationships between scientific and ordinary language concepts (see Margolis and Laurence 1999, Section 6).

²⁰ Schaffer (2008), p. 7, writes that “truthmaker commitments are what a theory says is fundamental”. On this view, the truthmakers for (*) are different for the different interpretations of quantum mechanics. If one now ties meaning to (possible) truthmakers, the meaning of (*) is radically different on the three interpretations and would not qualify as an instance of descriptive overlap between them. But few would wish to tie meaning to fundamental truthmakers. Indeed, one of the main motivations for introducing the truthmaker concept has been to free theorising about fundamental metaphysics from the shackles of linguistic analysis, cf. Heil (2003).

possible worlds which are exactly described by Everett, Bohm and GRW quantum mechanics, but deny that this means that the truth conditions vary according to the account one adopts. Overlap theorists could claim to know (*) without knowing what the truthmakers of (*) are, just as one can have knowledge of statues without knowledge of the truthmakers (Heil 2003, p. 53). Similarly, they could assert that certain statements about the wavefunction are true without committing themselves to the existence of the wavefunction. The scientific realism debate, as we characterised it, is not fundamentally about ontology but about whether we know that certain extra-empirical statements are true. Drawing on these sorts of resources, the overlap theorist could argue that they can supply meanings for interpretation-neutral statements without delving into quantum metaphysics.

We acknowledge that the overlap strategy is likely incompatible with some views of theory semantics and naturalist metaphysics, but our objective was never to convince all opponents. We have put forward two packages for semantics and meta-ontology that the overlap theorist might adopt, both of which have some currency in the contemporary scientific realism debate. This is sufficient to establish that the overlap strategy is a live option. Furthermore, while the overlap strategy may sit in tension with some projects in the metaphysics of science, Egg's effective ontology approach demonstrates that it is not intrinsically opposed to metaphysical theorising if one is happy to take non-fundamental ontology seriously. Rather, it conflicts with the idea that we can reliably read off statements about the fundamental metaphysical structure of reality from quantum theory. Some might view this as a heavy price, others (like us) a necessary concession.

4 The Overlap Strategy and Working Posits Realism

Much more work remains to be done on the overlap programme elaborated in the previous section. Still, we take our discussion to establish that this response to the quantum underdetermination problem is more powerful than is commonly thought. Since one cannot justifiably commit oneself fully to a particular interpretation on the basis of present evidence in our view, developing the overlap strategy further is the most promising route to defending knowledge claims about the quantum domain.

In this final section we consider some general epistemological questions that arise when we compare the overlap strategy to more general implementations of selective scientific realism. Our

discussion will lead us to flag two issues that need to be explored further before the viability of a realist epistemology of quantum theories can be fully assessed: whether the overlap strategy is vulnerable to an unconceived alternatives problem and whether it problematically weakens the explanatory component of scientific realism.

4.1 Unconceived Alternatives?

The overlap strategy is clearly allied to a broadly selective formulation of scientific realism, which takes us to be justified in believing only some of the claims about unobservables made by empirically successful theories. However, it appears to operate differently from the working posits approach that many selective realists have used to articulate their commitments.²¹ On the simplest definition, a working posit is a theoretical claim which cannot be removed from the derivation of a successful prediction without destroying the result (with removable assumptions being labelled “idle wheels”). Within each of our three ontic interpretations there appear to be theoretical claims which cannot be removed from the derivations of relevant empirical predictions, and therefore qualify as working posits by this definition, but do not lie within the descriptive overlap we started to sketch in section 3. If we remove the spontaneous collapse postulate from a GRW derivation of the probability of observing a particular outcome, for instance, the derivation no longer goes through.²² The working posits and the overlap strategies thus appear to conflict in this case, and it is the working posits approach that seems to get things wrong, since we arguably should not commit ourselves to the existence of a spontaneous collapse mechanism. Indeed, the quantum underdetermination problem can be read as a counterexample to a naive formulation of selective realism which takes us to be justified in believing claims about the unobservable which cannot be removed from successful predictions.

The question arises then – can the overlap strategy replace the notion of working posits as a master criterion for when we should make positive realist commitments? A number of issues arise if we simply believe the common claims made by the known theories which are compatible with

²¹ Variations of this approach are developed by Kitcher (1993), Psillos (1999), Vickers (2017), and Alai (2021), among many others.

²² One might object here that removing the spontaneous collapse mechanism from GRW quantum mechanics simply leaves you with Everett quantum mechanics so the derivation need not be ruined. This line of thinking is closely related to the “redundancy argument” for the Everett interpretation, which as we discussed in section 2.2 is controversial. Removing the collapse postulate from a GRW derivation of empirical predictions does not immediately yield an Everettian one, in our view, as the latter makes use of additional interpretive postulates, namely the identification of decoherent branches of the wavefunction with quasi-classical worlds.

relevant empirical data (a naïve believe-the-overlap procedure). Consider the results this would give in different actual, and counterfactual, historical scenarios. Of the three interpretations we have focused on here, the Bohmian interpretation is the oldest – the first version of this approach dates back to De Broglie’s 1927 Solvay conference presentation, and Bohm’s first papers on the subject appeared in 1952. Suppose we grant that in 1952 Bohm’s interpretation was the only way to understand non-relativistic quantum mechanics as a descriptive theory – that is, there was no underdetermination present. Employing a believe-the-overlap procedure in that year would apparently require one to believe all of the claims of Bohmian quantum mechanics!

This is problematic for a few reasons. An epistemic agent following this rule would have to radically strip back her commitments when the Everett interpretation came onto the scene and revealed that one could understand quantum mechanics as a descriptive theory without committing oneself to Bohmian hidden variables. Employing the overlap strategy as a master criterion for what we ought to believe would thus lead to retractions of scientific claims which were previously taken to be justified. If we frequently have to retract previous beliefs in the face of new evidence this suggests that the rules we are using to evaluate the support for scientific claims are unreliable. A second problem arises from the fact that we can easily imagine counterfactual histories in which the GRW interpretation was developed first, leading a follower of the overlap strategy to commit themselves to the existence of a real indeterministic collapse mechanism. A naïve believe-the-overlap procedure is thus too sensitive to contingent facts about the order in which independent theoretical discoveries are made.

Clearly then, the overlap strategy cannot be employed in such a flatfooted way. If it is to have a role in the scientific realist’s arsenal, it must be embedded in a broader epistemic framework that incorporates additional assumptions. Some reflection reveals that the plausibility of a believe-the-overlap move in the quantum underdetermination scenario is predicated on the assumption that physicists have done a good job of exploring the space of viable ontic interpretations. More generally, putting one’s faith in the overlap of underdetermined theories makes sense when the exploration of the space of theories compatible with relevant empirical data is assumed to be complete, or close to complete.

If one wants to hold that believing in the overlap of competing ontic interpretations was ill-advised in 1952 but is sensible some 70 years later, one is thereby committed to the claim that

physicists have made major progress in probing the space of viable interpretations in the interim. In support of this, we can point to the long period of time over which a concerted search for solutions to the measurement problem has taken place. We can also potentially put weight on results like the Bell and Kochen-Specker theorems which both constrain the space of viable interpretations and offer a means of categorising the existing approaches in terms of the assumptions of these theorems that they reject.

A sceptic with sympathies for Kyle Stanford's unconceived alternatives argument may be unconvinced by this, however. If there are unconceived interpretations of quantum mechanics which do not share the overlap we started to identify in Section 3, then we will be employing the overlap procedure prematurely, just as one would have in 1952. In addition to general arguments used to support the unconceived alternatives worry, there are local features of the quantum foundations context which might raise worries about the completeness of the current landscape of interpretations (cf. Vickers 2022, Chapter 6). The sceptic can point to the fact that the search for interpretations of quantum mechanics has been a somewhat marginal activity in the context of 20th-century physics, and to the appearance of new interpretations (and new variants of existing interpretations) in recent decades, to support the counterclaim that we have likely not yet exhausted the relevant space of theories.²³

In sum, the promise of the programme for defending knowledge claims about the quantum domain set out on this paper is conditional on a response to the unconceived alternatives problem – not just in its general form but also in the specific context of the interpretation of quantum theory, where it appears to have some additional intuitive bite.

4.2 Explanatory Deficit?

There is another issue that arises if we consider replacing the working posits approach with the overlap strategy that is important to consider. In addition to taking successful scientific theories to furnish knowledge of unobservables, the scientific realist typically also claims that the truth of these theoretical claims explains the theory's predictive accuracy. Even if it is admitted that there is a substantial descriptive overlap between the various ontic interpretations of quantum theory,

²³ The sceptic might also point to more general worries about claims of completely exploring a space of mathematised physical theories – see Dardashti (2019).

there is a remaining worry that this common content is insufficient to do that kind of explanatory work. The conflict between the overlap and working posits approach points in this direction: we seem to need the distinctive theoretical claims made by the various ontic interpretations in order to solve the measurement problem and explain the predictive efficacy of the Born rule.

A possible response to this problem is to concede that the knowledge and explanatory components of scientific realism come apart in this context. The trouble with giving up the explanatory thesis, however, is that it is implicated in the traditional no-miracles argument for scientific realism. Consequently, there is a danger that our reason for taking a positive epistemic attitude towards the theoretical claims of quantum theories in the first place will be undermined. Like Hoefer's response to the quantum underdetermination problem discussed in section 2.1, taking this route will require some careful recalibration of one's general epistemology of science.

There is another possibility: it may be that by adopting a more nuanced version of the working posits strategy one can dissolve the apparent conflict with the overlap strategy and rescue the explanatory component of scientific realism in the process.²⁴ Rather than simply looking for assumptions that cannot be removed from the derivations of empirical results, more sophisticated implementations of the working posits approach require us to check whether it is possible to replace the assumptions of a given derivation with more abstract assumptions that are also capable of deriving the same result (Vickers 2017; Alai 2021). The thought is that the abstract descriptive claims we identified as common to our three ontic interpretations in section 3 might equivalently be obtained by an abstraction procedure of this kind. Earlier we said that the dynamical collapse mechanism cannot be removed from a GRW model without spoiling its predictions. The existence of no-collapse interpretations suggests that this posit can be replaced by a more abstract statement that does not commit one to the existence of real stochastic collapse events. Specifically, it suggests that all that is needed to recover the empirical predictions of quantum theory is the weaker notion of an "effective collapse" discussed in section 3: a generic mechanism that selects outcomes to be realised for a particular observer in accordance with the Born rule. Thus, so this line of thought goes, we reach the same end-point from two different directions: (i) by asking if we can make the derivation go through with a more abstract

²⁴ Some implementations of the working posits concept are purely negative – that is they merely try to identify idle wheels which the realist should not commit themselves to as a defensive maneuver against anti-realist historical arguments (Vickers 2017). This is another way that sophisticating the working posits approach might dissolve the apparent tension with the overlap strategy.

assumption, and (ii) by asking what assumptions are shared by all of the underdetermined ontic interpretations.

In truth, it was never clear just *how* abstract a practitioner of this sophisticated working posits strategy should go. The overlap strategy can thus be understood as guiding the appropriate level of abstraction for our commitments to be epistemically secure. As we argued in section 3, the amount of abstraction required to reach interpretation-neutrality is not so dramatic that one is left making completely trivial claims. However, a question remains about the strength of the explanation of predictive success which might be salvaged in this way. Advocates of particular ontic interpretations will claim, with some plausibility, that their theories offer robust physical explanations of the efficacy of the Born rule while the explanation available to an overlap theorist is, at best, extremely thin.

A final question that remains for the programme sketched in this paper then is whether it wins the battle but loses the war by depriving an aspiring scientific realist about quantum theory of the kind of explanatory resources they need to make their epistemic position compelling in the first place.

Acknowledgements

Thanks to audiences at the European Philosophy of Science Association biannual conference and the Bristol philosophy of physics seminar for useful feedback on presentations of this work. Alasdair Wilson, Mattias Egg, Juha Saatsi and three anonymous referees gave comments on earlier versions of the manuscript which help us immensely. James Fraser would like to thank the British Academy for funding during the period in which this research was conducted.

James D. Fraser

University of Durham

jamesf09@hotmail.co.uk

Peter Vickers

University of Durham

peter.vickers@durham.ac.uk

References

- Acuña, P. (2021). 'Charting the landscape of interpretation, theory rivalry, and underdetermination in quantum mechanics', *Synthese*: 198,1711-1740.
- Alai, M. (2021). 'The Historical Challenge to Realism and Essential Deployment', in T. D. Lyons and P. Vickers (eds.) *Contemporary Scientific Realism: The Challenge from the History of Science*. Oxford University Press.
- Albert, D. Z. (1996). Elementary quantum metaphysics. In Cushing, J. T., Fine, A., and Goldstein, S., editors, *Bohmian Mechanics and Quantum Theory: An Appraisal*, pages 277–84. Kluwer Academic Publishers, Dordrecht.
- Albert, D. (2010). 'Probability in the Everett picture', in Saunders, S., Barrett, J., Kent, A., & Wallace, D. (Eds.). *Many worlds?: Everett, quantum theory, & reality*. Oxford University Press.
- Allori, V. (2020). 'Scientific Realism without the Wave Function', in S. French and J. Saatsi (eds.) *Scientific Realism and the Quantum*. Oxford: Oxford University Press, pp. 212-28.
- Asay, J. (2011): *Truthmaking, Truth, and Realism: New Work for a Theory of Truthmakers* (Dissertation, University of North Carolina at Chapel Hill).
- Bacciagaluppi, G. (2003). 'The role of decoherence in quantum mechanics', *Stanford Encyclopedia of Philosophy*.
- Berta, P. et al. (1990). 'Genetic evidence equating SRY and the testis-determining factor', *Nature* 348: 448–50.
- Belousek, D. W. (2005). 'Underdetermination, realism, and theory appraisal: An epistemological reflection on quantum mechanics', *Foundations of Physics*, 35(4), 669-695.
- Blackburn, S. (1986): 'Morals and modals', in *Fact, Science and Morality: Essays on A. J. Ayer's Language, Truth and Logic*, eds. G. Macdonald and C. Wright (Oxford: Basil Blackwell): 119-41.
- Brown, H. R., & Wallace, D. (2005). 'Solving the Measurement Problem: De Broglie--Bohm Loses Out to Everett.' *Foundations of Physics*, 35(4), 517-540.
- Callender, C. (2020). 'Can we Quarantine the Quantum Blight?', in S. French and J. Saatsi (eds.) *Scientific Realism and the Quantum*, Oxford: OUP, pp. 57-77.
- Cordero, A. (2001): 'Realism and underdetermination: Some clues from the practices-up', *Philosophy of Science* 68: S301-S312.
- Dardashti, R. (2019): 'Physics Without Experiments?', in R. Dardashti, R. Dawid, K. Thébault (eds.) *Why Trust a Theory? Epistemology of Fundamental Physics*, Cambridge: Cambridge University Press, pp. 154-72.
- Dawid, R., & Thébault, K. P. (2015): 'Many worlds: decoherent or incoherent?', *Synthese* 192(5): 1559-80.

- Douven, I. (1998): 'Truly Empiricist Semantics', *Dialectica* **52**(2): 127-51.
- Egg, M. (2019). 'Dissolving the measurement problem is not an option for the realist', *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* **66** 62-68.
- Egg, M. (2021). 'Quantum Ontology without Speculation', *European Journal for Philosophy of Science* **11**, 32.
- Esfeld, M. (2014). 'The primitive ontology of quantum physics: guidelines for an assessment of the proposals', *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*. **47**: 99-106.
- Esfeld, M., Hubert, M., Lazarovici, D., & Dürr, D. (2014). The ontology of Bohmian mechanics. *The British Journal for the Philosophy of Science*, **65**(4), 773-796.
- Fodor, J. (1978): 'Tom Swift and his procedural grandmother', *Cognition* **6**: 229-47.
- Fodor, J. (1981): 'The present status of the innateness controversy', in *Representations: Philosophical essays on the foundations of cognitive science* (Cambridge, MA: MIT Press): pp. 257–316.
- Fuchs, C. A. (2017). 'On participatory realism', in *Information and interaction*. Springer, Cham. pp. 113-134
- Fraser, J. D. (2018). Renormalization and the formulation of scientific realism. *Philosophy of Science*, **85**(5), 1164-1175.
- French, S. (2011). 'Metaphysical underdetermination: why worry?', *Synthese*, **180**(2), 205-221.
- Greaves, H. (2007). Probability in the Everett interpretation. *Philosophy Compass*, **2**(1), 109-128.
- Healey, R. (2016). 'Quantum-Bayesian and pragmatist views of quantum theory', Stanford Encyclopedia of Philosophy.
- Healey, R. (2020). 'Pragmatist quantum realism', in S. French and J. Saatsi (eds.) *Scientific Realism and the Quantum*. Oxford: Oxford University Press. 123-146.
- Heil, J. (2003): *From an Ontological Point of View*. Oxford: Oxford University Press.
- Hofer, C. (2020). 'Scientific realism without the quantum', in S. French and J. Saatsi (eds.) *Scientific Realism and the Quantum*. Oxford: Oxford University Press. pp. 19-34.
- Hofer, C. and Martí, G. (2020): 'Realism, reference & perspective', *European Journal for Philosophy of Science* **10**(38).
- Holland, P. R. (1995). *The quantum theory of motion: an account of the de Broglie-Bohm causal interpretation of quantum mechanics*. Cambridge university press.
- Kitcher, P. (1993). *The Advancement of Science: Science Without Legend, Objectivity Without Illusions*. Oxford: Oxford University Press.
- Klima, G. (2009): *John Buridan* (New York: Oxford University Press).

- Lakoff, G. (2008): *Women, Fire, and Dangerous Things: What Categories Reveal about the Mind* (Chicago: University of Chicago Press).
- Margolis, E. and Laurence, S. (1999): *Concepts: Core Readings* (Cambridge, MA: MIT Press): Chapter 1.
- Myrvold, W. (2016). 'Philosophical issues in quantum theory', *Stanford Encyclopedia of Philosophy*.
- Ney, A. (2012). 'Neo-positivist metaphysics', *Philosophical studies*, 160(1), 53-78.
- Norsen, T. (2013): 'The pilot-wave perspective on quantum scattering and tunneling', *American Journal of Physics* 81: 258-266.
- Norton, J. D. (2000). How we know about electrons. In *After Popper, Kuhn and Feyerabend* (pp. 67-97). Springer, Dordrecht.
- Psillos, S. (1999). *Scientific Realism: How Science Tracks Truth*. London; New York: Routledge.
- Rosaler, J. (2016). Interpretation neutrality in the classical domain of quantum theory. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 53, 54-72.
- Saatsi, J. (2019). 'Scientific realism meets metaphysics of quantum mechanics', in Cordero, A. (ed.) *Philosophers Look at Quantum Mechanics*. Cham, Switzerland: Springer. pp. 141-162
- Saatsi, J. (2021). '(In)effective realism?' *European Journal for Philosophy of Science*.
- Schaffer, J. (2008): 'Truthmaker Commitments', *Philosophical Studies* 141: 7–19.
- Seifert, V. A., & Franklin, A. (2020). The Problem of Molecular Structure Just Is The Measurement problem. Forthcoming in *British Journal for the Philosophy of Science*.
- Squires, E. J. (1991). Wavefunction collapse and ultraviolet photons. *Physics Letters A*, 158(9), 431-432.
- Stanford, P. K. (2006) *Exceeding Our Grasp*. Oxford: Oxford University Press.
- Van Fraassen, B. (1980): *The Scientific Image* (Oxford: OUP).
- Van Fraassen, B. (1989): *Laws and Symmetry* (Oxford: OUP).
- Van Fraassen, B. (1991): *Quantum Mechanics: An Empiricist View* (Oxford: OUP).
- Vickers, P. (2017). 'Understanding the selective realist defence against the PMI', 194(9), 3221-3232.
- Vickers, P. (2020). Disarming the ultimate historical challenge to scientific realism. *The British Journal for the Philosophy of Science*, 71(3), 987-1012.
- Vickers, P. (2022): *Identifying Future-Proof Science*. Oxford: Oxford University Press.
- Vision, G. (2003): 'Lest we forget 'the correspondence theory of truth'', *Analysis* 63: 136-42.
- Wallace, D. (2012). *The emergent multiverse: Quantum theory according to the Everett interpretation*. Oxford University Press.

Williams, P. (2019). Scientific realism made effective. *The British Journal for the Philosophy of Science*, 70(1), 209-237.

Wilson, A. (2020). *The nature of contingency: Quantum physics as modal realism*. Oxford: Oxford University Press.