

An Awkward Symmetry: The Tension between Particle Ontologies and Permutation Invariance*

Benjamin Jantzen^{†‡}

Physical theories continue to be interpreted in terms of particles. The idea of a particle has required modification for successive theories but remains central to scientific explanation. Particle ontologies also have the virtue of explaining basic epistemic features of the world and so remain appealing for the scientific realist. However, particle ontologies are untenable when coupled with the empirically necessary postulate of permutation invariance—the claim that permuting the roles of particles in a representation of a physical state results in a representation of the same state. I demonstrate that any permutation invariant theory is incompatible with a particle ontology.

1. Introduction. A prominent realist interpretation of physics consists of the following claims: (i) material objects are composed of particles, (ii) these particles belong to a finite number of types, and (iii) particles of the same type are perfectly indistinguishable. These assumptions are manifest in the explanatory language of physics, chemistry, and molecular biology. I maintain that they are mutually inconsistent.

At the scale of everyday experience, we perceive discrete objects to which

*Received November 2009; revised May 2010.

†To contact the author, please write to: Department of Philosophy, Baker Hall 135, Carnegie Mellon University, Pittsburgh, PA 15213; e-mail: jantzen@cmu.edu.

‡I am grateful to John Earman, Jeremy Butterfield, Bryan Roberts, Peter Spirtes, and Kristina Jantzen for valuable discussion and criticism of previous drafts of this article. I am particularly indebted to Clark Glymour for helping me to clarify the logical structure of my argument, to two anonymous referees for their careful reading and insightful critique, and to Mara Harrell for guiding the development of the ideas reported here.

Philosophy of Science, 78 (January 2011) pp. 39–59. 0031-8248/2011/7801-0003\$10.00
Copyright 2011 by the Philosophy of Science Association. All rights reserved.

we attribute properties; it is no great imaginative leap to posit similar entities at microscopic scales. The discrete nature of many phenomena—the ‘atoms’ in a graphite lattice discovered by scanning tunneling microscopy, the tracks in bubble chambers, the integer multiples of electric charge on oil droplets—are easily accounted for and imagined in terms of particles and their properties. Doing so has proven scientifically fruitful. Richard Feynman declared that the greatest scientific content one could pack into a single sentence was the claim that “all things are made of atoms—little particles that move around in perpetual motion, attracting each other when they are a little distance apart, but repelling upon being squeezed into one another” (Feynman, Leighton, and Sands 1963, 1–2). As the next section demonstrates, there is good reason for the scientific realist to take Feynman’s words at face value.

For Newtonian mechanics, both macroscopic objects and microscopic particles are distinguished by continuously varying properties of mass and shape. However, to be compatible with modern physics, the particles we posit have to differ from the Newtonian conception in one seemingly minor way—they must belong to a limited number of types. How many types depends on the level of analysis: in chemistry, the relevant types are the atomic elements; in particle physics, they are the quarks, leptons, and so on, of the Standard Model. But at any level, the particles of a type are presumed to be perfectly indistinguishable—interchanging the role of two such particles in a representation of a physical state is posited to result in a (possibly distinct) representation of the identical state. Such state representations are said to satisfy ‘permutation invariance’ (PI). If these physical state representations correspond to complete descriptions of the world, PI entails that two particles of a kind are exactly indistinguishable with respect to every possible observable property.¹

I show that the empirical need for PI is disastrous for particle ontologies. A demonstration of the mutual inconsistency of PI and particles is given in section 7. To get there requires a clear statement of the class of particle ontologies (sec. 3), a clarification of the way in which theories are presumed to be interpreted (sec. 5), and a rigorous formulation of PI in theory-independent terms (sec. 6).

2. The Appeal of Particle Ontologies. A basic claim of all particle ontologies is that the world is composed of objects, and each object is largely independent of the rest with respect to the attributes it bears. This coarse assertion is supported by an inference to the best explanation. The in-

1. I have indicated that PI is not a postulate of classical physics. I should note that it is formally possible, although anachronistic, to impose PI on a classical theory (see Bach 1997; Saunders 2006).

ference begins by noting a prominent epistemic fact: with limited knowledge about some piece of the world, we can project with success the future state of that piece. Not only is science predictively successful, but its predictions require knowledge of only a tiny fraction of all facts concerning the current state of the world. For example, we need only know the angle of a cannon's barrel relative to the ground and the velocity of the ball as it exits the muzzle to figure precisely where to look for the shot when it lands. We need not take account of, say, the relative position of every star in our galaxy or the density of seawater off the coast of Greenland. I will call the fact that the future of a piece of the world can be predicted from limited facts about the current state of the world 'Epistemic Divisibility' (EDiv).²

To continue the argument, we note that in a world composed of ontologically independent objects, predicting the future attributes of one object would not require knowledge of the state of most others in the world. In other words, EDiv would be true. Take, for instance, Newton's account of material corpuscles, which he posited to be "solid, massy, hard, impenetrable, moveable Particles . . . [that] are moved by certain active Principles, such as is that of Gravity" (1966, 102–3). Whether the interactions are thought to be mediated by direct contact or a force such as gravity, each particle in this view is independent of those far enough away from it to exert a negligible force. To predict the future of some limited region of the universe in some interval Δt , it is only necessary to know about the particles close enough to interact with the particles in that region during Δt . Thus, the number of particles we need to know about is some modest fraction of all the particles there are, and EDiv follows naturally.

However, if the world is not composed of independent entities—if, for instance, the world is such that the attributes of one object are linked to those of all or most others—then the truth of EDiv would be at best extremely improbable. We can thus conclude that an ontology of independent objects is the best explanation of the epistemic facts, at least at this coarse level of explanation.³ Of course, not all ontologies that entail EDiv are particle ontologies. However, as we will see in section 4, particle

2. As an empirical generalization, EDiv does not assert that for every piece of the world under all possible conditions we can make a reliable prediction from severely limited facts. EDiv is the weaker claim that we can do so in the overwhelming majority of cases.

3. This argument is a version of the 'miracle argument' for scientific realism (see, e.g., Musgrave 2007) and suffers similar weaknesses (see, e.g., Magnus and Callender 2004). In particular, it is implausible to assert a determinate probability for the truth of EDiv, given some particular ontology, but no more so than to assert the improbability of a theory's empirical success, given its falsity.

ontologies have featured prominently in the interpretation of scientific theories and thus have special appeal for the scientific realist.

3. A Minimal Particle Ontology. Those ontologies of discrete particles that entail EDiv are the collective target of my argument. In order to proceed further, we need a precise characterization of the ontologies in this class. I claim that what follows is the logically weakest particle ontology that can support EDiv:

minimal atomism (MA)

- i) Every material object can be divided into a collection of discrete objects (called ‘particles’) each belonging to one of a finite number of kinds; particles of the same kind have in common a nonempty set of state-independent properties.⁴
- ii) Every nonempty set of particles possesses at least one state-dependent monadic property.⁵
- iii) For every set S of particles and for every $\sigma \subset S$, it is the case that for most physically possible conditions, the state-dependent monadic properties of σ are approximately independent of the properties of most subsets of the complement of σ in S over a finite interval of time.

A few points of clarification are in order. First, by “state-independent” properties I mean those that do not change in value for a particle within or across physically possible worlds. Conversely, by “state-dependent” in MA ii, I mean those properties that can assume different values at different times for a given particle (or set of particles) in a single possible world or different static values for particles of the same type in different possible worlds—in other words, those properties that are not state-independent. Second, MA i does not entail the existence of fundamental mereological simples—it does not imply that there exists some ontologically fundamental set of indivisible particles. Everything I have to say about particles would hold even if these particles were discovered to have internal structure and so too the particles composing them, ad infinitum. By “physically possible” in MA iii, I mean roughly that, for whatever physical theory

4. Particle ontologies developed as interpretations of scientific theories typically include the additional assumption that there exist only countably many particles.

5. Here and throughout I refer to the monadic properties of particles and their aggregates. It is possible (although much more complicated) to spell out a notion of ‘minimal particle ontology’ that admits only the existence of relational properties. The strict relationist is invited to read ‘monadic’ as shorthand for whatever sort of relational properties we predicate of an isolated object, perhaps with respect to an observer.

we take to be true of the world, the relevant condition obtains in the possible world described by one of the theory's models. This notion will be made clearer in section 5. The notion of property independence, also mentioned in MA iii, is of central importance. To say that the properties of particle x do not depend on the properties of particle y is to say that it is physically possible for x to assume (almost) any of the properties it might have borne had y not existed and vice versa. More will be said about this notion in section 5. Finally, the quantification over "most" physically possible conditions and subsets is admittedly vague but necessary if MA is to entail EDiv. It will turn out that we do not need to make the idea of "most" very precise since PI precludes satisfying MA under any physically possible conditions for any subset of the complement of σ when S contains only particles of the same type.

It is straightforward to show that MA is sufficient for establishing the fact of EDiv. Let S be the set of all particles in the world, and let $\sigma \subset S$ be the set of particles composing some material object. From MA i we know σ is nonempty, and from MA ii we know that σ bears at least one state-dependent property. From MA iii we know that—under most conditions in which we might find the world—the properties of σ are approximately independent of most subsets in the complement of σ in S (denoted σ^c). This means that, to predict the future properties of σ , it is almost always sufficient to know about the properties of σ and at most a minority of the subsets of σ^c . That is, it is sufficient to know about the properties of a limited subset of all the objects in the world, and so EDiv obtains.

To give a more intuitive demonstration, suppose that we are interested in the whereabouts of the cannonball some time t after a cannon is fired. Assuming that MA obtains, then we know from MA i that the universe consists of discrete particles, some of which comprise the cannonball in question. From MA iii we know that under most physically possible conditions the properties of the cannonball (e.g., momentum) will not appreciably depend on the state of most other particles or composite objects in the universe. In fact, with the exception of the volume of air through which it eventually moves and the globe of the earth that attracts both air and ball, this is the case. Thus, we need only know about the state of the ball and a particular volume of air in order to predict the ball's next location. Of course, which objects will actually be independent of the cannonball depends on which physical theory is true. MA does not tell us what the right physics is, just that a tractable physics is possible. It guarantees that for many ways of slicing up most physically possible worlds into material objects (i.e., lots of different ways of partitioning the particles that make up material bodies), there is a lack of dependence between some small portion of the universe and the rest. This ontological

independence underwrites an epistemic independence and allows us to predict the state history of that small portion without having to know about everything else. In this way, EDiv follows from MA.

It is important to note that EDiv fails to follow if we drop any of the propositions in MA—the assumptions really are minimal. For instance, without MA i, there is no guarantee that the universe can be divided at all. Without this guarantee, the remaining propositions do not imply any sort of independence between parts of the material universe since the material universe might have no parts. Without MA ii, it might be the case that the only monadic properties possessed by particles or composites of particles are state independent, such as mass. If this were the case, MA iii would be trivially satisfied, but there would be no sense in talking about the changing state of a portion of the universe. Every material portion would in itself possess only trivial state-independent properties. The relations in which that portion stands might change, but EDiv asserts that there are features of each part of the world that can be said to change in time without reference to other parts of the world. Thus, without MA ii, EDiv does not follow. Finally, MA iii ensures that, under most conditions, most portions of the world are approximately independent of one another—an ontological fact that underwrites the epistemic independence asserted by EDiv.

It is also worth stressing that MA iii is not a locality condition. As a metaphysical claim, locality is the assertion that spacelike-separated space-time regions, particles, events, and so on, are independent of one another. But there are other ways in which the world might be divided into independent units. Consider, for example, one of Dalton's early theories in which the atoms of a gas are posited to be centers of repulsion and each atom exerts a repulsive force only on atoms of the same element (Nash 1950, 19–20). In our postquantum sophistication, we might imagine a related theory in which atoms only interact with atoms of the same element but for which nonlocal interactions (or nonlocal correlations) between atoms of the same element are the rule. Since there would be no interaction between distinct elements, the world would still divide into units that are independent in the sense that MA (and by extension EDiv) requires. One would not need to know what state the oxygen is in to predict the future of the carbon. The upshot is that a locality constraint is just one way of satisfying MA iii, but it is not the only way. In this sense, MA is much weaker than a particle ontology with a locality condition.

4. Scientific Atomism and MA. MA constitutes the core of the historical series of interpretations that might be called 'scientific atomism' and remains a viable interpretation of quantum theories as long as PI is not imposed. In section 2, I mentioned Newton's corpuscular theory and

indicated how, as an instance of MA, it entails EDiv. By the end of the nineteenth century, the impenetrable corpuscles of Newton had been replaced by centers of electromagnetic force. ‘Atoms’ in this later view are composite objects made of particles with positive and negative charge, and particle interactions are mediated by electromagnetic fields. The charged particles, like the corpuscular atoms before them, still exist independently of one another in the sense that the state of a given particle is only significantly affected by spatially nearby particles. To predict the future of a portion of the universe, it is still sufficient to know about only a relative handful of fundamental particles and their attendant properties. This revised interpretation of physical theory again provides an ontological basis for EDiv and, like Newton’s view, is an instance of MA.

Even quantum mechanics (QM) can be given an interpretation compatible with MA. On the face of it, this is an unlikely proposition. After all, QM does not permit the attribution of properties to particles in the same way that classical mechanics does—under the standard interpretation, relatively few properties are definitely possessed by a composite system at any one time, and it is not clear that any are possessed by its constituent parts.⁶ Furthermore, particles have no definite position. In essentially all physical states, particles like waves are ‘spread out’ through space. Many authors have argued from such nonclassical features of quantum physics to the conclusion that QM is not a particle theory. For instance, Erwin Schrödinger argued that putative quantum particles cannot be individuated by definite space-time trajectories and so cannot be particles in the sense of individuals with transtemporal identities.⁷ Other proponents of the ‘Received View’ (French and Krause 2006, 115) of quantum particles emphasize features of quantum statistics that—when coupled with the assumption that all physically distinct states are equiprobable—suggest there is no fact of the matter as to which quantum particle bears which properties. This implies that the putative particles fail to possess identities at any one time and thus cannot be particles in the sense of individuals with definite properties. Like all other prior arguments against quantum particles, however, both of these approaches assume more than is necessary to entail EDiv. As Steven French argues (see, e.g., French 1989; French and Krause 2006), consistent particle ontologies can be found that retain particle identities and yet comport with the physics. One need not assume that particles are individuated by trajectory or that all distinct configurations of the world are equiprobable

6. According to a rather conservative but standard interpretation of QM, “a system in the state W has a value for the observable F if and only if W assigns probability 1 to one of the possible values of F ” (Dickson 2007, 285).

7. See Bitbol (2007) for an overview of Schrödinger’s arguments.

in order to account for EDiv in terms of particles. The question of interest is whether the weakest particle ontology and the benefits it brings for the realist can be maintained for quantum theories or for alternative theories that respect PI.

I contend that—as long as we leave out PI—interpretations of QM that retain MA are still available. Such an interpretation would require us to adopt a more liberal notion of property or what it means to bear a property. For instance, if one understands particles as hangers for probability distributions over property values (as opposed to determinate property values), then it is still possible to conceive of an inventory of particles as a partial description of the world in accord with MA.⁸ The ‘particles’ in this case are just those things that bear distributions over property values—they are those ontological units that are found to manifest some particular property value when a measurement is made but for which no determinate value pertains. So, for instance, I might imagine a universe empty, save for two noninteracting particles with nearly exact momenta. Each of these ‘particles’ is as nonlocalized as can be—they are distributed evenly throughout all of space. But we can view each as a particle in the sense required by MA. We can hang property distributions on each, and we can ask whether any physical circumstances are possible in which the distributions over property values for one particle—which are in principle measurable—are independent of those of the other particle. We can sensibly ask whether the properties of a portion of the world described by a quantum system are independent of the properties (or property distributions in the proposed view) of other systems. This is only a cursory gloss of the sort of interpretation of QM that would be necessary to retain the particle ontology of MA and in this way explain EDiv.

The ultimate point I wish to make is that there is nothing about the structure of quantum theories without PI that precludes a particle ontology. However, assuming strict PI in any theory renders that theory incompatible with a particle interpretation.

5. Connecting Metaphysics with Physics. In the preceding sections, I argued that MA is an appealing ontology for the realist because it accounts for the fact of EDiv. I also argued informally that MA is a viable interpretation of particle theories that lack PI. In this section, I provide a more rigorous way to connect the theories of mathematical physics with metaphysical interpretations. The idea is to construct an exact means of deciding whether a given theory is compatible with a given interpretation.

8. This is a strong claim akin to asserting the existence of vague properties. The point is not to defend such an interpretation but to show that a plausible or at least consistent interpretation exists.

If our metaphysical accounts of the world are to have content, it must be the case that physical theories constrain the metaphysical possibilities; some ontologies must be false if certain physical theories are empirically adequate. There is a way to make this relationship of constraint precise.

Let me begin by saying a little about how I am conceptually carving up theories and their interpretations. By a ‘theory’ I mean a specification of a mathematical space along with one or more laws. The latter act as constraints that, when supplied with supplementary conditions that correspond to the particular physical system to be represented, pick out a particular subset of the mathematical space. So, for example, Hamilton’s equations of motion, given a specification of particle number n and initial positions and momenta, pick out a single trajectory in $6n$ -dimensional phase space (modulo some technical conditions).

A ‘model’ of a theory is a mathematical structure compatible with the laws and some particular set of boundary conditions. For instance, Newton’s laws of motion when coupled with suitable boundary conditions pick out a trajectory in a continuous phase space. The postulates of QM, given a Hamiltonian and initial state, pick out a ray within a Hilbert space in the time-independent case and a trajectory through such rays in the time-dependent case. What I am calling ‘models’ of a theory are what physicists often call ‘solutions’.

An ‘interpretation’ of a theory is a set of property attributions over a set of objects. The interpretation of a theory is supposed to be a description of the way the world is. Some of the elements of the interpretation refer to objects in the world, while others correspond to the properties of these objects. I do not mean to insist that interpretations be expressible in first- or second-order formal languages. Rather, I intend only that an interpretation be a structure for which distinct elements directly correspond to distinct entities and property instances in the world. I will call a part of an interpretation that refers to just one object and its associated properties a ‘specification’.

With this taxonomy in hand, we can talk about how one might go about binding various metaphysical theses to physical theories. The idea is to treat interpretations of a theory as descriptions of physically possible worlds. If we fix the manner in which interpretations are extracted from models of a theory, then the class of models represents the class of possible worlds. MA iii is now a constraint on physically possible worlds and thus on models of a theory. We now have a way of specifying mathematical features that any physical theory must have if it is to be compatible with MA iii. Loosely speaking, it must be physically possible under most configurations of the world for one particle to assume any of the physical states open to a lonely particle, irrespective of what most of the others are doing. That is, under most possible conditions, our physicist can

change the boundary conditions and produce another apparently lonely one-particle state, irrespective of what state the other particles in the universe happen to occupy.

In the case of just two particles, MA iii can be put as follows: the theory must support a pair of sets of models that correspond to nontrivial sets of one-particle specifications such that each specification in one set can be combined with any specification from the other set to get what are approximately the specifications corresponding to a two-particle model of the theory. For example, it is easy to produce a set of models in classical mechanics, each of which represents a molecule of gas moving with a different velocity within a small volume of space to my left, and another set of models that represent an exactly similar molecule flitting about to my right. Choose one model from the first set and one from the second. The union of the two corresponding interpretations contains specifications of molecules that are approximately the same as we would have gotten by interpreting a model of the theory that takes both molecules into account. This is because the molecules interact only weakly at such a distance (using realistic potentials), and so each has nearly (but not quite) the same properties it would have had in the absence of the other. In this case, it looks like the one-particle models account for all of the possible experiments I could do on either particle, although in fact there are two particles. That is precisely what EDiv demands.

There remain two vague notions in the conditions laid out above: that of a ‘nontrivial’ set of specifications and that of ‘approximate’ specification equivalence. What we need in order to clarify these notions is a measure of the difference or distance between any two specifications—we need to impose a metric on the space of possible one-particle specifications. It does not matter what metric, so long as it is a metric (i.e., it is positive-definite, symmetric, and obeys the triangle inequality). Given such a measure of distance, we can then choose a threshold separation, ε , a minimum distance at or below which we call two specifications approximately the same. Clearly there is arbitrariness in picking such a threshold. There are some plausible guides we could use, such as the practical distinguishability of states. But we need not settle on a particular value since it ultimately will not matter what value we choose.

The threshold ε and the metric to which it refers give us a way to clarify the relation of approximate equivalence: two specifications are approximately the same just if they are within a distance ε of one another. They also give us a way to cash out the notion of a nontrivial set of specifications: to say that a set of one-particle specifications is nontrivial is just to say that its diameter (the greatest distance between two specifications in the set) is large compared to ε . According to what we have taken to constitute approximate similarity, any set that is not large compared to

ε would contain ‘approximately’ one specification. One can think of ε as setting a natural scale with respect to which the rest of the properties of the space of possible specifications are spelled out. While the selection of ε is arbitrary, the remaining properties are not.

In summary, for a theory T to be compatible with MA iii, there must exist three sets of models of T —call them α , β , and γ —with the following properties:

1. The models in α are interpreted as representing the properties of a single particle. Let S_α be the set of all one-particle specifications extracted from the interpretations corresponding to the models in α . The set S_α is large with respect to the threshold ε .
2. The models in β are likewise interpreted as representing a single particle. Let S_β be the set of all specifications in the interpretations of the models of β . The set S_β is also large compared to ε .
3. The models in γ are interpreted as representing the properties of two particles. For every ordered pair of specifications in $S_\alpha \times S_\beta$, there exists an interpretation of a model in γ that contains approximately both of these specifications (that approximately attributes each of these specifications to one particle in the interpretation).⁹

In section 7, I show that insisting on symmetrization renders these conditions unsatisfiable. In this way, PI sharply precludes particle interpretations.

6. Permutation Invariance and Perfect Indistinguishability. It is my central claim that PI and MA are mutually inconsistent for any physical theory. To show this requires precisely formulating PI in a manner general enough to apply beyond the quantum mechanical framework in which the postulate was introduced. To construct such a statement, I will begin with the principle as it is given in QM and then excise the details pertaining specifically to that theory.

In the physics literature, a typical statement of PI runs as follows:

(\mathbf{PI}_{QM1}): Dynamical states represented by vectors which differ only by a permutation of [particles of the same type] cannot be distinguished by any observation at any instant of time. (Messiah and Greenberg 1964, 250)

The “vectors” in this statement refer to vectors in a Hilbert space. \mathbf{PI}_{QM1}

9. The sets S_α and S_β are subsets of a common set of all possible single-particle specifications. From the interpretation of each model in γ , it is possible to extract an ordered pair of specifications, the elements of which are members of the same common set of single-particle specifications.

is effectively a restriction on what counts as an observable. Which sets of vectors represent states is fixed by the following two assumptions: “(1) Two vectors representing the same state must give the same expectation value with respect to all observables; and (2) Two vectors representing distinct states must give different expectation values for at least one observable” (Hartle and Taylor 1969, 2045). From these assumptions, it follows that vectors represent the same physical state if and only if they are indistinguishable with respect to all observables. Coupled with PI_{QM1} , these assumptions entail that any two state representations differing only by a permutation are in fact representations of the same physical state. So we can rewrite PI_{QM1} in my terminology as follows:¹⁰

(PI_{QM2}): Physical states are invariant under permutations of those parts of the state representation that correspond to particles of the same type.

But PI_{QM2} no longer contains explicit reference to the particular formalism of QM. Noting that ‘state representations’ are what I have been calling ‘models’ of QM, and assuming there exists a function taking models to interpretations, we can now state the principle of PI in a precise, theory-independent fashion:

(**PI**): If the interpretation *Int* of a model *M* of a theory describes multiple particles of the same type, then permuting the parts of *M* that individually correspond to descriptions of those particles in *Int* results in a model *M'* with the identical interpretation *Int*.

10. I have glossed over a great deal of technical detail. Suppose that H^1 is the Hilbert space corresponding to the states of one particle. For N distinguishable particles (for which PI does not apply), the states of the system are given by the vectors of the outer product Hilbert space: $H^N = H^1 \otimes H^1 \otimes \dots \otimes H^1$ (N times). PI along with conditions 1 and 2 imply that ‘states’ of indistinguishable particles must be represented by the subspaces of H^N corresponding to the irreducible (unitary) representations of the permutation group. Operating on one of the vectors in such a subspace yields a (possibly distinct) vector in the same subspace. It can be shown that H^N decomposes into a direct sum of subspaces, each corresponding to an irreducible representation of the permutation group. The one-dimensional subspaces correspond to either completely symmetric or completely antisymmetric states (vectors that are either invariant or change sign under the action of a permutation operator). The former describe the states of so-called bosons (e.g., photons), while the latter describe the ‘fermions’ (e.g., electrons). In principle, the remaining higher-dimensional subspaces in the decomposition of H^N could describe the states of one or another type of particle. The states of such particles—called ‘paraparticles’—would exhibit more complex symmetries and thus, for instance, more complicated statistics in large ensembles. But only bosons and fermions are believed to exist. For an overview of the relevant formalism, see Messiah and Greenberg (1964) and Hartle and Taylor (1969).

Note that in all of the above formulations PI refers to the interpretation of a model in that it refers to permuting those parts of a model that correspond to particles.¹¹ There is thus an implicit but essential assumption contained within PI about the manner in which models are mapped to interpretations. In the most general process of interpretation, there is a function from the class of models of a theory to the class of interpretations—to every model there corresponds exactly one interpretation.¹² PI further assumes that the function taking models to interpretations is a composite of functions mapping pieces of each model to pieces of an interpretation. Specifically, each model is supposed to contain one or more mathematical parts that are taken to correspond directly to the descriptions of individual particles within the interpretation of the model, that is, each part corresponds to a specification. For instance, if a model of the hydrogen atom in QM contains a ket (a vector) from the Hilbert space used to represent a single electron and a ket from the space used to represent a proton, then these should correspond in the interpretation to the specifications of an electron and a proton. Importantly, it must be the case that the same electron ket corresponds to the same electron specification across all models and interpretations. So if the mathematical component corresponding to an electron produces a particular specification in the interpretation of one model, it should yield the same specification in the interpretation of any other model that contains that component.

An example from classical physics might help to clarify this idea. Suppose we are interested in the physics of pendulums. In the simplest case, the motion of a single pendulum can be represented in terms of two variables: an angular position and an angular momentum. A model of Hamilton's equations of motion for a single pendulum is a trajectory through the two-dimensional phase space built from these two variables. Interpretation of such a model is straightforward. At any given time, we consider a point on the trajectory and read one coordinate as yielding the position of the bob and the other as giving its momentum. The mathematical piece that corresponds to our pendulum is just the entire model—the trajectory in two-dimensional phase space. In classical physics without PI, it is easy to build models of composite systems. If we want to model two pendulums, whether coupled or not, we need only move to a four-

11. The 'parts' of a model include mathematical structures produced from the model as an intermediate in interpretation. For instance, in QM the standard approach to assigning properties to individual particles is to construct a reduced-density operator for each. These reduced-density operators would constitute 'parts' of the original model.

12. Multiple models may share the same interpretation.

dimensional phase space built from four coordinates: the position and momenta of the two pendulum bobs. Models of the theory are now trajectories through this phase space. In our interpretation, however, we have specifications for not one but two pendulums. The mathematical pieces that correspond to these specifications are just the projections of the four-dimensional trajectory onto each of the two-dimensional momentum-angle spaces. As long as that projection is the same as the trajectory we had originally (it will not be if the pendulums are coupled), then the specification for that pendulum—the description of its properties—should remain the same.

However one is inclined to interpret the postulate, PI is an empirically essential component of QM.¹³ Quite generally, states that are permutation invariant give rise to so-called interference terms when the Born rule is applied, terms that do not necessarily obtain for nonsymmetric states.¹⁴ These interference terms produce widely different predictions for such phenomena as the cross-sections of particle collisions, the structure of atoms, and the statistical mechanics of systems of large numbers of particles. By way of illustration, consider the simple example of two non-interacting particles in a one-dimensional box. If the particles are of different kinds (a neutron and a proton, say), then the probability density of finding both at the same point x in the box is given in terms of wave functions by $|\psi_{\text{proton}}(x)|^2 |\psi_{\text{neutron}}(x)|^2$. If the particles are of the same type such that PI applies, this probability density is either 0 (for particles that satisfy PI with antisymmetric states) or twice as great (for those that satisfy PI with symmetric states). PI imposes strong, empirically verified restrictions on measurement frequencies. Conversely, no states are observed for particles of the same type that do not conform to PI.

7. Proving an Inconsistency. A model is a mathematical structure. Parts of this structure are presumed to correspond to individual particles. If we call the part of a structure corresponding to one particle the ‘role’ of that particle, then corresponding to any model of a theory there is an indexed set of particle roles. When the model is interpreted, each role is mapped to a specification of the corresponding particle. So to any model there also corresponds an indexed set of particle specifications. If two models share the same interpretation, then the indexed set of particle specifications corresponding to each model must be identical (although the set of roles may not be).

13. It is unclear how the postulate could be eliminated from quantum field theory where it manifests in the commutation relations among creation and annihilation operators.

14. For an introductory overview, see Shankar (1994), 269–78.

As an example, consider two classical particles moving in one dimension such that in the chosen coordinate system they are each confined to the interval $[0, 1]$. The motion of the particles can be described in the two-dimensional configuration space of points (x_1, x_2) with $0 \leq x_1 \leq 1$ and $0 \leq x_2 \leq 1$. Let's suppose that models of the relevant mechanics of this system are just points in the configuration space, for example, $(0.5, 0.1)$. For each particle, there is a corresponding 'role' in this model: for the first particle, it is the real number 0.5, and for the second it is 0.1. Furthermore, the interpretation, *Int*, of this model contains two specifications, s_1 and s_2 . Specification s_1 asserts that the first particle is at a position 0.5, and s_2 asserts the second particle is at 0.1. If another model of the theory has an interpretation *Int'* containing specifications (s_3, s_4) that is identical to *Int*, then the specification of particle one in the first interpretation must be identical to the specification of particle one in the second (so $s_1 = s_3$) and likewise for the second particle (so $s_2 = s_4$).

To be a little more precise, let A be the set of particle roles for a given model M . Let n be the total number of particles represented by M , and let N be the set of natural numbers $\{1, 2, \dots, n\}$. Let $f: N \rightarrow A$ be a surjective function mapping the index of the j th particle in the system to its role in M . A model can thus be seen as giving an ordered n -tuple of particle roles, where the j th role is given by $f(j)$. Corresponding to each model, there is also an indexed set of particle specifications that can be obtained by interpreting the roles of the model. If we let S be the set of all one-particle specifications and *interp* : $A \rightarrow S$ be the interpretation function taking roles to specifications, then the j th specification in the ordered n -tuple of particle specifications corresponding to M is just *interp*($f(j)$). If M_1 and M_2 are models representing the same number of particles and f_1 and f_2 are the respective indexing functions labeling particle roles for each, then these models have the same interpretation only if for all j , *interp*($f_1(j)$) = *interp*($f_2(j)$).

What consequence does PI have? Let P_i be a permutation of N (a bijection from N to itself). Then a 'permutation of particle roles' for a model is just the ordered n -tuple for which the j th role is given by $f(P_i(j))$. Suppose that we are only concerned with models for which all particles are of the same type (e.g., all electrons). Then PI requires that any permutation of particle roles for such a model result in a model with the identical interpretation. In the formalism given here, that means that for all permutations P_i and for all j in N , *interp*($f(j)$) = *interp*($f(P_i(j))$). But this holds if and only if *interp*($f(j)$) = constant. That is, this holds only if all particle specifications are the same. Since the model is interpreted piecewise with each role corresponding to a distinct specification of a particle, PI forces every particle in the interpretation of a model to bear the same description.

Attributing to every particle of a type in a system the same set of properties is a problem if we are to satisfy MA iii. For the latter, we need two sets of one-particle specifications—call them S_1 and S_2 —each of which is large compared with a given threshold ε . We must be able to select an arbitrary specification from each of these sets—call these s_1 and s_2 —and be able to locate a pair of specifications in the interpretation of a two-particle model, one of which is within ε of s_1 and the other within ε of s_2 . Now, suppose that we choose s_1 and s_2 such that the distance between them—denoted $d(s_1, s_2)$ —is greater than 2ε . Such a choice must be open to us because both sets are large compared to ε . From PI we know that every interpretation of a two-particle model contains a pair of identical specifications, and we will call this shared specification s^* . By the triangle inequality for metrics, it must be the case that

$$d(s_1, s_2) \leq d(s^*, s_1) + d(s^*, s_2).$$

By supposition, $d(s_1, s_2) > 2\varepsilon$, so

$$2\varepsilon < d(s^*, s_1) + d(s^*, s_2).$$

But this means that there is no way to choose a single state s^* such that $d(s^*, s_1) \leq \varepsilon$ and $d(s^*, s_2) \leq \varepsilon$. Because PI forces every particle to receive the same specification, it is impossible to satisfy MA iii—symmetrization prevents particles from possessing the sort of independence of properties that could explain EDiv.

I have stated it in the abstract, but the argument can be formulated within specific theories. Take QM, for example. To see how the argument works there, we first need to be clear on what constitutes the theory and how models of the theory are to be interpreted. To be as inclusive as possible, I will consider the set of quantum state representations to be the full set of density operators on a given Hilbert space. Doing so admits as viable representations both the usual ‘pure’ states equivalent to rays in the Hilbert space and the ‘mixed’ states equivalent to convex combinations of pure states. The theory then consists of the von Neumann Equation ($i\hbar \partial\rho/\partial t = [\rho, H]$)—which specifies a trajectory through the set of density operators—along with the usual postulates for extracting expectation values and measurement probabilities. For simplicity, I will ignore time dependence and focus just on the time-independent solutions of the von Neumann Equation for suitably specified boundary conditions. In that case, models are just single-density operators.

How should these models be interpreted? What objects are specified by the models, and how should we assign properties to them? In the case of a single-particle system, we can interpret each density operator as specifying the properties possessed by a particle in the weak sense I suggested in section 4. In this view, particles do not necessarily possess definite values of properties but rather probability distributions over possible values for

each property (or equivalently, a definite expectation value with respect to every observable). Such a view relaxes the eigenvector-eigenvalue link and allows us to interpret mixed states as representing the properties of single particles.

With respect to the space of density operators on a Hilbert space, there are many metrics in the literature to which we might appeal. To make the argument easy to visualize geometrically, I will take the distance $d(\rho_i, \rho_j)$ between any two density operators ρ_i and ρ_j to be given by twice the standard ‘trace distance’ (see, e.g., Nielsen and Chuang 2000, sec. 9.2.1):

$$d(\rho_i, \rho_j) = \text{Tr}|\rho_i - \rho_j|.$$

The expression $\text{Tr}|A|$ stands for the trace norm on A (see Omnès 1994, 245). Each density operator specifies a unique set of distributions with respect to property values, so there is a one-to-one correspondence between one-particle density operators on the Hilbert space and specifications in the interpretations of one-particle models. This means we can transfer the trace-distance metric onto the space of specifications and define the distance between any two specifications as the distance between the corresponding density operators.

If we consider the case of two-dimensional particle states, using the trace distance to measure the difference between specifications makes the inconsistency argument easy to visualize. For instance, we might consider just the spin degree of freedom of a spin-half particle. In that case, we can represent all possible density operators with a vector \mathbf{r} in three-dimensional Euclidean space. Specifically, each density operator ρ can be uniquely represented in the form $(I + \mathbf{r} \cdot \boldsymbol{\sigma})/2$, where I is the identity operator, \mathbf{r} is a vector unique to ρ with $\|\mathbf{r}\| \leq 1$, and $\boldsymbol{\sigma}$ is a fixed vector (the vector of Pauli matrices). There is thus a one-to-one correspondence between the density operators and the points of a sphere called the ‘Bloch sphere’. Points on the surface of the Bloch sphere for which $\|\mathbf{r}\| = 1$ correspond to pure states, while the points in the interior of the sphere correspond to mixed states. In this representation, the distance $d(\rho_i, \rho_j)$ between two density operators ρ_i and ρ_j corresponds to the Euclidean distance between the two points in the sphere corresponding to those operators. That is, $d(\rho_i, \rho_j) = \|\mathbf{r}_i - \mathbf{r}_j\|$.

To visualize the argument in the Bloch sphere, pick a threshold distance ε small compared to its diameter. Now choose a pair of vectors \mathbf{r}_1 and \mathbf{r}_2 corresponding to operators in the space of one-particle models such that $\|\mathbf{r}_1 - \mathbf{r}_2\| > 2\varepsilon$. According to MA iii, we must be able to find a two-particle model that—when interpreted—yields two copies of a single-particle specification corresponding to a vector \mathbf{r}^* that is within ε of both \mathbf{r}_1 and \mathbf{r}_2 . This means that \mathbf{r}_1 and \mathbf{r}_2 must lie within a ball of radius ε centered on \mathbf{r}^* (see fig. 1). Obviously, if this is the case it cannot also be true that \mathbf{r}_1




Figure 1.

and r_2 are more than 2ϵ apart since the diameter of the ball is 2ϵ , and we arrive at a contradiction. Given PI and the trace-distance metric on the space of quantum models, MA iii is not satisfiable.

8. Objections. There are a number of claims in my argument with which one might take issue. I will start with the most specific and work to the most general. To begin with, one might complain that QM itself ought to belie any concern that symmetrization must always ruin the independence of systems. J. J. Sakurai puts the objection this way: “There is no need to antisymmetrize [the state vector] if the electrons are far apart and the overlap [between wave functions] is negligible. This is quite gratifying. We never have to worry about the question of antisymmetrization with 10 billion electrons, nor is it necessary to take into account the antisymmetrization requirement between an electron in Los Angeles and an electron in Beijing” (1994, 365–66). Sakurai is, of course, exactly right so long as we restrict our attention to the outcome of possible measurements, none of which can make reference to which electron is which. Empirically, the single-particle measurements one can gather are predicted just as well by a one-particle model of the theory as by a 10 billion-particle model under the conditions Sakurai describes. However, if one wants to attribute properties to individual particles on the basis of these possible measure-

ments, if one wants to interpret the quantum state in terms of a collection of particles satisfying MA, then it does matter whether there is an electron in Beijing as well as Los Angeles (LA). Suppose $\psi_B(x)$ is the wave function of an isolated electron in Beijing, and $\psi_{LA}(x)$ is the wave function of one in LA. Ignoring the spin degree of freedom for simplicity, the only two-electron state admitted by QM with PI is $\psi_B(x_1) \psi_{LA}(x_2) - \psi_B(x_2) \psi_{LA}(x_1)$. In the standard approach to attributing properties to particles individually, we compute the marginal probability distribution for the Beijing electron by integrating out the states of the other particle and vice versa. The result is, of course, a pair of identical distributions over space—both electrons are attributed the same probability of being found in any given volume of space. Importantly, both are attributed a probability of 0.5 to be found in either Beijing or LA, a property that would not have been attributed to, say, the LA electron if the Beijing electron did not exist. Because the properties that are attributable to particles depend on whether there are additional particles of the same kind in the world, one does not have the option to choose not to symmetrize on occasion—there would be an enormous difference in the properties attributed. If we want to be realists and think that the theory of QM does more than rescue the phenomena, then one cannot ignore the fact that the content of one's interpretation depends on the number of particles of the same kind supposed to exist.

Second, one might simply dismiss my result as old news. After all, there are plenty of folks who have provided good arguments to believe that neither QM nor quantum field theory can be interpreted in terms of individual particles. However, each of their arguments is grounded in a commitment to one of various mutually exclusive accounts of 'particle'. These are diverse arguments subject to diverse objections. Since they each posit incompatible metaphysical principles, the convergence of these arguments ought not to increase our confidence in the conclusion. The argument I have provided takes on minimal metaphysical baggage—just what is required to account for an undeniable epistemic fact. If you reject the metaphysics that has been brought on board, then a great many realist interpretations must go with it. Furthermore, my argument is not specific to QM; it applies to any theory in which PI can be stated and is presumed true.

Finally, one might object that MA is not as weak as it sounds or that it fails to capture one or more necessary conditions of the 'correct' notion of particle. To this line of argument, I say that I am more than willing to adjust my terminology. I am not interested in developing a conceptual analysis of the terms 'particle' or 'atomism', nor do I intend to exhaustively catalog the many ways in which these terms have been employed. Rather, I am interested in those ontologies—call them what you will—which claim that the material world is discretely divisible into entities that

bear properties in a manner that can sustain EDiv. I have given reasons to think that what I have called MA does in fact line up with various scientific conceptions of atomism. But even if this fails to be the case, the argument I have presented rules out a very large and very appealing class of ontologies.

9. Consequences. The general incompatibility between PI and MA means we have to give up one of the three propositions with which I began this article. PI itself appears unassailable, as without it the particle theories we have (e.g., QM) are empirically inadequate. That leaves MA, the view of the universe as a discretely divisible entity made up at one or more levels of particles belonging to a handful of types. It is no trivial affair to jettison this view, no matter how many have impugned it before. A great deal of scientific explanation in such fields as physical chemistry could no longer be taken literally, and it is difficult to see what should replace it.

Furthermore, giving up MA means giving up the most successful metaphysical account of EDiv at hand. But herein lies the positive value of the incompatibility result presented here. It is clear that we need an alternative to MA. For scientific realism to remain plausible, any successful alternative must comport with the epistemic facts. Any metaphysical account that fails to underwrite EDiv or that directly contradicts it is empty at best and patently false at worst. We thus have a guide for producing viable alternatives to a particle ontology. The imposition of PI leads the proponent of MA to contradiction partly because discrete particles were identified as the property bearers. One might instead ascribe properties to a different set of objects that can nonetheless be identified with components of the models of our best theories of matter and that satisfy conditions similar to MA. I do not have the space here to pursue these alternatives, but I can suggest what one possibility looks like in outline—it is a sort of space-time substantivalism. The relevant property-bearing substances, rather than particles, are taken to be regions of space-time. The properties once attributed to particles in quantum states would be attributed to one or more regions of space-time. PI would no longer obtain, at least not as a statement about the permutation symmetry of property bearers. Rather, it would have to be seen as a restriction on the possible properties attributable to space-time regions. As long as it could be shown that there are conditions under which the properties attributable to one region are independent of the others, EDiv will obtain. And Sakurai (among many others) already pointed out that such independence can obtain. As long as we treat the outcomes of putative one-particle measurements as properties of the space-time regions occupied by our instruments, then the electron in Beijing really does not matter for the

electron in LA. I do not know if this is the way to go; space-time substantivalism may be problematic for other reasons related to theoretical symmetry. Nonetheless, the result reported here provides a clear guide for developing such alternatives.

REFERENCES

- Bach, Alexander. 1997. *Indistinguishable Classical Particles*. Berlin: Springer.
- Bitbol, Michael. 2007. "Schrödinger against Particles and Quantum Jumps." In *Quantum Mechanics at the Crossroads: New Perspectives from History, Philosophy and Physics*, ed. James Evans and A. S. Thorndike, 81–106. Berlin: Springer.
- Dickson, Michael. 2007. "Non-relativistic Quantum Mechanics." In *Philosophy of Physics*, ed. Jeremy Butterfield and John Earman, 275–415. New York: Elsevier.
- Feynman, Richard P., Robert B. Leighton, and Matthew Sands. 1963. *The Feynman Lectures on Physics*. Vol. 1. Reading, MA: Addison-Wesley.
- French, Steven. 1989. "Identity and Individuality in Classical and Quantum Physics." *Australasian Journal of Philosophy* 67 (4): 432–46.
- French, Steven, and Décio Krause. 2006. *Identity in Physics*. Oxford: Clarendon.
- Hartle, James B., and John R. Taylor. 1969. "Quantum Mechanics of Paraparticles." *Physical Review* 178 (5): 2043–51.
- Magnus, P. D., and Craig Callender. 2004. "Realist Ennui and the Base Rate Fallacy." *Philosophy of Science* 71:320–38.
- Messiah, A. M. L., and O. W. Greenberg. 1964. "Symmetrization Postulate and Its Experimental Foundation." *Physical Review* 136 (1): 248–67.
- Musgrave, Alan. 2007. "The 'Miracle Argument' for Scientific Realism." *Rutherford Journal: The New Zealand Journal for the History and Philosophy of Science and Technology*, no. 2.
- Nash, Leonard K. 1950. "The Atomic-Molecular Theory." In *Harvard Case Histories in Experimental Science*, ed. James Bryan Conant. Cambridge, MA: Harvard University Press.
- Newton, Isaac. 1966. "Optiks." In *The World of the Atom*, ed. Henry A. Boorse and Lloyd Motz, 102–4. New York: Basic.
- Nielsen, Michael A., and Isaac L. Chuang. 2000. *Quantum Computation and Quantum Information*. Cambridge: Cambridge University Press.
- Omnès, Roland. 1994. *The Interpretation of Quantum Mechanics*. Princeton, NJ: Princeton University Press.
- Sakurai, J. J. 1994. *Modern Quantum Mechanics*, ed. San Fu Tuan. Rev. ed. Reading, MA: Addison-Wesley.
- Saunders, Simon. 2006. "On the Explanation for Quantum Statistics." *Studies in History and Philosophy of Modern Physics* 37:192–211.
- Shankar, R. 1994. *Principles of Quantum Mechanics*. 2nd ed. New York: Plenum.