

The Physical Salience of Non-Fundamental Local Beables

Matthias Egg

University of Bern, Institut für Philosophie, Postfach, CH-3000 Bern 9, Switzerland

Abstract

I defend the idea that objects and events in three-dimensional space (so-called local beables) are part of the derivative ontology of quantum mechanics, rather than its fundamental ontology. The main objection to this idea stems from the question of how it can endow local beables with physical salience, as opposed to mere mathematical definability. I show that the responses to this objection in the previous literature are insufficient, and I provide the necessary arguments to render them successful. This includes demonstrating the legitimacy of dynamical considerations in the derivation of local beables and responding to the threat stemming from the availability of different sets of local beables in the context of the GRW theory.

Keywords: Local beables, Primitive ontology, Quantum mechanics, Wave function realism

1. Introduction

A central issue in the debate on the ontology of quantum mechanics concerns the status of objects and events situated in three-dimensional space, the so-called *local beables*. According to one camp in the debate, to which I will refer as *primitivist*, local beables need to be postulated at the fundamental ontological level, thus specifying what has come to be known as the *primitive ontology* of the theory (see Allori 2015 for an up-to-date review of this approach). Its opponents, whom I will call *reductionists*², deny the need for such postulates, because they

Email address: matthias.egg@philu.unibe.ch (Matthias Egg)

To appear in *Studies in History and Philosophy of Modern Physics*, available on <http://dx.doi.org/10.1016/j.shpsb.2016.10.007>

²I introduce this label because all the other labels that have been attached to this camp in the previous debate strike me as somewhat inaccurate. For example, the position is often dubbed “wave function realism”, but this obscures the fact that also the primitivist approach is compatible with some varieties of realism about the wave function (or more precisely, about the quantum state represented by the wave function; see Belot 2012; Egg and Esfeld 2015). The term “wave function monism” would be more accurate, but it would exclude Albert’s (1996; 2015) version of Bohmian mechanics, which is explicitly opposed to primitivism. The recently proposed term “3N-fundamentalism” (Chen forthcoming) is almost perfect, but it fails to capture Wallace’s (2012) spacetime state realism, which is another important foil for primitivism.

hold that insofar as we need local beables to account for our experience of a three-dimensional world, they are reducible to a fundamental ontology that does not include them. This fundamental ontology is given by the quantum state of the world and its temporal evolution, and the reductionists claim that certain features of this evolution manifest themselves as three-dimensional structures.

Some reductionist approaches appeal to functional roles of (parts of) the quantum state (Albert, 2015, chaps. 6 and 7), some to real patterns in its evolution (Wallace, 2012, chap. 2), some to symmetries (Ney forthcoming), but these differences do not matter here. My purpose in this paper is rather to defend the general idea of non-fundamental local beables against some criticism from primitivists, in particular the critique by Maudlin (2007). I do not advocate any particular account of how to reduce local beables to the quantum state, but I aim to show that some of the principled worries voiced by primitivists against these accounts are unfounded.

The paper is structured as follows: In section 2, I will clarify the different positions in the debate and spell out what I take to be the main challenge for reductionism about local beables, namely the need to account for the physical salience of (non-fundamental) local beables. Sections 3 and 4 then discuss two responses to this challenge, the first one based on dynamical considerations, the second one on the empirical significance of local beables. In each of these two sections, I will first show why earlier responses along the same lines were incomplete, and then provide the necessary arguments to render them convincing. This will lead to the conclusion (in section 5) that the reductionist can successfully deal with the challenge issued by the primitivist.

2. Preliminary distinctions

2.1. *Informational vs. ontological completeness*

In order to analyze the ontological status of local beables in relation to the quantum mechanical wave function, Maudlin (2007) introduces the crucial distinction between *informational* and *ontological* completeness. A description is informationally complete “if every physical fact about the situation can be recovered from the description” (3151). By contrast, “an ontologically complete description of a physical situation should provide—in a relatively transparent way—an exact representation of all of the physical entities and states that exist. . . . [It] should say just what there is and no more” (3154). As an initial illustration of this contrast, Maudlin mentions the scalar and vector potentials of classical electrodynamics, which furnish informationally complete descriptions, but are usually not taken to directly describe the physical ontology of the classical electromagnetic theory (3152).

To further clarify the relevant distinction, Maudlin shows that it can also be framed in terms familiar from other philosophical debates, namely *nomic supervenience* and *ontological reduction*³ (3152-53): Let *A* be an informationally

³Neither Maudlin nor the other participants in the debate are very explicit about the

complete part of a description B of a certain situation. Then B nomically supervenes on A in the sense that there cannot be a difference in what is described by B without a difference in what is described by A . This is to be distinguished from the claim that the referents of B are ontologically reducible to the referents of A , which is to say that A is ontologically complete.

The relation between informational and ontological completeness (or, equivalently, between nomic supervenience and ontological reduction) is a complex one, as Maudlin (2007, 3153) illustrates by giving some examples of informationally complete descriptions which differ radically with regard to ontological completeness. On one end of the spectrum are cases where informational and ontological completeness clearly go hand in hand, as when the tables in a room are described in terms of the distribution of atoms in the room: not only are all the facts about tables recoverable from a complete description of the atoms, but the tables “are nothing *over and above* the atoms” (ibid.). On the other end of the spectrum is the case of a deterministic universe, where the entire history of the world supervenes on the global physical state of the world at one particular moment (given the dynamical laws), but no one would claim that this history is nothing over and above the state at that one moment. In other words, the description of the state of the world at one particular instant is informationally but not ontologically complete. Between these two extremes, there are controversial cases such as the description of the electromagnetic field in classical electrodynamics: it is informationally complete in the sense that all the facts about the charge distribution can be recovered from a full description of the electromagnetic field (given Maxwell’s equations), and this might suggest, but does not imply ontological completeness. As Maudlin puts it: “The attempt to somehow reduce charges or charged particles to nothing but states of the field does not seem crazy, but neither does it seem inevitable” (ibid.).

Applying this distinction to versions of quantum mechanics without additional variables, we may (neglecting some subtleties mentioned in Maudlin 2007, footnote 3) assume that the description given by the wave function is informationally complete. The central question then is whether it is ontologically complete as well. Maudlin’s negative answer to that question is based on two claims: (1) a reasonable ontology for a physical theory must contain local beables, and (2) it is not clear how local beables could be ontologically reduced to (or derived from) the quantum state.

2.2. *Eliminative vs. conservative reductionism*

The three basic options regarding the ontological status of local beables are summarized in the following table:

notion of reduction they presuppose (Ney (2013) being an exception). The common practice of using “reduction” and “derivation” interchangeably (which I will also adopt in what follows) indicates that some broadly Nagelian picture constitutes the background for the debate.

	Local beables exist:	They are fundamental:
Eliminative reductionism	no	no
Conservative reductionism	yes	no
Primitivism	yes	yes

The distinction between eliminative and conservative reductionism is not always properly drawn, because there is some ambiguity in the notion of “ontology”. On a restrictive reading, the ontology of a theory only includes what is fundamental. On a more liberal reading, it includes whatever is real, be it fundamental or derivative. At some points Maudlin clearly adopts the second reading, otherwise the very notion of “derivative ontology” (discussed on p. 3161 of his paper) would be oxymoronic. But then sometimes he also sympathizes with the restrictive reading, according to which, unless one introduces local beables as elements of fundamental ontology, they “do not really exist” (3163) and “all of this talk of local beables in ordinary space is just a fiction” (3165).

One might think that this is just a verbal issue: The primitivist’s core claim is that local beables are ontologically fundamental, and both brands of reductionism deny this. Does it then really matter that one of them still calls the local beables “real” (or “existing”), while the other one views them as fictions? I think it does, but I also think that there is a legitimate way of treating eliminative and conservative reductionism together in the present context. So let me first explain why we must acknowledge the difference between these two reductionist positions, before I argue (in the next subsection) that the difference can be neglected for the rest of this paper.

Notice first that the debate about the ontological status of local beables is a debate between scientific realists, who generally believe that the entities we posit in our successful scientific explanations really do exist (although they disagree sharply about which entities are actually posited by quantum mechanical explanations). Now since most of our scientific explanations seem to involve local beables, eliminativism about them has not been a very popular position in the debate.⁴ I am not here concerned with the question whether this lack of popularity is justified or not, but I want to highlight a consequence of this situation for the debate between primitivists and reductionists: under such circumstances, primitivism can be made to look much more convincing if it is presented as an alternative to eliminative reductionism than if the foil is conservative reductionism. But this would be an illegitimate move, because it would allow the primitivist to attack reductionism by means of arguments that only work against its eliminative variant.

Unfortunately, this is precisely what happens in the following passage, where Maudlin (2007, 3166) mistakes an argument in favor of the *reality* of local be-

⁴I know of only two instances in the literature where eliminativism about local beables is advocated, namely Albert (1996) and Ney (2015, section 7). But even these authors, in other parts of their work, at least implicitly acknowledge the relative attractiveness of conserving local beables instead of eliminating them (see, in particular, Ney forthcoming, p. 1 and fn. 10).

ables for an argument in favor of their *fundamentality*:

But it is clear what would solve the problem [of accounting for observations that seem to be of particularly situated objects in a three-dimensional space]: remove all the talk of fiction! If one believes that in addition to the wavefunction *there really is* an ordinary space that *really does* contain local beables that *really do* evolve in a specified way determined by the wavefunction, then you have something.

The (conservative) reductionist can agree with everything in this passage except the phrase “in addition to the wavefunction”. That phrase, she will (correctly, to my mind) insist, expresses an ontological commitment which is in no way warranted by the explanatory success arising from a commitment to the reality of local beables. The clearest way to see this is to consider a parallel case based on one of the examples mentioned above: I have a visual experience of a table standing in the corner of my room, and I explain this by saying that *there really is* a table in the corner of my room. Does that somehow commit me to the belief that the table is an ontologically irreducible object? Of course not. The explanation succeeds as long as the table is real, regardless of whether it is something over and above a collection of atoms.

2.3. The core challenge to reductionism

At the end of subsection 2.1, I mentioned two claims on which primitivism about local beables relies. By asserting the reality of local beables, agreement to claim (1) unites primitivism with conservative reductionism (against eliminative reductionism). By contrast, claim (2) marks the distinction between primitivism and reductionism (of either variant) by questioning the reduction of local beables to the quantum state. I will now explain why I take the disagreement about (2) to be much more substantial, which is why the rest of this paper will focus on the debate between primitivism and reductionism, disregarding the disagreement within reductionism.

First, while eliminativism’s above-mentioned lack of popularity depends on a particular conception of scientific realism, I have argued elsewhere that this is indeed a reasonable stance to take in the scientific realism debate (Egg 2014, 14-16; see also Psillos 2005). Accordingly, there are good philosophical reasons for the sociological fact that eliminativism has few advocates and that conservative reductionism first and foremost must be defended against primitivism, not against eliminative reductionism.

Furthermore, one motivation for entertaining eliminative reductionism comes from the conviction that conservative reductionism about local beables is not viable. By responding to what I take to be the most carefully worked out attack on this latter kind of reductionism, the following sections seek to go at least some way towards removing that motivation.⁵

⁵ Other (somewhat less well-developed) arguments against conservative reductionism can

Finally, although the argument to be discussed here is primarily directed against conservative reductionism, it can easily be transformed into an argument against eliminative reductionism and should therefore properly be regarded as an argument against reductionism *tout court*. Here is how the argument seeks to support claim (2):

The notion is that the dynamics of a very high-dimensional object in a high-dimensional space could somehow implicitly contain within it—as a purely *analytical* consequence—a description of local beables in a common low-dimensional space. This approach turns critically on what such a derivation of something isomorphic to local structure would look like, where the derived structure deserves to be regarded as *physically* salient (rather than merely mathematically definable). Until we know how to identify physically serious derivative structure, it is not clear how to implement this strategy. (Maudlin, 2007, 3161)

The worry expressed in this quote is that there is no obvious connection between the configuration space in which the quantum state evolves and the three-dimensional space of local beables. While it is possible to mathematically define such a correspondence, one would have to show that this particular correspondence is “physically salient”, distinguishing it from the plethora of other mathematically definable correspondences between the wave function and low-dimensional structures (most of them very bizarre and utterly different from what we experience). The fact that this particular correspondence strikes us as natural and simple may be seen as just an epistemic observation without any ontological import (Monton, 2002, 269).

This is certainly a problem for the conservative reductionist, who claims that the quantum state by itself somehow gives rise to local beables. Why then does it not also give rise to all these other bizarre entities? Or if it does, why do we not experience them? Let us dwell on this second question for a moment, as it allows us to see how the worry also affects eliminative reductionism. The eliminativist, of course, is under no obligation to look for a feature that marks out local beables as *real* among this host of merely mathematically “derivable” structures. But he nonetheless needs to account for the fact that our experiences are *as if* there were local beables. And if the quantum state is ontologically sufficient to ground these experiences (as it must be, since there is nothing else, according to the eliminativist), the very same question as above arises: why does the quantum state not also give rise to all these other bizarrely structured experiences? Hence, the challenge issued by the primitivist faces both types of reductionism with equal strength. Whether one seeks to reduce *real* local beables or only the *appearance* thereof to the quantum state, in either case one needs to demonstrate the *salience* of the derived structure.

be found in Lewis (2013), Monton (2013) and Ney (2015). I plan to address these in another paper.

3. Shapes vs. dynamics as the vehicle for reduction

Reductionists like Albert and Wallace address the foregoing objection by appealing to the dynamical behavior of the quantum state. They admit that mathematically, the quantum state could correspond to many different low-dimensional structures, but they point out that the dynamics of the quantum state singles out the three-dimensional structure of local beables. This is reflected in the fact that the Hamiltonian operator describing the evolution of a quantum state formally corresponds to a classical Hamiltonian describing a system of particles in three-dimensional space (Albert, 1996).

Primitivists are, however, not convinced by this response. They may acknowledge that the appeal to dynamics could, in principle, be used to reduce local beables to the quantum state, but they maintain that this kind of reduction falls short of the transparency and simplicity we find in other reduction relations, for example in the reduction of macroscopic bodies to collections of particles. Maudlin (2014, 798) explicitly emphasizes this contrast in his discussion of Wallace’s (2012) proposal to identify physical objects (such as cats) as dynamical patterns in the development of the quantum state:

Whether any of these individual [quantum] states contains a “cat-shaped object” cannot be determined from examination of the state alone. (Contrast this with a theory that postulates particles in space: the disposition of particles at a given time can unproblematically be described as “cat-shaped” without reference to any dynamics.) Rather, to determine the physical content of any one state, one has to refer back to all the other states, and how they are connected by law. Reflection on all of the dynamical connections is somehow supposed to reveal whether a certain situation contains a cat chasing a mouse, even though the individual states, in themselves, have nothing cat-shaped or mouse-shaped in them. How this bootstrapping could exactly occur I can’t quite tell.

The reductionist could, of course, simply insist that reductions referring to dynamics, though more complicated, are still legitimate, because this complication is more than compensated by the avoidance of postulating local beables as additional primitives. The debate would thereby threaten to degenerate into a mere exchange of expressions of ontological taste. What I want to show now, however, is that the reductionist can do better, by arguing that even the instances of reduction favored by the primitivist ultimately cannot do without dynamical considerations.

In fact, a first set of dynamical considerations is already needed before the discussion just mentioned can even start. In the above quotation, Maudlin takes for granted the physical salience of macroscopic bodies like cats (as opposed to arbitrary other collections of particles). What accounts for this salience is not the mere *shape* of these collections, but their *dynamical stability*. There is, after all, nothing physically salient about a cat-shaped collection of air molecules.

But there is also a more subtle way in which dynamical considerations are likely to be necessary even for the (supposedly) unproblematic cases of reduction Maudlin cites. The success of identifying macroscopic objects as particularly shaped collections of particles without reference to any dynamics depends critically on the physical theory describing the particles. It is well known that contemporary physics does not straightforwardly support the idea of particles as fundamental building blocks of material bodies (see, e.g., Ladyman and Ross 2007; Healey 2013), so what is needed to hold on to that metaphysical picture is some non-standard quantum theory of matter, along the lines of Bohmian mechanics. What makes Bohmian mechanics an appealing candidate for spelling out the physical details of how atoms form a table is that it can account for all the surprising quantum effects in the behavior of particles (entanglement, non-locality, spin etc.), but treats them as purely dynamical effects, while the underlying ontology is a classically atomistic one of particles in ordinary space.

However, everybody agrees that Bohmian mechanics (just like standard quantum mechanics) is an incomplete theory, not only because it is non-relativistic, but (more importantly for our purpose) because it treats interactions (insofar as it treats them at all) in a semi-classical way, so it is certainly not the final word on how particles form macroscopic bodies. There are various directions of research towards more complete Bohmian theories of matter seeking to recover the empirical predictions of ordinary quantum field theory (QFT) (see Struyve 2011 for a review). Some of them abandon the particle ontology of Bohmian mechanics and turn to a field ontology, thereby undermining the idea that macroscopic bodies are made of elementary particles. But even the approaches that retain a particle ontology render the task of identifying macroscopic objects without reference to any dynamics more difficult than the above quotation by Maudlin suggests.

There are essentially two particle-based approaches to QFT, and both of them seem to populate space with many more particles than the electrons and quarks of which we think as matter constituents. First, there is the Dirac sea pilot-wave model (Colin and Struyve 2007; Deckert et al. forthcoming), which postulates a large (but finite, due to a high-momentum cut-off) number of fermions occupying negative energy states. Second, according to the Bohmian QFT proposed by Dürr et al. (2004), particle creations and annihilations are viewed as actual processes in space and time. Within such a theory, vacuum fluctuations correspond to real particle-anti-particle pairs appearing and disappearing in space and time.

I do not know of any precise calculation of the density of these additional particles, but Struyve (2011, 8) estimates that it is of the same order for both types of theories, namely of order Λ^3 , where Λ is the inverse Planck length. If anything like this is true, then the quarks and electrons which constitute macroscopic bodies make up only an extremely small fraction of the complete particle content of any given region of space. And since there is no intrinsic difference between the fermions that constitute matter and the background fermions (either belonging to the Dirac sea or to the vacuum fluctuations), there seems to be no way of discerning a cat-shaped arrangement in this abundance of particles

unless one takes into account the behavior of the particles *in time*.⁶

These are admittedly speculative considerations, so one might think that they do not constitute a definitive argument for the inevitability of dynamical considerations in reductions. I agree that future theoretical developments may yet again create space for a non-dynamical reduction of macroscopic bodies to collections of particles, but the situation sketched here makes this a rather unlikely scenario. It is not just that, despite their best efforts, researchers in Bohmian QFT have not yet come up with a theory that would support the idea of such a reduction, but that even the research programs which offer the most hope for generating such a theory have turned out to render this idea problematic. Therefore, it is fair to say that neither current physics nor any reasonable expectations about future physics do support Maudlin's misgivings about using dynamics in ontological reductions.

4. Underdetermination of (physically salient) derivative ontologies

There is in the literature a surprisingly simple second response to the primitivist's question of what distinguishes local beables from all the other mathematically definable structures which have nothing to do with our experience: local beables are physically salient *precisely because* they do explain our experiences and therefore strike us as natural, in contrast to all the bizarre, contrived and explanatorily irrelevant structures one could define in addition to them. Huggett and Wüthrich (2013, 283) emphasize this point (in direct response to Maudlin (2007)):

There are two ways of looking at the idea of physical salience: 'from below', as Maudlin seems to, taking for granted that the theory has physical salience, and asking what formal derivations of higher-level structures preserve it. But one can equally legitimately consider the idea 'from above', taking for granted that the empirical realm has physical salience (which it certainly does), and asking how it can 'flow down' formal derivations to give physical significance to the underlying theory.

Assessing physical salience in that manner obviously provides a way of distinguishing local beables from all kinds of bizarre structures that could in principle

⁶Colin and Struyve (2007) argue that on a coarse-grained level, their model allows for a distinction between macroscopic states of sufficiently distinct fermion density (they give the example of a spherical region filled with either graphite or air, and they calculate that the distinction succeeds if the radius b of the region satisfies $b \gg 2.6 \times 10^{-6}$ m). But surely any satisfactory account of macroscopic bodies should also work for objects surrounded by matter of similar density, for example fish in water. Colin's and Struyve's account is not fine-grained enough for this, because it is unable to discern fermion arrangements corresponding to different types of atoms, which would be necessary to (non-dynamically) distinguish a fish from the water surrounding it.

just as well be “derived” from the wave function. However, a look at the main example used by Maudlin (2007, 3161) to illustrate the worry of surplus derivative structure shows that this response is, by itself, insufficient:

Until we know how to identify physically serious derivative structure, it is not clear how to implement this strategy [of deriving local beables from the wave function]. As we will shortly see, for example, the same GRW dynamics can be supplemented by *different* local beables: a mass density ontology and a flash ontology. Presumably, an account that treats these local beables as derived would determine whether one or the other (or neither) of these is the derivative ontology that is ‘really’ implicit in the GRW dynamics.

The case of GRW is a well-known example of a theory that comes in ontologically different versions sharing the same mathematical formalism (see Allori et al. 2008). The crucial point is that the two versions cited by Maudlin (usually called GRWm and GRWf, respectively) also share their empirical predictions, which immediately implies that the strategy of choosing between them “from above” will not work. This is a familiar problem of underdetermination, but it is here combined with a less familiar one: not only is the choice between GRWm and GRWf underdetermined by the empirical evidence, but it is also underdetermined by the underlying fundamental theory (GRW without local beables, called GRW0).

There are two principled ways to deal with problems of underdetermination: one can either try to look for a sense in which the seemingly contradicting theories are compatible after all, or one can accept their incompatibility and argue that one of the theoretical options (although we may never know which one) is in fact the right one. Since I do not want my defense of reductionism to depend on any one of these two methodologies, I will give two responses to the problem of underdetermination in GRW, each one attuned to one of the two methodologies. Taken together, they will show that, regardless of the more general stance one takes with respect to underdetermination, the GRW case does not present a problem for reductionism about local beables.

4.1. *Dissolving underdetermination by seeking compatibility*

The idea that different derivative ontologies can be compatible with each other has previously been defended by Ney (2015, 3115), but it is unclear whether her argument can be applied to the case of the different GRW ontologies. To illustrate her proposal, Ney gives the example of a fundamental ontology in a nine-dimensional space, giving rise to two different derivative ontologies in a three-dimensional and a six-dimensional subspace, respectively. Unfortunately, this example is significantly dissimilar to the GRW case, not only because our three-dimensional space is *not* a subspace of the configuration space in which the quantum state evolves,⁷ but also because, as we will see, there is

⁷Ney (2015, fn. 18) herself draws attention to this fact.

some strong *prima facie* evidence that the ontologies of GRWm and GRWf are actually *not* compatible with each other. I shall now attempt to overcome these problems.

The first obstacle to the compatibility of GRWm and GRWf is that the space-time pictures of physical reality painted by these two theories differ radically, as Maudlin (2007, 3167) points out:

If all one could see were space time and the local beables (but if one could see them in all details), the fine-grained pictures provided by the theories would have nothing in common. In Ghirardi's [GRWm] world, one sees an evolving continuous mass density that would usually change continuously but would sometime undergo discontinuous changes that result in the sudden clumping of the mass in small regions of space. In Bell's [GRWf] world, one would see almost nothing at all: just a relatively sparse distribution of flashes, one for each of the sudden jumps in the alternative theory.

How can I claim that these two ontologies are compatible with each other, when for almost all space-time points, one of them affirms and the other one denies that there is something at that point? As a first step towards an answer, notice that Maudlin's assertion about the space-time pictures of the two theories having "nothing in common" is plainly false, because the most striking features of the GRWm world—namely, the spontaneous clumpings of matter in small regions of space—occur precisely where the flashes are in the GRWf picture. The two ontologies therefore agree exactly about the set F of all space-time points at which a local maximum of matter density occurs. This commonality allows us to apply a strategy familiar from other debates on underdetermination, suggesting that we should be committed to only those features which the different ontological options have in common (French, 2011). In the present case, this is to assert that there definitely is matter at the points in F , leaving it open whether the rest of space-time $M \setminus F$ is empty (as GRWf claims) or not (as in GRWm). This latter question, after all, does not make a difference for the empirical content of the theory.

Not everybody will be satisfied with such a partial commitment, however. Some will insist that a serious ontological position must not simply duck the question of what there is outside of F . Before addressing that objection, let me briefly discuss two more worries raised by Maudlin (2007, 3167-8) against the view that GRWm and GRWf are compatible with each other.

First, he observes that "the research programmes that naturally arise from the alternative pictures are in some ways diametrically opposed": whereas GRWm suggests looking for a *continuous* spontaneous localization (CSL) dynamics instead of the discontinuous GRW collapses, this "is the *last* thing that one would pursue" in GRWf, because eliminating the GRW jumps would eliminate the local ontology altogether.

It seems to me that this view begs the question against the reductionist by treating GRWm and GRWf as fundamental theories (relative to the current

state of research), because otherwise they would not be regarded as helpful guides towards a more fundamental theory. But from the reductionist point of view, GRW0 is the fundamental theory which should guide further research, and as it turns out, the considerations motivating the CSL model (e.g., Bassi and Ghirardi 2003, 312) are indeed just as much part of GRW0 as of GRWm. Also the publication history suggests that it was not GRWm that motivated Ghirardi and his collaborators to develop CSL, but that they developed CSL on the basis of GRW0 (Ghirardi et al., 1990) and only later proposed the matter density ontology as a supplement for both the GRW and the CSL model (Ghirardi et al., 1995).

The second feature Maudlin adduces to emphasize the contrast between GRWm and GRWf is that the latter has a far better chance of being made compatible with special relativity than the former. As with the previous objection, the force of this argument is diminished by the fact that, in a framework within which space-time is regarded as derivative, issues of Lorentz invariance can hardly claim fundamental importance. Nevertheless, the question how local beables (even derivative ones) behave under Lorentz transformations is of some relevance. At the time Maudlin’s paper appeared (2007), it looked indeed as if there was a decisive advantage for GRWf as regards Lorentz invariance: Tumulka (2006) had just published his relativistic version of GRWf, while the available definitions of the matter density in GRWm depended on integration over a spacelike hypersurface, which seemed to rule out Lorentz invariant GRWm models from the start. In the meantime, this situation has significantly changed. We now not only have a relativistic GRWm model (Bedingham et al., 2014), but also an argument indicating that Tumulka’s GRWf model is Lorentz-invariant only in a somewhat restricted sense (Esfeld and Gisin, 2014, section 3). Lorentz invariance should therefore no longer be regarded as a crucial difference between GRWm and GRWf.

4.2. *Breaking the underdetermination*

The above attempt to establish compatibility between GRWm and GRWf will not have convinced those who think that there ought to be a fact of the matter as to whether there is something or nothing at the points in $M \setminus F$. The two ontologies definitely disagree about this issue, so at most⁸ one of them can be right. Therefore, under the assumption that the matter content of $M \setminus F$ is a substantial issue, the underdetermination between GRWm and GRWf cannot be dissolved, but must be broken by opting for one of the two ontologies.⁹

⁸In the following, I will disregard the possibility that both ontologies are false. In the absence of an empirically adequate proposal for yet another GRW ontology (and keeping in mind the remarks about eliminativism in section 2), GRWm and GRWf can be treated as exhausting the space of possibilities.

⁹There is also an intermediate position, which acknowledges the incompatibility of the two ontologies but still refrains from choosing one of them, remaining agnostic instead. This position does not need to be treated separately, because it differs from the one advocating a choice between the two options only on the epistemic level. On the ontic level, both these

At first sight, this situation does not seem to create a specific problem for reductionism about local beables. After all, the need to select one of the two ontologies arises for the primitivist (who thinks of GRWm and GRWf in terms of fundamental ontology) just as much as for the reductionist (who views GRWm and GRWf as alternative ways to derive local beables from GRW0). Therefore, one might say, insofar as underdetermination does not prevent us from realism about fundamental ontology, it should not prevent us from realism about derivative ontology either.

There is something unsatisfactory about this line of argument, however. While it is true that both the primitivist and the reductionist have to react to the problem of underdetermination by invoking ontological choices, the primitivist account now looks more natural, because its choice concerns only the fundamental level. Once this is settled, everything else follows. By contrast, the reductionist first needs to specify her fundamental ontology (the quantum state) and then, in the course of deriving local beables from that basis, needs to invoke a second choice, investing one particular set of local beables with ontological significance.

This second choice may raise a number of critical questions. First, one might ask how there can be a choice at all, if we assume that there are principles which, once the fundamental ontology is fixed, uniquely determine the derivative ontology as well.¹⁰ My reply is that there is a choice in the same sense as in the case of the primitivist's choice between GRWm and GRWf. In both cases, there are two mutually incompatible (but empirically equivalent) ways the world could be. The difference between the two cases is that in one case, the two ways differ on the level of fundamental ontology, while in the other case, they differ in the principles according to which the fundamental ontology gives rise to local beables. The reductionist may be ignorant of these principles, but in the same way, the primitivist is ignorant of the fundamental ontology, so they both need to make an educated guess about whether we live in a GRWm world or in a GRWf world. Furthermore, the reductionist's guess is regulated by the very same principles of metaphysical theory choice as the primitivist's. Let me briefly discuss the two principles that matter most in the present context.

On the one hand, a principle of *parsimony* speaks in favor of GRWf. As we have seen above, locating matter at the points in F suffices to account for the empirical content of the GRW theory. The most parsimonious hypothesis about the set $M \setminus F$ is then simply that it does not contain any matter at all. The resulting picture is precisely the one of GRWf.

On the other hand, there is a well-entrenched metaphysical principle *forbidding action at a distance*. This clashes with the GRWf ontology, as can be seen by considering the simplest possible particle physics experiment: a source

positions agree that there is one correct ontology, otherwise there would be nothing for the agnostic to be agnostic about.

¹⁰I owe this suggestion to an anonymous referee, who calls these principles "principles of emergence". I avoid this terminology, because emergence is sometimes understood as being incompatible with reduction.

emits a particle, which is subsequently recorded by a detector. According to GRWf, both the source and the detector consist of a multitude of flashes, but there is nothing that travels from one to the other, hence nothing to mediate the interaction between them (see Esfeld and Gisin (2014, section 5); Esfeld (2014a) makes the same argument, but erroneously takes it to threaten the *empirical adequacy* of GRWf, rather than just its conformity to a metaphysical principle). This principle thus favors GRWm, in which a matter field mediates the interaction between the source and the detector.

Obviously, invoking such metaphysical principles does not straightforwardly resolve the underdetermination, because metaphysical principles may pull in opposite directions (as is the case here). Furthermore, people disagree about the soundness of metaphysical principles and about the relative weight these principles should be given. But again, this problem confronts the primitivist and the reductionist alike.

As noted above, an advantage for the primitivist seems to consist in the fact that he can confine his ontological choice to the fundamental level, while the reductionist's choice between GRWm and GRWf takes place on the derivative level, in addition to her specification of the fundamental ontology. However, it is not the case that the primitivist can do without any ontological choices on the derivative level. For example, even if one opts for a primitive ontology of particles, this does not settle the further question as to which collections of particles constitute real objects and which ones do not. Just as in the case of GRWf and GRWm, considerations of parsimony may motivate antirealism about composite objects, whereas explanatory considerations speak in favor of recognizing some of them as real. This means that, even taking into account the underdetermination between GRWf and GRWm, the reductionist's methodology is no less respectable than the primitivist's: after having specified their respective fundamental ontologies, they both need to make additional ontological choices on the derivative level, based on the ontological principles they value.

As a final objection, one might question the analogy between the GRWm/GRWf choice and the kind of ontological disputes mentioned in the previous paragraph. The temptation to regard a dispute as unsubstantial (and accordingly to believe that one can do without choosing) seems stronger when one is asked whether some collection of particles “really” constitute an object than when one considers the matter content of $M \setminus F$. However, I do not see any robust justification for this difference of attitude. First, let us not be fooled by Maudlin's above-quoted counterfactual story about “seeing local beables in all details” into believing that the difference between GRWm and GRWf is somehow “observable in principle”. Given the way the GRW dynamics works, this difference is just as much a piece of unobservable metaphysics as the difference between a real cat and some cat-wise arranged particles. More importantly, suppose one tries to dissolve the dispute between realism and antirealism about cats by suggesting that the disagreement about the reality of cats is really just a pragmatic issue, whereas ultimately both parties agree that there are particles which sometimes arrange themselves cat-wise. This appeal to the fundamental level invites a similar story to be told in the GRW case, understood as a disagreement about derivative on-

tology: here as well, both parties ultimately agree about what there is, namely the quantum state. At that level, everything that constitutes their disagreement (not only the matter content of $M \setminus F$, but the space-time M itself) disappears. In sum, the underdetermination between GRWm and GRWf is really on a par with other cases of underdetermined derivative ontology and therefore need not worry the reductionist.

5. Conclusion

The attempt to reduce local beables to the quantum state faces the challenge of demonstrating the physical salience of the reduction relation between the quantum state and the local beables. I have looked at two ways in which the reductionist literature has responded to that challenge and found them both to be incomplete as they stand. I have then performed the necessary argumentative work to render these responses successful. As a result, the threat to the physical salience of derived local beables is now completely neutralized.

Of course, this only establishes that reductionism about local beables is *viable* (insofar as one accepts my promissory note in footnote 5 above), not that it is *preferable* to primitivism. Considerations of parsimony in fundamental ontology may motivate a preference for reductionism over a position that postulates primitive local beables alongside the quantum state, but this cannot be the whole story. First, ontological parsimony is not all that matters, and second, some recent versions of primitivism score quite well on parsimony by taking the opposite direction of the road discussed in this paper: instead of reducing local beables to the quantum state, they attempt to reduce the quantum state to the local beables (Esfeld, 2014b; Callender, 2015). Evaluating the advantages and disadvantages of reductionism (as defended in the present paper) compared to these other proposals is a project for another day.

Acknowledgments

I would like to thank Claus Beisbart, Michael Esfeld, Ward Struyve and two anonymous referees for helpful comments on earlier versions of this paper. I am also indebted to audiences in Manchester (BSPS 2015), Düsseldorf (EPSA 2015), Paris, Bern and Leeds for inspiring discussions.

Albert, D. Z., 1996. Elementary quantum metaphysics. In: Cushing, J. T., Fine, A., Goldstein, S. (Eds.), *Bohmian Mechanics and Quantum Theory: An Appraisal*. Vol. 184 of Boston Studies in the Philosophy of Science. Kluwer Academic Publishers, Dordrecht, pp. 277–284.

Albert, D. Z., 2015. *After Physics*. Harvard University Press, Cambridge MA.

Allori, V., 2015. Primitive ontology in a nutshell. *International Journal of Quantum Foundations* 1, 107–122.

- Allori, V., Goldstein, S., Tumulka, R., Zanghì, N., 2008. On the common structure of Bohmian mechanics and the Ghirardi-Rimini-Weber theory. *British Journal for the Philosophy of Science* 59, 353–389.
- Bassi, A., Ghirardi, G., 2003. Dynamical reduction models. *Physics Reports* 379, 257–426.
- Bedingham, D., Dürr, D., Ghirardi, G., Goldstein, S., Tumulka, R., Zanghì, N., 2014. Matter density and relativistic models of wave function collapse. *Journal of Statistical Physics* 154, 623–631.
- Belot, G., 2012. Quantum states for primitive ontologists. *European Journal for Philosophy of Science* 2, 67–83.
- Callender, C., 2015. One world, one beable. *Synthese* 192, 3153–3177.
- Chen, E. K., forthcoming. Our fundamental physical space: An essay on the metaphysics of the wave function, preprint, submitted January 15, 2016. <http://philsci-archive.pitt.edu/11863/>.
- Colin, S., Struyve, W., 2007. A Dirac sea pilot-wave model for quantum field theory. *Journal of Physics A: Mathematical and Theoretical* 40 (26), 7309.
- Deckert, D.-A., Esfeld, M., Oldofredi, A., forthcoming. A persistent particle ontology for QFT in terms of the Dirac sea, forthcoming in the *British Journal for the Philosophy of Science*. <http://philsci-archive.pitt.edu/12375/>.
- Dürr, D., Goldstein, S., Tumulka, R., Zanghì, N., 2004. Bohmian mechanics and quantum field theory. *Physical Review Letters* 93, 090402.
- Egg, M., 2014. *Scientific Realism in Particle Physics: A Causal Approach*. De Gruyter, Boston.
- Egg, M., Esfeld, M., 2015. Primitive ontology and quantum state in the GRW matter density theory. *Synthese* 192, 3229–3245.
- Esfeld, M., 2014a. The primitive ontology of quantum physics: Guidelines for an assessment of the proposals. *Studies in History and Philosophy of Modern Physics* 47, 99–106.
- Esfeld, M., 2014b. Quantum Humeanism, or: Physicalism without properties. *The Philosophical Quarterly* 64, 453–470.
- Esfeld, M., Gisin, N., 2014. The GRW flash theory: A relativistic quantum ontology of matter in space-time? *Philosophy of Science* 81, 248–264.
- French, S., 2011. Metaphysical underdetermination: Why worry? *Synthese* 180, 205–221.
- Ghirardi, G., Pearle, P., Rimini, A., 1990. Markov processes in Hilbert space and continuous spontaneous localization of systems of identical particles. *Physical Review A* 42, 78–89.

- Ghirardi, G. C., Grassi, R., Benatti, F., 1995. Describing the macroscopic world: Closing the circle within the dynamical reduction program. *Foundations of Physics* 25, 5–38.
- Healey, R., 2013. Physical composition. *Studies in History and Philosophy of Modern Physics* 44, 48–62.
- Huggett, N., Wüthrich, C., 2013. Emergent spacetime and empirical (in)coherence. *Studies in History and Philosophy of Modern Physics* 44, 276–285.
- Ladyman, J., Ross, D., 2007. *Every Thing Must Go: Metaphysics Naturalized*. Oxford University Press, Oxford.
- Lewis, P. J., 2013. Dimension and illusion. In: Ney and Albert (2013), pp. 110–125.
- Maudlin, T., 2007. Completeness, supervenience and ontology. *Journal of Physics A: Mathematical and Theoretical* 40, 3151–3171.
- Maudlin, T., 2014. Critical study: David Wallace, *The Emergent Multiverse*. *Noûs* 48, 794–808.
- Monton, B., 2002. Wave function ontology. *Synthese* 130, 265–277.
- Monton, B., 2013. Against 3n-dimensional space. In: Ney and Albert (2013), pp. 154–167.
- Ney, A., 2013. Ontological reduction and the wave function ontology. In: Ney and Albert (2013), pp. 168–183.
- Ney, A., 2015. Fundamental physical ontologies and the constraint of empirical coherence: A defense of wave function realism. *Synthese* 192, 3105–3124.
- Ney, A., forthcoming. Finding the world in the wave function: Some strategies for solving the macro-object problem, preprint, submitted September 19, 2016. <http://philpapers.org/archive/NEYFTW.pdf>.
- Ney, A., Albert, D. Z., 2013. *The Wave Function: Essays in the Metaphysics of Quantum Mechanics*. Oxford University Press, Oxford.
- Psillos, S., 2005. Scientific realism and metaphysics. *Ratio* 18, 385–404.
- Struyve, W., 2011. Pilot-wave approaches to quantum field theory. *Journal of Physics: Conference Series* 306, 012047.
- Tumulka, R., 2006. A relativistic version of the Ghirardi-Rimini-Weber model. *Journal of Statistical Physics* 125, 821–840.
- Wallace, D., 2012. *The Emergent Multiverse: Quantum Theory according to the Everett Interpretation*. Oxford University Press, Oxford.