Primitive Ontology or Primitive Relations?

Quentin Ruyant

November 23, 2015

Abstract

Primitive ontology is a program which seeks to make explicit the ontological commitments of physical theories in terms of a distribution of matter in ordinary space-time. This program targets wave-function realism, which interprets the high-dimensional configuration space on which wave-functions are defined as our fundamental physical space. Wave-function realism allegedly fails to account for a correspondence between the ontology it postulates and the 'manifest image' of the world in which experimental tests of the theory are performed, and therefore the wave-function must be completed with an additional structure which describes what fundamentally exists in ordinary space-time: the 'primitive ontology'. However primitive ontologies face some difficulties, in particular concerning the ontological status of the wave-function. In this paper, I defend a realist interpretation of the wave-function as describing objects in ordinary space-time, which does not require supplementing the theory with additional structure. The main difference between this proposal and other primitive ontology proposals is that the fundamental constituents of reality are not conceived of as localised objects, but as primitive relations. This interpretation purports to share the advantages of both wave-function realism and primitive ontologies without facing the same difficulties. I argue that the need for an additional structure stems from a commitment to the locality of the fundamental constituents of reality, and that such commitment is unnecessary for recovering the manifest image of the world.

Primitive ontology has been proposed as a requirement for proper realist interpretations of physical theories. The requirement goes along these lines: physical theories should tell us precisely what exists in the world, in terms of distribution of matter in space-time, or 'local beables' (Allori et al. 2008, Bell 2004). It is claimed that such commitments are necessary to make intelligible the correspondence between theory and experience, or between the 'scientific image' and the 'manifest image' of the world.

Experimental tests of our physical theories are usually performed in laboratories, with the help of measuring apparatus that we perceive as ordinary objects with determinate, macroscopic properties localised in our familiar $3 + 1$ dimensional space-time (or our 3 dimensional space evolving with time). These measuring apparatus are part of the 'manifest image of the world'. It is reasonable to assume that a realist interpretation of a physical theory should make explicit how this 'manifest image' is recovered from the theory, in pain of undermining its own experimental basis. One should be able to map the real entities posited by the theory to ordinary objects in space-time, and this requires that the interpretation of the theory provides a primitive ontology of fundamental objects located in space-time, to which properties might be attributed, and of which ordinary objects, including measuring apparatus, are constituted. Arguably, a theory that does not fulfil this standard could be deemed incomplete.

Newtonian mechanics or classical electromagnetism are quite easy to interpret as theories that attribute properties to localised objects: physical models are typically viewed as attributing charges, masses and velocities to particles, and the electromagnetic field can be viewed as attributing more complex properties, represented by vectors, to the space-time points themselves. Special and general relativity also fulfil this standard, even though the geometry of space-time involved is different from that of classical physics (and, in the case of general relativity, not independent of the physical properties themselves). Quantum theory, however, does not come equipped with such a primitive ontology in ordinary space-time that could be directly read off the models. According to standard quantum mechanics, the wave-function of the universe is not a field on our familiar 3 dimensional space, but on an abstract, $3N$ dimensional configuration space, where N is the number of particles in the universe. (This aspect is directly related to the notion of entanglement: a system of N particles is not generally reducible to the combination of N independent systems.)

From a realist perspective, what then could be the ontology of quantum theory? Two options are usually considered in the literature. The first one, dubbed 'wave-function realism', consists in biting the bullet and assume that the fundamental objects or properties quantum theory is about do not live in a 3 dimensional physical space, but in the $3N$ dimensional configuration space. Our familiar space would somehow emerge from this high dimensional space and from the dynamical aspects of the theory. However many authors doubt the viability of this solution. They suggest a second option, which consists in completing the formalism of the theory with additional structures in a 3 dimensional space that would play the role of the primitive ontology, i.e. of what exists in space-time at the fundamental level, and that would therefore fill the gap between the theory and the manifest image of the world. Proposals so far include matter density fields, particles or flashes.

The aim of this paper is to evaluate the primitive ontology program. I will argue that even though some of the worries raised by wave-function realism are well founded, one needs not complement the wave-function with additional structures to fulfil the desirata of primitive ontology proponents. One only needs to properly interpret the wave-function as an object living in ordinary space-time. As a demonstration of this claim, I will propose such an interpretation and show how the manifest image of the world can be recovered from it.

In section 1, the concept of primitive ontology and its motivations are introduced. Its difficulties are briefly reviewed. Section 2 is devoted to a realist interpretation of the wave-function in ordinary space-time, which demonstrates that no additional structure is required to fulfil the desirata of the primitive ontology program. Section 3 highlights the advantages of this alternative over other primitive ontology proposals. Possible objections are addressed. In conclusion, I will argue that the difficulties of the primitive ontology program and of wave-function realism stem from a prejudice against holism.

1 The Primitive Ontology Program

In standard, non-relativistic quantum mechanics, the wave-function of a system can be represented either as a vector in Hilbert space or as a complex field on a configuration space of $3N$ dimensions, where N is the number of particles of the system. Both representations are on a par, but the latter is viewed by wave-function realists as ontologically more significant (Albert 1996) in particular because it pictures the world as constituted of bearers (the configuration space points) of properties (the complex numbers), rather than just a single vector in a structured space.

The configuration space is a concept which finds its origin in classical physics, where each of its dimensions represents a degree of freedom of the system. If a composite system has N particles and each particle has 3 degrees of freedom (for example the three spatial coordinates of its position), then the whole system has $3N$ degrees of freedom. The state of a composite system can be represented either as a single point or as a probability distribution on configuration space (see figures 1 and 2 for illustrations of the configuration space of a composite system of two particles with only one degree of freedom each, e.g. their position along an axis). The high number of dimensions of the configuration space is simply a combinatorial effect due to the high number of particles in the system, and this space is generally interpreted as a mathematical tool rather than as a physical space. In any case one is free to return to a more intuitive representation of the system in ordinary space-time.

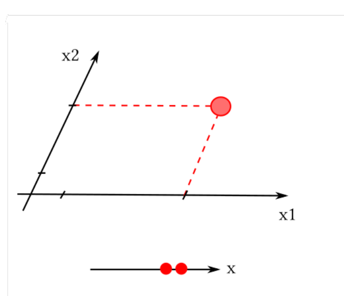

Figure 1: Illustration of the state of a composite system of two particles with one degree of freedom each

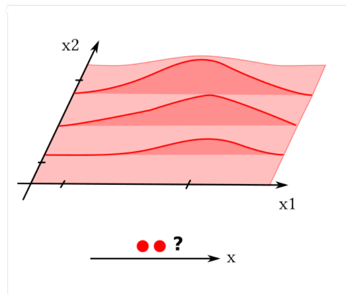

Figure 2: Illustration of a probability distribution for a composite system of two particles with one degree of freedom each

In quantum mechanics, systems are represented by a complex valued field on a configuration space: the wave-function. However in this case, it is often assumed that interpreting the configuration space as a mere mathematical tool is not possible, because the wave-function of a composite system is not generally reducible to a set of independent fields on ordinary space (in particular when the system is entangled. See figure 3).

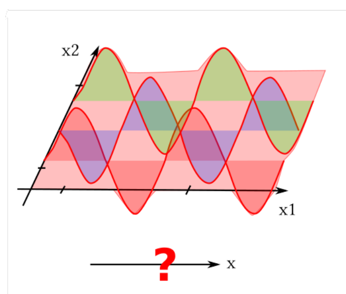

Figure 3: Illustration of a wave-function for a composite system of two particles with one degree of freedom each

If the wave-function represents all that exists, then, according to wave-function realists, our real physical space must be the configuration space. Without particles or other objects in ordinary space, the notion of configuration is no more intelligible: there is nothing to be configured, so the $3N$ dimensional configuration space must be the fundamental physical space we live in and ordinary space must be either emergent or illusory.

However wave-function realists must tell us how exactly objects in a 3-dimensional space, or the appearence thereof, would arise from the arrangement of objects in a $3N$ dimensional space. Albert proposes that the dynamic explains the appearence of familiar objects, but the prospects of reducing the three dimensionality of these objects to dynamical aspects rather than to the fact that their fundamental constituents already live in ordinary space-time is still controversial. Some authors argue that we have no theoretical resources at our disposal to map the dimensions of the configuration space onto three identifiable spatial dimensions in a non-ad-hoc way, and that something is missing in our ontology to recover our familiar $3D$ objects (see Monton (2002), Lewis (2004), Monton (2006), Ney & Albert (2013)).

This argument motivates the primitive ontology program, which seeks to complement standard quantum mechanics with an additional structure which would represent what 'really exists' in ordinary space-time at the fundamental level. The program includes three mature proposals: Bohmian mechanics, GRWm and GRWf, where the primitive ontology is constituted of respectively particles, matter density fields and bare instantaneous events, a.k.a. 'flashes'.

Note that the requirement for our theories to provide an ontology in ordinary space-time is distinct from the measurement problem in quantum mechanics. The latter is the problem of explaining the determinacy of measurement outcomes, whereas wave-functions generally do not assign determinate properties to the systems they describe. Realist solutions to the measurement problem involve either completing the structure of the theory with additional variables (e.g. position of particles in Bohmian mechanics), with an additional dynamic (spontaneous localisation in GRW and CSL) or denying that measurement outcomes are objectively determinate after all (many-world interpretation). However once the measurement problem is solved, the problem we are concerned with might remain. In particular, although Bohmian mechanics provides an ontology of particles, GRW, CSL and many-world theories do not provide an ontology for quantum mechanics apart from the wave-function, and it remains unclear, according to these theories, what exactly exists in ordinary space-time. Conversely, providing an ontology in ordinary space-time for quantum mechanics solves the measurement problem out of hand, in so far as the locality of the primitive ontology can ground the determinacy of measurement outcomes (assuming, following Bell, that all measurements reduce to position measurements). Thus the metaphysical problem that primitive ontologies aim to solve can be viewed as a stronger version of the measurement problem.

The program is not devoid of difficulties. There are difficulties in elaborating quantum relativistic versions of these new theories, but the main difficulties concern status of the wave-function. The wave-function cannot be dispensed with, for it does all the predictive job, so to speak, but its ontological status remains unclear (Belot 2012).

If the wave-function is accepted as an object of our ontology, on a par with the additional structure, then the primitive ontology seems derivative rather than truly 'primitive', and it seems that it could be reduced. In the case of GRWf and GRWm, the primitive ontology is deduced from the wave-function and laws, and although mathematical dependence does not always imply ontological dependence, the primitive ontology might seem superfluous in these particular cases. Furthermore, having causal relations between objects living in different spaces, one of which is a still mysterious high-dimensional space, is arguably not a very satisfying option (Egg & Esfeld 2015).

For these reasons, most authors fall back on a construal of the wave-function as a nomic entity. The wave-function would be a dynamical parameter figuring in a physical theory to describe the evolution of the primitive ontology. From this perspective, several metaphysical interpretations are possible. They correspond to the main metaphysical interpretations of the nomic aspects of theories in general: governing laws of nature, humeanism or dispositionalism.

Conceiving of the wave-function as a law seems hardly intelligible, since laws of nature are generally taken to apply in all places and times. But wave-functions can vary in space and time. Furthermore, assuming that physicists are actually measuring 'laws' when, say, they prepare a system in a superposition of state seems rather unintuitive. In addition, it is legitimate to wonder if we are not merely sweeping under the rug of 'laws' all the interpretational difficulties associated with the wave-function's multi-dimensional space.

One could have a humean conception of laws, whereby laws are part of the 'best system', i.e. the best description of the distribution of the primitive ontology in terms of simplicity and informativeness. Following this view, laws do not govern matter, but supervene on it. Humeanism is problematic at least for some of the primitive ontology proposals, such as GRWm, where the quantum state does not supervene on the distribution of the primitive ontology (Egg & Esfeld 2015). In any case there are objections against humeanism in general (Carroll 2008) (although some standard objections, such as objections from quidditism or humility, do not apply here (Esfeld 2014)).

The last option consists in assuming that the wave-function represents the dispositional properties of matter (Egg & Esfeld 2015). However Egg & Esfeld observe that these dispositions cannot be assigned to the individual parts of the primitive ontology (the local beables) directly. They have to be assigned to local-beable configurations instead. Ultimately under this view, the wave-function of the universe is just the dispositional structure of the whole universe. The difference between attributing a dispositional structure, or a state, is slim and it is not clear that we have made much progress so far.

Recall that the aim of primitive ontologies is to reconcile the scientific and the manifest image of the world. A central aspect of the manifest image of the world is that objects have localised properties. Arguably, some of (if not all) these properties are dispositional and one should be able to explain why ordinary, localised objects have localised dispositions. If all that exists is a disposition of the whole universe to evolve, then the account fails on this respect and the explanatory gap remains (see Dorato (2015) for related arguments).

Another concern one can have with the primitive ontology program is methodological. The program seeks to complement our theories with an *ad-hoc* mathematical structure for metaphysical purposes only. The primitive ontologies are thus inert, in the sense that they play no predictive role. They rather play an explanatory role as bearers, in ordinary space-time, of the nomic aspects of the theory. Albert (2015) notes that this contradicts our common intuitions regarding the causal role played by the fundamental constituents of reality and in any case, from this lack of predictive role, an underdetermination by experiment ensues.

This could raise some suspicions: aren't we merely trying to fit the theory into the mould of our pre-theoretical intuitions? After all, physicists do not seem to have much difficulties in making a connection between models and experiments (at least once a solution to the measurement problem is assumed). They do not postulate additional structures beyond those of the textbooks for that purpose. It seems reasonable to assume that the theory already has all the resources needed to relate to our familiar space-time. Shouldn't metaphysics be content with providing good metaphysical interpretations of the predictive, theoretical structure the physicists postulate for their empirical needs? Do we really need to build some formal machinery to cure our metaphysical worries? One could interpret the primitive ontology program as an attempt to domesticate science with metaphysics (Ross et al. 2007) and be sceptical about the whole approach. As Ross et al. argue, our pre-theoretical intuitions, such as the common idea that all physical objects are constituted of localised parts, are adapted to our familiar environment but nothing indicates that they are still relevant when applied to remote domains of experience. Our common intuitions might not be good guides to address the fundamental nature of the physical world, and modifying or completing our scientific theories on their behalf might not be the best strategy.

Personally I tend to share this scepticism, and while I agree that physics should tell us explicitly what exists in ordinary space-time, it seems to me that the primitive ontology program and wave-function realism both share the unwarranted assumption that an interpretation of the wave-function in ordinary space-time is not available. Rather than arguing furthermore, I would like to give a positive account of how the wave-function can be interpreted as representing objects living in ordinary space-time. If it works, then an alternative exists to the primitive ontology program and to wave-function realism which does not face the difficulties of both approaches.

2 An alternative to primitive ontologies

2.1 Interpreting the configuration space

Both the primitive ontology defenders and the wave-function realists implicitly assume that the wave-function alone cannot represent objects living in ordinary space-time, because the wave-function is a field on configuration space. However there are some reasons to think that the configuration space derives from ordinary space rather than the converse. First, the concept of configuration space originates in classical physics, and one could expect a certain continuity in the way it is interpreted and connected with experience. Second, it should be noted that quantum field theory assumes a relativistic space-time as a background, and that configuration spaces associated with bounded regions of space-time are derived from this background space. But even restricting the discussion to non-relativistic quantum mechanics, the fact that a mathematical object is defined on a high dimensional *mathematical* space does not entail that the objects it represents live in a high dimensional *physical* space.

Peter Lewis (2004) proposes an argument to the effect that the many dimensions of configuration space refer to the few dimensions of ordinary space (*pace* Monton (2006)). The argument rests on the observation that certain coordinate transformations leave dynamical laws invariant, and not others, and that the ones which do correspond to transformations of a three dimensional space. A similar and perhaps more general argument can be based on the fact that in quantum mechanics, particles of the same kind are believed to be indiscernible, which translates into the assumption that exchanging two coordinate variables changes the wave-function by only a plus or minus sign:

$$\Psi(...x_i...x_j...) = \pm\Psi(...x_j...x_i...)$$

The x_i in this formula correspond to three dimensional position vectors. The assumption means that the probabilities calculated from the wave-function of a composite system should be invariant under a permutation of the (ordinary) space-coordinates of each identical particle. Wave-functions which do not respect this constraints are considered unphysical.

If particles do not exist, but only the wave-function, an interpretation in terms of indistinguishability of particles is no more available, yet the assumption in its mathematical form remains valid.

At this point there are two possible approaches. Either one assumes that the 'real' space is the configuration space, and that this constraint on acceptable physical states is a brute fact about nature. Perhaps this brute fact would help us explain why physical space *seems* to have 3 dimensions, when it actually has $3N$. However Lewis' argument can be transposed to this case: it would force us to assume that not all coordinate transformations of the configuration space are on a par. Certain transformations (those which do not correspond to transformations of the coordinates of ordinary space) would change a physical wave-function into an unphysical wave-function. This is not acceptable from a scientific point of view, where coordinate choices in a physical space are generally considered arbitrary. Typically, a coordinate transformation can correspond to the same system observed from a different perspective. Then permutation invariance would entail that some perspectives in ($3N$ dimensional) physical space are forbidden by the theory, which doesn't seem to make much sense.

The second and less contentious option is to assume that physical space actually has 3 dimensions, and that different mathematical dimensions of the configuration space are actually referring to the same physical dimensions, which partly *explains* the constraint on possible states and preserves the arbitrariness of coordinate choices in such a way that any coordinate transformation in ordinary space changes a physical wave-function into another physical wave-function.

If this is the case then no additional structure is needed to interpret the wave-function as an object in ordinary space, since the physical dimensions on which the wave-function is defined are the 3 dimensions of ordinary space. All we need is to interpret the structure that is already there, and in particular, the configuration-space points.

Let us assume just this. This will be our starting point for developing an ontology in ordinary space-time from the structure of the wave-function alone.

2.2 Interpreting the wave-function

If physical space has three dimensions, then a configuration-space point is an ordered set (or N-tuples) of N space-time points (just as in figure 3 above, each configuration-space point is a couple of one-dimensional space points) and the wave-function assigns properties to these sets of space-time points. So far, we have what Belot (2012) calls a multi-field. (Note that the statistical features of quantum mechanics mentioned above, concerning particles indistinguishability, can help us simplify the picture and assume that, at least in the case of fermions of the same kind, configurations can be represented by simple, unordered sets of space-time points. See appendix A.)

The position I will defend is that configuration-space points directly represent the fundamental constituents of reality. Let us call these fundamental constituents 'configurations' for simplicity, although one should bear in mind that they are not configurations of anything, but fundamental objects.

What could be the ontological interpretation of a N-tuple (or a set) of space-time points?

Each configuration could represent a complete classical state of the world where N particles have well-defined positions, each point in the set representing one of these positions[1]. However one should resist this interpretation: that would require particles at the fundamental level, and that would contradict the idea that configurations are fundamental. The idea is to conceive of configurations as primitive objects.

The alternative I propose is to view configurations as primitive, concrete N-ary relations between space-time points. This proposal can be seen as an elaboration of ontic structural realism (Ladyman 1998), or the view that the fundamental constituents of reality are primitive relations which do not supervene on their relata. Configurations do not indeed supervene on space-time because they have their own dynamical properties, the complex-valued coefficient (a weight and a phase) that the wave-function assigns to them, and therefore they do not reduce to the structure of space-time alone.

These primitive relations are represented mathematically by sets of space-time points. Following this view, the configuration space should not be interpreted as a

[1] That would correspond to Sebens (2015)'s proposal to construe quantum mechanics as describing many interacting classical worlds.

physical space, but as a mere mathematical tool encoding the structure of spatial relations between different configurations. Two configurations are close in configuration space when they are similar, in that the points they relate are, two by two, in the same neighbourhoods.

Note that the weight of a configuration should be interpreted as a coefficient of existence of some sort (a degree of 'something being there'), otherwise the link with observations would be broken. State attribution to the world (hence knowledge) would be impossible: some states of the world more probably exist than others, in particular the ones which contain the wave-functions or state matrix we attribute to the systems we prepare and measure. This 'coefficient of existence' could be interpreted as a disposition to manifest in the world (which would fit with the view defended by Ladyman that the primitive relations are modal). What this 'manifestation' amounts to (and whether it is objective or not) will depend on the interpretation. I will leave aside these interpretational issues here.

In any case, attributing 'coefficients of existence' to configurations does not amount to attributing matter density to regions of ordinary space-time (and the proposal does not reduce to GRWm). There is more structure to it. It amounts to applying a density function to configuration space, but as just said, the configuration space should be viewed as a mere mathematical tool describing the relational structure of configurations in ordinary space, while configurations are the only fundamental constituents of reality. Moreover there is a one-to-many relation between density functions on configuration space and density functions on ordinary space. Density functions on ordinary space are less informative. In that sense, a matter density field on ordinary space would be derivative rather than primitive.

The Schrödinger equation describes the evolution of these complex-valued coefficients with time. One of its salient aspects is that it is local-deterministic, which means, loosely speaking, that the evolution of the coefficient of a configuration is influenced by other configurations and their coefficients only when all their points are in the same neighbourhoods. In other terms, configurations interact together when they are closely related in configuration space, which translates in their relating closely related points in physical space.

Thus we have an interpretation of the wave-function as attributing properties to objects living in a 3 dimensional space, the configurations, which demonstrates that we need not complement quantum mechanics with additional mathematical structures to obtain an ontology in ordinary space-time. However, in order to show that this interpretation is successful, one needs to show that the manifest image of the world, consisting in our familiar, localised objects with determinate properties, can be recovered from this picture. Only two more ingredients are needed for this task, and they are fully part of standard quantum mechanics.

2.3 Recovering familiar objects

In the remaining of this section, I will explain how one can recover familiar objects from an ontology of configurations. My purpose is not to provide a rigorous account, rather to give the reader an intuitive grasp of how it can be done. I hope this will suffice to convince that an alternative exists to primitive ontologies.

Our first ingredient to recover familiar objects is decoherence. If the decoherence program is sound, then one can imagine that configurations eventually form separate branches, *i.e.* clusters of configurations with high coefficients which no more interact significantly with configurations located outside their own cluster (because the actions of these external configurations cancel each other out), yet still interact with each other inside their own cluster. Further, all the configurations in a given branch will display similarities in terms of the location of the space-time points they relate: a branch corresponds, roughly, to a dynamically stable set of macroscopically indiscernible configurations with high coefficients (see figure 4 for an illustration)[2].

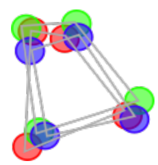

Figure 4: Illustration of a cluster of configurations

In collapse theories such as GRW and CSL, stochastic processes might smooth out some branches at a rate sufficiently high to remove all but a single branch at any time (this branch will then divide into new stable branches, but they will be smoothed out in turn by future processes). Then the wave-function would represent a single evolving branch, or cluster of macroscopically indiscernible configurations. Having such processes is not obviously required (what exactly is the trouble with having several non-interacting branches in the world?) but it can help recover the link between configuration coefficients and empirical probabilities, through an appropriate dynamic.

The second ingredient we need to recover familiar objects is the separability of a cluster of configurations into localised objects. Here we can simply follow the physicists, who have a well-defined notion of separability (or lack of entanglement) between different systems in terms of factorisability of the wave-function.

Separability ensures that objects will not display spooky correlations in their respective evolutions. The notion can be explained intuitively as follows: imagine we divide the whole physical space S in two sub-spaces $S1$ and $S2$. Then each configuration

[2] In traditional interpretations, decoherence results in a trade-off between the determinacy of the relative positions of particles and the determinacy of their relative momenta, which would translate here into a trade-off between the macroscopic 'similarity' of the configurations in the branch and the dynamical stability of the branch.

of N space-time points can be described as a combination of two sub-configurations, one of P points in $S1$ and the other of Q points in $S2$, where $P + Q = N$. If the configurations are similar enough, P and Q will be the same for all configurations in the branch. The wave-function can be viewed as attributing coefficients to pairs of sub-configurations instead of configurations. Now picture these coefficients as the probabilities that each sub-configuration is paired with the other, and imagine we find no specific correlation: a sub-configuration in $S1$ has the same probability of being paired with any sub-configuration of $S2$ than any other. In such a case, we might view our whole branch as the combination of two independent 'regional branches' living in sub-spaces $S1$ and $S2$, each with its own dynamical evolution. The separability condition ensures that these regional branches will evolve autonomously, at least in a short time range.

We can continue this decomposition process recursively until the separability condition fails. The recovery of separability is not always possible. It is impossible for entangled states, such as states in a EPR experiment. Moreover, separability is usually approximate, and the conditions for separability are a probabilistic matter. However we have good reasons to think that these conditions will only fail at a microscopic level.

Assuming this, a branch can be decomposed into a set of microscopic three dimensional objects, which are really very localised 'regional branches', or clusters of small sub-configurations each relating a few space-time points in the same location (or even a single space-time point for each), and this decomposition will be relatively robust in time (figure 5). There is a sense in which mereology is recovered: big clusters (and branches) can be seen as composed of smaller, regionalised sub-clusters, and the properties of the big cluster approximately supervene on the properties and arrangement of its parts. These small sub-clusters will look like atomic, autonomous classical objects. More importantly, as we saw that the law of evolution is local-deterministic and in virtue of separability, configurations will interact together just as if these atomic objects were directly interacting (i.e. objects will influence other objects nearby, independently of objects far apart). Thus the manifest image of the world can be recovered.

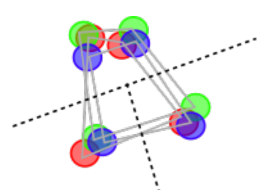

Figure 5: Illustration of a separable cluster of configurations

To summarise, we have one fundamental level and two levels of emergence in this picture:

1. At the fundamental level, configurations are concrete relations between space-

time points.

2. Branches emerge from decoherence; they are dynamically robust sets of macroscopically indiscernible configurations.

3. Our familiar localised objects and their constituents emerge from separability; they are stable sets of regional sub-configurations.

Note that the emergence of objects from configurations is top-down rather than bottom-up: parts emerge from the whole. The reason for this is that configurations are holistic objects. As we will see below, this holism is the main characteristic difference between this interpretation and other primitive ontology accounts.

3 Virtues and possible objections

As already noted this presentation does not pretend to be perfectly rigorous. For example, it should be noted that a certain notion of separability of systems from their environment (in terms of interaction hamiltonian) is a pre-condition of decoherence, and things are probably a little more entangled and complicated: the question of separability in quantum mechanics is still an active area of research (For more on this topic, see Dugić & Jeknić 2006, Jeknic-Dugic et al. 2013).

However we can leave out the details. Assuming that we can make sense of decoherence and separability in physics, it seems that there is no real obstacle in interpreting the wave function as a structure living in a low-dimensional space-time, from which familiar objects can be recovered. This might not be a complete, true picture of the world (not as far as quantum field theory is concerned), but if it is coherent, then the motivations for a primitive ontology in terms of local beables in complement to the wave-function are unfounded.

One virtue of this proposal, shared by wave-function realism, is its parsimony. One doesn't need to add new structures to the theory for metaphysical reasons, pointing at them as 'what really, primitively exists'. All the needed structures are already present in the formalism: they are the configuration space points.

Incidentally, we face no specific problem in combining these structures with special relativity. The view presented here can be easily extended to Fock spaces by assuming that configurations do not all relate the same fixed number of space-time points, which can account for particle creation and annihilation, and which is a first step toward a relativistic version of this interpretation[3].

Another virtue of the present interpretation is that the status of the wave-function is clear. As noted before, the primitive ontology program faces some difficulties concerning the status of the wave-function. It is generally assumed that it must be relegated to

[3]Although there might be complications for separability conditions and the identification of objects, in relation to the well known interpretational difficulties of quantum field theory (Earman & Fraser 2006).

the status of a nomic entity, but its exact ontological status remains to be clarified (see section 1).

In the alternative proposed here, the wave-function is not a nomic entity: it assigns properties to objects in ordinary space-time, the configurations. Furthermore, the fundamental objects of ontology are not surplus structure which play no predictive role, since their mathematical counterparts are part of the standard formalism, contrarily to other primitive ontology proposals. The way dispositions can be assigned to local objects, through separability, is also made clear.

These virtues of parsimony and clarity of the status of the wave-function are shared by wave-function realism, but the present interpretation does not face the difficulties of wave-function realism in accounting for the correspondence between the theory and ordinary space-time. For example, no *ad-hoc* assumption concerning the way the $3N$ dimensions of configuration space are mapped to the 3 dimensions of ordinary space is required, since the configuration space is a mere mathematical tool constructed from ordinary space-time, just as in classical physics. Thus we take the advantages of both primitive ontologies and wave-function realism, without the drawbacks.

Let us briefly review some possible objections to this proposal.

A first caveat is that I had to appeal to the notion of probability in my explanation of separability (interpreting coefficients as a probability measure). The association between probabilities and the coefficients is not obvious, but this is an independent issue, which is related to the measurement problem. Furthermore, the narrative proposed here was meant to give an intuitive grasp of what separability amounts to, but this narrative needs not be taken too seriously. The important point is that separated objects can be assigned a relatively autonomous dynamic. Different interpretations might then be envisaged. Even if this aspect were a problem, one can accept that the coefficients do indeed correspond to dispositional properties, associated with a probabilistic dynamic.

A second worry could concern the use of emergence. What kind of reduction is involved in talks of branches and 'regional branches' of 'sub-configurations'? (a functional reduction perhaps?) Should we be eliminativist with regards to ordinary objects and say they are illusory? I have no strong opinion on this subject, but it doesn't seem absurd to me to say that when we talk about ordinary objects, we are really talking about stable, local similarities between the instantiated configurations of the universe. Talk of ordinary objects is full of identity paradoxes, which the loose notions of similarity and stability might help circumvent.

Another possible objection is that I did not really do away with a primitive ontology, because I assumed space-time points. Indeed the present interpretation is not entirely neutral with regards to the metaphysics of space-time. Under a relational construal, i.e. if space-time only describes spatiotemporal relations between objects, we would have to assume that space-time describes spatiotemporal relations between configurations, which are themselves relations between space-time points. Perhaps there is a way to make sense of it (for example by assuming that configurations are simple objects but have many spatiotemporal relations to every other configurations instead of one), but such an account seems *prima-facie* circular[4]. However the interpretation is at least compatible with a substantialist metaphysics of space-time, or with a structuralist construal, according to which space-time points are identified by their positions in the relational structure of space-time only.

Finally I think that the present interpretation comes close to the way the formalism is naturally interpreted in practice. Experimentalists implicitly assume that the many dimensions of configuration space refer to the three spatial dimensions of our physical space. This assumption is involved for example when they calculate a detection probability in a particular position in space from the wave-function of a composite system, or when a wave-function is assigned to a composite system in a prepared initial state. It is also implicitly involved in the indistinguishability postulate, and perhaps also in the dynamical laws when they display symmetries corresponding to coordinate transformation in ordinary space-time. However experimentalists need not assume that there are objects beyond the wave-function to derive their predictions.

4 Conclusion: a prejudice against holism?

Wave-function realists and primitive ontologies both assume that taking the wave-function ontologically seriously implies that the configuration space must be our fundamental physical space. The former are willing to bite the bullet, while the latter think this is a problem and attempt to complement the wave-function with a primitive ontology in ordinary space-time and relegate the wave-function to the status of a nomic entity.

I think this shared assumption is unwarranted: first, it is doubtful that the configuration space is primitive, rather than derivative, and second, it is perfectly possible to interpret the wave-function as representing objects in ordinary space-time. In this paper, I attempted to develop such an interpretation by taking the points of configuration-space to represent the fundamental constituents of the world. There might exist other possibilities, but suffice to say that this account succeeds to show that the motivations of primitive ontology proponents are unfounded. Interpreting the wave-function as representing objects in ordinary space-time allows one to take the advantages of wave-function realism (parsimony and a clear ontological status for the wave-function) and that of primitive ontologies (an explicit correspondence to ordinary space-time objects) without facing the respective difficulties of both positions.

The account proposed here share some similarities with primitive ontologies: it also attempts to make clear the link between the ontology of quantum mechanics and the manifest image of the world, by taking seriously our ordinary space-time. The main difference with other primitive ontology proposals such as flashes, particles and matter density fields is that the fundamental objects of our ontology are not local, but

[4] I am grateful to Matthias Egg for raising this point.

relational. We have to accept a form of relational holism, the view that relations do not supervene on the intrinsic properties of their relata, and the view that ordinary objects supervene on relational aspects rather than the converse.

The assumption that an ontology should assign properties to local objects is shared by wave-function realists and primitive ontologists altogether: both camps are willing to interpret the wave-function as attributing properties to points in a physical space (be it our ordinary space or not). This is true at least following the prescriptions of some authors who put emphasis on the term ‘local beable’. This amounts to denying that entanglement, which is a central feature of quantum mechanics, is a fundamental aspect of reality: entanglement between particles would not mean that physical systems are themselves non local, but could be reduced to nomic aspects, or explained away by the fact that we do not actually live in a 3 dimensional space.

From this assumption that ‘what exists’ must be local (which probably has its roots in common sense intuitions) stem all the difficulties, including the remarkable fact that the dispositionalist cannot avoid a form of holism, with dispositions being assigned to configurations of local-beables rather than to the primitive ontology directly [5].

Why not, then, reject this assumption of locality and accept the main tenet of ontic structural realism that relations be primitive objects? The position is known to have other virtues, after all.

A Appendix

Permutation invariance can be formulated as follows:

$$\Psi(...x_i...x_j...) = \pm\Psi(...x_j...x_i...)$$

It can be shown from the formula above that the value of the field at any position in configuration space uniquely specifies the value at a position which corresponds to a different ordering of the space-time points.

In the anti-symmetric case $\Psi(...x_i...x_j) = -\Psi(...x_j...x_i)$ which corresponds to fermions, the wave function will assign a zero value in case two space coordinates have the same value ($x_i = x_j$).

It follows that assigning a complex value to every (unordered) set of (distinct) space-time points uniquely determines the wave-function of fermions of the same kind, and that the wave-function for fermions can be made equivalent to a field on sets of space-time points. This might not be the most useful description for the purpose of the physicist, but it captures the fact that we do not really need particles in our ontology, only space-time points (where the ordering and repeatability could have left a doubt).

[5]Esfeld (2015) argues that primitive ontologies can be reconciled with structuralist, holistic accounts by assuming that primitive objects are identified by their positions in a relational structure. However there is a difference between assuming that laws and structures of objects are holistic, and assuming that primitive objects are themselves non-local entities, and it is the latter assumption that makes all the difference.

Note that this picture can easily be extended to Fock spaces, if the number of space-time points in configurations is indeterminate: wave-functions are then fields on the power set of space-time points.

References

Albert, D. Z. (1996), Elementary Quantum Metaphysics, *in* R. S. Cohen, J. T. Cushing, A. Fine & S. Goldstein, eds, 'Bohmian Mechanics and Quantum Theory: An Appraisal', Vol. 184, Springer Netherlands, Dordrecht, pp. 277–284.

Albert, D. Z. (2015), *After physics*, Harvard University Press, Cambridge, Massachusetts.

Allori, V., Goldstein, S., Tumulka, R. & Zanghì, N. (2008), 'On the common structure ofBohmian mechanics and the Ghirardi–rimini–weber theory: Dedicated to Giancarlo Ghirardi on the occasion of his 70th birthday', *British Journal for the Philosophy of Science* **59**(3), 353–389.

Bell, J. S. (2004), *Speakable and Unspeakable in Quantum Mechanics: Collected Papers on Quantum Philosophy*, Cambridge University Press.

Belot, G. (2012), 'Quantum states for primitive ontologists', *European Journal for Philosophy of Science* **2**(1), 67–83.

Carroll, J. W. (2008), *Laws of nature*, Cambridge University Press, Cambridge.

Dorato, M. (2015), 'Events and the Ontology of Quantum Mechanics', *Topoi* .
URL: *http://link.springer.com/10.1007/s11245-015-9315-6*

Dugić, M. & Jeknić, J. (2006), 'What is "System": Some Decoherence-Theory Arguments', *International Journal of Theoretical Physics* **45**(12), 2215–2225.

Earman, J. & Fraser, D. (2006), 'Haag's theorem and its implications for the foundations of quantum field theory', *Erkenntnis* **64**(3), 305 – 344.

Egg, M. & Esfeld, M. (2015), 'Primitive ontology and quantum state in the GRW matter density theory', *Synthese* pp. 1–17. forthcoming.

Esfeld, M. (2014), 'Quantum Humeanism, or: Physicalism without Properties', *The Philosophical Quarterly* **64**(256), 453–470.
URL: *http://pq.oxfordjournals.org/cgi/doi/10.1093/pq/pqu030*

Esfeld, M. (2015), 'How to account for quantum non-locality: ontic structural realism and the primitive ontology of quantum physics', *Synthese* pp. 1–16.

Jeknic-Dugic, J., Arsenijevic, M. & Dugic, M. (2013), *Quantum Structures A View of the Quantum World*, LAP LAMBERT Academic Publishing, Saarbrücken.
URL: *http://arxiv.org/abs/1306.5471*

Ladyman, J. (1998), 'What is structural realism?', *Studies in History and Philosophy of Science* **29**, 409–424.

Lewis, P. J. (2004), 'Life in Configuration Space', *British Journal for the Philosophy of Science* **55**(4), 713–729.

Monton, B. (2002), 'Wave function ontology', *Synthese* **130**(2), 265 – 277.

Monton, B. (2006), 'Quantum Mechanics and 3nDimensional Space', *Philosophy of Science* **73**(5), 778–789.

Ney, A. & Albert, D. Z. (2013), *The Wave Function: Essays in the Metaphysics of Quantum Mechanics*, Oxford University Press.

Ross, D., Ladyman, J. & Spurrett, D. (2007), In Defence of Scientism, *in* 'Every Thing Must Go: Metaphysics Naturalized', Oxford University Press.

Sebens, C. T. (2015), 'Quantum Mechanics as Classical Physics', *Philosophy of Science* **82**(2), 266–291.
URL: *http://www.jstor.org/stable/info/10.1086/680190*