Identity, Superselection Theory and the Statistical Properties of Quantum Fields

David John Baker
Department of Philosophy, University of Michigan
djbaker@umich.edu

October 9, 2012

Abstract

The permutation symmetry of quantum mechanics is widely thought to imply a sort of metaphysical underdetermination about the identity of particles. Despite claims to the contrary, this implication does not hold in the more fundamental quantum field theory, where an ontology of particles is not generally available. Although permutations are often defined as acting on particles, a more general account of permutation symmetry can be formulated using superselection theory. As a result, permutation symmetry applies even in field theories with no particle interpretation. The quantum mechanical account of permutations acting on particles is recovered as a special case.

1 Introduction

The permutation symmetry of quantum theory has inspired a long debate about the metaphysics of identity and such related principles as haecceitism and the identity of indiscernibles. Since permutations are normally defined as acting on particles, this debate has largely proceeded by interpreting quantum physics as a theory of particles. As a result, there is presently no clear connection between this work and the latest foundational studies of quantum field theory (QFT), many of which have argued against particle ontologies.

It is not quite right to say that the debate concerning the metaphysics of statistics has ignored QFT. Rather, parties to this debate have engaged with QFT in a set of idealized special cases: Fock space QFTs on Minkowski spacetime, which can be treated as particle theories (albeit with the number of particles sometimes indeterminate). More realistic interacting theories and QFTs on curved spacetime–which resist a particle interpretation–have been ignored.

This is understandable, and hardly surprising. The metaphysical debate is supposed to concern the implications of permutation symmetry, and of the Bose-Einstein and Fermi-Dirac statistics obeyed by quantum systems. But permutation symmetry is normally understood in terms of the effect on the quantum state of permuting some particles. This extends to the usual mathematical definition, where one defines a representation of the permutation group acting on the labels for particles appearing in the states. Statistics is then understood in terms of which particle states are physically allowed; only states that transform a certain way under the action of permutations are deemed physically possible. How could any of this work, in the absence of a particle ontology? For QFTs with no particle interpretation, where does the philosopher interested in permutation symmetry even begin?

We need a generalization of permutation symmetry that remains meaningful even in QFTs with no particle interpretation. Luckily such a generalization already exists, as a little-known consequence of the mathematical theory of superselection sectors in the algebraic approach to QFT (AQFT). By applying this more general definition, we will see that in our most fundamental quantum theories, permutation symmetry has no implications whatsoever for the metaphysics of particles–because in the most fundamental permutation symmetric QFTs, there are no particles. The existing debate concerning identical particles in quantum mechanics and special relativistic free QFT disappears when more fundamental interacting and curved-spacetime theories are considered.

As with any philosophical work involving AQFT, any dependency of my argument on the legitimacy or eventual success of the controversial algebraic approach is potentially concerning. It may help to be upfront about the point of dependency. I have two theses here, a negative one and a positive one. My negative thesis is that the nonexistence of a particle interpretation in curved-spacetime and interacting QFTs dissolves the existing debate concerning the metaphysics of identity in quantum physics. As we'll see in §3, this thesis does not depend on the algebraic approach. My positive thesis is that even without an ontology of particles, we can make sense of permutation symmetry in interacting AQFT. Obviously this does depend on the algebraic approach; I will motivate this choice of framework when it becomes relevant, in §5.

2 Particle permutations and Fock space

I'll begin by laying out the standard definition of permutations as acting on particles. According to the orthodox statistical interpretation of non-relativistic quantum mechanics (QM), the basic objects are particles whose probabilistic propensities are described by state vectors in a Hilbert space of wavefunctions. These probabilistic properties of a particle in state ψ are given by the expectation values $\langle\psi|A|\psi\rangle$ of observables represented by self-adjoint operators A.[1]

Multiple particles are represented by a tensor product of single-particle Hilbert spaces. So to represent two identical particles, one would label two copies of the single-particle Hilbert space $\mathcal{H}$ and treat their tensor product $\mathcal{H}_1 \otimes \mathcal{H}_2$ as the two-particle state space. At this point something interesting becomes apparent. If ψ and ϕ are two distinct single-particle states, one might naively expect the vector $|\psi\rangle_1 \otimes |\phi\rangle_2$ to represent a possible state in the two-particle space–intuitively, one in which particle 1 is in state ψ and particle 2 is in state ϕ.

But here, as in so many places, quantum mechanics confounds our naive expectations. The only available two-particle states that involve both a particle in state ψ and a particle in ϕ are

$$\frac{1}{\sqrt{2}}(|\psi\rangle_1 \otimes |\phi\rangle_2 + |\phi\rangle_1 \otimes |\psi\rangle_2) \tag{1}$$

and

$$\frac{1}{\sqrt{2}}(|\psi\rangle_1 \otimes |\phi\rangle_2 - |\phi\rangle_1 \otimes |\psi\rangle_2), \tag{2}$$

with state (1) characterizing boson systems and state (2) fermions. What's most interesting for present purposes is the way these states transform when the particle labels are permuted. State (1) is symmetric under such permutations, while (2) is antisymmetric. Equivalently, each state has as a symmetry some representation of the permutation group on two objects,

[1] With reservations, I will follow the convention (in the philosophical literature on permutation invariance) of assuming the orthodox statistical interpretation for purposes of argument. But it deserves to be noted that this assumption is hardly innocent, given our concern with the metaphysics of particles. The GRW interpretation of QM, for example, is not fundamentally a particle theory in either its "mass density" or "flash" formulation. Nonetheless, the assumption is innocent for present purposes, since my ultimate conclusion is negative: even assuming the orthodox statistical interpretation, permutation symmetry does not have the metaphysical implications that have previously been claimed for it. Moreover, most of the existing debate (along with what I have to say) extends straightforwardly to the popular Everett interpretation as well as the orthodox one.

S_2: the trivial representation for bosons and the alternating representation for fermions. For three or more particles, multi-dimensional representations become available, allowing particles with so-called parastatistics which are (perhaps mysteriously) not observed in nature.

This behavior under permutations–and the resultant Bose-Einstein or Fermi-Dirac statistics–has been thought to entail some novel metaphysics. For example, if we assume that the probabilistic predictions of a state (its expectation values) exhaust the properties it ascribes to particles, it can be shown that bosons, fermions and parastatistical particles ("paraparticles") all violate the principle of the identity of indiscernibles (French and Krause (2006, 150-173)). It has also been suggested that permutation symmetry indicates quantum particles are not "individuals" meeting the usual modal or logical criteria (French and Krause, 2006, 140-148).[2] The fact that neither the individuality nor the non-individuality of quantum particles is logically entailed by the theory has been taken to imply a sort of theoretical underdetermination: the scientific facts provide no basis for accepting either metaphysical stance (van Fraassen, 1991, 434-482). In an attempt to dissolve this underdetermination (among other motivations), structural realists have tried to formulate a metaphysics without a fundamental ontology of objects (Ladyman and Ross, 2007).[3]

These questions about the metaphysics of QM are of great intrinsic interest. Since QM has enjoyed great success, and was for a time the most fundamental theory of matter available, I see considerable value in determining what a purely non-relativistic quantum universe would be like. But the correct metaphysics for our own universe is an even more pressing issue–and in the domains explored by experiment thus far, our universe appears to be relativistic. So it would be a mistake to ask whether (for example) there are individuals in our own world, and turn to non-relativistic QM for an answer. Of course, no final answer is available at present, since the project of physics remains far from complete. But we may hope that a more fundamental theory like QFT can provide a more reliable answer to our metaphysical question than QM.

In this spirit, philosophers concerned with the metaphysics of individuals have turned to QFT on Fock space. For example, Teller (1995, 16-52) has argued that particles in Fock space, while countable, are manifestly not individuals. If correct, this would resolve the issue

[2]For example, it is sometimes suggested that "individuals" must be potential bearers of names (French and Krause, 2006, 198-237).

[3]In my opinion, a more promising (and precisely formulated) approach along similar lines is that of Dasgupta (2009).

of metaphysical underdetermination. Teller reasons as follows: Written in the commonly-used occupation number basis, Fock space entirely avoids the use of particle labels. For example, if the vector states $|\psi_i\rangle$ form a basis of the one-particle space, a basis for Fock space is given by the set of states $|n_1, n_2, n_3, ...\rangle$, where n_1 is the number of particles in state ψ_1, and so on. So for example, a two-particle state might be written $|1, 1, 0, 0, ...\rangle$. Trivially, this way of representing states draws no distinction between a two-particle state and the same state after a permutation of the two particles. In particular, there is no way of representing a non-symmetric state like $|\psi\rangle_1 \otimes |\phi\rangle_2$ in the occupation number basis. Thus Teller concludes that "Fock space realizes the idea of quanta, understood as entities that can be (merely) aggregated, as opposed to particles, which can be labeled, counted, and thought of as switched." On this picture, Fock space particles or "quanta" are not individuals in the same sense as classical particles.

Teller's argument fails, however, for reasons pointed out by Huggett (1995, 73-75). Fock space QFT is not so different from many-particle QM, at least not in the ways most salient to Teller's position. Teller's use of the occupation number basis obscures this fact. There are other, equally legitimate, ways of representing Fock space states. Consider the fact that Fock space is given by the direct sum

$$\mathbb{C} \oplus \mathcal{H} \oplus S_{+/-}(\mathcal{H} \otimes \mathcal{H}) \oplus S_{+/-}(\mathcal{H} \otimes \mathcal{H} \otimes \mathcal{H}) \oplus ... \tag{3}$$

where $S_{+/-}$ is the projection operator onto the symmetric (resp. antisymmetric) subspace in the case of bosons (fermions).[4] This means there is nothing wrong with writing down a Fock space state (in this case a state of indeterminate particle number) as

$$0 \oplus \frac{1}{\sqrt{2}}[|\psi\rangle \oplus \frac{1}{\sqrt{2}}(|\psi\rangle_1 \otimes |\phi\rangle_2 + |\phi\rangle_1 \otimes |\psi\rangle_2)] \oplus 0 \oplus ...,$$

permitting us, Teller's claims aside, to label the particles in the two-particle component of the state and think of them as switched (by permuting them as usual). The apparent difference between Fock space and the labeled tensor-product Hilbert spaces of QM is simply an artifact of the occupation number basis.[5]

[4]The first term ($\mathbb{C}$) in the direct sum represents the vacuum or "zero-particle" subspace.

[5]Teller himself notes that Fock space has this direct sum structure, but claims it is significant that the multi-particle subspaces are (anti-)symmetrized by $S_{+/-}$, leaving out asymmetric states like $|\psi\rangle_1 \otimes |\phi\rangle_2$. But

Butterfield (1993) insists that states of indeterminate particle number, a genuinely new development of Fock space QFT, frustrate any interpretation in terms of individuals: "For it seems to me that according to any reasonable notion of an individual (even a very liberal notion...), the number of individuals must be definite." (Butterfield, 1993, 474) Although this is suggestive, I do not see it as a decisive point. In interpretations that take a superposition of eigenstates of an observable to indicate that the observable's value is indefinite, it seems quite likely that the number of individuals will be indefinite even in QM (at least on a common-sense understanding of composition). For example, whether a given pair of protons form a helium atom will in some cases be indefinite, as it is when they are in a superposition of a bound state and a non-bound state. Thus it will be indefinite whether the helium atom exists. But it is not obvious that this rules out understanding the atom as an individual.[6] Similarly, one might understand QFT as describing individuals which sometimes exist indeterminately.

It appears that QFT on Fock space cannot decisively resolve the metaphysical questions surrounding identity in quantum theory. But Fock space is not the most general setting for QFT, and indeed, realistic interacting field theories cannot be formulated on Fock representations of the algebra of observables. Meanwhile, free QFTs on some curved spacetimes admit an infinity of inequivalent Fock representations, each with its own definition of "particle," and none physically privileged over the others. These more realistic theories do have the power to resolve (or rather, dissolve) the controversy about identity, because the particles that caused all the trouble in the first place are not part of their ontology.

3 QFTs without particles

Fock space is used to describe linear or free QFTs, that is, quantum fields which evolve in the absence of a force law. Such description involves an obvious idealization: real-world matter is always interacting via some force or other, and although spatially distant systems can sometimes be treated as non-interacting for practical purposes, strictly speaking we can

it's hard to see what the difference is supposed to be from many-particle QM, where this also holds true of each n-particle Hilbert space. Moreover, there is no obstacle to constructing a non-(anti-)symmetrized Fock space; indeed, such a space will be necessary if we want to represent paraparticles.

[6] Although the atom is not a fundamental entity, if we can make sense of indeterminate numbers of non-fundamental entities it does not seem like much of a leap to posit indeterminate numbers of fundamental entities.

never really "turn the interaction off." Similarly, we normally idealize by setting QFT in Minkowski spacetime and assuming special relativity is exactly true. But it is also possible to formulate QFT in more general, curved spacetimes, as predicted by general relativity. This means metaphysical conclusions drawn from interpreting the special relativistic Fock space formalism are suspect unless they can be reproduced in more physically realistic QFTs, namely interacting theories and QFTs on curved spacetime.

The interpretation of quantum states as describing particles is one feature of Fock space that cannot be reproduced when interactions are introduced. Although particles do appear as an approximate or emergent phenomenon in certain idealized limits, they are not, properly speaking, part of the theory's ontology. Meanwhile, in many curved spacetimes it is not possible to construct a QFT with a physically privileged definition of "particle." For these reasons, I think the metaphysical implications of particle permutations in Fock space (as well as the resulting controversies) can be dismissed, at least where our own physical reality is concerned.[7]

Let's first review the case against a particle ontology for the interacting theory. Fraser (2008) has argued against the existence of particles in "constructive" interacting AQFTs, although as she makes clear elsewhere, her argument also applies to the most physically realistic non-algebraic interacting QFTs (Fraser, 2006, 63-64). Fraser argues first that we have no grounds for interpreting a QFT as describing particles unless we can represent it using a Fock space. Fock space is the only known way of representing a system with infinitely many degrees of freedom as containing countable objects and a Lorentz-invariant zero-particle vacuum state. Using the constructive ϕ^4 theory as a representative case, she then shows that attempts to represent an interacting theory on Fock space must fail.

Fraser notes that, due to a foundational theorem called Haag's theorem, ϕ^4 and other interacting theories cannot be represented on the same Hilbert space as the free theory. Free and interacting fields must necessarily be defined on different, unitarily inequivalent Hilbert space representations. It is therefore impossible to interpret an interacting field as describing superpositions of free particles (Fraser, 2008, 14-20). A seemingly more promising option may be to construct a Fock space from the interacting one-particle space instead of the free single-particle Hilbert space. But Fraser also shows that this method cannot succeed. In

[7]Of course, if a Bohmian particle interpretation of QFT were adopted this would not be the case, but recall that the present discussion is restricted to the orthodox statistical approach.

an interacting theory, there is no covariant way to divide the wavefunctions into positive- and negative-frequency classes–a crucial step in the construction of a single-particle Hilbert space. As a result, there is no way to define a one-particle space that can be used to construct a Fock space (Fraser, 2008, 20-23). And without Fock space, there is no way of representing a state of the interacting field as containing aggregable particles. The particle concept cannot be applied to interacting QFT.

Although Fraser's central example of an interacting theory is an AQFT, her argument also applies to non-algebraic QFTs unless they are highly idealized. For example, Wallace (2011) has advocated, as an alternative to AQFT, the interpretation of QFTs which implement Wilsonian renormalization using short-distance length cutoffs. As Fraser (2006, 63-66) shows, Haag's theorem applies to these theories as well, unless they also employ physically unrealistic long-distance cutoffs. And when Haag's theorem applies, the interacting theory cannot be formulated on a Fock space, and so the premises of Fraser's argument are satisfied. Moreover, even in theories with long-distance cutoffs, where the free Fock space will be unitarily equivalent to the Hilbert space of the interacting theory, the resulting "particles" will not possess the expected physical properties.[8] So even rejecting the algebraic approach will not open the way for an acceptable particle interpretation of interacting QFT.

Particle interpretations also fail to apply when free QFT is formulated on general curved spacetimes, but for quite different reasons. Here the problem is not that a Fock space representation is unavailable. Rather, there may be many different Fock representations–each with a different definition of 'particle'–and no way to single out one representation as describing *the* particle ontology of the theory.

These inequivalent particle concepts appear because when constructing a Fock space, as already noted, one must divide the one-particle states into positive- and negative-frequency classes. This method employs a group of timelike isometries of spacetime to determine which solutions have future-directed momentum (positive frequency) and which have past-directed momentum (negative frequency). But some curved spacetimes have several different groups of timelike isometries. Each will correspond to a different Fock space.[9] If a moving observer's

[8]Renormalization offsets the values of certain quantities, like mass, which indicates that the "particles" of an interacting theory should possess this renormalized mass rather than the mass characteristic of the free theory. But if we interpret the interacting states using the free theory's particle concept, the resulting particles will not possess the renormalized values. Thanks to David Wallace for pointing this out.

[9]These Fock space representations are "different" in the sense of being unitarily inequivalent, a concept

history is given by a family of hypersurfaces related by a given group of isometries, the Fock space for that group will plausibly correspond to that observer's definition of "particle." There are a couple of ways to understand the resulting problem for particle interpretations. On one understanding, since one observer's particle concept may ascribe non-zero particle number to the state another observer would call a zero-particle vacuum, there is no objective fact about the particle content of QFT states in curved spacetime (Wald, 1994, 166). Alternatively, we might want to say that there is no single particle concept that applies to all the possible states of a QFT on curved spacetime, ruling out a fundamental ontology of particles for that theory (Ruetsche, 2011, 204-239).

Either way, particle interpretations of QFT on curved spacetime are doomed to fail. And again, this result is not an artifact of the algebraic approach. The existence of inequivalent Fock representations can be derived (although perhaps with less mathematical rigor) without employing AQFT (Birrell and Davies, 1984).

Although free QFTs in Minkowski space can be understood as possessing a particle ontology, more fundamental QFTs incorporating interactions and curved spacetime cannot. This leaves us with no reason to suppose that metaphysical puzzles about the identity of quantum particles apply to actual objects in our own world. Among other things this undercuts any motivation, from the permutation symmetry of quantum theory, for structuralist approaches to the metaphysics of identity. But it also raises a new puzzle, which I will turn to now.

4 Permutation symmetry is a superselection rule

This way of dissolving the metaphysical puzzles raised by quantum statistics leads quite directly to a new puzzle. As in QM, physical systems in QFT can still be classified into bosonic, fermionic, and (in physically unrealized examples) parastatistical. Indeed, one of the most important results of axiomatic QFT, the spin-statistics theorem, would be meaningless if states could not be classified according to their statistics.

But evidently, the statistics of a quantum field system incorporating interactions or curved spacetime cannot be defined in terms of permutations acting on particles. So we are faced with the new problem of explaining exactly what it means for a given QFT to be a theory of bosons or fermions. Independent of concerns about statistics and identity,

explained in §5.

a solution to this problem will be a welcome sanity check for existing arguments against particle interpretations of QFT.

A natural place to begin looking for a solution is superselection theory. It is often pointed out that the permutation symmetry of quantum mechanics is equivalent to a particular superselection rule–that is, a restriction on which states can be superposed, or (equivalently) on which self-adjoint operators can count as representing observable quantities. One of the great successes of the algebraic approach to QFT is a detailed account of superselection rules which applies to mathematically rigorous interacting QFTs as well as free theories. So perhaps the statistical properties of quantum fields can be defined in terms of a superselection rule, even in the absence of a particle ontology.

Let's take this one step at a time, starting with ordinary QM (QFT will be treated in §5). Why, exactly, is permutation symmetry equivalent to a superselection rule?

Superselection occurs when we have a quantity that behaves like electric charge, in the sense that superpositions of states with different values of the quantity cannot occur. In nature we never see a pure state that assigns non-zero probability to both charge 0 and charge +1, for example. The notion of a superposition being forbidden by a superselection rule can seem a bit murky. If ψ is a vector state with charge 0, and ϕ a vector state with charge +1, does the superselection rule imply that the linear combination of these two vectors doesn't exist?

As a matter of linear algebra, of course the linear combination exists. What's going on is that in the presence of a superselection rule, the linear combination $\frac{1}{\sqrt{2}}(\psi + \phi)$ represents a mixed state rather than a pure state. This means that this state is indistinguishable from $\frac{1}{\sqrt{2}}(\psi + e^{i\theta}\phi)$, since the relative phase of the pure states in a mixture is not observable. This can occur only if we forbid observables that map vector states to states with different charges, so that (for example) if $A\psi$ is a state with a charge different from ψ's, A cannot count as an observable even if it is self-adjoint.[10] This condition is equivalent to the existence of an operator which commutes with all observables; each of its eigenspaces then corresponds to a particular charge. We call such an operator a superselection operator. In QM, as we shall see, the unitary operators representing particle permutations are superselection operators. States

[10]In algebraic terms, this means that the space of states is a direct sum (e.g. $\mathcal{H}_1 \oplus \mathcal{H}_2$) and that the algebra of observables is given by the bounded operators on each subspace in the direct sum, rather than on the whole Hilbert space (e.g., by $\mathbf{B}(\mathcal{H}_1) \oplus \mathbf{B}(\mathcal{H}_2)$ instead of $\mathbf{B}(\mathcal{H}_1 \oplus \mathcal{H}_2)$).

from different eigenspaces of these operators cannot be combined in a coherent superposition, so we call the eigenspaces *coherent subspaces.*

In QM, permutation symmetry is generally taken to require that no permutation acting on particles can change the expectation value of any observable. That is, for any permutation P, observable A and state vector ψ,

$$\langle P\psi|A|P\psi\rangle = \langle\psi|A|\psi\rangle. \tag{4}$$

This requirement is sometimes referred to as the "Indistinguishability Postulate," and sometimes simply as "permutation invariance." Interestingly, it is equivalent to the requirement that each permutation operator P commute with all observables, i.e. that the permutations are superselection operators.

What is the analogue of "charge" for the superselection rule (in n-particle QM) that imposes permutation symmetry? Recall that states of different charge cannot be superposed to form pure states, a phenomenon which breaks down the space of states into coherent subspaces of states with the same charge, which are the eigenspaces of the superselection operator(s). These subspaces are often called *superselection sectors.* In effect, the values of charge are labels for the different superselection sectors.

What distinguishes the different superselection sectors for the permutation superselection rule? The eigenspaces of the permutation operators on n particles will correspond to the irreducible representations of S_n, the permutation group on n elements. In the familiar case, one of these representations (the trivial one) corresponds to bosons, while the alternating representation corresponds to fermions. Clearly no permutation will transform a system of fermions into a system of bosons. So the analog of charge is the statistics parameter which labels the irreducible representations of S_n and corresponds one-to-one with the so-called Young tableau diagrams (Halvorson and Mueger, 2007, 853). In effect, bosons have one value of this "charge," another value corresponds to fermions, and additionally there is a value for each order of para-Bose or para-Fermi statistics.

As we've seen, the quantum mechanical account of permutation symmetry as acting on particles cannot apply to the fundamental ontology of interacting QFT. But permutation symmetry is equivalent to a particular superselection rule, and a rigorous framework exists for understanding superselection rules in a broad variety of QFTs, including some interacting and curved-spacetime theories. Therefore, as I will establish in what follows, the permutation symmetry of QFT can still be understood in terms of this superselection rule. To explain how this is possible, we must first take a look at the mathematical theory of superselection in algebraic QFT.

5 Superselection theory in algebraic QFT

Before starting in on the details of superselection, I will briefly introduce the reader to AQFT. I will eventually defend this choice of framework, but the defense will make more sense once some of its details are on the table. The most basic elements of AQFT are best understood as a simple generalization of the formalism of ordinary QM. As in QM, the basic elements of the theory are states and observables. Observables (or physical quantities) are given by the self-adjoint elements of an abstract C*-algebra $\mathfrak{A}$, often called the algebra of observables. A state (or physical possibility) is then taken to be an assignment of expectation values to every quantity–mathematically, a functional $\omega : \mathfrak{A} \to \mathbb{C}$ which is normed (so that all probabilities sum to one) and linear. This will be our definition of a state on $\mathfrak{A}$.

Since QFT is a relativistic theory, we will need a way of connecting our physical quantities with the structure of spacetime. In order to implement the spacetime symmetries of special relativity and impose the requirement of relativistic causation, we need a definition of which physical quantities take on values in which regions of spacetime. So to each open region O of spacetime, we assign a subalgebra of the algebra of observables, $\mathfrak{A}(O) \subseteq \mathfrak{A}$. The subalgebra $\mathfrak{A}(O)$ is often said to contain all the observables measurable within O, but a better, less operationalist definition might be: the physical quantities whose values can be instantiated within O.

The connection between this algebraic formalism and the more familiar picture of states as density operators in a Hilbert space, and observables as operators on that space, is well understood. A given Hilbert space, with its associated operators, captures (at least part of) the physically significant structure of the algebra of observables if it constitutes a *representation* of $\mathfrak{A}$. A representation is a Hilbert space with a distinguished algebra of operators that mirror the algebraic structure of $\mathfrak{A}$. Mathematically, a representation of $\mathfrak{A}$ is a pair $(\mathcal{H}, \pi)$ consisting of a Hilbert space $\mathcal{H}$ and a *-homomorphism π from $\mathfrak{A}$ into $\mathbf{B}(\mathcal{H})$ (the bounded operators on $\mathcal{H}$). The density operators on $\mathcal{H}$ will then correspond to a subset of the states on $\mathfrak{A}$. We call a representation *irreducible* iff no nontrivial subspaces of $\mathcal{H}$ are invariant under $\pi\mathfrak{A}$.[11] We can think of an irreducible representation as one in which the operators in $\mathfrak{A}$ "connect" or relate all of the state vectors by mapping between them.

Importantly for present purposes, any reducible representation can be decomposed into a direct sum of irreducible representations. In this case, we may ask whether a linear combination of state vectors from different irreducible representations denotes a mixed state or a pure state. If a linear combination denotes a mixed state, clearly the two representations will behave like different superselection sectors. Thus a systematic way of determining when states from different representations can be superposed will allow us to reproduce the behavior characteristic of superselection sectors in QM. And in fact, the AQFT formalism gives us a means of determining when two representations should be counted as different sectors.

A central theorem of AQFT states that every state of the algebra of observables has a unique "home" representation, where it is represented by a state vector ψ which is *cyclic*:

GNS Theorem. For each state ω of $\mathfrak{A}$, there is a representation $(\mathcal{H}_\omega, \pi_\omega)$ of $\mathfrak{A}$, and a vector $|\psi\rangle \in \mathcal{H}_\omega$ such that $\omega(A) = \langle\psi|\pi_\omega(A)|\psi\rangle$, for all $A \in \mathfrak{A}$, and the vectors $\{\pi_\omega(A)|\psi\rangle : A \in \mathfrak{A}\}$ are dense in $\mathcal{H}_\omega$. If ω is a pure state, $(\mathcal{H}_\omega, \pi_\omega)$ is an irreducible representation. This representation is unique in the sense that for any other representation $(\mathcal{H}, \pi)$ satisfying the previous two conditions, there is a unitary operator $U : \mathcal{H}_\omega \to \mathcal{H}$ such that $U\pi_\omega(A) = \pi(A)U$, for all A in $\mathfrak{A}$.

This is considered the natural representation for ω to live in, since it contains only states that can be related to ω by operators in the algebra of observables. The criterion of "same representation" employed in the GNS theorem is called unitary equivalence.[12] In effect, the theorem tells us that a state's home representation (often called its GNS representation) is unique up to unitary equivalence.

Now, it can be shown that two algebraic states of $\mathfrak{A}$ can be superposed iff their GNS representations are unitarily equivalent (Baker and Halvorson, 2010, 109fn22). This means that a superselection sector corresponds to a unitary equivalence class of irreducible representations of $\mathfrak{A}$. When representations $(\mathcal{H}, \pi)$ and $(\mathcal{H}', \pi')$ are unitarily inequivalent, a

[11] That is to say, if there is no nontrivial subspace $\mathcal{S} \subset \mathcal{H}$ such that $\pi(A)\mathcal{S} = \mathcal{S}$ for all $A \in \mathfrak{A}$.

[12] Since the unitary intertwiner U provides an isomorphism whose existence implies that the states of the two representations make the same probabilistic predictions for all observables, unitary equivalence is often considered sufficient for physical (or at least empirical) equivalence in quantum theory.

superselection rule forbids the superposition of states from $\mathcal{H}$ with states from $\mathcal{H}'$.

This means it is possible in principle to represent superselection rules in AQFT. But one would prefer something more: the ability to derive which superselection rules hold in a given theory from the algebra of observables for that theory. A goal of the algebraic approach is to derive the interesting physical structure of a quantum theory from properties of its algebra of observables (at least where possible). So rather than treating a theory's superselection rules as a brute posit, we would like to be able to derive them from the structure of this algebra.

In fact, such a derivation is possible. This is the lesson of the Doplicher-Haag-Roberts (DHR) theory of superselection, one of the landmark results in AQFT. This approach was first explored by Doplicher *et al.* (1969).

DHR theory begins by setting a selection criterion to narrow down which states of $\mathfrak{A}$ are counted as physically possible. Typically, the number of mathematically well-defined states on an algebra of observables includes many states with physically unrealistic properties (e.g., states of determinate position or momentum in QM). A selection criterion is a condition states must meet to count as possibilities. Even in the absence of physical justification, a selection criterion can be useful in constructing new theoretical frameworks. By pretending that only a narrow class of states are physically possible, we can develop a physical concept that applies to those states in the hope that it can be further generalized in the future.

In this spirit, the DHR approach relies on a selection criterion which is admittedly too restrictive to admit all of the known physical possibilities. They require that states meet the

DHR selection criterion: Let $(\mathcal{H}_0, \pi_0)$ be the GNS representation induced by the privileged vacuum state ω_0 of $\mathfrak{A}$. A representation $(\mathcal{H}, \pi)$ of $\mathfrak{A}$ is *DHR* iff (1) for each Minkowksi double cone O, the representations $\pi_0|_{\mathfrak{A}(O')}$ and $\pi|_{\mathfrak{A}(O')}$ are unitarily equivalent;[13] and (2) $(\mathcal{H}, \pi)$ possesses finite statistics, that is, a finite-dimensional representation of the permutation group. Here O' is the spacelike complement of O, $\pi|_{\mathfrak{A}(O')}$ is the restriction of the representation π to the subalgebra $\mathfrak{A}(O')$, and $\mathfrak{A}(O')$ is the C^*-algebra generated by $\mathfrak{A}(O_1)$ with O_1 a double cone spacelike separated from O. A state is DHR if it is representable by a density operator in a DHR representation.

The intuitive idea is that DHR states are states which differ from the vacuum only within a

[13] A double cone is the intersection (*not* the union) of two light cones.

finite spatial region. This requirement is not met by states in electrodynamics, since Gauss's law implies that these states are distinguishable from the vacuum everywhere in space.[14] On the other hand, it is possible for asymptotically free theories like chromodynamics and the electroweak theory to satisfy the DHR condition.

For now, let's impose the DHR condition as an admitted idealization and see what it gets us. Recall that our goal is to derive the superselection structure of a QFT from its algebra of observables. As we've seen, a unitary equivalence class of representations of $\mathfrak{A}$ behaves like a superselection sector. So when the DHR condition is applied, an equivalence class of DHR representations will act as a superselection sector for DHR states.

A useful notion in superselection theory is that of an *endomorphism*: effectively a way of shuffling around the operators in an algebra without changing its overall algebraic structure. Since any DHR representation $(\mathcal{H}, \pi)$ is unitarily equivalent to the vacuum outside of some region O, it is unitarily equivalent to $(\mathcal{H}_0, \pi_0 \circ \rho)$ where $\rho : \mathfrak{A} \to \mathfrak{A}$ is an endomorphism of the algebra of observables localized within O (Halvorson and Mueger, 2007, 800-803).[15] That is, acting on the vacuum representation with ρ (by composing π_0 with ρ) gives us the DHR representation $(\mathcal{H}, \pi)$. Whenever ρ and ρ' map the vacuum representation to the same DHR representation–physically, whenever they create the same quantity of charge–they will belong to the same unitary equivalence class $\tilde{\rho}$. It can thus be shown that the DHR representations correspond one-to-one to the unitary equivalence classes $\tilde{\rho}$ of localized endomorphisms.[16] So we can use localized endomorphisms to stand for DHR representations. For present purposes, this will be convenient.

Recall that the charges (or additive quantum numbers) of a quantum theory are effectively labels for superselection sectors. In particle physics, as has been well confirmed by experiment, a theory's charge quantum numbers are determined by its group of global internal symmetries (its gauge group). It can be shown that an AQFT's (unitary equivalence classes

[14] By contrast, the finite statistics requirement is easily met, since bosons, fermions and paraparticles of any finite order all satisfy it. States with infinite statistics are often considered almost pathologically strange.

[15] A *-endomorphism ρ is *localized* within O iff $\rho(\mathfrak{A}) = \mathfrak{A}$ outside O. We also require that endomorphisms be transportable in the sense that there exist unitary operators in $\mathfrak{A}$ which map between endomorphisms localized in different regions (Halvorson and Mueger, 2007, 786, Def. 152). For simplicity one usually deals only with endomorphisms localized within double cones (so O is only allowed to range over double cones), but this assumption is eliminable.

[16] Mathematically, this means that the category of DHR representations of $\mathfrak{A}$ is naturally isomorphic to the category Δ of localized, transportable *-endomorphisms of $\mathfrak{A}$ (Halvorson and Mueger, 2007, 800-801).

of) DHR representations, which act as superselection sectors, correspond one-to-one to the charge quantum numbers one would expect given its gauge group (Baker and Halvorson, 2010, 110-114).

This means that DHR theory successfully predicts the structure of the superselection sectors for any QFT satisfying the AQFT axioms and the DHR condition. Rather than simply stipulate which superselection rules hold, we may derive them from the structure of the algebra of observables. This constitutes one of the greatest explanatory successes of the AQFT framework. Unfortunately, the domain of this explanation is presently limited to QFTs satisfying the DHR condition, which we know does not apply to all physically possible states. But the DHR account has already been generalized to states meeting the weaker Buchholz-Fredenhagen condition, which requires equivalence to the vacuum in the causal complement of one spacelike cone, and there is considerable hope for generalizing the framework further (Doplicher, 2010, 728).

At the beginning of this section, I promised some remarks in defense of my choice of the AQFT framework. The trustworthiness of AQFT, for physics-based metaphysics, is the subject of lively debate (Fraser, 2006; Wallace, 2011). For present purposes, though, I think AQFT is clearly the best framework available, for two reasons. First, even if AQFT is not a trustworthy source of evidence for metaphysics, it is important that we establish the conceptual possibility of permutation-symmetric QFTs with no particle interpretation–and as we will see, AQFT can accomplish this. Second, superselection theory is one of the areas where AQFT is likely to be a trustworthy source of metaphysical evidence, since it concerns only the so-called infrared domain.

My first point is easy to grasp if we direct our attention to the dialectic. Recall that our main concern here is to resolve a conceptual puzzle: how can we make sense of QFTs without particles, when these theories are supposed to be permutation symmetric and the definition of permutation symmetry seems to depend on particles? Even if AQFT can teach us nothing about the metaphysics of the actual world, it can help resolve this puzzle by making clear why permutation-symmetric AQFTs without particles make conceptual sense. Thus the negative thesis of this paper–that permutation symmetry in QFT has no consequences for the metaphysics of identity, since there are no particles in interacting and curved-spacetime QFTs–escapes the charge of incoherence motivated by this conceptual puzzle.

That said, I also believe that superselection is one of the domains where AQFT can be trusted to inform metaphysics. To see why, we must examine the reasons that have been given *not* to trust AQFT. The novel feature of the theory, one not treated rigorously by alternative approaches, is the infinity of inequivalent representations, which result from its infinitely many degrees of freedom. Infinitely many degrees of freedom become relevant when one considers the limit of very short distance scales (UV degrees of freedom) or very large scales (IR degrees of freedom). In AQFT terms, these are associated with two sorts of unitary inequivalence of representations: local unitary inequivalence, in which representations are inequivalent even when restricted to finite volumes, and global inequivalence, which only appears in infinite regions of spacetime.

Wallace (2011) has argued persuasively that AQFT's infinitely many UV degrees of freedom are un-physical artifacts of its mathematical representation, and hence any interpretive implications of locally unitarily inequivalent representations can be disregarded. His argument, which I accept for present purposes, is that events at very short length scales (equivalently, very high energies) are outside QFT's domain of application, and renormalization theory gives us positive reason not to trust the theory's predictions in the UV domain. Fortunately, the only inequivalent representations relevant to DHR superselection theory are locally unitarily *equivalent*, despite being globally inequivalent. In short, Wallace has given us excellent reason to question AQFT's predictions concerning locally inequivalent representations, but is "happy to grant that *long-distance* divergences really should be tamed by algebraic methods," noting that "most quantum field theorists" can be expected to agree (Wallace, 2011, 123). Since DHR theory applies only in this long-distance domain where AQFT's applicability is not in question, we have excellent reason to treat it as our best, most fundamental theory for present purposes.

Now that we know how to represent superselection rules in AQFT, the way is open to investigate permutation symmetry. We will see that, as in QM, permutation symmetry in AQFT is equivalent to a superselection rule, where the charges correspond to bosons, fermions and the different orders of parastatistics. We will also see that the action of the permutation group on physical states in AQFT can be understood perfectly well even in the absence of a particle interpretation.

6 Permutation symmetry generalized

We've seen that every DHR representation (up to unitary equivalence) corresponds to a localized endomorphism of the algebra of observables. Although it isn't immediately obvious, these endomorphisms have a natural physical interpretation. They "create" charge, in the sense that applying an endomorphism localized within O to a state ω will map ω to a different state that contains some charge initially localized within O.

Endomorphisms are defined as acting on operators in the algebra $\mathfrak{A}$. What does it mean to apply one to a state? Recall that states themselves are functionals on $\mathfrak{A}$, which map operators to their expectation values. This means that $\omega \circ \rho = \omega(\rho(A))$ is also a state for any localized endomorphism ρ. Applying ρ to ω in this way will alter ω's expectation values within the region O where ρ is localized, leaving them unchanged elsewhere. It will also move ω to a different superselection sector, assuming ρ is non-trivial. In particular, if ω is a state in the vacuum representation $(\mathcal{H}_0, \pi_0)$, ρ will take ω to a state in the sector corresponding to its equivalence class $\tilde{\rho}$, that is, a state of the representation $(\mathcal{H}_0, \pi_0 \circ \rho)$. So ρ creates charge, in this case by mapping a state in the (neutral/un-charged) vacuum sector to a state in a sector of nonzero charge. In this sense, its effect on a state is somewhat analogous to that of a Fock space creation operator, the difference being that the latter introduces a single particle rather than a unit of charge. (From now on, we'll treat the vacuum as our "basic" state, setting $\omega = \omega_0$.)

The key to implementing permutation symmetry (and hence, statistics) in the DHR approach is the realization that a representation of the permutation group S_n can be constructed on the charge-creating morphisms. Suppose that $\rho_1, \rho_2, ...\rho_n$ are unitarily equivalent endomorphisms; physically, this means that acting on a state with ρ_i generates the same amount of charge in region O_i that ρ_j creates in region O_j when it acts on a state. Then, because the category Δ of endomorphisms is a tensor category, we have a well-defined notion of the tensor product of these morphisms, which amounts to composing them:

$$\rho_1 \otimes \rho_2 \otimes ... \otimes \rho_n = \rho_1 \circ \rho_2 \circ ... \circ \rho_n.$$

Our choice of a unitary equivalence class of morphisms, along with the symmetry property of the category Δ, uniquely determines a representation of S_n on this tensor product. If this is the trivial representation, we say that $(\mathcal{H}_0, \pi_0 \circ \rho_i)$, the sector associated with the ρ_i, is a bosonic sector. If it is the alternating representation, we say the ρ_i correspond to a fermionic sector. If the representation is higher-dimensional, the sector is parastatistical.

The physical meaning of permutations acting on endomorphisms isn't immediately obvious. To understand it, it helps to work within a Hilbert space representation of $\mathfrak{A}$ that is a direct sum of sectors. In such a "big" representation $(\mathcal{H}, \Pi) = (\bigoplus_i \mathcal{H}_0, \bigoplus_i \pi_0 \circ \rho_i)$, the sector $(\mathcal{H}_0, \pi_0 \circ \rho_i)$ for each charge $\tilde{\rho}_i$ will be a coherent subspace. And among the operators on $\mathcal{H}$, we can find unobservable field operators that create charge by moving states between these coherent subspaces. In this sense, these field operators act like the morphisms ρ_i. In particular, suppose ψ_0 is the vector in $\mathcal{H}$ corresponding to the vacuum state ω_0. Then we can define F_i to be an operator on $\mathcal{H}$ such that $F_i\psi_0$ is a vector state in the sector $(\mathcal{H}_0, \pi_0 \circ \rho_i)$.[17] This whole construction may seem pedantic (why talk about F_i instead of ρ_i itself?), but in fact it will allow us to see how the statistics of a sector is reflected in the vector states of that sector and their behavior under permutations.

First, suppose that $\psi_i = F_i\psi_0$ is a vector state that implements the state $\omega_0 \circ \rho_i$ in the representation $(\mathcal{H}_0, \pi_i) = (\mathcal{H}_0, \pi_0 \circ \rho_i)$ corresponding to $\tilde{\rho}_i$'s superselection sector. (This means that all of ψ_i's expectation values for the operators $\pi_i(A) \in \pi_i(\mathfrak{A})$ agree with $\omega_0 \circ \rho_i$'s expectation values for the operators $A \in \mathfrak{A}$.) Then, in the sector corresponding to $\rho_i \circ \rho_j$, there will be a state vector $F_iF_j\psi_0$ corresponding to $\omega_0 \circ \rho_i \circ \rho_j$. In effect, this will be a composite of the vector state ψ_i (corresponding to the abstract state $\omega_0 \circ \rho_i$) and the vector state ψ_j (corresponding to the abstract state $\omega \circ \rho_j$). So let's define the "product" vector $\psi_i \times \psi_j$ to be the vector state $F_iF_j\psi_0$.

Note that the "product" we just defined is *not* the tensor product on $\mathcal{H}_0$. If ψ_i and ψ_j are states in bosonic sectors, for example, $\psi_i \times \psi_j$ will be a symmetric state, even though the expression $\psi_i \times \psi_j$ changes when we permute the indices i and j. So in terms of the tensor product on $\mathcal{H}_0$, we may (for example) have $\psi_i \times \psi_j = \frac{1}{\sqrt{2}}(|\psi_i\rangle \otimes |\psi_j\rangle + |\psi_j\rangle \otimes |\psi_i\rangle)$.

Now we have the formal machinery in place to see how the statistics of a sector determine the statistical properties of its states. Suppose ψ_i and ψ_j are states in the same sector (which is to say, they correspond to unitarily equivalent morphisms, i.e. $\tilde{\rho}_i = \tilde{\rho}_j$). Then, if the sector they live in has Bose statistics, permutations will act symmetrically on the product state and we will have $\psi_i \times \psi_j = \psi_j \times \psi_i$. On the other hand, if they live in a Fermi sector, we

[17] In the case of Fock space theories, the field operators F_i will be creation and annihilation operators. There will not generally be a unique F_i for a given ρ_i.

will have $\psi_i \times \psi_j = -\psi_j \times \psi_i$.

In general, if a given sector is identified by the DHR method as bosonic, a vector state composed of n charges from this sector will transform according to the trivial representation of S_n under permutations of the order of the n charges. A vector state composed of n charges from a Fermi sector will transform according to the alternating representation of S_n. The same formula applies to parastatistics. A composite of n charges from a para-Fermi (resp. para-Bose) sector of order m will transform according to the mth-order para-fermionic (para-bosonic) representation of S_n. This may seem no different in substance from the standard quantum mechanical approach to statistics, and indeed (as we will see) in particle theories the two are identical. But there is a crucial difference: the entities being permuted are not necessarily single-particle states. Instead they are charge-creating morphisms.

This allows the DHR picture of statistics to apply even to interacting theories which lack the Fock space structure needed to sustain a particle interpretation. In her argument against particles, Fraser (2008) uses the constructive ϕ^4 theory (in two and three spacetime dimensions) as a representative example. Although as Fraser shows it lacks a particle interpretation, the states of this theory meet the DHR condition, and so the DHR picture of statistics applies to them. From the point of view of superselection theory, this is pretty uninteresting, since the superselection behavior of ϕ^4 theory is trivial. Since it treats only Bose fields and lacks any internal symmetries, all of the theory's states belong to a single superselection sector. But there are non-trivial examples as well.

The earliest example of an interacting theory with non-trivial sectors is the Yukawa interaction in two spacetime dimensions. Summers (1982) proves that this theory satisfies the DHR condition. It provides a simple model of an interaction between a scalar field and a spin-1/2 fermion field, approximating the scattering behavior of mesons and nucleons. Due to the theory's global $U(1)$ symmetry, its charge quantum numbers are given by the integers (the charge is carried by the Fermi field). In addition, the sectors break down into Bose and Fermi sectors, since the theory includes both types of statistics.

If Fraser's line of argument (recounted in §3) is correct, neither of these theories admits an ontology of particles. And although no mathematically rigorous form of the Standard Model has yet been constructed, we can expect these problems for particle interpretations to persist in future constructive QFTs. So the standard quantum mechanical account of statistics, as governing what happens to a state when we permute its particles, cannot apply to existing interacting theories and probably will not apply to their successors. Nonetheless, there is a perfectly well-defined sense in which the states of these theories can be classified into Bose, Fermi and parastatistical states, which correspond naturally to the expected representations of the permutation group. Because the DHR approach to superselection applies, the states break down into sectors which uniquely determine their statistical properties.[18]

The success of DHR is even clearer in curved spacetime. The approach has been generalized to apply to all globally hyperbolic spacetimes (Guido *et al.*, 2001). Free QFTs on many of these, e.g. Schwarzschild spacetime, exhibit multiple inequivalent particle concepts as described in §3. So in fact there are two classes of QFTs without particles where the DHR account nonetheless applies: constructive QFTs and free QFTs on curved spacetime.

In the absence of particles, the physical interpretation of this formalism is not obvious: what does it mean physically to permute the morphisms that generate units of charge, when the morphisms are mathematical objects and units of charge are values of a physical quantity rather than concrete entities? By analogy, it would make little sense to "permute" two values of the electric field. I see this as an outstanding question for future research. One thought: perhaps the permutation invariance of field theories consists in the fact that, when constructing a charged state from the vacuum, the order in which one generates the charges does not matter to the physical interpretation of the resulting state. Whether this is a satisfactory picture or not, the point remains that no particle-based picture will be satisfactory in interacting theories, and this is sufficient to establish my (negative) thesis.

It will be instructive to perform one further "sanity check" on the DHR approach to statistics, by verifying that the standard picture of permutations acting on particles emerges as a special case in Fock space QFTs which do exhibit particle behavior.

7 Particles as a special case

The two canonical examples of QFT on Fock space are the free spin-0 field (real Klein-Gordon theory) and the free spin-1/2 field (Dirac theory). Since the former is a theory of bosons and the latter a theory of fermions, these examples taken together encompass both experimentally observed forms of statistics. We will now see how the DHR approach

[18]A necessary caveat: of course it has not yet been established that DHR superselection theory can be generalized to encompass a rigorous version of the Standard Model, if in fact one can be constructed.

applies to both examples, reproducing the predictions of the standard particle-based picture of statistics. In short: when we permute particles in a Fock space QFT, from the DHR perspective we are permuting charge-generating morphisms.

On the standard picture, of course, every state of the Klein-Gordon (KG) theory obeys Bose-Einstein statistics. In DHR terms, this means there should be no Fermi or parastatistical sectors. In fact, the superselection theory for KG fields is completely trivial. The theory has only a single sector. This is the vacuum sector, of course, which must (as always) be a Bose sector. So for all states ψ_i, ψ_j, we will have $\psi_i \times \psi_j = \psi_j \times \psi_i$–every composite state will be a symmetric state, just as one would expect.

The case of Dirac theory is a little more interesting, because in this case the superselection theory is (barely) non-trivial. There are two sectors in Dirac theory. One is a Fermi sector, as one would naively expect. But since Dirac theory includes a vacuum state, there must also be a vacuum sector for that state to live in. And since every vacuum sector is a Bose sector, the Dirac theory has both a Bose and a Fermi sector.

This may seem surprising, at first. How can the theory of a fermion field include a sector with Bose statistics? This becomes clearer if we keep in mind what it means for a sector to have Bose statistics. It doesn't mean that the states of that sector, taken individually, will transform like symmetric wavefunctions when their particles are permuted. Instead it implies *only* that, if we form a composite state out of states from that sector, this composite state will transform symmetrically when we permute the states we composed to form it.

Now suppose that ψ_1 is a state of *two* fermions, and likewise for ψ_2. How should we expect the composite state $\psi_1 \times \psi_2$ to behave when we permute ψ_1 and ψ_2?[19] To permute ψ_1 and ψ_2 is, effectively, to permute two pairs of fermions. We switch the two fermions in the state ψ_1 with the two fermions in state ψ_2. But permuting two pairs of particles makes no physical difference, even to an antisymmetric state. While performing one particle permutation on an antisymmetric wavefunction will flip its sign, performing a second permutation flips it back again, and we are left with the same state we started with. Thus, if ψ_1 and ψ_2 are two-fermion states, we should expect that $\psi_1 \times \psi_2 = \psi_2 \times \psi_1$. But this is just what it means for ψ_1 and ψ_2 to belong to a Bose sector. In general, we should expect states containing even numbers of fermions to live in a Bose sector rather than a Fermi sector.

[19]Remember that this product state $\psi_1 \times \psi_2$ is the one we defined in the previous section, and is not the same as the tensor product $\psi_1 \otimes \psi_2$.

In fact, this is exactly how the sectors of the Dirac theory break down. To elaborate with a bit more formal detail: the state space for Dirac theory is the antisymmetric Fock space,

$$\mathcal{F} = \mathbb{C} \oplus \mathcal{H} \oplus S_{-}(\mathcal{H} \otimes \mathcal{H}) \oplus S_{-}(\mathcal{H} \otimes \mathcal{H} \otimes \mathcal{H}) \oplus ..., \tag{5}$$

where S_{-} denotes projection onto the antisymmetric subspace. The terms in this direct sum break down into two subspaces whose direct sum is the whole Fock space. That is, we have $\mathcal{F} = \mathcal{F}^e \bigoplus \mathcal{F}^o$, where $\mathcal{F}^e$, the even subspace, is the subspace of $\mathcal{F}$ with even particle number:

$$\mathcal{F}^e = \mathbb{C} \oplus S_{-}(\mathcal{H} \otimes \mathcal{H}) \oplus S_{-}(\mathcal{H} \otimes \mathcal{H} \otimes \mathcal{H} \otimes \mathcal{H}) \oplus ..., \tag{6}$$

while $\mathcal{F}^o$, the odd subspace, is given by

$$\mathcal{F}^o = \mathcal{H} \oplus S_{-}(\mathcal{H} \otimes \mathcal{H} \otimes \mathcal{H}) \oplus ..., \tag{7}$$

which includes only states with an odd number of particles.

Acting on the Dirac vacuum state with creation and annihilation operators (or equivalently, with the Dirac theory's field operators), we can generate the entire antisymmetric Fock space. However, the field operators of the Dirac theory are not elements of its algebra of observables.[20] When we act on the vacuum state with operators from the algebra of observables, the Hilbert space generated is the even subspace $\mathcal{F}^e$, not the full antisymmetric Fock space $\mathcal{F}$. This means that the even subspace, and not the full Fock space, is the GNS representation of the vacuum state. In the language of superselection theory, the even subspace is the vacuum sector, and thus the odd subspace must constitute a different superselection sector, which must be reachable from the vacuum by the action of charge-creating morphisms. The theory's charge-creating morphisms (or rather, the concrete operators that represent them, which we labeled F_i in §5) must act on a state in the even subspace to generate a state in the odd subspace, and vice versa. So clearly these are given by the creation and annihilation operators on the antisymmetric Fock space. In this sense, the n-particle states of Fock space are generated from the vacuum by (in DHR terms) "adding charges."[21]

[20]The reason for this is that the microcausality axiom of AQFT requires spacelike separated observables to commute, but spacelike separated spinor fields (like the Dirac field operators) anticommute.

[21]Thanks to Hans Halvorson for helping to verify these points about the superselection structure of the Dirac theory.

As noted above, states with odd numbers of fermions will obey Fermi statistics (in the DHR sense), while even-numbered states will obey Bose statistics. So we would expect the even subspace to count as a Bose sector, and the odd subspace as a Fermi sector. Since the vacuum sector always exhibits Bose statistics, the former is obviously true. And when we construct the representation of the permutation group on the class of charge-generating morphisms corresponding to the odd subspace's superselection sector, indeed we find that it is a Fermi sector.[22] The statistical predictions of the DHR picture reproduce the predictions of the standard particle picture for free QFTs. Thus the statistics of particles, in those QFTs that support a particle ontology in the first place, emerge as a special case of the more general DHR account of statistics for quantum fields.

8 Conclusions

We've now seen how to make sense of the notion that a quantum state obeys Fermi-Dirac statistics, or Bose-Einstein statistics, even if said state cannot be understood as describing a system of particles. Permutation symmetry can still be understood as a superselection rule, even in the absence of particles to be permuted. Moreover, we've seen that the statistical properties of particles, for those states that do possess a particle interpretation, emerge as a special case of this more general picture of statistics.

In the DHR picture, permutations act on the morphisms that connect states of differing charge. Permutation symmetry then amounts to the notion that the order in which these morphisms are applied to a state–the order in which charge is formally added to an algebraic state to generate a different state–does not matter to which state is generated. In the case of free QFTs admitting a particle interpretation, the charge-creating morphisms can be represented by particle creation and annihilation operators, and permutation symmetry has its usual physical interpretation. In the interacting and curved-spacetime field theories that fall under the domain of DHR, however, this physical interpretation is not available, since a particle ontology is ruled out by the arguments surveyed in §3.

What to make, then, of the notion that the statistical behavior of quantum particles

[22] To see why, consider what happens when we change the order in which we act on a state with the operators representing the morphisms, i.e. the creation or annihilation operators. Since these operators anticommute, the state is multiplied by -1 when we change the order of two of them. Thus the corresponding representation of S_n is the alternating representation.

undermines classical conceptions of individuality, or leaves it indeterminate whether classical notions of individuality apply? Under the assumption that our most fundamental theory of matter is non-relativistic QM, or free QFT on flat spacetime, this notion may be of great interest. But in fact, our most fundamental theories are interacting QFTs and QFTs on curved spacetime. There is no natural way for these theories to accommodate an ontology of particles, except as an approximation in the limited domain where they resemble Fock space QFTs. The reader may be tempted to understand the DHR approach to statistics as somehow re-casting the existing debate about identical particles in terms amenable to realistic QFTs. But to the contrary, my point is that there is no analogue of the existing debate in interacting or curved-spacetime QFTs. So puzzles about the statistical behavior of quantum particles would seem not to bear on the question of whether the actual world is made up of individuals. According to the QFTs that offer the best available approximation to reality, there are no quantum particles, and we have no particular reason to expect that they will be re-introduced by some later, more fundamental theory.

For those seeking evidence from physics about the nature of identity or individuality in our universe, there are better places to look than quantum mechanics. One potential source of insight is offered by those interpretations of quantum theory which do posit particles, namely Bohmian interpretations. For example, Durr *et al.* (2005) have advanced a Bohmian particle theory that can (they argue) reproduce the predictions of interacting QFTs. A study of whether these Bohmian particles conform to our classical concept of individuals would be of great value, especially to those tempted by such an interpretation.

The other promising avenue for potential research concerns the treatment of individuals in spacetime theories, a project which is already underway (see French and Krause, 2006, 65-80). Since QFT is probably best understood as describing the assignment of fundamental quantities to regions of spacetime (Wallace and Timpson, 2010), it is plausible that the best candidates for the "individuals" posited by the theory are spacetime points, or spacetime regions. For this reason it seems to me that spacetime theories, and not basic quantum mechanics, should be the locus of philosophical debate about the nature of identity and individuality in modern physics.

Acknowledgements

Hans Halvorson deserves considerable thanks for helping me develop these ideas. Many thanks to Gordon Belot, Dan Peterson and Chip Sebens for illuminating comments on a previous draft. Discussion of these and related issues with Noel Swanson was also crucial to the development of this paper.

This material is based upon work supported by the National Science Foundation under Grant No. SES-1127260.

References

Baker, David John and Hans Halvorson (2010), "Antimatter," *British Journal for the Philosophy of Science* 61:93–121.

Birrell, N.D. and P.C.W. Davies (1984), *Quantum fields in curved space*, Cambridge: Cambridge UP.

Butterfield, Jeremy (1993), "Interpretation and Identity in Quantum Theory," *Studies in History and Philosophy of Science* 24:443–476.

Dasgupta, Shamik (2009), "Individuals: An Essay in Revisionary Metaphysics," *Philosophical Studies* 145:35–67.

Doplicher, S. (2010), "Spin and Statistics and First Principles," *Foundations of Physics* 40:719–732.

Doplicher, Sergio, Rudolf Haag and John E. Roberts (1969), "Fields, observables and gauge transformations. I," *Communications in Mathematical Physics* 13:1–23.

Durr, Deltef, Sheldon Goldstein, Roderich Tumulka and Nino Zanghi (2005), "Bell-Type Quantum Field Theories," *Journal of Physics A* 38:R1–R43.

Fraser, Doreen (2006), *Haag's Theorem and the Interpretation of Quantum Field Theories with Interactions*, University of Pittsburgh: unpublished dissertation.

Fraser, Doreen (2008), "The fate of 'particles' in quantum field theories with interactions," *Studies in History and Philosophy of Modern Physics* 39:841–859.

French, Steven and Decio Krause (2006), *Identity in Physics*, Oxford: Oxford UP.

Guido, Daniele, Roberto Longo, John E. Roberts and Rainer Verch (2001), "Charged sectors, spin and statistics in quantum field theory on curved spacetime," *Reviews in Mathematical Physics* 13:125–198.

Halvorson, Hans and Michael Mueger (2007), "Algebraic Quantum Field Theory," in Butterfield, Jeremy and John Earman (eds.), *Philosophy of Physics*, North-Holland.

Huggett, Nick (1995), "What are Quanta, and Why Does it Matter?" *PSA 1994 (Proceedings of the 1994 Biennial Meeting of the Philosophy of Science Association* 2:69–76.

Ladyman, James and Don Ross (2007), *Every Thing Must Go*, Oxford: Oxford UP.

Ruetsche, Laura (2011), *Interpreting Quantum Theories*, Oxford: Oxford UP.

Summers, Stephen J. (1982), "Normal Product States for Fermions and Twisted Duality for CCR- and CAR-Type Algebras with Application to the Yukawa2 Quantum Field Model," *Communications in Mathematical Physics* 86:111–141.

Teller, Paul (1995), *An Interpretive Introduction to Quantum Field Theory*, Princeton, NJ: Princeton UP.

van Fraassen, Bas C. (1991), *Quantum Mechanics: An Empiricist View*, Oxford: Oxford UP.

Wald, Robert M. (1994), *Quantum Field Theory on Curved Spacetime and Black Hole Thermodynamics*, Chicago: University of Chicago Press.

Wallace, David (2011), "Taking particle physics seriously: A critique of the algebraic approach to quantum field theory," *Studies in History and Philosophy of Modern Physics* 42:116–125.

Wallace, David and Christopher Gordon Timpson (2010), "Quantum Mechanics on Spacetime I: Spacetime State Realism," *British Journal for the Philosophy of Science* 61:697–727.