

Quantum Mechanics on Spacetime I: Spacetime State Realism

David Wallace Christopher G. Timpson

May 2009

Abstract

What ontology does realism about the quantum state suggest? The main extant view in contemporary philosophy of physics is *wave-function realism*. We elaborate the sense in which wave-function realism does provide an ontological picture; and defend it from certain objections that have been raised against it. However, there are good reasons to be dissatisfied with wave-function realism, as we go on to elaborate. This motivates the development of an opposing picture: what we call *spacetime state realism*; a view which takes the states associated to spacetime regions as fundamental. This approach enjoys a number of beneficial features, although, unlike wave-function realism, it involves non-separability at the level of fundamental ontology. We investigate the pros and cons of this non-separability, arguing that it is a quite acceptable feature; even one which proves fruitful in the context of relativistic covariance. A companion paper discusses the prospects for combining a spacetime-based ontology with separability, along lines suggested by Deutsch and Hayden

1 Introduction

In this paper we address two main questions: Suppose one were a realist about quantum mechanics; and suppose, in addition, that in the version of quantum mechanics one was committed to, the state is to be understood as standing for some thing in the world. How is one to think of what the world is like in that case? What *is* it that the state stands for? This is our first question.

Our second follows on. Quantum mechanics is well-known for its nonlocal features; for its failure to satisfy a principle of separability. Quantum states might, generically, be non-separable; but given an answer to our first question, will it be true that the *thing* thereby represented (however we are to think of it) is non-separable too? That is, when we are more careful to specify what the ontology associated with realism about the quantum state is, will that ontology satisfy separability, or not?

There has been remarkably little discussion of what the correct way of thinking about the ontology of the quantum state is. Here we develop an approach

which we term *spacetime state realism*, in opposition to *wave-function realism*, which has so far been the default option; apparently for want of alternatives.

We begin (Section 2) by sketching some requirements on specifying an ontology for quantum mechanics, and go on to introduce wave-function realism (Section 3), indicating how it meets these requirements. But there are good reasons to resist wave-function realism: it is plausible only in the restricted setting of fixed-particle number non-relativistic quantum mechanics and it has no reasonable extension to more physically realistic situations (Section 4). We therefore develop a position which avoids these problems, taking states of spacetime regions as the primary elements (Section 5); and we enumerate its benefits over wave-function realism.

What of separability (Section 6)? The spacetime state ontology is unambiguously non-separable: the state of a given spacetime region will not be determined by the states of the sub-regions of which it is composed. By contrast, however, the wave-function ontology *is* separable at the fundamental level. This marks a significant difference between these ontological pictures, but we go on to argue that the failure of separability is by no means a pathological feature of our preferred account. Rather, it is a feature which has active benefits when one comes to consider relativistic covariance and an important phenomenon recently noted by David Albert, that quantum mechanics exhibits so-called *narratability failure* (Section 8).

Our primary concern here is with realist accounts of quantum mechanics that involve purely unitary evolution, thus the Everett interpretation will loom large; however a number of our remarks will apply equally to other kinds of theories which take the quantum state seriously as an element of reality, such as Bohm-like hidden variable theories and dynamical collapse theories. We consider the matter of non-unitary dynamics within our spacetime state ontology in Section 7.

In a companion paper (Wallace and Timpson [2009]) we explore the matter of separability in more detail; and investigate the remarkable possibility held out by Deutsch and Hayden ([2000]) that one may gain the benefits of a spacetime-based approach to quantum ontology while retaining separability at the fundamental level.

2 Prolegomena

What ontology does quantum theory prescribe? Historically, finding an answer has not proven easy, for a variety of reasons. One way at least to set foot on a path to an answer is by beginning with the question of what in *general* might be involved in specifying an ontology associated with a physical theory. We have no wish to be prescriptive, but it seems reasonable to agree that at least *one* way in which an ontology may be specified (at least minimally) is by identifying suitable bearers for the properties that the theory trades in, and then going on to say something more about those properties themselves. The austere mathematically framed description of these properties provided by theoretical physics might,

on its own, seem a little thin: one would also like to be in a position to say something about how they relate to the everyday world of moderately sized objects which surrounds us. It is, however, notoriously difficult to provide a satisfactory story about any of these various elements when discussing quantum theory. Potential bearers and properties that one might identify generally seem too rebarbatively alien to get a good grip on; and the problem of measurement casts a long shadow over the question of recovering appearances (see Wallace ([2008]) for a review).

But what might be required in order to gain at least a reasonable grip on a proposed ontology? We should be wary of setting our sights too high and insisting on too much when trying to ponder what fundamental physics presents us with. Consider, for example, a simpler case: that of electromagnetism.

We are happy enough that electromagnetism does not present us with an irredeemably obscure picture of the world, yet it is not as if we really have an intuitive grasp of what an electric or a magnetic field is, other than indirectly and by means of instrumental considerations (“A test charge would be accelerated *thus*.”, for example). But in the case of field theory, we do at least understand that “the electric field at spacetime point p ” denotes a property *of point p* ,¹ and this gives us substantial intuitive understanding of how to think about that field: namely, it describes certain, admittedly somewhat alien, properties *of spacetime regions*.

Does electromagnetism offer any *further* elucidation of its ontology, beyond an assignment of vectorial quantities to spacetime points? We think not: what could that understanding derive from? All we have to work with are the mathematical structure of the theory and the instrumental considerations alluded to above—and the latter, when one takes seriously the idea of physics as universal, ultimately collapse into interactions between the electromagnetic field and other comparably alien entities. (It may help here to recall that *matter*, in the sense of the macroscopic material bodies we observe, consist of, and manipulate in the lab, cannot ultimately be thought of as some non-electromagnetic entity on which electromagnetic forces act: modern physics makes clear that solid matter is made up of electromagnetic fields as well as fermionic matter. In any case, the latter is scarcely less alien than the former. Even if we naively regard fermions as classical point particles, “point particles” are shapeless, colourless, textureless entities with little in common with familiar macroscopic bodies, for all that schoolbook physics may have acclimatised us to them.)

Thus it seems that we gain a basic understanding of the electromagnetic field by seeing it as a property of spatial regions²; and our further understanding must be mediated by reflecting on its role in the theory (including, importantly, the instrumental considerations); beyond that there doesn’t seem to be much

¹Strictly speaking, assigning a vectorial quantity to a point will also involve making reference to relations with neighbouring points, c.f. Butterfield ([2006]).

²We do not wish, in making this claim, to make some sort of dogmatic commitment to spacetime substantivalism. It may well be, for instance, that talk of spatial *regions* is in principle eliminable in turn in favour of talk of the spatial *relations* between events, as in Barbour’s version of field theory (Barbour, Foster, and Ó Murchadha [2002]).

further to be grasped.

Now consider a converse example. In *classical* mechanics (of particles or fields) it is possible to construct very elegant and abstract formulations in terms of trajectories on a symplectic manifold. Such formulations are normally understood as a mere mathematical convenience—a shorthand for the real field-based or particle-based description—but how should we think about such a theory (the theory, say, of 10^{23} point particles) if we were given it *without being told* that it was intended in that way?

The theory would certainly look alien. It would present a single point evolving in an enormously high-dimensional (6×10^{23} -dimensional) symplectic space. The space should not be understood as featureless, though (or even featureless other than the *symplectic* structure): the Hamiltonian that guides the point's evolution will be highly non-invariant under generic symplectomorphisms (transformations only preserving symplectic structure), and so the symmetry group of the theory will be very small compared to the symplectic group. On the usual assumption (see, e.g., Earman ([1989])) that a spacetime must have a structure rich enough to formulate its dynamics, but no more³—and so can have no invariants or non-invariants under transformations that the dynamics do not have—this means that the space (unlike, say, 6×10^{23} -dimensional Euclidean space) is highly non-isotropic and non-homogeneous.

Nonetheless, if we were to think of our theory as just about the evolution of a single particle on that huge, exotically structured space, we would have failed to understand its structure fully. True understanding comes from the recognition that the theory is naturally isomorphic to—and so can be understood as representing—the positions and velocities of 10^{23} particles in ordinary, Euclidean, three-dimensional space. Only once we recognise that fact can we get a clear grip on the theory's structure, or relate it to our observed world of spatially located objects.

Our claim, in essence, is that thinking about quantum mechanics in terms of a wavefunction on configuration space is rather like thinking about classical mechanics in terms of a point on phase space. In both cases, there is a far more perspicuous way to understand the theory, one which is connected to spacetime in a more direct way.

Before moving on, a caveat is in order. In trying to address the question of how the world is according to quantum mechanics, we wish, so far as possible, to remain neutral between a metaphysically substantive and a deflationary reading; neutral, that is, between the (substantive) thought that we can successfully talk about the one correct reading of the quantum state (description of how the world is if QM is true) and the (deflationary) thought that perhaps the best we can do is compare the merits of different more or less perspicuous descriptions of something we can get no closer to (while remaining within the ambit of the theory). Thus while we argue for the adoption of spacetime state realism over wave-function realism, we wish to remain neutral on whether one of these (or

³The conclusion also follows from Brown's thesis (Brown [2005]) that a spacetime's structure is just a *description* of its dynamics.

perhaps some third) really does provide the One True Interpretation of the quantum state, or whether one is merely a more perspicuous description than the other; a description of something that we are ultimately unable to render unequivocally in intuitive terms.

Equally in classical physics, it seems clear to us that the more perspicuous description of the theory is in terms of N particles in three-dimensional space, not one particle in $6N$ -dimensional space, but we are neutral (here, at any rate!) as to whether this means that the N -particle story is closer to the truth, or just that the N -particle story is a more sensible way of describing the structure of the world.

We do, however, want to stress that we regard the one-particle story as a *coherent possibility* (whether or not it is just the same possibility as the N -particle story, differently described)—and that if the N -particle story were empirically adequate (which it isn't, of course) then so would the one-particle story be. For that story (by construction) is isomorphic to the N -particle story, and the emergence or otherwise of higher-level ontology from lower-level theories depends, to our minds, primarily on the structure of those theories and not on their underlying true nature (whatever that is). On the one-particle theory, three-dimensional space would turn out to be emergent, but it would be no less real for that: indeed, it would be emergent in such a strong and exact sense as to make us inclined to regard the N -particle story as to be preferred.⁴

One might put things in these terms: in our view, there is no guide to the ontology of a mathematically formulated theory beyond the mathematical structure of that theory (including that structure to do with observable quantities and experiment). But when trying to learn ontological lessons from the theory, one does well to prefer a representation which makes manifest the structure that the theory ascribes to the world. One of the difficulties with the single particle, high dimension story about classical mechanics is that it makes it less immediately obvious what structure is being attributed to the world. The physical world is very richly structured; a distribution of field amplitudes over 3-space, or an arrangement of particles in 3-space is very obviously richly structured; a single point in high dimensional space less obviously so. Despite the structure to that space implied by the Hamiltonian, one might conceivably miss the isomorphism to N 3-space particles.

3 Wave-function realism

Let us now consider how we might go about understanding the quantum state. A traditional view is that the state should be understood as somehow being a description or encoding of various classically describable or measurable properties. According to this view, the potential properties of the system are represented

⁴We have found, in conversation, that there are some for whom it is just intuitively obvious that the world isn't like that. But personally speaking, we have no pre-theoretic intuitions at all about what it would be like if the world were a highly non-isotropic 6×10^{23} -dimensional space in which one particle moved, and we wouldn't trust such intuitions if we did.

by projectors on Hilbert space (although typically not all such projectors will correspond conversely to *bona fide* classical properties), and the system determinately possesses a property if (and only if) it is in the subspace projected onto by the associated projector, and determinately does not possess a property if (and only if) it is orthogonal to that subspace. As such, many properties seem to be neither determinately possessed nor determinately not possessed.

This understanding is adequate for interpretations of quantum mechanics in which the theory is simply an algorithm for predicting measurement outcomes; it is adequate for interpretations which abandon classical logic. For more traditionally realist interpretations—in particular for hidden-variable theories like Bohm’s, dynamical-collapse theories like GRW’s, or for the Everett interpretation—it is hopeless. Such theories read the quantum state literally, as itself standing *directly* for a part of the ontology of the theory. To every different quantum state corresponds a different concrete way the world is. For Everett and for some readings of dynamical-collapse theories, the quantum state (perhaps together with some background space or spacetime) gives the *only* part of the ontology.

From this perspective, regarding the state as encoding properties of the system in the traditional way is at best unhelpful and incomplete—many properties, like “being in an entangled state” or “being in some eigenstate of energy” or “possessing an even number of zero amplitudes in configuration-space” cannot be expressed using the traditional approach⁵. Focusing on projectors to represent properties is too crude to capture all of the interesting properties of the world when the quantum state directly describes ways the world is. At its worst, the traditional approach can be actively misleading: for instance, it is responsible for much of the confusion about the “problem of tails” in dynamical-collapse theories and makes obscure the nature of the branches in the Everett interpretation. (These claims are defended *in extenso* in Wallace ([2008])).

In that case, how *should* we think about the quantum state? Mathematically it is a vector (or, strictly, a ray) rotating in a very high-dimensional Hilbert space. But this feels unsatisfactory: the physical universe is, as we have said, very highly structured, whereas Hilbert-space vectors seem pretty much alike. The problem is partially resolved when we recall that a quantum theory can only be specified if we also give some particular set of ‘observables’ (e.g., the position and momentum operators in one-particle QM, or the various field observables in QFT)—this breaks the symmetry of Hilbert space, allowing us to see how merely mathematically distinct vectors could represent physically distinct states of affairs. (In a somewhat similar way, the complicated structure of a classical-mechanical system is lost when that system is represented by a mere point on phase space, until it is recognised that phase space is a highly structured, non-homogeneous manifold.) The state vector can now be seen as a way of codifying the various rich properties of the physical state: in one-particle quantum me-

⁵E.g., the disjunctive property of being in some energy eigenstate or other would, in projectors language, be given by a sum of projectors onto a complete set of energy eigenstates, returning the projector onto the whole Hilbert space; thus every quantum state has this property. Clearly something has gone wrong.

chanics, for instance, to any self-adjoint function $f(\widehat{X}, \widehat{P})$ of the position and momentum operators it assigns an ‘expectation value’ $\langle \psi | f(\widehat{X}, \widehat{P}) | \psi \rangle$. Differences between states can now be understood in terms of their differing patterns of assignment of numbers to operators.

But how should such properties be understood? This is where perplexity sets in. As argued above, anyone who takes the quantum state seriously should avoid definitionally relating it to measurement: whatever “expectation value” means, it had better not mean the mean result of measuring the observable. Yet what other sense might we assign to it?

It is at this point that one may turn to the lessons illustrated in Section 2 for guidance. Can we follow the example that electromagnetism presented and seek to grasp the meaning of the otherwise alien quantum state by relating it to spatiotemporal properties?

At first blush it might seem not: the problem in quantum mechanics is that the quantum state does not trivially have a suitable spacetime representation; in particular, it is not representable as any ordinary sort of spacetime field. But a natural move is available: *wave-function realism*, which has been explored recently by (e.g.) Albert ([1996]), Monton ([2002], [2006]) and Lewis ([2004]). Albert argues, following the lead of Bell ([1981]), that the quantum state, if understood physically at all, should be understood in terms of its configuration-space representation: that is, as a complex-valued field on $3N$ -dimensional space, for an N -particle quantum theory. In this picture we have clear bearers and properties: the bearers are the points of a space at which certain properties—which happen to be represented by a complex number (or vector, if we include internal degrees of freedom)—are instantiated. Understood this way, there seems to be nothing more mysterious about the quantum state than about any other complex-valued field—except that it lives not on spacetime, but on a far-larger space—a space, furthermore, which (just as in the example of phase space considered in section 2) is highly non-homogeneous and non-isotropic, since its dynamical symmetry group is far smaller than $SO(3N)$. This seems a perfectly respectable and intelligible candidate for a quantum ontology. (Vaidman ([2002]) and Barbour ([1999]) describe their respective versions of the Everett interpretation from the perspective of wave-function realism, though they do not explore its implications.)

If wave-function realism is correct (and if it alone, and not some hidden variables, is the physical basis for observed reality), the world is really $3N$ -dimensional at its most fundamental level, and our 3-dimensional world is in some sense emergent from it.

Wave-function realism has not been without its critics, however. Monton ([2002], [2006]), for example, raises the questions of whether the wave-function realist ontology could underwrite requisite facts about three-dimensional objects and their behaviour; and whether it provides a setting in which suitable facts about mental goings-on could obtain.

The argument in (Monton [2002]) focuses on whether or not there is a unique mapping from a) points in the $3N$ -dimensional space on which the field repre-

mented by the wave-function lives, to b) configurations of N three-dimensional objects. Now, of course we agree that *purely mathematically* any number of such mappings can be constructed, but that is not of prime importance. Rather, in our view, the central question is whether the *structure of the field* that the wave-function represents could (both synchronically and diachronically) provide a supervenience base for facts about the three-dimensional objects and spatial relations among them of our experience; and it seems it can. Three-dimensional features will emerge as a consequence of the dynamics (in large part due to the process of decoherence): the form of the Hamiltonian will ensure that the wave-function decomposes into components evolving autonomously according to quasi-classical laws. That is, each of the decohering components will correspond to a system of well-localised (in 3-space) wavepackets for macroscopic degrees of freedom which will evolve according to approximately classical laws displaying the familiar three-dimensional symmetries, for all that they are played out on a higher-dimensional space. Three-dimensional quasi-classical structures are thereby recovered (just as in the case of classical mechanics which we mentioned in section 2) and mental events—should one be especially concerned about them—will find their usual home.

Monton might be concerned that this structure of the field could underwrite *more than one* set of emergent 3-spaces, but even if this were so (which seems highly unlikely⁶) it would mean only that more than one set of emergent 3-spaces existed at the same time, supervening on the wave-function ontology, not that none did. (The pattern of thought here is the same as that applied in decoherence-based solutions to the preferred-basis problem in Everett, cf. Gell-Mann and Hartle ([1990]); Saunders ([1995]); Wallace ([2003]).)

In his 2006, Monton changes tack slightly, arguing that the main problem with wave-function realism is its revisionary nature, rendering false reams of everyday judgements concerning the nature of ordinary objects (e.g., that they are three-dimensional); but again we demur. While the wave-function realist will deny that three-dimensional objects and spatial structures find a place in the fundamental ontology, this is not to say that the three dimensional objects surrounding us, with which we constantly interact, and which we perceive, think and talk about, do not exist; that there are not truths about them; it is just to maintain that they are emergent objects, rather than fundamental ones. But an emergent object is no less real for being emergent. It is also worth keeping in mind that many workers in quantum gravity have long taken seriously the possibility that our four-dimensional spacetime will turn out to be emergent from some underlying reality that is either higher-dimensional (as in the case of string theory⁷) or not spatiotemporal at all (as in the case of loop quantum

⁶Permuting the 1-dimensional variables in the expression of the wave-function into different groupings of three, as Monton envisages, may well not give rise to structures that are diachronically isomorphic to three-dimensional objects and relations. The ensuing structures certainly wouldn't obey the empirically correct quasi-classical laws.

⁷Popular accounts of string theory can sometimes give the impression that the emergence of four-dimensional spacetime from an underlying 11-dimensional space is purely a matter of coarse-graining, of ignoring sufficiently small distances; the actual physics is significantly

gravity). In neither case is it suggested that ordinary spacetime is *nonexistent*, just that it is *emergent*.

4 Problems with wave-function realism

But if wavefunction realism is a coherent metaphysical possibility, we feel it nonetheless sits uneasily with real physics. One of our worries can be understood even in the context of nonrelativistic quantum mechanics (NRQM): namely, physicists normally regard all bases as on a par (so that the configuration-space representation of the quantum state is just one representation among many); they also regard various different formulations of the dynamics (Heisenberg, Schrödinger, interaction-picture) as on a par. Both of these are violated by wave-function realism, which gives the Schrödinger form of the dynamics, and the position representation, a special status.⁸

To be fair, a case can be made in NRQM that both position-basis formulations of the theory, and the Schrödinger picture, do indeed have a special status; indeed, many students first learn NRQM in this form, and move on to the Hilbert-space formalism only later. If NRQM were the only quantum theory with which we were concerned, it would *perhaps* not be unreasonable to regard the configuration-space formalism as primary, and the Hilbert-space alternative just as a mathematical curiosity.

This brings us to quantum field theory (QFT), where things are not so straightforward. In our view, two related worries about QFT make wave-function realism an unattractive position at best, an unviable one at worst:

1. In quantum field theory, the particles are not fundamental and their positions are imprecisely defined; as such, no really satisfactory notion of “configuration space” is available for us to formulate wave-function realism.
2. Wave-function realism obscures the role of spacetime in the formulation of quantum theory.

(A third problem, how wave-function realism deals with nonlocality and relativistic covariance, will be postponed until Section 6 and following).

more complex, however, and makes essential use of dynamical principles. (For instance, it has been argued that the extra dimensions might have diameters up to several centimetres; and we certainly don't coarse-grain that much!) The correct criterion for whether a given higher-dimensional theory is empirically adequate is, instead, whether it can be approximated at low energies by an effective field theory which can be expressed on a four-dimensional background. We are grateful to an anonymous referee for pressing us on this point.

⁸Is this just a sociological observation? Not in the relevant sense. Of course, we do not want to claim that the philosophical commitments of physicists—even those demonstrated in their practice as well as their own philosophising—are *infallible* (the measurement problem is an obvious counter-example). But equally, we feel that philosophers of a naturalistic bent should be cautious in arguing for the metaphysical priority of any formulation of a physical theory that is unnatural by physics' own standards.

To begin with the first problem: in QFT, particle number is not conserved. If there is a “configuration-space” representation of the quantum state, it is given by assigning a (non-normalised) wave-function to each of the infinitely many $3N$ -dimensional configuration spaces. Wave-function realism now gives us a picture, not of a single field undulating in a high-dimensional space, but of an infinite number of such fields, each undulating in its own high-dimensional space and each interacting with the others.

Is this picture problematic? It is certainly unintuitive; but so was wave-function realism itself, and we have been at pains to stress that this is no criticism of it. But things are actually rather worse than this, in fact. For particles are not only non-conserved in QFT, they are non-fundamental: mathematically speaking they are emergent entities supervenient on an underlying field ontology. In fact, the properties of these entities (mass and charge, for instance) do not take precisely-defined values, but must be chosen according to the particular situation under consideration (see Wallace ([2006]) for a more detailed discussion). This makes particle configurations unattractive—technically as well as conceptually—as the basis for defining the ontology of QFT.

A superficially attractive alternative is to use a basis of *field* configurations: in vacuum quantum electrodynamics, for instance, the wave-function would assign a complex number to each possible state of the magnetic field. Such a wave-function would be defined on an infinite-dimensional space⁹ but at least we would only have one such wave-function.

The problem with this move is that there is no single preferred choice of fields by which a QFT can be specified. A number of results from QFT (from formal results like Borchers equivalence (Haag [1996]) to concrete models like the Sine-Gordon equation (Coleman [1985], pp. 246–252)) suggest that a single QFT can be equivalently described in terms of several different choices of field observable, with nothing in particular to choose between them.

We might sum up the objection thus: wave-function realism requires a metaphysically preferred basis. In NRQM, there is a single natural choice of such basis: the configuration-space basis is conceptually natural and dynamically special. There is no such natural choice in QFT: many different bases are dynamically on a par and the “natural” choice is situation-dependent. This makes any particular metaphysical preference appear somewhat ad hoc.

(This objection is probably most significant for Everettians, who generally regard it as a virtue of their preferred interpretation that it requires no additional formalism, and so are unlikely to look kindly on a requirement in the metaphysics for additional formalism. Advocates of dynamical-collapse and hidden-variable theories are already committed to adding additional formalism, and in fact run into problems in QFT for rather similar reasons: there is no longer a natural choice of basis to use in defining the collapse mechanism or the hidden variables. We are not ourselves sanguine about the prospects of

⁹We ignore details of renormalisation; if we regard the high-energy cutoff as real (as advocated by Wallace [2006]) and if the universe is spatially closed then the dimension of the space will be finite but extremely large.

overcoming this problem¹⁰; but if it *were* to be overcome, the solution might well also suggest a metaphysically preferred basis to use in formulating a QFT version of wave-function realism.)

What of the second objection? We stress that we have no *conceptual* objection to a view of spacetime on which it is emergent from some more fundamental theory; rather, we do not find this a good account of spacetime in physics as we find it. Whilst it is true that spacetime in nonrelativistic quantum mechanics is somewhat elusive, it is mathematically fundamental in QFT: the dynamical variables in QFT are explicitly associated with spacetime regions.

At least heuristically, these dynamical variables can be thought of as being built up from integrals over products of field operators located at a point, such as

$$\widehat{O} = \int_{\Delta} dx \widehat{\phi}(x) \widehat{\phi}(x) \quad (1)$$

where Δ is some spacetime region. (The point of Borchers equivalence is that \widehat{O} could be expressed in terms of different field operators whilst still being associated with the same region Δ .)

The objection is sharper still in the case of semiclassical quantum gravity (see, for instance, Wald [1994]). Here we define the quantum fields on a *curved* spacetime background and require that background to satisfy the classical Einstein field equations with respect to the expected distribution of stress-energy. Such models are used to address issues like black hole evaporation and early-universe cosmology; there are sound theoretical reasons to regard them as an appropriate limiting case to full quantum gravity in some regimes, though there is very little direct empirical evidence for their accuracy. It is opaque to us how semiclassical quantum gravity can be understood except via a quantum ontology which treats spacetime as fundamental.

If wave-function realism were “the only game in town”, these objections would have to be overcome somehow; by biting various bullets, perhaps they can be. But in fact an alternative ontology is available, one which respects the democracy of Hilbert-space bases and which gives a more central role to spacetime. This alternative, we believe, bears much the same relation to wave-function realism as the N -particles-in-space view of classical mechanics bears to the one-particle-in-phase-space view. We present it in the next section.

5 Spacetime state realism

Wave-function realism presented a clear ontology by preferring a particular set of operators and finding a spatial arena by enlarging the fundamental physical space from three dimensions to $3N$ dimensions; thereby providing sufficient property bearers. If we do not wish to prefer a basis, then we should stick to the characterisation of the quantum state of a system as a (positive, normalised) linear functional of the dynamical variables; that is, mathematically, as a density

¹⁰Extant proposals for Bohmian QFTs, in particular, seem to us to pay insufficient attention to the problem of renormalisation, though a full discussion lies beyond the scope of this article.

operator on the Hilbert space of the system. To go from this to a candidate ontology, though, we still require both clear bearers of properties and a clear view of the structure that a quantum state would ascribe to the world. We propose that what is crucial is an analysis of the total system—the Universe—into subsystems.

Suppose one were to assume that the Universe could naturally be divided into subsystems; assign to each subsystem a density operator¹¹. We then have a large number of bearers of properties—the subsystems—and the density operator assigned to each represents the intrinsic properties that each subsystem instantiates, just as the field value assigned to each spacetime point in electromagnetism, or the complex number assigned to each point in wave-function realism, represented intrinsic properties. While the property that having a given density operator represents may not be a familiar one, the case need be no different in principle from that of electromagnetism. In so far as one can continue to press for the physical meaning of the density operator, the theory in which these objects are postulated must provide the answer. (We can say such things, for example, as: the property is not a scalar one, in contrast to wave-function realism. We sketch how the relation to the macroscopic goings-on of our experience would be developed in Section 9.) To provide a simple model, imagine a Universe consisting of a great many interacting qubits (cf. Deutsch and Hayden ([2000]) or von Weizsacker’s *ur*-alternatives (von Weizsacker [2006])). The qubits each bear the property or properties represented by their two dimensional density operator; pairings of qubits bear properties represented by a four dimensional operator; and so on. There need be no reason to blanch at an ontology merely because the basic properties are represented by such objects: we know of no rule of segregation which states that only those mathematical items to which one is introduced sufficiently early on in the schoolroom get to count as possible representatives of physical quantities, for example!

Note that if, by contrast, we were to treat the Universe just as one big system, with no subsystem decomposition, then we would only have a single property bearer (the Universe as a whole) instantiating a single property (represented by the Universal density operator) and we would lack sufficient articulation to make clear physical meaning of what was presented (as with the one-particle-in-high-dimensional-space view of classical mechanics, one would struggle to see the structure being imputed to the world in this case).

Now, our proposal becomes more concrete when one connects this system-subsystem analysis of the quantum state with spacetime. To do this we need some notion of the spatial location of a physical system. In NRQM this is somewhat complicated: the systems are naturally taken to be particles and assemblies of particles, and a particle’s spatial location is one of its dynamical properties, not something to be specified *ab initio*.

It is somewhat simpler if we consider the Fock-space formalism of NRQM, where we allow the number of particles to vary.¹² On the one-particle Hilbert

¹¹We intend that if A and B are both subsystems of the Universe then $A \cup B$ is too. On the significance of this, see Section 6.

¹²Recall that if \mathcal{H} is any Hilbert space, the associated symmetric (bosonic) and antisym-

space of any non-relativistic particle we can define projectors \widehat{P}_Δ projecting onto those states with wave-functions having support in spatial region Δ . If $\{\Delta_i\}$ is a partition of real space into measurable subsets and if $\mathcal{H}_i = \widehat{P}_{\Delta_i}\mathcal{H}$ is the Hilbert space built from all states having support in Δ_i , then we have $\mathcal{H} = \oplus_i \mathcal{H}_i$ and moreover

$$\mathcal{F}(\mathcal{H}) = \mathcal{F}(\oplus_i \mathcal{H}_i) = \otimes_i \mathcal{F}(\mathcal{H}_i), \quad (2)$$

so the Hilbert spaces $\mathcal{F}(\mathcal{H}_i)$ may be taken as representing the possible states of the subsystem in, or comprising, region Δ_i .¹³ This means that we can take the regions of space (and their unions) as our basic bearers of properties: tensor products of states belonging to each region (in general, superpositions of such products) allow us to express the original total state of varying particle number.

Already our presentation is sounding somewhat field-theoretic. To illustrate what is going on, consider a particular region of space Δ_j . This region has a Fock space $\mathcal{F}(\mathcal{H}_j)$ whose (pure) basis states can be represented in the form $|n_1, n_2, \dots\rangle$, where n_1, n_2 etc. represent the occupation numbers of what we can think of as the available modes of Δ_j , that is, the number of excitations in each of some orthogonal set of states of \mathcal{H}_j . What we would normally think of in NRQM as a single particle localised in Δ_j will, in this setting, be represented by a singly excited state of $\mathcal{F}(\mathcal{H}_j)$, e.g. $|1, 0, 0 \dots\rangle$, tensor product with the vacuum state $|0\rangle$ for all the other regions' Fock spaces. In general, then, a single particle (which usually *won't* be localised in some particular region) will be represented by an entangled state composed of a superposition of states each differing from the vacuum only in a small region Δ_i :

$$\dots |1, 0, \dots\rangle|0\rangle|0\rangle \dots + \dots |0\rangle|1, 0 \dots\rangle|0\rangle \dots + \dots |0\rangle|0\rangle|1, 0, \dots\rangle \dots$$

and so on.

Things become simpler still when we move to full quantum field theory. In the algebraic formulation of QFT, we associate to each spacetime region \mathcal{R} a C^* -algebra $\mathcal{A}(\mathcal{R})$ of operators, representing the dynamical variables associated

metric (fermionic) Fock spaces are

$$\mathcal{F}_S(\mathcal{H}) = \oplus_{i=0}^{\infty} \mathcal{S}(\otimes^i \mathcal{H})$$

and

$$\mathcal{F}_A(\mathcal{H}) = \oplus_{i=0}^{\infty} \mathcal{A}(\otimes^i \mathcal{H})$$

where \mathcal{S} and \mathcal{A} are symmetrisation and antisymmetrisation operators. The distinction between $\mathcal{F}_S(\mathcal{H})$ and $\mathcal{F}_A(\mathcal{H})$ plays no part in the argument, and will be ignored in the text.

¹³Intuitive support for the result expressed in eqn. 2 can be seen once we recognise that the Fock-space operation is a sort of 'exponential' of the Hilbert space (this is clearest from the power-series expansion of the exponential), so that we can write $\mathcal{F}(\mathcal{H}) = \exp(\mathcal{H})$. Then equation 2 becomes $\exp(\oplus_i \mathcal{H}_i) = \otimes_i \exp(\mathcal{H}_i)$. More formally, for each i let $\widehat{a}_{i,k}^\dagger$ be a set of creation operators for states in \mathcal{H}_i . Then it is easy to see (via the observation that creation operators on different Hilbert spaces commute) that states of form

$$\widehat{a}_{i_1, k_1}^\dagger \dots \widehat{a}_{i_n, k_n}^\dagger |\Omega\rangle$$

(where $|\Omega\rangle$ is the vacuum) form a basis both for $\otimes_i \mathcal{F}(\mathcal{H}_i)$ and for $\mathcal{F}(\oplus_i \mathcal{H}_i)$.

to region \mathcal{R} . A state ρ of such a region is a positive linear functional on $\mathcal{A}(\mathcal{R})$ (often described in rather instrumentalist terms as giving the expectation value of each element of $\mathcal{A}(\mathcal{R})$) and by the Gelfand-Naimark-Segal construction we can associate ρ with a state in a Hilbert space \mathcal{H}_R , and represent $\mathcal{A}(\mathcal{R})$ as an algebra of operators on \mathcal{H}_R (see, e. g. , Haag ([1996], pp.122-124) for the details). \mathcal{H}_R can then be taken as the Hilbert space of the field in region \mathcal{R} .¹⁴ If preferred, one can even remain at the more abstract level, forego the representation theorems and just take the C^* -algebraic state itself as denoting the properties of a region.

This alternative ontology avoids the problems previously identified for wave-function realism: it is well-defined for any quantum theory with compositional structure (in particular, for any quantum field theory, and for any many-particle theory once it is expressed in field-theoretic terms); it respects the dynamical structure of QM, indicating no preference for Schrödinger over Heisenberg or interaction dynamics (as the state is just construed as a linear functional of the dynamical variables); it adds no additional interpretational structure (given that the compositional structure of the system is, *ex hypothesi*, already contained within the formalism); and it gives an appropriately central role to spacetime. For want of a better name, we call it *spacetime state realism*.¹⁵

6 Locality

How is one to visualise the quantum state, according to our proposal? Recall the discussion of Section 2: typically in classical physics we understand physical systems via the properties they ascribe to spacetime regions. Wave-function realism keeps this basic framework at the cost of greatly enlarging the background spacetime: the quantum state is a complex field much like the classical Klein-Gordon field, but it is defined on an extremely high-dimensional space. Our proposal is in one sense closer to the classical situation: we associate a set of properties (represented by a density operator) to each region of spacetime. In another sense, though, it is less classical than wave-function realism: it displays a certain sort of nonlocality.

To see this, let A and B be spacelike separated regions. The universal quantum state, like a classical field, associates a state to A , B , and $A \cup B$. But there is an important disanalogy. Classically, the separate states of A and B completely determine the state of $A \cup B$ —once the electromagnetic field, say, is known in A and B , it is known also in $A \cup B$. Healey ([1991]) calls this property *separability*: the principle that the properties of a spacetime region supervene on the properties of its spatiotemporal parts (see also (Earman [1987]) for a

¹⁴In the standard presentations of AQFT, the algebra $\mathcal{A}(\mathcal{R})$ is infinite-dimensional; as such it has multiple non-isomorphic representations, and so different states lead to different Hilbert spaces. This makes it less clear that we are licensed to talk about \mathcal{H}_R as “the” Hilbert space for region \mathcal{R} . One of us has argued elsewhere, however (Wallace [2006]) that this is an artefact of the formalism, which disappears when we properly understand the renormalisation process; as such, we ignore this complication in the text.

¹⁵Although if the theory has a compositional structure not induced by spacetime regions—as in, say, an abstractly specified quantum computer—this name is a misnomer.

discussion of various senses of locality). We might equally call it *Humean supervenience*, after Lewis’s thesis—which he attributes to Hume—that all properties of the world supervene on properties of spacetime points and on the relations between those points (Lewis [1986], pp.ix–xvii).

Thanks to the phenomenon of entanglement, separability is simply false for the quantum state according to spacetime state realism: if the state $\hat{\rho}_{A \cup B}$ is known, then via the partial trace operation we can learn the states of A and B , but of course the converse is not so.

We should not be surprised, perhaps: quantum mechanics is nonlocal; everyone knows that, don’t they? But notice that the description given of nonlocality is very different according to wave-function realism. There, the state *on configuration space* is completely separable: specify a complex number at each point and you specify the state. According to wave-function realism, non-separability on three-dimensional space is an illusion, or better, an emergent property—it is a consequence of our attempt to use three-dimensional language to describe an entity which fundamentally lives in $3N$ -dimensional space. By contrast, in spacetime state realism non-separability is fundamental.

Is this a reason in favour of wave-function realism (setting aside our previous arguments against its tenability)? Not in our opinion, for three reasons:

1. As nonlocal forms of behaviour go, non-separability is fairly mild. It does not imply any sort of action at a distance: the quantum state of spacetime region A is dynamically determined by the state of its past light cone (more precisely: by the state of any spacelike slice of its past light cone). The state of $A \cup B$ may indeed be changed by operations in the vicinity of either A or B , but the state of B is unaffected by operations performed at A . (Note also that it is milder than the nonlocality predicted by the Bell inequalities: those concern changes made to the local state at A as a consequence of operations performed at B . Since we are considering only unitary quantum mechanics—and so, by implication, presuming the Everett interpretation—some of the premises of Bell’s theorem aren’t in play. (See e.g. Timpson and Brown ([2002]) for a discussion of Bell’s theorem and locality in Everett.)

We are, in fact, unmoved by purely *metaphysical* objections to action at a distance. Our concern about it stems from its tension with relativistic covariance. But there is no such tension in the case of non-separability.

2. It is tempting to regard separability as part of our ordinary conception of space: arguably, if some putative spacetime has essentially nonlocal properties, or perhaps better, if the things in it (e.g., fields) have to end up being attributed non-separable properties, we ought not to call the arena “spacetime”.¹⁶ But there is nothing *mathematically* improper about these

¹⁶Essentially this idea may be found in Einstein: “It is...characteristic of...physical objects that they are thought of as arranged in a space-time continuum. An essential aspect of this arrangement of things in physics is that they lay claim, at a certain time, to an existence independent of one another, provided these objects ‘are situated in different parts of space’.

nonlocal properties. And even according to wave-function realism, the four-dimensional spacetime we live in (an emergent object, supervenient on the wave-function) itself involves non-separability in exactly the way described by spacetime state realism. The two ontologies are not in disagreement about whether ordinary spacetime allows *separability*; they are in disagreement about whether this spacetime is *fundamental*. And if our paradigm example of a space involves non-separability, we see no plausible argument that our conception of space *requires* separability. Rather, we have a mere intuition that it does—an intuition developed because space *looks* separable, unless you look very carefully.

3. In any case, there are respectable *classical* examples of non-separability. In particular, the Aharonov-Bohm effect is best understood as implying that electromagnetism is non-separable: specifying the connection on regions A and B does not suffice to determine it on $A \cup B$ unless the latter region is simply connected. If non-separability is reasonable for the A-B effect, why not for the quantum state?¹⁷

(Of course, the implications of the A-B effect are not uncontroversial; one might as easily argue that non-separability is unreasonable in both cases. We expand on our preferred reading of the A-B effect, and on the analogies between that effect and quantum non-locality, in Wallace and Timpson ([2009])).

In fact, once we consider the relativistic domain there are active *advantages* to spacetime state realism’s treatment of nonlocality, as we will see in Section 8. First, however, it will be useful to consider how the spacetime state ontology fares in the case of non-unitary dynamics.

7 Non-unitary dynamics on the spacetime ontology

So far we have been considering quantum states evolving under purely *unitary* dynamics. There are, however, some reasons to be interested in adding non-unitary processes to the dynamics—specifically, if we are unhappy with the Everett interpretation and with hidden-variable theories, we may wish to regard wave-function collapse as a real microphysical process and to model it non-unitarily.

Unless one makes this kind of assumption about the independence of the existence (the ‘being thus’) of objects which are far apart from one another in space...physical thinking in the familiar sense would not be possible.” (Einstein [1948])

¹⁷Why do we call the A-B effect classical? Because mathematically it is a classical consequence of the interactions between a $U(1)$ connection and a complex field. Inconveniently, the complex fields we see around us (the proton and electron fields) have no accessible classical limit, so we observe its effect on the quanta of those fields, but the effect is not *essentially* quantum-mechanical. See Wallace ([2009]) for more on this point.

It is reasonable to assume that wave-function collapse must be implemented by a non-Lorentz-covariant dynamics, for familiar reasons: if one prepares two particles in an entangled state, separates them widely, and measures one of them, the other is supposed to collapse *immediately*—and “immediately” seems to imply a preferred frame. But there is an alternative tradition (defended at various times by, e.g., Aharonov and Albert ([1980]), Fleming ([1988]), Myrvold ([2002], [2003]), and Tumulka ([2006])): *foliation-relative collapse* (FRC). According to FRC, a collapse event should be modelled as occurring not *at a point* but *on a hypersurface*.

It is difficult to make a really thorough evaluation of the merits of FRC, since it remains a framework only: there is at present no satisfactorily worked-through relativistic collapse theory, let alone one which reproduces the predictions of QFT. Yet two things seem clear. Firstly, FRC does seem to offer, in principle, a framework for a covariant quantum theory which nonetheless incorporates dynamical collapse. Secondly, it has some extremely counter-intuitive consequences.

To illustrate this latter point, suppose that we prepare two widely-separated particles in an entangled state (see Figure 1), so that on surface Λ_1 the combined state is

$$|\psi_1\rangle = \alpha |0\rangle |1\rangle + \beta |1\rangle |0\rangle; \quad (3)$$

and we go on to measure the spins of both of them (with the measurement on particle i occurring at point X_i).

If we analyse the situation using a foliation containing Λ_1 , Λ_2 and Λ_4 —that is, on a foliation in which particle 1 is measured first and particle 2 second—the history of the system will be something like: first there is a stochastic process (between Λ_1 and Λ_2) caused by the measurement of particle 1, in which particle 1 collapses into (say) $|0\rangle$ and particle 2 into $|1\rangle$; then there is a deterministic process (between Λ_2 and Λ_4) where we passively determine that the state of particle 2 is indeed in state $|1\rangle$. However, the history looks very different on a foliation containing Λ_1 , Λ_3 and Λ_4 , so that 2 is measured before 1. According to this second foliation, the stochastic process (between Λ_1 and Λ_3) is triggered by the measurement of particle 2, and it is the measurement of particle 1 (between Λ_3 and Λ_4) that is the passive observation.

Note, however, that the *observed consequences* of the two measurements are the same (particle 1 has spin up; particle 2 has spin down), and that the end state on Λ_4 is the same in both foliations,

$$|\psi_4\rangle = |0\rangle |1\rangle. \quad (4)$$

This makes it a somewhat subtle business to say just what, if anything, is unacceptable about FRC. Indeed, some have argued that in fact there is nothing wrong with it, and that its consequences—though indeed counter-intuitive—are inevitable given relativity and non-locality, and should not be regarded as unacceptable. (See Myrvold ([2002], [2003]) and references therein.)

We are unpersuaded, but we do not want to defend this here. Our goal in this paper is more modest: to use the spacetime state framework to clarify the

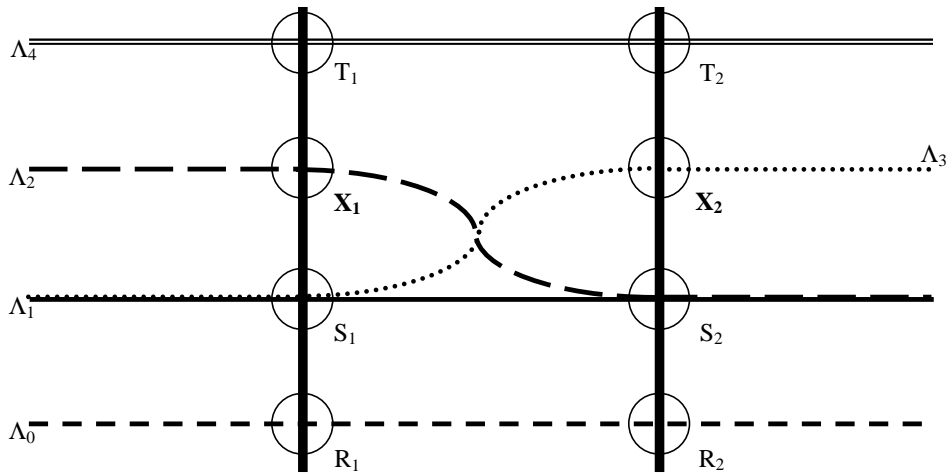


Figure 1: A spacetime diagram of the histories of two particles; time runs vertically. Λ_0 to Λ_4 label distinct hypersurfaces of simultaneity. Particles 1 and 2 are spacelike separated and begin (on Λ_0) in an entangled state. Regions R_j, S_j, T_j lie on the worldline of particle j (1 or 2). At X_1 , particle 1 is measured; at X_2 , particle 2 is measured. The story one tells about the effects of these measurements is very different depending on whether one uses a foliation in which X_1 precedes X_2 ($\{\Lambda_0, \Lambda_1, \Lambda_2, \Lambda_4\}$) or in which X_2 precedes X_1 ($\{\Lambda_0, \Lambda_1, \Lambda_3, \Lambda_4\}$). Note that in the spacetime state framework, we may replace talk of states of particles relative to elements of a foliation with direct talk of the states of spacetime regions, such as R_j, S_j and T_j .

sense in which non-unitary dynamics has consequences going beyond those of unitary dynamics. To this end we shall now allow ourselves to talk directly of the states of the pertinent spacetime regions shown in Figure 1; considering first the unitary case, before seeing how things differ under non-unitary dynamics.

Thus suppose that in Figure 1 the transformations performed at X_1 and X_2 were unitary—say,

$$\begin{aligned} |0\rangle &\longrightarrow |1\rangle \\ |1\rangle &\longrightarrow -|0\rangle; \end{aligned} \tag{5}$$

a spin-flip rather than a measurement. The states of the various relevant regions are easily calculated:

$$\begin{aligned} \widehat{\rho}_{R_1} &= \widehat{\rho}_{S_1} = |\alpha|^2 |0\rangle\langle 0| + |\beta|^2 |1\rangle\langle 1| \\ \widehat{\rho}_{R_2} &= \widehat{\rho}_{S_2} = |\beta|^2 |0\rangle\langle 0| + |\alpha|^2 |1\rangle\langle 1| \\ \widehat{\rho}_{T_1} &= |\beta|^2 |0\rangle\langle 0| + |\alpha|^2 |1\rangle\langle 1| \\ \widehat{\rho}_{T_2} &= |\alpha|^2 |0\rangle\langle 0| + |\beta|^2 |1\rangle\langle 1| \end{aligned}$$

and $\widehat{\rho}_{\Lambda_i} = |\psi_i\rangle\langle\psi_i|$ where

$$\begin{aligned} |\psi_1\rangle &= \alpha |0\rangle |1\rangle + \beta |1\rangle |0\rangle \\ |\psi_2\rangle &= \alpha |1\rangle |1\rangle - \beta |0\rangle |0\rangle \\ |\psi_3\rangle &= -\alpha |0\rangle |0\rangle + \beta |1\rangle |1\rangle \\ |\psi_4\rangle &= -\alpha |1\rangle |0\rangle - \beta |0\rangle |1\rangle. \end{aligned}$$

We observe that, as expected, the state on both Λ_2 and Λ_3 is changed from Λ_1 , since either X_1 or X_2 lies in the past light cone of each region. The states of the individual particles, however, are unambiguously fixed by the dynamics, and neither particle's *individual* state is affected by a transformation outside its light cone — X_1 lies outside the past light cone of S_2 , for instance, so $\widehat{\rho}_{S_2} = \widehat{\rho}_{R_2}$ even though S_2 is part of Λ_2 and the state of Λ_2 is changed by X_1 .

Things are otherwise in the event of collapse, where we make a measurement rather than a spin flip at X_i . The (hypothetical) stochastic dynamics produce perfectly reasonable (non-deterministic) predictions for all the joint states—assigning a non-zero probability to, say, $\widehat{\rho}_{\Lambda_i} = |\phi_i\rangle\langle\phi_i|$ where

$$\begin{aligned} |\phi_1\rangle &= \alpha |0\rangle |1\rangle + \beta |1\rangle |0\rangle \\ |\phi_2\rangle &= |\phi_3\rangle = |\phi_4\rangle = |0\rangle |1\rangle. \end{aligned}$$

But there is no unique prescription for the one-particle states. Take S_2 , for instance—if it is regarded as part of Λ_1 then it has state

$$|\beta|^2 |0\rangle\langle 0| + |\alpha|^2 |1\rangle\langle 1|; \tag{6}$$

if it is regarded as part of Λ_2 then it has state $|1\rangle\langle 1|$. We might interpret this (at the risk of confusion with a terminology adopted elsewhere in philosophy of QM) as a *contextuality* of local states—the state of a region depends on which larger region it is regarded as a part of. We might also interpret it as denying the existence of states for many regions of spacetime—regions only have states at all if they are assigned the same state by all foliations. We might call this *nihilism* about local states, a name which will be particularly apt if, as we are inclined to suspect (but have not proved), in realistic models virtually no regions smaller than entire foliations are assigned states.

To conclude this section: there is an “hierarchy of counterintuitiveness” here. We have

1. **Full locality** (classical physics). All spacetime regions have states; the state of a region supervenes on the states of its subregions.
2. **Non-separability** (unitary quantum physics). All spacetime regions have states but the state of a region does not supervene on the states of its subregions.
- 3a. **Contextuality** (non-unitary quantum mechanics, first interpretation). Spacetime regions smaller than entire hypersurfaces have states which are dependent on which larger region they are regarded as part of.
- 3b. **Nihilism** (non-unitary quantum mechanics, second interpretation). Some or all spacetime regions smaller than entire hypersurfaces have no states at all.

8 Relativity and Poincaré covariance

Relativistic covariance is straightforward to understand according to spacetime state realism: if A is a spacetime region and \mathcal{L} is a Poincaré transformation (represented on Hilbert space by some operator $\widehat{U}_{\mathcal{L}}$) then the state of region A after performing \mathcal{L} as a global active transformation is determined by the state of $\mathcal{L}^{-1}A$ before the transformation. This is easy to see: the state of A is just given by the expectation value of all operators localised in A . So after the active transformation, to specify the state we need terms like

$$\langle\psi|\widehat{U}_{\mathcal{L}}^{\dagger}\widehat{\Phi}(x)\widehat{\Phi}(y)\widehat{U}_{\mathcal{L}}|\psi\rangle \quad (7)$$

where x and y are points in A . But if $\widehat{\Phi}$ is a scalar field, by the construction of a QFT we have

$$\widehat{U}_{\mathcal{L}}^{\dagger}\widehat{\Phi}(x)\widehat{U}_{\mathcal{L}} = \widehat{\Phi}(\mathcal{L}^{-1}x), \quad (8)$$

and so

$$\langle\psi|\widehat{U}_{\mathcal{L}}^{\dagger}\widehat{\Phi}(x)\widehat{\Phi}(y)\widehat{U}_{\mathcal{L}}|\psi\rangle = \langle\psi|\widehat{\Phi}(\mathcal{L}^{-1}x)\widehat{\Phi}(\mathcal{L}^{-1}y)|\psi\rangle; \quad (9)$$

considering vector or spinor fields adds only technical complexity.

The situation is rather different for wave-function realism. There is a sense in which the theory is Poincaré-covariant—namely, that there is a representation of the Poincaré group as unitary operators on the space of wave-functions on the configuration space¹⁸ such that a sequence of wave-functions $\Psi(t)$ satisfy the Schrödinger equation only if the sequence $\widehat{U}_{\mathcal{L}}\Psi(t)$ does for all \mathcal{L} . (Put another way: the Schrödinger equation is form-invariant under the action of this representation of the Poincaré group.) But since the wave-function is a function on a space of dimension far higher than three, it is somewhat opaque what the connections are between this result and the usual interpretation of the Poincaré group in terms of the symmetries of Minkowski spacetime¹⁹.

The problem may be sharpened via a remarkable observation due to David Albert (Albert [2007]). If wave-function realism is correct, then we might naturally expect that the entire history of the universe should be given by the sequence $\Psi(t)$ for all t . As in the classical case, an alternative sequence like $\widehat{U}_{\mathcal{L}}\Psi(t)$ should be a mere redescription of the first sequence. But in fact, this is not the case, as can be shown by a simple example. Consider again the situation of Figure 1: two distinguishable spin-half particles 1 and 2 are prepared in a singlet state²⁰

$$|\psi\rangle = \frac{1}{\sqrt{2}}(|0\rangle|1\rangle - |1\rangle|0\rangle). \quad (10)$$

We now consider:

Scenario A: The spin degrees of freedom of the particles remain unaltered for all time.

Scenario B: At X_1 and X_2 we perform a spin flip operation on 1 and 2 respectively:

$$\begin{aligned} |0\rangle &\longrightarrow |1\rangle \\ |1\rangle &\longrightarrow -|0\rangle. \end{aligned} \quad (11)$$

If we describe the evolution of the state in a foliation where X_1 and X_2 are simultaneous (one containing $\Lambda_0, \Lambda_1, \Lambda_4$, say), then scenarios A and B have the same effect on the quantum state: nothing. That is, the state remains a singlet at all times.

This is not the case in other foliations. Consider a foliation where X_1 precedes X_2 (i.e., we include Λ_2 of Figure 1). In scenario A, of course, the state

¹⁸Those readers who recall our criticisms of wave-function realism in Section 4 may reasonably ask: which configuration space? Two natural answers are available: if we consider a fixed number of particles which do not interact, we can use ‘ordinary’ N -particle quantum mechanics; more generally we can use a field configuration space.

¹⁹A similar concern is raised by Lewis ([2004]), who queries the apparent mismatch between the symmetry of the Hamiltonian and the structure of the space on which the wavefunction is defined in NRQM.

²⁰This is a simplified version of Albert’s own example, which involves four particles. Albert’s example, however, has the advantage of being manifestly covariant, demonstrating that the non-covariance of Scenario B is not important.

remains unchanged at all times. But in scenario B, there is a period during which 1 but not 2 has undergone the spin flip, leading to a global state

$$|\phi^+\rangle = \frac{1}{\sqrt{2}}(|1\rangle|1\rangle + |0\rangle|0\rangle). \quad (12)$$

It follows from this that the sequence $|\Psi(t)\rangle$ does not actually fix the Lorentz-transformed sequence: two worlds may be described by the same sequence but have different Hamiltonians, and so will have different descriptions following a Lorentz transformation. Two observations immediately follow:

1. For a Poincaré-covariant quantum theory, the sequence of states $\widehat{U}_{\mathcal{L}}|\Psi(t)\rangle$ cannot be interpreted as a mere redescription of the sequence $|\Psi(t)\rangle$.
2. For a Poincaré-covariant quantum theory, the sequence of states $|\Psi(t)\rangle$ cannot be regarded as completely specifying the properties of the system.

Albert calls a theory *narratable* if specifying a system's state at all times is sufficient to specify all properties of a system. Poincaré-covariant quantum mechanics is not narratable: if we give the state at all times on a given foliation, we have given something less than the complete description of the system.²¹

One of the common arguments used in favour of the Everett interpretation over other interpretations is that it is fully compatible with special relativity. The conclusion Albert draws from narratability failure is that this presumed advantage is overstated. On the one hand, as we have seen in Section 7, there are Lorentz-covariant frameworks for dynamical-collapse theories—they have counter-intuitive consequences, but Albert interprets these as special cases of the failure of narratability. If he is correct, then if narratability failure is also acceptable in the Everett interpretation, why not in dynamical-collapse theories?

On the other hand, if narratability failure is *not* acceptable, the only alternative is to give up on Lorentz covariance as fundamental and accept a preferred (albeit undetectable) foliation. But if *this* is acceptable in the Everett interpretation, why not in other interpretations (notably, hidden-variable theories like the Bohm theory, or non-covariant collapse theories)?

We reject the first horn of Albert's dilemma. As we have tried to show in Section 7, the consequences of covariant non-unitary dynamics are relevantly worse than those of covariant unitary dynamics. But we wish to draw a different conclusion from Albert's argument: namely, that it provides a further argument for spacetime state realism over wave-function realism.

To elaborate: according to wave-function realism the world is fundamentally local, but the dimensionality of space is very large ($3N$ -dimensional)—the state of the universe at an instant is represented by a complex function on that space. Then it is hard to see how to think about the entire history of the universe other

²¹The phenomenon of narratability failure in Poincaré-covariant quantum theories relates to the fact that, unlike in the Galilean case, the operators generating the Lorentz-boosts depend on the Hamiltonian of the system; see for example (Fleming [2002]). We thank Jeremy Butterfield for the observation and Wayne Myrvold for the reference.

than as a complex function on a $3N + 1$ -dimensional spacetime. And *that* is what narratability failure rules out.

On the other hand, spacetime state realism is fundamentally *non*-local (non-separable, to be more exact). However, suppose that we have *any* spacetime theory which is (a) non-separable, so that there can be simultaneous spacetime regions A and B such that the state of $A \cup B$ is not determined by the states of A and B separately, and (b) is also covariant. Covariance entails that there can also be *non*-simultaneous spacetime regions whose joint state is not fixed by their separate states. This opens up the possibility of failure of narratability: specification of global states on elements of one foliation on their own will not in general fix the joint states of non-simultaneous regions (regions which do not lie on some single hypersurface in the foliation), hence will not fix the global states on another foliation (as the latter foliation will involve linking up hypersurfaces which are non-simultaneous, according to the first).²²

From this perspective, we can explain any initial surprise at narratability failure as due to our failure to take non-separability sufficiently seriously. In the Good Old Days of classical physics, we thought that to specify the state of a spacetime region it sufficed to specify the state of all arbitrarily small subregions of that region. When confronted with nonseparability, we accepted that this was insufficient, but assumed without justification that it would suffice to specify the state of all elements in an arbitrary *foliation* of the region. In fact, in a nonseparable theory there are no shortcuts. In principle at least, the state of any region must be specified *directly*: no attempt to specify it in terms of any non-trivial decomposition into subregions will work.

This is not to downplay the importance of foliations, of course. For a variety of reasons²³ we are often only interested in local properties—properties of spacetime points, in the limit. And all of *these* properties are indeed fully specified by the quantum state on each element of a foliation. Furthermore, they are dynamically determined—indeed, all properties, however nonlocal, are dynamically determined—by the quantum state on a *single* element of this or any other foliation, once the Hamiltonian is taken as fixed. But this dynamical observation should not be confused with the entirely different—and quite false—claim that Poincaré-covariant QM is narratable.

9 Finding the appearances

A final point should be addressed briefly before concluding. We noted amongst the prolegomena that when presenting an ontology, one hopes to be able to

²²It is an interesting question exactly to what extent narratability failure is generic in non-separable relativistic theories, quantum or otherwise. Note that if one were able to infer the dynamics by perusing the sequence of states on a given foliation then this would allow one to calculate the state on any hypersurface from a previous one. So examples illustrating narratability failure will need to exploit the presence of symmetries precluding such inference. We thank Wayne Myrvold and Tim Maudlin for discussion of this point.

²³Perhaps best understood in terms of the structure of the decoherence-defined branching in the Everett interpretation, which is in turn determined by the local nature of the dynamics.

recount, even if only sketchily, how the property bearers and the properties one postulates in the ontology relate to appearances. How might such a story go in the case of spacetime state realism?

The *general* story will be the same as that in any approach which takes unitary quantum mechanics to be a complete theory: decoherence will provide an emergent classicality; more precisely, a superposition of effectively autonomous quasi-classical sequences of events. Observers able to monitor and interact with their surroundings emerge as inhabitants within such independent quasi-classical histories; corresponding to particular complex arrangements of fundamental particles; evolved to take advantage of the existence of enduring records that decoherence allows (Gell-Mann and Hartle [1990]; Saunders [1993]).

The only addition to this general picture in our proposal is that the states of spacetime regions are the primary elements. Determinate quasi-classical goings on will correspond to certain *special* states of large regions (which will themselves correspond (in general) to highly entangled states of the component sub-regions). In the absence of dynamical collapse, though, the states of large regions will typically correspond to no (single) quasi-classical situation. Rather, decoherence will guarantee that the states of these regions are convex *sums* of the special “quasi-classical” states:

$$\hat{\rho}_A = \sum_i p_i \hat{\rho}_A^i \quad (13)$$

where each $\hat{\rho}_A^i$ is a quasi-classical state of A . Each quasi-classical component evolves independently of the others, so $\hat{\rho}_A$ encodes not one, but a great many simultaneously present quasi-classical situations. And the spatial extension of the quasi-classical goings-on is encoded in the entanglement between states: if B is some largish region adjacent to A , then the state of $A \cup B$ might be

$$\hat{\rho}_{A \cup B} = \sum_i p_i \hat{\rho}_A^i \otimes \hat{\rho}_B^i; \quad (14)$$

the local states of A and B each encode many simultaneously present quasi-classical situations, and the nonlocal information in $\hat{\rho}_{A \cup B}$ encodes the connections between them that justifies the “many-worlds” language customarily used to describe the Everett interpretation²⁴. This brings home the point that the true state of a spatial region is very far from being directly accessible to any realistic agent. An observer in region A (present in the quasi-classical situation encoded by $\hat{\rho}_A^i$, say) might very well speak of *the* state of A being $\hat{\rho}_A^i$ and *the* state of B as being $\hat{\rho}_B^i$, but these would be emergent and approximate notions (somewhat akin to Everett’s original ‘relative states’). The true, ontologically primary, state of A would still be $\hat{\rho}_A$.

²⁴The state given by (14) is actually what would be termed a ‘classically correlated’ or separable state of $A \cup B$, but this shouldn’t mislead one into thinking that the regions are unentangled. Whether or not reduced states of subsystems of a N -party entangled state are themselves entangled depends on the type of N -party entanglement that obtains. If the global state is pure, then a state like (14) for subregions can only arise because of the existence of entanglement.

To help relate this rather abstract account to the more familiar, if somewhat specialised, examples one is used-to from quantum mechanical measurement theory, we can consider how an idealised measurement situation might look. In general, such cases of measurement would involve the establishment of entanglement between the regions where the object system may be thought of as located and the regions corresponding to the possible positions of macroscopic pointer variables. (The positions of pointers, and later, the states of observers, would constitute parts of the quasi-classical situations described by certain of the quasi-classical states, in the terminology of the preceding paragraph.) To take a very simple example in illustration, consider the Coleman-Hepp model of a spin measurement (Hepp [1972]), discussed by Bell (Bell [1975]).

The model consists of a single moving spin-1/2 system which interacts with a one-dimensional (semi-infinite) array of spatially fixed qubits as it moves along the array (one can think of this as a one-dimensional version of a Stern-Gerlach measurement). The qubits fixed at their locations $x_1, x_2, \dots, x_n, \dots$ all begin with spin-up in the z -direction; when the moving spin passes by one of the fixed qubits, the two interact by a controlled-not operation; that is, the qubit's value is flipped if the moving spin is spin-down and is unchanged if it is spin-up (the flip will be completed by the time the moving spin moves on to the next position in the chain). At t_0 , the moving spin starts at x_1 ; we can represent the states of each of the regions x_i by number states $|n_\uparrow, n_\downarrow\rangle_i$ for spin in the z -direction. If the moving spin begins in an equal superposition of spin-up and spin-down, then the state of the array at t_0 will be:

$$\frac{1}{\sqrt{2}}(|2, 0\rangle_1 + |1, 1\rangle_1)|1, 0\rangle_2|1, 0\rangle_3 \dots |1, 0\rangle_n \dots,$$

that is, a product state of all the regions in the array. At t_1 , however, as the moving spin has passed on to position x_2 , the state becomes:

$$\frac{1}{\sqrt{2}}(|1, 0\rangle_1|2, 0\rangle_2 + |0, 1\rangle_1|1, 1\rangle_2)|1, 0\rangle_3 \dots |1, 0\rangle_n \dots;$$

regions 1 and 2 are now entangled; and so it continues, until at t_n all the regions up to and including the n th are entangled together:

$$\frac{1}{\sqrt{2}}(|1, 0\rangle_1|1, 0\rangle_2|1, 0\rangle_3 \dots |2, 0\rangle_n + |0, 1\rangle_1|0, 1\rangle_2|0, 1\rangle_3 \dots |1, 1\rangle_n)|1, 0\rangle_{n+1} \dots$$

For sufficiently large n , the magnetisation of the array (the average of the dipole moments of the spins) will count as a macroscopic variable recording the value of spin possessed by the measured system. If we were to take the union of the regions x_1 to x_{n-1} at t_n as our region A , for example, and \hat{P}_1 and \hat{P}_{-1} to be projectors onto states with magnetisation of A close to 1 and -1 respectively, then the state of A will be:

$$\hat{\rho}_A = \frac{1}{2}\hat{P}_1 + \frac{1}{2}\hat{P}_{-1}; \tag{15}$$

corresponding to separate, determinate, values of the pointer variable.

10 Conclusion

We have seen that a superior alternative to wave-function realism is available when trying to understand the nature of the quantum state. While both the approaches we have discussed exploit the gains in intelligibility that finding a spatial arena confers, by choosing *spacetime* as the arena, we gain a univocal ontological picture for all standard types of quantum theory (non-relativistic fixed particle number; non-relativistic variable particle number; relativistic; field-theoretic) and we avoid artificially picking out certain classes of observables as preferred. While the kinds of property that will be attributed to a spacetime region by the assignment to it of a density operator are not ones with which we are terribly familiar, we have suggested that this kind of feature is generic in fundamental physics: the case, we have suggested, does not differ in principle from the unfamiliarity of the electromagnetic field. Our understanding of the nature of these properties will, in the end, have to be mediated by our becoming familiar with their role within the theory. We sketched, finally, an—admittedly highly schematic—account of the kind of relation these properties will have to experience; a form of account which is familiar from recent discussions of Everettian quantum mechanics.

On the matter of separability, we have suggested that the gains for the wave-function realist position of being separable at the fundamental level are not substantial. Non-separable theories seem perfectly intelligible; and when it comes to considering relativistic covariance, the fact that wave-function realism fails narratability looks mysterious and problematic. By contrast, with a non-separable spacetime picture, narratability failure is rendered innocuous as a perfectly natural consequence of adopting a covariant non-separable theory.

On a final note, recall the hierarchy of counterintuitiveness presented in Section 7. Sticking with purely unitary quantum mechanics keeps us in the top-half of the hierarchy. But given all of our experience with the notion of entanglement, we surely can't expect to move all the way to the top and recover the full locality of classical physics, while sticking with a theory empirically equivalent to quantum mechanics. Or can we? This is the question we take up in Part II.

Balliol College and Faculty of Philosophy,
University of Oxford,
OX1 3BJ, UK.

david.wallace@balliol.ox.ac.uk

Brasenose College and Faculty of Philosophy
University of Oxford,
OX1 4AJ, UK.

christopher.timpson@bnc.ox.ac.uk

Funding

Arts and Humanities Research Council, CGT.

Acknowledgements

We would like to thank Jeremy Butterfield, David Deutsch, Tim Maudlin, Wayne Myrvold and the anonymous referees for helpful comments and discussion. CGT would also in particular like to thank Joseph Melia for helpful discussion.

References

- Aharonov, Y. and D. Albert [1980]: ‘States and observables in relativistic quantum field theories’, *Physical Review D*, **21**, pp.3316–3324.
- Albert, D. Z. [1996]: ‘Elementary quantum metaphysics’, In J. T. Cushing, A. Fine, and S. Goldstein (eds.), *Bohmian Mechanics and Quantum Theory: An Appraisal*, Dordrecht:Kluwer Academic Publishers, pp. 277–284.
- Albert, D. Z. [2007]: ‘Physics and Narrative’, Unpublished manuscript, available from <fas-philosophy.rutgers.edu/philosophy/science/papers/>.
- Barbour, J., B. Z. Foster, and N. Ó Murchadha [2002]: ‘Relativity without relativity’, *Classical and Quantum Gravity*, **19**, pp.3217–3248.
- Barbour, J. B. [1999]: *The End of Time*. London: Weidenfeld and Nicholson.
- Bell, J. S. [1975]: ‘On wave-packet reduction in the Coleman-Hepp model’, *Helvetica Physica Acta*, **48**, pp.93–98. Reprinted in Bell ([1987]), pp. 45–51.
- Bell, J. S. [1981]: ‘Quantum Mechanics for Cosmologists’, In C. J. Isham, R. Penrose, and D. Sciama (eds.), *Quantum Gravity 2: A Second Oxford Symposium*, Oxford: Clarendon Press. Reprinted in Bell ([1987]), pp. 117–138.
- Bell, J. S. [1987]: *Speakable and Unsayable in Quantum Mechanics*, Cambridge: Cambridge University Press.
- Brown, H. R. [2005]: *Physical Relativity: Spacetime Structure from a Dynamical Perspective*, Oxford: Oxford University Press.
- Butterfield, J. N. [2006]: ‘Against *Pointillisme* about Mechanics’, *British Journal for the Philosophy of Science*, **57** (4), pp.709–753.
- Coleman, S. [1985]: *Aspects of Symmetry: Selected Erice lectures*, Cambridge: Cambridge University Press.
- Deutsch, D. and P. Hayden [2000]: ‘Information flow in entangled quantum systems’, *Proceedings of the Royal Society of London*, **A456**, 1759–1774.
- Earman, J. [1987]: ‘Locality, non-locality, and action at a distance: A skeptical review of some philosophical dogmas’, in R. Kargon and P. Achinstein (eds.), *Kelvin’s Baltimore Lectures and Modern Theoretical Physics*, Cambridge, MA: The MIT Press, pp. 449–490.

- Earman, J. [1989]: *World Enough and Space-Time: Absolute versus Relational Theories of Space and Time*, Cambridge MA: The MIT Press.
- Einstein, A. [1948]: ‘Quantum mechanics and reality’ *Dialectica*, **320**. Reprinted in *The Born-Einstein Letters*, tr. I. Born, London: Macmillan [1971] pp. 168–173.
- Fleming, G. N. [1988]: ‘Hyperplane-dependent quantized fields and Lorentz invariance’, In H. R. Brown and R. Harré (eds.), *Philosophical Foundations of Quantum Field Theory*, Oxford: Clarendon Press, pp. 93–115.
- Fleming, G. N. [2002]: ‘The Dependence of Lorentz Boost Generators on the Presence and Nature of Interactions’, available online at <philsci-archive.pitt.edu/archive/00000663/>.
- Gell-Mann, M. and J. Hartle [1990]: ‘Quantum Mechanics in the Light of Quantum Cosmology’, in W. Zurek (ed.), *Complexity, Entropy and the Physics of Information*, Redwood City Ca.: Addison-Wesley, pp. 425–458.
- Haag, R. [1996]: *Local Quantum Theory: Fields, Particles, Algebras*. Berlin: Springer-Verlag.
- Healey, R. [1991]: ‘Holism and nonseparability’, *Journal of Philosophy*, **88**(8), pp.393–421.
- Hepp, K. [1972]: ‘Quantum-theory of measurement and macroscopic observables’, *Helvetica Physica Acta*, **45**, p.237.
- Lewis, D. [1986]: *Philosophical Papers, Vol. II*, Oxford: Oxford University Press.
- Lewis, P. J. [2004]: ‘Life in configuration space’, *British Journal for the Philosophy of Science*, **55**, pp. 713–729.
- Monton, B. [2002]: ‘Wave function ontology’, *Synthese*, **130**, pp. 265–277.
- Monton, B. [2006]: ‘Quantum mechanics and 3N-dimensional space’, *Philosophy of Science*, **73**(5), pp. 778–789.
- Myrvold, W. [2002]: ‘On peaceful coexistence: Is the collapse postulate incompatible with relativity?’, *Studies in the History and Philosophy of Modern Physics*, **33**, pp. 435–466.
- Myrvold, W. [2003]: ‘Relativistic Quantum Becoming’, *British Journal for the Philosophy of Science*, **54**, pp. 475–500.
- Saunders, S. [1993]: ‘Decoherence, Relative States, and Evolutionary Adaptation’, *Foundations of Physics*, **23**, pp. 1553–1585.
- Saunders, S. [1995]: ‘Time, Decoherence and Quantum Mechanics’, *Synthese*, **102**, pp. 235–266.
- Timpson, C. G and H.R. Brown [2002]: ‘Entanglement and Relativity’, in R. Lupacchini and V. Fano (eds.), *Understanding Physical Knowledge*, Bologna: University of Bologna CLUEB, pp. 147–166. Available from <arxiv.org/abs/quant-ph/0212140>

- Tumulka, R. [2006]: ‘A Relativistic Version of the GRW Model’, *Journal of Statistical Physics*, **125**, pp. 821–840.
- Vaidman, L. [2002]: ‘The Many-Worlds Interpretation of Quantum Mechanics’, In the Stanford Encyclopedia of Philosophy (Summer 2002 edition), Edward N. Zalta (ed.), <plato.stanford.edu/archives/sum2002/entries/qm-manyworlds>.
- von Weizsäcker, C. F. [2006]: *The Structure of Physics*, Berlin: Springer.
- Wald, R. M. [1994]: *Quantum Field Theory in Curved Spacetime and Black Hole Thermodynamics*, Chicago: University of Chicago Press.
- Wallace, D. [2003]: Everett and Structure. *Studies in the History and Philosophy of Modern Physics*, **34**, pp. 87–105.
- Wallace, D. [2006]: ‘In defence of naïveté: The conceptual status of Lagrangian quantum field theory’, *Synthese*, **151** (1), pp.33-80.
- Wallace, D. [2008]: ‘The interpretation of quantum mechanics’, In D. Rickles (ed.), *The Ashgate Companion to Contemporary Philosophy of Physics*, London: Ashgate [2008].
- Wallace, D. [2009]: ‘QFT, Antimatter, and Symmetry’, forthcoming. available online at <philsci-archive.pitt.edu/archive/>.
- Wallace, D. and C. Timpson [2009]: ‘Quantum mechanics on spacetime II: Quantum gauge freedom’, Forthcoming.