

Investigating quantum mechanics

Tushar Menon

Department of Linguistics and Philosophy,
Massachusetts Institute of Technology*

ABSTRACT

This paper rejects a widely-held but ultimately untenable semantic presupposition in the philosophy of quantum mechanics: that the significance of the modal vocabulary of quantum mechanics (including quantum state ascriptions) is ultimately to be understood in terms of denotation. By instead invoking a Sellarsian approach to modal vocabulary, and a Brandomian approach to meaning-conferral, this paper reconciles three apparently incompatible readings of the quantum state—nomological-modal, pragmatic-advisory, and descriptive-denotational. In doing so, it allows us to hold on to the significant insights that each of these readings provides into the worlds modeled by quantum mechanics.

- 1 *Introduction*
 - 2 *Two familiarization strategies*
 - 3 *Denotational vocabulary and modal vocabulary*
 - 4 *The modal classical state*
 - 5 *The modal quantum state*
 - 6 *Meaning-conferral and Q-vocabulary*
 - 7 *The world according to quantum mechanics*
 - 8 *For inimical to irenic*
-

1 Introduction

Quantum theories are difficult, and insights into the worlds they model are hard-won. The quantum state tells us about values that some physical magnitudes can and cannot take, which suggests a nomological-modal reading. The quantum state is the central dynamical entity of quantum theories, which suggests a descriptive reading. The quantum state is a repository of objective statistical facts about physical systems, which suggests an advisory reading. Frustratingly, these readings sit uncomfortably together, and interpretations of the quantum state typically give one of them up.

The goal of this paper is to take seriously all of these insights about the quantum state, set up a framework for their reconciliation, and in doing so offer a diagnosis of the source of their apparent incompatibility. That diagnosis is that incompatibility is forced on us only if we adopt an interpretative strategy that I call *derivative*. A diagnosis is most helpful when it is accompanied by a remedy. I propose one: we should instead adopt an interpretative strategy that I call *investigative*.

To interpret is to familiarize. An interpretation of a physical theory might tell us what its models purport to describe. But that endeavor is only successful to the extent that we're already familiar with whatever the objects of those purported descriptions are. The standard familiarization strategies are derivative: they attempt to derive, or otherwise inherit, familiarity from antecedently familiar theories.

An alternative, and in my view superior, familiarization strategy is investigative: to figure out how we made familiar a theory we now treat as familiar, and then to do *that* to the unfamiliar theory. Familiarization strategies can be given either a metaphysical or a semantic gloss; contrast 'chances are hypothetical relative frequencies' with "chance' means (or is synonymous with) hypothetical relative frequencies. This paper treats the project of interpreting QM as a semantic project: to interpret QM is to familiarize its claims, as articulated in its own distinctive vocabulary, call it *Q-vocabulary*.

The most important standard of adequacy for an interpretation of QM is that it deal with the Measurement Problem. There is a well-established tradition of understanding the Measurement Problem as primarily a semantic problem: see e.g. (Barrett [2019], p. 3), (Myrvold [2022], p. 62–63), (Muller [2023], p. 33). The interpretative strategy in this paper, accordingly, is to familiarize the unfamiliar elements of the Q-vocabulary using an investigative strategy. This includes, but is not limited to, making sense of superpositions of macroscopically distinct states. On this approach, metaphysical considerations are not dismissed, but they are downstream: claims about what the world would be like if QM were true will follow from claims about what Q-vocabulary can be used to assert. It will take some work, and the trajectory will be somewhat non-standard, but we will end up with an answer to the question: what is the world

like according to QM?

The argument is structured as follows. In §2, I argue that to familiarize unfamiliar elements of Q-vocabulary, we should adopt an investigative familiarization strategy. In particular, a *pragmatist* version, according to which we do with the unfamiliar vocabulary whatever it was we did to familiarize the familiar vocabulary. In §3 I claim that to understand the semantic role of the quantum state—which mediates probabilistic, and therefore modal claims—on the pragmatist investigative strategy, we must understand the broader semantic norms of modal vocabulary deployment, and their relationship to denotational vocabulary deployment.¹ According to the Sellarsian strategy advocated in this section, modal vocabulary is treated not as itself denotational, but as making explicit facts about the semantic framework that allows for the denotational deployment of other vocabulary.

In §4, I focus on the classical state which, like the quantum state, codifies modal structure. But it can *also* be given a denotational reading. This is a contingent fact about classical mechanics. If we prise apart its framework-explicitating function from its denotational function, we can avoid the pitfalls of a denotational reading of the quantum state inherited uncritically from the habit of reading the classical state denotationally. In §5, I argue that we can import the norms of denotational deployment of some vocabulary element from classical mechanics to QM without thereby inheriting their meanings. I demonstrate that we can understand the quantum state as meaningful insofar as it codifies (only) the modal content of denotationally-deployed claims such as ‘for system S in quantum state ρ , the value m of magnitude M is x units.’ without, in addition, having to treat the quantum state as itself denotational.

Having first established the denotational significance of some expressions of Q-vocabulary, independently of their meanings, in §6, I explore options for meaning-conferral norms. On a Sellars-Brandom inferentialist approach, the meaning is understood in terms of the inferential profile of the practice of Q-vocabulary deployment. Modal structure is part of the inferential structure, so the quantum state *does* play a central role in meaning-conferral to all Q-vocabulary. But claims about quantum states do not on their own satisfy the norms of denotational deployment, so the quantum state does not acquire a denotational reading.

All of which finally puts us in a position to answer the question: what is the world like according to QM? §7 provides that answer. I conclude in §8 by demonstrating how the investigative strategy can play an irenic role in reconciling some seemingly incompatible positions, each of which, captures something important about QM. In particular, this approach allows us to simultaneously hold on to key insights from (i) nomological-modal approaches to interpreting the quantum state, (ii) pragmatist-advisory interpretations of QM and (iii) the (denotational) Everett interpretation of QM.

¹ I use the term ‘denote’ and its cognates instead of the more common ‘refer’ or ‘represent’ because (i) denotation is not tied to model-theoretic semantics in the way reference and extension are and (ii) it has a relatively clear, non-metaphorical intuitive reading, unlike representation.

2 Two familiarization strategies

Q-vocabulary includes all the expressions one needs in order to articulate the claims that quantum theories make. Quantum theories make claims about all sorts of things, including observable measuring devices and everyday phenomena. Q-vocabulary thus includes not only mathematized elements such as $|\psi\rangle$, \in and \mathcal{H} , and scientific English (or other natural language) terms such as ‘hadron’, ‘electron’, and ‘quark-gluon plasma’ but also elements we recognize from ordinary natural languages such as ‘ammeter’, ‘cat’ and ‘chair’. Many of these expressions are inherited from other, more familiar vocabularies. Just as different theories can be articulated in an English vocabulary, different theories can be articulated in a Q-vocabulary. To interpret QM is then to answer a semantic question, of the form ‘what do the assertions of the theory of QM, *as articulated in an unfamiliar Q-vocabulary*, mean?’

Q-vocabulary inherits elements from familiar vocabularies. As a result, not all Q-vocabulary claims are entirely unfamiliar. We know, or are at least sufficiently familiar with, what $|\psi\rangle = |\text{live cat}\rangle$ and $|\psi\rangle = |\text{dead cat}\rangle$ mean. They mean, respectively, that the system consists of a live cat, and a dead cat (and we take those assertions to be familiar). By contrast, we do not know what is meant by $|\psi\rangle = \alpha|\text{live cat}\rangle + \beta|\text{dead cat}\rangle$. This is an instance of the semantic rendering of the measurement problem. It is a demand for a familiarization strategy for a particular class of Q-vocabulary assertions, the superposed quantum state ascriptions. The broader quantum mechanical interpretative project is the demand for a familiarization strategy for *all* Q-vocabulary articulations of quantum mechanical claims.

To see how a derivative familiarization strategy plays out in a well-understood theoretical context, consider the classical Maxwellian electromagnetism. There, the Faraday tensor F_{ab} is taken to denote the electromagnetic field, which is itself understood as a familiar sort of material field on spacetime.² The unfamiliar-vocabulary term F_{ab} denotes whatever the familiar-vocabulary term ‘electromagnetic field’ does. As a result, F_{ab} becomes a familiar vocabulary element because, for whatever reason, (i) we take F_{ab} to be the sort of expression that denotes, and (ii) the sort of thing that it does denote (call this its *denotatum*) is itself denoted by a familiar vocabulary element, in this case ‘electromagnetic field’. We can now define an inter-vocabulary relation of reference between the vocabulary of which F_{ab} is an element and the vocabulary of which ‘electromagnetic field’ is an element. Using quotation, we can then establish a standard reference schema in the latter vocabulary, which now functions as a metavocabulary: F_{ab} refers to an electromagnetic field. The reference relation is a conduit for denotational familiarity, so insofar as we are familiar what ‘electromagnetic field’ denotes, we

² There is some debate about whether the Faraday tensor, as opposed to some other mathematical expression, for example the electromagnetic potential A^a or a closed-loop holonomy, is best treated as denoting the electromagnetic field. What matters for this section is not the specific choice of mathematical expression, but the near-universal conviction that some mathematical expression denotes the electromagnetic field.

become familiar with what ' F_{ab} ' denotes.

An analogous move in QM would be to say that ' $|\psi\rangle$ ', which is an unfamiliar element of Q-vocabulary, refers to a quantum state. This is unhelpful if we have no idea what a quantum state is, i.e. if the vocabulary element 'quantum state' is itself unfamiliar. More promising, although this is not a view I endorse, is to say that ' $|\psi\rangle$ ' refers to a 2-component material field over a high-dimensional arena (Albert [2013]; Ney [2021]). Given our familiarity with vocabulary elements 'material field in a $(3 + 1)$ -dimensional arena', the leap to the interpretation of ' $|\psi\rangle$ ' as a field in a $(3n + 1)$ -dimensional arena is, in principle, intelligible. Here, we familiarize an unfamiliar vocabulary element by highlighting the similarities in its use as a denoting term with a familiar vocabulary element ('material field'), and parametrizing the differences (' $(3n + 1)$ -dimensional', rather than ' $(3 + 1)$ -dimensional'). In these cases, the unfamiliar vocabulary element inherits familiarity from the familiar vocabulary element to which it is tethered, making this a derivative familiarization strategy.³

An alternative familiarization strategy, and the one adopted in this paper, is based on Brandom's Analytic Pragmatism (R. Brandom [2008]; R. B. Brandom [2015]). It begins by asking of some familiar vocabulary element what makes it familiar. On this strategy, the familiarization proceeds, not by establishing a semantic familiarization relation between the unfamiliar object vocabulary and the antecedently familiar denotation-conferring metavocabulary, but instead by doing to an unfamiliar vocabulary whatever was done to render familiar what was previously an unfamiliar vocabulary.⁴

As with derivative strategies, the framework here is one of object vocabularies and metavocabularies. But unlike with derivative strategies, (i) the metavocabularies here are not denotation-conferring, in the sense that they are the sources of denotational significance to an object vocabulary (they will, nonetheless, contain denoting terms) and (ii) the relations between unfamiliar object vocabularies and metavocabularies are not semantic or logical (i.e. they do not invoke either the semantic machinery of e.g. reference and truth or the logical machinery of e.g. reduction or definition). The original familiar vocabulary is treated as another *object* vocabulary from whose usage we can extract patterns of vocabulary deployment. These patterns can then be encoded in a *second* familiar vocabulary, which we call a 'pragmatic metavocabulary' and with which we can specify the practices linguistic agents engage in that allow them to use a vocabulary to do different sorts of things. The goal of this sort of familiarization strategy is to establish what we might call 'pragmatically-mediated deployment relations' between pragmatic metavocabularies and unfamiliar object vocabularies, rather than to establish semantic or

³ In what follows, I will simply refer to these as derivative strategies. Other examples of derivative strategies are Davidson's T-schematic account of meaning (Davidson [1967]), The Canberra Plan (Chalmers [2012]; Jackson [1998]; D. Lewis [1970]), and the logical empiricists' semantic reduction (Carnap [1936]).

⁴ The derivative vs. investigative familiarization strategy distinction reduces to (Price [2011])'s distinction between object and subject naturalism, when the vocabularies under consideration are scientific. Precursors of this view include (Dewey [1926]; Sellars [1950]; Wittgenstein [2010]).

logical familiarization relations between object vocabularies and denotation-conferring metavocabularies.⁵

Here is an example of the differences between a denotation-conferring and a pragmatic metavocabulary, as it might apply to everyday discourse. Consider a denotation-conferring metavocabulary which contains elements of everyday discourse, and treat, e.g. ‘room’, ‘chair’ and ‘table’ as antecedently denotationally familiar. A sentence like ‘the table and chair are located in the room’ can then be analyzed in terms of vocabulary elements that name particulars and properties (and relations), using the syntactic categories of singular terms (‘table’, ‘chair’, ‘room’,), predicates (‘...located in...’), and logical connectives (‘and’). This strategy does not work for syntactically similar sentences like ‘Priya and Stephen are in cahoots.’ The standard strategy is to first seek some appropriate paraphrase that trades in some denotationally familiar terms. So such sentences will be treated as synonymous with e.g. ‘Priya and Stephen sat at a table yesterday and devised a plan to...’ If that paraphrase still trades in unfamiliar, non-denotational terms (e.g. ‘plan’), then it is paraphrased further.

A pragmatic metavocabulary on the other hand, does not attempt to be a source of denotational significance to an unfamiliar vocabulary. Instead, it uses some familiar vocabulary to tell us what to do—what practices to engage in, what abilities are sufficient, or what norms to abide by—in order *both* to denotationally deploy ‘chair’ and ‘table’, *and* to non-denotationally deploy ‘sake’ and ‘cahoot’, without forcing us to paraphrase sentences in which these expressions occur into ones that trade only in denotational terms (and other familiar vocabulary).⁶ We can then turn to the unfamiliar vocabulary, for example one with the element ‘cahoots’, and try to construct a pragmatic metavocabulary for the deployment of that vocabulary, and decide whether or not to make our use of ‘cahoots’ abide by the norms of denotational or of non-denotational deployment that we extracted from our familiar everyday vocabulary and encoded in a pragmatic metavocabulary.

At this point, I have not made any suggestions about what these norms are; for that, we need a little more metasemantic machinery, which I will provide in §6. One thing I can say is that these norms need not say anything about what ‘cahoots’ means; the denotational deployment norms need not, in addition, be meaning-conferring. This is in stark contrast to the derivative strategy, in which the characterization of a deployment of an expression as denotational relied on our familiarity with its purported denotatum. The investigative strategy prisms these two sorts of norms apart: one can adhere to norms of denotational deployment, without having to specify a denotatum.

⁵ See (R. Brandom [2008], p. 10-29) for some illuminating graphical depictions of this approach.

⁶ Note that the cracks already start to appear for the strict denotationalist when they encounter logical vocabulary. Even the Tractarian Wittgenstein was forced to treat logical vocabulary as meaningful without inheriting any denotational significance. In a sense, the pragmatist strategy tries to do for all vocabulary what the denotationalist was forced to do for logical vocabulary.

3 Denotational vocabulary and modal vocabulary

The quantum state allows us to infer, via the Born Rule, certain sorts of probability judgments. A probability judgment is a quantitative modal claim.⁷ So, whatever else we might say about the quantum state, we have to acknowledge its role in capturing at least some aspect(s) of the modal structure of QM.

This raises the question of what we've captured when we've captured the modal structure of QM. To help answer that question, in this section, I discuss two approaches to understanding the modal elements of a theory expressed in *any* vocabulary: the Humean and the Kantian. derivative versions of both approaches can be found in the literature. Following a Sellarsian suggestion (Sellars [1948], [1957]), I will demonstrate how the Kantian approach plays out in the context of an investigative strategy. The next section will then use a Kantian investigative strategy to arrive at the standard interpretation of classical particle mechanics, as a precursor to arriving at an interpretation of the quantum state.

One of Hume's enduring skepticisms is over our ability to discover modal connections as constituents of the structure of the external world. For Hume, the challenge is couched in the empiricist language of sense impressions, and knowledge of the world. But this challenge can be given a semantic articulation as follows: how can we make modal discourse familiar? Hume's assumption, articulated anachronistically in our framework, is that non-modal denotational discourse is familiar. The further assumption that non-modal denotational discourse can play the required familiarizing role for modal discourse in a non-circular way, is captured in what, following Brandom, I will call the *modal Hume–Quine thesis*.

The modal Hume–Quine thesis: The ability to use ordinary empirical denotational terms does not presuppose the ability to use modal vocabulary.

The derivative strategist solution to Hume's challenge is to discover the correct semantic relations between the unfamiliar modal object vocabulary and a familiar denotational *non-modal* metavocabulary. A canonical example of an attempt at this sort of reduction is Lewis' modal realism, according to which the intensional modal vocabulary of necessity and possibility is semantically reduced to the extensional non-modal denotational vocabulary whose denotata are taken to be elements in concrete but spatiotemporally disconnected possible worlds. But Lewis is not unique in tacitly accepting the modal Hume–Quine thesis in proposing a derivative familiarization strategy for modal discourse; other examples include (Cresswell [1972]; Plantinga [1978]; Stalnaker [2012]).

⁷ To say this and no more is to remain neutral both over what sort of modality (metaphysical or epistemic) is at play, and what exactly is being quantified (degrees of nomic possibility, degrees of metaphysical possibility, degrees of belief, strength of practical advice etc.).

In keeping with the exploratory, rather than critical, intentions of this paper, I will not argue against the plausibility of such a model. Instead, at this point, I will introduce an alternative approach to the Humean challenge, based on what, again following Brandom, I will call the *modal Kant–Sellars thesis*:⁸

The modal Kant–Sellars thesis: The ability to use ordinary empirical denotational terms presupposes an ability to use modal vocabulary.

The derivative strategist will struggle to construct a familiarization of modality that accepts the Kant–sellars rather than the Hume–Quine thesis, if they follow the same path as the Humean by trying to *reduce* modal vocabulary to denotational vocabulary; circularity beckons. But there is another path open to the derivative strategist: flip the script, and argue that *denotational* vocabulary is unfamiliar, and modal vocabulary is familiar. One can then argue that denotational concepts are, in some sense, irreducibly modal without having to worry about the circularity of such a familiarization strategy. This is a common strategy among metaphysicians: trope theorists and bundle theorists (e.g. (Armstrong [1997]; McDaniel [2001])), for example, argue that particulars are just agglomerations of properties, some of which are irreducibly modal.

While this is a perfectly respectable rendering of the modal Kant–Sellars thesis, it fails to harness its real power. To do that, it needs to be combined with an investigative strategy. In particular, with a strategy that seeks to understand the deployment of a vocabulary in terms of a pragmatic metavocabulary. Brandom’s own statement of the modal Kant–Sellars thesis does this definitionally: ‘the ability to use ordinary empirical descriptive terms such as ‘green’, ‘rigid’, and ‘mass’ already presupposes a grasp of the kinds of properties and relations made explicit by modal vocabulary.’ (R. Brandom [2008], pp.96–97).

The pragmatic metavocabulary specifies how to deploy modal vocabulary, but it does so in a very particular way: it allows us to say (i.e. to ‘make explicit’) what it is that we do implicitly (i.e. what we ‘grasp’) when we correctly denotationally deploy ‘ordinary empirical descriptive’ (i.e. denotational) terms. We are familiar with the idea that we can use a vocabulary to say what we did when we engaged in some activity. To ride a bicycle, for example, I do not need to know how to describe the various physical, bodily, and psychological facts that conspire to make me ride a bicycle. What I do need is a good amount of practical knowledge that allows me to ride a bicycle. I can then tell someone about, say, my training regimen for a big race, by deploying a vocabulary that draws attention to certain aspects of my practical knowledge that allows me to ride a bicycle. According to Sellars, the deployment of a piece of vocabulary is just as much an activity performed by an agent in the world as riding a bike is. Just as I have practical

⁸ As with the Hume–Quine thesis, there are grades of strength, depending on the scope of the denial. Sellars, for example, is open to the idea that some denotational vocabulary could be non-modal, but that class of vocabulary is not itself sufficiently robust to ground a familiarization strategy of the sort defended by derivative strategists.

knowledge of how to ride a bike, I have practical knowledge, as a speaker of English, of how to deploy English vocabulary denotationally. The role of a pragmatic metavocabulary is to allow speakers to make explicit different aspects of the practical knowledge that constitutes their discursive abilities.

Our discursive abilities and goals are multifarious, and different vocabularies make explicit different aspects of the norms. What, then, is special about what is made explicit by our *modal* vocabulary? According to Sellars ([1957], §80), '[t]o make first hand use of these [modal] expressions is to be about the business of explaining a state of affairs, or justifying an assertion.' That is to say, in describing some part of the world, one tacitly makes commitments to the propriety of certain sorts of counterfactually-robust inferences. To say 'the cat is on the mat', on this account, is to commit to further claims such as 'the cat would still be on the mat if the mat were an inch thicker, but probably not if I threw a ball at it'. Now, you might not commit to that specific claim, but if you do not commit to the propriety of *any* such inferences, then you have not, in fact, deployed the concept of 'cat' (or 'mat', 'ball', etc.). This is analogous to the more quotidian fact that, if you don't know how to appropriately manipulate the physical structures that constrain the working of your body and a bicycle, you do not know how to ride a bike. Being able to talk about riding a bike (or deploying a vocabulary) is not essential to being able to do it, but when you do talk about it, you are talking about something that you need to know how to do in order to ride a bike (or deploy a vocabulary).

On this view, modal content is an inseparable, though distinguishable, component of any denotational content. This much is common ground between the derivative and the investigative strategists. Where they diverge is over what the modal vocabulary captures. For the former, it denotes the (perhaps *sui generis*) modal constituent of the denotata of our successfully denotationally deployed familiar vocabulary. For the latter, modal vocabulary captures that component of implicit denotational content which, when made explicit, allows us to make claims *about the discursive framework within which the descriptive claims acquire their content*. Note that if we were to follow Sellars' own view, at this stage, we would be committed not only to the claim that modal vocabulary is explicative, but also to the additional claim that it is explicative of *inferential* relations of counterfactual robustness of judgments made using denotational vocabulary. This is strictly optional. But it is a commitment that is very much germane to the view about meaning-conferral that will form the basis of the second part of this paper, and one that I will adopt at that stage.

With this in mind, we can turn our attention to the quantum interpretative project. The interpretation of the quantum state as capturing modal structure has a long and distinguished history ((Dieks [1991]; Dieks and Vermaas [1998]; Healey [1989]; Kochen [1985]; Margenau [1963])). All those modal views are built on the tacit acceptance of a derivative familiarization strategy, and as a result, read the modal content of quantum state ascriptions as denoting some

properties of, or relations between, entities in the world. In §5, I advance an interpretation of the quantum state built instead on the combination of a pragmatist investigative familiarization strategy and the modal Kant–Sellars thesis. In some important ways, this interpretation picks up on some underappreciated aspects of the standard interpretation of the classical state, so it is worth spending a little time there first.

4 The modal classical state

The impulse behind the Hamiltonian formalism for classical mechanics⁹ is that, with judiciously chosen geometrical structure on the set of classical states, we can specify a deterministic set of evolution equations such that, for any choice of initial state at time t , the geometric structure uniquely associates, for any other moment in time t' , some classical state in the phase space.¹⁰ On this way of setting things up, a classical state performs the role of what van Fraassen ([1991], p. 275) calls a *dynamic state*, which is ‘fully specified by stating how [a] system will develop if isolated, and how if acted upon in any definite given fashion’.

It is clear that, on such a setup, the modal structure of a classical theory is given top billing. It allows us to say things like ‘if the state of the classical system at time t were ω_t , then at time t' it would be $\omega_{t'}$ ’. (Rickles [2007], p. 10) highlights that these spaces ‘function as ‘possibility spaces’ for the theories we are interested in.’ So whatever else we might want to say about a classical state, one thing is uncontroversial: it picks out the (diachronic) modal structure associated with a classical theory. What we have here is a (partly mathematical, partly natural-language) classical-mechanical vocabulary, call it a *C-vocabulary*. Different theories can be articulated using C-vocabulary, depending on the precise modal structure captured by the geometry of the associated classical state space.

But what is a classical state? It is decidedly unhelpful to say that it is an element of a symplectic phase space; ‘symplectic phase space’ is an unfamiliar bit of C-vocabulary. Luckily, we can familiarize it relatively easily. One of the primary goals of a classical theory is to allow us to predict the values of (ideally all) magnitudes associated with the class of systems modeled by the theory at some arbitrary point in time, given some information about the system at some initial time t . The classical state allows us to do this because (i) every magnitude of interest is a function of the dynamic state and (ii) a dynamic state at any time, in particular some initial time t , can be inferred from some collection of privileged *base* magnitudes at that time. The conception of a state in terms of magnitudes rather than dynamics is what van Fraassen (Van Fraassen [1991], p. 275) calls a *value state*, ‘which is fully specified by stating which

⁹ Also known as the ‘(symplectic) phase space formalism’ or simply the ‘classical state space formalism.’

¹⁰ Here, we assume a non-relativistic split of space and time. But this is only for illustrative purposes—the Hamiltonian framework generalizes easily to relativistic contexts. For details, see e.g. (H. Goldstein [1950], Ch. 12).

observables have values, and what [those values] are’.

In classical physics, the two roles are played by the same expression, the classical state. Call the existence of an expression that plays both roles the *contingent identity condition (CIC)*. For a classical particle theory, the base magnitudes are six quantities: the three components of position, expressed with respect to some coordinate system (q_1, q_2, q_3) , and three components of linear momentum along those directions (p_1, p_2, p_3) . All other quantities of interest are functions of these base quantities. Kinetic energy, for example, is proportional to the square of the momentum. These magnitudes are extremely familiar sorts of denotata: they are usually denoted by predicates in a C-vocabulary, which might be further familiarized by establishing semantic relations to some everyday or experimental denotational metavocabulary; for example, ‘position’ can be made sense of in terms of an operational procedure involving measuring tape and some fiducial reference body.

The interpretation of the classical state looks like a straightforward implementation of the derivative familiarization strategy: we had some unfamiliar C-vocabulary elements, ‘classical state’ and ‘ ω_t ’, and we inherited the denotational familiarity from antecedently denotationally familiar vocabulary elements ‘position’ and ‘linear momentum’. Of course, what is really going on here, semantically speaking, is that we are familiarizing ourselves with the classical *value* state. Given the satisfaction of the contingent identity condition by classical particle mechanics, it might look as if the project we engaged in is a derivative familiarization of the classical *dynamic* state.

But we should not be fooled by this contingency into thinking that the interpretative project, in general, is the project of derivatively familiarizing the *dynamic* state with the help of a denotation-conferring metavocabulary. The denotational component of interpretation, in classical as much as in quantum mechanics, is concerned only with the value state. The universe need not be so cooperative as to always allow that the same expression play the roles of both dynamic and value state in a successful theory. It is crucial, therefore, to understand the relationship between the (not necessarily denotational) modal structure encoded in the dynamic state and the (necessarily denotational) magnitude structure encoded in the value state for a given theory.

According to the pragmatist investigative strategy in classical mechanics, what the classical dynamic state does is make explicit the implicit modal content of our denotationally deployed C-vocabulary. For example, when we describe some system as having base magnitudes of x units, we are tacitly committed to modal claims of the form: ‘if this system were to be placed in a force field of y units, then its base magnitudes would be z units.’ There is no requirement that the modal structure tacit in our original base magnitude ascriptions is itself to be understood as denoting some quantifiable magnitudes. The pragmatic metavocabulary is flexible enough to allow us to talk about denotational deployment of a vocabulary without

being forced to understand modality denotationally. One can treat modal vocabulary as an explicating metavocabulary without committing to its having any denotational significance.

But here's where things become especially interesting for the pragmatist. There is nothing that stands in the way of understanding modal vocabulary as an explicating metavocabulary which, *in addition* stands in synonymy (or other appropriately tight semantic) relations to elements of Q-vocabulary which themselves satisfy the criteria for a denotational deployment of a vocabulary. So even though ' ω ' is a metavocabulary element, we can propose its synonymy with the object vocabulary element ' $(q_1, q_2, \dots, p_1, p_2, \dots)$ '. In making this defeasible suggestion, we rely on the world to vindicate (or censure) our decision to do so, in exactly the same way as we did when assessing our decision to deploy 'cat' denotationally, or 'sake' non-denotationally. As it happens, classical theories do vindicate this decision in domains in which they are successful.

The upshot of this vindication is the observation that just because we are not *forced* to understand modality denotationally, it does not mean that we cannot accommodate a denotational rendering of modal structure. The two characterizations of vocabularies—as explicating and as (synonymous with vocabulary that is) denotationally deployed—are independent and mutually consistent. The pragmatist does not object to a denotational account of modal vocabulary. They object only to the insistence that the *only* way to make sense of modal vocabulary, in any context, is denotationally.¹¹ This is how, on a pragmatist investigative strategy, one can hold the position that the modal structure of classical mechanics is captured by a vocabulary that is *both* explicating and denotationally deployed.

Having prised apart those two aspects of modal vocabulary deployment, we can explore the possibility of an interpretation of QM which treats the quantum state as capturing modal structure, but which does not treat quantum state or its various cognates, as *denoting* (or synonymous with Q-vocabulary elements that denote)¹² modal properties or magnitudes. On this account, the quantum state is treated as a device for making explicit what was implicit in the correct use of other denotationally deployed Q-vocabulary.

Of course, this is not the end of the quantum interpretative project. By denying the default semantic significance of the quantum state is denotational, the investigative pragmatist sets themselves a new challenge: to make familiar the content of the elements of Q-vocabulary, whether or not they are deployed denotationally. This is a two-stage project: (i) to understand the relationship between the quantum dynamic state and denotationally-deployed Q-vocabulary (§5) and (ii) to understand the content of all Q-vocabulary claims (§6).

¹¹ To continue highlighting the affinities between my framework and Huw Price's, I note that this observation is an alternative rendering of Price's insight that '[s]ubject naturalism is theoretically prior to object naturalism, because the latter depends on validation from a subject naturalist perspective' (Price et al. [2013], p. 6).

¹² To avoid unnecessary repetition, in the rest of this paper, when I talk about modal metavocabulary elements as denoting, what I mean is that these vocabulary elements are synonymous with object vocabulary elements that are denotationally deployed.

5 The modal quantum state

The goal of this section is to introduce and argue for the intelligibility of an investigative familiarization strategy for the quantum state (and Q-vocabulary more generally) based on the modal Kant-Sellars thesis. In the exploratory spirit of this paper, I will not argue against more standard readings of the quantum state, but instead concentrate on introducing the alternative strategy.

The pragmatist investigative strategy attempts to familiarize new Q-vocabulary elements like ‘quantum state’ and ‘ $|\psi\rangle$ ’ by trying to understand the role that they play with respect to the denotationally deployed elements of Q-vocabulary inherited from C-vocabulary. Recall that, on the pragmatist picture, one can prise apart the denotational deployment of a vocabulary element from its semantic content. The process of updating the meanings of C-vocabulary expressions can proceed while the norms around the deployment of many of the denotational C-vocabulary elements that are now also Q-vocabulary elements remains constant. For example, in going from classical electrodynamics to QM, we continued to deploy ‘atom’ denotationally, even though we changed what we meant by ‘atom’. The question of how to understand this change is the question of how to understand meanings; this is the topic of §6.

According to the investigative pragmatist, we were vindicated in the denotational deployment of some scientific C-vocabulary terms (e.g. ‘physical system’, ‘atom’, ‘momentum’) from successful pre-quantum theories. In addition, in those theories, we understood the content of those denotationally deployed vocabulary elements in such a way that the implicit modal structure, when made explicit by articulating a classical dynamic state, could also be given a denotational reading as a maximally determinate classical value state. We then developed a new vocabulary, not by jettisoning and then replacing our old C-vocabulary, but instead by updating it to eventually form a new, but related, Q-vocabulary. Q-vocabulary thus inherited a number of elements from C-vocabulary, many of which were treated as denotationally deployed. The default was *both* (i) to continue to deploy those elements (‘system’, ‘atom’, ‘momentum’) denotationally and (ii) to continue to deploy ‘dynamic state’ denotationally. In a theory that satisfies the contingent identity condition, this conjunctive default is unimpeachable. But in a theory where that condition fails, while the first conjunct is still sound, the second leads to disaster. The failure of QM to satisfy that condition is occluded by derivative strategies, but emphasized by investigative strategies. We should therefore understand the deployment of ‘quantum state’ as playing the same framework explicitating role as ‘classical dynamic state’ did in classical mechanics, but not the same denotational role.

In making this switch, however, the investigative strategist encounters a new problem: although they have a way of identifying when a claim purports to denote, by checking for the satisfaction of norms of denotational deployment, they do not thereby have a way of establish-

ing what those denoting expressions mean. This is not a problem for the derivative strategist. On establishing the denotational deployment of ‘quantum state’, the derivative strategist immediately has an account, or a strategy for developing an account, of what the world is like according to QM: it is a world of whatever is denoted by the denotationally deployed Q-vocabulary elements, including but not necessarily limited to, ‘quantum state’. While their opinions vary over precisely what the relevant denotata are—a field in configuration space (Albert [2013]), non-separable properties of spacetime regions (Wallace and Timpson [2010]), subjective degrees of belief of an agent (Caves et al. [2002]), perfectly localized corpuscles in spacetime whose evolution laws are encoded in either pure quantum states (Bhagal and Perry [2017]; Callender [2015]) or mixed quantum states (Chen [2021])—derivative strategists are united in their conviction that there is some familiar denotatum to be had. The investigative pragmatist does not share this conviction, so needs to address the question of meaning in a different way. We turn to that now.

6 Meaning-conferral and Q-vocabulary

The first stage of an investigative strategy for Q-vocabulary is to adapt, from the practice of deploying C-vocabulary, the norms of explicitation for modal C-vocabulary. This allows us to do to ‘quantum state’ one of the things we were able to do to ‘classical state’, namely to understand its role as a piece of explicating modal metavocabulary. To begin the second stage, we need an alternative to an uncritical derivative strategy, which would have us replicate what we did in classical mechanics, by attempting to locate some denotata for modal vocabulary. Instead, let us begin with deployments of Q-vocabulary elements that are uncontroversially denotational. For example, claims articulated using terms inherited from C-vocabulary that purport to ascribe magnitudes to observables. Since the question about denotational deployment is separate from the question of semantic content, we can say of ‘system’ or ‘chair’ or perhaps even ‘atom’, that these terms are denotationally deployed as elements of both C-vocabulary and Q-vocabulary, even if they mean different things in different theories. Call denotational claims of the form ‘the value of magnitude M on system S is m ’ *canonical magnitude claims (CMCs)*. The relationship between the quantum state and a quantum CMC is one-many, and captured by the Born Rule: a rule for associating with a particular dynamic quantum state a Kolmogorov axiom-obeying distribution over denotational but classically non-maximal canonical magnitude claims.¹³

Having laid the groundwork for the possibility of an account of meaning-conferral that

¹³ The term ‘Kolmogorov axiom-obeying’ rather than ‘probability’ is chosen advisedly here, since, as Healey ([2012a],b, [2017]) points out, the mere satisfaction of those axioms by some aspect of the mathematical formalism does not entail that the associated magnitudes perform the same functions as a probability. In particular, they do not, by default, establish what (Papineau [1996]) calls the inferential and the decision-theoretic links. The areas of sub-surfaces of my desk can be modeled by a functions that obey the Kolmogorov axioms, but that does not make them probabilities.

imbues the same expression with different meanings depending on the theory it is used to articulate, we can get down to the business of establishing those meanings. In this paper, I endorse a Sellars-Brandom-style material-inferentialism. After setting up the broad contours of how this account works, in §6.1, I discuss how Q-vocabulary, both denotational and non-denotational, gets its semantic content in §6.2.

6.1 Brandomian inferentialism

Pragmatism about discursive practice begins by asking: given the aims of our discursive practices, 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, denotation is a good candidate, and we should develop a metasemantics that explains how words come to denote. 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 inheritance of entitlement to commitments.

To assert a claim, on Brandom's view, is to 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 regulated by some norms: assertions, and other discursive performances, 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.¹⁴ And these inferential norms are understood

¹⁴ One consequence of this view is that the primary bearer of semantic significance are judgments, or assertions or sentences, rather than words, terms or vocabulary elements. After all, words don't stand in inferential relations, assertions do. In this paper, I have been treating the latter as semantically significant, in part because I assume that vocabulary elements can inherit their semantic significance from the judgments which they

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 thereby committed to the claim that that ball was visible at t . What makes that inference good or valid is that its conclusion's deontic status is appropriately normatively related to the premise's deontic status. Call such inferences 'material inferences'. That a material inference is good—i.e. that the deontic statuses of premises and conclusions are sanctioned by the norms of the practice—is simply a primitive fact about a discursive practice. A set of material inferences might eventually be treated as bad, and replaced, if they sanction performances that are censured by the world in some way. But all that means is that the discursive practice now primitively treats a different set of material inferences as good.

Brandom's central metasemantic claim, following (Sellars [1953]), is that these sorts of material inferences are meaning-conferring. An assertion is meaningful in virtue of the other assertions (and acts)¹⁵ that it licenses, and that it is licensed by. The meaning-conferring inferences are not to be 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 discursive norms that are, often tacitly, assented to by all the speakers. Indeed, the role of certain classes of speech acts is to make explicit using linguistic expressions, different aspects of what is implicit in the practice of asserting. 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. We have seen this sort of explicitation move already. In §4, the commitments made explicit are commitments to the propriety of certain counterfactually-robust material inferences. The explicating vocabulary elements there are the modal ones, such as 'classical state' and ' $|\psi\rangle$ '.

The deontic scorecard model also allows us to say something concrete about the norms of denotational deployment. According to Brandom, we can understand denotational deployment 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 other propositional attitude 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 set of doxastic commitments and entitlements to S , as well as (ii) an undertaking of a doxastic commitment by me (e.g. that Big Ben exists and

articulate. For details on a concrete proposal for such a 'sub-sentential semantic scalpel', see (R. Brandom [1994], Ch. 6).

¹⁵ In order to incorporate the interaction 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.

has some properties, such as being a bell). The role of such locutions is to coordinate different social perspectives, by nailing down the meaning of my attribution to *S*, in order that I can then, for example, challenge *S*.¹⁶ Notice that the norms of denotational deployment mention only the fact that the deontic scorecard is updated, but says nothing about *how* they should be updated. This is what has allowed us so far in this paper to speak of denotation-conferring norms whilst remaining agnostic over the nature of meaning conferring norms.

There are three important, cascading choice points here: (i) the choice to understand discursive norms in terms of a deontic scorecard in the first place; (ii) the choice to define a deontic scorecard in terms of material inference and (iii) the choice to understand the norms of denotational deployment in terms of Brandom's specific proposal for how the deontic scorecard is updated by *de re* attitude ascriptions. These are all presented here as one set of options among many for making precise the central claim of this paper: that a pragmatic investigative familiarization strategy for QM is attractive. One can accept that claim without committing to any of the further Brandomian specifications, although one would then be on the hook for an alternative specification of an appropriate investigative strategy.

6.2 What is the world like according to everyday discourse?

The quantum state, in its capacity as a dynamic state, is implicated in the inferential profile of canonical magnitude claims. This just means that the quantum state is implicated in meaning-conferral onto CMCs (and various other Q-vocabulary claims). Importantly, the specific choices of meaning-conferral norms are not themselves implicated in the choice of denotational deployment of those claims—we might choose something other than material-inferential norms as meaning-conferring, and still adopt a deontic scorecard model of denotational deployment norms. So if we stipulate, or otherwise establish, that at least some quantum CMCs are denotational and meaningful, we can turn to the task of characterizing what those CMCs describe. We can take our cues, once again, from the relationship between modal content and denotational content in some familiar bit of discourse: this time, let us use everyday discourse about medium-sized dry goods.

There is very little unfamiliar about my asserting 'the chair in my office is in the corner.' It passes the denotational-deployment test from §6.1: you can ascribe to me the belief, *of* my chair that it is in the corner, and undertake the appropriate commitments in your deontic scorecard. 'Chair' is thus the paradigm of a denotationally-deployed vocabulary element. The assertion is implicated in various inferences—counterfactual and actual—that I am licensed to

¹⁶ By contrast, suppose I say 'S believes of glumbles that they are voriant.' Although this has the surface grammar of a *de re* attribution, it is not one; it fails the condition of making me undertake a particular doxastic commitment. I do not have any antecedent views about the properties of glumbles that can be updated by *S*'s assertion, because I do not know what follows from the deployment of 'glumble' and 'voriant'.

make. Together with my background folk-theoretic physics, I might also be committed to the claim that my office chair would have fallen over if it did not have the support of a pair of back legs. And I can then examine my deployment of the term ‘back legs’ and check whether the claims that I articulate using those terms pass the denotational deployment test. Suppose they do.

Then, according to our best theory of everyday objects, the denotational deployment of ‘chair’ entails the denotational deployment of ‘back leg’. The norms of everyday material inferences, which include inferences about what our everyday theories say about counterfactual situations, vindicate the move from the denotational deployment of ‘chair’ to the denotational deployment of ‘back leg’. To commit to this is not to commit to the idea that back legs are somehow necessary constituents of chairs, but rather, to the idea that the norms around the deployment of ‘chair’, theory-laden as they are, vindicate adopting norms of descriptive deployment of ‘back leg’. To speak metaphorically, the concept of a back leg is part of the everyday-theoretic concept of a chair. One might, of course, be wrong to have this commitment, perhaps because one has the wrong background theory, or because of some other bad commitments. But conditional on accepting a particular set of denotational deployment norms, and the vindication of that acceptance, one is vindicated in the denotational deployment of ‘back leg’.

To know what the part of the world described by assertions about back legs is like, I need to know not only the meaning of ‘back leg’, but also the meanings of a whole slew of other vocabulary elements of our everyday discourse. ‘Back leg’ is a piece of semantic machinery that allows me to encode the propriety of a certain class of material inferences, for example the counterfactual inference from ‘if the chair did not have back legs’ to ‘it would have been perched at a angle.’ To know what ‘back leg’ means is to have some grasp of the inferential profile associated with the class of claims articulated in everyday vocabulary. Equipped with a sufficient grasp of this profile, which includes an adherence to sufficiently vindicated norms of denotational and non-denotational deployment, one has an account of what the world is like according to (whatever theory is expressed in) everyday vocabulary.

7 The world according to quantum mechanics

To interpret QM, we need to start with quantum CMCs, and the *familiar* denotationally-deployed Q-vocabulary elements that we inherited from some extant familiar vocabulary. We use those elements as anchors, whose meanings and whose denotational character are both understood in terms of the material-inferential profile associated with some extant, but for whatever reason inadequate, theory.

It is tempting to think of the quantum CMC in terms of classical value states, and to char-

acterize quantum dynamic states as giving rise to distributions over non-maximal CMCs. But the quantum CMCs are only non-maximal with respect to (somewhat artificially abstracted) classical standards of magnitude determinacy. Given that the meanings of predicates that denote magnitudes are themselves determined by the inferential structure of the theory being interpreted, it is preferable to think of quantum CMCs as specifying, by the standards of QM, determinate magnitudes to systems.¹⁷ It is our job, as interpreters of QM, to coordinate these values to the values associated with experimental, and other observational, outcomes. To do that is to establish at least some relations to denotationally familiar vocabulary elements (such as ‘pointer’, ‘display’, ‘observable’).

The first interpretational step is to render, within denotational Q-vocabulary, familiar denotational claims from classical mechanics. For example, ‘the position of the chair is r ’. Classically, this CMC is associated with a relatively simple dynamic (and value) state ascription. Quantum mechanically, for a system of as many degrees of freedom as a chair together with its environment (which, in the limit, is the entire universe), the dynamic state ascribed to the whole system is a *pure quantum state*, while the dynamic state ascribed to the chair in particular is a *mixed state*.¹⁸ The pure quantum state of the total system evolves linearly and unitarily; the so-called *decohered* mixed state of the chair evolves non-unitarily and quasi-classically. The details of this dynamical fact can be found in e.g. (Schlosshauer [2007]).

A suitably decohered mixed state ascription therefore mimics, in some important ways, the classical dynamics of everyday systems. This is why a quantum CMC associated with a suitably decohered mixed state can be treated as denotationally deployed; the norms around the denotational deployment are preserved in going from C- to Q-vocabulary. But, crucially, the modal structure—and therefore the meaning-conferring norms—made explicit by the quantum state is *vastly* different from the modal structure encoded in the classical state. And so the meanings of the denotationally deployed vocabulary elements in classical CMCs are vastly different from those in quantum CMCs. ‘Chair’, as it appears in a classical CMC, exhibits a modal structure regimented by the mathematics of Hamiltonian vector fields on a symplectic manifold. ‘Chair’, as it appears in a quantum CMC, exhibits a modal structure regimented by the mathematics of operator algebras on Hilbert spaces. The machinery of decoherence is what allows us to establish conditions of sufficient similarity between the modal structures of quantum and classical mechanics that we can justify importing some, but not all, of the norms from C- to Q-vocabulary. For example, the norm that ‘chair’ is to be denotationally deployed, but not the norm of material-inference from ‘I place a cat in a closed box with a vial of poison and a uranium emission-sensitive vial-smasher’ to ‘the cat is dead’.

¹⁷ (Menon [2024]) highlights the theory-relativity of standards of determinacy, albeit in a non-inferentialist framework.

¹⁸ For a definition and characterization of pure and mixed states, see e.g. (Hughes [1989], Ch. 3)

In the context of the everyday-theoretic denotational deployment of terms like ‘chair’ and ‘back leg’, we were able to justify the treatment of ‘back leg’ as denotational by invoking the norms of counterfactually robust material inferences in our everyday theory. In doing so, we were able to understand ‘back leg’ as conceptually inseparable, by those normative standards, from ‘chair’. Let us explore what happens when we make the same move in the quantum mechanical inferential context.

We inherit from C-vocabulary the treatment of ‘chair’ as a denotationally deployed Q-vocabulary element. The material-inferential norms of meaning-conferral onto denotationally deployed *quantum* CMCs tells us that the only way to recover the robust empirically-vindicated dynamics of the chair within the inferential structure of QM is for the dynamical quantum state of the chair, as a subsystem of the universe, to be a highly decohered mixed state.

But a highly decohered mixed state such as this has a lot more structure in it than the (modal content associated with the) mere denotation of a single chair system. It has a dynamical structure within which the structures that encode the denotational significance of elements like ‘chair’ treat that denotational significance as inseparable from a *branching modal structure*.

This is a subtle and crucial point: we have not arrived at the denotational significance of vocabulary elements that purport to denote objects in other branches by insisting that the only way to make sense of modal structure is denotationally. To do so would be to fall back on the derivative familiarization strategy that invoked a contingency about the inferential structure of classical mechanics that QM lacks. Instead, what we have done is investigate the conceptual structure of the Q-vocabulary elements that we use *in hooking up QM to the familiar everyday world*.¹⁹ We have a view on which the branching structure of the quantum state is not itself treated as independently denotational, but rather as an explicitation of the modal structure of the denotationally deployed Q-vocabulary elements.

A pragmatist-inferentialist investigative strategy applied to (unmodified) QM results in a many-worlds view.²⁰ The modal structure of the quantum mechanical dynamics of everyday objects is branching. The modal structure of any theory, on an inferentialist approach, is meaning-conferring. So the meanings of the quantum concepts deployed to denote everyday objects commit us to the denotational content of other quasi-classically evolving branches. The inferentialist about Q-vocabulary as it articulates an unmodified quantum theory is, by default, a type of Everettian. To deploy ‘chair’ according to the norms of Q-vocabulary deployment is to deploy a concept whose modal structure entails the propriety of the denotational deployment of ‘chair in another branch’. This is the exact analogue, in Q-vocabulary, of the vindication of the denotational deployment of ‘back leg’ in an everyday vocabulary in which ‘chair’ is

¹⁹ This is a clear instance of using a Sellarsian methodology to engage in the Sellarsian project of reconciling the manifest image with the scientific image (Sellars [1963b]).

²⁰ I demonstrate how the investigative strategy allows us to articulate a robust justification of the decision not to modify QM in §8.3.

denotationally deployed.

Healey and the other semantic anti-representationalists are spot on when they say things like ‘the function of the quantum state... is not to represent properties of a physical system to which it is assigned, nor anyone’s knowledge of its properties’ (Healey [2017], p. 207). But to say that is not to deny that the quantum state *is implicated* in the denotational significance of other claims made by QM. After all, the quantum state is not, even according to Healey, an idle wheel in the inferential structure of QM. So the inferentialist cannot simply deny the representational (denotational) significance of the quantum state, dust their hands and move on. They are still on the hook for an account of the role that the quantum state (and other elements of Q-vocabulary) plays in the denotational deployment of quantum mechanical claims, that is, in accounting for what the world is like according to QM.

The investigative strategy gives us exactly the tools we need here. The world according to some new, unfamiliar, perhaps successor theory *S* consists of the denotata of the denotationally deployed elements of *S*; we work out what those denotata are by inheriting, from some familiar theory *T*, some denotational vocabulary, and then adjusting its place within the new material-inferential web of *S*. So the world according to Newtonian cosmology is the world according to Aristotelian cosmology, but with some new fields, understood as the denotata of some new concepts, and without a privileged center of the universe, understood as the denotatum of some previously denotationally-deployed concept. And the world according to unmodified QM is the world according to classical mechanics, but with a large but indeterminate²¹ number of other approximately classically-evolving branches characterized in terms of the denotata of all our denotationally-deployed Q-vocabulary elements. In that sense, perhaps surprisingly, the view I propose here bridges two formerly diametrically opposed views: Oxford-style Everettianism and Healey-style pragmatism.

8 For inimical to irenic

This paper began with a bold claim: that three mutually inimical views—nomological, pragmatist and Everettian—could not only be reconciled, but also that their insights could also be acknowledged as valuable, and accommodated simultaneously. In this final section, I highlight one key insight from each approach that I take to be valuable, worth holding on to, and which a rejection of the approach would force us to give up. I then demonstrate how an investigative strategy allows us to hold on to all of them. This section is not intended to exhaust the advantages of going investigative. The goal is instead to motivate an interest in exploring the investigative resources that philosophers of physics have left largely unexplored.

²¹ See (Wallace [2012], Ch. 3).

8.1 Investigating nomological-modal approaches

Nomological-modal approaches to the quantum state share a commitment to the claim that, roughly speaking, the quantum state expresses, or performs a role sufficiently analogous to, a law of nature. The underlying insights are sound: (i) the quantum state is a dynamic state, so does not demand a denotational reading in the way that a value state does; (ii) the quantum state encodes paradigmatically nomological modal content over non-maximal value states: the values that particular magnitudes could take, according to the laws of nature.

The derivative strategy does not allow the proponent of the nomological view to cash in their chips there. It demands a familiarization of the nomological quantum state in terms of a denotation-conferring metavocabulary. And this fact is borne out in the literature, where views split along the divide between Hume-Quineans and Kant-Sellarsians. The derivative Hume-Quineans ((Bhogal and Perry [2017]; Callender [2015])) treat the quantum state as encoding the modal-nomological facts about the behaviour of entities described in entirely non-modal terms; this is easiest to see in the case of Humean Bohmians, who denotationally and non-modally deploy ‘Bohmian corpuscle’ and render the quantum state as a Lewisian Best System (D. Lewis [2013]). In this case, the non-denotational insight is swamped by the demand to construct a Humean Best System that reproduces QM. To date, this has remained a promissory note, one whose plausibility remains dubious (see e.g. (Dewar [2016])).

The derivative Kant-Sellarsians (Allori [2015]; Chen [2021]; S. Goldstein and Zanghi [2013]) treat the quantum state by analogy (roughly) with the Hamiltonian vector field on a classical symplectic phase space—a distinguishable component of the theoretical structure that they identify as capturing modal structure which is inseparable from the denotational component of that theoretical structure. The quantum state is not treated denotationally, but as capturing the inseparable modal properties instantiated by the entities denoted by the denotationally-deployed vocabulary. Hamstrung by the derivative strategy, they are forced to render nomological modal structure in denotational terms, and thus undermine the original non-denotational insight. So on derivative versions of both the Hume-Quine and the Kant-Sellars approaches, the nomological approach is on shaky ground.

The investigative approach fares considerably better. The inferential (and therefore meaning-conferring) structure is one that supports robust counterfactuals. Insofar as a pragmatist investigative approach treats the quantum state as making explicit features of the quantum mechanical inferential framework within which we denotationally deploy CMCs, it eschews a denotational reading of the quantum state. Moreover, since these inference licenses are scientific, the modality at play is nomological (this allows us to make sense of Sellars’ claim that concepts involve laws of nature and are inconceivable without them (Sellars [1948])). Both of the key insights of the nomological approach are therefore accommodated by the investigative view, but there is no further demand to understand insights in directly or indirectly denotational terms.

8.2 Investigating pragmatist approaches

The key insight of pragmatist approaches to QM is that QM is very good at giving us advice about how to cope with the world, by telling us what we ought to infer from what we are committed to. The question of what we ought to infer is a normative question, so a purely denotational interpretation of QM would leave us open to a version of Moore's naturalistic fallacy (Moore [1903]): that normative claims can be reduced without residue to, or derived without further assumptions from, descriptive non-normative claims. Pragmatist approaches to QM such as Healey's are immune to this worry, because there are no purely descriptive, non-normative claims: to describe is to be subject to the norms of denotational deployment. For proponents of Sellarsian inferentialism, such as Healey, it is to locate oneself 'in the logical space of reasons' (Sellars [1963a]).

Immunity to the naturalistic fallacy is a key selling point of pragmatist approaches. But this immunity is only as strong as its characterization of the semantic content of denotational claims about the world. Healey offers the following:

[u]nlike theories of classical physics, the novel mathematical structures in quantum theory's models are not applied directly to offer representations of physical situations... they are [used to make] statements about entities and magnitudes acknowledged by the rest of physics, not statements about physical entities and magnitudes newly represented in models of quantum theory. (Healey [2017], pp. 136–137)

Peter Lewis and David Wallace raise important, and related, concerns for this aspect of Healey's view. For 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.' (Wallace [2020]) demands a sharper characterization of these magnitudes that are 'acknowledged by the rest of physics', which he calls *non-quantum physical magnitudes (NQPMs)*.

Wallace canvasses a few options for the origins of NQPMs. One option is that these magnitudes are determined by non-quantum physical theories. Another is that they are imported uncritically from classical theories. Wallace proceeds, correctly in my view, to argue against the plausibility of these readings, on the grounds that, strictly speaking, specific theories within the classical framework are *really* approximations of better theories in the quantum framework. He finally suggests that the right way to understand Healey's position is that the magnitudes are picked out by quantum decoherence. But, according to Wallace, it is unclear that Healey's non-representational reading allows him those resources.

It is easy to see where both Lewis' and Wallace's unease comes from: the presupposition that denotational magnitudes as expressed in Q-vocabulary must attain their denotational significance and content via a derivative familiarization relation to C-vocabulary. If, instead, we

adopt an investigative strategy, together with my proposed view about the relationship between decoherence and *denotation* rather than meaning, we avoid the naturalistic fallacy without having to denude Q-vocabulary of its rich inferential and denotational significance. We can make sense of Healey's claim in such a way that Lewis' and Wallace's worries are assuaged. We can (and should) say something like: an NQPM is the denotatum of some piece of denotationally deployed C-vocabulary (or other non Q-vocabulary). We import that expression into Q-vocabulary, together with (i) a (defeasible) commitment to continuing to adopt the norms of denotational deployment and (ii) a commitment to updating the meaning of that expression by adhering to the *new* norms of meaning-conferral in QM. This sidesteps Wallace's worries because (i) there is no claim that classical magnitudes are, in any sense, conceptually prior to and independent of (the inferential structure of) QM and (ii) quantum CMCs are meaningful both in undecohered as well as decohered contexts, so Healey can help himself to the idea that decoherence *does* pick out the circumstances in which quantum magnitudes have representational (i.e. denotational) significance. Decoherence acts as a demarcation criterion for circumstances under which claims articulated within a normative-inferential framework also have denotational significance. The investigative pragmatist thus remains immune to the naturalistic fallacy.

8.3 Investigating Everettian approaches

The astute reader might have noticed a sleight of hand between the introduction and this section. The introduction spoke of a 'descriptive reading' of the quantum state as a reading worth taking seriously. This section speaks of an Everettian approach to the quantum state. But, as is well known, there is a slew of purely descriptive but non-Everettian readings available. So why the subterfuge?

There is no subterfuge. While there are other brands of descriptive approach available, the Everettian view is the only descriptive approach to *unmodified* QM. The key insight of Everettian QM is that an unmodified QM is conceptually coherent. Everettians usually invoke some kind of parsimony argument to support their adoption of an unmodified QM.

The investigative strategy allows us to see why it makes sense not to modify QM: by shifting the denotational burden away from the quantum state without thereby absolving it of all metasemantic responsibility, the pragmatist investigative strategy makes modifications to QM superfluous. The standard motivations for introducing hidden variables, or supplementary dynamics, come from the assumption that the quantum state has to play this dual role: modal structure codifier and (sole) ontology specifier. In other words, the assumption is that QM should also satisfy the contingent identity condition. Unmodified QM does not, so according to the standard derivative approach, its inferential structure should be modified, either by supplementing the modal structure by adding hidden variables or altering the modal structure by

modifying the dynamics. But if you drop the insistence on satisfying the contingent identity condition, as the investigative strategist recommends, you lose a key motivation to modify QM.

In addition to justifying beginning with an unmodified QM, the investigative pragmatist lowers the barriers of entry for Everettian QM. The party line at Oxford on the Everett interpretation is that ‘Everettian quantum mechanics reads the quantum state literally, as itself standing *directly* for a part of the ontology of a theory’ (Wallace [2012], p. 295). In order to make sense of this claim, Wallace adopts three substantial commitments: (i) derivative familiarization standards from some pre-quantum theory, for example classical field theory (‘[w]hile the kinds of property that will be attributed to a spacetime region by the assignment to it of a density operator are not the ones with which we are terribly familiar, we have suggested that this kind of feature is generic in fundamental physics: the case, we have suggested, does not differ in principle from the unfamiliarity of the electromagnetic field.’(Wallace and Timpson [2010], p. 724));(ii) a relatively robust form of scientific/structural realism (‘most defences of the Everett interpretation (including some of my own) do still adopt that ‘fairly uncritical scientific realism’ (Wallace [2020], p. 382)); and (iii) a Dennettian functionalism about entities (‘Dennett’s criterion: A macro-object is a pattern, and the existence of a pattern as a real thing depends on the usefulness—in particular, the explanatory power and predictive reliability—of theories which admit that pattern in their ontology.’ (Wallace [2012], p. 50)), in order to cash out this notion of literal interpretation.

I do not intend to suggest that these moves are unjustified. But I do want to highlight that none of these additional moves are necessary on the investigative strategy: the metasemantics that imbues QM claims with meaning in the first place is already sufficient to make sense of a denotational deployment of concepts whose meanings entail an Everettian commitment; no standards illicitly imported from classical mechanics, no across-the-board scientific realism, and no objecthood criteria that float above metasemantics.

It would be disingenuous (and wrong!) to claim that inferentialist-pragmatism entails the Everett interpretation. Indeed, if reading the quantum state as denotational is partly constitutive of the Everett interpretation, then what I’ve elaborated here is better described as ‘Everett-like’. But I hope to have demonstrated that views in the vicinity of nomological approaches to the quantum state, quantum pragmatism and Everettian QM have more in common than previously suspected, and insights from all three schools can be simultaneously acknowledged as part of a more comprehensive metasemantic approach to interpreting quantum mechanics.

References

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

- Allori, V. [2015]: “Primitive Ontology in a Nutshell”, *International Journal of Quantum Foundations*, **1**, pp. 107–22.
- Armstrong, D. M. [1997]: *A world of states of affairs*, Cambridge University Press.
- Barrett, J. A. [2019]: *The conceptual foundations of quantum mechanics*, Oxford University Press.
- Bhogal, H. and Z. Perry [2017]: “What the Humean should say about entanglement”, *Noûs*, **51**, pp. 74–94.
- Brandom, R. [1994]: *Making It Explicit: Reasoning, Representing, and Discursive Commitment*, Harvard University Press, ISBN: 978-0-674-54330-0.
- Brandom, R. [2008]: *Between saying and doing: Towards an analytic pragmatism*, Oxford University Press.
- Brandom, R. [2009]: *Articulating Reasons: An Introduction to Inferentialism*, Harvard University Press.
- Brandom, R. B. [2015]: *From empiricism to expressivism: Brandom reads Sellars*, Harvard University Press.
- Callender, C. [2015]: “One World, One Beable”, *Synthese*, **192**, pp. 3153–77.
- Carnap, R. [1936]: “Testability and Meaning”, *Philosophy of science*, **3**, pp. 419–71.
- 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](https://arxiv.org/abs/quant-ph/0106133).
- Chalmers, D. J. [2012]: *Constructing the World*, Oxford: Oxford University Press.
- Chen, E. K. [2021]: “Quantum Mechanics in a Time-Asymmetric Universe: On the Nature of the Initial Quantum State”, *British Journal for the Philosophy of Science*, **72**, pp. 1155–83.
- Cresswell, M. J. [1972]: “The world is everything that is the case”, *Australasian Journal of Philosophy*, **50**, pp. 1–13.
- Davidson, D. [1967]: “Truth and Meaning”, in *Philosophy, Language, and Artificial Intelligence: Resources for Processing Natural Language*, Springer, pp. 93–111.
- Dewar, N. [2016]: “La Bohème”, *Synthese*, **197**, pp. 1–19.
- Dewey, J. [1926]: *Experience and Nature*, Chicago: Open Court Publishing Co.
- Dieks, D. [1991]: “On some alleged difficulties in the interpretation of quantum mechanics”, *Synthese*, pp. 77–86.
- Dieks, D. and P. E. Vermaas [1998]: *The modal interpretation of quantum mechanics*, Vol. 60, Springer Science & Business Media.
- Goldstein, H. [1950]: *Classical mechanics*, Addison-wesley Reading, MA.
- Goldstein, S. and N. Zanghì [2013]: “Reality and the role of the wave function in quantum theory”, *The wave function: Essays on the metaphysics of quantum mechanics*, pp. 91–109.

- Healey, R. [1989]: *The philosophy of quantum mechanics: An interactive interpretation*, Cambridge University Press.
- Healey, R. [2012a]: “Quantum Decoherence in a Pragmatist View: Dispelling Feynman’s Mystery”, *Foundations of Physics*, **42**, pp. 1534–55, ISSN: 1572-9516.
- Healey, R. [2012b]: “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.
- Hughes, R. I. G. [1989]: *The structure and interpretation of quantum mechanics*, Harvard university press.
- Jackson, F. [1998]: *From Metaphysics to Ethics: A Defence of Conceptual Analysis*, Clarendon Press.
- Kochen, S. [1985]: “A new interpretation of quantum mechanics”, in P. Lahti and P. Mittelstaedt (eds.), *Symposium on the foundations of modern physics*, World Scientific, pp. 15–170.
- Lewis, D. [1970]: “How to Define Theoretical Terms”, *The Journal of Philosophy*, **67**, pp. 427–46, ISSN: 0022-362X, JSTOR: [2023861](#).
- Lewis, D. [2013]: *Counterfactuals*, John Wiley & Sons.
- Lewis, P. J. [2020]: “Quantum Mechanics and Its (Dis) Contents”, in S. French and J. Saatsi (eds.), *Scientific realism and the quantum*, Oxford: Oxford University Press, pp. 168–82.
- Maher, C. [2012]: *The Pittsburgh School of Philosophy: Sellars, McDowell, Brandom*, Vol. 42, Routledge.
- Margenau, H. [1963]: “Measurements and quantum states: Part II”, *Philosophy of Science*, **30**, pp. 138–57.
- McDaniel, K. [2001]: “Tropes and ordinary physical objects”, *Philosophical Studies*, **104**, pp. 269–90.
- Menon, T. [2024]: “On algebraic naturalism and metaphysical indeterminacy in quantum mechanics”, *Studies in History and Philosophy of Science*, **105**, pp. 1–16.
- Moore, G. E. [1903]: *Principia Ethica*, ed. by T. Baldwin, Mineola, N.Y.: Dover Publications.
- Muller, F. A. [2023]: “Six measurement problems of quantum mechanics”, in *Non-reflexive logics, non-individuals, and the philosophy of quantum mechanics: Essays in honour of the philosophy of Décio Krause*, Springer, pp. 225–59.
- Myrvold, W. [2022]: “Philosophical Issues Raised by Quantum Theory and its Interpretations”, in O. Freire, G. Bacciagaluppi, O. Darrigol, C. Joas, and A. Kojevnikov (eds.), *The Oxford handbook of the history of quantum interpretations*, Oxford University Press.
- 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.

- Plantinga, A. [1978]: *The nature of necessity*, Oxford University Press.
- 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.
- Rickles, D. [2007]: *Symmetry, structure, and spacetime*, Vol. 3, Elsevier.
- Schlosshauer, M. [2007]: *Decoherence and the quantum-to-classical transition*, Springer.
- Sellars, W. [1948]: “Concepts as Involving Laws and Inconceivable Without Them”, *Philosophy of Science*, **15**, pp. 287–313.
- Sellars, W. [1950]: “Language, Rules and Behavior”, in S. Hook (ed.), *John Dewey, philosopher of science and freedom: a symposium*, Greenwood Press, pp. 289–315.
- Sellars, W. [1953]: “Inference and Meaning”, *Mind*, **62**, pp. 313–38, ISSN: 0026-4423, JSTOR: [2251271](https://www.jstor.org/stable/2251271).
- Sellars, W. [1954]: “Some Reflections on Language Games”, *Philosophy of Science*, **21**, pp. 204–28.
- Sellars, W. [1957]: “Counterfactuals, Dispositions, and the Causal Modalities”, in H. F. M. Scriven and G. Maxwell (eds.), *Minnesota Studies in The Philosophy of Science, Vol. II*, University of Minnesota Press.
- Sellars, W. [1963a]: “Empiricism and the Philosophy of Mind”, in *Science, Perception, Reality*, Humanities Press, New York.
- Sellars, W. [1963b]: “Philosophy and the Scientific Image of Man”, in W. Sellars (ed.), *Science, Perception and Reality*, Routledge & Kegan Paul.
- Stalnaker, R. [2012]: *Mere possibilities: Metaphysical foundations of modal semantics*, Princeton University Press.
- Van Fraassen, B. C. [1991]: *Quantum mechanics: An empiricist view*, Oxford University Press.
- 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. 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>.
- Wittgenstein, L. [2010]: *Philosophical Investigations*, John Wiley & Sons, ISBN: 978-1-4443-0797-9.