The inferentialist guide to quantum mechanics

Tushar Menon

ABSTRACT

This paper introduces a new approach to understanding quantum mechanics (QM) called 'pure-inferential quantum mechanics.' The animating thought behind this approach is that, while theories can be about the world without describing it, we would be just as wrong to assume, as anti-realists such as QBists do, that quantum states *never* describe the world, as we would be to assume, as realists such as Everettians do, that they *always* do. I marshal some resources from pragmatist philosophy of language to delineate the circumstances under which QM should be understood as descriptive from those under which it should not be. By constructing a novel inferentialist-pragmatist interpretation of QM, I demonstrate, *pace* Healey, that the inferentialist about QM has the resources to ground an autonomous, descriptive interpretation of QM in its inferential profile, under the appropriate circumstances. I argue that, if we should be inferentialists about QM (and there are good reasons why we should be), then we should be pure-inferentialists.

1 *Introduction*

2 *Pragmatism and the descriptive capacity of physics*

3 *Brandomian inferentialism*

4 *Anchored-inferentialist quantum mechanics*

- **4.1** *Healey's approach*
- **4.2** *Assessing Healey's approach*

5 *Pure-inferentialist quantum mechanics*

- **5.1** *The decoherent pure-inferential interpretation of quantum mechanics*
- **5.2** *Decoherent pure-inferential quantum mechanics vs. Healey's interpretation*

6 *Conclusion*

1 Introduction

There is no research programme in philosophy that seeks to understand what the world is like according to dishwasher user-manuals. And there is no mystery about why. User-manuals are not in the business of describing the world, they are in the business of guiding people in their use

of particular machines to cope with the world. To expect a world-picture to emerge from user-manuals is to misunderstand what they can deliver. But user-manuals are still *about* the world. They just do not describe it.

Newtonian mechanics as applied to everyday objects, by contrast with user-manuals, is in the business of describing the world. It works by ascribing properties to objects in such a way as to predict the dynamical evolution of those properties. Now as a consequence, some Newtonian theories do, in fact, turn out to be helpful in guiding people in their use of particular pieces of mathematical machinery to cope with the world. But that utility is grounded in the fact that Newtonian theories describe the world.

There is a near industrial-scale collection of research programmes in philosophy that seek to understand what the world is like according to quantum mechanics (QM). A majority of these programmes take for granted that QM is in the business of describing the world. But is it? QM is clearly about the world, but is it like a user-manual or a Newtonian theory? This is the animating question behind this paper. I will marshal some resources from pragmatist philosophy of language to argue that the answer, perhaps surprisingly, is both.

Now, admittedly, user-manuals are not expected, on their own, to be about the world. There is a tacit acceptance that there is some other descriptive theory which populates the world with the objects that the user-manual teaches us how to use. So dishwasher user-manuals come to be about the world because they guide our use of machines which are themselves described by some theory that goes beyond the user-manual; our everyday theory of medium-sized white goods. Interpretations of QM that have followed the user-manual model have similarly denied descriptive capacities to QM whilst outsourcing that responsibility to some other theory (see e.g. (Bohr [2010]; Caves et al. [2002]; Mermin [2014])). Healey's ([2012], [2017]) inferentialist-pragmatism about QM is one such interpretation: QM is a non-descriptive user-manual about magnitudes that are themselves part of some descriptive *non-quantum* theory. On this view, there is no such thing as 'the world according to QM (alone)'; QM is not descriptively autonomous, but is instead part of a larger set of discursive scientific practices that collectively describe the world.

In this paper, I demonstrate that the inferentialist about QM is not in the impoverished position of having to rely, for descriptive labour, on some other theory. They have the resources to ground an autonomous, *descriptive* interpretation of QM in its inferential profile. In other words, I will demonstrate that, in virtue of its action-guiding, user-manual character, under certain well-delineated circumstances, QM can come to be descriptive on its own. Those circumstances include all of our ordinary empirical discourse. I call this the 'pure-inferential' approach to quantum mechanics to contrast it with Healey's 'anchored-inferentialist' approach.

I begin, in §2, by laying out the criteria for when a practice can be understood as using elements of its discourse to represent parts of the world. In §3, I present the semantic machinery that underpins both Healey's approach to QM and mine: Brandom's inferentialism. In §4, I introduce Healey's interpretation of QM, and identify two substantial points of disagreement between us: (i) Healey reads Brandom as mandating what I will call a 'semantic' anti-representationalism,'

whereas I read Brandom's view as mandating a 'metasemantic' anti-representationalism; and consequently (ii) Healey's endorses an anchored-inferentialist view, when in fact there is a preferable pure-inferentialist option available. In §5 I present and defend this pure-inferentialist approach to quantum mechanics.

2 Pragmatism and the descriptive capacity of physics

In this section, I address what we might call the descriptive capacity question: when, and to what extent, can we claim that a theory of physics can be used to describe the world? I compare two approaches: pragmatism and anti-pragmatism. Pragmatism is something of a moving target. As with many philosophical isms, there is a host of claims that, in some intuitive sense, are naturally allied with pragmatism, and would be assented to by many pragmatists. It is important to be explicit about which of these claims are central to the pragmatist, and which follow from some additional commitments that are perhaps friendly, but ultimately peripheral.

In this paper, I take the central pragmatist tenet to be: our assessment of a claim's validity is conceptually inseparable from how that claim impacts our agency within the world. The assessment of a claim's validity is the object of study of semantics, and our agency within the world is the object of study of pragmatics. So the central tenet of pragmatism, as I understand it in this paper, says that semantics is wholly answerable only to pragmatics. Semantics is not conceptually autonomous from pragmatics.There are no semantic facts that are not pragmatic facts. Anti-pragmatism is just the denial of pragmatism.

The pragmatist needs to specify two commitments, since there are two aspects to assessing a claim. The first is establishing what the content of a claim is. The second is establishing whether or not that contentful claim is valid. Pragmatism is therefore a relational claim about both semantics and pragmatics, not an absolute claim about semantics. So one can adopt a wide variety of semantic positions, as long as one can suitably adjust one's view of pragmatics (or vice versa). For example, according to (Davidson [1967]), the semantic content of sentences (i.e. their meaning) is exhausted by a specification of truth conditions, via the semantic machinery of satisfaction and reference (which is a kind of representation). However, for Davidson, truth is an unanalysable primitive, characterised not by any sort of correspondence between word and world, but instead entirely by its role in our best theory of the behaviour of rational agents (Davidson [2001], Ch. 10). Davidson demonstrates, *pace* quantum pragmatists like Bächtold ([2008]), Healey ([2012], [2017]), Friederich ([2014]), and Ruyant ([2023]) that one can be a pragmatist even if one believes in the theoretical indispensability of representation for understanding QM on its own. One just needs to find a suitable pragmatics to ground representation.

Brandom's pragmatism embraces the converse strategy. It begins with pragmatics, which then guides the construction of a semantic theory, and in particular, informs the choice of what semantic machinery our theory will formalise. For Brandom, the sort of agency that demarcates discursive practices is the use of concepts. And conceptual activities are characterised as moves in

the game of giving and asking for reasons while communicating with other agents; in other words in making inferences. Inferences, on this view, ground meaning. On top of that, under certain circumstances, inferences can also establish relations of representation between words and parts of the word. But this is contingent and epiphenomenal. We should not mistake the appearance of representation relations in the vindication of claims, in certain contexts, as demonstrating that it is those representation relations that are doing the work of vindicating those claims.

The reason that Brandom counts as a pragmatist is that, for him, pragmatics grounds semantics because inference grounds meaning (and in some circumstances, representation). On the other hand, the reason that Davidson counts as a pragmatist is that, for him, pragmatics grounds semantics because (pragmatically characterised) truth grounds meaning, via representation. Pragmatism is, therefore, consistent with most interpretations of QM: the realists, on the one hand, make representational claims, as Davidson does, but there is no prohibition on grounding these representational claims in pragmatics, while the anti-realists, on the other hand, foreground the use to which QM is being put and consequently, as Brandom does, accept a certain semantic revisionism. While the latter might strike some as perhaps being more in the spirit of pragmatism, they are no more pragmatist than the former.

The observation that vindication of claims can come apart from representation is the basis of Price's ([2013]) distinction between subject- and object-naturalism.[1] The object-naturalist treats representation as necessary for (i) the content of claims: the meaning of a word is determined by the object in the world that it picks out, and its properties; and (ii) the vindication of claims: a sentence is vindicated if it is true, and a sentence is true if represents the world as being the way it is. As Price ([2013], p. 10) notes, '[o]bject naturalism... rests on... the assumption that substantial 'word-world' semantic relations are part of the best scientific account of our use of the relevant terms.' Understood in this way, object naturalism is a form of anti-pragmatism, because it grounds pragmatics in semantics, not the other way around. For the object-naturalist, the descriptive capacity question is answered as follows: a theory is meaningful only if it is descriptive. And QM is meaningful. Therefore it is descriptive. So to understand what QM says, we need to understand what it describes. Bell, for example, is an object-naturalist. For him, the key to understanding quantum mechanics lies in understanding 'beables', which are 'those entities in [a theory] which are, at least tentatively, to be taken seriously, as corresponding to something real' ([2004], p. 234).

By contrast, the subject-naturalist denies that representation is necessary for either vindication or content-conferral. Instead, the subject-naturalist looks to 'account for the use of various terms—among them, the semantic themselves—in the lives of natural creatures in a natural environment' ([2013], p. 14). Subject-naturalism is, therefore, a form of pragmatism. It seeks to provide an account of when the use of (among other things) representation-talk is appropriate, by looking at the vindication of such talk by its pragmatic consequences. A subject-naturalist can still believe in representation, and moreover, they can believe that representation is sufficient for

1 Nothing in Price's setup requires the subject-/object- distinction to apply only to naturalism. But in this context, where our interest is squarely in the domain of quantum mechanics, there is no reason to move to a more general context than Price's naturalism.

either content-conferral or vindication, as long as that claim follows from an appropriate account of how pragmatics grounds semantics.

To articulate the subject-naturalist's response to the descriptive capacity question, we need one final piece of machinery, from (Price [2011]) . Until now, I have relied on an intuitive understanding of 'representation' as something like a relation between an expression or a model, a part of the world, and perhaps an agent using the expression or model. If such a relation exists, then there are at least two distinct roles that it can perform, and it is important to prise them apart. The first is the 'environment-tracking' role: some linguistic parameter covaries with some worldly parameter, in the same way that, for example, the readings on a speedometer on a car varies with its speed. Price calls this 'e-representation'. The second is the 'content-conferring' role: the content associated with some expression is determined by the internal role played by that expression within some discourse. Price calls this 'i-representation.' For an object naturalist, these two roles coincide entirely: a word is meaningful (i-representationally) in virtue of the relation that tracks how some object in the environment covaries with the use of that word (e-representationally). But for a subject-naturalist, these two roles can come apart. A word can be meaningful in virtue of its place within some internal theoretical architecture, without having to reach out and track something in the world.

The subject-naturalist can assess the descriptive capacity of a theory of physics by asking the following question: to what extent does i-representation of terms coincide with e-representation? From this perspective, anti-representationalism is just the view that e-coordination does not ground i-coordination, while representationalism is the view that it does; both are consistent with subject-naturalism. In the next section, I will cash out i-representation in Brandomian inferentialist terms.This will allow me to state, and then assess, Healey's inferentialist interpretation of QM.

3 Brandomian inferentialism

Brandom's inferentialism is central to Healey's account, To understand Brandom's position, it is helpful to begin by distinguishing between three levels of philosophical discourse about language (I borrow the terminology from (Dasgupta [2020])):

Semantics: The theory of how semantic properties distribute through a language.

Metasemantics: The theory of what imbues elements of a language with those semantic properties.

Metametasemantics: The theory of why it is those semantic properties, rather than any others, that we are interested in.

Pragmatism, in general, begins by understanding the metametasemantic question as: given the aim of our discursive practice, which semantic concepts are best placed to achieve these aims? Different views of the aims of a discursive practice will deliver different verdicts on what the

appropriate semantic machinery is. If, for example, the aim of all discourse is description, then, plausibly, reference is a good candidate, and we should develop a metasemantics that explains how words come to refer. For Brandom, the aim of discursive practice is communication and justification of commitments ([2009], p. 165):

> [O]ne essential aspect of [the Brandomian] model of discursive practice is *communication*: the interpersonal, intra-content inheritance of entitlements to commitments... [A]nother... is *justification*: the intrapersonal, intercontent inheritence of entitlement to commitments.

Having thus identified the aim of discursive practice, Brandom answers the metametasemantic question by positing that *representation* is the central piece of semantic machinery that allows for the coordination central to communication (Brandom [2009], pp. 167-168):

> Talk about representation is talk about what it is to secure communication by being able to use one another's judgements as reasons, as premises in our own inferences, even just hypothetically, to assess their significance in the context of our own collateral commitments.

Wait. It seems as if we have taken a step back. If it is representation that grounds our ability to communicate, then surely that favours the object-naturalist, and consequently the representationalist. This is where Price's distinction is so crucial. If we do not separate e-representation from i-representation, and instead treat all representation as e-representation, then the pragmatist's subject-naturalism is a nonstarter. So the pragmatist's reading of the above Brandom quote must treat 'representation' as 'i-representation': communication is established by each speaker recognising the i-representational profile of the other's use of an expression. The Brandomian answer to the metametasemantic question is: i-representation is the semantic machinery that best accounts for the success of our discursive practice. This answer invites an immediate metasemantic follow-up question: what imbues an expression with an i-representational profile? For Brandom, it is the inferential profile of that expression.

To understand how what an expression's inferential profile is, we first need to understand the model of discursive practice that underpins Brandom's view. To assert a claim, on this view, is set oneself up to respond appropriately when challenged. On receipt of a challenge from another speaker, the original speaker can choose to defend their claim, modify it, defer to someone else or withdraw it entirely. Each of these actions is governed by some norms: assertions are understood in terms of how these norms regulate what speakers are allowed or required to do. These norms govern what Brandom calls a 'deontic scorecard', which is simply a way for interlocutors to keep track of their own, and each other's commitments and entitlements (call these 'deontic statuses'). Each speaker tacitly keeps track of everyone's deontic status, by keeping track of which commitments they attribute to others and which they acknowledge themselves. Note that, as Maher ([2012], p. 70) highlights, 'there is no "super-scorecard"'; since each speaker attributes different

commitments to different speakers, each speaker will, in general, ascribe a different deontic status to a speaker than the other speakers do.

The deontic scorecard is updated by adherence to some set of norms. Which norms? According to Brandom, it is the norms of inference. And these inferential norms are understood not in terms of compliance with any formal schemata, but instead in terms of the deontic statuses encoded in the scorecard. So, for example, if I had asserted that, and therefore committed myself to accepting that, a particular cricket ball was red at time t, then I would have been equally committed to the claim that that ball was visible at t. What makes that inference good or valid is that is conclusion preserves the premise's deontic status. Call such inferences 'material inferences'. That material inferences are good is a primitive fact about a discursive practice. It is not grounded in the fact that it could be rendered as an enthymeme.

Brandom's central metasemantic claim, then, is that these sorts of material inferences are meaning-conferring. An assertion is meaningful in virtue of the other assertions (and acts)[2] that it licenses, and that it is licensed by. Understood in terms of preservation of entitlement, material inferences might, although they do not need to, take the form of inferences that we ordinarily classify as deductive, inductive or abductive. But this model is capacious enough to incorporate inferences that might be given other formal characterisations. Note that the meaning-conferring inferences are not understood as inferences that a speaker is disposed to make, or that they actually make. Instead, they are the inferences that are licensed by the norms that are, perhaps tacitly, assented to by all the speakers. Indeed, the act of speaking is the act of making explicit using linguistic expressions, different aspects of what is implicit in the practice in which the speakers are embedded. And this explicitation of commitments is what brings them into the game of giving and asking for reasons, and moves in this game are kept track of with the deontic scorecard.

According to the Brandomian model, we can understand i-representation by understanding what is going on when speakers make *de re* ascriptions to other speakers. By contrast with the standard view, according to which a *de re* attribution to S is an attribution to S of a belief or commitment about some worldly object, on the Brandomian view, to attribute a *de re* commitment to S is make a specific claim about S's deontic status. Brandom argues that a *de re* ascription from, say, me to S (e.g. I say 'S believes of Big Ben that it is a clock') expresses (i) an attribution of a doxastic commitment to S, as well as (ii) an undertaking of a doxastic commitment by me (e.g. that Big Ben exists and has some properties, such as being a bell). The role of such locutions is to coordinate different social perspectives, by nailing down the conceptual content of my attribution to S, in order that I can then, for example, challenge S. Note that, in making a *de re* attitude ascription, even though I undertake a sort of existential commitment, I do not have to undertake a commitment to a relation of e-representation between the word and some part of the world.

2 In order to incorporate the interation between language and world, we need to broaden the concept of inference to include what Sellars ([1954]) calls 'language entry- and exit-transitions'. These are rules for how a worldly event or stimulus gets incorporated into inferences (I see a fast-moving projectile and say 'Look out! A cricket ball!') and how inferences can include intentional actions (you hear my claim and duck). These moves are also treated as inferences.

In other words, I can just as easily make sense of '*S* believes of Big Ben that it is a clock' as '*S* believes of unicorns that they can fly'. As long as claims involving them update the deontic scorecard as described above, both 'Big Ben' and 'unicorn' i-represent some semantic content. But only 'Big Ben' e-represents some part of the world.

So it is clear that, in all cases, an expression's inferential profile picks out its content, and if (but not only if) the associated deontic status is understood as making a *de re* attitude ascription, then an i-representation relation is established between an expression and its semantic content. Now, in some cases, an expression's inferential profile might also be sufficient to pick out an e-representational relation. But that latter relation is metasemantically inert: it plays no role in explaining why expressions mean what they do. This observation allows us to distinguish between two families of views that stake a good claim to the label 'representationalism'. The first is *semantic representationalism*, according to which some linguistic expressions stand in a relation of representation to the appropriate relatum. Note that there are distinct versions of these representationalisms that correspond to i- and e-representation, so for the e-representationalist, the appropriate relatum is a proper part of the world, while for the i-representationalist, the appropriate relatum is meaning or semantic content. A semantic anti-representationalist is a *representational eliminitavist*: expressions *never* stand in relations of representation to anything.

The second sense of representationalism operates at the metasemantic level: according to this *metasemantic representationalism*, the semantic content of a linguistic expression is determined by the relation of representation it stands in to its relatum. Brandom's inferentialism is metasemantically anti-e-representationalist but not semantically anti-e-representationalist. Indeed, two long chapters of (Brandom [1994]) are devoted to demonstrating how inferentialism can ground relations of reference, and *de re* attitude ascriptions, at least with respect to ordinary empirical discourse. Brandom himself is not an e-representational eliminativist, he is what we might call an 'e-representational reductionist.' In §4.2, I will argue the quantum pragmatism of people like Healey ([2017]) and Friederich ([2014]) is semantically anti-representationalist, when in fact pragmatism is more usefully aligned with the sort of metasemantic anti-representationalism that Brandom defends.

On the basis of this brief summary of the Brandomian view, let me demonstrate how our discursive practice latches onto the world, even though we have not first established word-world links of reference or representation. It is useful to think of the world constraining our discursive practices in largely the same way as it constrains any of our other practices. I might want to teleport my body just by thinking about the beach, but no matter how hard I try, the world will not cooperate. Similarly, every set of inferences that I make will ultimately be confronted with my experience of the world, and if I make bad inferences—i.e. if I am mistaken about whether an inference is deontic status-preserving— then the world will penalise me. For example, if I accept that a cricket ball is hard, but do not accept that it will injure me if it hits me at a high velocity, then I will be penalised for not protecting myself appropriately. So the world, being the way it is, is constantly impinging on our practices, discursive or otherwise. With respect to the discursive

practices, these impingement are constantly making us re-assess which material inferences are good and which are not. And this is what establishes the requisite link between the world and our discursive practices, even if this link cannot always be factorised into a set of more atomic word-world relations.

I will not say much about the semantic question for the Brandomian: how does i-representation/inferential valence distribute through a language? Ultimately, this question is beyond the scope of this paper, since my concern here is to set out a view about what it is that makes QM come to be about the world. This is squarely in the domain of metametasemantics and metasemantics; I adopt a pragmatist metametasemantics, and an inferentialist metasemantics. Now, although I believe that inferentialism provides a compelling metasemantic account of i-representation, it is not clear to me that this uniquely picks out an obviously appropriate semantics across the board. As will become inmportant in what follows, metasemantic anti-representationalism does not entail semantic anti-representationalism. As we have seen, in some cases i-representation coincides with e-representation. In the further subset of cases where e-representation is best understood in terms of reference, our canonical model-theoretic semantics (the generalisations of Tarski-semantics) is ideal. So an inferentialist metasemantics for, say ordinary empirical discourse, is perfectly well suited to a model-theoretic referentialist semantics (which is, of course, semantically representationalist).

But there may good reason to doubt that such a semantics works for, say, quantum field theory (QFT) . Even if we are able to establish that the practice of QFT is best understood (semantically) representationally, it seems unlikely that that representation relation should be understood in terms of reference (this is one of the motivations behind the 'math-first' approach to structural realism in (Wallace [2022])). My personal view is that the 'implication phase space semantics' developed by Dan Kaplan, and discussed in (Brandom [2021]), is a promising avenue to explore. But I will leave details of this project to future work. Luckily none of those details impinge on the metasemantic account of QM that I defend in this paper.

4 Anchored-inferentialist quantum mechanics

The central challenge to understanding what QM says about the world is the measurement problem, which Maudlin ([1995], p. 7) presents as an inconsistent triad:

1. The [quantum state] is complete, i.e. the [quantum state] specifies (directly or indirectly) all of the physical properties of a system.

2. The [quantum state] always evolves in accord with a linear dynamical equation (e.g. the Schrödinger equation).

3. Measurements of e.g., the spin of an electron always (or at least usually) have determinate outcomes.

In the literature on the philosophy of QM, 'interpretation of QM' generally means 'resolution of the measurement problem', as opposed to its broader meaning in metaphysics or semantics, of something like 'an account of what is described/represented/picked out by the expressions of the theory.' So, for example, the Everett interpretation resolves the measurement problem by denying Maudlin's third premise. But it does so without saying anything specific about what the ontology of quantum mechanics is—Everettian QM is consistent with, among many other descriptive ontologies, a wavefunction-realism (Albert [2013]; Ney [2021]), a density-operator realism (Wallace and Timpson [2010]), and a spacetime-qubit realism (Deutsch and Hayden [2000]).

Nonetheless, most Everettians are committed to the claim that there is some descriptive account that underpins their interpretation. But this descriptive assumption is one of the assumptions that underpins the measurement problem in the first place: drop this assumption, and the measurement problem does not even arise. This is, of course, precisely idea behind anti-descriptivist views such as those advanced in (Caves et al. [2002]; Friederich [2014]; Healey [2017]). All of these views therefore count as interpretations of QM, even though some of them do not purport to describe a world according the QM. So with this in mind, I distinguish between descriptive and non-descriptive interpretations of QM. Note that an interpretation of QM might be motivated by a metametasemantic or a metasemantic commitment, but it is ultimately underwritten by a semantic commitment: semantic representationalists take the quantum state to have e-representational content, while semantic anti-representationalists do not.[3]

I introduce Healey's non-descriptive interpretation in §4.1. Then, in §4.2, I assess how this approach fares in light of some criticism from Lewis ([2020]) and Wallace ([2020]). Together with my methodological disagreements, these criticisms motivate the alternative inferentialist view that I present in §5.

4.1 Healey's approach

Healey's interpretation is a prescriptive, pragmatist, non-descriptive interpretation of QM. According to this view, the role of the quantum state is to prescribe a course of action for agents using quantum mechanics, rather than to describe either the world or some agents' belief states. Of course, in general, nothing stands in the way of accepting that the quantum state is action-guiding *in virtue of* its descriptive capacity; what is central to prescriptive interpretations is that the prescriptive function of the quantum state is primary, and exhausts its role. Any ontic or epistemic residue, if present, is epiphenomenal (Healey denies that there is any ontic residue, but, as we saw in §2, other pragmatists might not).

The starting point of Healey's approach is the Born Rule, which is sometimes understood as an algorithm for generating a probability distribution over measurable magnitudes. Healey prefers an alternative reading of the Born Rule, which he attributes to Einstein, according to which it is understood as associating a probability distribution over magnitudes of some dynamical variable M.

3 Description is a form of semantic representation, so the denial of semantic representationalism entails the denial of descriptivism.

The difference is subtle, but important. On the former view, one first picks out a set of measurable magnitudes, and then applies the Born Rule indiscriminately to any and all of them. On the latter view, one reads the Born Rule as applicable to any dynamical magnitudes (which are not picked out in advance), but only delivering what we might call 'physical probability magnitudes' for a proper subset of those claims. For Healey, then, the Born Rule delivers a probability distribution over a proper subset of what he calls 'canonical magnitude claims' ([2017], p. 80):

Canonical magnitude claim: The value of a dynamical variable M on a physical system s lies in Δ.

Based on how QM gives rise to canonical magnitude claims, we can identify the formal structures and properties that QM trades in:

1. Pure quantum states: Each isolated quantum system s is assigned a pure quantum state ψ, which is a (normalised) ray in a Hilbert space equipped with a preferred basis. A superposition of pure quantum states is a pure quantum state.

2. Quantum Dynamics: Pure quantum states evolve linearly and unitarily.

3. Mixed quantum states: Each system, whether isolated or not, is assigned a quantum state ρ which is a density operator on a Hilbert space.

4. (Generalised) Born Rule: The probability of a canonical magnitude claim being true is a specific function of the (possibly mixed) quantum state associated with s, and the interval Δ.[4]

The basic idea here is that a physical system is assigned a quantum state, which is subject to some dynamical evolution law. Through the Born Rule we can associate, with a quantum state, a mathematical distribution that obeys the Kolmogorov axioms (call this a K-distribution) over canonical magnitude claims.

Any interpretation of QM needs to (i) demonstrate how to understand these structures in a way that avoids the measurement problem; and (ii) explain how, in virtue of these structures, QM comes to be about the world. There is a host of proposals for how to deal with (i)—these are just the interpretations of QM. Regarding (ii), though, the default view in the literature on the philosophy of QM is descriptivist: a Bell-style object naturalism about beables. On this view, quantum states are about beables: they ascribe properties to the entities represented by beables. So superpositions of states, and their dynamical evolution, are to be understood as describing how these properties distribute and evolve, in terms of how the entities represented by beables evolve. And so too is the Born rule: it makes claims about probabilities (however you choose to interpret them) associated with properties of the entities represented by beables described by the quantum state. The importance of beables to the canonical understanding(s) of QM can be seen by considering two plausible-sounding claims:

[4] For a statement of the Generalised Born Rule, see e.g. (Healey [2017], p 267).

1. If quantum theory had no beables, it could not make predictions or explain anything.

2. Quantum mechanics makes predictions and explains some things.

Contemporary approaches tend to accept both (1) and (2), and then use (2) to deny the consequent of (1), thus establishing the existence of (entities represented by) beables. By contrast, Healey accepts (2) but denies (1); we will identify the specific move that allows him to do this in §4.2. On this basis, he offers us an alternative way to make sense of the central components of the formalism of QM: quantum state (1-4), its dynamics (5) and the Born rule (6):

1. Non-Representation: Quantum states are never descriptive; they do not ever describe entities in some domain of the world.

2. No Quantum Beables: Beables exist; but these beables are not part of quantum theory.

3. Relational Objective assignment: The assignment of a quantum state to a system is relational, i.e. it depends on the situation of an actual or possible observer. But it is not subjective, i.e. given a choice of observer, there is a subject-independent fact about the correct quantum state assignment.

4. Decoherent Metasemantics: Putative magnitude claims derived from undecohered quantum state assignments are meaningless. But putative magnitude claims derived from suitably decohered quantum state assignments are meaningful. Decoherence—i.e. the delocalisation of the phase of a quantum state via dynamical interactions—is the mechanism for semantic content-conferral on to quantum mechanical claims.[5]

5. Unitary Dynamics: Pure states evolve in accordance with a unitarity-preserving dynamical equation.

6. Born Rule Applicability: The Born Rule is the correct specification for the probability of a dynamical variable P on a system s to have a value in some real interval Δ only when s is in a situation where application of quantum theory justifies the conclusion that the correct state assignment is a sufficiently decohered state.

I take these six claims to characterise Healey's approach, at least with respect to his deviation from the more standard (metasemantically) representationalist approaches. These claims are not all independent. No quantum beables, for example, follows from Non-representation, which in turn follows from Decoherent metasemantics. Similarly, Born Rule applicability follows from Decoherent Metasemantics. But I think there's much to be gained by being explicit when discussing Healey's assumptions, even if it risks being a little redundant.

[5] The actual mechanics of decoherence play no significant role in this paper, so I do not go into any detail about decoherence. All that matter in this paper is that decoherence is a mechanism for demarcating certain sorts of quantum states, whicha re given a special semantic treatment. For an overview of decoherence and its different roles in QM, see (Bacciagaluppi [2020]).

As presented, there is an immediate problem for Healey: canonical magnitude claims are descriptive—they assign properties to systems in the world—but quantum states are not. So quantum states cannot give rise to canonical magnitude claims, *a fortiori* to probability distributions over canonical magnitudes. What we need is a further claim about how canonical magnitude claims can be imbued with descriptive content. This is where Decoherent Metasemantics comes in. The Born rule can be used to associate a K-distribution over the values of such dynamical variables, but that does not immediately qualify that distribution as what we might call a 'physical probability distribution'. Healey is clear about this ([2017], pp. 214-215):

> Even though a thoughtless application of the Born rule to [an undecohered] quantum state would associate a number between 0 and 1 with [the canonical magnitude claim], an agent should not base partial belief in [that claim] on this number.

Healey's stance regarding what counts as a physical probability distribution, as opposed to a mere K-distribution, thus incorporates what Papineau ([1996], p. 238) calls the decision-theoretical link:

> We base rational choices on our knowledge of objective [(physical] probabilities. In any chancy situation, a rational agent will consider the difference that alternative actions would make to the objective probabilities of desired results, and then opt for that action which maximizes objective expected utility.

For Healey, the decision-theoretical link is established only by adding to a K-distribution the further semantic condition that the canonical magnitude claims over which the distribution is defined are *meaningful*. And a canonical magnitude claim about a system s is meaningful only if the quantum state ρ associated with s is suitably decohered. To see how this works, let us borrow an example from Healey ([2017], p. 211-212). Contrast the following two sets of canonical magnitude claims:

$C_1(t)$: The x-component of the moon's position at time t lies in D.

$C_2(t)$: The x-component of the particle's spin at time t is r.

The Born Rule, applied to the underlying states, gives rise to K-distributions over both sets of canonical magnitude claims. But, for Healey, the interpretation of these claims depends heavily on whether or not the underlying states are decohered. It is wildly implausible that the mixed state correctly attributed to the moon will be undecohered. Roughly speaking, this is because the position of the moon is constantly being 'measured' by its environment. More precisely, its state is constantly becoming entangled with its environmental state in such a way that Born Rule assigns an appropriately peaked probability distribution to its x-position that is dynamically robust according to some approximately classical equations of motion; this is the precise sense in which Healey talks of quantum states being 'suitably decohered.'[6]

The same cannot be said of a particle, in general. There will be some circumstances under which its associated quantum state assignment is an undecohered state, others in which, like the moon, it is decohered by its environment. In both cases, there is a K-distribution over a set of canonical magnitude claims, but only in the latter case should that distribution be taken to be action-guiding. To quote Healey ([2017], p. 214), '[b]lunty put, an agent should regard C_2 as devoid of empirical content when the particle has [an undecohered] quantum state.' For Healey, empirical content just is semantic content: e-representation just is i-representation. So to deny e-representation is to deny content.

Ultimately, our discursive practice has to make contact with the world, and, in particular, some of that discourse needs to be descriptive. We need to be able to say things that state facts by describing parts of the world, things like 'the moon is roughly a quarter of a million miles from the Earth'. On Healey's view, that descriptive role is played by the bits of vocabulary that exist outside quantum theory. This is a delicate and subtle point, and one that comes straight out of the pragmatist's playbook: our starting point, when understanding the world according to science, is not some pristine, abstract theoretical structure presented to us by God, that we need to somehow glue onto the world. As contemporary philosophers, we enter a world replete with discursive practices that have been evolving for centuries. Those practices had been making meaningful claims, even if, as a result of sometimes being mistaken about, and subsequently correcting, how we use these words, we have often had to change their meanings.[7]

QM, according to Healey, is a way of updating the meanings of the terms we already had a reasonably good grip on from previous physical theories, and moreover, which we understood as meaningful in virtue of *those theories'* descriptive capacities. QM is unique in the history of science, on this view, in not introducing new beables—recall, No Quantum Beables—but instead, in systematically updating the meanings of our pre-existing beables. Healey identifies this as the revolutionary aspect of QM. Ironically, what is methodologically revolutionary about QM is just how metaphysically conservative it is.

Decoherence plays an important semantic, but not metaphysical role. It is a physical process

6 Even more precisely, Healey ([2017], p. 214) identifies the following three conditions as sufficient for a state to count as suitably decohered: [Every mixed state $\rho(t)$] will define a set of vector [(pure)] states $|\psi_x(t)\rangle (x \in \mathbb{R}$ with several special features:

1. It is stable—if $|\psi_x(t_1)\rangle$ is an element of the set $\rho(t_1)$ defines at t_1, then $|\psi_x(t_2)\rangle$ is an element of the set $\rho(t_i)$ defines at t_2;
2. $|\psi_x(t)\rangle$ approximates a classical mechanical state in the following sense: the Born probability distributions it yields for x-components of position and momentum are each concentrated around precise values (x, p_x respectively) and are consistent with the corresponding single probability distribution derived from a *joint* probability distribution on a space of classical states for a system of precise but unknown position and momentum;
3. The classical state with values x, p_x obeys classical mechanical equations of motion.

7 As Ruetsche ([2012], p. 7) puts it: 'It is essential to understand that, in the present sense of interpretation, the vast majority of the theories philosophers talk about are already partially interpreted. Otherwise they wouldn't be theories of physics. These theories typically come under philosophical scutiny already having been equipped, by tradition and by lore, with an interpretive core almost universally acknowledged as uncontroversial.

that happens in the world, on the basis of which we can specify exactly how QM changes the meanings of these beables. There is no spooky metaphysical transition between quantum goings-on and classical goings-on. And there is no sense in which pre-quantum beables need to be derived from quantum beables. The material inferences that confer meaning onto these beables are derived from the practice of pre-quantum theories and updated by QM. These inferences are thus better described as 'meaning-modifying' than 'meaning-conferring.'

This observation allows us to characterise the specific brand of inferentialism that Healey's view embraces. Recall that, as a metasemantic thesis, inferentialism says that it is the inferential profile of expressions determines their meaning. But how should we understand the scope of this claim? In other words, is there a restriction on the sorts of vocabulary that should be given this metasemantic treatment, or should we be committed to this metasemantics across the board? Broadly speaking, there are two types of inferentialist position, distinguished by the amount of vocabulary that is treated inferentially (The terminology is based on, but not identical to (Chalmers [2012], p. 2)):

Strongly anchored inferentialism: '[T]he contents of some primitive concepts are determined noninferentially,... and that the content of all other concepts are determined at least in part by their inferential relations to these concepts.'

Pure inferentialism: '[T]here are no [non-inferentially meaningful] concepts...: the content of every concept is determined only by its place in the web, and the web as a whole is characterized only by its abstract structure.'

Strongly anchored inferentialism, gives the anchors a completely representationalist metasemantics, and the other concepts a inferentialist metasemantics. On this view, the anchors have a self-standing representational directedness which grounds their meaning. This meaning then flows to the other, non-anchored, concepts via material inferential links. In Pricean terms, the anchors stand in meaning-conferring e-representation relations to parts of the world, and, in virtue of this, confer meaning, and an i-representational profile to the non-anchored concepts.

Healey's view requires two modifications: (i) the anchors need to be less brittle and (ii) the non-anchors are meaningless. After all, on his view, the main semantic consequence of QM is a systematic updating of the meanings of the beables that we already had a grip on from pre-quantum scientific practice. The strongly anchored inferentialist does not have the resources to allow the non-anchors to semantically influence the anchors, but not vice-versa. So Healey's view is form of what we might call a 'weakly anchored inferentialism', according to which (i) the representational directedness of anchors towards worldly entities is not self-standing, and is only intelligible in combination with considerations of their inferential links to other anchors as well as non-anchors and (ii) non-anchors are themselves meaningless. Anchored inferentialisms, sit at various points on a spectrum, one end of which is a pure metasemantic representationalism, the other of which is a pure metasemantic inferentialism.

The order of semantic explanation is something like the following: pre-quantum beables describe parts of the world, but they do this imperfectly. QM comes along and prescribes, via the Born Role, an objective but relational K-distribution over the properties of these beables. This distribution is action-guiding, therefore a physical probability distribution when the quantum state is suitably decohered. But QM also has some structure that looks like it might be descriptive even though it does not say anything about the pre-quantum beables. These claims merely modify the e-representational profile of the anchors, via the material inferences encoded in QM. Decoherence is the mechanism for identifying when an inference is e-representation modifying.

4.2 Assessing Healey's approach

In this subsection, I discuss some of the criticism that has been directed towards Healey's approach. In particular, I introduce two challenges that it struggles with. This is not intended as a way of dismissing his approach, whose central proposal—that we should understand the content of quantum-mechanical claims in terms of their inferential, not their representational, significance—I strongly endorse. Instead its dialectical purpose is to motivate my own alternative approach.

4.2.1 The descriptivist challenge

In §3, I presented pragmatism as a metametasemantic position which informs a specific choice of semantic machinery, whose mechanics is specified by the appropriate metasemantics. Healey understands inferentialism as entailing a negative semantic claim: whatever the appropriate semantic machinery for pure QM is, it is not representation. Lewis ([2020], p. 172) raises the question of how such a claim can be justified. Here is a quick reconstruction of what I take Healey's justification to be: (i) inferentialism entails the denial of metasemantic representationalism, and (ii) semantic representationalism entails metasemantic representationalism. Together with an endorsement of inferentialism, (i) and (ii) constitute a modus tollens against semantic representationalism.

Healey's weakly-anchored inferentialism requires a distinction between the presecriptive content of QM and the descriptive content of other elements of the broader practice. This gives rise to a new problem, also noted by Lewis ([2020], p. 173): '[t]he worry, then, is that the distinction between the prescriptive content of quantum claims and the descriptive content of non-quantum claims is not supported by the inferentialist account of content.' Our scientific discursive practice of using QM, at least some of the time, includes predictive claims. That's how we know that QM is a good theory. But predictions of experimental outcomes are descriptions of possible (changes in the properties of) experimental apparatuses. So any account of a physical theory needs to account, at the very least, for the descriptive nature of confirmation; call this the descriptivist challenge.

Healey openly acknowledges the descriptivist challenge ([2012], pp. 740–741):

> It is critical for the present approach to have available non-quantum descriptions of outcomes of quantum measurements... If one could not express the result of a measurement [using pre-quantum beables], then the Born rule could acquire no empirical

> support from measurements and should have little or no reason to believe quantum theory.

Lewis' point is that Healey-style inferentialist struggles because semantic anti-representationalism cannot meet the descriptivist challenge. So Healey's view is either inadequate, or it must deny the descriptivist challenge in the first place and offer in its place a theory about why we erroneously thought confirmation requires description. We can infer from the quote above what Healey's response would be: this dilemma only arises under the further assumption that the descriptive claims associated with our scientific discourse which includes QM must come entirely from the purely quantum mechanical component of that discourse. And Healey denies this.

On Healey's view, if we were to extract from our scientific practice just the pure quantum mechanical claims, i.e. claims about undecohered quantum states, we would find that they do not latch on to the world, because the inferences that they give rise to are not licensed by the norms of our practice: if we were to try to use them, the world would eventually penalise us. So the descriptive capacity of our scientific practice, which includes QM, has to come from somewhere else. Healey is thus led to anchored inferentialism, with the pre-quantum beables serving as anchors. But while this solves one problem, it does so at the expense of creating an even bigger one.

4.2.2 The beables challenge

The descriptivist challenge sets up an inconsistent triad for the inferentialist:

`SAQS`: Semantic anti-representationalism about quantum states: quantum states do not stand in relations of representation to parts of the world.

`MAQS`: Metasemantic anti-representationalism about quantum states: the meanings of quantum states are not grounded in relations of representation.

`QMDA`: Quantum mechanical practice is discursively autonomous.

Healey resolves this inconsistency by denying `QMDA`. His denial of QM's discursive autonomy, and subsequent adoption of anchored inferentialism shifts the metasemantic focus from QM states on to the anchors. The immediate question, raised by both Lewis ([2020]) and Wallace ([2020]) is about the beables: since the discourse around QM states, whether decohered or not, involves the same sorts of expressions (recall $C_2(t)$ from §4.1), why should an expression be treated as a beable (or as naming a property of what is represented by a beable) when it arises from a decohered quantum state assignment, given that, under an undecohered assignment, that same expression is not to be treated as a beable? Call this the beables challenge.

Lewis ([2020]) suggests a Pricean answer: i-representation endows all canonical magnitude claims, whether they are made on the basis of decohered or undecohered quantum state assignments, with semantic content. But that semantic content only becomes descriptive content once decoherence has endowed a quantum state with an e-representational profile. So the beables are

just picked out by the contingent coincidence of i-representation and e-representation, where that coincidence is encoded in the appropriate quantum state assignment. I think this is the correct way to think about inferentialist approaches to quantum mechanics, and will have more to say about it in §5. Unfortunately for Healey, this option is not available to the anchored-inferentialist.

As we saw in §4.1, for Healey, semantic content is grounded entirely in e-representation. A complete theory of the world is descriptive, and at least at this point in our scientific development, our best descriptive theory of the world is one whose beables are entirely pre-quantum. Although this may change in the future, it will not change in virtue of any features of our extant QM because QM merely modifies the e-representational profile of pre-quantum beables. Indeed, the central difference between pre-and post-quantum theories is the e-representational profile of these beables. The material inferences that QM brings have no self-standing authority to confer meaning. If they did, then there would be the possibility that QM had its own beables. But Healey explicitly and repeatedly denies this.

With this in mind, one can see that Healey does not have the resources to separate meaning conferral by i-representation from meaning conferral by e-representation in the way that Lewis suggests. Meaning conferral is *always* metasemantically e-representationalist, it is just that the e-representeds are systematically modified by QM. Of course, this is wholly compatible with Healey's metasemantic anti-representationalism, because the latter applies only to quantum states, and QM is not discursively autonomous. To put things another way, on Healey's view, the pre-quantum beables have their e-representational profile established before considerations of QM. On being incorporated into our discursive practice, quantum states only come to be about those non-quantum magnitudes when they are sufficiently decohered; otherwise quantum state assignments are meaningless.

The problem with Healey's view regarding pre-quantum beables is not that he has no way of demarcating them. He does: they are just the ordinary beables of classical mechanics. It is, rather, that there is no explanation of how our practice of QM, which, historically speaking, is continuous with the practice of classical physics, somehow ended up, despite appearances, being entirely non-e-representational on its own terms. The beable challenge is therefore one of explanation, not demarcation. I think the descriptivist and beable challenges are the most pressing challenges to inferentialist interpretations of QM. Healey's view deals with the former, but only at the expense of foundering on the latter.

5 Pure-inferentialist quantum mechanics

An interpretation of QM is only complete once it makes a commitment to either semantic representationalism or semantic anti-representationalism. Healey's view counts as an interpretation because his metasemantic anti-representationalism motivates his semantic anti-representationalism. But, as I stressed in the last section, *qua* pure-inferentialist, I deny this link. So, an austere version of the pure-inferential approach, which embraces only a metasemantic anti-representationalism, is

strictly speaking neutral with respect to semantic representationalism. And as a result it does not qualify as an interpretation of QM on its own.

The specific version of pure-inferentialism that I adopt takes a stand on semantic representationalism: it embraces the *additional* commitment to decoherence as a mechanism for e-representation conferral. This is not mandatory. Adopting this additional commitment in §5.1, I introduce my preferred *interpretation* of QM, the 'decoherent pure-inferentialist interpretation of QM' (DPIQM).[8] In §5.2, I compare DPIQM with Healey's interpretation with respect to the challenges identified in §4.2.

5.1 The decoherent pure-inferential interpretation of quantum mechanics

It is difficult to decide how much one needs to say in order to have engaged in the grandiose project of Proposing a New Interpretation of QM. Most other interpretations have accumulated book-length defences which demonstrate how the interpretation deals with not only the measurement problem, but also a host of other topics, including quantum probabilities, no-go theorems like Bell's theorem, the Bell-Kochen-Specker theorem, and even crucial experiments to test the interpretation. But constrained as I am by space, I will simply present the core tenets of my view here here, and demonstrate (i) how adopting these tenets dissolves the measurement problem, and (ii) how it improves upon Healey's interpretation. I will have to delegate the responsibility of showing how these tenets entail (or, more broadly, motivate) accounts of probability, non-locality and contextuality, and other no-go theorems to future work.

The pure-inferential approach to QM starts with a pragmatist metametasemantics, and follows Brandom in identifying i-representation as the piece of semantic machinery to focus on. Brandom identifies the circumstances under which the inferential profile of an ordinary empirical claim can be understood, by interlocutors, as grounding a *de re* attitude ascription, and consequently an i-representational profile; this profile can be compared systematically with e-representational claims until, in some cases, these profiles are discovered to co-vary. In such cases, the semantic content of a claim is descriptive. Nothing in this account precludes its extension to QM. In fact, QM is even better-suited to this view than ordinary empirical discourse, because decoherence explains why certain expressions have the e-representational profile that they do: the content of those claims is sufficiently systematically linked to physical magnitudes that those claims can ground the decision-theoretical link between K-distributions and actions. I will say more about this in §5.2.2.

The core of the pure-inferentialist approach consists in the following commitments:

[8] Even at this stage, one might argue that a unique interpretation is not fixed, given, as Bacciagaluppi ([2020]) stresses, the different roles ascribed to decoherence by different interpretations. But I think this is actually a delicate question, since Bacciagaluppi's analysis takes place within a metasemantically representationalist paradigm. It is unclear to me that these different roles come apart for the metasemantic inferentialist. However, I will not make that assessment in this paper. Readers who think that choosing decoherence as a mechanism for e-representation-conferral still underdetermines an interpretation should read my interpretation as one, rather than the only, option.

Metasemantic Anti-representationalism without Semantic Anti-representationalism: Quantum claims do not derive their meanings from a relation of representation (although some expressions might still be representational).

Inferentialist Metasemantics: Quantum claims derive their meanings (i.e. i-representational profile) from the material-inferential web in which they are embedded. Quantum claims are the explicitation of the tacit inferential commitments of agents who use quantum mechanics.

To turn this into an interpretation, we need a further claim about how to understand semantic e-representation. I adopt the following:

Decoherence is a Matching Criterion: If a system is assigned a highly decohered quantum state, then that signals an alignment of i-representation and e-representation of the central terms of its canonical magnitude claims. In these circumstances, the quantum state is descriptive, and its central terms are (quantum) beables.

With these commitments on the table, we can focus on what the DPIQM says about the central components of the formalism of QM:

1. i-representation: Quantum states can be descriptive; but they earn their descriptive credentials by having their i- and e-representational profiles align.

2. Quantum Beables: Beables exist; some of these beables are novel posits of quantum theory, and all beables are ultiamtely quantum.

3. Relational Objective Assignment: The assignment of a quantum state to a system is relational, i.e. it depends on the situation of an actual or possible observer. But it is not subjective, i.e. given a choice of observer, there is a subject-independent fact about the correct quantum state assignment.

4. Decoherent Representational Profile: Putative magnitude claims derived from any quantum state assignments are meaningful because of the material-inferential web in which those claims are embedded. In addition, magnitude claims derived from suitably decohered quantum state assignments can be e-representational. Decoherence is the mechanism for e-representation-conferral on to quantum mechanical claims.

5. Unitary Dynamics: Pure states evolve in accordance with a unitarity-preserving dynamical equation.

6. Born Rule Applicability: The Born Rule is the correct specification for the probability of a dynamical variable P on a system s to have a value in some real interval Δ only when s is in a situation where application of quantum theory justifies the conclusion that the correct state assignment is a sufficiently decohered state.

My denial of semantic anti-representationalism without a consequent denial of metasemantic anti-representationalism underpins my central points of disagreement with Healey. Thus, for me, (1) i-representation replaces Non-Representation, (2) Quantum Beables replaces No Quantum Beables and (4) Decoherent Representational Profile replaces Decoherent Metasemantics. But there are still some points of agreement between us: (3) Relational Assignment, (5) Unitary Dynamics and (6) Born Rule Applicability.

DPIQM counts as an interpretation of QM because it resolves the measurement problem. But the mechanics of this resolution is very different from Healey's, which dissolves the measurement problem by denying semantic representationalism. DPIQM does not deny semantic representationalism. Instead, it specifies the circumstances under which one should be a semantic e-representationalist about the quantum state: when the correct quantum state assignment is a suitably decohered one.

What this means is that, according the the DPIQM, if an undecohered quantum state is assigned to a system, then it is not e-representational. For such states, the measurement problem is dissolved just as it is on Healey's interpretation. If, on the other hand, the system is assigned a state that evolves in such a way that the Born Rule spits out a K-distribution over canonical magnitude claims whose peaks are sufficiently robust under some classical dynamical evolution equations that one can establish the decision-theoretical link, and thus use QM to guide one's actions (i.e. it is suitably decohered), then that state can be treated as semantically e-representational. And, as is well-known from the literature on decoherence (see e.g. (Wallace [2012], Ch. 3), the dynamical equations of QM, for appropriately decohered quantum states, do give rise to precisely these sorts of K-distributions. In other words, the dynamics of QM, as it applies to the states of rapidly decohering systems, does e-represent a plurality of approximately dynamically-isolated, classically evolving magnitudes. Worlds, if you will. So the measurement problem is avoided by denying that all experiments have determinate single-outcomes. This is a version of the Everett interpretation. Given its alternative metasemantic foundations, it is not the version of Everett defended by the contemporary Oxford School (as exemplified in (Saunders [1993]; Wallace [2012]). I will, unfortunately, have to leave an assessment of the relative merits of these two versions of the Everett interpretation to future work, since the goal of this paper was to build a pure-inferentialist interpretation and compare it to its prominent anchored-inferentialist alternative.

5.2 Decoherent pure-inferential quantum mechanics vs. Healey's interpretation

5.2.1 The descriptivist challenge

In §4.2, I presented the descriptivist challenge as posing an inconsistent triad for the inferentialist. Healey resolves it by denying that the practice of QM is autonomous. But this leads to an unattractive form of weakly-anchored inferentialism. I think there is a better way. Recall that I mentioned, somewhat in passing in §3, that Brandom himself, although metasemantically anti-

representationalist is, in fact, semantically representationalist about ordinary empirical discourse. He vindicates the use of reference-talk by grounding such talk in inferential practice. What this means is that there is no reason to believe that metasemantic anti-representationalism is inconsistent with semantic representationalism.

So there is another route to resolving the inconsistent triad: to accept both `MAQS` and `QMDA`, but to deny `SAQS`. I think that this is methodologically preferable, because it opens the door to the possibility of a thoroughgoing, pure inferentialism applicable to any self-standing scientific practice, not just QM. Of course, in denying `SAQS`, I now have to account for how a metasemantic anti-representationalist can ground the semantic representational profile of quantum states, i.e. how we can ground the descriptive capacities of QM without having to rely on external referential crutches like pre-quantum beables. Our new, achievable, goal is to demonstrate how representation can be grounded in inference. If we can establish this, then we will meet descriptivist challenge.

My (Peter) Lewisian strategy here comprises two steps. The first is to establish that quantum state assignments are i-representational (therefore meaningful) by demonstrating that they encode *de re* propositional attitudes of speakers. The second is to explain how decoherence supplies a i-representational claim with an e-representational profile. Together, these steps demonstrate how semantic representationalism about quantum states (when appropriate) is grounded in their inferential profile.

Recall from §3, that what it is for me to make a *de re* attribution to an interlocutor, on the basis of a claim that they make, is for me to adopt a specific view about how that claim updates the deontic score I attribute to them, on the basis of the material-inferential norms I take to be a tacit component of our discursive practice. So if my interlocutor S says 'the quantum state of system s is ρ', then I adopt a commitment to the claim 'S believes/claims[9] of s that its quantum state is ρ'. I then attribute to S a new deontic status, that includes (i) commitments of the form 'S believes/claims of s that the probability of the canonical magnitude claim C is x'; (ii) entitlements of the form 'S is rational to wager y pounds on the truth of C. I myself might adopt an updated commitment about s, for example, if I take S to be a competent user of QM, and use S's claim as a premise in my own reasoning. Versions of these three moves (i.e. attributing commitments to S, attributing entitlements to S and undertaking commitments myself) are available to anyone whose interlocutors make an explicit quantum state assignment. So, on the Brandomian view, any quantum state ascription allows for a *de re* attitude ascription. As we saw in §3, this is sufficient to establish a relation of i-representation between an expression (in this case a quantum state) and its semantic content (i.e. its meaning). But just as with 'unicorn', the presence of i-representational content does not guarantee e-representation.

So in the second step, we use decoherence to establish the e-representational profile of a quantum state ascription. The account I provide here is imported similar to Healey's. But the crucial difference between our views is that, in denying that decoherence is content-conferring,

9 Talk of belief ascriptions makes some people queasy; however nothing turns on the ascription being of a belief, rather than, for example, a claim or an expression of a commitment.

and providing, as I have, an alternative account of content-conferral, my view does not have to rely on external referential crutches to establish the beables' e-representational profile. Recall that, for Healey, only suitably decohered quantum states assignments are about antecedently e-representational pre-quantum beables. Decoherence thus establishes a sort of matching, between quantum expressions and descriptive pre-quantum beables. For Healey, undecohered quantum states are meaningless.

Part of the problem with views that say that certain sentences are meaningless is that, often, they seems to fly in the face of an accepted and useful conception of 'meaning'. I can perfectly well (or certainly well-enough) understand what claims such as 'God is omniscient' or 'unicorns can fly' mean, even if I do not believe that the central terms refer. Similarly I can understand what 'the x-component of the particle's spin at time t is r' means, even if I subsequently trust that QM gives me reasons not to assent to it.

Recall the canonical magnitude claim $C_2(t)$ that we discussed in §4.1. According to Healey, $C_2(t_0)$ is meaningless if the associated state ρ_{t_0} is undecohered, because the inference from ρ_{t_0} to $C_2(t_0)$ is materially invalid. This material invalidity is the result of the Born Rule not assigning a K-distribution over canonical magnitude claims which is sufficiently peaked at one of them. But this is too quick. Compare Healey's meaninglessness claim with the following: 'The horse is brown' is meaningless because the inference from 'the horse is green' to 'the horse is brown' is materially invalid. Something has gone wrong here. That 'the horse is brown' is a conclusion of a bad inference does not mean that it is not part of several other good material inferences. And it is *those* inferences that are meaning-conferring. In other words, we must consider *all* the material-inferential links associated with a claim in order to ascribe to it a meaning. The only circumstance under which a claim should be treated as meaningless is when the *only* material-inferential links it stands in to other claims are invalid. And this is plainly not the case for $C_2(t_0)$, and indeed most honestly-asserted canonical magnitude claims. So Healey's view, as stated, requires, but cannot sufficiently justify the radically revisionary semantics that it posits.

DPIQM does not posit a revisionary semantics: all quantum states are meaningful, though they might not be semantically e-representational. But under some circumstances, a mixed state assigned to some subsystem might evolve in such a way, relative to the environment with which it is entangled, that the probability distribution over certain canonical magnitude claims associated with that subsystem becomes dynamically robust in the sense elaborated in footnote 4.1. The decoherent state ascription can then be read as describing (a multiplicity of) quasi-classically evolving parameters. Crucially, these parameters co-vary with some worldly magnitudes to which we have empirical access. This is precisely the covariation that characterises e-representation. So not only does decoherence delineate the circumstances when certain quantum magnitudes are e-representational, it explains why this is the case: the underlying quantum dynamics attributes to certain sorts of special quantum states, parameters which covary with directly empirically accessible, macroscopic magnitudes.

Decoherence is, therefore, not a mechanism of content-conferral. Content is conferred by the

material-inferential norms which manifest themselves in the way that interlocutors update each others' (and their own) deontic scorecards. As a result, all dynamical claims about the evolution of quantum states (in particular mixed states) are meaningful. Decoherence is then the quantum-theoretic mechanism of alignment of i-representation (i.e. the meanings of magnitude assignments to quantum parameters) and e-representation (i.e. the covariation of these parameters and empirically accessible magnitudes). Note that decoherence does not establish an alignment, or matching, between independently-established relations. Instead, decoherence takes i-representation and adds to it an e-representational profile. This is how the DPIQM meets the descriptivist challenge.

5.2.2 The beables challenge

In the course of meeting the descriptivist challenge, we invoked exactly the right machinery to solve the beables challenge as well: decoherence as a matching mechanism. For Healey, decoherence serves to make quantum state assignments meaningful only to the extent that it modifies the meanings of the pre-quantum beables. This was, of course, the source of the beables challenge.

The DPIQM, on the other hand, reverses the direction of matching. There are no antecedently e-representational pre-quantum beables. If we treat QM as complete, then decoherence establishes the circumstances under which beables, indeed *quantum* beables, exist: whenever the contingent quantum dynamics gives rise to parameters which match appropriately with empirically accessible magnitudes. Nothing in this account relies on those empirically accessible magnitudes being, in any sense, non-quantum. From an epistemological perspective, we might make a dynamical assumption that their quantum nature (i.e. deviations in the magnitudes which characterise those beables independently of their entanglement with the system s under consideration) is negligible. But ultimately, QM still applies to them: they can be assigned a mixed state with respect to some even bigger environmental system, and their (quasi-classical) dynamics can also be explained by QM.

The DPIQM's resolution of the beables challenge is as follows: there are only quantum beables, so there is no beables challenge. More importantly, DPIQM explains why there is no beables challenge: QM on its own has the resources (in principle), for any choice of physical system, to explain how the quantum dynamics of that system gives rise to the parameters whose robust quasi-classical evolution characeterises them as (for all practical purposes) classical/non-quantum beables. We were correct to treat those beables as e-representational. We were wrong to think that those beables had to be treated as an e-representationalist metasemantic starting point.

6 Conclusion

The current state of the dialectic between the pragmatist and the realist over quantum mechanics pits two extreme positions against each other: either the quantum state is never descriptive or it always is. The goal of this paper was to demonstrate, by construction, the existence of a pragmatist middle-ground: pure-inferential QM. According to this approach, we should judge the descriptive

capacities of QM by determining whether description is an inelimnable component of what makes it the case that QM is about the world. If quantum states should be seen as non-descriptive, that should be the conclusion, not the premise, of an argument.

I began by setting out the standards for what makes a discursive practice descriptive. To do so, I invoked Price's distinction between i-representation and e-representation, and argued that a descriptive practice was one whose (central) linguistic expression had an e-representational profile. On the standard metasemantically representationalist views that underpin many realist interpretations of physical theories, e-representation grounds i-representation.

I introduced Brandomian inferentialism as a metasemantically anti-representationalist proposal, according to which linguistic expressions could have an i-representational profile without necessarily having an e-representational one. I used this to develop a pure-inferentialist alternative to Healey's anchored inferentialist view. I introduced a specific version of the view, the DPIQM, which uses the quantum mechanical phenomenon of decoherence to establish the requisite link, in the appropriate circumstances, between i-representation and e-representation for quantum states. This is what allows the view to occupy the middle-ground between Healey's anti-descriptivism about the quantum state, and, say, the contemporary Everettian's descriptivism. I then identified two challenges, the descriptivist challenge and the beables challenge, both of which pose problems for Healey's view, but not the DPIQM.

References

Albert, D. Z. [2013]: "Wave Function Realism", *The wave function: Essays on the metaphysics of quantum mechanics*, pp. 52–7.

Bacciagaluppi, G. [2020]: "The Role of Decoherence in Quantum Mechanics", in E. N. Zalta (*ed.*), *The Stanford Encyclopedia of Philosophy*, Fall 2020, Metaphysics Research Lab, Stanford University.

Bächtold, M. [2008]: "Interpreting Quantum Mechanics According to a Pragmatist Approach", *Foundations of Physics*, **38**, pp. 843–68, ISSN: 1572-9516.

Bell, J. S. [2004]: *Speakable and Unspeakable in Quantum Mechanics: Collected Papers on Quantum Philosophy*, Cambridge university press.

Bohr, N. [2010]: *Atomic Physics and Human Knowledge*, Courier Dover Publications.

Brandom, R. [1994]: *Making It Explicit: Reasoning, Representing, and Discursive Commitment*, Harvard University Press, ISBN: 978-0-674-54330-0.

Brandom, R. [2009]: *Articulating Reasons: An Introduction to Inferentialism*, Harvard University Press.

Brandom, R. [2021]: "On the Structure of Reasons: Pragmatics, Semantics, and Logic".

Caves, C. M., C. A. Fuchs, and R. Schack [2002]: "Quantum Probabilities as Bayesian Probabilities", *Physical Review A*, **65**, p. 022305, ISSN: 1050-2947, 1094-1622, arXiv: quant-ph/0106133.

Chalmers, D. [2012]: "Nineteenth excursus: Inferentialism and analyticity", *Constructing the world. Extended edition. Internet: David Chalmers*.

Dasgupta, S. [2020]: "Undoing the Truth Fetish: The Normative Path to Pragmatism", *Unpublished manuscript*.

Davidson, D. [1967]: "Truth and Meaning", in *Philosophy, Language, and Artificial Intelligence: Resources for Processing Natural Language*, Springer, pp. 93–111.

Davidson, D. [2001]: *Inquiries into Truth and Interpretation: Philosophical Essays Volume 2*, Clarendon Press.

Deutsch, D. and P. Hayden [2000]: "Information Flow in Entangled Quantum Systems", *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, **456**, pp. 1759–74.

Friederich, S. [2014]: *Interpreting Quantum Theory: A Therapeutic Approach*, Springer.

Healey, R. [2012]: "Quantum Theory: A Pragmatist Approach", *The British Journal for the Philosophy of Science*, **63**, pp. 729–71, ISSN: 0007-0882, JSTOR: 23356442.

Healey, R. [2017]: *The Quantum Revolution in Philosophy*, Oxford University Press.

Lewis, P. J. [2020]: "Quantum Mechanics and Its (Dis) Contents", *Scientific realism and the quantum*, pp. 168–82.

Maher, C. [2012]: *The Pittsburgh School of Philosophy: Sellars, McDowell, Brandom*, Vol. 42, Routledge.

Maudlin, T. [1995]: "Three Measurement Problems", *Topoi*, **14**, pp. 7–15.

Mermin, N. D. [2014]: "Physics: QBism puts the scientist back into science", *Nature*, **507**, pp. 421–3.

Ney, A. [2021]: *The World in the Wave Function: A Metaphysics for Quantum Physics*, Oxford University Press.

Papineau, D. [1996]: "Many Minds Are No Worse than One", *The British Journal for the Philosophy of Science*, **47**, pp. 233–41, ISSN: 0007-0882, 1464-3537.

Price, H. [2011]: *Naturalism without Mirrors*, Oxford University Press.

Price, H., S. Blackburn, R. Brandom, P. Horwich, and M. Williams [2013]: *Expressivism, Pragmatism and Representationalism*, Cambridge University Press.

Ruetsche, L. [2012]: *Interpreting Quantum Field Theory*, Oxford University Press, Oxford.

Ruyant, Q. [2023]: "Consistent histories through pragmatist lenses", *Studies in History and Philosophy of Science*, **98**, pp. 40–8, ISSN: 0039-3681.

Saunders, S. [1993]: "Decoherence, Relative States, and Evolutionary Adaptation", *Foundations of Physics*, **23**, pp. 1553–85, ISSN: 0015-9018, 1572-9516.

Sellars, W. [1954]: "Some Reflections on Language Games", *Philosophy of Science*, **21**, pp. 204–28.

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

Wallace, D. [2020]: "The Quantum Revolution in Philosophy", *Analysis*, **80**, pp. 381–8.

Wallace, D. [2022]: "Stating Structural Realism: Mathematics-First Approaches to Physics and Metaphysics", *Philosophical Perspectives*.

Wallace, D. and C. G. Timpson [2010]: "Quantum Mechanics on Spacetime I: Spacetime State Realism", *The British Journal for the Philosophy of Science*, **61**, pp. 697–727, available at <https://doi.org/10.1093/bjps/axq010>.