La Bohume

May 30, 2017

1 Introduction

David Lewis, famously, suggested a certain kind of picture of what the world is like. He called that picture *Humean supervenience*, and described it as follows:

> Humean Supervenience [...] says that in a world like ours, the fundamental relations are exactly the spatiotemporal relations: distance relations, both spacelike and timelike, and perhaps also occupancy relations between point-sized things and spacetime points. And it says that in a world like ours, the fundamental properties are local qualities: perfectly natural intrinsic properties of points, or of point-sized occupants of points. Therefore it says that all else supervenes on the spatiotemporal arrangement of local qualities throughout all of history, past and present and future.[1]

However, there is a concern that Humean Supervenience is inconsistent with our best physical theories.[2] More specifically, there is a concern that the kind of world described by Lewis above—one which is fully and exhaustively characterised by the assignment of intrinsic qualities to points of spacetime—could not be a world described by quantum mechanics.[3] More specifically still, the concern is that the characteristic quantum-mechanical phenomenon of *entanglement* rules out the possibility of giving

[1](Lewis, 1994, p. 474)

[2]Teller (1986), Maudlin (2007)

[3]To be strictly accurate, there are good reasons for thinking that Humean supervenience, at least on the letter of the above, is inconsistent with *classical* physics. However, one can get around this by (roughly speaking) taking "local qualities" to be intrinsic properties of infinitesimally small spacetime regions, rather than spacetime points *per se* (see Butterfield (2006)). I will ignore this subtlety for the purposes of this essay.

an exhaustive description of the world by describing it point-by-point. So (according to these arguments), insofar as we take quantum mechanics to be *true* (i.e., insofar as we take the actual world to be accurately described by quantum mechanics), we should not take Humean Supervenience to be true either.

More recently, however, there has been a fightback on behalf of Humean Supervenience: it has been argued that, at least if one is a Bohmian about quantum mechanics,[4] then Humean Supervenience remains a consistent option after all. This paper seeks to resist this most recent defence of Humean Supervenience. First, I introduce the relevant pieces of Bohmian mechanics, and indicate the *prima facie* tension between entanglement and Humean Supervenience. Second, I discuss the argument that Bohmian Humeans (from here on out, "Bohumeans") make to render their ontology compatible with Humean Supervenience. I then raise three problems for this argument: a problem concerning the status of the newly-introduced ontology; a problem concerning determinacy of quantities; and a problem concerning scientific practice.

Before I start, I want to clarify that this paper is *not* about whether some suitably modified version of Humean Supervenience is compatible with quantum mechanics. For instance, Loewer and Albert have observed that quantum mechanics, standardly formulated, is straightforwardly compatible with the requirement that qualities be local in *configuration* space, rather than physical space;[5] whilst Darby has argued that we can preserve the "spirit" of Lewis' proposal by allowing that there are fundamental relations *besides* the spatiotemporal relations.[6] All three note that doing so is consistent with Lewis' broader Humean goal of recovering all else (all mental and nomological facts, in particular) from the categorical world, i.e., from a particular distribution of non-modal properties and relations. This is all well and good, but not my concern here. I am exclusively attending to the question of whether the specific variety of Humean Supervenience defended by Lewis (that requiring the locality of *all* fundamental properties in *physical* space) can be rendered consistent with quantum mechanics.

[4]Esfeld (2014) claims that the proposed rescue of Humean Supervenience is available to any "primitive-ontology" approach to quantum mechanics, not just Bohmian mechanics. It's not my intention to examine this claim in this essay: in the interests of brevity, I will focus on the specific case of Bohmianism (though see fn. 12).

[5]See e.g. Loewer (1996), Albert (1996).

[6]Darby (2012)

2 Entanglement in Bohmian mechanics

The fundamental entities of Bohmian mechanics are the *particles*: pointlike objects, which have definite positions at all times, and which are held to be the fundamental constituents of macroscropic matter.[7] Thus, if we have N particles, then their collective state at any given time may be represented by an N-tuple $(Q_1, \ldots, Q_N)$ of points of X, where X is the space representing physical space—that is, by a single point $\mathbf{Q}$ in the N-fold configuration space X^N (the N-fold direct product of X with itself). The behaviour of the particles is determined by the *wavefunction*, a function $\Psi : X^N \to \mathbb{C}$, via the *guidance equation*

$$\frac{dQ_i}{dt} = \frac{\hbar}{m_i} \mathrm{Im} \left(\frac{\nabla_i \Psi}{\Psi}(\mathbf{Q}) \right) \tag{1}$$

where m_i is the mass of the ith particle, and ∇_i denotes the gradient associated with the ith productand of X^N. (This is all in the absence of spin: for the purposes of this essay, we need only consider spinless particles.) The wavefunction itself evolves according to the usual Schrödinger equation,

$$i\hbar \frac{d\Psi}{dt} = H\Psi \tag{2}$$

where H is the Hamiltonian.

The challenge for the aspiring Bohumean may now be stated quite succintly: the wavefunction cannot be any part of a Humean Supervenience basis, and hence cannot (for one attracted by Lewis' picture) be interpreted as a fundamental physical component of the world. For, the wavefunction assigns values (complex numbers) to N-tuples of points of space, not to individual points of space. But the Humean Supervenience basis was required to include only local qualities, i.e., those comprising the assignment of intrinsic properties to individual spacetime points (or to point-sized occupants of spacetime points). So the wavefunction is not the kind of local property with which Lewis would be happy, unless there is some way of showing that any given wavefunction can be reduced to (uniquely specified by) some collection of suitably local qualities.

Certainly, there are specific circumstances in which such a reduction is possible: those in which the wavefunction is not *entangled*. For simplicity, let $N = 2$; now sup-

[7]My presentation of Bohmian mechanics here follows that of Dürr et al. (1992) and Dürr and Teufel (2009).

pose that the wavefunction $\Psi(x_1, x_2)$ is a *product* wavefunction,

$$\Psi(x_1, x_2) = \psi_1(x_1)\psi_2(x_2) \tag{3}$$

for some $\psi_1 : X \to \mathbb{C}$ and $\psi_2 : X \to \mathbb{C}$. Then since

$$\frac{\nabla_1(\psi_1\psi_2)}{\psi_1\psi_2} = \frac{\nabla_1\psi_1}{\psi_1} \tag{4}$$

and similarly for particle 2, we find that the general guidance equation (1) decomposes into the two individual guidance equations

$$\frac{dQ_1}{dt} = \frac{\hbar}{m_1}\mathrm{Im}\left(\frac{\nabla_1\psi_1}{\psi_1}(Q_1)\right) \tag{5a}$$

$$\frac{dQ_2}{dt} = \frac{\hbar}{m_2}\mathrm{Im}\left(\frac{\nabla_2\psi_2}{\psi_2}(Q_2)\right) \tag{5b}$$

So in a case such as this, where the joint wavefunction is simply a product of single-particle wavefunctions, we can make the joint wavefunction Humeanistically acceptable by regarding it as a "conjunction" of duly local individual wavefunctions.

The problem, though, is that generic wavefunctions are entangled, i.e., are not expressible as a product of single-particle wavefunctions. Still with $N = 2$, consider as an example

$$\Psi(x_1, x_2) = \frac{1}{\sqrt{2}}\left(\psi_1(x_1)\psi_2(x_2) + \psi_1'(x_1)\psi_2'(x_2)\right) \tag{6}$$

where $\int_X \psi_1^*\psi_1'\, dx_1 = 0 = \int_X \psi_2^*\psi_2'\, dx_2$. The sum (6) cannot be factorised into a single product, and so we cannot treat it as simply arising from some pair of assignments to the points of X individually.

The nearest proxies for individual wavefunctions, in a case such as (6), are the *conditional* wavefunctions.[8] The conditional wavefunction of particle 1, relative to particle 2's being in location Q_2, is given by

$$\Psi_1^{Q_2}(x_1) := \Psi(x_1, Q_2) \tag{7}$$

and similarly for the conditional wavefunction of particle 2, relative to particle 1's being in location Q_1. More generally, given an N-particle joint wavefunction $\Psi(x_1, \ldots, x_N)$, if we select (say) the first $M < N$ particles as a subsystem, then the *conditional wavefunc-*

[8]The below follows (Dürr and Teufel, 2009, chap. 11).

tion of that subsystem (relative to the configuration of the remaining $N - M$ particles) is given by

$$\Psi^{\mathbf{Y}}_{1\ldots M}(\mathbf{x}) := \Psi(\mathbf{x}, \mathbf{Y}) \tag{8}$$

where $\mathbf{x} := (x_1, \ldots, x_M)$ and $\mathbf{Y} := (Q_{M+1}, \ldots, Q_N)$. That is, the conditional wavefunction of the subsystem is obtained by "saturating" the joint wavefunction with the actual locations of the remaining particles.

The importance of the conditional wavefunction is as follows. Suppose that the joint wavefunction is of the form

$$\Psi(\mathbf{x}, \mathbf{y}) = \phi(\mathbf{x})\psi(\mathbf{y}) + \Psi_{\perp}(\mathbf{x}, \mathbf{y}) \tag{9}$$

where $\mathbf{y} = x_{M+1}, \ldots, x_N$ and $\Psi_{\perp}$ and ψ have macroscopically disjoint $\mathbf{y}$-supports; moreover, suppose that the actual configuration $\mathbf{Y}$ of the environment is in the support of ψ (so that $\Psi_{\perp}(\mathbf{x}, \mathbf{Y}) = 0$ for all $\mathbf{x}$). It then follows that the conditional wavefunction $\Psi^{\mathbf{Y}}_{1,\ldots,M}$ is given by the wavefunction ϕ—and furthermore, that the guidance equation for the subsystem's configuration $\mathbf{X} := (Q_1, \ldots, Q_M)$ reduces to

$$\frac{d\mathbf{X}}{dt} = \frac{\hbar}{\mathbf{m}} \mathrm{Im} \left(\frac{\nabla_{\mathbf{x}} \phi}{\phi} (\mathbf{X}) \right) \tag{10}$$

where $\mathbf{m} = (m_1, \ldots, m_M)$. In such a case, we say that ϕ is an *effective* wavefunction for the subsystem. If the subsystem is sufficiently decoupled from its environment, then the effective wavefunction will also abide by Schrödinger's equation; if there is interaction, however, then it will not evolve in this unitary fashion.

It is, however, important to note that although the conditional wavefunctions of the subsystems can be computed from the "universal wavefunction" Ψ and the actual configuration $\mathbf{Q}$, the reverse is not true: the conditional wavefunctions associated to subsystems underdetermine the joint wavefunction. For example, in the two-particle case, one can easily have a distinct pair of joint wavefunctions $\Psi(x_1, x_2)$ and $\Phi(x_1, x_2)$ such that $\Psi(x_1, Q_2) = \Phi(x_1, Q_2)$ and $\Psi(Q_1, x_2) = \Phi(Q_1, x_2)$: that agreement only requires that they coincide on certain surfaces within configuration space. Moreover, in the case where the subsystem and the environment are coupled to one another, it is not just that the conditional wavefunction does not evolve according to the Schrödinger dynamics—in general, there will not be any autonomous dynamics for the conditional wavefunction at all.

One more remark. In the above, I have followed orthodoxy by supposing that the best way to interpret the wavefunction "ontologically" is as a field of some sort (i.e., as assigning properties to points of configuration space). But as Belot (2011) points out, Bohmians have a reasonably natural alternative: that of interpreting the wavefunction as representing a collective property of the particles. Each possible wavefunction, on this view, would be a kind of dispositional property which specified, for each possible configuration of the collective of particles, how the particles would behave if they found themselves in that configuration. However, this interpretation would be subject to the same problem as the more mainstream interpretation in terms of fields: at any given time, a collective of N particles is an occupant of N spacetime points, not an occupant of a (single) spacetime point, and so the wavefunction is not the kind of local property that can be safely admitted into the Humean Supervenience basis.

3 The Humean response

As mentioned in §1, I am not going to consider responses that modify the Lewisian statement of Humean Supervenience; my interest in this essay is in responses which preserve the letter as well as the spirit of Humean Supervenience. Doing that requires that everything in the supervenience basis——everything that comprises the fundamental ontology—is local in space and time. As we have just seen, though, the wavefunction is *not* spatiotemporally local in the required sense. So that leaves only one option: deny that the wavefunction is part of the supervenience basis.

The natural next question, then, is what the status of the quantum state is on this picture. If standard Bohmian mechanics is indeed to be recovered, then it had better be the case that the wavefunction—like everything else—supervenes upon the supervenience basis, i.e., upon the motions of the Bohmian particles. The proposal by a number of recent authors[9] is that the Bohmian trajectories determine *both* the quantum dynamics *and* the wavefunction, through the same "best-system" method that Humeans take to determine what the laws of nature are. That is, the idea is that of the candidate dynamics-plus-wavefunction packages, precisely one will maximise simplicity and strength (under some appropriate weighting); and *this* package is the one which the Humean takes to be the correct characterisation of what's going on.

In general, these authors seem more or less sympathetic to the idea that the supervenience basis be extremely austere: that it be constituted by nothing other than the

[9]Miller (2014), Callender (2014), Bhogal and Perry (2015), Esfeld et al. (2014); Dickson (2000)also prefigures some of the relevant ideas.

Bohmian trajectories. Such austerity may not be necessary, however. The Bohumean could include other data in their supervenience basis, provided only that such data are appropriately local. (The advantage of doing so is that the richer the supervenience basis, the more plausible it is that the full Bohmian dynamics really will supervene upon it.) For instance, they could perhaps include such particle properties as mass or charge, or (total) spin[10]—provided that such properties are construed as intrinsic properties of the Bohmian particles, rather than characteristics of the wavefunction.[11] They could even include the conditional wavefunction of each particle (relative to the other $N-1$ particles), although this might need some explanation of why the conditional wavefunctions get to be part of the fundamental ontology but the joint wavefunctions do not.[12]

So, the picture is as follows. We take as given our supervenience basis, which certainly includes the Bohmian trajectories, and may or may not include other local data (e.g. particle-properties or the conditional wavefunction). In order to specify the best system, we need to then introduce a new piece of theoretical vocabulary: that of the wavefunction, Ψ. The Humean should then claim that the best system for codifying H is one which asserts the following:[13]

- That Ψ is a complex-valued field on $T \times X^N$
- That Ψ satisfies the Schrödinger equation

[10]That is, the spin quantum number of the particles; not the projection of the spin along some axis, which cannot plausibly be interpreted as a property of the particle rather than the wavefunction (see e.g. (Dürr and Teufel, 2009, §8.4)).

[11]Note that doing so is not entirely straightforward: see Brown et al. (1996).

[12]Esfeld (2014) observes that other primitive ontologies could be used to provide alternative austere supervenience bases; it's not so clear that other primitive ontologies are so amenable to forming the richer bases discussed here, however. For example, if the mass of a particle is to be localised by being taken as a property of the particle, then the primitive ontology for that particle will have to be a point-sized occupant of some spatial point (at each time), as is the case in Bohmian mechanics. In GRWm or GRWf, by contrast, the primitive ontology of the particle is either a region-sized occupant, or else a point-sized occupant of multiple spatial points at each time (and in the case of flashes, sometimes an occupant of no point)—so treating mass as a property of a particle with that primitive ontology would not mean that mass was a local quality.

[13]Bhogal and Perry (2015) use a best system which postulates a space Q (with the structure of X^N) and a particle ω moving around within Q (whose location at any time is exactly correlated with the configuration of the N particles); the wavefunction is then postulated as a function assigning a complex number to each point of Q, which then acts on ω via the guidance equation. If Q here is intended to simply be defined as X^N (i.e., as the space consisting of N-tuples of points of X), then I take these systems to be essentially the same. If not—that is, if the idea is to stipulate Q's structure separately and then put it into appropriate correspondence with N-tuples of points of X—then it seems to me that the system outlined here will be considerably simpler, at no cost in strength.

- That the location of each particle, together with Ψ, satisfies the Bohmian guidance equation
- That the Quantum Equilibrium Hypothesis (QEH) holds:[14] i.e., that the probability distribution of particles[15] is given by $|\Psi|^2$
- Perhaps: that $\Psi(0, \mathbf{x})$ has such-and-such a value at $\mathbf{x}$, for each $\mathbf{x} \in X^N$

We will discuss later whether the Bohumean should think that the final ingredient mentioned here should be included in the best system.

Esfeld et al. (2014), Miller (2014) and Callender (2014) don't characterise their position as involving a non-standard form of Humeanism. For these authors, it remains the case that only the nomological facts arise from a best-systems analysis; so for them, making this Humean move requires treating the wavefunction as nomological rather than ontological. Although they recognise that this treatment of the wavefunction may require some revision of our usual conception of laws,[16] I think that the more significant novelty is that we are utilising a best system whose vocabulary is not confined to terms referring to individuals and properties in the supervenience basis.[17] Bhogal and Perry, however, do discuss this departure from more standard presentations of Humeanism:[18]

> The way we do this is by expanding the language that candidate systems can be formulated in. As before [i.e., in standard Humeanism], systems can use vocabulary that refers to perfectly natural properties (the properties that make up the mosaic)—what we've called the "base language." But in addition to this they can introduce and use any other vocabulary so long as it comes in uninterpreted.
>
> How does such uninterpreted vocabulary come to have content? It can have content if a system links the novel vocabulary to the base language;

[14]See Dürr et al. (1992) for a more detailed discussion of the role of the QEH in Bohmian mechanics. I assume that the Bohumean should include the QEH in the best system, given the role that it plays in the deriving the Born rule within Bohmian mechanics; I thank an anonymous referee for suggesting it.

[15]Presumably, this will involve understanding probabilities in a Humean fashion; as a referee pointed out, this will plausibly require expanding the criteria for best-system-hood to include fit. I will put this complication aside.

[16]Callender, in particular, discusses this in detail.

[17]If the wavefunction Ψ did refer to anything in the basis, then we would instead be dealing with something like the Albert/Loewer/Darby strategy.

[18]Albeit one which—as they observe—is prefigured by Lewis (1994)'s discussion of chance, and Hall (2009)'s discussion of mass and charge.

> that is, if the system contains sentences that contain both novel vocabulary and the already interpreted vocabulary of the base language.[19]

This is not the only place such a liberalised form of Humeanism has been entertained. In the debate over the foundations of spacetime theory, Huggett (2006) has suggested that one can take as fundamental just a basis of Leibnizian distance relations, by using liberalised Humeanism of this kind to introduce inertial structure:

> My proposal is that there are a wider range of strategies that can be employed in the service of systematizing a Humean (in this case, relational) history; the strongest-simplest system might involve laws formulated in terms of natural properties and supervenient properties.[20]

And Stevens,[21] following a suggestion of Pooley,[22] has looked at how such an approach could be extended to a defence of the "dynamical approach" to relativity:[23] here, the suggestion is that only the topological or differential structure of spacetime is taken as fundamental, with all other spatiotemporal structure constructed via a best-systems analysis.

One could also imagine yet further extensions. For instance, one could maintain that the only things which fundamentally exist are phenomenological experiences, with everything else showing up as aspects of the best system for codifying and summarising those experiences; or, one could even apply the same strategy to a fundamental basis consisting only of one's *own* experiences, with a best system constructed from that. Such positions raise a stability question for liberalised Humeans, at least if they don't want to advocate Berkumeanism or Hulipsism (as I propose to call these doctrines): if Bohumeanism is a compelling position, then why not Berkumeanism? After all, by construction these bases are more epistemically accessible than either a basis of Bohmian trajectories. However, I don't intend to pursue this line of thought here.

Instead, in the remainder of this paper, I raise three problems for the aspiring Bohumean (which are not faced by the Humean about laws). The first is a general problem with extending Humeanism from nomic to ontic matters (and so applies equally

[19](Bhogal and Perry, 2015, p. 5)
[20](Huggett, 2006, p. 50)
[21]Stevens (2015), Stevens (forthcoming)
[22]Pooley (2013)
[23]Brown (2005), Brown and Pooley (2006)

to regularity relationalism); the second and third are specific to the Bohumean proposal (and so, are reasons why even those sympathetic to regularity relationalism should resist Bohumeanism).

4 The status of Humean ontology

The first problem is this: what is the status of the new structure (the wavefunction, or the inertial spacetime structure)? Specifically, what is involved in committing to such structure?[24] Prima facie, there are two options. It could be that the new structure does not represent any further ontological commitments beyond those made by commitment to the supervenience basis; or, it could be that the new structure *does* represent a further ontological commitment, albeit to structure that is not metaphysically fundamental. Let us consider these options in turn.

4.1 Ontological innocence

First, the claim that the new structure represents no further commitment. I assume that Bohumeans will want to claim that our discourse about wavefunctions is in good standing, in the sense that when we make claims about wavefunctions we say things which are meaningful and, under the standard success-conditions accepted by the discourse, true. In other words, I put aside the claim that our talk of wavefunctions should be understood as a wholesale or creative fiction, with no more bearing on the actual condition of the world than the work of a Tolkien or a Trump.

Thus, if wavefunction-discourse is not to represent further ontological commitments, then there would seem to be two options: either we are eliminativist about it (so it represents no commitment at all), together with some kind of paraphrasing strategy for understanding wavefunction-talk; or we claim that it is an "ontological free lunch" (so it represents no *further* commitment, beyond what we have committed to in the supervenience basis).[25]

This would make our attitude to the wavefunction parallel to, say, van Inwagen's view of composite objects:

[24]Note that this is a question that only liberalised Humeans must address. Since classical Humeanism doesn't use a best-system analysis to introduce new ontology, it does not need to settle the issue of our commitment to that ontology: classical Humeanism merely delimits the range of nomically possible worlds, rather than modifying the structure of those worlds.

[25]The phrase is due to Armstrong (1997).

> My position vis-à-vis tables and other inanimate objects is simply that there *are* none. Tables are not defective objects or second-class citizens of the world; they are just not there at all. But perhaps this wretched material mode is a part of the difficulty. Let us abandon it. There are certain properties that a thing would have to have to be properly called a 'table' on anyone's understanding of the word, and nothing has all of these properties.[26]

Or to David Lewis' view of composite objects:

> ...Mereology is ontologically innocent.
>
> To be sure, if we accept mereology, we are committed to the existence of all manner of mereological fusions. But given a prior commitment to cats, say, a commitment to cat-fusions is not a further commitment. The fusion is nothing over and above the cats that compose it. It just is them. They just are it. Take them together or take them separately, the cats are the same portion of Reality either way. Commit yourself to their existence all together or one at a time, it's the same commitment either way. If you draw up an inventory of Reality according to your scheme of things, it would be double counting to list the cats and then also list their fusion. In general, if you are already committed to some things, you incur no further commitment when you affirm the existence of their fusion. The new commitment is redundant, given the old one.[27]

Or to the early Carnap's view of non-observational language:

> Quite generally, everything that we talk about must be reducible to what I have experienced. Everything that I can know refers either to my own feelings, representations, thoughts and so forth, or it is to be inferred from my perceptions. Each meaningful assertion, whether it concerns remote objects or complicated scientific concepts, must be *translatable* into a statement that speaks about contents of my own experience and, indeed, at most about my perceptions.[28]

The problem, however, is that in all these cases a very specific kind of connection is posited between the "base" ontology (parts of objects, observational structure) and the

[26](Van Inwagen, 1990, pp. 99-100)
[27](Lewis, 1991, pp. 81–82)
[28](Carnap, 1929, p. 12); quoted and translated in (Coffa, 1991, p. 227).

"superstructure" ontology (composite objects, non-observational structure), which undergirds the claim that the latter encodes no further ontological commitments beyond the former: namely, that claims about the superstructure can be *translated* into claims about the base. Thus, both van Inwagen and Lewis agree that any statement about chairs could (in principle, and assuming a suitably powerful language) be replaced by an appropriately ontologically hygienic statement about simples arranged chairwise,[29] whilst Carnap argues that any statement must admit of a translation into some observational statement.

I claim that the possibility of such a translation is a precondition of either eliminativism or "free-lunch-ism" about a category of ontological objects.[30] If we are really eliminativists about some domain of discourse A, then we are committed to thinking that it is (in principle) dispensable, in the sense that any A-statement is merely a *façon de parler* for some B-statement, and so could (in principle) be replaced by it. Alternatively, if we are free-lunchers about A, then we need to think that the subjects of the A-statements *just are* the subjects of the B-statements; if you tally up your commitments by counting the As and the Bs, then you'll have counted the same thing(s) twice. But it follows from that that anything which is said about an A could be rephrased as an equivalent statement about some Bs—namely, whatever Bs are the ones that *just are* the A.

For example, in the context of relativistic theories, the Levi-Civita connection, and hence the Ricci tensor and Ricci scalar, are definable in terms of the metric (in the sense that a proper statement of the theory, if given in a language including terms for the connection and the Ricci fields, should include the relevant definitions). Consequently, although the Einstein Field Equations are typically phrased in terms of the stress-energy tensor, metric, Ricci tensor and Ricci scalar, it is entirely natural to regard our ontological commitments as exhausted by commitment to the metric and the stress-energy tensor: after all, in principle one could write the EFEs purely in terms of those two quantities. Thus, we can (depending on our tastes and inclinations) either eliminate the metric and regard connection-talk as an elliptical way of talking about

[29]Assuming that there are indeed such simples. For the purposes of the point being made here, we could equally well note that statements about chairs are, in principle, translatable into statements about arbitrarily small parts of chairs.

[30]Note that the projection of articulating and defending such a translation occupies a significant part of van Inwagen's work on mereology; and in his autobiographical reflections, Carnap noted that a key reason for his move away from stricter forms of empiricism was the recognition that "we must abandon the earlier view that the concepts of science are explicitly definable on the basis of observation concepts" (Carnap, 1963, p. 59).

the metric, or take connection-talk as talk of an entity that, though real, is nothing over and above the metric.

However, in the cases of liberalised Humeanism, no such translation is to be had: it will not be the case that we can translate an arbitrary statement about inertial structure (say, that the Sun is moving inertially) into a statement about topological or Leibnizian relations, nor that an arbitrary statement about the wavefunction (say, that its value at this point is such-and-such) into a statement about Bohmian particles. This should not be surprising: if such a translation were on the table, then there would be no need for the Humean manoeuvre at all. One could postulate a set of laws for the Leibnizian or Bohmian structure directly (by translating the laws mentioning the inertial or wavefunction structure into Leibnizian or Bohmian terms), and just taking that theory to be the fundamental theory.

To see that there cannot be such a translation, observe first that whenever we do have a translation of A-discourse into B-discourse, then a given distribution of B-ontology is consistent with at most one distribution of A-ontology.[31] For, suppose that a distribution β of B-ontology were consistent with two distinct distributions α_1 and α_2 of A-ontology. Then α_1 and α_2 must disagree on the truth-value of some A-proposition P: suppose that α_1 makes P true, whilst α_2 makes P false. But if there is a translation of A-discourse into B-discourse, we can translate P into a B-proposition Q. And Q's truth-value will be settled one way or the other by β: either β makes Q true, or β makes Q false. In the first case, β is consistent only with α_1, and the in the second case, β is consistent only with α_2. The examples above illustrate this relationship between definability and uniqueness: a distribution of cat-parts is consistent with just one distribution of cats; had the early Carnap been right, then a distribution of observational structure would have been consistent with just one theoretical structure; and a given metric is consistent with just one Levi-Civita connection.

But in general, a given distribution of the liberalised Humean's basic ontology (topological or Leibnizian relations, Bohmian trajectories) is consistent with various different distributions of the new ontology (inertial structure, wavefunctions). For example, a pair of globes joined by a rod and maintaining a fixed distance from one another is consistent with an inertial structure that takes them to be rotating, and with an inertial structure that takes them to be non-rotating. In the case of Bohmian mechanics, a given trajectory through configuration space is consistent with any two wavefunctions

[31] This is just a metaphysically-dressed-up version of the well-known model-theoretic fact that explicit definability entails implicit definability.

which agree on that trajectory.[32] Indeed, agreement on the trajectory is sufficient but not necessary. Consider a Bohmian particle in a box:[33] that is, a particle with one positional degree of freedom, which is confined to the unit interval $[0, 1]$ (but is otherwise free). Then the energy eigenfunctions of the system are of the form

$$\phi_n(x) = \sin(n\pi x) \tag{11}$$

for $n = 1, 2, 3, \ldots$. As an eigenfunction, ϕ_n evolves under the Schrödinger equation only into states which are equivalent to ϕ_n (up to phase). But by the guidance equation, $dQ/dt = 0$ if the wavefunction is ϕ_n, or if it is any wavefunction equivalent to ϕ_n . So any pair of such eigenfunctions are associated to the same Bohmian trajectory: namely, that of the particle remaining at rest.

Now, this isn't directly a problem for the liberalised Humean: after all, although there are multiple distributions of the full ontology consistent with the distribution of the basic ontology, the Humean will simply claim that at most one features in the best systematisation. ("At most one", because it might be that the best systematisation of a pair of globes or a single static particle in a box makes no mention of the inertial structure or wavefunctions at all.) But it does mean that there cannot be a direct translation of wavefunction-talk into particle-talk; and correlatively, that the prospects for regarding wavefunction-talk as elliptical for particle-talk, or for regarding wavefunctions as nothing over and above particles, look dim.[34]

4.2 Ontological commitment

So now consider the alternative: that the new structure does represent a further ontological commitment, but to structure which is metaphysically non-fundamental. There are two problems with this stance. One is that it starts to become unclear whether anything has really been gained: if it remains the case that we are committed to objectionably non-local or unobservable entities, and if our commitment to those entities is a further commitment (beyond our commitment to the unobjectionable entities in the Humean supervenience basis), then what does it matter whether such entities have been treated to the honorific "fundamental" or not? Presumably, the

[32] Slightly more carefully, which agree on an arbitrarily small region around that trajectory.

[33] My presentation of this example follows Belot (2011), although the example goes back (at least) to Einstein (1953): see Myrvold (2003).

[34] It also leads into some difficulties about the relationship between models of the laws and nomic possibility; see §5 below.

attractiveness of the thesis of Humean supervenience is either due to some principle of parsimony, or else a feeling that entities not admissible to the basis (by virtue of their non-locality, for instance) are somehow metaphysically dubious. But any plausible parsimony principle will surely enjoin us to reduce our ontological commitments in general, not merely to reduce our ontological commitments to *fundamental* objects in particular.[35] And a commitment to something metaphysically dubious is still queasy-making, independently of whether that something is fundamental or non-fundamental.

The second problem is that if the new ontological structure is a further commitment, then truth is not grounded in being (or at least not in the way one would normally expect). For consider a proposition about the new structure; say, about the wavefunction. First, suppose that this proposition is one included in the best system (e.g. the proposition that there is a wavefunction). The question is then: what makes this proposition true? What is the ground of the truth of this proposition? Normally, one would claim that this proposition is true in virtue of the way in which the wavefunction is, that is, in virtue of the state of the wavefunction. But according to the liberalised Humean, what makes it the case that there is a wavefunction at all is simply the fact that the proposition *that there is a wavefunction* is a part of the best systematisation of the fundamental facts. It is by virtue of being a part of the best system, together with the Bohumean analysis of what a wavefunction is, that that proposition gets to be true; and it is (seemingly) in virtue of the truth of that proposition that a wavefunction exists. Thus, the existence of the wavefunction is grounded in the truth of the proposition *that the wavefunction exists*, and not vice versa.

Of course, the Bohumean can (and should) claim that the truth of thsi proposition about the wavefunction is grounded in the facts about the distribution of the Bohmian particles over time, so its truth is grounded in being *in general*. But it remains the case that its truth is not grounded in the being *of its subject matter*, which is a rather odd state of affairs—especially given that (*ex hypothesi*) propositions about the wavefunction are not to be understood as elliptical claims about the Bohmian particles.

Could the liberal Humean claim that the grounding goes both ways: i.e., that the truth of the proposition *that a wavefunction exists* is grounded in the existence of the wavefunction, and vice versa? Only by advocating a peculiar view of grounding, according to which the grounding relation is not asymmetric. [***expand?***] Note that

[35] Except, that is, insofar as our non-fundamental commitments are nothing over and above our fundamental commitments. But clearly, that just takes us back to the first option discussed already.

commitment to circular grounds is something that Humeans have generally sought to avoid: for example, it is precisely in order to avoid circular grounds that Humeans regarding laws have argued for a distinction between *scientific* grounding and *metaphysical* grounding (the idea being that laws are metaphysically grounded in their instances, as per Humean orthodoxy, but that instances are scientifically grounded in laws, as per the usual standards of scientific explanation).[36]

5 The problem of determinate quantities

The second problem concerns the extent to which the wavefunction is a determinate quantity on the Bohumean picture. Again, we need to consider two possible ways the Bohumean could make their case: either it is the case that the best system includes a full specification of $\Psi(x, 0)$—and so settles every proposition about the wavefunction's value at all points of spacetime—or it does not. I'll call a best system which does include such a specification a *full system*, and a best system which does not a *partial system*.

5.1 If the best system is a full system

If the best system does specify the value of the wavefunction at all points, then there are two problems. First, there is the fact that treating the wavefunction as lawlike is somewhat counterintuitive; since this is a well-known issue (which arises on other treatments of Bohmian mechanics), and one already discussed at length in the literature,[37] I don't intend to dwell on it here.

The second problem is specific to Bohumeanism: it is just that it is prima facie highly implausible that a specification of the wavefunction's value everywhere in space could possibly be a part of the best system, since it is such a massive quantity of information. In David Albert's vivid parable, a Humean summary is supposed to be "something meaty and pithy and helpful and informative and short that [God] might be able to tell you about the world which (you understand) would not amount to everything, or nearly everything, but would nonetheless somehow amount to a lot";[38] and (to put it mildly) it is hard to see how the specification of a complex number for every point of

[36] See Loewer (2012), Hicks and van Elswyk (2015).

[37] See, in particular, Dürr et al. (1997) and Callender (2014) for defences of the claim that the wavefunction may be considered lawlike.

[38] (Albert, 2015, p. 23)

a vastly high-dimensional configuration space is supposed to fit that bill.

If one was really wedded to the idea that the best system should be a full system, then the only plausible response would be to rely on the fact that specifying a wavefunction will give much more information about what the trajectories are going to do, i.e., will generate a stronger system. Perhaps, with the right way of balancing simplicity against strength, the former factor will win out and the facts about the wavefunction get counted as part of the best system. In that case, however, a different problem rears its head. If the gains in strength from specifying the wavefunction could be paid for in the coin of simplicity, why would the same not be true of specifying the (initial) positions of the particles? After all, much less data is involved in specifying where the particles are than in specifying what the wavefunction is up to: the former is just $3N$ real numbers, rather than uncountably many complex numbers (one for every point of space). But the gains from this extra data are enormous, since (together with the wavefunction) one obtains a perfect prediction of everything to happen at every moment. That would then, *ad absurdum*, make the initial conditions a lawlike matter.[39] So the Bohumean who wants a full system faces a dilemma, neither horn of which is palatable. If their rules for assessing the best system do not give strength enough weight, then they cannot make it plausible that a direct specification of the wavefunction (at a time) will be included. And if their rules do give strength enough weight to avoid this, then they cannot make it plausible that a direct specification of the particle positions (at a time) will be excluded.

5.2 If the best system is not a full system

If the best system is not a full system, then there is a further choice point: either the wavefunction does have a determinate value everywhere in spacetime, or it does not. If it does not, then it becomes a vague and indeterminate sort of entity. On that analysis, it would be true that the wavefunction exists; given that the wavefunction is defined as a map from times to complex-valued fields over configuration space, it would presumably follow that the proposition *at t, either the phase difference between x and x' is $\pi/2$ or it is not $\pi/2$* is true (since that proposition is true for any such complex-valued field); yet neither disjunct has a determinate truth-value. So it would turn out that quantum mechanics exhibits just the kind of indeterminacy and disrespect for classi-

[39] cf. (Hall, 2009, §5.6)

cal logic that Bohmians are apt to deny![40]

On the other hand, supposing that the wavefunction is everywhere determinate is hardly more palatable. For one thing, the problem of ungrounded propositions (raised in the last section) becomes even more acute. If the proposition that *the phase difference between* $\psi(t, x)$ *and* $\psi(t, x')$ *is* $\pi/2$ is true but not settled (even via a best-system analysis) by the Humean mosaic of Bohmian particles, then it starts to look awfully as though it could only be true primitively, in a manner that is not grounded in being at all.

Moreover, even if we suppose that the wavefunction is determinate in the actual world, then there is still a problem regarding its determination in worlds that are nomically possible relative to the actual world. It is worth noting that in this regard, the Bohumean is worse off than the regularity relationalist.[41] For, consider again the "Newton's globes" world: i.e., a world consisting of a pair of globes, joined by a rod, which remain a fixed distance from one another. Now, the regularity relationalist can (and should) simply deny that this is a world which defines any inertial frames at all: its motions are too impoverished to admit of interesting Newtonian laws, of the kind that would require the introduction of inertial frames. The problem, however, is that this world can also be analysed *from the perspective of a world where Newton's laws do hold*—that is, in which Newton's laws (and the attendant commitment to inertial frames) are part of the best system. That system is supposed to yield an account of nomic possibility, which can be used for assessing scientific counterfactuals, explanations, etc. But, at least on the face of it, that best system admits distinct nomic possibilities (distinct models of the laws) which correspond to the same relational facts: e.g. the possibility according to which the spheres are rotating, and the possibility according to which they are not. So, in effect, there are not enough metaphysically possible worlds to account for all the nomically possible worlds; as a result, the laws fail to determine a state of rotation for the globes world, as that single world has to do double duty for possibilities with different states of rotation.

Huggett's response is to point out that a relational world will not (he claims) typically be best systematised *merely* by Newtonian mechanics. Rather, the best system will consist of Newtonian mechanics together with specific force laws, specifying (for example) how much a rod would need to stretch in order to exert a given restoring

[40]If one is some kind of fictionalist about the wavefunction, then this isn't so bad: the Sherlock Holmes canon, after all, fails to settle the proposition *that Holmes has a mole on his left shoulder*. But Bohumeanism advertises itself, at least, as a means of reconciling a commitment to locality with realism about quantum mechanics.

[41]See (Huggett, 2006, §3.2)

force. And that system will not admit both the rotating- and stationary-globes worlds as nomic possibilities: the force laws will specify how much restoring force is present in the rod, which in turn determines what the acceleration (and hence state of rotation) of the globes. This response depends, of course, on the assumption that there are no absolutely rigid bodies; but that seems eminently reasonable.

However, this response does not extend to just any liberalised Humean proposal, as Huggett notes:

> One can imagine accounts of supervenient dynamical quantities akin to mine, but with the feature that although the laws determine the values of the quantities in worlds in which they are the laws, they do not in all worlds in which they are true. That would mean that there were models of the laws—distinguished by different possible values of the quantities—that corresponded to the same nomically possible world. And that would count against the account it seems.[42]

One can probably guess the punchline: the Bohumean case turns out to be just such an account. Consider again the particle in a box discussed in §4.1. Here again, we have a number of distinct models of quantum mechanics, which should therefore be counted as distinct nomic possibilities (according to the laws of any world which is best systematised by Bohmian mechanics). But because they all agree on the Bohmian trajectories, according to the Bohumean, they correspond to the same possible world. And unlike the case of the rotating globes, these distinct models are all models of the same set of laws, even if we include "specific" laws such as force laws. Thus, the newly introduced quantities fail to have determinate values at all nomically possible worlds, and so Bohumeanism fails a test proposed by fellow liberal Humeans (according to Huggett, "we should reject accounts that lead to laws which fail to determine the supervenient quantities in all nomically possible worlds").[43]

6 Scientific practice

6.1 The direct argument from scientific practice

The third problem concerns how we are to verify—or at least, have reason to believe—the claim that the best systematisation of the Humean mosaic is indeed Bohmian me-

[42](Huggett, 2006, p. 60)
[43](Huggett, 2006, p.70)

chanics (or something like it). The most explicit way to do so would be to first fix some set of trajectories which is a plausible candidate to represent the actual evolution of the world; then determine some way of measuring the combined simplicity-plus-strength, relative to those trajectories, of candidate "packages" of differential equations and wavefunctions (or other wavefunction-like things, appropriate to differential equations different from Bohmian mechanics); and then show that the package consisting of the guidance equation, Schrödinger equation, and some universal wavefunction are maximal with respect to that measure.

This is an *insanely* difficult problem. First, we need to overcome the formidable hurdles of finding an appropriate means of evaluating candidate packages. Second, even given such a means, it would be extraordinary if the project of delineating those sets of trajectories for which a Bohmian package is indeed the best system proved to be even remotely mathematically tractable. Third, it is rather opaque what would be involved in showing that a given Bohmian distribution is "a plausible candidate to represent the actual evolution of the world"; but given that at least a necessary condition would be that the distribution contain an unbelievably large number of particles, the prospects for doing so do not look good. In other words, if the Bohumean is going to convince us that their supervenience basis is best systematised by Bohmian mechanics, they aren't going to do so by direct computation.

In this regard, of course, they are in the same boat as standard Humeanism about laws of nature. In general, Humeans have not sought to show directly that such-and-such a theory is the best codification of such-and-such primitive categorical facts. (Although it is worth noting that the direct computation is even less possible for the Bohumean than for the standard Humean, given that we're now allowed to introduce new theoretical vocabulary into the best system. This means that the available systems of equations to consider are not limited to just those equations employing only a fixed stock of variables and parameters (i.e., those ranging over the supervenience basis); rather, we must consider any equations whose variables and parameters include that fixed stock.) Instead, Humeans have usually taken the practice of science itself to provide some reason for thinking that our actual scientific theories—or some extension thereof—are plausible candidates for being the best systematisation of the physical facts.

The reason this manoeuvre has legs is because the Humean standards of simplicity and strength are not simply plucked from thin air: rather, they are exactly the standards to which (it is claimed) science typically adverts in determining what the laws are, given the evidence. Thus, Lewis remarks that he "[takes] a suitable system to be one that has the virtues we aspire to in our own theory-building, and that has them to the greatest extent possible given the way the world is";[44] and that "the standards of simplicity, of strength, and of balance between them are to be those that guide us in assessing the credibility of rival hypotheses as to what the laws are."[45] Hall (2015) identifies this as the "second guiding idea" of Humean supervenience:[46]

> What the second guiding idea really needs to assume is that there are, implicit in the practice of physics, evidential standards for determining what the fundamental physical laws are that induce a mapping from possible *total bodies of evidence* to something like a probability distribution (or perhaps a family of such distributions) over propositions about the fundamental laws of nature. [. . .] we can summarize the second guiding idea this way: the Humean reductionist is taking standards that *both* sides endorse—but that his anti-reductionist opponent views solely as epistemic standards—and elevating them to the status of standards constitutive of the laws of nature.[47]

Call this argument the *direct argument from scientific practice*.

Now, merely liberalising our Humeanism will not cut off support from this argument. Indeed, Huggett makes use of essentially the same idea in defending his regularity relationalism:

> Consider Newton's *Principia*. In the final part of his *Scholium* to the definitions Newtona acknowledges that part of the problem facing him is to determine the 'true motions' from the 'merely relative'; to work out the absolute accelerations given that only relative motions are observable. How does he do this? In large part, by showing that his dynamics [. . .] provides the strongest-simplest axiomatisation of the relative motions of the planets.[48]

[44](Lewis, 1983, p. 367)

[45](Lewis, 1986, p. 123)

[46]The first guiding idea is the proposal that the relevant standards are, in particular, those of simplicity and strength. Hall argues convincingly that we should take the second guiding idea as the core proposal, with the first guiding idea being a substantive (epistemic-cum-sociological) proposal about what the standards of science in fact are.

[47](Hall, 2015, p. 266)

[48](Huggett, 2006, p. 48)

Can the Bohumean make a similar appeal to scientific practice, to justify their claim that Bohmian mechanics is the best systematisation of the Bohmian trajectories? Doing so would mean defending the following claims: that the evidence available to scientists is (some portion of) the facts in the supervenience basis; and that in seeking to systematise this evidence, scientists have come up with Bohmian mechanics. Unfortunately for the Bohumean, however, both of these claims are false.

First, the claim that the evidence basis for scientists comprises certain facts about the Bohmian trajectories. The idea here will be (to paraphrase Bell) that all measurements are measurements of position;[49] and so—on the view of the world advanced by Bohmians—any experimental data can be characterised in terms of the positions of the Bohmian particles. The problem, however, is that our evidence for quantum mechanics is (famously) *statistical* in nature. It is not that we have direct access to some small number of the Bohmian trajectories, and have successfully stitched those together by overlaying a wavefunction governed by quantum dynamics. What we have instead are individual but imprecise measurements of positions at particular times. By making many such measurements of identically prepared systems, and looking at the frequency distributions of the results, we can obtain high confirmation of the probability densities over such trajectories (on the Bohmian picture). So what we have really woven together into a quantum tapestry are those probability densities, rather than the trajectories themselves; and *on the Bohmian's own account*, those probability densities represent all that can ever be known for sure about the trajectories.[50]

Now, this might seem like nit-picking: surely probabilistic evidence about Bohmian trajectories is still evidence. The trouble is that in order to establish an argument from scientific practice analogous to that available to other Humeans (including regularity relationalists), the Bohumean needs to establish that the particle trajectories are epistemically prior to the wavefunction, just as the relational motions are epistemically prior to the structure of spacetime. Such a relation of epistemic priority would ensure that any disputant would have to be inferring the laws from an evidential basis coinciding with the Humean's supervenience basis: in the spacetime case for instance, Huggett observes that "the only evidence available to Newton—to anyone, whatever spatiotemporal ontology they adopt—is that of a relational history".[51] But because

[49] "[...] in physics the onl yobservations we must consider are position observations, if only the positions of instrument pointers." (Bell, 1987, p. 166)

[50] Of course, even our knowledge of the probability densities is somewhat indirect. But the point being made here is just that it is nevertheless *more* direct than whatever knowledge we might have about trajectories.

[51] (Huggett, 2006, p. 49)

our access to the trajectories is only via the probability densities, no such priority relation can be established. Yes, we can use the probability densities to make (partial) inferences about the trajectories; but we can equally well use them to make inferences about the wavefunction directly. So what scientific practice (and theory) shows us here is that the probability densities are epistemically prior to *both* the trajectories and the wavefunction.

Nevertheless, one could still argue that a weaker form of the argument from scientific practice holds: if scientific practice has indeed used Bohmian mechanics to codify the (partial) information about trajectories represented in the probability densities, then that would still be some evidence in favour of Bohmian mechanics' credentials as the best systematisation of the trajectories. However, just as a matter of sociological fact, it is false that the scientific community has alighted on Bohmian mechanics as the preferred theory for explaining and systematising quantum phenomena. What they have in fact come up with is what we might call "ordinary" or "orthodox" quantum mechanics:[52] a messy, foundationally unclear, and yet incredibly empirically successful combination of systematic dynamics, particular models, and pragmatic rules for extracting empirical content. The exact content of this theory is still a matter of some debate (for instance, whether the projection postulate or the eigenvalue-eigenstate link should be included);[53] but I take it to be reasonably clear that no matter how precisely one draws the boundary, the theory does not coincide with Bohmian mechanics.

One might hold that this is irrelevant, given the empirical equivalence between Bohmian mechanics and textbook quantum mechanics:[54] doesn't that show that Bohmian mechanics and textbook quantum mechanics are *equally* capable of systematising the relevant data, and hence that it makes no odds (so far as the argument from scientific practice is concerned) whether scientists have adopted one or the other? But whether or not something is the best system is not invariant under empirical equivalence. After all, textbook quantum mechanics is empirically equivalent to the theory consisting of all and only its observational predictions—but no-one is going to think that *that* theory is a serious candidate for best system. So we cannot use the empirical equivalence of Bohmian and textbook quantum mechanics to argue that they are equally well (or poorly) qualified to be best systems.

Alternatively, one might just think it is *obvious* that no such hodge-podge as textbook

[52]Note that I take this to include both non-relativistic and relativistic quantum mechanics.
[53]See Wallace (unpublished).
[54]My thanks to an anonymous referee for raising this concern.

quantum mechanics could possibly be the best systematisation of the empirical data.[55] But that's just a reason to think that the argument from scientific practice is not a good argument: if it's clearly false that textbook quantum mechanics is the best system, then that shows that actual working scientists do not always converge upon the best system, not that they have not converged upon textbook quantum mechanics.

6.2 The indirect argument from scientific practice

Nevertheless, following this line of thought suggests a different move that the Bohumean could make. The Bohumean could concede (at least for the sake of argument) that orthodox quantum mechanics is a better systematisation of the empirical data than Bohmian mechanics, but argue that this is irrelevant: the question at hand is what the best systematisation of the Bohmian trajectories is, not what the best systematisation of the empirical data is. Now obviously, just asserting that claim by itself is question-begging. But Bhogal and Perry outline a way of using scientific practice to indirectly support this claim:

> This worry, that mere positional facts wouldn't be complicated enough to distinguish something like Bohmian Mechanics as the best system of that world, strikes us as far too pessimistic. One of the key motivating thoughts behind the best system account is that whatever an ideal scientist, if she was fully rational and knew everything about the state of the mosaic, would take to be the best overall theory given the evidence is the best system of that world.
>
> Actual scientists are not ideal reasoners and they do not have access to the entirety of the facts about the mosaic. Of the elements of the mosaic, actual scientists only have direct access to facts about positions. [...]
>
> If we look to actual scientific practice, we see that physicists, even with access to only a tiny slice of the position facts, have a great deal of confidence that the world is quantum mechanical (and consider this position very well

[55] Indeed, in the literature on the philosophy of quantum mechanics one sometimes sees the view that orthodox quantum mechanics is so muddled and incoherent as to fail to have any content at all—which would surely guarantee that it is not the best system, regardless of whether scientists have adopted it or not. But this just seems to me to be a reason to be very sceptical of that claim about the content of orthodox quantum mechanics; indeed, one of the most striking features of quantum physics is the extraordinary empirical success it has enjoyed *despite* the absence of any widely agreed account of its foundations. (Compare the discussion in Wallace (2012) of the role decoherence plays in making the measurement problem "a philosophical rather than a practical problem" (p. 4586).)

confirmed). If this, in the grand scheme of things, meager set of position facts is enough to satisfy non-ideal working scientists, then we see very little reason to be skeptical that the ideal scientist, with access to all the position facts at our Bohmian world, would settle on a Bohmian Mechanical physical theory.[56]

We could summarise this argument as follows:[57]

1. Non-ideal actual scientists have empirical access to coarse-grained data about the trajectories.

2. Those scientists have come up with quantum mechanics as the best systematisation of that data.

3. If an ideal scientist were already committed to quantum mechanics, then Bohmian mechanics would be the best way to extend their commitment to include the claim that particles follow exact trajectories.

4. So if an ideal scientist could somehow come to know that particles have trajectories, then they would advocate Bohmian mechanics.

The argument is indirect because it still uses a counterfactual premise (premise 3) about what an ideal scientist would do with access to certain data (whereas the direct argument from scientific practice seeks to avoid making assumptions of this kind). So inevitably, the support for this argument will be somewhat weaker, given that counterfactuals like premise 3 can't be as conclusively verified; of course, it also makes the argument harder to decisively refute. So what I want to do here is just offer some plausibility considerations against premise 3.

First, we need to be a little careful about the sense in which Bohmian mechanics and quantum mechanics are empirically equivalent. That empirical equivalence means that *over those situations where both theories apply*, they will generate the same predictions. However, at least as things currently stand, there are many situations to which quantum mechanics, but not Bohmian mechanics, can be successfully applied. Most notably, although it is an ongoing (and important) frontier of research,[58] there is currently no Bohmian version of quantum field theory capable of fully replicating the

[56](Bhogal and Perry, 2015, p. 18)

[57]I'm grateful to an anonymous referee for suggesting this way of formulating the argument, and for suggesting that Bhogal and Perry should be interpreted as making this argument rather than the direct argument.

[58]See Struyve (2010) for a survey.

standard formalism; that cuts off support from the predictive success of high energy physics. Thus, Bohmian mechanics is less strong than textbook quantum mechanics. Of course, this isn't to say that this will always remain the case: extending the scope and range of Bohmian analyses is an important and ongoing research project. The point being made here is just that whilst that project is still ongoing, comparing the two theories on strength will favour textbook quantum mechanics.

Second, there is a question about whether the (counterfactual) scientists' knowledge of trajectories is itself consistent with Bohmian mechanics. After all, it is part of the content of Bohmian mechanics that the trajectories cannot be precisely known by any physical means: scattering a photon off the particle, say, will perturb the particle to at least the degree of precision given by the shortness of the photon wavelength. So to posit Bohmian mechanics as the best systematisation of the trajectories would mean characterising one's own knowledge of those trajectories as supernatural. This means that we are dealing with a kind of epistemic instability, analogous to that in the case of Boltzmann brains: the very theory that is supposed to codify and account for a certain kind of evidence also points to that evidence being unreliable or inadmissible.

Third, recall that (by hypothesis) the particles are distributed in such a way that the best Bohmian systematisation is one according to which the Quantum Equilibrium Hypothesis holds. But that means that the wavefunction by itself suffices to determine the probability distribution of the Bohmian particles. Hence, by characterising the dynamics of the wavefunction, textbook quantum mechanics also serves to characterise the dynamics of that probability distribution. And, as is well-known, there is no in-principle barrier to thinking that the best system might be one that merely characterises the probabilities for how things are distributed (rather than their categorial distribution): provided that the gains in simplicity are enough to offset the loss in strength, the probabilistic system can win out over the deterministic one.

Now, the point I'm trying to make here is not that textbook quantum mechanics is, necessarily, the best systematisation. Indeed, I am confident that one could, with a little creativity, cook up plausibility arguments that point to some form of GRW dynamics, or a stochastic particle theory, or something different yet again, as the best system. The aim of the above is to illustrate that, until and unless the Humean tells us considerably more about the criteria for evaluation of systems, any guess at what the best systematisation of some data might be will be highly contestable, and constrained only by the ingenuity of philosophers. And of course, these difficulties are only increased when—as here—that data is itself empirically inaccessible.

7 Conclusion

To conclude, let's take a brief step back, and think about how Bohumeanism compares to Humeanism about classical physics. The reason the problems in the previous section arise is that the Bohumean is, in one crucial respect, worse off than her classical cousin: the latter could, at least, identify the kind of structure in the supervenience basis (i.e., intrinsic properties of points or pointlike things) with the experimental data that (idealised) science collects, and hence argue that the vast parallel-processor of the scientific enterprise has in fact systematised that data into an optimally simple and strong codification. By doing so, the classical Humean can relieve some of the pressure to make precise the nature of the best systematisation they envisage, or to show that such a thing is even possible, since science itself could be taken as demonstrating a proof of principle. The experimental basis for quantum mechanics, on the other hand, is a poor fit with the supervenience basis of the Bohumean. On the one hand, it is too big: it covers many more situations than those to which Bohmian mechanics (at its current state of development) is readily applied. On the other, it is too small: the proposed supervenience basis (even over some local region) goes far beyond what could be gathered by empirical investigation (even in principle). Without this tight fit between the supervenience basis and the empirical basis, I don't see how empirical practice can be a source of optimism that Bohmian mechanics is, indeed, the best systematisation of the supervenience basis.

But this prompts a further question. Classically, a significant component of the motivation for Humean Supervenience has been taken to be *epistemic*: since what we have direct epistemic access to (the thought goes) are facts about intrinsic properties of individual spacetime points or pointlike entities, we should seek a metaphysics founded upon those facts. (This isn't to claim that Humeans are *committed* to this claim about the nature of scientific evidence; it's just that without such a claim, it's not obvious what the advantage is of insisting that all physical facts be local facts.) Now, one can certainly criticise this move, from a premise about what is epistemically available to a conclusion about what is metaphysically acceptable.[59] If, though, the practice of quantum physics does not help the Humean, then we should start to question the antecedent claim too. After all, we do in fact perform entanglement experiments, which (at least on some interpretations of quantum mechanics) constitute the observation of non-Humean facts! So what is going on?

[59] See Maudlin (2007) for a particularly biting critique.

The answer, I contend, is that although individual observations are indeed (somewhat) localised, it just does not follow that those observations cannot provide information about or evidence for irreducibly global goings-on. Prima facie, at least, the way in which one does so is about the simplest imaginable: we simply make multiple local observations, and then *aggregate* those observations. Suppose, for example, that mass was not locally conserved, but was conserved on some larger scale—let's say, on the scale of the Earth. It is straightforwardly possible to accumulate evidence for this hypothesis, by making continuous observations at different points of space, and then comparing the results. Mass disappears here, we find; but we then find that just the same quantity reappeared elsewhere, at exactly the same time. Obviously, no one observer could simultaneously verify the reappearance and the disappearance of the mass. But that's not a problem, given that they can write down their results and compare them, at a later date, with other observers. And clearly, this kind of process is somewhat more involved than the experimental processes needed to confirm or disconfirm a purely local phenomenon—and were the non-local phenomena more widespread (either covering a larger, or concerning more kinds of phenomena), then it might well move beyond our capacities to verify it. But less pervasive non-locality seems like something well within our confirmatory capacities.

The point of this little parable, of course, is that it's more or less exactly what we do to verify the non-local aspects of quantum mechanics:[60] we make simultaneous local measurements in multiple locations, and then bring the results together to compare them. So the simplistic picture of scientific evidence that seemingly motivates the doctrine of Humean Supervenience is long due retirement—and with it, the insistence that our best scientific theories be made to fit that doctrine, at whatever price.[61]

References

Albert, D. Z. (1996). Elementary Quantum Metaphysics. In *Bohmian Mechanics and Quantum Theory: An Appraisal*, volume 184 of *Boston Studies in the Philosophy of Science*, pages 277–284. Kluwer Academic Publishers, Dordrecht.

Albert, D. Z. (2015). *After Physics*. Harvard University Press.

[60] That is, according to those accounts of quantum mechanics in which there are non-local physical phenomena.

[61] Of course, this isn't to say that locality isn't still desirable, and worth pursuing where there is not a significant cost to so doing; it is just that we should not contort ourselves or our theories in order to obtain it.

Armstrong, D. M. (1997). *A World of States of Affairs*. Cambridge University Press.

Bell, J. S. (1987). On the impossible pilot wave. In *Speakable and Unspeakable in Quantum Mechanics: Collected Papers on Quantum Philosophy*, pages 159–168. Cambridge University Press, Cambridge.

Belot, G. (2011). Quantum states for primitive ontologists. *European Journal for Philosophy of Science*, 2(1):67–83.

Bhogal, H. and Perry, Z. (2015). What the Humean Should Say About Entanglement. *Noûs*.

Brown, H. R. (2005). *Physical Relativity: Space-Time Structure from a Dynamical Perspective*. Oxford University Press, Oxford.

Brown, H. R., Elby, A., and Weingard, R. (1996). Cause and Effect in the Pilot-Wave Interpretation of Quantum Mechanics. In *Bohmian Mechanics and Quantum Theory: An Appraisal*, volume 184 of *Boston Studies in the Philosophy of Science*, pages 309–319. Kluwer Academic Publishers, Dordrecht.

Brown, H. R. and Pooley, O. (2006). Minkowski Space-Time: A Glorious Non-Entity. In *The Ontology of Spacetime*, volume 1, pages 67–89. Elsevier.

Butterfield, J. (2006). Against Pointillisme About Mechanics. *British Journal for the Philosophy of Science*, 57(4):709–753.

Callender, C. (2014). One world, one beable. *Synthese*, pages 1–25.

Carnap, R. (1929). Von Gott und Seele. Unpublished lecture. Item number RC 089-63-02, Archives of Scientific Philosophy in the Twentieth Century, Department of Special Collections, University of Pittsburgh.

Carnap, R. (1963). Intellectual Autobiography. In Schilpp, P. A., editor, *The Philosophy of Rudolf Carnap*, number 11 in The Library of Living Philosophers, pages 3–84. Open Court, La Salle, IL.

Coffa, J. A. (1991). *The Semantic Tradition from Kant to Carnap: To the Vienna Station*. Cambridge University Press, Cambridge.

Darby, G. (2012). Relational Holism and Humean Supervenience. *The British Journal for the Philosophy of Science*, 63(4):773–788.

Dickson, M. (2000). Are There Material Objects in Bohm's Theory? *Philosophy of Science*, 67(4):704–710.

Dürr, D., Goldstein, S., and Zanghí, N. (1992). Quantum equilibrium and the origin of absolute uncertainty. *Journal of Statistical Physics*, 67(5-6):843–907.

Dürr, D., Goldstein, S., and Zanghì, N. (1997). Bohmian mechanics and the meaning of the wave function. In Cohen, R. S., Horne, M., and Stachel, J., editors, *Experimental Metaphysics: Quantum Mechanical Studies for Abner Shimony*, volume 1. Kluwer Academic Publishers, Boston.

Dürr, D. and Teufel, S. (2009). *Bohmian Mechanics: The Physics and Mathematics of Quantum Theory*. Springer Science & Business Media.

Einstein, A. (1953). Elementäre Überlegungen zur Interpretation der Grundlagen der Quanten-Mechanik. In *Scientific Papers Presented to Max Born*, pages 33–40. Hafner, New York.

Esfeld, M. (2014). Quantum Humeanism, or: Physicalism without properties. *Philosophical Quarterly*.

Esfeld, M., Hubert, M., Lazarovici, D., and Dürr, D. (2014). The Ontology of Bohmian Mechanics. *The British Journal for the Philosophy of Science*, 65(4):773–796.

Hall, N. (2009). Humean reductionism about laws of nature. Unpublished MS. Available at http://philpapers.org/rec/HALHRA.

Hall, N. (2015). Humean Reductionism about Laws of Nature. In Loewer, B. and Schaffer, J., editors, *A Companion to David Lewis*. John Wiley & Sons.

Hicks, M. T. and van Elswyk, P. (2015). Humean laws and circular explanation. *Philosophical Studies*, 172(2):433–443.

Huggett, N. (2006). The Regularity Account of Relational Spacetime. *Mind*, 115(457):41–73.

Lewis, D. (1983). New work for a theory of universals. *Australasian Journal of Philosophy*, 61(4):343–377.

Lewis, D. (1986). *Philosophical Papers, Volume II*. Oxford University Press, Oxford.

Lewis, D. (1991). *Parts of Classes*. Basil Blackwell, Oxford.

Lewis, D. (1994). Humean Supervenience Debugged. *Mind*, 103(412):473–490.

Loewer, B. (1996). Humean Supervenience. *Philosophical Topics*, 24(1):101–127.

Loewer, B. (2012). Two accounts of laws and time. *Philosophical Studies*, 160(1):115–137.

Maudlin, T. (2007). Why Be Humean? In *The Metaphysics Within Physics*. Oxford University Press, Oxford.

Miller, E. (2014). Quantum Entanglement, Bohmian Mechanics, and Humean Supervenience. *Australasian Journal of Philosophy*, 92(3):567–583.

Myrvold, W. C. (2003). On some early objections to Bohm's theory. *International Studies in the Philosophy of Science*, 17(1):7–24.

Pooley, O. (2013). Substantivalist and relationalist approaches to spacetime. In *The Oxford Handbook of Philosophy of Physics*, pages 522–586. Oxford University Press, Oxford, UK.

Stevens, S. (2015). The Dynamical Approach as Practical Geometry. *Philosophy of Science*, 82(5):1152–1162.

Stevens, S. (forthcoming). Regularity Relationalism and the Constructivist Project. *The British Journal for the Philosophy of Science*.

Struyve, W. (2010). Pilot-wave theory and quantum fields. *Reports on Progress in Physics*, 73(10):106001.

Teller, P. (1986). Relational Holism and Quantum Mechanics. *The British Journal for the Philosophy of Science*, 37(1):71–81.

Van Inwagen, P. (1990). *Material beings*. Cornell University Press, Ithaca, NY. OCLC: 231198754.

Wallace, D. (2012). Decoherence and its role in the modern measurement problem. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 370(1975):4576–4593.

Wallace, D. (unpublished). What is orthodox quantum mechanics? *arXiv:1604.05973 [quant-ph]*.