

Finding ordinary objects in some quantum worlds

CIAN DORR

Draft of March 9, 2009

Comments welcome

1 Introduction

When metaphysicians pay attention to theories in fundamental physics, they are usually interested in their implications for fundamental metaphysics: they want to know what the theories imply or suggest concerning the ultimate structure of reality. This essay, by contrast, is concerned with the “derivative metaphysics” of quantum mechanics. I will be considering a certain hypothesis about the fundamental structure of reality, which I will call H . I am interested in H because it strikes me as embodying the most obvious, flat-footed way to read an account of the fundamental structure of reality off the mathematical apparatus of elementary quantum mechanics. But I won’t be defending this claim here, or trying to argue that the members of H are more plausible than other accounts of the fundamental structure of reality inspired by quantum mechanics. Instead, my question is: if H were true, could there be (non-fundamentally speaking, of course) such ordinary objects as plants, planets and people? More specifically: could there be such ordinary objects as ourselves and the things we are familiar with, having more or less the ordinary properties we take them to have? And more generally: given that the fundamental structure of a possible world is such as to make H true, what is its correct description in ordinary terms?

Of course, I care about these questions in part because I want to explore the extent to which we might have good reason to believe H , or something akin to it. But even if H turned out to be *consistent* with our ordinary knowledge about the world of ordinary objects, it would

obviously take a lot more work to show that it has any serious claim on our credence. At a minimum, one would have to say something about how *likely* we should regard our evidence as being, given H ; one would want to make some sort of survey of alternative hypotheses; and one would want to take into account our specialised scientific knowledge as well as our ordinary knowledge, including the facts that make trouble for elementary quantum mechanics and force us to look beyond it to quantum field theory and ultimately towards some kind of theory of quantum gravity.

Before saying what H is, I should announce a couple of additional assumptions I will be making. The first is that the existence of ordinary objects of the kinds we are familiar with—including human beings with properties like knowledge, consciousness and ethical virtue—doesn’t, in general, require positing any distinctive extra kinds of fundamental structure over and above the kinds of fundamental structure one would need to posit to make true the kinds of descriptions of the world we find in physics. If you reject this kind of physicalism, you may or may not have anything to learn from my discussion: it depends on what kind of story you want to tell about the constraints the “physical” fundamental facts place on the rest of the fundamental facts, and about the ways in which the two kinds of fundamental facts co-operate in determining the complete truth about the realm of ordinary objects. But you’ll have to sort this out by yourself.

I will also be assuming the falsity of an extreme version of scientific essentialism. A natural moral to draw from the work of Kripke

and Putnam is that it is impossible for there to be water at a world without hydrogen, and impossible for there to be hydrogen at a world where the fundamental physical facts are too dissimilar from those at the actual world. If so, there is no hope of answering questions of the form “if the fundamental physical facts were thus and so, would there be any water?” from the armchair, prior to an investigation of the physical facts at the actual world. The extreme scientific essentialist I am thinking of is someone who generalises this claim about “water” to all our ordinary predicates, so that questions like “if the fundamental physical facts were thus and so, would there be any plants/planets/people. . . ?” are equally incapable of being answered from the armchair. I am inclined to think that the adoption of this view would require me to make only fairly superficial changes to the paper, involving judicious substitutions of indicative conditionals for subjunctives, and of claims about a priori knowledge for claims about metaphysical necessity. But since I am not an extreme scientific essentialist, and since I appreciate the expository benefits of possible-worlds talk, I won’t try to speculate about how this might go.

2 The hypothesis

Suppose you are a metaphysician with a physicalistic bent opening a quantum mechanics textbook for the first time in search of interesting suggestions about the fundamental structure of the universe. The obvious first place for you to look (since the universe is by definition an isolated system) is the book’s treatment of *isolated, many-particle systems*. This means you’ll ignore all the stuff about “measurement”, which is standardly discussed as something that happens when non-isolated systems interact with their environments. This is unsurprising: *prima facie* you’d expect the physics of measurement to be about the ways in which certain non-fundamental facts depend on the fundamental ones rather than about the fundamental facts themselves.

To explain what *H* is, and why it is the natural first thought about the fundamental structure of the world that might be inspired by quantum mechanics, I need to say a little about the role played by *configuration space* in the mathematical formalism of (elementary) quantum mechanics.

In classical mechanics, “configuration space” is normally intro-

duced as a name for some sort of mathematical abstraction. In the central case of a system of finitely many point-particles moving in Newtonian absolute space, one natural candidate is the set of all functions from particles to points of space. This set is naturally regarded as a “space” in its own right, since there are natural ways to extend geometrical notions like betweenness and distance from the points of space to the members of this set.¹ The history of the universe can then be exhaustively represented as a function from times to the points of configuration space, so understood. To claim that a given function *f* accurately represents the history of the universe is to claim that at each time *t*, each particle *p* is located at *f(t)(p)*. In the Lagrangian formalism for classical mechanics, the basic dynamical laws are expressed as constraints on this function.

In quantum mechanics, configuration space takes on a new sort of life. The basic dynamical law, Schrödinger’s equation, is a constraint on a function—“the wavefunction”—that describes the history of the universe by associating each time not with a single *point* in configuration space, but (in the simplest case, where we ignore spin) with a *function from configuration space to the complex numbers*.²

But what is it for such a function to accurately represent the history of the universe? It is hard to see how to answer this question while continuing to think of points of configuration space as abstract entities constructed from particles and points of space. That a certain function from particles to points of space is mapped to a certain complex number at a certain time is not the sort of thing that could itself be a primitive fact: such facts must, somehow, obtain in virtue

¹The definition of betweenness is straightforward: *f*₁ is between *f*₂ and *f*₃ iff for each particle *p*, *f*₁(*p*) is between *f*₂(*p*) and *f*₃(*p*). The standard definition of distance is a bit more complicated: the distance between *f*₁ and *f*₂ is

$$\sqrt{\sum_p \text{Mass}(p) |f_1(p) - f_2(p)|^2}.$$

The naturalness of this particular definition only becomes apparent once one starts to think about how to use configuration space in succinctly formulating the laws of classical (or quantum) mechanics.

²While it is customary to speak of “the wavefunction” as if a unique such function had been singled out, in fact it is universally agreed that there is some slack in the description ***

of facts about the relations among concrete objects. And it is hard to see how this could be true, if the only fundamental concrete objects we have to work with are particles and points of space, time and/or spacetime, and the fundamental relations among these objects don't have gigantic numbers of argument places.³ The space of possibilities generated by the permutations of such relations among such objects isn't big or varied enough to constitute the space of distinct possible wavefunctions over configuration space. We need to supplement our ontology with some new category of fundamental concrete entities. And the most natural candidate for this role are *points of configuration space*, understood now not as mathematical abstractions but as concrete particulars, as real and fundamental as the points of ordinary space.

The thought here is familiar from attempts to make ontological sense of classical field theories. Classical field theories are naturally taken to be about points (or regions) of *space* (or spacetime); and the task of explaining what in the fundamental structure of the world could make such a theory true seems to require us to take such entities to be fundamental, concrete particulars rather than logical constructions out of some "sparser" domain of objects such as particles. But the wavefunction stands to configuration space as classical fields stand to ordinary space; so *prima facie*, the way to make ontological sense of descriptions of the world in terms of a wavefunction seems to involve some kind of realism about configuration space.

The idea that quantum mechanics requires realism about configuration space is not new; it has most recently been defended by Albert (1996). But Albert wants to combine the claim that configuration space is fundamental with the claim that ordinary space is *not* fundamental. Descriptions of the world in terms of ordinary space—say, the claim that configuration space point c differs from point c' only as regards where they put seventeen particles—must, on Albert's account, be

³Well, actually you can do it if you think it's OK to help yourself to non-symmetric relations with very many (in the order of the number of particles) argument places (see Forrest 1988: chapter ?). For reasons I won't go into here (see Dorr 2004) I don't think this is OK. But I don't think that the difference between this sort of fundamental structure and those discussed below could matter from the point of view of the existence of ordinary objects.

made true by facts entirely about the points of configuration space and the pattern of relations among them summarised in the wavefunction. This strikes me as an audacious eliminativist programme, worryingly reminiscent of other programs to reduce domains of facts which *prima facie* seem to play a role in explaining dynamics to the dynamical facts they purport to explain—I am thinking, for example, of the program to reduce facts about the (non-topological) geometric structure of spacetime to facts about the distribution of matter, or of the proposal to reduce facts about putatively intrinsic properties of particles like mass and charge to facts about their dynamical role. In any case, even if Albert's pared-down ontological scheme turns out to be defensible, we have little hope of assessing its implications for ordinary objects without first assessing the implications of a more richly structured fundamental ontology, which helps itself to standard characterisations of points of configuration space in terms of where they put particles in ordinary space.

I've been glossing over one important point. If we take points of space and instants of time as fundamental, we will have to recognise a structural definition between a world where the wavefunction is at absolute rest (in the sense that its amplitude at the point of configuration space corresponding to any given function from particles to points of space is the same at every instant) and one where the wavefunction is in uniform rectilinear motion. The fact that this distinction has no further physical consequences constitutes a defeasible reason to prefer an ontology in which no such distinction can be made. What might this look like? Well, instead of separate categories of points of space and instants of time, we can take as fundamental the points of Galilean (or "Neo-Newtonian") spacetime, in which there is a natural notions of simultaneity and a natural distinction between inertial and accelerated trajectories, but no natural way to characterise certain trajectories as being at absolute rest. And instead of taking points of configuration *space* as fundamental, we will have to take points of configuration *spacetime*, which correspond one-one with functions from particles to points of spacetime which are all simultaneous with each other. Points of configuration spacetime work like ordered pairs of points of configuration space and instants of time, but he ques-

tion whether two points of configuration spacetime correspond to the same point of configuration space will only make sense relative to the choice of an arbitrary frame of reference. Instead of taking a wavefunction as a function that assigns complex numbers to points of configuration space relative to instants of time, we can understand it as in the first instance assigning complex numbers once and for all to points of configuration spacetime.

4

Thus, my Hypothesis H will feature three categories of fundamental objects: *particles*, *points of spacetime*, and *points of configuration spacetime*. These entities are structured by four families of physically basic relations.

- (i) Fundamental properties of and/or relations among the particles, which serve to ground assigning them to different kinds, with different masses, charges, etc.
- (ii) Some fundamental relation or system of relations among the points of spacetime, which confer on them a Galilean geometry.
- (iii) A single ternary relation '*c* puts *p* at *x*', holding among points of configuration space, particles and points of spacetime, and subject to the law that for every function *f* from particles to points of space all of which are simultaneous, there is exactly one point *c* of configuration space such that *c* puts *p* at *x* whenever $f(p) = x$.
- (iv) Some system of fundamental relations among the points of configuration spacetime, which single out certain functions

⁴There is a complication: the question whether a given function from configuration spacetime to the complex numbers is admissible as a wavefunction doesn't make sense absolutely, but only relative to the choice of a frame of reference. To take a wavefunction that is good from the standpoint of one frame of reference and make it good from the standpoint of some other frame of reference, we have to multiply it by a certain function from configuration spacetime to complex numbers of modulus one, depending on the velocity of the new frame relative to the old frame. Only in that case is there any prospect of stating a law for the evolution of the wavefunction in a way that doesn't give special privilege to any one frame of reference.

from them to the complex numbers as "admissible wavefunctions".⁵ The admissible wavefunctions so determined all satisfy Schrödinger's Equation.

And that is *all* the fundamental structure of the world, according to *H*. There is nothing more to the world beyond the facts that constitute the wavefunction and govern its evolution. There is, for example, no further *location* relation assigning each particle to a point of space at each time.⁶ If there is any sense in which particles at a *H*-world can be described as having "locations", these descriptions will have to be made true somehow by the pattern of fundamental relations of categories (i)–(iv). Also, the evolution of the wavefunction is entirely deterministic, in accord with Schrödinger's equation: there is nothing in the fundamental structure of a *H*-world corresponding to the "collapse of the wavefunction". Thus, "collapse" and "hidden variables" interpretations of quantum mechanics are both false at *H*-worlds.

⁵What could these relations look like? One not particularly elegant approach is use separate relations to ground the assignments of amplitudes and phases. For amplitudes, we can take the four-place relation '*c*₁ and *c*₂ are further apart in amplitude than *c*₃ and *c*₄' as fundamental. For phases, we need something more complicated, since as we saw in note 4, comparisons of assignments of phase make sense only relative to a frame of reference. The obvious approach is to use the six-place relation '*c*₁ and *c*₂ are further apart in phase than *c*₃ and *c*₄ relative to a frame of reference in which the inertial trajectory through *x*₁ and *x*₂ counts as being at rest', where *x*₁ and *x*₂ are two non-simultaneous spacetime points. (This role could equally well be played by points of configuration spacetime.) It would be nicer, though, if we could find a way to see the rule relating the wavefunctions admissible relative to different frames of reference not as a mere law of nature but as flowing from a definition of frame-relative notions in terms of some common underlying reality. I suspect that when we take spin into account, and allow for the possibility that configuration spacetime has a nontrivial topology, the most attractive approach will involve replacing each point of configuration spacetime with a miniature space—the fibre over that point—and grounding the facts about the wavefunction in the distribution of some monadic property attaching to one point in each fibre (see Grabowska et al. MS).

⁶To some, this will make "particle" seem an inappropriate term for the category of fundamental objects that I have called by that name. But nothing hangs on the name: *H* purports to be an account of the fundamental *structure* of reality, so I take it that all the names I have given to the fundamental relations and categories are really just variables bound by an initial existential quantifier.

H, with its three categories of fundamental entities, is not the most elegant account of the fundamental structure of reality one can imagine. There are many ways in which one might attempt to simplify it while preserving the basic spirit of configuration space realism. But the task of exploring these options is best left to another occasion: once we figure out how the language of ordinary objects works at *H*-worlds, we will be better-placed to figure out what difference, if anything, the various possible simplifications might make.

3 An argument that ordinary objects don't exist at *H*-worlds

Many (most?) interpreters of quantum mechanics have thought that no world exhaustively characterised by a wavefunction evolving in accordance with Schrödinger's equation—and thus no world where *H* was true—could contain ordinary objects remotely like ourselves and our surroundings, as we find them. Some who hold this conclude that while our world *is* exhaustively characterised by a wavefunction, it occasionally evolves in a manner quite different from the one described by Schrödinger's equation: the notorious “collapse of the wavefunction”. Others prefer the view that the wavefunction evolves in accordance with Schrödinger's equation, but only partially characterises the world: an exhaustive characterisation of the world at the fundamental level would have to involve some additional physically basic relations over and above those involved in grounding the truth about the wavefunction, and perhaps also some additional fundamental entities to instantiate them; and the facts about these so-called “hidden variables” play a crucial role in constituting the existence of ordinary objects.

Why do people think this? The reason has to do with the following characteristic of the evolution determined by Schrödinger's equation (given any remotely any remotely realistically complex Hamiltonian). Even if the wavefunction at one time is tightly bunched, so that only a small and compact region of configuration space has any amplitude to speak of, it will not stay bunched. It will quickly spread out, gradually becoming smeared across more and more of configuration space. And its shape will come to be quite complex, with many peaks and troughs. So any reasonably complicated *H*-world will spend almost all of its time with its wavefunction spread widely and

messily across configuration space.

And why should this show that such worlds don't contain ordinary objects like ourselves and our surroundings? Abstracting from the customary focus on stylised “measurement interactions” involving “measuring devices” equipped with “pointers”, the standard argument goes something like this. First, what Albert (2005: 849) calls a “truism”: ordinary objects, “whatever else they may be, are physical systems, and so are subject to the same universal physical laws as other such systems are”. But, just as the complete truth about any *particle* at a typically messy *H*-world is not such as could credibly ground the attribution to it of even an approximately determinate *location* or *velocity*; and just as the complete truth about any *collection* of particles at such a world is not such as could credibly ground the attribution to them of even an approximately determinate *shape*; so, more generally still, the complete truth about *any* physical system at such a world is not such as could ground the attribution to it of even an approximately determinate shape, size, location, velocity, colour, hardness, etc., on any credible account of what it is to have properties like these. But ordinary objects that were anything like the ones posited by common sense would have to have reasonably determinate shapes, sizes, etc. So such objects don't exist at *H*-worlds.

There is a big problem with this argument. For it to work, “physical system” needs to mean something very like “mereological fusion of particles”. It is indeed true that on any remotely credible theory of what it is to be *roughly spherical*, for example, it will turn out that there are no roughly spherical mereological fusions of particles at a typical time at a typical *H*-world. (Or, to be cautious in a way I don't actually think is necessary: that there are no *determinately* roughly spherical mereological fusions of particles at a typical time at a typical *H*-world.) Likewise, on any remotely credible theory of what it is to be a table or a tree or a person, it will turn out that no mereological fusion of particles is a table or a tree or a person at a typical time at a typical *H*-world. But it is far from being a truism that ordinary objects are mereological fusions of particles. It is, indeed, obviously false. For one thing, ordinary objects, unlike mereological fusions of particles, frequently gain and lose parts.

OK, that was a bit fast. It is true that almost all presentations of the “measurement problem” make use of a notation that seems to make sense only if ordinary objects (such as “pointers”) *are* mereological fusions of particles.⁷ But it is also true that one doesn’t need a premise quite as strong as this to run the argument for the non-existence of ordinary objects at *H*-worlds. It would be enough if we could assume that each ordinary object is composed *at any given time* of particles, and that two ordinary objects composed of the same particles at some time can’t differ at that time in respect of certain temporally localised, categorical properties and relations—a class that includes, at least, spatial properties like *being roughly spherical* or *being table-shaped*. For it is hard to see how ordinary objects could exist if they never (determinately) have properties like these.

This weaker assumption about the relation between ordinary objects and collections of particles is a natural one to make if we are confining our attention to worlds where classical mechanics is true. But I don’t think it’s a necessary truth. For one thing, ordinary objects can be composed of things other than particles. For another thing—and more interestingly—even at worlds where ordinary objects *are* composed entirely of particles, it can happen that two objects composed of the same particles at some time are radically dissimilar at that time, even in respect of properties like shape and size.⁸

⁷The notation in question is one in which it is assumed that the Hilbert space of the universe can be represented as the tensor product of a Hilbert space representing to the nomologically possible states of an ordinary object and a Hilbert space representing the nomologically possible states of the rest of the universe. Strictly speaking, this assumption doesn’t quite require ordinary objects to correspond to sets of particles—it allows us, for example, to countenance an object whose state is completely specified just by giving the state of the *centre of mass* of some collection of particles. But I don’t think this additional freedom does much to make the assumption any more plausible.

⁸Fine (2003) also denies that objects composed of the same particles have to be alike in shape and size, on the grounds that ‘a loaf of bread and the bread that composes it are materially yet not spatially coincident’. I assume he’s thinking that the region of space occupied by the loaf includes the holes inside the bread, whereas the region of space occupied by the quantity of bread does not. Fine’s argument has the advantage of not requiring us to consider exotic possible worlds; but the argument below has the advantage of generalising much more naturally to the worlds I am interested in.

Consider a world where, fundamentally speaking, there are just particles and spacetime points, with a physically basic relation *L* holding between the particles and the spacetime points. Suppose for concreteness that the facts about *L* are governed by Newton’s laws, in the sense that some appropriate formulation of Newtonian mechanics is true when ‘is located at’ is interpreted as expressing *L*. Presumably such a world could be host to many ordinary objects, including people. Now take one such world and add a new physically basic relation *L**, also holding between particles and spacetime points. Let the facts about *L** also be governed by Newton’s laws, and let the pattern of holding of *L** be such that if we deleted *L* altogether and just kept *L**, we would have a world with many ordinary objects, including people. It seems absurd to suppose think that in *adding* the *L** relation to the original world, we could turn it into a world *without* ordinary objects or people. Surely we must say that the world with both relations is a world where both sets of ordinary objects exist. If we tinker with the laws to allow some circumstances in which the *L*-facts make a difference to the *L**-facts, we can even imagine that there is communication between the *L*-people and the *L**-people. For example, we could imagine that there is a special region of space, such that so long as a particle bears both *L* and *L** to points within that region, its *L*-trajectory and its *L**-trajectory tend (other things being equal) to approach one another.

If we say that the ordinary objects at this world are composed of particles, we will have to conclude that being composed by the same particles at a given time need not prevent ordinary objects from being not only distinct, but dissimilar in all sorts of manifest, categorical, temporally localised respects. The very same particles might happen to compose a person in virtue of bearing *L* to certain distinctively-arranged spacetime points, and compose some other ordinary object—a chair, say—in virtue of bearing *L** to some other spacetime points. This is a situation in which two dissimilar ordinary objects have an unusual form of metaphysical kinship. It is certainly not a situation in which one and the same thing is both a person and a chair, or even one in which one and the same thing is neither determinately a non-person nor determinately a non-chair. Merely listing the

particles that are parts of a given ordinary object and describing their various properties and relations doesn't entail anything much about the object's ordinary properties, such as its shape, size or location. One needs to know, in addition, whether the object is an "L-object" or an "L*-object".⁹

It's not obvious how the distinction between L-objects and L*-objects is to be understood, metaphysically speaking. It is tempting to relativise the notion of composition, saying that certain particles compose one thing relative to L and another thing relative to L*. But we need a story about what the adverbial clause 'relative to L' is doing here. If it is saturating an argument place, what sense can we make of ordinary claims about composition in which no argument for that place is explicitly supplied? What we want is a way of talking about the constitution of ordinary objects out of fundamental ones flexible enough to handle all possible cases, not just this one. Some progress towards this goal has been made by proponents of broadly "hylomorphic" conceptions of material objects, like Kit Fine (2000) and Mark Johnston (2005). In these systems, we get to speak of objects as consisting of some things (the matter, or 'basis') *related in a certain way* (the form, or 'gloss'): so we could take L-objects to be those whose form or gloss is a relation specifying the arrangement of the spacetime points to which the relata bear L, and similarly for L*-objects.¹⁰

⁹One might attempt to hold on to the principle that objects that are composed of the same fundamental entities at a time must be alike in respect of properties like shape at that time by claiming that ordinary objects at the L-L* world are composed of spacetime points instead of, or in addition to, particles. But even on this view, there will still be surprising failures of the ordinary properties of an object to supervene on the truth about which fundamental entities are parts of it. For example, it seems wrong to say that an L-object can *touch* an L* object, even if they happen to occupy adjoining regions of space. So the facts about whether two objects touch at a time don't depend on where they are at that time, or on which particles are parts of them: we need to know whether they are L-objects or L*-objects.

¹⁰This is only a first approximation, as Fine and Johnston realise: something more flexible is needed to characterise the structure of objects with temporally and modally variable constitution. Here is a general scheme that I think might well be adequate: to specify the essence or metaphysical structure of an object, we specify a single time-dependent property of [or relation between times and] objects; necessarily, the immediate constituents of an object at a time are all and only those things that have its associated property at that time. Fine's object with basis $a_1 \dots a_n$ and gloss R will

I introduced the L-L* world to help undermine an argument that ordinary objects can't exist at H-worlds. But we can also use this world as the basis for a positive argument that suggests that ordinary objects *can* exist at H-worlds. For the verdict that the L-L* world is host to several collections of ordinary objects does not seem to be affected if we make various changes which take us closer to a H-world:

- (i) It is hard to see how it could matter if, instead of having each of the "location" relations be a physically basic relation in its own right, we posited two (or more) special fundamental *objects*—call them "locators"—with a physically basic, ternary relation of "putting" holding between locators, particles, and points of spacetime, subject to the law that each locator puts each particle at exactly one spacetime point from each time. In a world like that, each locator would have its own family of ordinary objects.
- (ii) It is hard to see how it could matter if, instead of having locators endure through time, putting each particle at many non-simultaneous spacetime points, we let each locator be associated with a single instant of time, putting each particle at exactly one spacetime point belonging to that time. The role played by locators in the old framework could be played in this new framework by equivalence classes of locators under a new, physically basic, equivalence relation G (for 'genidentity'), whose equivalence classes contain exactly one locator for each instant of time. In a world like that, each G-equivalence class of locators would have its own family of

then be the special case where the essence-specifying property is *being one of* $a_1 \dots a_n$ and such that $Ra_1 \dots a_n$. Given this generalised framework, we can define inductively what it is for something to be an L-object at a time as follows: an object is an L-object at t iff its immediate constituents at t are all particles or L-objects, and those of them that are particles at t have the object's essence-specifying property at t *in virtue of the facts about which spacetime points they bear L to*. This even lets us allow for objects which are L-objects at some times and L*-objects at others—a welcome result, since if the laws allow exchange of information between the "L-realm" and "L*-realm", travel between the two realms should also be nomologically possible, at least on the kind of liberal view of personal identity which allows people to survive teleportation (Parfit 1986: 199).

ordinary objects.

- (iii) It is hard to see how it could matter if there were many more than two locators for each time—even continuum many of them, fitting together in such a way that there was at least one locator corresponding to each function from particles to simultaneous spacetime points. Provided the laws took the right form, it would still be true that each of the continuum many G -equivalence classes of locators had its own family of ordinary objects. In a world like that with N particles, every temporally local property that a family of ordinary objects composed of N particles could possibly instantiate would be actually instantiated at each time.
- (iv) It is hard to see how it could matter if, instead of always containing just one locator for each time, the G -equivalence classes of locators sometimes contained more than one locator for a time, provided that the simultaneous spacetime points assigned to any particle by G -equivalent locators were always very close to one another. Claims about the geometric properties of an ordinary objects at a world like would, arguably, be slightly vague even if it were not at all vague which particles composed the object. But such a small dose of “fuzziness” wouldn’t prevent the objects from being tables, trees, people, etc.

These “locators” are structurally so like what we have been calling “ c -points” that I think we can naturally draw a parallel conclusion about what ordinary objects at a H -world would be like, if there were any. Namely: there is something in the metaphysical structure of any given ordinary object that lets us think of it as ‘occupying’ or ‘belonging to’ a certain set of c -points. Perhaps we will still want to say that ordinary objects are composed of particles. But if we do, we shouldn’t expect objects that are composed of the same particles at a time to be alike, even in respect of properties like shape and size. If we want to know what shape or size an object is, it is not enough to know which particles compose it and what the fundamental facts about those particles are. We also need to know something

more about *how* the object in question is built up from its constituent particles, something that gives us an answer to the question *where in configuration space the object is located* at the time in question.

Friends of compositional uniqueness (who hold that any things whatsoever compose at most one object) won’t be able to agree that there are ordinary objects *composed of particles* at H -worlds. But even they should be able to agree that these worlds contain ordinary objects, by regarding them as composed at least in part of fundamental entities of other sorts. For example, one could think of an ordinary object as being composed entirely of c -points and spacetime points, subject to the constraint that each point of spacetime that is part of an ordinary object must be one at which one of that object’s constituent c -points puts some particle or other. This gives us a very direct way of making sense of the idea that an object has a location in configuration space as well as a location in ordinary space.¹¹ True, this is a bit artificial. But friends of compositional uniqueness already have good reason to put up with this sort of artificiality rather than dogmatically ruling out whole ranges of possible fundamental ontologies as inconsistent with what we know about ordinary objects.

If there *are* ordinary objects at H -worlds, there are clearly very many of them—so many that every sufficiently interesting point of configuration spacetime is occupied by some appropriate collection of ordinary objects. For it would be objectionably arbitrary to claim that there is some positive threshold—even some vague threshold—such that ordinary objects only ever occupy points of configuration spacetime whose amplitude is above that threshold. And this means that any given region of the *spacetime* of a H -world will likewise be occupied by an immense variety of ordinary objects. But if these ordinary objects are to be anything like those to which common sense is committed, ordinary relations like seeing, touching, holding, kicking, and loving had better not hold too indiscriminately among them. The natural assumption is that these relations will turn out to hold only among objects that occupy the same region of configuration space.

¹¹However, it doesn’t let us agree with the natural thought that being spread out in configuration space at a single time leads to vagueness or “fuzziness” in claims about one’s geometric properties at that time.

Objects that occupy disjoint regions of configuration space will be cut off from one another in the same way that L -objects are cut off from L^* -objects.

In a sense, then, there will be many different “worlds” of ordinary objects at a H -world, just as there are two “worlds” at the L - L^* world. But I won’t be using the word “world” in this sense, for two reasons: first, because I often need to use it to mean “possible world”; and second, because some advocates of “many worlds interpretations of quantum mechanics” have used “world” to stand for some putative extra bit of *fundamental* structure or ontology over and above whatever is involved in grounding facts about the wavefunction, and thus over and above the sort of fundamental structure and ontology that is present at H -worlds.¹²

4 Thick objects and the future

If ordinary objects at H -worlds occupy regions of configuration spacetime as well as regions of ordinary spacetime, what shape are the regions they occupy? In particular: are they *thin*, occupying only one point of configuration spacetime per time, or *thick*, extended across configuration space as well as through time?

The present section will be devoted to an overview of the range of possible answers to this question. One thing that will emerge from this overview is that unless the regions of configuration spacetime occupied by ordinary objects at H -worlds are *thin*, in the sense that they contain at most one c -point per time, the hypothesis that the world is a H -world is inconsistent with some very central commonsense beliefs about ordinary objects and the future—e.g. my belief that there is a determinate fact of the matter as regards whether I will ever write a book that sells more than a million copies.

A useful way to taxonomise the range of possible views about the shape in configuration spacetime of the ordinary objects at a H -world is to consider what happens to the ordinary objects in a case where the wavefunction goes from being bunched up in one small region of

¹²I am thinking here especially of Deutsch (1985). It is not so clear whether DeWitt (1970), the originator of the term ‘many worlds interpretation’, has this sort of thing in mind or not. Recent work in the many-worlds tradition by Saunders (1998) and Wallace (2002) is clearly not in the business of positing new fundamental structure.

configuration space to being bunched up in several disparate regions. This sort of thing happens all the time at H -worlds, most dramatically in “measurement interactions” where the states of very many particles change so as to reflect some aspect of the initial states of a few particles. Suppose, for example, that the wavefunction is initially bunched up in a small region of configuration space which has whatever features a region would need to have to be occupied by a person sitting alone in a darkened room staring at a photon detector which is separated by a half-silvered mirror from a light source poised to emit a single photon. Given the way Schrödinger’s equation works, a world with this sort of wavefunction will soon evolve to one whose wavefunction assigns substantial and equal amplitude to two regions. And it will be clear that, if these regions are occupied by any people at all, one of them will be occupied by a person who is seeing a flash on a detector, and the other will be occupied by a person seeing no flash.

Which regions of configuration spacetime at this world are the ones occupied by people? If we assume that there is only *one* person occupying any part of the high-amplitude region of configuration space at the initial time, there are three possible views about where in configuration space this person is at the later time. *Discrimination*: she is either entirely in the “seeing a flash” region or entirely in the “not seeing a flash” region. *Bilocation*: she is in both regions. *Cessation*: she is in neither region, presumably because she is no longer anywhere at all. If we drop the assumption that there is initially only one person in the high-amplitude region, further combinations of these options open up. By far the most interesting of these *Overpopulation*: at least one of the people initially in the high-amplitude region ends up in the “seeing a flash” region, at least one ends up in the “not seeing a flash” region, and none of them ceases to exist or ends up in both regions.

This space of options has already been extensively explored: it is essentially the same space of options that we face in ordinary cases of “fission”, such as the division of amoebae, and the much-discussed science-fiction scenario in which a person’s brain is bisected and transplanted into two different bodies. Here too, there are four views worth taking seriously. *Discrimination*: the amoeba that existed before the fission continues to exist, and to occupy a relatively compact

("amoeba-shaped") region of space. *Bilocation*: the amoeba that existed before the fission continues to exist, and occupies a scattered region of space. *Cessation*: the amoeba that existed before the fission ceases to exist, or at least ceases to have a spatial location, when it divides.¹³ *Overpopulation*: there was more than one amoeba to begin with, and all of them continued to exist and to occupy a compact region of space; no new amoebae came into existence.¹⁴

It is controversial which of these is the correct view about the amoeba case. But there is much less controversy on the following point: if you know you are going to undergo amoeba-style fission, and you know what each of your "fission products" will be like, you are in a position to know all there is to know about what your future will be like. For example, if you are an intelligent amoeba about to divide, and you know that one of your fission products will be well-fed and the other one will starve, it makes no sense for you to wonder whether *you* will be the well-fed one or the starving one. There is no subject matter for ignorance or uncertainty here: no set of mutually exclusive and jointly exhaustive hypotheses, compatible with what you know, to which you could assign degrees of belief summing to one. If you thought that your existence involved some sort of non-physical facts, about "soul pellets" or whatever, your knowledge of the physical facts would of course leave you free to entertain different hypotheses about those facts. But on the assumption that it doesn't, your knowledge of these physical facts leaves no room for any relevant kind of uncertainty, given that they are symmetric in all important respects.

This claim is intuitively plausible. When we encounter science-fiction stories of fission, we have very little temptation to think that the fission products should be *uncertain* which of them existed before the fission. But if there is no sense in a fission product wondering 'Am I really who I seem to be, or am I a new person who only recently came into existence', there should likewise be no sense in a person who is soon to undergo fission wondering 'Will I be the well-fed

fission product or the starving one?' It is also supported by theoretical considerations, which differ depending on which of the four available views of the metaphysics of the situation we adopt. In the case of Cessation, the point is obvious: if the original amoeba is in a position to know that it won't be around at all after the fission, what could it have to wonder about? Likewise for Bilocation: if it knows that it will be starving *over here* and well-fed *over here*, what sense could there possibly be in wondering whether it will be starving or well-fed *simpliciter*? In the case of Discrimination, the relevant theoretical consideration comes from the pressure to say, on pain of arbitrariness, that it is *indeterminate* whether the original amoeba will be the well-fed one or the starving one.¹⁵

In the case of Overpopulation, the theoretical considerations are more subtle and merit a bit more discussion. *Prima facie*, one might expect that if Overpopulation were true for dividing amoebas, it would *have* to make sense for the pre-fission amoebae to be uncertain whether they will be well fed. Shouldn't each of the two amoebae that exist before the fission be uncertain about whether *it* is the amoeba that will be well fed or the one that will starve, even if they know all the relevant objective facts? Not according to the most prominent defender of Overpopulation, David Lewis. On Lewis's view, temporally extended objects like people and amoebae are not capable of referring directly to themselves. There is such a phenomenon as "self-locating uncertainty", but it is always based in uncertainty about the location and properties of *one's current time-slice*: the objects of belief can be identified with *properties of time-slices* (Lewis 1979a). And since, for

¹⁵Epistemicists about vagueness will want to resist this, at least if 'indeterminate' is understood in such a way that granting that it is indeterminate whether *P* (while knowing all the relevant underlying facts) rules out being uncertain whether *P*. But the kind of ignorance about the future that would be provided by epistemicism together with a Discrimination-style view of ordinary objects at *H*-worlds is very different from our ignorance of the future as we normally imagine it: the epistemicist can allow that I don't know what result I'll get when I perform a certain measurement, but there is no prospect of my overcoming this ignorance: after the measurement, I'll see one result, but I'll know that if I'm the person who existed before the measurement a new person with all my apparent memories has just come into existence and is seeing the other result; knowing this, I surely won't be able to know that I am not the person who has just come into existence.

¹³This view is defended by Parfit (1971), at least in the case of people.

¹⁴Overpopulation is defended for splitting people by Lewis (1976) and for amoebae by Robinson (1985).

Lewis, the two amoebae share their pre-fission time-slices, there is no room for them to have any self-locating uncertainty about the future.

I think Lewis is on to something important and true here, something even those who like Overpopulation but don't like the framework of temporal parts have reason to agree with. Here is an attempt at a more neutral formulation: say that two entities are *fundamentally indiscernible* at a time iff they are exactly alike in their relations to all fundamental entities belonging to that time. Then we can put the point as follows: thinkers who are fundamentally indiscernible at a time cannot differ in respect of their *de re* thoughts at that time. So it can't be true, for example, that each of the amoebae hopes that *it* will be the well-fed one but doesn't have the same hope concerning its twin; since if there are two amoebae, they are fundamentally indiscernible before the split. Differences between thinkers can provide the basis for self-locating uncertainty at a certain time only when they are grounded in that time.

Let me attempt to argue for this principle without simply appealing to the intuition that there is nothing to be uncertain about in fission cases. Suppose that it is true that if a person's brain were divided in two and the halves transplanted into donor bodies, there would have been two people all along. Still, there could be a linguistic community the members of which use 'person' in such a way that 'no person could survive brain bisection' is true in their language. When members of this community who are about to undergo brain bisection utter the words 'I will not survive this brain bisection', they speak the truth, since the meaning they assign personal pronouns like 'I' fits with the meaning they assign 'person'. But the linguistic community to which a person belongs makes no difference to whether that person will survive a given operation.¹⁶ So when these people say 'I will not survive any brain bisection', they are not saying of *themselves* that they will not survive any brain bisection, since if they were doing that, they would be saying something false.

The members of this community will not be tempted by the thought that there is something to be uncertain about when they face a brain bisection. If there is something they should be uncertain about, they

are incapable of expressing it in their language. But their language isn't in any important sense objectively worse than ours.¹⁷ If we think that there are self-locating questions with right answers that they can't even think to themselves in their language, then we must think something parallel about our own language. As they stand to us, so we must stand to the members of other imaginary communities, in which the word 'person' is used in such a way that temporary coincidence of "persons" is more common than it is for us. For example, we could imagine a community which uses 'person' in such a way that for any set of people no two of whom exist at the same time, they can truly say 'there is a unique person who exists just when any of these entities exist, and coincides at each time with whichever of them exists at that time'. We will have to say that when a member of this imaginary community utters a sentence like 'I fought on the side of the French in 1814', he expresses a thought concerning which he ought to be uncertain, no matter how much he might learn about how the world is impersonally speaking, and no matter how much he might learn about his own location in the world. And we should be uncertain in the same way as him, although of course we can't use these words to express it. But this is absurd. There just isn't any subject matter for uncertainty in this vicinity.

So much for amoebae. As far as I can see, all these considerations carry over to the case of "thick" ordinary objects at *H*-worlds. If the right thing to say about the person in our measurement situation is that she won't exist at all after the measurement, or will exist in both the 'seeing a flash' region of configuration space and the 'not seeing a flash' region, there isn't even *prima facie* anything more for her to be uncertain about. If she will exist in one of these regions and not the other, it must be indeterminate which, so again there is no room for uncertainty. The only way to provide for uncertainty is to say that there are many people there all along. But if we want to allow for uncertainty about the future by embracing Overpopulation, we had better not say that the people who exist before the measurement share their pre-measurement temporal stages; or more generally, that they

¹⁶Pace Johnston (1989).

¹⁷For a defence of the claim that our "concept of personal identity" isn't objectively better than other possible concepts, see Sider 2001.

are fundamentally indiscernible at times before the measurement. Given our picture of ordinary objects as occupying regions of configuration spacetime, this will mean that the many people must occupy different sets of c -points even before the experiment. The only way I can see for this to be the case is for them to be thin, occupying only one c -point each per time. For the only temporally local fact about the high-amplitude region of the wavefunction that could provide any principled basis for saying that it is occupied by many distinct people is the simple fact that it is larger than a single point.

So far I have been confining my attention to a case where the wavefunction spreads in an essentially symmetric way. But there is good reason to think that whatever we say about fate of ordinary objects in this sort of case will carry over to cases of asymmetric spreading—e.g. an experiment like the previous one with a mirror that is more or less than half-silvered. We can't complain that it would be *arbitrary* to say that the person who ends up seeing a flash in this case is, determinately, the unique original person, while the person who doesn't see the flash did not exist before the experiment. Still, it is awfully hard to believe that this is how things are with ordinary objects at H -worlds. If I had good reason to believe that this view were true and that the actual world is a H -world, I should believe that *every time* I do the experiment with a less-than-half-silvered mirror, I will see the flash. I should confidently expect the front page of tomorrow's newspaper to contain reports of physicists' utter amazement at having obtained utterly freakish runs of the outcomes counted as "most likely" by the standards of ordinary quantum mechanics. By the same token, if I ever find myself seeming to remember doing the experiment without seeing a flash, reading reports of the empirical success of quantum mechanics, etc., I will have decisive evidence that if the world is a H -world, I have only very recently come into existence, with apparent memories of a lifetime's worth of experiences that were actually had by someone else.¹⁸

¹⁸I think we can reasonably draw from this the further conclusion that if we were to conclude that the ordinary objects at H -worlds follow the "highest amplitude" trajectories, we should conclude that it is very unlikely that the actual world is a H -world. But I cannot defend this claim here.

If we set the possibility of this sort of view aside, there is no important difference between symmetric and asymmetric splittings of the wavefunction. So if we want to say that ordinary objects at H -worlds are thick, we will have to say that there is never any room for uncertainty about what one will be seeing after a measurement for which one knows the range of possible outcomes. For, knowing that, you know enough to know that each of the possible outcomes will soon be witnessed by *someone*; and you know that each of these people will have an equally good claim to be you.

And indeed, the case is much worse than that. For "measurements" are just an especially clear case of a kind of spreading that the wavefunction undergoes all the time. If we are thick objects at H -worlds, we are *constantly* undergoing fission. Indeed, there is excellent reason to think that this fission will be so pervasive that a rather large proportion of the fates that could conceivably befall us will in fact befall one of our fission products. But I won't attempt to argue for this here.

The idea that the world works like that seems to me to be as deeply undermining of our ordinary pattern of thought about ourselves and our futures as any doctrine that has ever been worth taking seriously on scientific grounds. Take fear and hope, for example. When I hope that things will go well for me in the future, or fear that they will go badly, these attitudes make sense only if I am genuinely uncertain about whether things will go well for me in the future. The thought that there is nothing to be ignorant about here—that there will inevitably be people for whom things are going well *and* people for whom things are going badly, and that these people will inevitably have an equal claim to be me—undermines a presupposition of my hope and fear. To expunge this kind of hope and fear from my practical thought would make it unrecognisably different.

5 Bohmian threads

To sum up the main conclusion of the previous section: ordinary objects at H -worlds differ radically as regards their relation to the future from ordinary objects as we normally conceive them, unless they are thin. And even if we say that ordinary objects are thin, we will end up with the same conclusion unless there is some non-arbitrary principle by which, given a point of configuration space c

and two times t, t' , we can pick out a unique point of configuration space as the one occupied at t' by any ordinary objects that occupied c at t (and still exist at t'). Well, *is* there any such uniquely natural way to divide the configuration spacetime of a H -world up in this way? You could be forgiven for assuming, on the basis of the literature on “many worlds” interpretations of quantum mechanics, that it is well known that there isn’t.¹⁹ But you would be quite wrong to assume this. It turns out, delightfully, that there *is* a well-known, and uniquely natural, way to partition the configuration spacetime of any H -world (or indeed, any world describable in terms of a wavefunction) into a system of “threads” with one member per time.

One way to specify this uniquely natural system of threads is to give a recipe for reading a uniquely natural vector field on configuration space, \mathbf{v}^{ψ_t} , off the wavefunction ψ_t at an arbitrary time t . Our threads can then be defined as the *flow curves* of this field—i.e. as maximal functions f from times to points of configuration space which are such that, for each time t at which $f(t)$ is defined, $df(t)/dt = \mathbf{v}^{\psi_t}(f(t))$.²⁰ If we allow ourselves to continue to ignore spin, this vector field can be defined as the *gradient of the phase of the wavefunction*. Equivalently,²¹

$$(GE) \quad \mathbf{v}^{\psi_t} = \text{Im} \frac{\nabla \psi_t}{\psi_t}.$$

Talk of gradients makes sense, of course, only if we endow configuration space with a geometry. The one we need is the standard one,

¹⁹For example, Wallace (2002: 9) asserts, concerning a family of views of which the claim that each point of configuration space contains its own “world” of ordinary objects at each time is a special case, that ‘we may decompose the universal state into worlds at some given instant of time, but we cannot track the individual worlds when we evolve the state forward in time in any satisfactory way.’

²⁰This makes sense only relative to a choice of temporal unit; but we get the same threads no matter what temporal unit we choose. There isn’t any uniquely natural definition of a vector field on configuration space at a time in terms of the facts about the wavefunction at that time; but there *is* a uniquely natural definition of a function from temporal durations to vector fields.

²¹To see why the expression on the right hand side of GE is equivalent to the gradient of the phase, decompose ψ_t as $R_t e^{iS_t}$. Then $\text{Im} \nabla \psi_t / \psi_t = \text{Im} \frac{R_t \nabla e^{iS_t} + e^{iS_t} \nabla R_t}{R_t e^{iS_t}} = \text{Im} \frac{R_t e^{iS_t} i \nabla S_t + e^{iS_t} \nabla R_t}{R_t e^{iS_t}} = \text{Im}(i \nabla S_t + (\nabla R_t / R_t)) = \nabla S_t$.

touched on briefly in section 1 above, according to which the distance between two points of configuration space which differ only as regards where they put a single particle is equal to the product of the distance between the points of space where they put the particle and the square root of the particle’s mass.^{22,23}

\mathbf{v}^{ψ_t} is known as the *Bohmian velocity field*, after its use in the theory known variously as Bohmian Mechanics, the de Broglie-Bohm interpretation of quantum mechanics, and the pilot-wave theory. (Bohm 1952a; Goldstein 2002). As we have already seen, this theory posits facts about the *locations* of particles, over and above the facts which ground the description of the world in terms of its wavefunction. The location-facts pick out a single point of configuration space as special at each time: it is the point that puts each particle exactly where it is in fact located. Thus, the law governing the way particles move can be expressed as a claim about the shape of the path through configuration spacetime traced by the special point (the “world particle”). The answer, according to Bohmians, is that the velocity of the special point is always equal to the value of \mathbf{v}^{ψ_t} at the special point. Thus, each of the threads through configuration spacetime defined by the Bohmian velocity field corresponds to a distinct Bohmian world.

The fact that everyone who has accepted the general framework of the Bohmian approach has taken the velocity of the world-particle to be given by \mathbf{v}^{ψ_t} is thus indirect evidence that \mathbf{v}^{ψ_t} is indeed the uniquely most natural vector field on configuration space associated with ψ_t . For if there were several roughly equally natural ways to read vector fields on configuration space off wavefunctions, this agreement among advocates of Bohm-like theories would be rather surprising.

In fact the agreement is not surprising: Bohmians have given con-

²²Given units for mass and distance that fit with our chosen unit of time, in the sense that $\hbar = 1$.

²³The threads defined by \mathbf{v}^{ψ_t} do not always extend all the way into the past and the future. Sometimes they approach the limits of the configuration space in finite time; sometimes they bump into points where the wavefunction is zero, where \mathbf{v}^{ψ_t} is undefined. But it has been shown (Berndl et al. 1995) that the set of points in configuration space at any time that are on threads that do not extend all the way into the past and future is always of measure zero, so for most purposes we can ignore this possibility.

vincing arguments (Dürr et al. 1992) which show that at worlds where the wavefunction evolves in accordance with Schrödinger’s equation, \mathbf{v}^{ψ_t} is indeed a uniquely natural vector field on configuration space (modulo a choice of temporal unit). These arguments come in two different kinds. The first kind suggest that the Bohmian velocity field is extremely, and perhaps uniquely, *absolutely* simple, in a sense that could very crudely be measured by the length of the definition of “belongs to the same thread as” in terms of physically basic predicates. The second kind suggest that the Bohmian velocity field is *simple relative to the laws*—it fits naturally with Schrödinger’s equation in a way that would make it stand out from its competitors even if there were no difference in the length of their definitions in physically basic terms. (Consider the sense in which the relation ‘occurs seventeen instants of time after’ is more “natural” than the relation ‘occurs thirteen instants of time after’ at a world where time is discrete, and the laws say that objects randomly change some of their properties every seventeen instants.)

The arguments for the absolute simplicity of the Bohmian velocity field turn on the fact that the assignment of wavefunction-values to points of configuration space at a time involves various kinds of arbitrariness. Once we take this into account, we will see that the brevity of the formula defining a vector field is frequently misleading as a guide to the question how complex the vector field really is, when defined in terms of physically basic relations. Definitions of vector fields that might initially look not to be any more complicated than \mathbf{v}^{ψ_t} turn out to depend on aspects of the wavefunction ψ_t whose definition in fundamental terms requires us to make some sort of arbitrary choice. This can happen in at least three different ways.

The first kind of arbitrariness in the assignment of wavefunction-values is in the assignment of global phases: the physically basic relations don’t allow for any remotely natural way to give sense to the claim that ψ , rather than $\gamma\psi$, is the wavefunction at t , where γ is some complex number of modulus one.²⁴ Because of this, any

²⁴The *amplitude* of the wavefunction can be fixed by stipulating that the integral of $|\psi|^2$ over the whole of configuration space is 1. This makes sense only given a metric on configuration space, but a metric for configuration space is determined by

definition of a vector field which yields different results for ψ and $\gamma\psi$ really only defines a large family of different vector fields, no one member of which is simpler or more natural than any of the others. This means that, although \mathbf{v}^{ψ_t} might initially look less natural than the fields defined by $\text{Im}\nabla\psi_t$ or $\text{Re}\nabla\psi_t$, for example, the reverse is in fact true.

The second kind of arbitrariness in the assignment of wavefunction-values is in the choice between a function ψ_t and its complex conjugate ψ_t^* . This choice makes sense only relative to a choice of temporal direction: if we say that the wavefunction at t is ψ_t when we are thinking of one direction as the “forwards” direction in time, we will say that it is ψ_t^* when we are thinking of the other direction as the “forwards” direction. Thus, any recipe for reading vector fields off wavefunctions that yields different results for ψ_t and ψ_t^* is really a recipe for reading *unordered pairs* of equally simple and natural vector fields off wavefunctions. Moreover, while the definition of the *pair* may be simple, picking out any one member of the pair will require a definition that either contains an completely arbitrary element (e.g. one that arbitrarily nominates the direction from t_1 to t_2 as the forwards direction) or is extremely complex (e.g. one that precisifies and spells out in fundamental terms the idea that the forwards direction in time is the direction of increasing entropy). Because of this, the vector fields we might attempt to define using formulae like $\nabla|\psi_t|$ or $\text{Re}(\nabla\psi_t/\psi_t)$ are not actually anything like as simple as \mathbf{v}^{ψ_t} .

The third kind of arbitrariness in the assignment of wavefunction-values is present only if the underlying spacetime is Galilean (Neo-Newtonian) rather than Newtonian. As we saw in section 1, we will assign different wavefunction-values depending on which Galilean reference frame we think of as being at rest. A genuinely simple definition of a division of configuration spacetime into threads should not require us to make this choice. The definition of \mathbf{v}^{ψ_t} has this property. We can see this easily in the case of a single particle of unit mass moving in one spatial dimension. Suppose we have two Galilean co-ordinate systems that assign the same co-ordinates to each point of space at time t , but such that the origin of the second system

our choice of unit for time.

is moving at a velocity of v relative to the origin of the first system. Then, if the wavefunction at t relative to the first co-ordinate system at the point of co-ordinate space that puts the particle at the point of space with co-ordinate x is $\psi_t(x)$, the wavefunction relative to the second co-ordinate system is

$$\psi_t^v(x) = e^{-ivx} \psi_t(x).$$

Thus, relative to the second co-ordinate system, $\mathbf{v}^{\psi_t^v}(x) = \text{Im}(\nabla \psi_t^v(x) / \psi_t^v(x)) = \text{Im}((-iv e^{-ivx} \psi_t(x) + e^{-ivx} \nabla \psi_t(x)) / e^{-ivx} \psi_t(x)) = \text{Im}(-iv) + \text{Im}(\nabla \psi_t(x) / \psi_t(x)) = \mathbf{v}^{\psi_t} - v$.²⁵ This is exactly what it needs to be for the two co-ordinate systems to generate the same division into threads. Other candidate definitions of vector fields in terms of ψ_t which may initially look not much less simple than \mathbf{v}^{ψ_t} , and which respect the arbitrariness of the assignment of phases and temporal orientations—for example, the fields given by $2\mathbf{v}^{\psi_t}$, or $-\mathbf{v}^{\psi_t}$, or $|\psi| \mathbf{v}^{\psi_t}$, or $|\psi|^2 \mathbf{v}^{\psi_t}$ —don't have this property. This means that any attempt to use one of these definitions as the basis for a definition of a “belongs to the same thread as” relation in fundamental terms would have to incorporate a recipe for picking out a single one Galilean frame of reference at each possible world. Any such recipe would require a considerable amount of complexity or arbitrariness, which is not needed in the definition “belongs to the same Bohmian thread as”.²⁶

So much for the *absolute* simplicity of \mathbf{v}^{ψ_t} . Besides all this, there is also the matter of its simplicity relative to the laws. It would take us too far afield to survey all the senses in which the Bohmian velocity field “fits naturally with” Schrödinger's equation.²⁷ I will just talk about one central kind of “fit”, which will turn out to be quite important for us later on. It consists in the fact that the squared-amplitude measure ($\int |\psi_t(q)|^2 d^{3N}q$) on configuration space at a time

²⁵Here we are relying on the fact that the mass of the particle is 1, so that the gradient on the configuration space is the same as the gradient on ordinary space.

²⁶I suppose the least complex way to specify a privileged frame of reference, at a world with finitely many particles, is as the frame in which the expectation of the centre of mass of the universe is at rest. But even this seems pretty complex.

²⁷See, for instance, the discussion in Bohm 1952b on the relation between Schrödinger's equation and the classical Hamilton-Jacobi equation.

is *equivariant* under the evolution defined by the Bohmian velocity field (Dürr et al. 1992: 14–15). What this means is that the squared-amplitude measure of some region R_t of configuration spacetime all of whose points belong to t will always be the same as the squared-amplitude measure of any other region $R_{t'}$ which comprises all and only those c -points belonging to t' that belong to the same Bohmian thread as some point in R_t .

This fact suggests an analogy with fluid mechanics: if we think of the wavefunction as a fluid of density $|\psi|^2$ moving through configuration space, then we can think of the Bohmian threads as the paths of the particles of which the fluid is composed. The preservation of the squared-amplitude measure under the Bohmian evolution is analogous to the preservation of the mass of a given volume of fluid under the evolution given by the particle velocities.²⁸ This analogy is quite helpful for conveying an intuitive sense for the shape of the Bohmian trajectories. For example, we can immediately see that at a world where the wavefunction is concentrated in one region at an earlier time and in two regions at a later time, most of the threads that start in the single high-amplitude region must end up in one of the two high-amplitude regions. In general, paths that start out in high-amplitude regions will tend to be “swept along” by the wavefunction in such a way as to stay in high-amplitude regions.

Even if we completely set aside the considerations springing from the need to respect the Galilean symmetries, this fact would make the threads defined by \mathbf{v}^{ψ_t} stand out as far more natural than those defined, say, by $2\mathbf{v}^{\psi_t}$ or $-\mathbf{v}^{\psi_t}$. For these other velocity fields don't preserve *any* even remotely natural measures on configuration space. True, the Bohmian evolution isn't the only one under which $|\psi|^2$ is equivariant: if \mathbf{w}_t is any time-dependent vector field on configuration space for which $\nabla \cdot |\psi_t|^2 \mathbf{w}_t = 0$, then $\mathbf{v}^{\psi_t} + \mathbf{w}_t$ will define another evolution under which $|\psi|^2$ is equivariant. But I don't know of any way to define such a nonzero \mathbf{w}_t without having to make all sorts of arbitrary choices or injecting a large measure of complexity.

²⁸The standard terminology reflects this (Dürr et al. 1992: 15): the Bohmian velocity is the “probability current” ($\text{Im}(\psi_t^* \nabla \psi_t)$) divided by the “probability density” ($|\psi_t|^2$), just as in fluid mechanics, velocity is current divided by density.

6 Bohmian threads and the identity of ordinary objects through time

My thesis, then, is that there are ordinary objects at H -worlds, and that they follow Bohmian threads. This is a bit vague: let me attempt to state it a bit more precisely. Where w is any H -world and S is any set of c -points that exist at w , let w_S be a “classical” world represented by S —that is, a world whose fundamental ontology comprises just the particles and spacetime points that exist at w , where there is a new physically basic binary “location” relation that a particle p bears to a spacetime point x iff, at w , there is a c -point in S that puts p at x . (w_S need not, of course, be a world where the laws of classical mechanics are true.) Armed with this definition, we can state a more precise version of the thesis as follows:

- (*) Let w be any H -world; f any function from the c -points at w to sets of particles, and K any kind of ordinary object (people, tables, trees. . .). Then, at w , there is a K composed, relative to any c -point c , of just the particles belonging to $f(c)$ iff
 - (i) The c -points for which $f(c)$ is nonempty all belong to the same Bohmian thread B .
 - (ii) At w_B there is a K that is composed, at any time t at which it exists, of exactly the particles in $f(c_t)$, where c is the member of B belonging to time t .

In the previous section I argued, essentially, that (*) cannot fairly be convicted of arbitrariness. In this section I will give two arguments that (*) is true.

The first argument builds on the discussion from section 4. Let’s assume a plenitudinous ontology on which every function from c -points to sets of particles has a corresponding object, so that our goal is to say which, if any, of these objects which of these objects, if any, fall under predicates like ‘person’, ‘tree’, ‘table’, etc. In trying to answer this question we will of course be looking for interpretations on which certain sentences that express commonsense beliefs about ordinary objects come out true. But not just every sentence we are inclined to assert carries equal weight when we are trying to settle the extensions

of our predicates at an arbitrary possible world. Some seem to carry no weight at all: if the world we are dealing with consists entirely of diffuse hydrogen gas, the fact that an interpretation makes ‘there are dogs’ come out false is no objection at all to that interpretation. Some sentences, perhaps, are “absolutely analytic”, in the sense that any assignment of extensions to predicates at any possible world on which they turn out to be false can decisively be ruled out. But there are also many sentences that are *somewhat* analytic, in the sense that interpretations on which they are true at a given world are to be preferred, *ceteris paribus*, to interpretations on which they are false. This category includes, I think, claims like the following: ordinary objects often persist through appreciable stretches of time; there is often a fairly determinate fact of the matter about what will happen to a given object in the future; people are often capable of referring to themselves and other things in ways that leave it open what will happen to them in the future. Other things being equal, we should therefore favour an interpretation on which the ordinary objects at a H -world are confined to Bohmian threads over any of the interpretations on which these somewhat-analytic claims about ordinary objects and the future come out false at a such a world.

Are other things equal? They wouldn’t be if (*) forced one to deny some other comparably important piece of common sense. For example, if it turned out that the apparent memories of people who inhabited Bohmian threads would be systemically unreliable as a guide to what had happened to them in the past, or that ordinary inductive methods would be systematically unreliable as a guide to what would happen to them in the future, that would be a strong consideration against (*). But nothing like this is the case. Two different sorts of results that Bohmians have established can help to reassure us on this point, by showing that the lives of objects that follow Bohmian threads at the right sort of world are very much like our lives as we normally take them to be. First, that Bohmian threads approximate possible classical histories under just those circumstances—i.e. potentials that vary sufficiently gradually—in which classical mechanics has actually been found to be a good predictive tool (Bohm 1952a). This means that under those kinds of circumstances, memory and induction will be

just as reliable for the inhabitants of Bohmian threads at H -worlds as they would be for the inhabitants of worlds where classical mechanics is true. Second, that when one sets out to perform a series of idealised measurements, one's expectations should be exactly what orthodox quantum mechanics says they should be, provided that one's credences about the relation between the wavefunction and the locations of the particles in the systems to be measured conform to the squared-amplitude measure (Dürr et al. 1992: section ?). This means among other things that at most threads (in the squared-amplitude sense of 'most'), extrapolation from the frequency with which measurement-results have been obtained is a reliable guide to the future. A full exploration of the significance of this fact is beyond the scope of this paper. Suffice it to say that while H -worlds will, if (*) is true, contain *some* people who live lives within which memory and induction are misleading, they are atypical in the most natural sense of 'typical' that applies to H -worlds, and there is no reason to think that any other theory about the kinds of ordinary objects that exist at such worlds does any better than (*) in this regard.

The strategy of the second argument is to start with a world where it is clear that each Bohmian thread has a corresponding collection of ordinary objects, and subtract structure in several steps until we are left with a H -world, arguing at each step that the subtraction couldn't make a difference to the facts about ordinary objects.

The starting point—call it World 1—is a world where, in addition to Neo-Newtonian spacetime and configuration spacetime, there are many disjoint families of particles bearing a three-place location relation to times and points of space. There is a physically basic relation C (for "counterpart") which puts each family in one to one correspondence with any other family. Points of configuration space don't distinguish between counterpart particles; thus it makes sense to ask whether a point of configuration space correctly represents the locations of any given family of particles. And there are so many families of particles that at each time, each point of configuration space correctly represents the locations of at least one family. The velocity of each particle is determined, in accordance with the Bohmian formula, by the facts about the wavefunction and the positions of all the

particles in its family. Surely at this world, each ordinary object is composed of particles belonging to a single family, and the objects within each family are just as they would be if the particles from that family were the only particles in the universe. For the particles at an ordinary Bohmian world with just one family compose ordinary objects; and the mere addition of any number of new families could hardly affect the existence of these objects.

World 2 is like World 1 except that its particles are not simple or fundamental: instead, they are composed out of instantaneous temporal slices, bound together by a new physically basic relation G (for "genidentity"). Each particle at World 2 corresponds at World 3 to a G -equivalence class of particle-slices; the counterpart relation C now holds between particle slices, in such a way that G -equivalent particle-slices are always C -equivalent. This could hardly have any significance as far as ordinary objects are concerned, unless it were impossible in general for the existence of ordinary objects to be constituted by facts about fundamental entities all of which are themselves instantaneous. I suppose some endurantists do regard this as impossible. But I can't see how such a view could be motivated from the standpoint of the view of ordinary objects which I have been presupposing in this paper, according to which the existence of a given ordinary object just *consists in* certain facts about relations holding among fundamental entities which are not themselves ordinary objects. And those who reject this view are already committed to denying that there could be ordinary objects if H were true.

World 3 is like World 2, except that it lacks the physically basic relation G . At World 3, the division of particle-slices into classes following Bohmian threads is only as natural as the definition of the relation 'belongs to the same Bohmian thread as' in terms of the physically basic relations which constitute the wavefunction. But, as we saw in section 5, this definition *is* natural—much more natural than any other definition of a partition of configuration spacetime that cuts across the partition into times. Could it really matter, as far as ordinary objects are concerned, whether the relation that generates the partition is itself physically basic, or merely definable in some uniquely natural way out of relations at some lower level? I don't think it could,

any more than it could matter to the facts about ordinary objects whether the property *being a proton* is itself physically basic or merely definable in some natural way in terms of physically basic properties and relations (e.g. among quarks). In general, if some structure of natural properties and relations is sufficient to constitute a domain of ordinary objects, it will still be sufficient to do so even if it turns out to be grounded in turn in some even more basic structure.

World 4 is like World 3, except that we add some new entities, “super-particles”: each particle-slice “belongs to” exactly one super-particle. The relation *C* is no longer primitive, but is analysed as “belonging to the same super-particle”. Moreover, the physically basic “putting” relation now relates *c*-points and *super*-particles to points of spacetime, and applies to particle-slices only in a derivative sense. Surely none of this ontological tinkering could make a difference to the facts about ordinary objects.

World 5 is like World 4, except that we get rid of the particle-slices altogether. We lose no information in doing this, since at World 4 the facts about particle-slices could be summed up very simply by the following law: for each *c*-point *c*, there is a family F_c of simultaneous particle-slices such that for any super-particle *p* and spacetime point *x*, *c* puts *p* at *x* iff a particle-slice in F_c belonging to *p* is located at *x*. In other words, the particle-slices at World 4 were working just like tropes of the “putting” relation; the transition from World 4 to World 5 works like the transition from a world where trope theory is true to one in which it isn’t. Could this sort of information-preserving deletion matter to the facts about ordinary objects? For my part, I don’t see how it could.

Note that *H* is true at World 5: the only difference is that I have been calling the occupants of a certain structural role “super-particles” rather than “particles”. But that’s a merely terminological difference.

The only principled reason I can think of for resisting the move from World 4 to World 5 would be a commitment to the principle of uniqueness of composition. An adherent of this principle has no problem finding entities at World 4 sufficient to compose all the ordinary objects that I claimed exist there; at World 5 it’s harder to do this. As I noted in section 3, one can do it if one is willing to count *c*-points and

perhaps also spacetime points among the parts of ordinary objects; but some will reject this as artificial.

But uniqueness of composition seems to me needlessly restrictive and completely unmotivated if it is taken as theory of how facts about what there is in a non-fundamental sense might be constituted by fundamental facts.²⁹ Think of entities like sets, propositions, events, political movements, musical compositions, holes, afterimages. . . . Surely at least some of these things exist only in a non-fundamental sense; but it seems hopeless to try to construct them as mereological sums of fundamental entities. And if these things aren’t mereological sums of fundamental entities, why should ordinary spatiotemporal objects be?

7 Do threadbound objects lack causal unity?

Bohmian threads of *c*-points at a *H*-world are not causally isolated from one another. Each thread is shaped as it is because of facts about the wavefunction in its neighbourhood—in other words, because of facts about certain relations among its *c*-points and those belonging to other nearby threads. But if it is to be true that ordinary objects can only see, touch, kill, etc. . . one another when they belong to the same thread, there had better be a really big difference between the kinds of causal relations that hold among threadmates and the kinds of causal relations that hold among objects from different threads. It would be disastrous for the proposed view if we had to conclude that many different teapots, belonging to different threads, play a roughly equal role in causing my current visual experience as of a teapot: for if that were so, it could not be true that the only teapot I see is the one that belongs to the same thread as I do. More generally, if there were no substantial difference in the teapots’ causal relations to other things, it would be hard to seriously maintain that there really are many teapots here rather than one; and likewise for people.³⁰

²⁹Of course, proponents of classical mereology generally *don’t* think of it in this way, since they generally reject the distinction between fundamental and non-fundamental quantification that I have been taking for granted. For a defence of the intelligibility of this distinction see Dorr 2005.

³⁰Thus Wallace (2002: p. 10) argues against thinking of people as confined to threads on the grounds that ‘if mental facts are supervenient on facts about brain

An adequate response to this worry would require a fully-fledged theory of causation. Here I shall just take a few tentative steps towards a defence of the claim that the causal relations among objects belonging to different threads are utterly different in character from the causal relations among threadmates, in a way that is adequate to explain why causally-based relations like seeing can only hold within a thread, and to undercut the argument that objects bound to a single thread would not be causally independent enough to be the ordinary objects of our acquaintance.

One way to approach questions about the causal relations among ordinary objects is to think about relations of counterfactual dependence. One might even claim that the truth of a counterfactual ‘if object *A* had been different, object *B* would have been different’ is sufficient for the existence of a causal relation between distinct objects *A* and *B*. At any rate, the conceptual relations between causal and counterfactual dependence are clearly intimate enough that if we could show that cross-thread relations of counterfactual dependence are very different in character from intra-thread relations of counterfactual dependence, that would be very good news for the claim that there is a similar difference as regards causal relations.

Suppose, then, that *H* is true, and that ordinary objects—among them the famous assassin Gavrilo Princip and his victim, the archduke Franz Ferdinand—are confined to a single thread. What would the world have been like, fundamentally speaking, if Princip hadn’t shot the archduke? What different pattern of physically basic relations among fundamental entities would count as a situation in which Princip didn’t shoot? Bear in mind that, thanks to the constant “teasing apart” of the threads under the influence of Schrödinger’s equation, the actual world surely already contains at least one person who lives a life extremely like Princip’s up to the moment of the shooting, but who refrains from shooting at the last minute.³¹ Moreover, the considerations that lead us to say this kind of thing would have been

configurations in single worlds’—here Wallace is using “worlds” as a name for entities playing much the same role as *c*-points—‘then each of us remains a thinking being only because of constant interference from the particles comprising the brains of countless neurologically identical parallel-world copies of ourselves.’

³¹*This needs more careful argument*

the same whether or not Princip had shot. So if *Princip* hadn’t shot, there would still have been at least one person who lived a life very like Princip’s actual life, including the shooting. Thus, the crucial difference between the actual world and the counterfactual one is a *haecceitistic* one: it has to do not with the overall pattern of qualitative properties and relations but with the identities of the objects that play the roles in the pattern. Actually, Princip is one of the shooters; counterfactually, he is one of the non-shooters. There is thus no immediately obvious motivation for positing any qualitative difference between the two possible worlds. Perhaps some merely *haecceitistic* difference at the fundamental level would be enough to constitute its being the case that Princip didn’t shoot. For example, we might decide to say that if Princip hadn’t shot, some of the *c*-points on the thread to which Princip actually belongs would have put certain particles at different spacetime points. Numerically the same *c*-points would have been represented by different functions from particles to spacetime points.³²

Whether the difference between the counterfactual situation and the actual one turns out to be *merely* haecceitistic will turn on what we decide to say about the “displaced” *c*-points—the ones that *actually* put particles at the points of spacetime where the *c*-points actually occupied by Princip *would* have put them, if Princip hadn’t shot. Would these *c*-points have continued to put particles where they actually put them, thereby violating the actual law which guarantee that no two *c*-points put all the same particles at all the same spacetime points? No. For if ordinary objects are constitutively bound to Bohmian threads, there is no basis for positing two distinct collections of ordinary objects corresponding to the doubled-up thread through configuration spacetime at such a world: there is nothing to make it be the case that one of two coincident *c*-points rather than the other is occupied

³²If we accept some sort of origin essentialism for ordinary objects, we will probably have to say that a world where Princip didn’t shoot would have to be haecceitistically different from the actual world all the way back to the beginning of time. For such a world would have to be one in which Princip still had his actual grandparents; and it is hard to see how a merely haecceitistic difference at the fundamental level that only involved times after their death could make a difference to whether his grandparents instantiate the qualitative property *having a grandson who is an assassin*.

by the same ordinary objects as any given earlier *c*-point. So we must either say, that the displaced *c*-points do not exist in the counterfactual situation, or perhaps that they switch places with Princip's *c*-points (thereby freeing us of the need to add new *c*-points to plug up the gap). Either way, I can see no motivation for disturbing the qualitative facts about the pattern of the phase and intensity of the wavefunction across configuration space (qualitatively described), or for relocating any *c*-points other than those on the thread to which Princip belongs and the ones they displace.

(There is a view, sometimes called "anti-haecceitism", according to which the haecceitistic facts supervene on the qualitative facts, so that things in a counterfactual situation where Princip didn't shoot the archduke would have to be different in some qualitative respect from things as they actually are. I consider this view to have been refuted by Robert Adams (1979). But note that although Lewis 1986a contains a section called 'Against Haecceitism', Lewis's doctrine is perfectly compatible with all I have said. By allowing, as he does, for objects to have counterparts distinct from themselves at their own possible world, Lewis allows that it is possible for the qualitative facts to be exactly as they actually are while certain objects have different properties from those they actually have. Given Lewis's theory of possible worlds, the "counterfactual situations" I have been talking about are not possible worlds, but something more like possible worlds taken together with functions onto their domains from the domain of the actual world.)

Let's suppose, then, that the differences at the fundamental level between the actual world and the world as it would have been if Princip hadn't shot are entirely, or at least mostly, haecceitistic. How would things have gone differently at the level of ordinary objects? Well, if a difference as regards where certain specific *c*-points put certain specific particles is enough to constitute a difference in what Princip is doing, it must certainly also be enough to constitute corresponding differences in the other ordinary objects that occupy those *c*-points. So, if Princip hadn't shot, Franz Ferdinand would not have been shot. And since Princip and Franz Ferdinand and all their threadmates would have continued to be ordinary objects, and thus entities

whose persistence through time is constitutively tied to the structure of Bohmian threads, their future histories would have been different too. Typically, threads which start out very close ours but in which the Princip-counterpart does not shoot evolve rather differently: the Serbia-counterpart is not invaded by the Austria-counterpart, and so on; pretty soon these threads become so different from ours that they cease to contain close counterparts of many actual people and events.³³ Thus, it doesn't seem like it should be any harder to find an account of counterfactuals that can vindicate our ordinary judgments in this case than it is to find one that can vindicate our ordinary judgments given a more "classical" set of assumptions about the way the world works.

And what would have happened to ordinary objects that are not threadmates of Princip's—for example, his many doppelgängers from nearby threads, whose lives are just like his in all manifest respects? A big attraction of the idea that the difference between the actual world and the world as it would have been if Princip hadn't shot is exclusively (or primarily) haecceitistic is that it allows us to say that none of those things would have been different in any way if Princip hadn't shot. For we surely don't need to "relocate" the *c*-points to which *they* belong in order to make it be the case that Princip doesn't shoot. While it is plausible enough to think that it is essential to Princip that he belongs to a thread made up of certain specific *c*-points—we *have* to think something like this to explain the difference between Princip and his many doppelgängers—there is no motivation at all for thinking that it is essential to Princip that he belongs to a thread which is *close* to certain other specific *c*-points. So we are left with the view that if Princip hadn't shot, *c*-points that are not on his thread (with the possible exception of those displaced by the relocation of Princip's *c*-points) wouldn't have been different, either in respect of where they put particles, or in respect of their wavefunction-constituting relations. If so, there would have been no difference whatsoever in

³³The sense of "typicality" on which this claim is plausible is typicality according to the squared-amplitude measure; this measure will probably have to figure somehow in the truth-conditions for counterfactuals (insofar as we can talk about truth-conditions for counterfactuals!) in order to generate results that fit with our ordinary intuitions. I don't see why it shouldn't, but I can't argue for this claim here.

the careers of the ordinary objects that belong to these *c*-points. And this strongly suggests that Princip's choice whether or not to shoot is *causally irrelevant* to the fate of these objects. For similar reasons, there is only one teapot whose shape and location is causally relevant to my visual experiences. If any of the teapots from neighbouring threads had been different in respect of shape or location, my visual experiences would have been just the same: the thread containing the teapot in question would no longer have been a neighbour of my thread.³⁴

It's not clear that this view lets us say that ordinary objects belonging to one thread are completely causally isolated from ordinary objects belonging to other threads. If the facts about the phase and amplitude of the wavefunction at *c*-points in the neighbourhood of the one to which we belong had been different, different things would have happened to us. And *occupying a c-point with such-and-such phase and amplitude* is a physical property of ordinary objects—and not a straightforwardly extrinsic one, given the way *c*-points enter into the structure of ordinary objects—albeit one which plays only a very indirect role in constituting any of their manifest properties. So perhaps ordinary objects from other threads can make a difference to us; but if they can, it is only by being different in these obscure respects while remaining the same in ordinary respects like shape and location. Moreover, to have any effect on us, a difference of this kind must involve a *continuous infinity* of ordinary objects, including the complete populations of all the *c*-points in some open neighbourhood of a *c*-point to which we belong. No change involving just a single ordinary

³⁴The only worry has to do with the objects belonging to the “displaced” *c*-points. If we say that those *c*-points would have disappeared or swapped places with the ones belonging to our thread, it seems we'll have to say that if Princip hadn't shot, someone who actually lives a Princip-like life but doesn't shoot either wouldn't have existed at all, or would have shot. This seems awkward; although we can perhaps avoid drawing surprising consequences about Princip's causal powers to influence goings on at other threads by going on to say that there is no *specific* Princip-doppelgänger who wouldn't have existed, or would have shot, if Princip hadn't shot. More generally, I'm inclined to think that the kind of reasoning about counterfactuals which leads one to take seriously questions about the fate of the displaced *c*-points is not a good guide to the facts about causation: see the final paragraph of this section for more discussion.

object belonging to another thread, or even any countable collection of such objects, can make any difference to what happens to us. So it's clear that the causal relations across threads are at best extremely attenuated in comparison to causal relations within threads.

So far, I have been thinking through the implications of the idea that the differences at the fundamental level between the world as it actually is and the world as it would have been if Princip hadn't shot are exclusively (or at least primarily) haecceitistic ones. This is not the only possible view, however. The main alternative that I can see is a view that hews more closely to Lewis's (1979b). On this view, a counterfactual situation in which Princip doesn't shoot would be one in which the qualitative facts about the wavefunction were initially the same as they are in the actual world, but in which some small exceptions to Schrödinger's equation occurred in 1914, in such a way as to make the guiding equation send the thread whose pre-1914 portion was just like the pre-1914 portion of our thread in the actual world off into a region of configuration space corresponding to the absence of shooting.

If *that's* how things would have been at the fundamental level if Princip hadn't shot, we will have to say that not just our thread, but all the others—or at least those close to ours in configuration space—would have evolved differently. For any difference in the wavefunction that could affect the evolution of our thread via the guiding equation would have to involve some open neighbourhood in configuration spacetime, and hence also affect the evolution of neighbouring threads. This would mean that if Princip hadn't shot, many of Princip's actual doppelgängers from other threads wouldn't have shot either, and many of Franz Ferdinand's doppelgängers wouldn't have died that day. Far from allaying the worry that causal relations between objects from different threads are not as different in character as they would have to be for it to be correct to think of ordinary objects as belonging to different threads would not be allayed, this account of the relations of counterfactual dependence tends to exacerbate it.

How are we to decide between the two accounts? Other things being equal, the haecceitistic account seems preferable, as the difference it requires between the actual and counterfactual situations seems so

much smaller. However, a proponent of the qualitative account could argue that the haecceitistic account requires us to say, counterintuitively, that if Princip hadn't shot, the past would have been different all the way back to the beginning of time. For if Princip hadn't shot, he and all the rest of us would still have been the sort of thing whose persistence in time is constitutively tied to the threads defined by the guiding equation. Since such threads never overlap, any difference in the configuration-space location of our thread at one time without a difference in the qualitative facts about the wavefunction requires a difference in the configuration-space location of our thread at all times.

Well, as a matter of fact, I don't think that it would be so bad to have to say that the entire past would have been different if Princip hadn't shot. For there is every reason to think that the phenomenon of *chaos* is just as pervasive in Bohmian mechanics as it is in classical mechanics. In many ordinary complex systems, tiny differences very rapidly blow up into huge differences. Since the dynamics are continuous, this means that tiny *enough* initial differences will stay tiny for an arbitrarily long period, and then very suddenly blow up into huge differences. Human decision-making seems eminently like the sort of arena in which this sort of extreme sensitivity to initial conditions should prevail. If so, there will be threads that are extraordinarily close to ours all the way up to just before the time of the shooting, but which then suddenly diverge in such a way as to make the counterpart of Princip refrain from shooting; so the differences we need to posit in the past can all much too tiny to matter to any ordinary or observable matters. This sort of proposal seems in many ways more attractive, both in the present context and in that of classical mechanics, than Lewis-style "miracles".

But even if we end up thinking that we have to endorse the qualitative account of counterfactuals at *H*-worlds, I think we will still have a good shot at defending the view that the causal relations between threadmates are very different and more extensive than those between objects belonging to different threads. For even if we find ourselves forced to agree that a miracle would have had to occur for Princip not to have shot, or for me not to have raised my arm just

now, I don't think we'll want to be committed to the causal relevance of the various surprising counterfactuals we'll find ourselves committed to by thinking through what these miracles will have to be like to have the desired effect. Suppose, for example, that the world is almost deterministic, but that every trillion years or so a particle will give a little Epicurean swerve. There is a god who loves order and regularity so fervently that whenever there is a swerve, he waits a few seconds to collect his energies, and then annihilates the entire planet upon which the swerve occurred. If we take Lewis's theory of counterfactuals seriously, we'll find ourselves having to say that if I hadn't raised my arm, the planet would have been annihilated. And if we take the causal relevance of these counterfactuals seriously, we'll find ourselves saying that my raising my arm sustained the planet in existence. This seems bad. Moral: if a "miracles" theory of counterfactuals under determinism is true, counterfactual dependence does not always entail causal dependence; in particular, the entailment fails for counterfactuals that are true only because of the inevitable side effects of any miracle sufficient to bring about the truth of the antecedent.³⁵

I conclude that worries about the causal independence of threads do not sink the proposed view of ordinary objects at *H*-worlds.

Let me finish by saying what I have done and pointing out what I

³⁵Lewis (1979b: 40) attempts to block the conclusion that there is widespread backwards causation by claiming that there generally isn't any *specific* way that things would have been different if any given event hadn't occurred. As a general method for blocking implausible conclusions about causation, this seems insufficient. It is not hard to describe cases in which Lewis's criteria for closeness of worlds entail that all the worlds where a certain event fails to occur that are closest to actuality contain the same divergence miracle, in which case a counterfactual theory of causation will count the event in question as a cause of the earlier non-occurrence of this miracle. However, Lewis's move might be enough to block the derivation of conclusions about inter-thread causation from the "miracles" account of counterfactuals about ordinary objects at *H*-worlds. Since the open regions of configuration space in which Schrödinger's equation is miraculously violated can be made arbitrarily small, it is arguable that there won't be any *specific* archduke other than Franz Ferdinand who would have stayed alive if Princip hadn't shot, despite the fact that if he hadn't shot, continuously many archdukes who actually die would have stayed alive. If so, Princip's actions will no longer count as a cause of the death of any particular archduke other than Franz Ferdinand.

have not done. I have argued that H —and a fortiori, the hypothesis that the actual world is exhaustively characterised by a wavefunction that never collapses—is *consistent* with the facts about ordinary objects as we know them. I have not tried to argue that we have good reason to believe (to any non-negligible degree) that anything like H is true; I haven't even tried to argue that we would have good reason to believe this if it weren't for those recalcitrant aspects of our experience that force us to go beyond straight quantum mechanics to quantum field theory and ultimately to some kind of quantum gravity. Nor have I explained how to construe orthodox quantum mechanics—in particular, its talk of “probability”—in such a way as to be consistent with H . So what I've said doesn't yet amount to an “interpretation of quantum mechanics”, as this phrase is usually understood. But it's a start.

References

- Adams, Robert M. (1979). 'Primitive Thisness and Primitive Identity.' *Journal of Philosophy* 76: 5–26.
- Albert, David (2005). 'The Foundations of Physics.' In Jackson and Smith (2005), 848–880.
- Albert, David Z. (1996). 'Elementary Quantum Metaphysics.' In James T. Cushing, Arthur Fine and Sheldon Goldstein (eds.), *Bohmian Mechanics and Quantum Theory: An Appraisal*, 277–284. Dordrecht: Kluwer.
- Berndl, Karin, Detlef Dürr, Sheldon Goldstein, Giulio Peruzzi and Nino Zanghì (1995). 'On the Global Existence of Bohmian Mechanics.' *Commun. Math. Phys.* 173: 647–673. Available online at <http://arXiv.org/ps/quant-ph/9503013>.
- Bohm, David (1952a). 'A Suggested Interpretation of Quantum Theory in Terms of “Hidden” Variables, I and II.' *Physical Review* 85: 166–193.
- (1952b). 'A Suggested Interpretation of the Quantum Theory in Terms of “Hidden” Variables, I.' *Phys. Rev.* 85: 166–179.
- Deutsch, David (1985). 'Quantum Theory as a Universal Physical Theory.' *International Journal of Theoretical Physics* 24: 1–41.
- DeWitt, Bryce (1970). 'Quantum Mechanics and Reality.' *Physics Today* 23: 30–35. Reprinted in DeWitt and Graham 1973.
- DeWitt, Bryce and Neill Graham (eds.) (1973). *The Many-Worlds Interpretation of Quantum Mechanics*. Princeton: Princeton University Press.
- Dorr, Cian (2004). 'Non-Symmetric Relations.' In Dean Zimmerman (ed.), *Oxford Studies in Metaphysics*, volume 1, 155–192. Oxford: Oxford University Press. Available online at <http://www.pitt.edu/~csd6/papers/nsr.final.pdf>.

- (2005). ‘What We Disagree About When We Disagree About Ontology.’ In Mark Kalderon (ed.), *Fictionalist Approaches to Metaphysics*. Oxford: Oxford University Press. Available online at <http://www.pitt.edu/~csd6/papers/ontology.pdf>.
- Dürr, Detlef, Sheldon Goldstein and Nino Zanghì (1992). ‘Quantum Equilibrium and the Origin of Absolute Uncertainty.’ *Journal of Statistical Physics* 67: 843–907. Available online at <http://arXiv.org/ps/quant-ph/0308039>.
- Fine, Kit (2000). ‘Things and Their Parts.’ In Peter French and Howard K. Wettstein (eds.), *Midwest Studies in Philosophy XXIII: New Directions in Philosophy*, 61–74. Oxford: Blackwell.
- (2003). ‘The Non-Identity of a Material Thing and Its Matter.’ *Mind* 112: 195–234.
- Forrest, Peter (1988). *Quantum Metaphysics*. Oxford: Blackwell.
- Goldstein, Sheldon (2002). ‘Bohmian Mechanics.’ In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, Winter 2002 edition. Available online at <http://plato.stanford.edu/archives/win2002/entries/qm-bohm/>.
- Grabowska, Katarzyna, Janusz Grabowski and Paweł Urbański (MS). ‘The Schrödinger Operator in Newtonian Space-time.’ Available online at arXiv:math-ph/0611044v1.
- Jackson, Frank and Michael Smith (eds.) (2005). *The Oxford Handbook of Contemporary Philosophy*. Oxford: Oxford.
- Johnston, Mark (1989). ‘Relativism and the Self.’ In Michael Krausz (ed.), *Relativism: Interpretation and Confrontation*, 441–72. Notre Dame, Indiana: Notre Dame University Press.
- (2005). ‘Constitution and Identity.’ In Jackson and Smith (2005), 636–677.
- Lewis, David (1976). ‘Survival and Identity.’ In Amelie Rorty (ed.), *The Identities of Persons*, 17–40. Berkeley: University of California Press. Reprinted in Lewis 1983: ???
- (1979a). ‘Attitudes *De Dicto* and *De Se*.’ *Philosophical Review* 88: 513–543. Reprinted in Lewis 1983: 133–160.
- (1979b). ‘Counterfactual Dependence and Time’s Arrow.’ *Noûs* 13: 455–476. Reprinted in Lewis 1986b: 32–66.
- (1983). *Philosophical Papers*, volume 1. Oxford: Oxford University Press.
- (1986a). *On the Plurality of Worlds*. Oxford: Blackwell.
- (1986b). *Philosophical Papers*, volume 2. Oxford: Oxford University Press.
- Parfit, Derek (1971). ‘Personal Identity.’ *Philosophical Review* 80: 3–27.
- (1986). *Reasons and Persons*. Oxford: Oxford University Press.
- Robinson, Denis (1985). ‘Can Amoebae Divide Without Multiplying?’ *Australasian Journal of Philosophy* 63: 219–319.
- Saunders, Simon (1998). ‘Time, Quantum Mechanics, and Probability.’ *Synthese* 114: 373–404.
- Sider, Theodore (2001). ‘Criteria of Personal Identity and the Limits of Conceptual Analysis.’ In James Tomberlin (ed.), *Philosophical Perspectives 15: Metaphysics, ?* Cambridge, Mass.: Blackwell.
- Wallace, David (2002). ‘Worlds in the Everett Interpretation.’ *Studies in the History and Philosophy of Modern Physics* 33: 637–661. Available online at <http://users.ox.ac.uk/~mert0130/papers/everett.pdf>.