On algebraic object naturalism and metaphysical indeterminacy in quantum mechanics

Tushar Menon

Dianoia Institute of Philosophy,
Australian Catholic University, Melbourne

Abstract

I propose a technique for identifying fundamental properties using structures already present in physical theories. I argue that, in conjunction with a particular naturalistic commitment, that I dub 'algebraic naturalism', these structures can be used to generate a standard of metaphysical determinacy. This standard can be used to rule out the possibility of a virulent strain of 'deep' metaphysical indeterminacy that has been imputed to quantum mechanics.

Keywords: Quantum mechanics, Metaphysical Indeterminacy, Properties, Kochen-Specker Theorem.

Contents

1 Introduction

How opinionated is a physical theory with respect to its metaphysics? This question is central to the concerns of both the scientific realist and the naturalist. The scientific realist, to the extent that they believe in the literal truth of scientific claims including claims about unobservables, has to deal with the worry that physics is not opinionated enough; call this the problem of *ontological underdetermination*. The naturalist, to the extent that they believe that all there is to the world is what is studied by science (Price [25] calls this 'object naturalism'), also has to deal with the worry that physics is not opinionated enough, but in the sense that worldly entities might have properties (or stand in relations) that are invisible to our best physics. This, too, is a form of underdetermination, but is distinct from the ontological underdetermination that plagues the scientific realist.[1] This form of underdetermination, call it *ideological*

[1]The scientific realist and the scientific anti-realist can both agree that the properties ascribed by our best physics to, say, electrons completely exhausts the properties of electrons, whilst nonetheless disagreeing over how to characterise the referent of 'electron'. Similarly, the naturalist and the anti-naturalist might both agree that electrons are unobservable, but disagree over whether that electron has properties over and above those ascribed to it by physics.

underdetermination, is the central concern of this paper.

Accepting the force of ideological underdetermination, the *weak naturalist* argues that physics determines some constraints on metaphysical theorising about properties, but is not unique in doing so. Other considerations can, and some argue should, be appealed to. Considerations like explanatory power, ontological parsimony, and so on. In this paper, I develop a strong version of naturalism, according to which fundamental physics does not merely constrain our choices of properties and relations, but actually uniquely specifies them. Weak naturalism might initially have seemed more attractive than strong naturalism because the latter might have seemed unattainable. The project of this paper is to show that strong naturalism is a perfectly attainable, well-defined norm: physics can identify its own metaphysical commitments regarding properties and relations. Physics, accordingly, is maximally opinionated and the strong naturalist can deny ideological underdetermination.

Concerns around ideological underdetermination form the basis of the debate over so-called *metaphysical indeterminacy*, which roughly speaking, is a form of indeterminacy that attaches itself to the world, independently of how we may go about describing the world.[2] Although the locus classicus of this debate traces its way back to disputes over the status of vague or indeterminate everyday objects (see e.g. [11, 13, 21]), more recently, it has been suggested by a number of commentators, mostly philosophers [4, 8, 7, 9, 10, 29, 31, 35], but also some physicists [5, 26], that quantum mechanics (QM) provides more robust, scientifically grounded examples of metaphysical indeterminacy that do not rely on the vagaries of natural language or intuitions about objects.

Although these commentators agree that models of QM exemplify metaphysical indeterminacy, there is a wide range of precise diagnoses of its failures of determinacy, based on competing suggestions for how to cash out metaphysical indeterminacy in the first place. For example, Calosi and Wilson [7, 8] argue that determinables exist over and above the disjunction of their possible determinates; quantum metaphysical indeterminacy is then a matter of states of affairs having determinable properties with no unique determinate. Darby and Pickup [10], on the other hand opt for a situationalist account, according to which 'situations' can provide partial representations of states of affairs; quantum metaphysical indeterminacy is then a matter of propositions that report quantum states of affairs having different truth values in

[2] Contrast this with *semantic* indeterminacy, which is an expression of some deficiency in the language we use to describe or represent some part of the world, and *epistemic* indeterminacy, which similarly points to a deficiency in our knowledge of, or our ability to know about, some part of the world.

different situations. Torza [31], too proposes a situationalist understanding; quantum metaphysical indeterminacy is then a matter of propositions reporting quantum states of affairs not having truth values in 'adequately actual' situations, i.e. situations that contain exactly those sentences that are true of reality.

The diversity of approaches to understanding indeterminacy is matched by the diversity of alleged examples from within QM of metaphysical indeterminacy. It is important, therefore, to highlight that not all forms of putative metaphysical indeterminacy in QM are of the same type. We can distil from the literature at least two[3] putative types of quantum metaphysical indeterminacy (at this point, I eschew jargon to be introduced later, and provide rough, intuitive characterisations; these positions are more precisely rendered in §4):

`Superposition indeterminacy`: Generic quantum states are in superpositions of quantum states. Therefore the values of certain properties are metaphysically indeterminate (e.g. position, momentum, spin).

`Deep indeterminacy`: For some collections of properties, it is not possible to ascribe values to the complete collection consistently. Therefore the values of certain properties are metaphysically indeterminate.[4]

I demonstrate that these two types of indeterminacy are distinct in §4. In this paper, I argue that QM is not deeply metaphysically indeterminate. I remain neutral regarding the purported metaphysical indeterminacy of QM viz. `Superposition indeterminacy`; I agree with Calosi and Wilson's claim that positions with respect to `Superposition indeterminacy` will likely be indexed to particular interpretations of QM. But I argue that deep metaphysical indeterminacy can be ruled out in an interpretation-agnostic manner.

As I mentioned above, there are competing suggestions in the literature for how to cash out metaphysical indeterminacy. In this paper, I follow Barnes and Williams' characterisation. Their model (henceforth the 'BW model') [2] asserts that metaphysical indeterminacy corresponds to a new form of modality—*precisificational possibility*.

[3]There are other forms of indeterminacy implicit in the literature. For example, `Entanglement`: One cannot ascribe determinate properties to one out of the collection of 'constituents' of an entangled state. Therefore it is metaphysically indeterminate what properties those constitutents instantiate. I do not discuss these because they are not relevant to the broader claim about deep metaphysical indeterminacy.

[4]Readers familiar with Skow's work on deep metaphysical indeterminacy will recognise `Superposition indeterminacy` as a special case of what Skow calls 'shallow metaphysical indeterminacy'.

If there is metaphysical indeterminacy, Barnes and Williams claim, then '[t]here will be more than one world in the space of precisifications'. I describe the BW model in more detail in §4.

I choose this model for three reasons. First (and most importantly) I think it is a plausible account. Second, it has been argued by commentators like Darby [9], Fletcher and Taylor [14], and Skow [30], that the BW model is inadequate to model the indeterminacy that arises from the BKS theorem. I disagree. I demonstrate that the BW model is perfectly adequate to model states of affairs arising from BKS theorem-type considerations. Third, the BW model has a close natural association with bivalent model-theoretic approaches to semantic indeterminacy, in terms of which I cash out strong naturalism using what I call the definability thesis:

`Definability thesis:` For any given physical theory, if one uses that theory (and only that theory) to extract ontological commitments, then one ought to be committed only to properties and relations that are definable in that theory's models,

where 'definable' is used in a strict model-theoretic sense that I will elaborate on below.[5] The definability thesis entails Glick's [16] 'sparse view' at least insofar as the latter applies to deep indeterminacy. The central goal of the first part (§2–3) of this paper is to defend the definability thesis, and its consequences for metaphysical indeterminacy. In physical theories that admit of a formulation based on an algebra of quantities or observables, the definability thesis leads to a specific form of strong naturalism that I call 'algebraic naturalism'. Algebraic naturalism has a number of interesting consequences, and I choose to focus on one in particular: that it can be used to counter the threat of deep metaphysical indeterminacy. The goal of the second part (§4–5) of this paper is demonstrate how.

In more detail, in §2, I spell out the definability thesis, using Newtonian particle mechanics as an example. In doing so, I demonstrate the coherence of algebraic naturalism as a norm governing interpretations of classical physics. In §3, I discuss how algebraic naturalism can be applied to QM. In §4, I introduce the machinery required to make precise claims about metaphysical indeterminacy, where I discuss the BW model's ersatzist characterisation [2] of metaphysical indeterminacy as a *sui generis* form of modality. Finally, in §5, I trace the implications of algebraic

[5] The caveat in the parenthesis is motivated by the fact that we know that they physical theories we currently interpret are not complete and final. In practice, the ontological commitments we derive from theories are only valid with respect to what we take to be those theories' restricted domains of applicability. I am grateful to an anonymous reviewer for highlighting this point.

naturalism for claims of metaphysical indeterminacy in QM. I identify a significant assumption that goes into Skow's claim that QM might be deeply metaphysically indeterminate, as a result of the Bell-Kochen-Specker (BKS) theorem [3, 20]. By denying this assumption, which is derived from what I conjecture is too close an adherence to the metaphysics of Newtonian mechanics, I make available a simple alternative reading of QM on which, regardless of interpretation, it is necessarily not deeply metaphysically indeterminate.

I do not intend to claim that Skow's argument is incoherent or his position misguided; I merely claim that there are reasons to deny Skow's central assumption, and there is an alternative coherent position that does away with that assumption. As I elaborate in §5, I contend that, according to strong naturalism, physical theories come with their own standards of determinacy. Concluding that QM is metaphysically indeterminate by the standards of Newtonian mechanics is rather like concluding that Newtonian mechanics is metaphysically indeterminate by the standards of Aristotelian teleological cosmology because Newtonian mechanics does not specify the location of the centre of the universe. This, of course, does not rule out the coherence of such a view.

2 Properties in Newtonian mechanics

The ordinary-language meaning of 'property' is something like 'any attribute, characteristic or quality of an object'. Philosophical analyses of the concept of a property often begin with the impulse to capture this ordinary-language meaning. They then proceed to identify attributes, characteristics and qualities of properties themselves—for example, what is the metaphysical status of properties? What are the semantic characteristics of predicates that express properties? And so on.

Most treatments are primarily concerned with accounting for our everyday use of property talk. So discussions typically focus on domains of discourse that include coloured balls, or creatures with kidneys, or numbers. The formal structure of most of our everyday property ascriptions is relatively straightforward: we begin with entities which have or 'instantiate' properties but are not themselves properties (henceforth 'particulars') and analyse property ascription as a predication of the singular terms which refer to particulars. 'The ball named a is red' could thus be rendered as, say, Ra. Of course, not all properties are predicated of particulars—properties themselves can instantiate (higher-order) properties.

While it continues to be appropriate to think of property ascriptions in physics—call such properties 'physical properties'—in terms of predication, what is less clear is the manner in which we ought to restrict the class of properties that are apt for ascription to the domain of particulars of a physical theory. Not all well-formed predications count. For example, classical electromagnetism tells us that colour ascriptions are tied to the wavelengths of light reflected off an object. So if the object's dimensions are below a certain length, and thus incapable of such reflections, then it does not have a colour: redness is not a property that can be instantiated by, say, a classical point particle.

The general analysis of property ascriptions to particulars can be split into two components:

`The ontological question`: What kind of entity is a property?

`The identification question`: Given a collection of particulars (or properties), which first- (or higher-) order properties can be instantiated by them?

As the example above demonstrates, my concern in this paper is primarily with the identification question: from the perspective of physical theories, which properties are we licensed to ascribe to particulars?[6]

2.1 The identification question

Let us make a distinction between two types of property: determinables and determinates. These are, according to Wilson [34], 'type-level properties that stand in a distinctive specification relation: the 'determinable-determinate' relation.' Determinables are properties like colour, shape and momentum, while determinates are specifications of these properties, like 'red', 'square' and '5m/s'. Determinables can be predicated of either particulars or of other determinables, while determinates stand in a *sui generis* relation of 'specification' to determinables. In what follows, I will speak both of particulars instantiating determinables as well as their instantiating determinates, the latter merely a shorthand for a particular instantiating a determinable which is specified by the determinate in question.[7]

Physical theories (more broadly, frameworks for physical theories) such as Newtonian mechanics and quantum mechanics consist in the following entities or structures:

[6] For a classic survey of the literature that addresses the ontological question, see [24].

[7] Note that I have not, at this stage, switched to endorsing Wilson's determinables-based approach to metaphysical indeterminacy. I am simply helping myself to some general-purpose metaphysical machinery.

particulars (a first-order domain of discourse), first-order properties of these particulars (these will characterise the states), functions of these first-order properties (these will characterise the determinables) and value spaces for these determinables. On the standard formulations of Newtonian $n-$particle mechanics, for example, the particulars are the n particles, the states are ordered tuples of a privileged collection of first-order properties (i.e. tuples of positions and momenta), the determinables are quantities like kinetic energy, potential energy, angular momentum and so on, and the value space for each determinable is the space of determinates for each determinable.

A state is an ordered tuple of properties. Let us distinguish between what we might call an 'ontic' state from a 'kinematic' state:

`Ontic state`: A tuple whose entries exhaust the first-order properties of the particulars.

`Kinematic state`: A tuple whose entries exhaust the first-order *physical* properties of the particulars.

The kinematic state will be determined by details of the physical theory (how this is to be done is the subject of the next subsection). The ontic state could, in addition or instead, depend on antecedent metaphysical commitments. We can now refine our characterisation of strong and weak naturalism:

`Weak naturalism*:` The ontic state should not be inconsistent with the kinematic state.

`Strong naturalism*:` The ontic state is identical to the kinematic state.

2.2 Two routes to properties in Newtonian particle mechanics

In this section, I shift formal perspective on Newtonian mechanics, and re-express it in a form that facilitates comparisons with other physical theories. To that end, I introduce and discuss the algebraic approach to Newtonian particle mechanics. This formalism is sufficiently general as to apply to quantum mechanics as well, so it will allow me to discuss a proposal for how to extract the physically relevant properties and relations directly from a theory, rather than postulating them *a priori* and hoping for the best. This procedure can be demonstrated using a very simple system: an n-particle Newtonian system (described with respect to some fiducial spatiotemporal structure that constitutes an appropriate frame of reference).

Newtonian particle mechanics is a theory about the time-evolution of certain properties of a collection of particulars—the particles. Which properties? As it turns out, (relative) particle positions and linear momenta and functions thereof. One might ask whether this set exhausts all the physical or ontic properties one can associate with the particulars. For Newtonian dynamics, it does: all physical properties (energies, angular momenta, forces) are expressible as functions of positions and momenta.

The minimal structures we need to specify a physical theory are states, determinables (and their associated value spaces) and dynamics—call this a 'dynamical triple'. The kinematic state of a system is a tuple whose entries are first-order physical properties of the system. We can now ask the identification question with respect to physical systems: what are the kinematic states of an n-particle Newtonian system? I discuss two options.

2.2.1 Route 1

On the basis of commonsensical intuitions, we might propose that the kinematic states of a $n-$particle system are completely specified by the collection of the positions and velocities of the particles with respect to some fiducial body or structure (depending on your antecedent metaphysical commitments, this may or may not include points or regions of spacetime). But we would soon discover that Newtonian mechanics is not deterministic in such states. This is not unexpected: '...is deterministic' is not a monadic predicate. It is, rather, an m-place predicate whose inputs include the laws as well as particular determinables.[8] The question is: which determinables? Since Newtonian mechanics is not deterministic with respect to just velocities, we might consider adding another determinable, say linear momentum with respect to the same fiducal structure as we made our position ascriptions. Lo and behold, Newtonian dynamics is completely deterministic in these determinables. Further, it is completely deterministic in determinables like angular momentum, kinetic and potential energy, since all these determinables can be deduced from positions and linear momenta. Denote the value spaces of physical determinables as $\mathbb{D}$; let the value space of Newtonian determinables be $\mathbb{D}_N$.

At this point, we might wonder if there is a more systematic way of identifying physical determinables than iterative guesses. Notice that, while guessing, we did not make use of all of the structure at our disposal from the theory. In particular, we did not use any information about the determinables and their values until we started to test

[8] i.e. Newtonian mechanics is deterministic in positions and momenta and energies and angular momenta and...

whether positions and momenta were enough to characterise all physical determinables. We were able to verify that this was the case by treating the determinables as functions that mapped states onto values. For example, given some state, which specified the position coordinates, q_i and the momentum coordinates, p_i, we could calculate its kinetic energy using the formula $\sum \frac{p_i^2}{2m}$. Understood set-theoretically, a determinable in an n-particle Newtonian theory is an ordered $6n + 1$-tuple, where the first $6n$ characterise the kinematic state and the last entry is the value of the determinable (i.e. the determinate that specifies the determinable).

The idea that determinables are functions from a state space to some value space forms the basis of the Hamiltonian approach to Newtonian mechanics. The time evolution (i.e. dynamics) of states is represented by time-indexed paths through a space of kinematic states—by analogy with how a moving picture can be composed of time-indexed snapshots. If one equips the set of kinematic states with some further structure—topological, smooth and symplectic, and a privileged determinable known as the Hamiltonian function—then one can define a deterministic (i.e. complete and non-intersecting) family of paths in this kinematic state space, each of which represents the time evolution of a system from some initial to some final state. Since all physical determinables are just functions of these determinables, we have enough to specify the states, determinables and dynamics: we have a dynamical triple.

2.2.2 Route 2

Perhaps unexpectedly, the very same dynamical evolution of states can be encoded by structure on the space of determinables. In other words, since all we need is a particular structure on the dynamical triple, it is possible to associate structure either with the space of determinables or with the space of states, as long as between them, they give rise to the same structure on the dynamical triple.

We treated each determinable as a function from instantaneous 'snapshot' states to determinates: a device that takes in $6n$ inputs—$3n$ components of position and $3n$ of linear momentum—and outputs a determinate. Since these devices are mathematical objects, we can ask questions about the mathematical structure of the space of these objects, independently of their inputs. It turns out that if the set of these determinables—mathematically speaking, these are smooth functions on a manifold—is equipped with the structure of a so-called *Poisson algebra*, then the dynamics can be encoded using structure on the space of determinables, rather than kinematic states. Let $\mathfrak{A}$ denote, in general, an algebra of determinables, and let $\mathfrak{A}_N$ specifically denote Poisson algebras of Newtonian determinables.

Formally speaking, we can propose a new type of 'state' (this choice of terminology is justified below): a function from the Poisson algebra of determinables to the value space. Call these states 'algebraic states.' Algebraic states, qua functions, are ordered pairs, the first entry of which is a determinable, and the second of which is an element of the value space. We now have a new dynamical triple: a Poisson algebra, a set of algebraic states, and dynamics. These triples characterise the same theories as our previous triple. We can contrast algebraic states with the kinematic states we worked with on route 1:

`Algebraic state`: Properties of elements of an algebra of determinables.

`Kinematic state`: Physical properties of the particulars to which a theory is committed.

Note that, in identifying a state as kinematic, we do not have to commit to any specific vision of what those particulars are: they could be localised particles, or fields in spacetime, or a state space or a configuration space or indeed any other objects that can stand in an instantiation relation to physical properties (recall the orthogonality of naturalism and realism as discussed in footnote 1). Now prima facie, the algebraic and kinematic properties need not have anything to do with each other; they are predicated of different entities. And it is certainly far from clear how algebraic states might be properties of the world: whatever the particulars of our world are, they are certainly not Poisson algebra elements!

So if we wish to determine the ontic states via the collection of algebraic states, we need to figure out how to understand algebraic states' claims as claims about properties of worldly particulars. In other words, we need to associate a unique kinematic state with each algebraic state. We can then treat the latter as a means of using our physics to identify the former. This is what motivates my claim that the unmodified formalism of quantum mechanics can naturally be associated with an abstract characterisation of property ascription.

I added the qualifier 'unmodified formalism' advisedly. In the quantum mechanical context, 'interpretation' usually means 'a way to resolve the measurement problem', so Bohmian mechanics, GRW, Everett and Copenhagen are treated as interpretational variants of the same theory. I would like to hold on to this standard terminology, but distinguish from within this class, those interpretations which modify the formalism by adding some structure, either itself formal or metaphysical (or both), and those that do not. So Bohmian mechanics and dynamical collapse theories count as modified, while the Everett interpretation qualifies as an unmodified interpretation.

The account of how property-ascriptions work can be given an abstract characterisation in terms of the algebra of determinables for all interpretations. However, each interpretation might then further characterise what properties, either by modifying the algebraic account of properties (as in GRW) or by supplementing it (as in Bohmian mechanics). Let us use Bohmian mechanics as an example: the conception of 'position property' that the Bohmian works with is different from the one that the non-Bohmian works with. Let us distinguish between a determinable as a property represented by an element of the algebra of observables (call this property$_1$), and a property as an ascription of a determinable that is not an element of that algebra (call this property$_2$). There is an equivocation, in the Bohmian context, between those two: the 'position' of the particle is generally taken to be determinable in the second, but not the first, sense. Take any generic, uncollapsed quantum state. There is a non-zero probability that the 'position$_1$' of that particle, in the sense of 'eigenvalue of the position observable' is distinct from the position$_2$ of that particle. So for the Bohmian, the complete ascription of determinables to objects extends beyond those determinables that are represented by elements of the algebra of determinables. It is therefore an interpretation-specific question whether the algebraic property ascriptions perhaps conflict with some metaphysically deeper ascriptions, as they arguably do in the Bohmian case. What I present here is an account of how to understand property ascriptions in cases where: (i) the formalism is unmodified and (ii) modifications do not lead to a conflict between the algebraic and non-algebraic method of property ascription.[9]

The strategy involves the following two moves:

`Step 1`: Establish the definability of algebraic states.

`Step 2`: Link algebraic states to kinematic states.

On establishing this link, we can fully embrace the shift in perspective, and understand the state of the world as being a property of an algebra of quantities.

2.3 Establishing definability

On a standard account of how to provide a semantics for a first-order language $\mathcal{L}$, with signature Σ, one starts by defining a *model* $\mathfrak{S}$, which consists of a domain of discourse, D and an interpretation function, I, which is a function that maps constants in Σ to

[9]I am grateful to an anonymous referee for highlighting to me the importance of this distinction.

elements of D, and $n-$place predicate letters and $(n-1)-$place function letters in Σ to $n-$tuples over D.

For the calculus to be rich enough to be useful, some more machinery is needed. In particular, one needs to be able to speak unambiguously about properties and relations. This is where the concept of *definability* comes in. Consider a model $\mathfrak{S}_1$, and the open formula Gx, where x is a variable and G a predicate letter. A variable assignment, α is function from the set of variables in that language to the domain of discourse. The open formula Gx is satisfied by any variable assignment α that maps x to elements of tuples picked out by the interpretation function acting on G. For example, if G is supposed to capture, say, the property of being green, then the formula is satisfied by the variable assignment that maps x to any green element in the domain of discourse.

If there is more than one green element, then it does not matter which one we pick, since the variable assignment will satisfy the formula if it picks out any of them. So permutations of the green elements will not affect whether or not the variable assignment satisfies the formula. Roughly speaking, then, in order to speak unambiguously about the property of greenness, the model $\mathfrak{S}_1$ needs to be invariant under transformations that permute the green elements. But it should not be invariant under transformations that permute green and non-green elements. This means that greenness, as a property, is defined only up to permutations of green elements.

It is not enough to talk only about permuting elements of the domain D; we also need to think about how this might affect the interpretation function, I. Consider two models, $\mathfrak{S} = \langle D, I\rangle$ and $\mathfrak{S}' = \langle D', I'\rangle$, and a function $h : D \to D'$. Call such a function a *homomorphism* of the model $\mathfrak{S}$ iff for each $n-$tuple of D and,

(1) for each $n-$place predicate P: $\langle a_1, a_2...a_n\rangle \in P^{\mathfrak{S}}$ iff $\langle h(a_1), h(a_2)...h(a_n)\rangle \in P^{\mathfrak{S}'}$.

(2) for each $n-$place function f: $h(f^{\mathfrak{S}}(a_1, a_2...a_n)) = f^{\mathfrak{S}'}(h(a_1), h(a_2)...h(a_n))$.

If $\mathfrak{S} = \mathfrak{S}'$, then h is called an *endomorphism*. A bijective endomorphism is called an *automorphism*. We can now speak more precisely of what it is to be able to define an $n-$place relation (adapted from [12, p. 98-99]):[10]

[10]This is not generally taken to be a full characterisation of definability in a model, since it only provides a necessary condition. It is, in fact, a lemma that follows from the more general definition in terms of satisfaction of open formulae (see, e.g. [12, p. 90-91]. Since the argument in this paper hinges on a failure of this necessary condition, the full characterisation of definability in a model is not required.

`Definability in a model` (relations): An $n-$place relation R is **definable in a model** $\mathfrak{S}$ only if the set of n-tuples of elements of the domain of discourse corresponding to R that is picked out by the interpretation function I is invariant under the automorphisms of the model, $\mathfrak{S} = \langle D, I \rangle$.

2.3.1 Undefinability and indeterminacy

Definability is a semantic notion, while metaphysical indeterminacy is (at least allegedly) a metaphysical one. So in linking these two notions, we need to be confident that any alleged metaphysical indeterminacy is, in fact, metaphysical, and not, as is sometimes argued (see e.g. [11, 27, 28, 32]) the residue of our semantic decisions. In this subsection, I provide an example to illustrate how the two notions are linked.

In short, the proposal is as follows. The domain of discourse of any model is a set whose elements are (putatively) worldly entities. Therefore, a specification of two models defined over the same domain immediately links those two models to the same collection of worldly entities. Where the models differ is in their attribution of properties to the domain elements. Each model is the semantic correlate of a possible world, where the correlation is established via the model's interpretation function. So each of the two models defined over the same domain is correlated with a possible world, i.e. a way the domain elements could be. We describe 'a way those elements could be' via an attribution of properties and relations to those elements. So if one model attributes certain properties to some elements while another does not, then the two models disagree over the way those elements could be. Now, if we stipulate that one of the models (the one with a richer attribution of properties) determines the standards of determinacy *in the world that it represents*, then by those standards, the other model leaves certain determinables indeterminate in the world that it represents. This leads to what we might call the definability-determinacy proposal:[11]

`Definability-determinacy proposal`: a model fails to attribute a determinate to a determinable in the world represented by that model if that determinable is undefinable.

Definability is always specified relative to some model. We can see clearly how antecedent commitments might play into a choice of model: start with an intuitive idea of what a collection of objects, some of which are, say, green, might look like and then reverse-engineer the appropriate model. The freedom in choice of variable

[11]See [19] for a discussion of how the determinacy-definability proposal can be used to argue against the existence of spatial points in noncommutative geometries.

assignment, together with the freedom in choice of permutations of domain elements must conspire to leave the tuple of green elements untouched.

Consider, now, a model $\mathfrak{S}_2$ on the same domain as $\mathfrak{S}_1$, in which some of the green elements have the further property of being emerald coloured, and the rest are seafoam coloured. Assume that the predicates E : '...is emerald green' and S:'...is seafoam green' are definable in $\mathfrak{S}_2$. Intuitively speaking, $\mathfrak{S}_2$ is better at resolving greeness than $\mathfrak{S}_1$; there will be permutations on the domain that leave the predicate '...is green' invariant in $\mathfrak{S}_1$ that do not leave '...is seafoam' invariant in $\mathfrak{S}_2$. Since the permutations are defined on the same domain, let us assume that '...is seafoam' is not definable in $\mathfrak{S}_1$.

But just because it is not definable in $\mathfrak{S}_1$, it does not mean that our metaphysics should immediately jettison the property of being seafoam green. If $\mathfrak{S}_1$ is a model of our best physics, one might simply argue that our best physics is not maximally resolving when it comes to all genuine properties of objects. Perhaps being seafoam green is a property of the ontic state but not the kinematic state, and it is the physical theory that needs to be modified or replaced. In this case, with respect to this particular physical theory, one might propose that it is indeterminate whether some object is seafoam green or not: on some variable assignments, seafoam greenness is ascribed to some particular, while on others it is not. If we believe that $\mathfrak{S}_2$ is a model of a metaphysically more privileged, or more highly-resolving theory, then properties that are undefinable in $\mathfrak{S}_1$ might nonetheless be properties of ontic states to which physical theory of which $\mathfrak{S}_1$ is a model is blind.

In our example here, $\mathfrak{S}_2$ is playing two roles. First, it is making intelligible questions regarding certain properties (e.g. 'is object a seafoam green?') and second, it is providing an answer to those questions. But, according to the proponent of metaphysical indeterminacy, it is possible for a model to play first of those roles without playing the second. It is possible, accordingly, to use $\mathfrak{S}_2$ to pose a question, but insist that answers can only come from $\mathfrak{S}_1$.

If $\mathfrak{S}_1$ then represents some possible world, W_1, in the standard way (i.e. a interpretation function maps constants to domain objects, predicates to tuples of domain objects, and so on) then it is metaphysically indeterminate, by the standards of $\mathfrak{S}_2$ whether a is seafoam green in that world. Undefinability is thus the semantic correlate of metaphysical indeterminacy. More precisely, if there exists some property that is undefinable in $\mathfrak{S}_1$, but definable in $\mathfrak{S}_2$, then by the standards of W_2, i.e. the world possible according to $\mathfrak{S}_2$, certain properties are metaphysically indeterminate in W_1. Further (and this is the underlying impulse of the Barnes-Williams model discussed in §4) the world W_2 can be seen as a 'precisificationally possible' world of W_1, and metaphysical indeterminacy can be defined in terms of the correspondence between multiple precisificationally possible worlds and a particular nomologically possible world.

In what follows, I argue that one way in which the strong naturalist impulse manifests itself is in resolving an arbitrariness charge of the following form: why should the standards of metaphysical determinacy for W_1 be given by $\mathfrak{S}_1$, rather than $\mathfrak{S}_2$, or, indeed, any other $\mathfrak{S}_n$? The strong naturalist breaks this underdetermination of standards by stipulating that standards of determinacy should be internal to the theory in question. Indeed, one might even take this view to be definitional of strong naturalism.

However, one might argue that there is a plausible competitor for standards of metaphysical determinacy: the macroscopic world. One might argue that referential standards are determined by the macroscopic world, in the sense that we use the same semantic machinery for reference to unobservables as we do for reference to observables. So there is at least a *prima facie* plausible reason to consider doing the same for standards of determinacy. And, of course, since on this view, we are determining our standards of determinacy by appeal to a domain of the physical world that we take to be described by quantum mechanics, this view undoubtedly earns its naturalist stripes.[12]

Although this position might sound appealing, on further analysis, it reveals itself to be somewhat unattractive, if one assumes that our quantum theories are better than (i.e. more fundamental, or predictively or descriptively accurate than) our classical theories. Under this sort of assumption, the mere fact quantum theories give rise to emergent quasi-classical dynamics strongly suggests that the semantics appropriate to classical dynamics (i.e. the macroscopic world) should be derivable from, or at least consistent with, the semantics appropriate to the better theory. At the very least, it suggests we should adopt a norm according to which no theory should be taken to license ontological claims that are in conflict with those licensed by what we take to be a better theory. Call this the 'scientific progress norm'.

Now, of course, this is a defeasible suggestion. And, indeed, if we could not make sense of the semantics of quantum theories without appeal to the classical, then the algebraic naturalist would be in trouble. But, as I demonstrate in what follows, the algebraic naturalist has a perfectly sensible semantics. Given this, it seems ill-advised to violate the scientific progress norm without some extremely compelling reason to

[12] I am grateful to an anonymous reviewer for suggesting this alternative.

do so.

2.3.2 Definability of functions

Recall the second route to the dynamics of an n-particle Newtonian system, via algebraic states. Suppose we treat the algebra of determinables as the domain of discourse. In that case, there are some relations that need to preserved under a combination of a change of variable assignment, interpretation function and permutation of elements. The physics tells us that this domain has the structure of a Poisson algebra. From a mathematical perspective, all this means is that the elements of the domain stand in certain well-defined relations to each other and if we permute these elements in such a way as to preserve these relations, we end up with the same Poisson algebra.

Since we are interested in more than just the algebra of determinables (in particular, we also care about the values of these determinables), we need to enlarge our domain to include all the elements of all the value spaces of the determinables of interest. Call this new domain D_p. We can then extend our analysis of definable relations to encompass definable functions:

`Definability in a model` (functions): An $n+1$-place function f is **definable in a model** $\mathfrak{S}_p$ only if the set of $n+1$-tuples of elements of the domain of discourse corresponding to f that is picked out by the interpretation function I is invariant under the automorphisms of the model, $\mathfrak{S}_p = \langle D_p, I \rangle$.

A little care is required at this point. Not all Poisson algebra automorphisms leave the outputs of algebraic states unchanged. Consider, for example, the Poisson algebra automorphism induced by $q \to -p$ and $p \to q$. In general, this will change the form of the Hamiltonian function $H \to H'$, as a result of which the same algebraic state might map H and H' to different determinates. Unless the algebraic states take account of this, in general, they will map determinables to different determinates before and after an automorphism.

In order to correct for this, we need to ensure that, for every automorphism performed on the algebra of determinables, there is a corresponding automorphism on the space of algebraic states such that the each determinable is mapped to the same determinate before and after the automorphism, i.e. if ρ, ρ' are algebraic states, then if some transformation $s(H) = H'$, then there exists some transformation $t(\rho) = \rho'$ such that $\rho(H) = \rho'(H')$. Call the pair of such transformations $\langle s, t \rangle$ a 'paired automorphism'.

This is standard practice when dealing with active dynamical symmetries in physics. We can think of an automorphism of the Poisson algebra s as *generating* an automorphism t on the algebraic state space. All paired automorphisms are automorphisms of the model, but not vice versa. Thus algebraic states, qua functions on D_p, are definable only if the pairs picked out by the interpretation function are invariant under the automorphisms of the model. In what follows, 'automorphisms of the model' will always mean 'paired automorphisms', and I will assume that for every algebra of determinables, it is possible to find a space of algebraic states invariant under a set of paired automorphisms generated by the algebra.

2.4 Linking algebraic and kinematic states

The procedure behind linking kinematic and algebraic states can be demonstrated using a classical simple harmonic oscillator. This is a Newtonian one-particle system about which we can make determinate claims about position, velocity, kinetic and potential energy (among many other properties). The centre-of-mass of a pendulum, for example, is a simple harmonic oscillator. Imagine taking a snapshot of the system when it is at point halfway between its extreme and equilibrium points. The kinematic state of this system is specified by its position, q and momentum p relative to, say, its equilibrium position, q_0. If we wanted to know its kinetic energy, for example, we would act on this state, $S_k = (q, p)$ with the kinetic energy function K: $\frac{p^2}{2m}$. This function (determinable) maps us to some number, say 5 (determinate). Schematically, $K(S_k) = 5$, and the natural translation of this is 'the state has the property that the kinetic energy maps it to 5 units'(less awkwardly, we say that the system has a kinetic energy value of 5 units).

But we can switch the notation around in a revealing way. Consider functions of the form $S_a : \mathfrak{A}_N \to \mathbb{D}_N$. Writing $S_a(K) = 5$ suggests that we can think of the state as mapping a function to a real number (i.e. $\mathbb{D}_N \subset \mathbb{R}$), a natural interpretation of which is 'the kinetic energy has the property that the state S_a maps it to 5 units.' (In this case, it is no less awkward to say that the kinetic energy has a state value of 5 units). We can now immediately define a bijection f from the set $\{S_a$: S_a is an algebraic state$\}$ to the set $\{S_k : S_k$ is a kinematic state$\}$. If, for all determinables, F, $S_a(F) = F(S_k)$ then $f(S_a) = S_k$ and $f^{-1}(S_k) = S_a$. Thus, to each algebraic state, we link the kinematic state that agrees on the determinates specified by all determinables. Note that, since this definition makes reference to all determinables, the bijection f is invariant under automorphisms of the algebra of determinables.

2.5 Summary

I can now state precisely the positive proposal for how the algebraic naturalist can identify ontic states directly from physical theories: begin with the appropriate algebra of physical determinables, construct a space of definable algebraic states and then establish a bijection between algebraic and ontic states.

Let $v_i : \mathfrak{A} \to \mathbb{D}$ be functions from the algebra of determinables to the value space; call them 'value functions'. In Newtonian mechanics, there is an algebra isomorphism between the space of value functions and the space of algebraic states.[13] Therefore, in Newtonian mechanics, an algebraic state is uniquely associated with a value function and vice versa. Consequently, the bijection, f, from the space of algebraic states to kinematic states can be composed with this algebra isomorphism to render isomorphic all three of these spaces. Consider the following positions regarding states.

`Ontic-kinematic`: The ontic state is the kinematic state.

`Kinematic-algebraic`: The kinematic state is an algebraic state.

`Algebraic-value`: The algebraic state is a value function.

`Ontic-value`: The ontic state is a value function.

I propose that `ontic-algebraic`, i.e. the conjunction of `ontic-kinematic` and `kinematic-algebraic`, adequately captures `strong naturalism`*, across both Newtonian and quantum mechanics. This conjunction is the precise realisation of algebraic naturalism. In what follows, I explore the consequences of understanding `strong naturalism`* in this way. What I end up with is not so much a knock-down argument in favour of `ontic-algebraic`, as a demonstration that `ontic-algebraic` is a consistent and plausible position that provides an alternative to the standard `ontic-value`. Whether or not it is more attractive will depend on further background commitments. If `strong naturalism`* is one of those commitments, then, as I demonstrate using the case study of quantum mechanics, `ontic-algebraic` is a more attractive norm, in general, than `ontic-value`. But in order for us to even make that assessment, and understand how it might relate to further background commitments, we need a precise articulation of the position. That is the primary goal of this part of the paper.

In Newtonian mechanics, the coincidence (up to isomorphism) of the space of value functions with the space of algebraic and kinematic states obscures the fact that these are all conceptually distinct spaces. In that domain, the truth of `algebraic-value`

[13] Note that the space of value functions is different from the space of determinates.

means that `ontic-algebraic` entails `ontic-value`. So `ontic-value` also captures **`strong naturalism`***, and we can use any one out of the kinematic state, the algebraic state, or value function, to deduce a candidate ontic state of a Newtonian system, and we will not encounter a conflict. What this means is that, from the perspective of the metaphysics of properties, one can afford to be a little sloppy in asserting which of these spaces counts as the space of determinables. In a sense, they all do. In the next section, we will discover that quantum mechanics is not quite so forgiving: in that domain, `algebraic-value` is false, so the strong naturalist commitment is best captured by `ontic-algebraic`, not `ontic-value`.

3 Quantum mechanics and the Bell-Kochen-Specker theorem

3.1 Two routes to properties in quantum mechanics

There are striking similarities between the algebraic approaches to quantum mechanics and Newtonian mechanics. As with Newtonian mechanics, all we need to specify is a dynamical triple: states, determinables and dynamics. The difference is that the states[14] form a Hilbert space rather than a symplectic manifold, the determinables form a von-Neumann algebra of self-adjoint operators instead of a Poisson algebra of smooth functions, and the dynamics is given by Schrödinger's equation instead of Newton's laws. In addition, because of the probabilistic nature of quantum mechanics, there is an extra step, mediated by the so-called 'Born Rule', between identifying determinates of determinables and determining (the statistics of) experimental outcomes. But none of these differences affect the procedure for identifying algebraic states. As before, algebraic states are properties of the algebra of determinables, but this time, their outputs are interpreted as expectation values of determinates rather than determinates themselves.

Let us shift focus to a simple two-particle system: a non-relativistic quantum mechanical system with two independent particles called *2-spinors*. Consider a determinable whose value space is discrete. In quantum mechanics, spin, a form of angular momentum, is one such determinable. Unlike in Newtonian mechanics, this angular momentum is not a function of any other determinables. We treat it as a primitive property of 2-spinors.

[14]We restrict attention to so-called 'pure' states.

3.1.1 Route 1

Consider a quantum system whose kinematic state is $|\psi\rangle$. Let the (self-adjoint) operator for spin in the ith direction be σ_i, and let us restrict our attention to worlds with Euclidean three-dimensional spatial geometries. To determine which quantum states correspond to the in-principle exactly measurable spin properties, we need to solve the following equation, known as an *eigenvalue equation*:

$$\sigma_x|\psi\rangle = \lambda|\psi\rangle \tag{1}$$

where λ is a real number. We can deduce that according to QM, λ, can take one of two values, known as *eigenvalues*: ±1, with respect to some determinable σ_x, i.e. σ_x is specified by ±1. The states for which eq. 1 is true for $\lambda \in \{-1, +1\}$ are called *eigenstates* of σ_x. Every quantum mechanical state can be expressed as the *convex sum* of eigenstates of some determinable. That is to say, if B is some determinable whose eigenstates are $|\psi_B^i\rangle$, then any state $|\psi\rangle$ can be expressed as:

$$|\psi\rangle = \sum_i \lambda_i|\psi_B^i\rangle \tag{2}$$

where $\sum_i \lambda_i = 1$ (this is the 'convexity' condition).

Suppose the system is in an eigenstate of σ_x, call it $|\psi_{\sigma_x}\rangle$. If that state happened also to be an eigenstate of some other operator, call it B, then the following equation would be satisfied by some $\lambda_B \in \mathbb{R}$:

$$B|\psi_{\sigma_x}\rangle = \lambda_B|\psi_{\sigma_x}\rangle \tag{3}$$

If we act on this new state $B|\psi_{\sigma_x}\rangle$, with σ_x, it would yield the following:

$$\sigma_x \circ B|\psi_{\sigma_x}\rangle = \lambda\lambda_B|\psi_{\sigma_x}\rangle \tag{4}$$

where $\circ$ represents the algebraic product on the algebra of quantum determinables.

This is an eigenvalue equation for the new operator $\sigma_x \circ B$. The algebraic product of two determinables corresponds to a determinable, as does a linear combination of determinables. We can deduce the eigenvalue equation for the operator $B \circ \sigma_x$:

$$B \circ \sigma_x|\psi_{\sigma_x}\rangle = \lambda_B\lambda|\psi_{\sigma_x}\rangle \tag{5}$$

Since λ and λ_B are both real numbers, and multiplication of real numbers is blind to the order of multiplication, we can infer that the order of performing the measurements on σ_x did not affect the outcome. Operators that have this property are known as *commutative* operators. Since quantum states are linear, and expressible as the convex sum of eigenstates of any determinable, we can generalise the inference that the order of measurement does not affect the outcomes of measuring commuting determinables on all states, not just eigenstates. A measurement of a determinable on a state yields a probability distribution over eigenvalues. The special case with which we began, of a measurement on an eigenstate yields a trivial probability distribution that assigns 1 to the corresponding eigenvalue and zero to all others. This special case underpins the so-called *eigenstate-eigenvalue link*, which elevates this mathematical claim to a metaphysical one:

`Eigenstate-eigenvalue link` (EEL): A quantum system has a determinate value of a determinable if and only if its state vector is an eigenstate of the determinable's associated operator with eigenvalue equal to its determinate.[15]

The EEL provides a precise realisation, in the quantum mechanical context, of the circumstance in which `ontic-value` can be satisfied: the determinate (eigenvalue) corresponding to some determinable (operator) acting on a quantum state exists only when that state happens to be an eigenstate of that determinable (operator).

3.1.2 Route 2

Analogously to the Newtonian case, we can begin with an algebra of quantum determinables, call it $\mathfrak{A}_Q$. Once again, a value function is a linear map, this time from the algebra of *quantum* determinables to a codomain $\mathbb{R}$, which we can intuitively think of as mapping a determinable to its determinate.[16] Now, according to the EEL, a value function for a quantum determinable exists only if the state is an eigenstate of that determinable. In the more general quantum context, we can define a distinct object, call it an *expectation value function*: a linear map, from the algebra of quantum determinables to a codomain $\mathbb{R}$, intuitively to be thought of as mapping a determinable to the expectation value of its determinable. This function exists for all quantum states, regardless of whether they are eigenstates of any operator.

Define a quantum algebraic state as map from a determinable to its expectation

[15]A quick note: this equality is relative to a choice of units of measurement. Eigenvalues are just pure numbers, whereas determinates have units. So, for example, the claim that some state is spin $1/2$ in the x direction is made relative to a system of units where $\hbar = 1$

[16]In this paper, I restrict attention to algebraic states on the domain self-adjoint operators, so the codomain is $\mathbb{R}$.

value; in other words as an expectation value function.[17] The distinction between value functions and expectation value functions is absent in the Newtonian context; every value function there is trivially also an expectation value function, with respect to the probability distribution that assigns 1 a specific value of the determinable and 0 to all others.

We thus establish a direct link between algebraic and kinematic states, by associating, with each kinematic state an algebraic state that maps all determinables to the same expectation values as those determinables map kinematic states to; in other words a bijection between algebraic states and *expectation value function*. As in the Newtonian case, this bijection is not hard to find. Consider functions of the form $S'_a : \mathfrak{A}_Q \to \mathbb{D}_Q$. Define a bijection g from the set $\{S'_a$: S'_a is an algebraic state$\}$ to the set $\{S'_k : S'_k$ is a kinematic state$\}$. If, for all determinables, F, $S'_a(F) = F(S'_k)$ then $g(S'_a) = S'_k$ and $g^{-1}(S'_k) = S'_a$. In other words, the bijection maps algebraic states to kinematic states that agree on the expectation values of all determinables.

Clearly, we can define analogues of the Newtonian value functions, which will act on elements of the algebra of quantum determinables and map them to elements of the space of eigenvalues. For a spin system, the space of eigenvalues is a subset of the rational numbers, $\mathbb{Q}$. Note the difference in codomain compared to algebraic states.

3.2 The Bell-Kochen-Specker theorem

So far, everything we have done in the Newtonian domain has transferred relativley unscathed to the quantum domain. In Newtonian mechanics, algebraic states are just value functions (consequently, the codomains of algebraic states and value functions coincide, so we use the same label, $\mathbb{D}_N$). In QM, we can demonstrate that value functions and algebraic states are distinct. As we just saw, the codomain of value functions is the space of eigenvalues, which is usually a subset of the codomain of algebraic states. The BKS theorem demonstrates that some interesting value functions are not invariant under any paired automorphisms; these functions are therefore not definable, so are not candidates for algebraic states.

We have been working under the assumption, so far, that attributing a vectorial property like spin to a two-particle quantum system is as straightforward as the equivalent attribution of the vectorial property of angular momentum to a Newtonian system. But the BKS theorem makes this assumption untenable in QM. In the metaphysics literature, the BKS theorem has often been presented in the context of

[17] In what follows, I drop the qualifier 'quantum' or 'Newtonian' when talking about states; it will be clear from context.

a three-dimensional state space a single spin-1 particle [4, 30]. Whilst technically unimpeachable, this example is not as straightforward and clear as Mermin's example of the 4-dimensional state space of a pair of uncorrelated spin-1/2 particles [23]. Thus, Mermin's example applies to all state spaces of dimension higher than three. The original theorem is slightly more general since it also applies to state spaces of three dimensions. But what we lose in generality, we make up in clarity.

The demonstration that certain value functions are not definable proceeds as a reductio ad absurdum. We assume that, for some collection of determinables, there exist definable value functions onto appropriate determinates. We then demonstrate, using the BKS theorem, that these functions are not invariant under any paired automorphisms. From this, we conclude that these value functions are not definable.

The states of a single 2-spinor can be represented by a column vector with two entries. The $+1$ state with respect to, say, the $z-$ direction (i.e. the eigenstate with eigenvalue $+1$) can be represented by $\begin{pmatrix}1\\0\end{pmatrix}$; the -1 state as $\begin{pmatrix}0\\1\end{pmatrix}$. An arbitrary state can be represented by a convex linear combination of these states. The states of a system that consists of two independent 2-spinors can, accordingly, be represented as convex linear combinations of column matrices with four entries: the top two entries correspond to one of the particles, the bottom two to the other. Thus, a state in which, for example, the first spinor is spin $+1$ and the second -1 can be rendered as: $\begin{pmatrix}1\\0\\0\\1\end{pmatrix}$.

In general, we can construct operators corresponding to spin measurements along any direction (we restrict our interest to three orthogonal directions since any other direction can be expressed as a linear combination of these directions): σ^i_μ, where $i = 1, 2$ labels each spinor, and μ is the spatial direction (i.e. the x- or $y-$ or $z-$axis). Each σ^i_μ has eigenvalues ± 1. Each set of σ^1_μ and σ^2_μ are mutually noncommuting. We can represent these spin determinables using 4×4 matrices. For example, the 'first spinor's spin in the $z-$direction' determinable is represented by $\sigma^1_z = \begin{pmatrix}1 & 0 & 0 & 0\\0 & -1 & 0 & 0\\0 & 0 & 0 & 0\\0 & 0 & 0 & 0\end{pmatrix}$.
We can check that this is an appropriate representation by solving the eigenvalue equations. For example:

$$
\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix} = +1 \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix}
$$

This tells us that the spinor is in an eigenstate of spin, with eigenvalue $+1$, as expected.

We can construct other determinables from the spin determinables with which we began; after all, they form an algebra, so any algebraic products of these determinables are also determinables. In particular, we can construct the following nine:

$$
\begin{array}{ccc}
\sigma_x^1 & \sigma_x^2 & \sigma_x^1\sigma_x^2 \\
\sigma_y^2 & \sigma_y^1 & \sigma_y^1\sigma_y^2 \\
\sigma_x^1\sigma_y^2 & \sigma_x^2\sigma_y^1 & \sigma_z^1\sigma_z^2
\end{array}
$$

It is easy to demonstrate that the triple of operators in each row constitute a mutually *commuting* set. And similarly each column. Consequently, according to `Commuting`, the order of measurement of these determinables does not affect the measurements. So all determinables in each triple share eigenstates. Consider a system in an eigenstate of the operators in one triple.

According to the axioms of QM, if a state is in an eigenstate of some operator, the measurement outcome *for that one specific property* will be the corresponding eigenvalue with probability 1. Let us try to ascribe determinates to each of the nine operators. Notice that, if the possible eigenvalues of σ_μ^i are ± 1, then the product of eigenvalues for the top row cannot be anything other than $+1$, whatever the state of the system (after all the eigenvalues associated with an operator are state-independent facts about that operator). And similarly the other two columns' and the first two rows' operators' eigenvalues, when multiplied yield $+1$. The product of eigenvalues of the right-most column cannot be anything other than -1.

Since the three rows' products are all $+1$, the product of all nine eigenvalues must be $+1$. But the similar inference about the three columns' products is that the product of the nine eigenvalues must be -1, which is impossible. There is no assignment of eigenvalues that can satisfy these competing requirements. Therefore there is no way to consistently assign determinates to this entire set of determinables. Therefore, since a quantum value function maps each determinable to an eigenvalue, there is no possible spin value function for this collection of determinables.

We relied on the algebraic structure of the algebra of determinables to define the 'product determinables' like $\sigma_x^1\sigma_x^2$. These determinables are, therefore, invariant under the automorphisms of the algebra of determinables. We can identify the paired automorphisms to which these algebra automorphisms give rise using a similar prescription as in the classical case: each automorphism on the von-Neumann algebra of determinables is paired with an automorphism on the Hilbert space of algebraic states in such a way that expectation values of determinables remain unchanged.

In the classical case, for each automorphism s of the Poisson algebra, we could easily find an automorphism, t, of the space of algebraic states such that $\langle s, t \rangle$ was a paired automorphism. And each of those algebraic states was uniquely associated with a value function, because of the truth of `algebraic-value`. The BKS theorem demonstrates that `algebraic-value` is false in QM: it is impossible to find automorphisms on the space of value functions that give rise to paired automorphisms generated by von-Neumann algebra automorphisms.

The absence of the appropriate paired automorphisms renders certain properties undefinable. As we saw in §2.3.1, undefinability is the semantic correlate of indeterminacy. The BKS theorem demonstrates that value functions are not definable, so any states deduced from value functions will also be undefinable. This suggests the following two strategies:

`Ontic-algebraic strategy`: Deduce the (determinate) ontic state from the (definable) algebraic state.

`Ontic-value strategy`: Deduce the (indeterminate) ontic state from the (undefinable) value function.

In §5, I argue in favour of the ontic-algebraic strategy over Skow's ontic-value strategy. That argument relies on a precise characterisation of metaphysical indeterminacy, to which we now turn.

4 Quantum metaphysical indeterminacy

Let us characterise more precisely the two putative forms of metaphysical indeterminacy associated with quantum mechanics using the machinery set up in the introduction:

`Superposition indeterminacy`: Generic states are in superpositions of eigenstates of various determinables. Therefore it is metaphysically indeterminate what the values of those determinables are.

`Deep indeterminacy`: For some collections of commuting determinables, it is not possible to ascribe determinates to the complete collection consistently. Therefore it is metaphysically indeterminate which properties are instantiated.

In this section, I set up the machinery to make precise `deep indeterminacy`. In particular, I introduce the BW model which models metaphysical indeterminacy as the claim that the actual world is or could be unsettled between multiple possibilities. Now, this intuition needs some unpacking involving, as it does, at least two distinct notions of modality. So I begin, in §4.1 by discussing the relationship between modality and indeterminacy. This allows me to introduce both the Barnes-Williams model of metaphysical indeterminacy, and Skow's notion of 'deep metaphysical indeterminacy'. I then discuss Skow's argument for why the BKS theorem demonstrates the possibility that QM is deeply metaphysically indeterminate in §4.2.

In that section, I argue that Skow is led to his claim about deep metaphysical indeterminacy in QM in virtue of his accepting `Ontic-value`. In §5, I argue that the alternative, `Ontic-algebraic`, is coherent and attractive, and has the additional feature that it renders QM immune to the possibility of deep metaphysical indeterminacy.

4.1 Modality and indeterminacy

4.1.1 The first modal

The invocation of modality in introducing metaphysical indeterminacy typically indicates a hedge: it seems to reek of epistemic bravado to assert the existence of metaphysical indeterminacy. But weaker claims such as 'given a consistent interpretation of our best physics, there could be metaphysically indeterminate facts' are still highly nontrivial. With an eye on QM, one might read Skow's first modal claim nomologically: given the laws in the actual world, there are nomologically possible worlds that are metaphysically indeterminate.

> There are many... interpretations of quantum mechanics...that make no use of the notion of metaphysical indeterminacy. If we reject the orthodox interpretation and accept one of those instead, then we do not have to say that there is *actually* any deep metaphysical indeterminacy. But it would still be true that the metaphysical indeterminacy in the orthodox interpretation of quantum mechanics is a *possible* kind of metaphysical indeterminacy. [30, p. 856]

I read this as asserting that while one has a choice to reject an interpretation of the laws of the actual world on which QM presents possibilities which are metaphysically indeterminate, QM itself is consistent with interpretations on which some possible worlds are metaphysically indeterminate. Thus, it is nomologically possible that QM is metaphysically indeterminate.

4.1.2 The second modal

The BW model treats metaphysical indeterminacy as an irreducible aspect of the world by making precisificational possibility a theoretical primitive. Barnes and Williams thus provide a non-reductive analysis of metaphysical indeterminacy, and locate the indeterminacy in what they call the 'non-representational world'. They work with an ersatzist conception of precisificationally possible worlds (PPWs), where a possible world is 'an abstract object which represents a (classically complete) way the [non-representational] world might be'. According to ersatzist views, there exists a concrete actual world that stands in a determinate relation of actualisation to some ersatz possible world. If the actual world is metaphysically determinate, then this actualisation relation is a two-place relation between the actual world and the ersatz actual world. On the BW model, if the actual world is metaphysically indeterminate, then this actualisation relation is a $n-$place relation, where n might itself not be determinate, between the actual world and all the PPWs that do not determinately fail to represent the actual world.

For an ersatz PPW to determinately fail to represent the actual world, it must be the case that there is a determinate first-order matter of fact about the actual world that is contradicted by a first-order matter of fact about the possible world. So, if, for example, Schrödinger's cat is actually alive, then any possible world in which it is dead determinately fails to represent the actual world. If, on the other hand, there is no matter of fact in the actual world about whether the cat is alive or dead, then a possible world in which the cat is alive does not determinately fail to represent the actual world. Although the possible world does fail to represent the fact that the cat does not have the determinate property of being alive, this is a second-order failure, and thus a wholly separate matter.

Clearly, a great deal hinges on how exactly these ersatz possible worlds are characterised. As with the characterisation of properties in §2, we can consider two distinct questions regarding these PPWs:

`Ontological`$_{\text{PPW}}$: What kind of entity is a PPW?

`Identification`$_{\mathrm{PPW}}$: Given a collection of ersatz possible worlds, which of those worlds is appropriate to stand in an actualisation relation to the actual world?

Once again, our interest in this paper is with the identification question. Therefore, like Barnes and Williams, I will not (and do not have to) answer the ontological question.[18]

Their suggestion for how to answer the identification question is that each PPW represents a *maximally classical* way the world might be. 'Classical' refers to classical logic (not classical (Newtonian) mechanics!): there is no violation of standard axioms of bivalent logic. 'Maximality' is cashed out as the claim that each precisification will be such that, for any proposition p, the precisification will opt for either p or its negation. Two points worth higlighting: (i) the use of the term 'opt for': this is a placeholder for whatever relation propositions stand in with respect to worlds, which, in turn, depends on how the ontological question is answered; (ii) the 'maximally' in 'maximally classical' does not modify 'classical'; the characterisation of an ersatz world is perhaps better rendered as 'maximal *and* bivalent'

What is the domain of propositions over which this notion of maximality quantifies universally? Luckily, given our interests in this paper, we do not need an answer to the most unrestricted reading of the question (which is, in effect, 'what is the set of all propositions?'). It will suffice, for us, to consider all propositions that are about ontic states, and the subset of those propositions that are about kinematic states. Our views about how to characterise metaphysical indeterminacy, therefore, are heavily dependent on our views about ontic and physical properties.

On the BW model, then, the actual world is metaphysically indeterminate if and only if it actualises more than one ersatz PPW, where each such world is classical and determinately opts for p or for not-p so long as p is a proposition about ontic properties. This is the notion of metaphysical indeterminacy that often thought to apply to ordinary QM, and in particular to `Superposition indeterminacy`: each eigenstate is taken to correspond to a set of PPWs: i.e. a maximal and classical way the world could be. But quantum mechanics does not provide determinate truth values to questions about what property is instantiated in the actual world.

4.1.3 Deep metaphysical indeterminacy

Skow claims that it is (nomologically) possible that there is a form of metaphysical indeterminacy that outstrips what can be encoded in the BW model, thus rendering

[18]Their term 'ways the world might be' is consistent with at least four standard ersatzist ontologies: abstract states of affairs, propositions, sentences, and sui generis abstract objects.

the latter inadequate. To quote Skow:

> Maybe metaphysical indeterminacy runs so deep (or can run so deep) that *reality cannot be completely precisified*. I shall call this kind of metaphysical indeterminacy 'deep metaphysical indeterminacy'. [30, p. 852] (emphasis in original)

`Deep indeterminacy` thus includes a claim about consistent ascriptions of determinates. This is what distinguishes it from shallow forms of indeterminacy like `Supersposition indeterminacy`. Consider a quantum state in an $+1$ eigenstate of σ_x. According to the EEL, its value of σ_y is indeterminate. However, there is no prohibition, according to the BW model, on particular precisifications ascribing determinate values to both spin in the $x-$direction and spin in the $y-$direction; if there is metaphysical indeterminacy, then the actual world will correspond to multiple PPWs which consistently ascribe either ± 1 to spin in the $y-$direction. Clearly none of the PPWs misrepresent the value of σ_x; they all agree that it is $+1$. But importantly, none of these PPWs determinately misrepresent the value of σ_y since that property is, by hypothesis, metaphysically indeterminate.

They key here is to remember that any nomologically possible world can correspond to multiple PPWs, and *it is not a requirement that each PPW is nomologically possible*. So, in particular, the EEL does not need to be satisfied at each PPW, rather, it only needs to be satisfied by the collection of PPWs that together correspond to a nomologically possible world.[19] There is, nonetheless, a hidden constraint on each PPW: that its ascriptions be logically consistent, by the standards of the appropriate logic; in the case of the BW model, this logic is standard classical bivalent logic. There is nothing logically inconsistent about ascribing a determinate value to the σ_y operator in this example; to do so is to say that, despite not being in the appropriate eigenstate, the property of σ_y has a determinate value. But that's just another way of saying that the EEL is violated. And that's perfectly acceptable for a PPW.

So, in the case of shallow metaphysical indeterminacy, the appropriate *nomologically* possible world (which now *does* satisfy the EEL) corresponds to two PPWs, both of which agree on the determinate value of σ_x, but disagree on the determinate value of σ_y.

According to Skow's claim about deep indeterminacy, situations like those involving BKS operators are different. Here, any PPW that ascribes determinates to the all

[19]It is this observation that undermines a recent claim, by Fletcher and Taylor [14], that the BKS model is rendered ineffective by the EEL on its own, without recourse to BKS-type considerations. I discuss this in more detail in §5.3.

nine BKS determinables *determinately misrepresents* their values in the nomologically possible world to which it corresponds. This is because as a matter of mathematics, there is no consistent way to ascribe determinates to each of the BKS determinables such that the appropriate algebraic relations are satisfied. Thus, any putative PPW which does ascribe determinates to all nine BKS determinables violates the condition of mathematical consistency, *a fortiori* logical consistency. This is the source of the deep metaphysical indeterminacy that Skow claims is conceptually distinct from shallow indeterminacy like `Superposition indeterminacy`.[20]

Skow takes issue with the 'maximality' component of the BW model. He spells out two ways in which reality might be deeply metaphysically indeterminate. On one hand, it might be the case that a particular cannot determinately instantiate one out of a class of possible ontic properties on its own; on the other it might be the case that for a given collection of particulars, even though each can individually instantiate one out of a collection of ontic properties, the collection is such that no combination of those ontic properties can be simultaneously instantiated. Skow claims that the BKS theorem demonstrates that QM might engender the second kind of deep metaphysical indeterminacy.

4.2 Skow on deep metaphysical indeterminacy in quantum mechanics

Let me characterise three claims that we might want to posit about QM:

`Commuting`: If two determinables commute, then the order of their measurement does not affect the probabilities with which eigenvalues are measured (i.e. determinates).

`Noncommuting`: If two determinables do not commute, then the order of their measurement does affect the probabilities with which eigenvalues are measured.

`Faithful`: If a system is subjected to a measurement procedure to specify the determinate associated with some determinable, then that determinate was a specification of a determinable that was already instantiated by the particular before measurement.

Skow invokes the BKS theorem, and its application to the type of spin system discussed in §3.1, to argue that QM might be metaphysically indeterminate. He begins

[20] Skow is not the only person to have suggested the BW model is inadequate to model deep indeterminacy; others include [6, 8, 9, 14]).

by restricting to 'orthodox interpretations of QM' which, paraphrasing Skow, we can characterise as:

`Orthodox`: If the complete physical state of a system does not assign probability 1 to one out of a set of candidate determinable values, then, before we look, it is indeterminate what determinable value that particular instantiates.

`Orthodox` is thus the 'only if' component of the EEL. For the spin system, `Orthodox` says that spin properties are indeterminate when the probability of measuring them to have a particular value is not 1.

Recall, we characterised a spin measurement in §3.1 in terms of eigenvalues and eigenstates: if (and only if) a spin system is in an eigenstate of σ_z, say $\begin{pmatrix}1\\0\end{pmatrix}$ then it will, with probability 1, be measured to have a value of $+1$. If not, then, by `Orthodox`, it will not have a determinate. In particular, when the system is in an eigenstate of σ_x, it is determinately not in an eigenstate of σ_z, so according to `Orthodox`, the system has an indeterminate value of spin in the z direction. In other words, according to `Orthodox`, `Noncommuting`$\rightarrow \neg$`Faithful`

This fact on its own poses no problem for the BW model. Consider, once again, the example from §4.1.3: a spinor in an eigenstate of σ_z. Insofar as it is in a superposition of eigenstates (i.e. perfectly determinate states) of σ_z, two (sets of) ersatz worlds are actualised by the actual world: one in which σ_x and σ_z eigenvalues are $+1$, another in which σ_x has eigenvalue $+1$ and σ_z has eigenvalue -1.

So the BW model works perfectly well for representing the putative metaphysical indeterminacy of the orthodox interpretation of noncommuting determinables. But now recall that we can construct, out of these noncommuting determinables, a set of pairs of mutually commuting ones: the nine operators in §3.2. According to `Commuting`, these determinables can be measured in any order to yield the same set of determinates. And from `Faithful`, we can infer that the particulars already instantiate these properties before they are measured.

But, as we saw in §3.2, the BKS theorem imposes an obstruction on `Commuting`: there is no way to ascribe the appropriate determinates to all nine operators such that the repeated measurements are guaranteed to generate the same outcomes. Indeed, this violation of `Commuting` is a consequence of the violation of `Faithful`: the particulars cannot instantiate all nine of these properties simultaneously.[21] And this violation

[21] It is worth noting that violations of `Faithful` are sufficient, though not necessary for violations of `Commuting`.

is down to the fact that there is no mathematically consistent PPW where all nine determinables have determinate values.

Skow takes this to establish the possibility of deep metaphysical indeterminacy. He argues that since there are no PPWs where all nine determinables have determinate values, there are no appropriate PPWs: every PPW is non-maximal, and this conflicts with BW's stipulation. Thus, the BKS obstruction prevents the construction of the very worlds that the BW model quantifies over. Skow concludes that the BW model is thus inadequate to model the sort of indeterminacy that the BKS theorem implies. But this conclusion only follows if something like `ontic-value` is true.

5 Quantum mechanics is not deeply metaphysically indeterminate

In this section, I propose that Skow's conclusion, that QM is possibly deeply metaphysically indeterminate, can be avoided by denying `ontic-value` and instead adopting algebraic naturalism. In particular, I argue that, as with '...is deterministic', '...is metaphysically indeterminate' is not monadic, but rather n-place with respect to physical determinables. Algebraic naturalism suggests which physical determinables we should choose: those definable in the models of the theory. I illustrate this procedure explicitly using the example of intrinsic angular momentum (i.e. spin) first in Newtonian mechanics (§5.1) and then in quantum mechanics (§5.2). I then address the question of how we should interpret quantum mechanical models, at least with respect to property ascriptions, in §5.3.

The BKS theorem is a general claim about the non-embeddability of certain non-commuting sub-algebras into commuting algebras. For pedagogical clarity, I illustrated the BKS theorem in §3.2 using Mermin's intuitive example of quantum mechanical spin. As a result of that choice, we now find ourselves in a situation where we have to make a small compromise: we will have to invent a Newtonian quantity that is the analogue of quantum mechanical spin. For our purposes, it is enough to just define a quantity that has the same units as angular momentum, and christen that quantity 'Newtonian spin in the $x-$direction'. So let us choose the simplest function that delivers the appropriate units: $\sigma_x := q_x p_x$, where q_x is the $x-$coordinate associated with some particle and p_x its linear momentum in the $x-$direction. We can define Newtonian spin similarly for all other directions in space.

5.1 Newtonian 'spin' as a property of maximal and bivalent Newtonian worlds

The determinable σ_i is expressed in terms of q_i and p_i; the algebraic state is then a straightforward value function:[22] abstractly, it is a function from determinables to determinates, i.e. a pair $\langle \sigma_i, v_i \rangle$, where $v_i \in \mathbb{D}_N$. The value function is definable because it is invariant under the automorphisms of the Poisson algebra of determinables. Consider, for example, the paired automorphism $\langle s, t \rangle$: $s(q_i) = p_i$, $s(p_i) = -q_i$ and v_i is mapped to $t(v_i)$. Since each determinable σ_i is individually preserved under the paired automorphism, each value function is definable. As a result, it is easy to see that the Poisson algebra of determinables is invariant under the paired automorphism: we can think of the s map as acting directly on the algebra of determinables, including each σ_i, mapping each element to some element of the algebra in such a way that its algebraic structure is preserved.

Consider a putative PPW representing all of these properties of a Newtonian particle. In particular, consider the (Newtonian analogues of) the nine BKS spin determinables we met in §3.2. The algebraic state is just a value function defined on the set; since each value function is individually invariant under the paired automorphism, each determinable is mapped to a determinate. So each PPW corresponds to a mathematically consistent assignment of a determinate for each and every BKS determinable. It therefore meets Skow's criterion of being maximal and bivalent: every proposition purporting to report the the value of a determinable has a determinate truth value.

5.2 Quantum spin as a property of maximal and bivalent quantum worlds

When we try to play the analogous game for quantum spin defined on a quantum state, a number of subtle changes need to be made. As mentioned earlier, certain mathematical structures are different: Poisson algebras are replaced by von Neumann algebras, symplectic manifolds by Hilbert spaces and so on. Crucially, we are in a context in which value functions and expectation value functions are distinct, so now a choice needs to be made regarding our definition of ontic states, and their relationship to algebraic states. This is where Skow and I disagree.

Skow proposes that the ontic state continues to be a value function (i.e. he insists

[22] Which, in this context is also trivially an expectation value function.

on the truth of `ontic-value`). This immediately leads to a problem: as noted in §3.1, a (determinate) value function exists only on the subset (note, this is not a subalgebra) of von Neumann algebra operators for which the EEL is satisfied by the ontic state. The ontic state does not map any other operators to determinates. For the nine BKS operators, for example, there is no joint eigenstate, so no value function can be defined. The indeterminacy arises because, perhaps surprisingly, even for a subset of pairwise commuting operators, if the EEL (*a fortiori* `Orthodox`) is true, then for at least one determinable, there is no determinate. Thus the only way to hold onto `ontic-value` is to insist that the ontic state is itself metaphysically, indeed *deeply* metaphysically, indeterminate.

My proposal is that we instead infer the ontic state directly from the algebraic state (which, in this context is an *expectation value function*), without having to attempt a detour through value functions. After all, nothing in the structure of a physical theory requires that ontic states have anything to do with value functions. The value functions are not invariant under paired automorphisms, *but the expectation value functions (i.e. the algebraic states) are.* So if we accept `ontic-algebraic`, then we end up with an ontic state that is perfectly determinate with respect to the appropriate determinables.

Of course, now the question arises as to what 'appropriate determinates' are. Algebraic naturalism means that what we previously thought of as uncontroversial determinables like 'spin in the $x-$direction' are no longer appropriate determinables; the system simply lacks that determinable (*a fortiori* a determinate). This suggestion resonates very closely with Glick's 'sparse view' [16], at least as it applies to deep metaphysical indeterminacy:

> `Sparse view`: when the quantum state of A is not in an eigenstate of $\hat{O}$, it lacks both the determinate and determinable properties associated with $\hat{O}$.

Indeed, the sparse view follows from algebraic naturalism and, in particular, the `definability thesis`. Recall according to the thesis, if one accepts a physical theory, then one ought to be committed only to properties and relations that are definable in that theory's models.

The crucial insight that allows us to deny `ontic-value` and thus resist Skow's conclusion is that physical theories come with their own standards of definability, as encoded in the algebra of determinables and the value space. This standard of definability entails a standard of metaphysical determinacy via its characterisation of maximality for ersatz possible worlds. Recall, from §2.3.1, the example of emerald green and seafoam green as being definable in one model, $\mathfrak{S}_2$, but not $\mathfrak{S}_1$, even though both models were defined over the same domain. Each model came with its own standard of definability, and hence, via the definability-determinacy proposal, its own standard of metaphysical determinacy. Problems only arose when we made cross-model assessments of standards of determinacy.

So quantum mechanics is (possibly) deeply metaphysically indeterminate by the standards of the determinables appropriate to Newtonian mechanics, but there is no reason to hold quantum mechanics to that standard in the first place! We need to be careful to disentangle essential features of the algebra of determinables from contingent coincidences about how that algebra might be realised in particular contexts.

In the Newtonian context, the abstract algebra of determinables was realised as a Poisson algebra of smooth functions on a manifold. That was lucky! As a result of this piece of good fortune, we could interpret each quantity as being, in some sense, constituted by the basic quantities of position and momentum, which we take ourselves to have a relatively good intuitive grip on. But recall that, according to algebraic naturalism, our interpretative standards regarding determinacy should be guided by the abstract algebra, not any particular concrete realisation, and its associated intuitions. So when, in the quantum context, we discover that the algebra of determinables is realised by a noncommuting von Neumann algebra of self-adjoint operators on a Hilbert space, we shouldn't hold our interpretative practices to standards left over from when we thought that the algebra of determinables was a Poisson algebra of smooth functions. `Ontic-value` is one such leftover; it should be discarded. In its place, we should use the definability thesis to determine our standards of determinacy.

Given the distinction between the two types of putative metaphysical indeterminacy (i.e. `Superposition indeterminacy` and `Deep indeterminacy`), one might wonder whether the definability-determinacy proposal is up to the challenge of dealing with deep indeterminacy. It is. Recall 'maximality' was cashed out as the claim that each precisification will be such that, for any proposition p, the precisification will opt for either p or its negation. But not all determinables that are definable in the Newtonian theory are definable in the quantum theory (and, indeed, vice versa). So a maximal precisificationally possible quantum ersatz world *need not* pick out propositions about any of the nine quantum BKS spin determinables, or any other propositions about undefinable properties, even though a maximal precisificationally possible Newtonian ersatz world might have done so, for the Newtonian equivalent of the BKS spin determinables. As I mentioned in the introduction, Newtonian mechanics is metaphysically indeterminate by the standards of Aristotelian teleological cosmology (i.e. it is possible that no Newtonian ersatz possible world with a determinate privileged 'centre of the universe' determinately misrepresents a world governed by Newtonian mechanics), but I believe it is difficult to justify an insistence on imposing Aristotelian metaphysical standards on Newtonian mechanics; Skow's proposal is the exact analogue of this.

Now, with respect to those Newtonian ersatz worlds, the BW model does, indeed, fail, since each of those worlds determinately misrepresents the actual quantum mechanical world. But this failure is completely harmless: algebraic naturalists should not care about judgements about quantum semantics or metaphysics made with respect to Newtonian standards. And with respect to standards appropriate to quantum mechanics, the BW model does just fine.

Compare the PPWs understood in this way to the PPWs for the Newtonian spin system we discussed in §5.1. The algebraic quantum state is now an *expectation* value function defined on the set; each expectation value function is individually invariant under the paired automorphism, so each determinable is mapped to a well defined expectation value. So each PPW corresponds contains an assignment of a determinate for each determinable. If a putative determinable does not have a determinate, it fails to be a determinable: this is just Glick's sparse view. So, exactly as in the Newtonian case, these PPWs meet Skow's criterion of being maximal and bivalent: every proposition purporting to report the the value of a determinable has a determinate truth value.

There is an important methodological difference between Skow's proposal and mine. For Skow, the set of properties of worldly objects is assumed, or somehow pretheoretically given. It is then the job of physics to tell us how nature has ascribe values to these properties. And the BKS theorem shows us that QM does not ascribe determinate values to all assumed properties. Hence deep metaphysical indeterminacy. My proposal is that we do not start with a pre-theoretic list of properties and then ask physics to fill in their values. Rather, begin with a theory of physics, and use its models to determine whether certain questions are intelligible by determining whether the properties they ascribe to worldly entities are definable in those models. As a consequence, it is trivial that worlds come out as metaphysically determinate, but it's a triviality that has been earned by the work put in to establishing the coherence of algebraic naturalism.

Of course, Skow and other weak naturalists could always respond by denying that `ontic-algebraic` is a better norm than `ontic-value`. After all, Skow does not argue that QM is definitely deeply metaphysically indeterminate, only that it is nomologically possible that it is. In order to definitively rule out this possibility, I would need one final argument: that `ontic-algebraic` trumps `ontic-value`. It is not clear to me that an argument of that sort determinately exists; I certainly cannot provide one here. But I hope that, in showing that the conclusion of Skow's argument is not 'it is nomologically possible that quantum mechanics is deeply metaphysically indeterminate' but '*if* `ontic-value` is true, *then* it is nomologically possible that quantum mechanics is deeply metaphysically indeterminate', I have highlighted an important choice point in the metaphysical cost-benefit analysis regarding properties. In particular, I hope to have made the conjunction of `ontic-algebraic` and 'if `ontic-algebraic` is true then it is not nomologically possible that quantum mechanics is deeply metaphysically indeterminate' a position worth defending.

Ultimately, I think that the burden of proof with respect to deep metaphysical indeterminacy should fall on its proponents, not its opponents. And, to be fair, this is precisely how the debate has played out: the default view is that there is no deep MI and Skow's gambit is to use the BKS theorem to argue otherwise. It is an interesting and thought-provoking (indeed, paper-provoking) move. My view demonstrates that adopting what I take to be a plausible position, namely algebraic naturalism (or strong naturalism more generally), renders Skow's move dialectically inert. Algebraic naturalism entails that deep MI is *conceptually* impossible, not merely false. Of course, this leaves me open to a Moorean shift: if I'm presented with a suitably compelling argument for deep MI, then so much the worse for my algebraic naturalism. What I hope to have demonstrated in this paper is that the BKS theorem is not strong enough to mandate a Moorean shift against the algebraic naturalist. But I am open to the possibility that other considerations might be.[23]

5.3 The metaphysics of quantum properties

How should we think about property ascriptions in quantum mechanics, if we accept alegbraic naturalism? The opponent of algebraic naturalism might argue that there is an important disanalogy between the Newtonian 'centre of the universe' and quantum 'spin-up in the $x-$direction': No Newtonian system ever instantiates 'being at the center of the universe.' So it is appropriate to say that this is not an ontic property at all. But this is not true of putative properties like 'spin-up in the $x-$direction'. Those are properties that a system can determinately instantiate and that many systems in

[23]I also hasten to add that algebraic naturalism does not entail that MI is conceptually impossible. Just that deep MI is. The view that there might be shallow forms of MI is perfectly consistent with algebraic naturalism.

fact do.[24] So if algebraic naturalism entails that this is not, in fact an ontic property, then how should we think about quantum properties?

One way to answer this question is to invoke and defend a particular interpretation of QM. But I believe there is an interpretation-agnostic position one can embrace as well:

`Holism`: Since ontic states derived from algebraic states will eventually be such that some intuitively uncontroversial determinables of subsystems (e.g. such monadic determinables as 'spin in the $x-$direction') will not be ascribed determinates, the only option for a choice of quantum particulars is an entire world that is not factorisable into subsystems that are themselves particulars to which we can ascribe intuitively uncontroversial determinables.

The algebraic naturalist can easily accept `Holism`. Indeed, this is pretty much par for the course in QM. We don't even need to invoke the BKS theorem to argue that QM states are irreducibly global. Just consider a rough characterisation of a quantum particle like a fermion (for example, an electron): it is a 'constituent' of an n- particle state such that the pairwise interchange of any two such constituents of the state results in an equivalent state. Thus, to be sure that we can treat a particle as a fermion, we need to know the behaviour under pairwise permutation of every other particle in the universe.[25]

We can, of course, characterise subsystems of global quantum systems approximately, although how to make sense of this notion of approximation is tricky. For example, for some practical (and some impractical) purposes, we can speak of individual fermions as particulars. Indeed, if we were only interested in, say the first row of BKS determinables, then the system would look like it factorises into what we could legitimately call 'two independent 2-spinors'. As long as we as keep track of the appropriate contexts, it is perfectly harmless, then, to speak of independent 2-spinors, and of properties of two distinct particulars. What the BKS theorem adds to the mix is the fact that if we expand our interest to enough other quantities, even this characterisation is no longer available. This is what makes the property ascription 'contextual'.

[24] I am grateful to an anonymous reviewer for this raising this worry.

[25] This is a very rough characterisation that relies on us being able to talk of individual fermions in the first place. This is not necessary for their definition. Really, there is no such thing as a state of n individual fermions in 3-dimensions, but rather a $3n$-dimensional fermionic state. More careful quantum field theoretic treatments of the question of defining bosons, fermions and paraparticles use more sophisticated machinery, but are no less global in ascriptions of labels like 'fermionic' to a state. For a discussion of this subtle problem, see [1].

The contextuality of quantum mechanics means that we have to give up talk of 'spin-up in the $x-$direction' as describing fundamental, context-independent ontic properties of the world according to QM. But we might, nonetheless often find ourselves in situations where our interest determines a context, and in such situations, continue to make sense of such property ascriptions. Lewis [22] has recently suggested that although the problem of property ascriptions in quantum mechanics might bear some resemblance to issues of metaphysical indeterminacy, they are perhaps best understood as a *sui generis* quantum concept that does not map neatly onto either our everyday or our pre-quantum theoretic descriptions of the world. Perhaps, then, the best we can do is endeavour to keep track of the contexts in which this mismatch is not detrimental to inferences we would like to make.

What I propose here is, in a sense, revisionary. But not radically so; as the fermion example shows, there are precedents. Nonetheless, as a revisionist I should provide an error theory. about how we come to believe that such properties exist. Luckily for us, we can simply piggyback on a standard sort of quantum-theoretic error theory, at least for unmodified interpretations of QM.

Consider the case of two isolated systems A and B that we allow, through dynamical evolution to get entangled. One might have the intuition that, even after this entanglement, the two systems have individual states with some monadic properties; this is what is captured by taking the partial trace of one of the systems, and thus defining a mixed state. It is natural to interpret this mixed state (of say, system A) as attributing properties to A, but it is important to remember that at best this gives us information about A that can be detected without interaction with or information about B. Does this mean that A does not have the property of, say, spin-in-the-x-direction? Yes, if that property is understood as a monadic property of system A. But for some purposes, there is a property on the joint system, represented by a tensor product of some operator, call it S_x, on (the Hilbert space of) system A together with the identity on (the Hilbert space of) B, that behaves as if it ascribes the property of spin-up-in-the x-direction only to system A. And as long as we keep track of the situations in which the inferences we draw do are blind to whether we think of the property as monadic or dyadic, we can continue to speak (loosely) of the property as monadic; this, I claim is what accounts for our use of 'spin-in-the-x-direction' as a monadic predicate.

Suppose, for example, we are interested in the expectation value of spin-in-the-x-direction on A. When we take take the partial trace of A and calculate the expectation value S_x, which is defined on A, it looks as if we're performing operations only on (the state of) A. What we are, in fact doing, is calculating a trace over the product of two operators, one of which is S_x, but the other of which is the partial trace over system B. Clearly, then, the expectation value of S_x does take into account information about the entangled system B; it does not correspond to a monadic property of system A. However, since the precise details of the state of the system B does not impact the expectation value of S_x, those details about system B are irrelevant to inferences drawn from the expectation value of S_x. And for the details of some system to be irrelevant to a particular question is just for that system to be treated as independent; so S_x appears to be a monadic property of system A. Our error theory can then be summarised as: we thought S_x was a property of system A, even when entangled with system B, because the inferences we drew on the basis of the expectation value of S_x, were agnostic with respect to information about system B. The property of the joint system appeared to factorise into properties of individual systems. The lesson to draw from this, from the perspective of this paper, is that we do not require an account of properties that vindicates a belief that 'spin-in-the-x-direction' is fundamental. All we need to vindicate is the appropriateness of our talk of 'spin-in-the-x-direction' as useful in certain contexts, and our standard error theory does that.

Finally, let me assuage the worry, articulated by Fletcher and Taylor, that the BW account is rendered inadequate at a much earlier point in the dialectic than the BKS theorem.[26] Here is their central cricitism of the BW account ([14, p. 11202-11203]):

> The problem with applying BW's theory to the quantum case is that it cannot jointly model four propositions that it ought to be able to:
>
> (a) EEL holds determinately.
>
> (b) Mathematics... holds determinately.
>
> (c) It's indeterminate whether [the electron] e is [spin-up-in-the-x-direction].
>
> (d) It's determinate that [the electron] e is [spin-up-in-the-y-direction].

So, according to Fletcher and Taylor, if (a)-(d) ought to be modelled by an account of indeterminacy, then the BW account is inadequate. And all of this is established without having to mention the BKS theorem. Their conditional claim is true. But the claim only becomes interesting as a criticism of the BW account if we accept that it ought to be able to model (a)-(d). Fletcher and Taylor use the conditional claim to generate a *modus ponens* against the BW account, but it can just as easily be Moore-shifted to generate a *modus tollens* against the EEL.

[26] Once again, I'm grateful to one of the anonymous reviewers for raising this issue.

The Moorean shift, on its own, is not decisive. To tip the balance in favour of the BW account, we need to do one (or both) of two things: either argue that we ought to have high confidence that the BW account is correct, or argue against our confidence in the EEL. In the last part of this section, following Wallace [33], I do the latter.

Wallace has, over the years, argued against the EEL in many different ways. Here, I invoke his argument from Hegerfeldt's theorem [17, 18], which I quote at length [33, pp. 299-300]:

> `Hegerfeldt's theorem (indefiniteness form)`: Given a system evolving unitarily over some interval of time under a Hamiltonian whose spectrum is bounded below, a given property is either (a) definitely not possessed at every time in that interval, or (b) not definitely not possessed at almost every time in that interval.
>
> Put another way: suppose there is some property that, at some time in the indefinitely distant future, the system might have some probability to be found to possess. Then, according to the [EEL], it is immediately — that is, withinan arbitrarily short window of time — indefinite whether the system has that property.
>
> Put yet another way: anything that might at some future point be indefinite will be indefinite immediately. This seems to render the [EEL] fairly useless as a description of ontology. We might have imagined that systems begin having some definite value of a given quantity, then gradually evolve so as to be indefinite across several values of that quantity, and in due course become completely indefinite with respect to that quantity (perhaps until some wavefunction collapse restores definiteness). But dynamically, that can't happen: indefiniteness is immediate if it is going to happen at all.
>
> [...]
>
> The underlying problem here is a radical mismatch between the [EEL] and the way quantum mechanics actually handles the idea of a system's becoming more spread out (speaking loosely) with respect to a given quantity. QM handles the latter through probabilities: the likelihood of a particle localised at x being found very far from x is initially negligibly small and only gradually increases — and, depending on the dynamics, may never increase beyond negligible levels. But the [EEL] is all-or-nothing: as soon as the system has any probability, even [1 in $10^{10^{20}}$] , of being found in some region, it is completely indefinite whether it is in that region.

Fletcher and Taylor do not acknowledge this argument or cite the paper in which it appears. They do very briefly engage with some of Wallace's other arguments, but their relevance is dismissed in a footnote,[27] on the grounds that 'it is not our aim to argue for [the EEL] here. Instead our goal is to see what sort of theory of quantum indeterminacy arises by taking the EEL as our starting point. And this seems justified: EEL is a central part of orthodox quantum theory, and it's common ground among most recent accounts of quantum indeterminacy' [14, p. 11182].

This move is insufficient: they dismiss Wallace's argument as irrelevant because they claim that are only interested in justifying the conditional claim, but not its antecedent. But in the very next sentence, they argue that accepting the EEL is justified because of its centrality to many accounts of quantum indeterminacy. At best, this sociological observation might be used as motivation to *look for* an argument against Wallace's conceptual claim; it does not itself *constitute* that argument. Wallace's argument invokes nothing more than a mathematical fact about QM; Fletcher and Taylor, if anything, invoke an argument from consensus. And this is enough to ground the required Moorean-shift against Fletcher and Taylor. The BW account survives.

6 Conclusion

Strong naturalism, as a broad norm on metaphysical theorising, is a position according to which (i) physical theories are equipped with standards of determinacy for physical properties and (ii) judgements of determinacy should only be made with respect to those standards. The aim of this paper was to argue that strong naturalism, and in particular its algebraic realisation, is coherent and plausible in both Newtonian and quantum mechanics. I went about this by first demonstrating that, under an algebraic reformulation of Newtonian mechanics, we could identify plausible and intuitive standards of determinacy, by applying the definability thesis to the Poisson algebra of Newtonian determinables. I invoked the definability-determinacy proposal to argue that the undefinability of a property according to a model that represents a possible world should be understood as representing genuine metaphysical indeterminacy of that property in that world.

I then applied the same procedure to the von Neumann algebra of quantum deter-

[27] They also point to Gilton [15] as a rebuttal to Wallace. But Gilton only rebuts Wallace's historical claim that the EEL has played no role in QM, not in the conceptual claim that the EEL is not part of the conceptual structure of QM. Further, Gilton's paper does not comment on Wallace's paper which sets out the argument from Hegerfeldt's theorem, having been published before the latter paper.

minables, to arrive at the conclusion that certain properties, definable and therefore determinate by quantum standards, were undefinable therefore indeterminate by Newtonian standards. In particular, I argued that the inconsistency of determinate property ascriptions to systems described by the Bell-Kochen-Specker theorem should be understood not to entail any particularly deep metaphysical indeterminacy. According to strong naturalism, Newtonian standards are inappropriate for assessments of quantum determinacy. And with respect to the appropriate standards, quantum mechanics is not deeply metaphysically indeterminate.

References

[1] David John Baker. Identity, superselection theory, and the statistical properties of quantum fields. *Philosophy of Science*, 80(2):262–285, 2013.

[2] Elizabeth Barnes and J Robert G Williams. A theory of metaphysical indeterminacy. In D. Zimmerman and K. Bennett, editors, *Oxford studies in metaphysics*. Oxford University Press, Oxford, 2011.

[3] John S Bell. On the problem of hidden variables in quantum mechanics. *Reviews of Modern Physics*, 38(3):447, 1966.

[4] Alisa Bokulich. Metaphysical indeterminacy, properties, and quantum theory. *Res Philosophica*, 91(3):449–475, 2014.

[5] P Busch, PJ Lahti, and P Mittelstädt. The quantum measurement theory (lecture notes in physics), 1991.

[6] Claudio Calosi and Cristian Mariani. Quantum relational indeterminacy. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 71:158–169, 2020.

[7] Claudio Calosi and Cristian Mariani. Quantum indeterminacy. *Philosophy compass*, 16(4):e12731, 2021.

[8] Claudio Calosi and Jessica Wilson. Quantum metaphysical indeterminacy. *Philosophical Studies*, 176(10):2599–2627, 2019.

[9] George Darby. Quantum mechanics and metaphysical indeterminacy. *Australasian Journal of Philosophy*, 88(2):227–245, 2010.

[10] George Darby and Martin Pickup. Modelling deep indeterminacy. *Synthese*, pages 1–26, 2019.

[11] Michael Dummett. Wang's paradox. *Synthese*, 30(3):301–324, 1975.

[12] Herbert Enderton. *A mathematical introduction to logic*. Elsevier, 2001.

[13] Gareth Evans. Can there be vague objects. *Analysis*, 38(4):208, 1978.

[14] Samuel C Fletcher and David E Taylor. Quantum indeterminacy and the eigenstate-eigenvalue link. *Synthese*, 199(3):11181–11212, 2021.

[15] Marian JR Gilton. Whence the eigenstate–eigenvalue link? *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 55:92–100, 2016.

[16] David Glick. Against quantum indeterminacy. *Thought: A Journal of Philosophy*, 6(3):204–213, 2017.

[17] Gerhard C Hegerfeldt. Causality, particle localization and positivity of the energy. In *Irreversibility and Causality Semigroups and Rigged Hilbert Spaces: A Selection of Articles Presented at the 21st International Colloquium on Group Theoretical Methods in Physics (ICGTMP) at Goslar, Germany, July 16–21, 1996*, pages 238–245. Springer, 1998.

[18] Gerhard C Hegerfeldt. Instantaneous spreading and einstein causality in quantum theory. *Annalen der Physik*, 510(7-8):716–725, 1998.

[19] Nick Huggett, Fedele Lizzi, and Tushar Menon. Missing the point in noncommutative geometry. *Synthese*, 199(1):4695–4728, 2021.

[20] Simon Kochen and Ernst P Specker. The problem of hidden variables in quantum mechanics. In *The logico-algebraic approach to quantum mechanics*, pages 293–328. Springer, 1975.

[21] David Lewis. Many, but almost one. *Ontology, causality and mind*, pages 23–42, 1993.

[22] Peter J Lewis. Explicating quantum indeterminacy. In *Quantum Mechanics and Fundamentality*, pages 351–363. Springer, 2022.

[23] N David Mermin. Hidden variables and the two theorems of john bell. *Reviews of Modern Physics*, 65(3):803, 1993.

[24] Alex Oliver. The metaphysics of properties. *Mind*, 105(417):1–80, 1996.

[25] Huw Price. *Expressivism, pragmatism and representationalism.* Cambridge University Press, 2013.

[26] Eduard Prugovecki. *Quantum geometry: a framework for quantum general relativity*, volume 48. Springer Science & Business Media, 2013.

[27] Bertrand Russell. Vagueness. *The Australasian Journal of Psychology and Philosophy*, 1(2):84–92, 1923.

[28] Nathan Salmon. Reference and essence. 1981.

[29] David Schroeren. Quantum metaphysical indeterminacy and the ontological foundations of orthodoxy. *Studies in History and Philosophy of Science Part A*, 90:235–246, 2021.

[30] Bradford Skow. Deep metaphysical indeterminacy. *The Philosophical Quarterly*, 60(241):851–858, 2010.

[31] Alessandro Torza. Quantum metaphysical indeterminacy and worldly incompleteness. *Synthese*, pages 1–14, 2017.

[32] Alessandro Torza. Quantum metametaphysics. *Synthese*, 199(3):9809–9833, 2021.

[33] David Wallace. What is orthodox quantum mechanics?

[34] Jessica Wilson. Determinables and determinates. *Stanford Encyclopedia of Philosophy*, 2017.

[35] Johanna Wolff. Spin as a determinable. *Topoi*, 34(2):379–386, 2015.