Particles, Objects, and Physics

Justin C. Pniower

Christ Church, Oxford University

Submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy

Trinity Term 2005

Particles, Objects, and Physics

Abstract

Justin C. Pniower
Christ Church, Oxford University

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy
Trinity Term 2005

This thesis analyses the ontological nature of quantum particles. In it I argue that quantum particles, despite their indistinguishability, are objects in much the same way as classical particles. This similarity provides an important point of continuity between classical and quantum physics. I consider two notions of indistinguishability, that of indiscernibility and permutation symmetry. I argue that neither sort of indistinguishability undermines the identity of quantum particles. I further argue that, when we understand indistinguishability in terms of permutation symmetry, classical particles are just as indistinguishable as quantum particles; for classical physics also possesses permutation symmetry.

The permutation symmetry of classical physics does not conflict with classical statistics. While permutation symmetry is necessary for quantum statistics, it is not sufficient. The difference between classical and quantum statistics is therefore not due to a difference in permutation symmetry. It is rather due to the nature of the quantum state-space and the specific manifestation of permutation symmetry in the quantum formalism. Because both classical and quantum physics possess permutation symmetry, permutation symmetry is not related to a lack of particle identity. I argue that it is instead related to the modal thesis of anti-haecceitism, and that this connection reflects the qualitative nature of both classical and quantum physics.

I also examine the relation between a particle and field ontology. I argue that, while quantum particles only provide an emergent quantum ontology, they are still objects. I also consider and dismiss other arguments against the individuality of quantum particles based on surplus structure and statistical correlations. I further consider whether quantum particles provide a clear example of ontic vagueness that might undermine their individuality. I conclude that quantum theory does not provide any clear examples of ontic vagueness, and that any vagueness involved is essentially the same as we find in classical examples.

For Opa

Acknowledgments

One incurs many debts in the completion of a thesis. I would like to thank the American Friends of Christ Church and the British CVCP for freeing me of those that would be financial. I still, however, am burdened by many intellectual debts. The greatest is to my supervisor Simon Saunders. But there are also many other I must thank. Harvey Brown has not only provided guidance, but also an example of what a philosopher of physics should be. There are still others deserving of thanks. My conversations with Paul Mainwood, Oliver Pooley, Chris Timpson, and David Wallace have been invaluable. And of course I would like to thank my examiners, Jeremy Butterfield and Steven French.

While the philosophy department, and specifically the philosophy of physics group, has exercised the most direct intellectual influence my work, my college, Christ Church, has provided for me in countless ways. I would specifically like to thank Hugh Rice, my college moral tutor.

One of the advantages of Oxford's collegiate system is that you meet many students outside your chosen discipline. I have met many great people in my time here. They have not only been friends, but also intellectual influences. I would like to thank Joseph DiNunzio, Ronald Fitzpatrick, Ian Flanagan, Diana Gibson, Ashley Hern, Christopher Jones, Erland Koch, Paul Lettow, Sinead Mooney, Jörg Peltzer, Rosalind Porter, Gilles Rico, Alexander Shaw, Heidi Strebel, Catherine Stoodley, Melissa Terras and countless others who have offered friendship and guidance.

Finally, I must thank my family, although no thanks would ever be sufficient.

Contents

Chapter 1

Introduction

What first captivated me about physics was quantum physics. It was high school, I was about fifteen or sixteen, and my sole introduction to quantum physics was the briefest of chapters at the very back of the shabby textbook provided by the school. I think that was the first chapter I actually read in the book. There I was first told of a world very different from the one I grew up in, where nothing was certain, things could be in more than one place, interact in strange ways, and even meld together. It was a world of entanglement, uncertainty, and complementarity. This radical shift in world view, even though it was but outlined, complemented my young rebellious temperament. Not only could I reject my parents' tastes, views, and prejudices, but I could also reject, through my study of quantum physics, their understanding of the physical world around them.

But youth always gives way to age. And while I am by no means old, the rebellious spirit of my youth has been tempered with the conservative tendencies of maturity. And with this change in disposition, my view of a radically different quantum world has given way to a more measured interpretation that recognizes the continuities and similarities between the classical and quantum. Of course the quantum world is and always will be dramatically different from the classical world, but I have come to believe that many differences are overstated or simply incorrect, and many continuities and similarities are ignored or missed.

This brings us to the topic of this thesis—the nature of quantum particles. In the quantum revolution, particles are often taken to be the most fanatical members of the revolutionary party. Quantum particles differ drastically from their classical counterparts. Quantum particles, unlike classical particles, are indistinguishable; and because of this indistinguishability, quantum particles, unlike classical particles, lack identity, individuality, and objecthood.

"Indistinguishability" is a term used in many ways by physicists. There are, however, two important notions of indistinguishability that appear to have implications for the individuality and identity of quantum particles. One notion understands indistinguishability in terms of indiscernibility. Quantum particles, unlike classical particles, are indistinguishable because they can be qualitatively indiscernible. It is therefore impossible to ground the identity of quantum particles by use of qualitative properties and relations, and this leads us to doubt whether they have an identity at all. The other notion understands indistinguishability in terms of permutation symmetry. Quantum particles, unlike classical particles, are indistinguishable because a permutation of them does not lead to a distinct physical situation. We can understand this in terms of a lack of identity and objecthood. It because quantum particles are not objects that a permutation of them does not lead to a distinct physical situation.

This difference between the ontology of classical and quantum particles appears to have empirical consequences, for indistinguishability is intimately connected with such things as quantum entanglement and quantum statistics. A lack of particle identity therefore appears to be an important way in which the quantum world differs from the classical world.

It is impossible to fill in the details of these views with any brevity. That will be the task of the following chapters. But we can safely say that the predominate opinion among physicists and physically minded philosophers is that quantum particles, unlike classical particle, lack identity. They present us with a new ontology. This new ontology opens new philosophical horizons; for we need new philosophy to properly understand this new category of thing. We need new logics, like Schrödinger logic, and new set theories, like quaset and quasi-set theory. And there are also other philosophical ramifications. Indistinguishable quantum particles have been used to illustrate a controversial theories of vagueness—ontic vagueness—and bolster an unconventional view of scientific realism—ontic structural realism. And broadly speaking, if we are to successfully interpret quantum theory, our interpretation must capture or explain away the peculiar features of indistinguishable quantum particles and the radical shift in ontology between classical and quantum particles.

It is, however, the contention of this thesis that there is no radical shift in ontology. Quantum particles are objects in much the same way as classical particles. Instead of providing a point of departure from classical to quantum theory, the particle ontology of quantum theory actually provides an important point of continuity, a bridge that allows much of our traditional conceptual apparatus—our logic, our set theory, our semantics—to cross from one theory to the other unperturbed.

To be clear, this study of quantum ontology will mostly remain neutral with respect to the measurement problem. While we cannot claim to fully elucidate the nature of quantum ontology without providing a solution to the measurement problem, we can still elucidate, without directly addressing the measurement problem, the relations between indistinguishability, entanglement, quantum statistics, and the nature of quantum particles. We will argue that, insofar as we can understand a particle ontology independently of any specific solution to the measurement problem, quantum particles are objects in much the same way as classical particles.

Certainly quantum particles are not like classical particles in all respects. We will highlight some important ways in which they differ throughout this thesis. We therefore need to be clear about what notion of objecthood is at work here. The notion of objecthood that we will use in this thesis is a very basic one. It is one that is suggested by Quine's two main slogans: "No entity without identity," and "To be is to be the value of a variable." As suggested by the first slogan, we will hold that an object possesses self-identity: it is identical to itself and numerically distinct from everything else. And as suggested by the second slogan, we will also hold that an object is that which we describe by use of standard first-order quantificational logic. An object is that which we can quantify over and predicate properties and sentences to. Objects have to be able to enter into the domain of standard first-order quantifiers, domains that we describe by use of sets. We therefore further hold that objects are described by standard set theory. Quine no doubt associated more extensive views with his two slogans. We will, however, simply content ourselves with this basic notion of objecthood.

This notion of objecthood is very general. It does not distinguish physical objects from abstract objects. It does not require the possession of determinate spatiotemporal properties, or any specific sort of properties. And it also does not require that objects persist through time. But, as we will see, even this very minimal notion of objecthood is at odds with the prevalent view about the nature of quantum particles.

Our goal will be to show that, contrary to popular opinion, quantum particles do satisfy this notion of objecthood and are objects without problem. Rather than inciting revolt, the nature of quantum particles provides a conservative element in an otherwise revolutionary theory. The fact that quantum particles are objects, in the sense just stipulated, allows for the application of our standard philosophical views of identity, logic, set theory, and semantics to the quantum realm; for these philosophical views are intertwined with this notion of objecthood.

Before we move on to outline the contents of this thesis in greater detail, let us clarify some of our terminology. We will use term "particle" in the most generic way possible to refer to molecules, atoms, and fundamental particles. We therefore will not assume that particles lack further internal structure. And of course we will use the term in a way that is neutral with respect to the question of whether particles are objects or not.

When we talk about the identity of an object, what we will mean is that the object is identical to itself and numerical distinct from everything else. We are not talking about any act of identification through demonstrative or descriptive reference. Often physicists call particles that possess all the same state-independent intrinsic properties identical particles. To avoid any confusion with questions of numerical identity, we will only talk about particles of the same kind.

Throughout we will use the terms "individual" and "object" interchangeably. This use of the term "individual" differs from standard usage. Often individuals are taken to be persisting physical objects. Our notion of objecthood presumes neither that an object is physical or persisting. This slight difference in usage has rhetorical advantage and should not lead to any confusion.

We will use the term "non-individuals" to refer to particles that fail to be objects. The use of this term does not presuppose how the particles fail to be objects. For example, it might be that non-individual particles are a whole new category of entity that lacks self-identity, or it might be the case that particles are just properties of underlying quantum fields.

Now that all of that is out of the way, let us return to our topic. In order to establish our conclusion that quantum particles are objects, we will examine the philosophical arguments to the contrary and point out where they go wrong. In the process, we will remove the motivations for rejecting the identity of particles. We will also dispute the attempts to explain how quantum particles fail to be objects. While French and Redhead (1988) have already argued that it is possible, despite the reasons to the contrary, to take quantum particles as objects, we will go further and argue not only that we can take quantum particles as objects, but that we should.

One important group of arguments about the nature of quantum particles centers on indiscernibility and Leibniz's principle of identity of indiscernibles (PII). As we have already mentioned, one way we can understand indistinguishability is in terms of indiscernibility. This indiscernibility appears to have important consequences for the identity of quantum particles because of the plausibility of PII, which holds that no objects are qualitatively indiscernible.

Now there are a few relevant questions here. First, what motivations ground PII? Second, what is the precise relation between PII and identity and individuality? And third, to what extent are quantum particles qualitatively indiscernible?

The debate surrounding these questions is varied and active. There are arguments that hold that (at least some) quantum particles are indiscernible, and arguments that hold that these particles are therefore not objects. Margenau (1944), French and Redhead (1988), van Fraassen (1991), and Saunders (2003) are just some of the people who have made important contributions to the debate. It will be our goal in chapter 2 to separate the various threads of argument and understand what implications any indiscernibility might have for the individuality of quantum particles. We will present, as best we can, the various motivations for PII and establish to what extent quantum particles are actually indiscernible. We will conclude that some quantum particles can be indiscernible, but that this indiscernibility does not undermine the individuality of quantum particles. In the process we will present and defend certain views about individuality and identity. We will argue that the identity of indiscernible objects is primitive, and that there is no need for any additional individuators, such as unique spatiotemporal trajectories or non-qualitative properties.

Another important group of arguments about the nature of quantum particles centers on permutation symmetry and quantum statistics. As we have already mentioned, another way we can understand indistinguishability is in terms of permutation symmetry. Permutation symmetry appears to have important consequences for the identity of quantum particles because it seems to follow from a lack of particle identity. And since quantum statistics follows from permutation symmetry, quantum statistics is an empirical manifestation of this lack of particle identity.

In chapter 3, we will present the details and historical background behind this putative connection between permutation symmetry and a lack of particle identity. We will also point out a critical flaw in the position. This view assumes that there is a difference between the permutation symmetry of classical and quantum physics. But this is not so. Classical physics is just as permutation symmetric as quantum physics, as we will demonstrate through an analysis of classical dynamics and statistical mechanics. We will further bolster this claim by showing that the permutation symmetry of classical physics is necessary in order to ensure a correct expression for the statistical mechanical entropy.

In chapter 4, we will return to the topic of quantum statistics. As we have mentioned, it is often held that quantum statistics follows from the indistinguishability of quantum particles, and that the difference between quantum and classical statistics is due to a difference in indistinguishability. In chapter 4, we will look at the relation between indistinguishability—understood in both terms of indiscernibility and permutation symmetry—and quantum statistics. Quantum statistics follows from the symmetrization requirement. We will show that the neither the indiscernibility nor permutation symmetry of quantum particle implies symmetrization, and with it quantum statistics. We will argue that symmetrization is simply an empirical hypothesis.

We will then consider the difference between classical and quantum statistics. In the quantum formalism, a lack of permutation symmetry leads to classical statistics. The difference between classical and quantum statistics therefore seems to be due to some difference in permutation symmetry (even if permutation symmetry does not imply symmetrization). There are certain derivations in the phase space formalism that suggest the same is true in classical physics. In these phase space derivations, permutation symmetry leads to quantum statistics, while a lack of permutation symmetry leads to classical statistics. These phase space derivations conflict with our claim that classical physics, which includes classical statistics, is permutation symmetric. But we will show that the conflict is only apparent. Permutation symmetry is consistent with classical physics and classical statistics in the classical formalism.

In chapter 5, we will return to permutation symmetry itself. Since both classical and quantum physics possess permutation symmetry, we cannot take permutation symmetry to follow from any lack of individuality. But permutation symmetry still has important philosophical implications. In chapter 5, we will show how the permutation symmetry of physics relates to the modal thesis of anti-haecceitism, and how physics gives support to this metaphysical thesis. Anti-haecceitism holds that the representations *de re* of a possible world supervene upon the qualitative character of that possible world. The permutation symmetry of both classical and quantum physics ensures that physics is anti-haecceitistic. That is the permutation symmetry of physics essentially implies that physics is a qualitative theory. It makes use solely of qualitative properties and relations to describe its subject matter. There is no essential use of names or labels. As we will see, instead of following from a lack of particle identity, permutation symmetry actually follows from the qualitative nature of physics. And instead of providing an important difference between classical and quantum physics, permutation symmetry provides an important point of continuity, for both classical and quantum physics are qualitative in this anti-haecceitistic sense.

Now particles do not provide the only quantum ontology. There is also a field ontology. A field ontology appears to be more fundamental than a particle ontology; for only fields offer an adequate ontology for relativistic quantum field theory. A particle ontology therefore only provides an emergent or effective ontology for a phenomena in a certain regime. In chapter 6, we will examine the relation between a particle and field ontology and argue that even if a particle ontology supervenes upon a field ontology, this does not undermine the individuality of quantum particles. Quantum particles are still objects, even if they are only emergent objects.

In chapter 6, we will also look at another well known argument against the individuality of quantum particles: Redhead and Teller's surplus structure argument (Redhead and Teller, 1992). The basic outline of this argument is that if we take quantum particles as objects, we are left with some unwelcome surplus structure in our physical theory. Since we can easily excise this surplus structure when we deny that quantum particles are objects, Redhead and Teller claim that we have, at the very least, good methodological grounds for rejecting the individuality of quantum particles. In chapter 6, we will examine in detail this argument. We will build on some of the standard criticism leveled against it by Huggett (1994) and van Fraassen (1991), and conclude that there is nothing about the surplus structure argument that suggests that quantum particles are not objects.

There is also another ontological argument that we will consider in chapter 6. This one follows from the work of Reichenbach (1998) and Dieks (1990) and focuses on the statistical correlations between quantum particles. This argument holds that, since there is no apparent causal explanation for these statistical correlations, we should abandon a particle ontology in favor of a field ontology, where the correlations do not exist. Following van Fraassen (1998), we will argue that since similar correlations pervade quantum theory, regardless of whether or not we adopt a field ontology, these sorts of correlations do not undermine a particle ontology.

Having addressed the main arguments against the individuality of quantum particles, we turn, at the end of chapter 6, to the question of metaphysical underdetermination. French and Redhead (1988) have claimed that, even if we have good reason to doubt the individuality of quantum particles, it is still possible to regard them as objects. French and Ladyman (2003b) have built upon this position and claimed that it is metaphysically underdetermined whether particles are or are not objects. They propose ontic structural realism as a way to alleviate this underdetermination. Both Cao (2003) and Pooley (forthcoming), in their criticisms of ontic structural realism, have rejected that there is such underdetermination for quantum particles. While it might be possible to take quantum particles as objects, the philosophical evidence is overwhelmingly in favor of an ontology where particles lack individuality, or so they argue.

But it is the purpose of this thesis to establish the contrary, that we have good reason to regard quantum particles as objects. This does not mean that we will strengthen the claim of metaphysical underdetermination. Just the opposite. Instead of alleviating any underdetermination by rejecting individuals in favor of non-individuals, we will claim that we should reject non-individuals in favor of individuals. There is thus no metaphysical underdetermination of the nature of quantum particles, and thus no support here for ontic structural realism.

In chapter 7, we end our study by looking at whether quantum particles are ontically vague. There is a growing debate about whether vagueness, instead of just having its origins in our language, concepts, or knowledge, actually exists in the world. Lowe (1994) has claimed that quantum particles provide a good example of such ontic vagueness. The identity of quantum particles is vague because the particles themselves possess indeterminate identity. Lowe's actual examples of ontic vagueness are questionable, but

French and Krause (1995) have supplemented his views and bolstered these claims of ontic vagueness.

If it is true that quantum particles are ontically vague, then there is the worry that this might be due to their lack of individuality. In chapter 7, we will examine putative examples of ontic vagueness in quantum theory. We will conclude that quantum examples are no more persuasive than classical examples in establishing ontic vagueness. We will further conclude that if it turns out that objects are ontically vague, this will no more undermine the individuality of quantum particles than it will other classical objects.

As the outline of this thesis has made clear, our study will be one that often crosses the border, however ill defined, between physics and metaphysics. In our pursuit to understand the nature of quantum particles, we will engage with several subjects of traditional metaphysics. We will explore not only objecthood, individuality, and identity in general, but also, for example, delve into those perennial and important philosophical problems of possibility and vagueness. Because of this, an additional product of our study will be to illustrate first hand the important interaction between metaphysics and physics. As we will see, the two are often equal partners, both offering criticism and insight to the other. With this said, let us now commence with our study.

Chapter 2

The Principle of Identity of Indiscernibles

2.1 The Indistinguishability of Quantum Particles

As we have mentioned in the introductory chapter, we often say that quantum particles are indistinguishable. At the very least this means that particles of the same kind possess all the same state-independent intrinsic properties, such as mass and charge. But often a stronger claim is being made, the claim that quantum particles of the same kind are indiscernible, that they possess all the same properties and relations within a system, so that there is absolutely no way we can tell them apart. Indistinguishable particles, in this sense, violate the Principle of Identity of Indiscernibles (PII), the principle that no objects can differ numerically without differing qualitatively, that no object can be indiscernible from another.

Some hold that if quantum particles are indistinguishable in this strong sense, then we cannot regard them as objects. These ontological arguments are the subject of this chapter. Below we will examine how PII relates to quantum particles, and whether PII leads to any conclusion about the ontology of quantum physics. Our conclusion will be that, even though quantum particles can be indiscernible, they are still objects in much the same sense as classical particles. PII has no bearing on the ontology of quantum physics.

Before we continue, we should note that PII is not only relevant to the ontology of quantum theory. It was first introduced by Leibniz in order to argue for a relational view of space and time (Leibniz and Clarke, 1956). There is some debate as to Leibniz's actual use of the principle, but the basic argument is that substantival space and time violate PII, and that space and time are therefore relational. What we say here will have implications for this argument. In this chapter, however, we shall concentrate on the arguments of quantum theory. We will return to the topic of space and time in chapter 5.

2.2 The Traditional Formulation of the Principle of Identity of Indiscernibles

Let us start by considering the traditional formulation of PII and some well known counterexamples to the principle. Loosely PII holds that no two objects can possess all of the same properties and relations. We can state the principle with more precision in second-order logic. In its standard second-order formulation, PII states:

$$\forall x \forall y((\forall F(Fx \equiv Fy)) \supset (x = y)) \tag{2.1}$$

The strength of the principle depends upon the domain of the second-order quantifier. If the quantifier is unrestricted, so that it includes properties like "$= a$", where a is a singular term referring to some particular object, then PII is trivially true.

PII is not trivial if we restrict the domain to qualitative properties and relations. Roughly, a qualitative property is a property that does not imply or refer to particular objects. "$= a$" is not a qualitative property, but "having mass m" is—even if there is only one object with mass m.[1]

[1]There is no consensus on a more precise definition of qualitative properties. (The two main candidates are given by Adams (1979) and Rosenkrantz (1979).) Luckily we can remain neutral on this philosophical issue. All we need to do is admit some serviceable distinction between qualitative and non-qualitative properties. Provided such a distinction exists, our arguments will not be affected by any particular analysis of the distinction.

We can arrive at different versions of PII if we place further restrictions on the domain of properties (see French, 1989b). We can specify a version of PII where we limit the domain of the second-order quantifier to intrinsic properties. Roughly, an intrinsic property, as opposed to an extrinsic property, is a property that an object possesses in virtue of its own nature, independently of anything else. For example, "having mass m" is an intrinsic property, while "having a mass greater than Jupiter" is an extrinsic property. Possession of the former property only depends upon the nature of the given object, while possession of the latter depends upon the nature of both the given object and Jupiter.[2]

A version of PII that only admits intrinsic properties seems closest to Leibniz's original proposal, but it is too strong to be of much interest. Classical particles provide a good counterexample to its validity. Two classical particles at different positions are indiscernible by this version of PII if they possess the same intrinsic properties, like mass and charge, yet we do not identify them. So it seems like we must also include extrinsic properties of some sort within the domain of the quantifier. We can come up with any number of other versions. For the moment let us consider the version that includes all qualitative properties—both intrinsic and extrinsic—in its domain.

This version of PII has several well known counterexamples. These counterexample usually consist of highly symmetric possible worlds. The first to present a counterexample of this sort was Kant (1965, A263–264/B319–320), but the example given by Black (1952) is better for our purposes.[3] Black asks us to consider a possible world that consists solely of two iron spheres at some distance from each other. He further stipulates that the two spheres have the same shape and uniform constitution, and therefore share all intrinsic properties. The two spheres also appear to share all extrinsic properties, like

[2]As with qualitative properties, there is no consensus on precise definitions of intrinsic and extrinsic properties. Once again we can remain neutral on this issue, for our arguments will not be affected by any particular analysis of the distinction. All that is required is that a serviceable distinction between intrinsic and extrinsic properties exists.

[3]Similar counterexamples are also given by Ayer (1954), Strawson (1959, ch. 4), Adams (1979), and Wiggins (2001, p. 188).

being a given distance from a sphere. The two numerically distinct spheres therefore appear to be indiscernible, violating PII.

We can overcome counterexamples of this sort if we admit absolute position as a qualitative property. Since the absolute position of each sphere differs, the spheres are discernible. It is, however, difficult to understand the existence of absolute positions in a symmetrical universe; for in such universes, we cannot analyze absolute positions in terms of a unique reference frame. In such symmetrical universes, the relations between objects or spacetime points (if spacetime points exist) do not uniquely define any such frame.

But even those willing to grant the existence of absolute positions cannot overcome the further counterexamples provided by quantum mechanics. In quantum mechanics, systems that consist of particles of the same kind must be symmetrized: for bosons, the permutation of two particles must produce a symmetric state, and for fermions, an antisymmetric state. Because of the symmetrization requirement the reduced density matrix and the conditional and unconditional probabilities of measuring some physical magnitude (like a given momentum) are the same for each particle in the system.[4] If quantum mechanics is complete, then a particle's state-dependent properties supervene upon the reduced density matrix and conditional and unconditional probabilities of measurements. Each particle therefore possesses the same state-dependent properties, and, since the particles are of the same kind, the same state-independent properties, making them indiscernible. And the particles are still indiscernible even if we grant absolute position; for all the spatial properties of a particle—including its position—also supervene upon its reduced density matrix and conditional and unconditional probabilities (again assuming quantum mechanics is complete). Regardless of absolute position, quantum systems provide plausible counterexamples to PII.

[4]This point has been made in different ways by Margenau (1944), French and Redhead (1988), van Fraassen (1991), and Butterfield (1993).

2.3 A More General Principle of Identity of Indiscernibles

The counterexamples of the last section point out the deficiencies of the traditional formulation of PII. If the principle is to serve any interesting philosophical role, we must focus on a weaker, yet non-trivial, version of PII, a version that does not succumb to these counterexamples. Such a version exists. It originates in the work of Hilbert and Bernays (1934), and has been elaborated and defended by Quine (1960, 1976) and Saunders (2003).

Instead of focusing on a second-order formulation, let us consider open sentences. Using open sentences we can define three grades of discernibility. Let us call two objects *absolutely discernible* if there is a sentence in one free variable that is true of one of the objects, but not both. To put it more formally, x and y are absolutely discernible if and only if there is a an open sentence F in one free variable, such that $Fx\&\neg Fy$. For example, two particles of different mass are absolutely discernible. Let us call two objects *relatively discernible* if there is an open sentence in two free variables that is true of both objects, but only in one order. x and y are relatively discernible if and only if there is an open sentence F in two free variable, such that $Fxy\&\neg Fyx$. Some ordinal numbers provide examples of relatively discernible objects that are not strongly discernible (Quine, 1976). The sentence "$x < y$" is true of two ordinal numbers in only one order. Points of time in an empty Newtonian spacetime provide another example (Saunders, 2003). The sentence "x is earlier than y" is true of two points of time in only one order. Finally, let us call two objects *weakly discernible* if there is an open sentence in two free variables that is true of both of them, but not true of one of them. x and y are weakly discernible if and only if there is an open sentence F in two free variables, such that $Fxy\&(\neg Fxx \vee \neg Fyy)$. We will consider examples of weak discernibility shortly.

Now that we have specified these three grades of discernibility, we can formulate a more general version of PII: **Let PII be the principle that if x and y are not absolutely, relatively, or weakly discernible, then x = y.** Once again, in order to avoid a trivial version of the principle, we need to restrict our attention to sentences composed solely of variables and qualitative properties and relations—no proper names or non-qualitative properties allowed.

The traditional second-order formulation of PII stated above only concerns sentences that are open in one free variable: it states that if x and y are not absolutely discernible, then they are identical. By expanding to consider sentences open in not just one, but two free variable, we can also consider relatively and weakly discernible objects.

This more general version of PII overcomes the counterexamples leveled against the traditional formulation. Consider again Black's two spheres. If we do not appeal to absolute position, then the two spheres are not absolutely discernible. But this does not mean that they are identical. According to the more general PII (from now on just called PII), they are weakly discernible. There is a symmetric irreflexive relation—"distance d apart, measured center to center"—that weakly discerns between the two spheres. Black's spheres do not provide an example of *indiscernible* objects, only *weakly discernible* ones. The same is true of other highly symmetric possible worlds; there are similar symmetric irreflexive relations that allow us to weakly discern between the objects in those worlds.

Now consider quantum systems. Let us take the simple example of two non-interacting spin-1/2 fermions of the same kind. Let each particle be described by the same single-particle spatial state-function, so that the spin component of state-function is an antisymmetric superposition of orthogonal spinors. As we noted above, each particle in the system will have all the same intrinsic state-independent properties, the same reduced density matrix, and the same unconditional and conditional probabilities. The particles therefore do not appear to be absolutely discernible. But again there is a symmetric irreflexive relation that makes the particles weakly discernible, as Saunders (2003, p. 294) points out. While the superposition of spinors does not assign a specific spin direction to either fermion, it does state that they will always have opposite components of spin. In any direction, the probability that both particles are spin up or that both are spin down is zero, while the probability that one particle is spin up and the other is spin down is one.

We therefore can use the symmetric irreflexive relation "has spin component in the opposite direction from" to weakly discern between the two particles. The particles are not indiscernible, only weakly discernible. This is true for any system of fermions of the same kind: the antisymmetry of the state-function ensures that there are always symmetric irreflexive relations that allow us to weakly discern between fermions.

In the early days of quantum theory, it was commonly held, most notably by Weyl (1949, p. 247), that fermions satisfied PII because of the Pauli exclusion principle. The thought was that the Pauli exclusion principle forbids two fermions of the same kind from being in the same state, and therefore ensures that they have different properties. As Margenau (1944) was the first to point out, this is not strictly true since the two fermions will have the same reduced state. But as we have just noted, a similar view is true. The Pauli exclusion principle, which is just another statement of antisymmetrization, ensures that fermions of the same kind satisfy PII, not by preventing them from being in the same state, but by providing symmetric irreflexive relations that weakly discern between the fermions. Weyl is at least partially vindicated.

What about systems of bosons of the same kind? The state-functions for these systems must be symmetric, not antisymmetric. This allows a bosonic system to be in a state where each boson is described by the same single-particle state-vector. For such a state there will be no symmetric irreflexive relations that allow us to weakly discern between the bosons. Bosons of the same kind in such a state are indiscernible and therefore violate PII.

However we can deny that this provides a counterexample to PII, for we can deny that fundamental bosons are objects. PII therefore would not apply. In analogy with classical physics, we consider fermions, like electrons, to be the constituents of matter, and fundamental bosons, like photons, to be related to the interactions between matter. According to this view, fundamental boson are not objects, but are rather discrete excitations of modes of the interaction fields. Composite bosons, like certain atoms, are still objects, but they are never indiscernible; we discern between them by use of their fermionic constituents. This is a plausible view, and it has been defended by Saunders (2003, p. 294–295). We will examine it in more detail below, but for the moment let us simply note that PII appears to overcome all the standard counterexamples leveled against it.

2.4 Motivations for the Principle of Identity of Indiscernibles

Despite the recent discussion of PII in the philosophy of science and the philosophy of physics, a majority of the philosophical community take PII either to be a false or philosophically uninteresting principle. This is partly because of the putative counterexamples to the principle (which, as we now see, fail), but it is also because of the failed attempts to employ PII as a logically or metaphysically necessary principle that serves in a theory of reference, or as a criterion of identity, or even as a definition of the identity relation "=" itself. Wittgenstein (1961, 5.5302), Black (1952), Strawson (1959, ch. 4), Dummett (1973, ch. 16), and Wiggins (2001) have convinced a majority of the philosophical community, this author included, that PII is not a logically or metaphysically necessary principle, and that it fails in these roles.

However this does not mean that PII is false or philosophically uninteresting, for there are further ways we can employ PII. We can accept that PII is not logically or metaphysically necessary, but still regard it as a principle related to our physical theories. Since traditional counterexamples to PII fail, nothing seem to bar us from taking this position. In fact there are three main motivations for doing so: an epistemological one, a metaphysical one, and a methodological one. Let us take them in turn.

First let us consider the epistemological motivation. We do not appear to have any epistemological access to indiscernible objects. We never directly observe numerically distinct yet indiscernible objects, and we have no criterion of identity that allows us establish whether indiscernible objects are identical or not. Because of this lack of epistemological access, indiscernible objects do not appear to add to the empirical content of a theory. Their presence within a theory is arbitrary and unnecessary. We should therefore always prefer a simpler theory that obeys PII.

Now let us consider a metaphysical motivation for PII. The idea is that PII provides the only acceptable principle of individuation for physical objects. A principle of individuation tells us why an object is the object it is and no other object. When we take PII as a principle of individuation, what makes an object the object it is and no other object is that the object is discernible according to PII. The object either possesses a unique set of properties and relations or enters into individuating irreflexive relations with other objects.[5] We might further combine this view with a property-bundle or trope theory of objecthood (although it is not necessary to do so).

Some might balk at the use of individuating relations. Both Russell (1956) and Armstrong (1978, p. 94–95) claim that it leads to a vicious circle. They argue that objects must first be individuated before they can enter into extrinsic relations, such as the symmetric irreflexive relations we have cited above. We can, however, hold without contradiction that an object can simultaneously enter into and be individuated by a relation, just as we can hold that it can simultaneously instantiate and be individuated by a property. But if an object is individuated by extrinsic properties and relations, then even though the object will only be identical with itself, its identity—the fact that it is an object identical to itself and distinct from other objects—will not be an intrinsic feature of the object. It is only by relation to other objects that it is the object it is and no other. To some this is no doubt an unappealing position.

Further worries arise when we consider the relations that weakly discern between objects. These relations are not only extrinsic, but also external. An external relation (sometimes called a non-supervenient or inherent relation) is one that does not supervene upon the intrinsic properties of the relata (see

[5]Castañeda (1975) argues that we should separate the question of what makes an individual an individual as opposed to a universal from the question of what makes an individual different from other individuals. We will however ignore this distinction since all of the principles of individuation we will consider answer both of these questions in mostly the same way.

Lewis, 1986, p. 62).[6] Both the relations "distance d apart measured center to center" and "has spin component in the opposite direction from" are external relations. If an object, such as a Black sphere or a fermion, is individuated by the external relations it enters into, then the identity of that object not only fails to be an intrinsic property of that object, it also fails to supervene upon any set of intrinsic properties.

This might not bother us in the case of a spatial relation like "distance d apart measured center to center". For it is easy to grant that spatiotemporal relations enjoy a special status of some sort that allows them to provide an organizational framework for phenomena in a universe (see Lewis, 1986, p. 74–76). From here it is only a small step to admit that this special status allows spatiotemporal relations to individuate objects. (All of this is of course very Kantian.) But it might bother us that external relations like "has spin component in the opposite direction from" individuate objects, for these relations do not supervene upon spatiotemporal relations.

But even with extrinsic and external relations, PII is superior to the two main alternative principles of individuation. One of these alternatives individuates objects by use of special non-qualitative properties called haecceities or primitive thisnesses (Adams, 1979). An haecceity is the property of being identical to a particular object, like Bertrand Russell. An object is the object it is and no other object because it uniquely possesses an haecceity. The other main alternative appeals to some sort of individuating substance or bare particular. An object is the object it is and no other object because it uniquely consists of a given substance.[7]

While each of these alternative proposals might have its supporters, we should hold them under suspicion.[8] Neither haecceities nor individuating substances appear explicitly in our physical theories. We also do not have any epistemological access to them, at least for those associated with physical

[6]Not all extrinsic relations are external. For example, the extrinsic relation "more massive than" is not external. Whether or not it holds between two objects depend upon the intrinsic masses of the objects.

[7]See French (1989b) for more details on these alternative principles of individuation.

[8]Wiggins (2001, p. 125–126) presents a good criticism of haecceities. Van Fraassen (1991, p. 462) gives a brief summary of the problems of individuating substance.

objects.[9] For these reasons we naturally question whether we understand the identity of haecceities or substances, and further question the clarity of these concepts themselves. We cannot hope to provide an analysis of the identity of objects in terms of haecceities or substances if the identity of haecceities or substances is more mysterious than the identity of objects. And even if they are coherent and well understood, we should question whether they contribute to the empirical content of a theory, for we have no epistemological access to them.[10]

And even if they do contribute to the empirical content of a theory, we should question if haecceities and individuating substance threaten to make physics either non-qualitative or incomplete. We generally consider physics as a qualitative theory that at least aspires to be complete. It does not appear to rely upon non-qualitative properties, or recognize non-qualitative differences. Physics appears to apply to a physical system independently of considerations about what particular objects are involved in that system. But if we admit that haecceities or substances have empirical content, then physics must either include haecceities or substances and loose its qualitative nature, or exclude them and fail to describe important aspects of the empirical world.

For all these reasons, and most likely many more, PII is the only one of these alternatives that provides a suitable principle of individuation for physical objects, even with individuating relations. It is therefore a principle intimately connected with our physical theories.

[9]Some, such as Swinburne (1995), claim that every conscious being possesses an haecceity of which he is aware. Even if this is true, it does not give us reason to assign haecceities to physical objects lacking consciousness, nor to suppose that we have epistemological access to haecceities of physical objects.

[10]Even admitting these empirical worries, we can still accept an innocuous version of haecceities. For example, we can accept that "identical to Bertrand Russell" is a property, so long as we realize that it is a gerrymandered property that is dependent upon the independently established reference of the proper name "Bertrand Russell". Such an haecceity might be acceptable, but it is also incapable of doing any interesting metaphysical work, like providing a principle of individuation, since it is completely dependent upon the reference of other terms. Another option, following Lewis (1986, p. 225), is to take "identical to Bertrand Russell" to be a property that is equivalent to the singleton set that contains only Bertrand Russell. Such an haecceity is equally innocuous, but also equally incapable of doing any interesting metaphysical work since it is individuated by the element of the set and not the other way around.

Finally let us consider a methodological motivation for PII, specifically Simon Saunders's argument for PII (Saunders, 2003). For Saunders, PII serves as a methodological principle that guides the interpretation of the experiments and mathematical formalism of a theory. He takes an interpretation of a theory, which assigns properties to objects, as determined by the declarative sentences that we are willing to assert in light of the theory. As he writes:

> I suggest it is through *talk* of objects, in the light of mathematical theories and experiments, that we achieve a clear interpretation of these theories and experiments in terms of physical objects—our understanding of what objects there are, I am suggesting, is clearest in our use of simple declarative sentences. (Saunders, 2003, p. 290)

When making declarative statements in light of the theory, what is most evident are the physical properties and relations of the theory like charge and momentum. These provide us with an initial interpretation of the theory. This initial interpretation also includes the putative objects to which these properties apply, like bosons and fermion. "[W]e may read off the predicates of an interpretation from the mathematics of a theory, and, because theories are born interpreted, we have a rough and ready idea of the objects that they are predicates of" (Saunders, 2003, p. 290–291). However the identity of physical objects is not among the evident properties of the theory. It is neither a measurable physical property nor a property that the formalism assigns to objects. Identity does appear in the mathematical formalism, but only as a relation between mathematical expressions, not as a relation of physical objects. And so identity statements about the objects of a theory are not part of the initial interpretation of the theory. Therefore, while our initial interpretation does provide us with putative objects of the theory, like bosons and fermions, we still do not know what the correct objects of the interpretation are, and what are the proper identity relations describing them.

Saunders proposes that since the identity of the objects is not evident, and therefore the ontology of the theory is not clear, we should define identity by use of PII. This would determine the correct objects for the evident physical properties of our initial interpretation. As he writes:

> The proposal, rather, is that in a situation in which we *do not know* what physical objects there are, but only, in the first instance, predicates and terms, and connections between them, then we should tailor our ontology to fit; we should admit no more as entities than are required by the distinctions that can be made out by their means. (Saunders, 2003, p. 292)

PII here serves as a methodological principle that allows us to settle upon the correct ontology of a theory.

2.5 The Principle of Identity of Indiscernibles and Quantum Ontology

While PII is not necessary, it still might be an important principle related to our physical theories, with direct implications for the ontology of quantum theory. As we pointed out above, some quantum particles do not obey PII. If we accept PII, this leads to one of two conclusions: *i*) either these particles are not objects of the theory or *ii*) quantum mechanics is incomplete and there are further properties and relations not included in the theory that allow us to discern between the objects.

Consider the first alternative—a change in ontology. Fundamental bosons seem to be the only particles that violate PII. We do not need to consider these particles as objects of the theory. As we mentioned above, it is natural to take them as discrete excitations of modes of the field.

Now consider the second alternative—a change of the theory. Some proposed solutions to the measurement problem—such as GRW, the modal interpretation, and de Broglie-Bohm—take quantum theory to be incomplete. Some of these solutions posit additional properties. In the modal interpretation, these additional properties are provided by the value state (at least in the case of measurements). In the de Broglie-Bohm interpretation, these additional properties are provided by the spatiotemporal trajectories of particles.[11] We might think that the failure of PII in quantum theory gives us a reason to accept one of these solutions to the measurement problem, in the hope that the additional properties that they provide will allow quantum particles to satisfy PII. However the additional properties of these interpretations might not always ensure that bosons satisfy PII. In the modal interpretation, the value state does not always provide a unique sets of properties for bosons (van Fraassen, 1991, p. 419). And in the de Broglie-Bohm interpretation, bosons might share spatiotemporal trajectories unless they are always impenetrable. So while PII might indicate that quantum theory is incomplete, it is not clear how to amend the theory. But what is important from our point of view, is that PII, as a physical principle, is leading us to weighty conclusions. It sanctions either a change in quantum ontology or a change in quantum theory.

2.6 Bose-Einstein Condensates

At this point PII seems to be a well motivated and philosophically useful principle. It does not succumb to traditional counterexamples and it leads to interesting conclusions about quantum theory. But all is not well. To see that all is not well, we will look at the specific example of confined atomic Bose-Einstein condensates (BECs). In this section we will briefly outline the theory of confined atomic BECs, and in the next section we will consider the philosophical implications.

There are a few reasons we want to consider confined atomic BECs. First, they provide a nice realistic example with a straightforward formal description. Second, the condensates consist of atoms. We expect atoms to be objects if any particles in quantum theory are objects. Third, even though

[11] Here we are assuming a particle ontology instead of a field ontology. For a field ontology, the additional properties of de Broglie-Bohm are provided by the field configuration, not the spatiotemporal trajectories of particles. But since our arguments are concerned with the ontological status of particles in non-relativistic quantum mechanics, we will limit our attention to the particle version of the de Broglie-Bohm theory.

atoms can be bosons, they consist of fermions. As we have already mentioned above, this internal structure might let us discern between otherwise indiscernible atoms.

There are several popularized accounts of atomic BECs.[12] They generally state that when atoms are cooled to a few microkelvins or less, the momentum of the atoms approaches zero. As the momentum approaches zero, the position of the atoms expands and they become wave-like because of the Heisenberg uncertainty principle. When the waves of each atom start to overlap they form a BEC, where all the atomic waves meld into one coherent wave of a super-atom. Since the BEC seems to consist of nothing more than one big matter wave, it is hard to maintain that the original constituent atoms are numerically distinct objects. PII would therefore not apply. However it would be dangerous to draw any philosophical conclusions from this metaphorical description. We must look at the theory and the formalism.

We start by considering the ground-state of a non-interacting BEC at zero temperature.[13] A simple harmonic potential is an accurate approximation of the confining potential of most experiments. A non-interacting BEC confined by a simple harmonic potential is just a simple harmonic oscillator in three dimensions. The single-particle ground-state, $\psi_0(\mathbf{r})$, for this system is:

$$\psi_0(\mathbf{r}) = (\frac{m}{\pi\hbar})^{3/4}(\omega_x\omega_y\omega_z)^{1/4}\exp[-\frac{m}{2\hbar}(\omega_x x^2 + \omega_y y^2 + \omega_z z^2)] \tag{2.2}$$

The many-body state-function of the ground-state is:

$$\psi(\mathbf{r}_1, \mathbf{r}_2, \ldots, \mathbf{r}_N) = \prod_{i=1}^{N} \psi_0(\mathbf{r}_i) \tag{2.3}$$

where N is the number of particles in the gas. Neglecting for the time being their fermionic constituents, the bosons described by this state-function are clearly indiscernible. Not only is each boson described by the same single-particle pure state, but there are also no symmetric irreflexive relations that allow us to weakly discern between them.

[12]For example see Cornell and Wieman (1998) and Ketterle (2002).

[13]The theoretical description of BECs presented below follows Dalfovo et al. (1999).

However this non-interacting description provides a poor approximation of experimental BECs. In order to accurately describe the density, the critical temperature, and collective excitations of BECs, we must consider atomic interactions. A convenient description of interacting BECs with a large number of atoms is given by mean-field theory. In mean-field theory we solve for $\Phi(\mathbf{r}, t)$ instead of $\Psi(\mathbf{r}_1, \mathbf{r}_2, \ldots, \mathbf{r}_N, t)$, where $|\Phi(\mathbf{r}, t)|^2 = n(\mathbf{r}, t)$. $n(\mathbf{r}, t)$ is the density of the condensate. $n(\mathbf{r}, t)d^3r$ gives the number of atoms within a differential volume.

To describe the atomic interactions we introduce a potential V_{int} that is a function of the density of the gas. It describes the interaction that one atom has with the rest of the gas. When the atomic gas is sufficiently cold and dilute, the only interactions we need to consider are elastic binary collisions with no exchange of angular momentum. The potential for this s-wave scattering, as it is called, is:

$$V_{\text{int}}(\mathbf{r}, t) = \frac{4\pi\hbar^2 a}{m}|\Phi(\mathbf{r}, t)|^2 \tag{2.4}$$

where a is the scattering length constant. At higher temperatures or greater densities we also have to consider inelastic and three particle collisions.

The governing dynamical equation for s-wave scattering within mean-field theory is the Gross-Pitaevkii (GP) equation:

$$i\hbar\frac{\partial}{\partial t}\Phi(\mathbf{r}, t) = (-\frac{\hbar^2\nabla^2}{2m} + V_{\text{ext}}(\mathbf{r}) + \frac{4\pi\hbar^2 a}{m}|\Phi(\mathbf{r}, t)|^2)\Phi(\mathbf{r}, t) \tag{2.5}$$

This describes the BEC at zero temperature. To arrive at a time independent equation we plug the following solution into the GP equation:

$$\Phi(\mathbf{r}, t) = \phi(\mathbf{r}) \exp(-i\mu t/\hbar) \tag{2.6}$$

where μ is the chemical potential and $\phi(\mathbf{r})$ is real and normalized so that $\int \phi^2(\mathbf{r})d\mathbf{r} = N$. We then have the following:

$$(-\frac{\hbar^2\nabla^2}{2m} + V_{\text{ext}}(\mathbf{r}) + \frac{4\pi\hbar^2 a}{m}\phi^2(\mathbf{r}))\phi(\mathbf{r}) = \mu\phi(\mathbf{r}) \tag{2.7}$$

From this equation we can numerically calculate the density of the BEC at zero temperature.

In equation (2.4) for V_{int}, we see that the scattering length a mostly governs the atomic interactions. The value of a depends on the internal fermionic structure of the kind of atoms used in the BEC. For example, in a BEC that consists of cesium, the Coulomb interactions between electrons and nuclei determine the scattering length, which is a function of the total electron spin of the atom (Tiesinga et al., 1995). The atomic interactions described by the scattering length a describe symmetric irreflexive relations that hold between the interacting atoms. Each atom interacts with every other atom of the BEC, but not with itself. These irreflexive relations allow us to weakly discern between the atoms, even when the BEC is in its ground-state.

It is the interaction potential within the dynamical equation governing the system, and not the state-function, that expresses these symmetric irreflexive relations. In our consideration of fermionic systems, we saw that the kinematical requirement of symmetrization provided symmetric irreflexive relations that allowed us to weakly discern between fermions. We now see that the dynamics can also provide symmetric irreflexive relations. There is nothing special about BECs. The same goes for other sorts of interactions between other sorts of particles. And since all particles of the actual world interact with each other in an irreflexive way, all actual particles (when taken as objects) are at least weakly discernible.[14] Therefore, if nothing else, PII is contingently true.

Quantum theory is not the only place where we find interactions, nor the only place where we interpret these interactions as symmetric irreflexive relations. For example in a Newtonian spacetime theory, the gravitational interaction between two massive particles also expresses a symmetric irreflexive relation, a relation that weakly discerns between the two particles. The same is true of electromagnetic interactions between particles in classical electrostatics.

Of course we might try to understand interactions within a BEC in an-

[14]This includes photons; for even photons interact in quantum electrodynamics.

other way. If we take the mean-field theory seriously, we can regard the BEC as a whole with no constituent particles. The nonlinear term in the GP equation would not describe interactions between particles, but would be just a dynamical feature of the field, an interaction of the field with itself. But such an interpretation would ignore what brought us to the GP equation in the first place and our interpretation of the scattering length. The reason the non-linear term takes the form it does in the GP equation is because we understand it as an interaction between atoms.

We should not take the mean-field theory seriously. It is simply introduced for calculational ease. Talk of fields here is not literal. The ontology is still one of atomic particles. The cesium atoms are particles, not discrete excitation of some cesium field. And it is as particles that we can understand and calculate the interactions between them expressed within the mean-field theory.

2.7 A Displeasing Discontinuity

Interacting atoms in a BEC obey PII. The interactions provide symmetric irreflexive relations that weakly discern between the bosons. What about a non-interacting BEC? In a non-interacting BEC there are no similar symmetric irreflexive relations. The bosons seem to be indiscernible. Does this mean that non-interacting BECs provide us with a counterexample to PII?

We have already looked at bosonic counterexamples above (see page 17). We claimed that only fundamental bosons violate PII, and it was best not to view these as objects. Other bosons, such as atoms, can be discerned by use of their fermionic constituents.

It is true that atoms actually have fermionic constituents. These constituents are responsible for the interactions in a BEC. But a description of a non-interacting BEC abstracts away from any consideration of constituent fermions. Such a description presents us with a model that is silent as to the composite nature of the particles involved. Therefore, when we take the non-interacting BEC simply as a self-contained model (which may or may not describe actual phenomena), the bosonic particles in the BECs violate

PII.

So do we have a counterexample to PII after all? Well we can maintain that bosons in an interacting BEC are objects, but deny that bosons in a non-interacting BEC are objects because they violate PII. According to this position, the atoms in actual BECs are objects even though the atoms described by non-interacting models are not. But if we take this position, we arrive at a displeasing discontinuity.

The non-interacting BEC is just a limiting case of the interacting BEC. We continuously approach the non-interacting case as the scattering length decreases or the gas becomes more dilute. If we maintain that bosons in the interacting case are objects, but those in the non-interacting case are not, then as we approach the non-interacting case, the interaction potential's mere existence, no matter how weak it is, allows us to consider the bosons as objects, but as soon as we reach the non-interacting case the bosons are no longer objects. Along the continuous approach to the non-interacting case, we have a discontinuous change in ontology.

This certainly is a displeasing discontinuity. The limiting case is not so different as to demand a different ontology, yet a different ontology is what we have. And it is unclear how the ontology can even change. If two models are connected in this way, we expect both to have the same ontology, to describe the same sort of objects. This displeasing discontinuity applies not only to BECs, but to any system where interactions provide the only discerning relations.

We can eliminate the discontinuity, while maintaining PII, by denying that all bosons, even the interacting variety, are objects. Here we turn the limiting case around and claim that interacting bosons are not objects because non-interacting bosons are not objects. But this would force us to reject a very simple and straightforward interpretation of interacting BECs. It would also force us to deny that the atoms that compose these BECs are objects.

And so the problem is this. The non-interacting BEC provides us with a model in non-relativistic quantum mechanics where particles violate PII. It seems best to regard the particles in a non-interacting BEC as objects because the non-interacting model is a limiting case of the interacting model where the particles are objects. And even though a non-interacting BEC might not provide an accurate description of actual phenomena, it still provides a counterexample to PII.

2.8 In Defense of Indiscernible Objects

In light of this new counterexample to PII, we have reason to doubt both the principle and the ontological arguments we have outlined that follow from the principle. In order to judge the merit of the PII, let us revisit the main motivations behind the principle: the epistemological, metaphysical, and methodological motivations.

2.8.1 Epistemology

Consider again the epistemological motivation for PII. As we pointed out above, we do not appear to have any epistemological access to indiscernible objects, and it therefore does not seem that they add to the empirical content of a theory. But this is just not the case. While it is true that we never directly observe two indiscernible objects, their number need not be arbitrary. Consider again our example of a non-interacting BEC. The total number of indiscernible particles affects the empirical properties of the BEC, such as its mass and density. So even though the particles are indiscernible, they still contribute to the empirical content of a theory, and are not completely unobservable.

Of course some indiscernible objects will not contribute to the empirical content of a theory. It does not take much imagination to dream up alternative theories that are inferior because of their inclusion of indiscernible objects. However our simple interpretation of the quantum case, which takes every particle as an object, shows that indiscernible objects do not always lead to a bloated ontology. They can contribute to the empirical content of our best theory. Quantum mechanics shows us that we cannot appeal simply to simplicity in order to justify PII.

2.8.2 Metaphysics

Our metaphysical motivation for PII holds that PII provides the only suitable principle of individuation for physical objects. If we reject PII and accept indiscernible objects, we must resort to other dubious principles of individuation that invoke haecceities and the like.

But luckily things are not as bad as that. A third option is open to us and always has been: take the identity of an object as primitive. What is it that makes the object the object it is and no other object? Simply that it is identical to itself and numerically distinct from other objects. No further analysis or principle of individuation is necessary. There is no unique individuator for each object, nor are there further facts that constitute the identity of each object. Two objects can differ only numerically, and it is solely this numerical difference that is responsible for their individuality.

If we grant that indiscernible quantum object have primitive identity, then we should grant that all quantum objects have primitive identity, and possibly all physical objects, for we expect the nature of individuality to be the same for all quantum, and possibly all physical, objects. So even though the identity of discernible quantum objects supervenes on some set of facts, we deny that these facts constitute the identity of these objects. This supervenience base might provide us with a criterion of identity for those objects, allowing us to establish the validity of certain identity statements; but such a criterion will not serve as a principle of individuation.

Can identity serve the role we have set for it? French and Krause claim that "since every being has it, identity is useless for the purposes of individuation" (French and Krause, 2003, p. 110). Obviously the self-identity of an object does not provide a unique individuator since every object is identical to itself. But when we take the identity of an object as primitive, we deny that there are unique individuators for each object. Objects just are uniquely the objects they are. It is a mistake (not necessarily one that French and Krause make) to think that a lack of an individuator indicates a lack of an individual.

There is nothing about the concept of objects, which we have outlined in the introductory chapter, that requires unique individuators for each object. If an object is simply that which we describe by use of standard quantifiers and predicates, then all that is required is that each object is numerically identical to itself and distinct from other objects; and this requires no analysis, for identity is our best understood relation. If we can take anything in philosophy as primitive, it is the identity relation.

When we do so, the identity of an object is both qualitative and intrinsic. Primitive identity, in contrast with haecceities and the like, does not introduce anything non-qualitative into our account of individuality. This is because the identity relation "=", which establishes individuality, is qualitative; as French and Krause point out, it applies to every object. And as Rosenkrantz (1979) correctly notes, the defining feature of non-qualitative properties is that their instantiation implies the existence of particular objects. For example, the instantiation of "identical to Bertrand Russell" implies the existence of a particular object—Bertrand Russell. There is no such implication for a qualitative property, like "possessing mass m". It matters little for our purposes whether we elaborate this distinction in a modal way, like Rosenkrantz (1979), or in a syntactic way, like Adams (1979). The important point for us is that the identity relation "=" taken by itself is qualitative. While its instantiation implies the existence of an object, as the instantiation of any property does, unlike "identical to Bertrand Russell", "=" does not imply the existence of any particular objects. It is a mistake to think that PII is necessary in order to ensure a qualitative world. Indiscernible objects with primitive identity are just as qualitative as their discernible cousins.

And unlike PII, primitive identity ensures that the identity of an object is intrinsic. As we pointed out above on page 19, if we take PII as our principle of individuation, then the identity of objects can depend on extrinsic and external relations. We have no such problem when we take the identity of objects as primitive. The intrinsic nature of the identity of each object is just a trivial consequence of taking the identity of each object as a primitive feature of each object.

Just to be clear, we should also note that the primitive identity of objects does not imply either primitive cross-world identity or primitive cross-temporal identity. While we might hold that these are really just instances of numerical identity, and are therefore primitive as well, we can also claim that they are analyzable in terms of other relations. For example cross-world identity might be based upon a qualitative counterpart relation and cross-temporal identity might be based on a qualitative genidentity relation. Neither of these relations are inconsistent with primitive numerical identity.

No doubt all will not be satisfied with primitive identity. In order to allay some fears, let us point out that primitive identity does not make the nature of quantum objects mysterious. Even without a principle of individuation, we can still elucidate the nature of a quantum particle ontology by delineating kinds and pointing out supervenience relations. Quantum particle kinds are still defined by state-independent intrinsic properties, like mass and charge. And we can still claim, if we so choose, that quantum particles, even though they are objects, supervene on modes of the quantum field. (But once again, even if we admit that facts about quantum particles supervene upon facts about the modes of the field, facts about the field do not constitute the identity of quantum particles. The identity of quantum particles is still primitive, even if quantum particles are not. Simply the fact that a quantum particle supervenes upon a mode of a field will not establish why this particle is the particle it is and no other particle, for it will say nothing about how it differs from another particle that supervenes upon the same mode.)

But this still may not convince, for many hold that an object must be able to possess a proper name.[15] This does not mean that every object must actually have a name (there might not be enough names to go around), only that each object is capable of bearing one. If this view is true, then it is natural to claim that it is an object's unique individuator that allows it to possess a proper name. For how else can the reference of a proper name be guaranteed in all circumstances?

This cannot be our view, nor should it be. Not every object can be referred to uniquely by use of a proper name. A failure of reference does not necessarily indicate that the intended referent is not after all an object.

[15] For example, this view appears to be endorsed in (French and Redhead, 1988) and (Redhead and Teller, 1992).

It might just indicate the limited referential capacity of proper names. For example, if proper names refer through unique causal relations, we have no reason to believe that every object is capable of entering into the necessary causal relations. Indiscernible objects certainly cannot, for no causal relation will be unique. But this need not indicate that an indiscernible object is any less of an object. Even if we cannot refer to it by use of a proper name, it still can be something we can quantify over and predicate properties to. It still can be something that is identical to itself and numerically distinct from other objects.

But particle labels do appear in our quantum formalism. If we regard them as proper names, how do they refer to potentially indiscernible objects if not through haecceities and the like? It is better to avoid the question altogether and say that the particle labels in a Hilbert space are not proper names. Consider for example the two-particle Hilbert space state $\Phi(\mathbf{r}_1, \mathbf{r}_2)$. Instead of interpreting the indices 1 and 2 as names of particles, we should interpret them as two existentially bound variables that cannot take the same value. Instead of representing proper names, like a and b, they represent bound variables, like x and y.[16] If we view particle labels in this way, then there is no need to resort to haecceities and the like. We can, once again, take the identity of the particles as primitive.

To sum up, indiscernible objects do not lead to unsavory principles of individuation. Their identity, along with the identity of every other quantum particle, can be taken as primitive. There is therefore no metaphysical reason to hold that PII is true for physical objects.

2.8.3 Methodology

Let us finally consider the methodological motivation for PII. We have already pointed out a reasonable interpretation of quantum mechanics that includes indiscernible objects. We might think that the mere existence of such an interpretation is enough for us to dismiss PII as a methodological principle. After all how can we argue that theories without indiscernible ob-

[16] Cheung (forthcoming) expresses a similar view about the nature of particle labels.

jects are better or more successful if we have an example of a successful theory with indiscernible objects? But along with providing a counterexample, we can also directly address the argument for regarding PII as a methodological principle. As before, the argument we have in mind is the one given by Saunders.

As we mentioned above (see page 22), Saunders regards PII as a methodological principle intimately connected with the way we interpret the experiments and mathematical formalism of a theory. We start with an initial interpretation of the experiments and mathematical formalism of the theory. This gives us certain evident physical properties and the putative objects to which those properties apply. We then use PII in order to establish the proper ontology of the theory. The actual objects are those that are discernible according to the properties of the theory. Here PII is a methodological principle, but it does not serve merely as a constraint on an acceptable interpretation of a theory. It also provides a method of interpretation that allows us to determine the ontology of a theory from its experiments and mathematical formalism.

We can question whether this is how we should interpret physical theories. But instead we will question whether we can even apply this method to quantum mechanics. If we cannot, then there is reason to doubt whether we can or should apply it to other theories.

We can read off some properties from the experiments and formalism of a theory without considering the identity of the objects to which those properties and relations apply. Charge is a property. We know what charge is; we know how to measure it; we know how it is governed by the formalism. We know all of this without knowing what actually possesses charge. Do particles or fields possess charge?

But what about relations like "interacts with"? We cannot read these relations off from the formalism independently of specifying the objects to which they apply. It is because we postulate that there are multiple bosons in a BEC that we consider the interaction potential as expressing an irreflexive relation between bosons. As we mentioned above, if we took the ontology of the system to consist of a quantum field, we would interpret the interaction potential as expressing simply an intrinsic property of the field.

Many properties and relations are only evident once we have decided upon an ontology. We might grant that we can translate expectation values and conditional probabilities into predicates without considering the the objects to which these properties apply. However, we cannot translate the other relations that we have read off from the kinematics and dynamics of the system, such as "opposite spin from" and "interacts with", without first knowing the ontology and objects of the system. We therefore cannot use these properties and relations, along with PII, to determine the correct ontology. Nor can we define the identity of objects in terms of other physical properties by use of PII; for those properties are not evident until we decide on a specific ontology, an ontology that will already include statements of identity.

Our objection is different from Russell and Armstrong's objection to PII given above (see page 19). They argue that we must first be able to discern between objects before we can assign extrinsic properties to them. Unlike them, we can grant, if we so choose, that relations individuate. Our point is simply that we must specify the type (not number) of objects involved before we can determine the properties assigned to them by the mathematical formalism of the theory. Like van Fraassen we deny "that a theory wears its content on its sleeve, written unambiguously into the shape of its formalism" (van Fraassen, 1991, p. 435).

There do not seem to be any good arguments for PII. With a clear counterexample to it, the best thing to do is abandon the principle. While PII maybe contingently true, it is not a principle related to our physical theories. While quantum particles can be indiscernible, the failure of PII does not establish that quantum particles are not objects. The principle does not warrant the claim that there is a fundamental difference between quantum and classical ontology. It still seems best to take particles as objects in both theories. Indiscernible objects are not inherently bad. They can be of good character and do not deserve any prejudice we might harbor.

Chapter 3

Permutations and Ontology

3.1 Indistinguishability and Permutation Symmetry

We have already mentioned that there are different notions of indistinguishability. One way to understand indistinguishability is in terms of indiscernibility. In the last chapter we explored to what extent quantum particles were indistinguishable in this sense, and what consequences this had for their individuality. We found that quantum particles of the same kind can be indiscernible, but concluded that this did not undermine their individuality.

But there is another relevant notion of indistinguishability we must consider. This notion is connected to permutation symmetry. The idea here is that particles of the same kind are indistinguishable if a permutation of them never leads to a distinct physical situation. That is they are indistinguishable if they possess permutation symmetry.

These two notions of indistinguishability can be related, but they are distinct. Just because two particles are indiscernible does not imply that they possess permutation symmetry. A permutation of them still might lead to a distinct physical state. And as we will see, the permutation symmetry of two particles does not imply that they are indiscernible.

It is widely held that quantum particles are indistinguishable in this permutation symmetric sense, while classical particles are not. That is quantum particles and quantum physics possess a permutation symmetry that classical particles and classical physics do not. This difference in permutation symmetry is often taken to ground a difference in ontology. Because quantum particles, unlike classical particles, are indistinguishable in this permutation symmetric sense, quantum particles, unlike classical particles, are not objects. So here we have another ontological argument that endeavors to show that quantum particles are not objects because they are indistinguishable. But instead of citing indiscernibility, this argument cites permutation symmetry.

This ontological argument from permutation symmetry is the subject of this chapter. Here we will sketch the historical origins of this view, and present its structure in greater detail. We will conclude that while quantum particles are permutation symmetric, this does not undermine their individuality, identity, or objecthood. We will establish this conclusion by demonstrating that, contrary to popular belief, classical particles and classical physics are just as permutation symmetric as quantum particles and quantum physics. And since classical particles are clearly objects, permutation symmetry cannot lead us to conclude that quantum particles are not.

3.2 Historical Roots

The ontological argument from permutation symmetry has its origins in the first derivations of quantum statistics. As is well known, the historical development of quantum theory starts with Planck's statistical mechanical treatment of blackbody radiation.[1] Planck initially studied blackbody radiation in the hopes that it might provide some reconciliation between the irreversibility of physical processes, which he believed followed from the second law of thermodynamics, and classical dynamics. With this goal in mind he developed, by the use of classical electrodynamics, a theory for electromagnetic radiation that was analogous to Boltzmann's H-theorem. From this electrodynamic H-theorem, he explained the irreversible approach to equilibrium in a similar way as Boltzmann initially did with his H-theorem. Of

[1]For details of the historical development of Planck's views, see Jammer (1966, sec. 1.2) and Kuhn (1978).

course Planck's explanation of irreversibility eventually succumbed to the same sort of objections as Boltzmann's explanation, such as those raised by Poincaré reoccurrence and Loschmidt's paradox. But from his electrodynamic H-theorem he derived the entropy for blackbody radiation, which was at a maximum for Wien's radiation law. This was at the time a great achievement. Wien's radiation law was empirically well supported by the measurements that were available at short wavelengths, such as the Paschen series. The law, however, did not have a very good theoretical basis before Planck's work on the subject. Planck was therefore celebrated for providing a firmer theoretical foundation for the law.

But when further experimental evidence at longer wavelengths demonstrated the inadequacy of Wien's radiation law, Planck sought another expression for the entropy associated with blackbody radiation. He turned to the combinatorial approach that Boltzmann employed in (Boltzmann, 1877).

In this paper, Boltzmann starts by considering a collection of N molecules with a discrete energy spectrum.[2] The possible energies for each molecule are: $0, \epsilon, 2\epsilon, 3\epsilon, \ldots P\epsilon$, where $P\epsilon$ is the total energy of the system. Boltzmann considers the number ways of distributing energy elements over the molecules, so that there are ω_j molecules with energy $j\epsilon$. He calls each of these ways a complexion. A complexion assigns n_i energy elements to the ith molecule. A permutation of molecules with different energies leads to a distinct complexion. By combinatorics, the number of complexions Z associated with a distribution of energy elements is:

$$Z = \frac{N!}{\prod_j \omega_j!} \tag{3.1}$$

Boltzmann then claims that the most probable distribution is the one where Z is at a maximum (subject to constraints), which he is able to calculate using variational calculus for the case where $P \gg N$.

Boltzmann then moves on to consider a continuous energy spectrum for the molecules. By analogous calculations, he arrives at the Maxwell-

[2]The following summary of Boltzmann's work closely follows that given by Kuhn (1978, p. 46–54).

Boltzmann distribution for particles in two, not three, dimensions. To arrive at the correct distribution in three dimensions, Boltzmann considers the three-dimensional velocity space of molecules, instead of the energy continuum. It is this last derivation that was eventually popularized by Ehrenfest and Ehrenfest (1959, p. 26–31).[3]

In (Planck, 1900), in which he first outlines the derivation of his radiation law, Planck models blackbody radiation by considering a group of nondescript resonators. He takes the energy spectrum for each resonator to be a discrete quantity, so that the energy ϵ_ν of a resonator vibrating at frequency ν is $\epsilon_\nu = nh\nu$, where n is some integer greater than or equal to zero, and h is of course Planck's constant. Planck then considers the number of ways W—which, following Boltzmann, he calls complexions—of distributing P discrete energy elements of magnitude $h\nu$ among N oscillators of frequency ν. The number of ways is given by the following combinatorial equation:

$$W = \frac{(N + P - 1)!}{(N - 1)!P!} \tag{3.2}$$

Following Boltzmann, Planck takes it that "the entropy of a system of resonators with given energy is proportional to the logarithm of the total number of possible complexions for the given energy." As with Boltzmann, the equilibrium distribution, which is the most probable state, is the one where the entropy and W are at a maximum. Planck calculates this equilibrium distribution by the same sort of variational techniques as Boltzmann, and

[3] According to Kuhn, this combinatorial approach was not central to Boltzmann's thought. He developed it simply as an attempt to explain Loschmidt's paradox. This paradox points out that for every dynamical evolution toward equilibrium, there is a dynamical evolution away from equilibrium. By the combinatorial approach, equilibrium is the most probable distribution for the system. Boltzmann believed that this switch to a probabilistic view provided an explanation for the approach to equilibrium that was consistent with Loschmidt's paradox (Kuhn, 1978, p. 46–54). But this was only tangential to most of Boltzmann's work, which concentrated on a mechanical treatment of his H-theorem. Kuhn also points out that Boltzmann's work on this combinatorial approach was not widely known until Planck took up the approach to derive the entropy of blackbody radiation (Kuhn, 1978, p. 70–71).

from this arrives at his radiation law. This law states:

$$u_\nu = \frac{8\pi\nu^2}{c^3}\frac{h\nu}{e^{h\nu/kT}-1} \tag{3.3}$$

where u_ν is the energy density for radiation with frequency ν, c the speed of light, k Boltzmann's constant, and T the temperature.

Planck is here following Boltzmann very closely, especially his initial treatment of a discrete energy spectrum. The parallels are obvious. While Boltzmann considers the number of ways of distributing energy elements over molecules, Planck consider the number of ways of distributing energy elements over resonators of frequency ν.

Equation (3.2) is of a different form than equation (3.1), but this is here of little importance. Lorentz (1910) presents a derivation of Planck's radiation law that follows Boltzmann's calculation more closely, and makes use of a combinatorial equation of the same form as equation (3.1).[4] And it is this derivation that Planck presents in the second edition of his *Lectures on the Theory of Heat Radiation* (Planck, 1913). The differences between the two derivations is only due to the type of distributions under consideration. In Lorentz's derivation, he calculates the number of ways of distributing energy elements over resonators of a given frequency so that there are ω_j resonators with energy $jh\nu$. In Planck's first derivation, he considers resonators at different frequencies, and calculates the number of ways of distributing energy elements over resonators so that the total energy for resonators at a given frequency is E_ν. The maximum value of both of these leads to Planck's radiation law. And if we sum equation (3.1) over all of the allowable values of ω_j, we arrive at equation (3.2).

But there is an essential difference between the derivations of Planck and Boltzmann. For Boltzmann, discrete energy elements are a mathematical fiction. Boltzmann introduces a discrete energy spectrum simply to motivate and elucidate his approach. When he moves on to consider the more physical example of an ideal gas, he stipulates a continuous energy spectrum divided

[4]For an analysis of Lorentz's derivation, see Bach (1990, p. 23–24) and Kuhn (1978, p. 102–105).

into cells, and calculates the limit where the size of those cells approach zero. In Planck's derivation, the energy elements are not simply mathematical fictions. The energy spectrum of the resonators is actually discrete. And as we will see, while we can replicate Planck's derivation by considering a continuous energy spectrum divided into cells, these cells must have a fixed size. We cannot take the same continuous limit that Boltzmann does. But that will come later. Right now let us emphasize once again the important parallel between Boltzmann's combinatorial approach and Planck's derivation.

At this stage in the development of quantum theory there is no reason to think that quantum particles are indistinguishable in any sort of permutation symmetric sense. In both Boltzmann and Planck's derivations, the energy elements are indistinguishable according to the combinatorial equations. A permutation of them does not lead to a distinct complexion. But the energy elements are not particles or entities of any sort in these derivations. They are simply quantities of energy possessed by the actual particles under consideration—the molecules and resonators. Both of these are distinguishable in a permutation symmetric sense; for in both Boltzmann and Planck's derivations, a permutation of two molecules or resonators leads to distinct complexion if they possess different single-particle energies.

The important thing to note in Planck's initial derivation, is that the difference in distinguishability reflects a difference in ontology. The true objects of the theory, the resonators, are distinguishable, while their properties, the energy elements, are indistinguishable. It seems that this connection between distinguishability and ontology was first pointed out by Natanson (1911) and Ehrenfest (1911), and later Ehrenfest and Onnes (1914).

In these papers, the authors argue that we cannot identify Einstein's light quanta, which Einstein first discussed in (Einstein, 1905), with Planck's energy quanta. To illustrate this view and its relation to the question of indistinguishability, let us consider Ehrenfest and Onnes presentation of this argument. They write:[5]

> As a matter of fact Planck's energy-elements were in that case

[5]The emphases are theirs.

almost entirely identified with Einstein's light-quanta and accordingly it was said, that the difference between Planck and Einstein consists herein that the latter assumes the existence of mutually independent energy-quanta also in empty space, the former only in the interior of matter, in the resonators. The confusion which underlies this view has been more than once pointed out. Einstein really considers P similar quanta, existing *independently of each other*. He discusses for instance the case, that they distribute themselves irreversibly from a space of $N_1\,\mathrm{cm}^3$ over a larger space of $N_2\,\mathrm{cm}^3$ and he finds using Boltzmann's entropy-formula: $S = k \log W$, that this produces a gain of entropy:

$$S - S_0 = k \log(\frac{N_2}{N_1})^P \quad (\alpha)$$

i.e. the same increase as in the analogous irreversible distribution of P similar independent gas-molecules, for the number of ways in which P quanta may be distributed first over N_1, then over N_2 cells in space, are to each other in the ratio

$$N_1P : N_2P \quad (\beta)$$

If with Planck the object were to distribute P mutually independent elements ϵ over N resonators, in passing from N_1 to N_2 resonators the numbers of possible distributions would in this case also increase in the ratio (β) and correspondingly the entropy according to equation (α). We know, however, that Planck obtains the totally different formula

$$\frac{(N_1 - 1 + P)!}{(N_1 - 1)!P!} : \frac{(N_2 - 1 + P)!}{(N_2 - 1)!P!} \quad (\gamma)$$

(which only coincides approximately with (β) for very large values of P) and a corresponding law of dependence of the entropy on N. This can be simply explained as follows: Planck does not

> deal with really mutually free quanta ϵ, the resolution of the multiples of $[E]$ into separate elements ϵ, which is essential in his method, and the introduction of these separate elements have to be taken "cum grano salis"; it is simply a formal device ... The real *object which is counted* remains the number of all the different distributions of N resonators over the energy-grades $0, \epsilon, 2, \ldots$ with a given total energy P. ...
>
> We may summarize the above as follows: Einstein's hypothesis leads necessarily to formula (α) for the entropy and thus necessarily to Wien's radiation formula, not Planck's. Planck's *formal device* (distribution of P energy-elements ϵ over N resonators) *cannot be interpreted in the sense of* Einstein's *light quanta*. (Ehrenfest and Onnes, 1914, p. 872–873)

So according to Ehrenfest and Onnes, we cannot identify Einstein's light quanta with Planck's energy quanta, for the former, unlike the latter, exist independently of each other. Einstein's light quanta are actual physical particles, while Planck's energy elements are mathematical fictions used to describe the properties of the true particles of the theory, the resonators. This difference leads to a difference in combinatorial equations. Because light quanta exist independently of each other, they are distinguishable. A permutation of them leads to a distinct complexion. Because energy quanta do not exist independently of each other, they are indistinguishable. A permutation of them does not lead to a distinct complexion. If we were to hold that energy quanta are legitimate particles possessing independent existence, then we would have to substitute a combinatorial equation that counted a permutation of them as a distinct complexion. This would lead to Wien's radiation law instead of Planck's radiation law.

If we interpret Ehrenfest and Onnes's talk about the independent existence of objects to mean that objects possess their own identity, then they are here expressing the view that distinguishability is connected to particle identity. Einstein's light quanta are distinguishable because they possess identity, while Planck's energy quanta are indistinguishable because they lack identity. This difference in ontology prevents us from identifying the two.

Now of course after the work of Bose (1924) and Einstein (1924), it was no longer tenable to hold that there was a difference between light quanta and energy quanta in the theory of blackbody radiation; for Bose and Einstein derived Planck's law by explicitly considering a gas of light quanta, without any consideration of resonators. These papers marked a significant change in the understanding of quantum particles; for the combinatorial equations involved in their derivations are the same as those in Planck's, but they have a different interpretation. Instead of considering the number of ways of parsing out energy elements among resonators consistent with some distribution, they consider the number of ways of distributing light quanta over phase space cells. There is a switch in the particles of the theory. The resonators, which were the previous particles, now become phase space cells, and the energy elements are now the particles. Since the combinatorial equations have not changed, the energy elements are still indistinguishable in a permutation symmetric sense, but now they are acknowledged as legitimate particles. They are no longer simply mathematical fictions associated with the properties of resonators.[6]

And of course it is not just photons (as we will now call light quanta) that are indistinguishable. The statistical derivations of Bose and Einstein apply to all bosons. And a similar result applies to fermions. As we will see in greater detail below, the combinatorial equations we employ in Fermi-Dirac statistics also do not count permutations of particles as distinct complexions. All quantum particles are therefore indistinguishable in this permutation symmetric sense.

[6]Because of the similarities between the derivations of quantum statistics and some of Boltzmann's derivations, Bach (1990) argues that Boltzmann was actually the first to derive Bose-Einstein statistics. It is, however, important to keep in mind the two important ways in which Boltzmann's derivations differ from the standard derivations of quantum statistics, both of which we have already mentioned above. First, for Boltzmann discrete energy elements are simply a mathematical fiction that allow him to apply combinatorial techniques. Unlike Planck, Bose, and Einstein, Boltzmann takes the limit where the size of these energy elements goes to zero. As we will see in greater detail below, this limit leads to classical instead quantum statistics. Second, unlike Bose and Einstein, the particles of Boltzmann's derivations are the molecules and not the energy elements. While the combinatorial equations are similar in both sorts of derivations, they have very different interpretations.

So now there is a difference between the distinguishability of classical particles, like the molecules in Boltzmann's derivation, and quantum particles, like the photons in Bose and Einstein's derivations. If we still hold that distinguishability follows from particle identity, and indistinguishability follows from a lack of identity, as Ehrenfest and Onnes appear to claim, then we must conclude that quantum particles lack identity, and further, and that this difference in identity leads to a difference in statistics, for it is the indistinguishability of quantum particles that leads to quantum expressions like Planck's radiation law instead of classical expressions like Wien's radiation law. We then have a clear connection between a lack of particle identity, indistinguishability, and quantum statistics.

This connection between the three appears to exists not only in the quasi-classical derivations of Bose and Einstein, but also in the full Hilbert space formalism. As we will see in detail below, in the Hilbert space formalism indistinguishability in terms of permutation symmetry appears to be connected to the symmetrization requirement, which in turn leads to quantum statistics. If we continue to maintain that indistinguishability follows from a lack of particle identity, then quantum statistics, indistinguishability, and a lack of particle identity are connected as before.

It is unclear when this view of particle identity and indistinguishability takes explicit hold in the minds of physicists. It is clear, however, that this view did eventually take hold in the minds of several eminent physicists. For example, Born writes:

> If . . . photons are treated as genuine particles, having an individuality of their own, Planck's law would not obtain. One has instead to assume that two states which differ only by the exchange of two photons are physically indistinguishable and have statistically to be counted only as one state. In other words, photons have no individuality. (Born, 1943, p.27)

Here Born takes quantum statistics (specifically Planck's law) to follow from indistinguishability, and indistinguishability to in turn follow from a lack particle individuality, which we have interpreted as lack of objecthood and identity.

Weyl presents us with another example of someone who holds this view. He asks us to consider two twin quantum particles Mike and Ike in single-particle states E_1 and E_2. As Weyl writes, "the possibility that one of the identical twins Mike and Ike in the quantum state E_1 and the other in the quantum state E_2 does not include two differentiable cases which are permuted on permuting Mike and Ike; it is impossible for either of these individuals to retain his identity so that one of them will always be able to say 'I'm Mike' and the other 'I'm Ike.' " (Weyl, 1950, p. 241).

But it is perhaps Schrödinger that gives the clearest and most detailed expression of this connection between individuality, permutation symmetry, and physical statistics. He ask us to consider the following example:

> Three schoolboys, Tom, Dick, and Harry, deserve a reward. The teacher has two rewards to distribute among them. Before doing so, he wishes to realize for himself how many different distributions are at all possible. This is the only question we investigate (we are not interested in his eventual decision). It is a statistical question: to count the number of different distributions. The point is that the answer depends on the nature of the rewards. Three different kinds of reward will illustrate the three kinds of statistics.
>
> (a) The two rewards are two memorial coins with portraits of Newton and Shakespeare respectively. The teacher may give Newton either to Tom or to Dick or to Harry, and Shakespeare either to Tom or to Dick or to Harry. Thus there are three times three, that is nine, different distributions (classical statistics).
>
> (b) The two rewards are two shilling-pieces (which, for our purposes we must regard as indivisible quantities). They can be given to two different boys, the third going without. In addition to these three possibilities there are three more: either Tom or Dick or Harry receives two shillings. Thus there

are six different distributions (Bose-Einstein Statistics).

(c) The two rewards are two vacancies in the football team that is to play for the school. In this case two boys can join the team, and one of the three is left out. Thus there are three different distributions (Fermi-Dirac statistics).

> Let me mention right away: the rewards represent the particles, two of the same kind in every case; the boys represent states the particles can assume. Thus, "Newton is given to Dick" means: the particle Newton takes on the state Dick.

Schrödinger goes on to claim that the different statistics follows from the fact that quantum and classical particles "are of different categories." As he puts it:

> Memorial coins [which represent classical particles] are individuals distinguished from one another. Shillings [which represent bosons], for all intents and purposes, are not, but they are still capable of being owned in the plural. It makes a difference whether you have one shilling, or two, or three. There is no point in two boys exchanging their shillings. (Schrödinger, 1998, p. 206)

Once again, we have the view that quantum statistics follows from the indistinguishability of quantum particles, which in turn follows from a lack of individuality or identity.

Schrödinger's analogy with money is very informative, for it clearly expresses how permutation symmetry is suppose to follow from a lack of individuality.[7] It is because quantum particles are not like particular memorial coins that they possess permutation symmetry. They are something else, something that lacks identity, something that is not properly an object, like discrete quantities of money. And it is because they possess a different nature from classical particles, that they are indistinguishable, that they possess permutation symmetry. This manifests itself, among other ways, in quantum statistics.

---

[7] Teller (1998, p. 118–120) makes use of a similar analogy, as I am sure many authors do.

3.3 Responses to the Received View

French and Rickles (2003) have referred to this connection between permutation symmetry, ontology, and statistics as the "received view". The view is not limited to deceased physicists. As any quick survey of introductory textbooks will show, it is widely held throughout the physics community. And it also has contemporary proponents within the philosophy of physics community. Teller (1983) and Post (2000) offer similar, albeit more sophisticated, versions of the views expressed by Born, Weyl, and Schrödinger. We will discuss some of these in greater detail below.

There are some important responses to this view in the literature that we should review before we move on to our own response. One comes from van Fraassen (1991, ch. 11). He correctly points out that quantum statistics does not necessarily follow from permutation symmetry, nor does it necessarily imply that particles are not objects; for it is possible to construct statistical theories of classical particles that lack permutation symmetry, but still exhibit quantum statistics. Van Fraassen gives two examples of such theories. One posits correlations between the classical particles and the other makes use of de Finetti's representation theorem. Both these theories differ from *standard* statistical mechanics. The first posits primitive correlations without any dynamical or kinematical explanation. The second does not assign an equal probability to every allowable state of the same energy, as standard statistical mechanics does. But both theories do demonstrate that an ontology of objects, like classical particles, that lacks permutation symmetry is consistent with quantum statistics.

It still, however, might be the case that standard statistical mechanics leads to quantum statistics because of the indistinguishability and lack of identity of the particles involved. It makes no difference to this view whether or not other statistical theories lead to the same statistics for distinguishable classical particles. So there are still some important views that we need to address.

Another important response to the received view is given by French and Redhead (1988) (also see French (1989a) and French and Rickles (2003)).

French and Redhead correctly argue that it is possible to hold that quantum particles are still objects. Permutation symmetry, in the form of symmetrization, does not necessarily imply that quantum particles lack identity. It is possible to claim that the permutation symmetry of quantum theory simply expresses a restriction on available states for particles that are objects. There is nothing incoherent about this.

But while French and Redhead hold this to be a logically coherent position, they make it clear that they are not arguing for this position. So we are still left with an important question. While permutation symmetry does not necessarily imply that quantum particles are not objects, is permutation symmetry still best explained as following from the fact that quantum particles are not objects?

So there are still some questions we need to address, for the relation between permutation symmetry, particle identity, and quantum statistics is still obscure. In the remainder of this chapter we will attempt to throw some light on the matter by pointing out the permutation symmetry that exists in classical physics. Because classical physics is permutation symmetric in much the same way as quantum physics, we cannot explain permutation symmetry by saying that it follows from a lack of particle identity; for permutation symmetric classical particles clearly possess identity and are objects. The permutation symmetry of classical physics also prevents us from claiming that the difference between classical and quantum statistics is due to permutation symmetry alone.

The permutation symmetry of classical physics will leave us with further questions. If not because of permutation symmetry, why is there a difference between classical and quantum statistics? And if we cannot explain permutation symmetry as following from a lack of particle identity, can we offer any explanation? These are questions we will address in the next two chapters. But right now let us turn to the permutation symmetry of classical physics.

3.4 Classical Dynamics

If we limit ourselves to dynamics, it is easy to show that classical physics is just as permutation symmetric as quantum physics. This follows from the fact that the classical Hamiltonian, like the quantum Hamiltonian, is permutation invariant.[8] Let us illustrate this point by considering the simple example of two particles of the same kind. Let the sole potential acting upon the particles be a function of the distance between the particles. Further let ξ be a point in μ-space that gives the state of a single particle at a given time. The classical Hamiltonian for this system is:

$$H(\xi_1, \xi_2) = \frac{\mathbf{p}_1^2}{2m} + \frac{\mathbf{p}_2^2}{2m} + V(|\mathbf{r}_1 - \mathbf{r}_2|) \tag{3.4}$$

This Hamiltonian is clearly invariant under particle permutation:

$$H(\xi_1, \xi_2) = H(\xi_2, \xi_1) \tag{3.5}$$

The Hamiltonian for more complicated systems consisting of particles of the same kind is also invariant under particle permutation if the interaction potential is invariant. And the interaction potential is invariant for systems solely composed of particles of the same kind; for the interaction potential supervenes upon the dynamical properties of the system, such as the charge and mass of the particles. Since the particles are of the same kind, a permutation of particles produces no change in the dynamical properties of the system and therefore no change in the interaction potential. We can state the same point in another way by saying that particle labels alone are not dynamical properties, and so a permutation of particle labels does not affect the interaction potential. True, if we are not dealing with particles of the same kind, a permutation of particle labels can affect the interaction potential, but this is not because particle labels are dynamical properties, but because they are associated with other dynamical properties (like mass) that

[8] Of course we are not the first to point out that classical dynamics is permutation symmetric. Messiah (1962, p. 582–583) is just one example of someone who has already done so.

differ between the particles. All of this is as true for classical Hamiltonians as it is for quantum Hamiltonians.[9]

Because the Hamiltonian is permutation invariant, for a system composed of particles of the same kind, the single-particle trajectories will not change upon the permutation of particles. The dynamical evolution of the system is therefore permutation symmetric. Which particle is traveling upon which trajectory does not affect how the system evolves.

We often describe the dynamical evolution of a system by use of a trajectory in phase space. Assuming that the particles are impenetrable, so that no two particles can occupy the same spatial point at the same time, a permutation of particles always leads to a distinct trajectory in phase space. But even though these trajectories are distinct for a system of particles of the same kind, they describe the same dynamical evolution of the system, that is they describe the same set of single-particle trajectories. They only differ in which particle they assign to which trajectory.

In fact, we can adequately describe the dynamical evolution of a system without mentioning at all which particle is moving along which trajectory. Instead of taking points of phase space as describing the states of the system at a given time, take particle distributions functions in μ-space. Following Huggett (1999a), let us call this a Z-space (pronounced zeta space) description of the states. This Z-space description will tell us how many particles are in each single-particle state, but not which particle is in which state. From the Z-space description of the system, we can determine the set of particle trajectories for the system, which will remain the same regardless of which particle is in which trajectory. The lesson to draw is that single-particle trajectories are not affected by our assignment of particle labels, nor are the dynamical equations governing trajectories. Classical dynamics is permutation symmetric.

[9]De Muynck (1975) also notes that particle labels do not serve any dynamically relevant role, but he holds out the possibility that we might discover that they label some hitherto unknown dynamical property.

3.5 Classical Statistical Mechanics

But the real debate about permutation symmetry and its connection to ontology does not center around dynamics, but rather statistical mechanics. For even if a permutation of particles does not lead to a distinct set of single-particle trajectories, it still might be the case that it leads to a distinct physical situation; for there still might be a difference about which particle is traveling along which trajectory. While this would not have any consequences for the dynamical evolution of the system, it would affect the statistical mechanical description of the system; for statistical mechanics assigns a distinct probability to each distinct physical situation. So in order to fully demonstrate that classical physics is permutation symmetric, we need to consider not only classical dynamics, but also classical statistical mechanics.

In the standard Gibbsian approach to classical statistical mechanics, we describe the statistical behavior of a system by considering an ensemble of duplicate systems that satisfy certain macroscopic requirements.[10] Each system in the ensemble represents a distinct possible evolution of the system we wish to describe. For a system with a fixed number of particles, we can represent the ensemble at a given time as a distribution of points in phase space. We characterize this distribution by use of a density function $\rho(p_j, q_j, t)$, where p_j are the canonical momenta, q_j are the conjugate coordinates, and t is time. $\int_V \rho(p_j, q_j, t)\, d^{3N}p\, d^{3N}q$ gives the number of points of the distribution in the phase space volume V.

If the system we wish to describe is in equilibrium, then the distribution of points in phase space does not change with time and the density function is independent of time. Because of the time independence, we only have to consider a set of phase space points instead of a set of phase space trajectories. Each phase space point in the distribution represents a possible evolution of the system. For simplicity, we will only consider equilibrium statistical mechanics.

The probability that the system is in some state is given by the following

[10] We can also construct an analog to the following argument in Boltzmann's approach to statistical mechanics.

probability density $D(p_j, q_j)$:

$$D(p_j, q_j) = \frac{\rho(p_j, q_j)}{\int \rho(p_j, q_j)\, d^{3N}p\, d^{3N}q} \tag{3.6}$$

The mean value or ensemble average of a measurable property f is:

$$\langle f \rangle = \int f(p_j, q_j) D(p_j, q_j)\, d^{3N}p\, d^{3N}q \tag{3.7}$$

Now for a system that consists solely of particles of the same kind, there can be distinct phase space points in the distribution of its ensemble that only differ by a permutation transformation. That is they only differ in which particles they assign to which single-particle states. These permuted points appear to violate permutation symmetry, for they are assigned separate probabilities by the probability density function even though they only differ by a permutation transformation.[11]

But this need not be the case. We can regard permuted points in the distribution as representing the same state without any false empirical results. One way to do this is to change our formalism. Instead of representing the ensemble as a distribution of phase space points, we can represent it as a set of particle distributions over μ-space, that is as a set of Z-space states. In this description there is no distinction as to which particle is in which single-particle state. Huggett (1999a) shows that this formalism provides an adequate description of classical statistical mechanics. And his argument is further endorsed by Albert (2000, p. 45–48).

But we do not need to consider the details of a Z-space description, for we need not abandon a phase space description in order to express the permutation symmetry of classical statistical mechanics. We can still use

[11] Both Bach (1985) and Costantini (1987) argue that classical particles are as indistinguishable as quantum particles since both are described by symmetric probability functions, which are invariant under particle permutation. It is true that a permutation of particles does not change the probability assigned to each state in both cases. But the important question is not whether a permutation affects the probability function, but whether it leads to a distinct state in the domain of the probability function. It is in this way that classical particles appear to be distinguishable and quantum particles indistinguishable. Symmetric probability functions therefore do not settle the issue of indistinguishability or permutation symmetry.

the same formalism, provided we correct for the redundant states in the distribution. This is completely equivalent to a Z-space description. To illustrate this corrected phase space description in more detail, let us consider the canonical ensemble. For a canonical ensemble, we keep the number of particles N and temperature T of the system fixed, but allow the energy of the each state in the ensemble to vary. In phase space, the density function for the canonical ensemble is:

$$\rho(p_j, q_j) = e^{-H(p_j, q_j)/kT} \tag{3.8}$$

where H is the Hamiltonian of the system. Again there are distinct points in this distribution that are connected by a permutation transformation.

The thermodynamic properties of the canonical ensemble follow from the partition function, which gives the volume of the canonical ensemble in phase space. From the partition function, we can calculate the distribution function, the thermodynamic equation of state, and the entropy, among other quantities. Given the density function (3.8), the partition function Q is:

$$Q = \frac{1}{h^{3N}} \int e^{-H/kT} d^{3N}p\, d^{3N}q \tag{3.9}$$

The constant h appears in the expression in order to ensure that the partition function is a dimensionless quantity. In quantum statistical mechanics, h must be equal to Planck's constant. But in classical statistical mechanics, only the dimensions of h are important. It can take any value without affecting the thermodynamical properties of the system described by the ensemble (see Huang, 1963, ch. 8).

If we take permuted phase points in the distribution to represent the same state, then there are redundancies in our density function. Assuming that classical particles are impenetrable, so that no two particles can share the same single-particle state, there are $N!$ distinct representations of each state in the distribution. We can account for these redundancies by dividing the density function ρ by $N!$.

This correction does not affect the probability density given in equation (3.6), for the correction to the numerator cancels the correction to the denominator. The mean values of observable quantities, such as the distribution function, therefore do not change. What does change is the partition function. With the correction to the density function, the partition function is now:

$$Q = \frac{1}{N!h^{3N}} \int e^{-H/kT} d^{3N}p\, d^{3N}q \tag{3.10}$$

This correction to the partition function does not lead to any false empirical results. Dividing the partition function by $N!$ leave nearly all the thermodynamic equations that follow from the partition function unchanged (Huang, 1963, p. 154). The only effect this correction has is on the statistical mechanical entropy function. But, as we will see shortly, this change to the entropy function does not lead to any incorrect results.

So even within the phase space formalism, we can regard permuted phase space points as representing the same state. We need not recognize any distinction about which particles are in which trajectories. Classical statistical mechanics, and with it classical physics, can be permutation symmetric. Two points in an ensemble that only differ by a permutation transformation need not describe distinct evolutions of the system. Rather they can describe the same evolution, but with different names.

3.6 Permutation Symmetry and Particle Trajectories

Before we continue, we should address another relevant view of indistinguishability. It is summarized well by Landau and Lifshitz. They write:

> In classical mechanics, identical particles (electrons, say) do not lose their "individuality", despite the identity of their physical properties. For we can imagine the particles at some instant to be "numbered", and follow the subsequent motion of each of these in its path; then at any instant the particles can be identified. . . . In quantum mechanics the situation is entirely different,

> as follows at once from the uncertainty principle. We have already mentioned several times that, by virtue of the uncertainty principle, the concept of the path of an electron ceases to have any meaning. ... Thus, in quantum mechanics, there is in principle no possibility of separately following each of a number of similar particles and thereby distinguishing them. We may say that, in quantum mechanics, identical particles entirely lose their "individuality". (Landau and Lifshitz, 1958, p. 204)

This is a typical view, and we find very similar statements expressed by others such as Blokhinstev (1964, ch. 19), Schiff (1968, ch. 10), and Jauch (1968, sec. 15.3), to name but three.

The idea is that quantum particles, unlike classical particles, are indistinguishable because they, unlike classical particles, do not possess spatiotemporal trajectories that allow us to distinguish between them.[12] And because of this indistinguishability, quantum particles are not objects.

Several of the authors we have just cited combine this notion of indistinguishability with the other notions we have delimited: those of indiscernibility and permutation symmetry. The three notions are distinct, but that does not mean that there are no connections between them. We have already mentioned in the previous chapter how spatiotemporal trajectories are related to indiscernibility. The spatiotemporal trajectories of classical particles, coupled with impenetrability, ensure that the particles are at least weakly discernible. But we have not yet addressed whether there is a connection between spatiotemporal trajectories and permutation symmetry. Nor have we directly addressed whether the lack of a spatiotemporal trajectory indicates a lack of individuality.

First question first. Whether or not particles possess permutation symmetry has nothing to do with whether or not they possess spatiotemporal trajectories. We have just shown that it is possible for both classical

[12] Of course whether or not quantum particles possess spatiotemporal trajectories depends upon our interpretation of quantum theory. For example if we accept a de Broglie-Bohm interpretation, then quantum particles do have spatiotemporal trajectories. But for the sake of argument, we will in this section assume that quantum particles do not possess clear spatiotemporal trajectories in all circumstances.

and quantum particles to possess permutation symmetry. This is true even though (under some interpretations of quantum theory) only classical particles have clear spatiotemporal trajectories. And as we will see in the next chapter, it is logically possible—although not physically possible—for quantum particles to lack permutation symmetry even though they do not possess clear spatiotemporal trajectories.

Now some might take a permutation of classical particles to be a physical process where the particles actually switch positions over time, so that the initial position of one particle is the final position of the other. Obviously this situation is distinct from one where the particles do not switch positions over time. This classical case differs from the quantum case. If quantum particles lack distinct trajectories, then there is not a similar sort of process where the quantum particles switch positions. In the quantum case, unlike the classical case, there is thus no way to distinguish between the process where the particles are permuted and when they are not permuted. There thus appears to be a difference in permutation symmetry between the classical and quantum cases that follows from a difference in the existence of particle trajectories.

But this understanding of permutation symmetry is problematic; for while it makes sense to talk about particles exchanging position over time in classical physics, it does not make any sense in quantum physics. Because quantum particles do not possess clear spatiotemporal properties, we cannot understand a permutation transformation as a physical process where they switch positions.

And even in the classical case, this is not the sort of permutation transformation we are interested in. In the classical case, what is at issue is whether a permutation of particle labels associated with a set of particle trajectories reflects a distinct physical situation. That set of trajectories can describe the particles switching positions, or it can describe them staying still. It does not matter. What matters is whether or not it makes any sense to say which particle is traveling along which fixed trajectory in the set. If it does not, then classical physics is permutation symmetric.

What is of interest in both the classical and quantum case, is not the results of some physical process where the particles switch position, but whether a permutation of particle labels in the respective classical and quantum state descriptions describes a distinct physical situation. This question we can ask and answer without considering trajectories.

This leads us to our second question. Does the absence of a spatiotemporal trajectories in any way indicate a lack of individuality? There is nothing about objecthood, as we have characterized it, that requires spatiotemporal trajectories. Particles can still be physical objects even though they possess no precise spatiotemporal location. They can still be identical to themselves and numerically distinct from other objects. We can still quantify over them. And we can still predicate other physical properties to them. So unless we define "individual" in a dogmatic way that requires spatiotemporal trajectories, the lack of spatiotemporal trajectories does not lead us to reject the individuality of quantum particles. The only apparent reason why we would think that spatiotemporal trajectories were necessary for individuality, is if we thought they were part of a necessary principle of individuation. We have already discussed in the last chapter that no such principle of individuation is necessary. So while we can maintain that classical particles are distinguishable in the sense of possessing spatiotemporal trajectories and quantum particles indistinguishable in the sense of lacking spatiotemporal trajectories, this difference in distinguishability has no ontological implications.

3.7 Entropy and Permutation Symmetry

3.7.1 Extensive Entropy

Now back to permutation symmetry. So far we have only shown that we can regard classical statistical mechanics, and with it classical physics, as permutation symmetric if we so choose. We can deny that there is any difference as to which particle is traveling along which single particle trajectory without contradicting experiment. But we have not yet given any reason why we should regard this distinction as physically meaningless, why we should take classical physics to be permutation symmetric.

One simple reason follows from parsimony. The distinction as to which particle is traveling upon which trajectory has no empirical content, for we can do away with it and still correctly describe the dynamical and statistical mechanical behavior of a system. It is therefore an unnecessary distinction that should be cut from our theory by use of Ockham's razor.

This is a good reason to claim that classical physics is permutation symmetric, but we need not content ourselves with it alone. There is another reason which follows from a consideration of entropy. We noted in the last section that when we divide out the redundant states from the phase space distribution describing the ensemble, the only thing we affect is the entropy function. Let us see how.

The statistical mechanical entropy S is related to the partition function by the following equation:

$$S = k \frac{\partial}{\partial T} T \ln Q \bigg|_{V,N} \tag{3.11}$$

The $N!$ correction to the partition function changes the form of the entropy function; for the corrected partition function ensures an extensive entropy, while the uncorrected partition function does not.

Extensive functions are ones that scale with the size of the system. For a gas, the size of a system is determined by its volume V and number of particles N. An extensive function must therefore take the form $F(cN, cV) = cF(N, V)$ for a gas. The internal energy is a paradigmatic example of an extensive quantity in thermodynamics. Intensive functions are ones that are identical for scaled copies of a system. For a gas, they must take the form $F(cN, cV) = F(N, V)$. The density of a gas is a paradigmatic example of an intensive quantity.

To see how the $N!$ correction affects the entropy, let us calculate the entropy for the simple example of an isolated classical gas of particles of the same kind that is confined in a fixed volume. The Hamiltonian for this system is:

$$H = \sum_{i=1}^{N} \frac{\mathbf{p_i}^2}{2m} + \sum_{1\leq i<j\leq N} \phi(|\mathbf{r_i} - \mathbf{r_j}|) \tag{3.12}$$

The corrected partition function is:

$$Q = \frac{1}{N!\lambda^{3N}} \int \cdots \int_V d\mathbf{r_1} \cdots d\mathbf{r_N} \exp\left[(-kT)^{-1} \sum_{i<j} \phi(|\mathbf{r_i} - \mathbf{r_j}|)\right] \tag{3.13}$$

where

$$\lambda = h(2\pi mkT)^{-1/2} \tag{3.14}$$

For typical interatomic potentials:

$$Q = \frac{V^N}{N!\lambda^{3N}} (f(T))^N \tag{3.15}$$

where $f(T)$ is an intensive function (Mandl, 1971, sec. 7.8). (When $f(T) = 1$ there is no interatomic potential and the gas is ideal.) By equations (3.11) and (3.15), the entropy for this system is:

$$S = Nk \frac{\partial}{\partial T} T \left(\ln \left[\frac{Vf(T)}{N\lambda^3} \right] + 1 \right) \Bigg|_{V,N} \tag{3.16}$$

where we have made use of the Stirling approximation ($\ln N! \approx N \ln N - N$ for large N). This entropy is extensive, $S(cN, cV) = cS(N, V)$.

If the partition function does not include this $N!$ correction, we do not arrive at an extensive entropy. If we use the uncorrected partition function the entropy is:

$$S = Nk \frac{\partial}{\partial T} T \ln \left[\frac{Vf(T)}{\lambda^3} \right] \Bigg|_{V,N} \tag{3.17}$$

From this we see that $S(cN, cV) \neq cS(N, V)$.

3.7.2 Generic Phases

So the $N!$ correction ensures an extensive statistical mechanical entropy. Gibbs, in his seminal work on statistical mechanics, took this as reason to regard classical statistical mechanics as permutation symmetric (Gibbs, 1902, ch. 15). In that work, Gibbs draws a distinction between specific and generic phases. Specific phases can differ by a permutation of particles of same kind, whereas generic phases cannot. A description by use of generic phases is equivalent to our corrected phase space description. Gibbs points out that generic phases, unlike specific phases, lead to an extensive entropy function. He argues that the entropy function should be extensive, and that, therefore, generic phases provide the correct description of statistical mechanical systems.

Gibbs points out the importance of the extensivity that follows from general phases by use of the following example.[13] Consider an airtight box split by a divider into two chambers of equal volume, where each chamber is filled by N particles of an ideal monotonic gas. The thermodynamical entropy function for an ideal monotonic gas is:

$$S(P,T) = \frac{5}{2}Nk\ln T - Nk\ln P + CN \tag{3.18}$$

where C is some constant (see Planck, 1927, §119). If the gases in each chamber are of different kinds, then, by use of equation (3.18) and Dalton's law of partial pressure, we can calculate that the change in thermodynamic entropy ΔS when we remove the divider and allow the gases to mix is $\Delta S = 2NkT\ln 2$.

If the gases in each chamber are of the same kind, then the system does not change its equilibrium state when we remove the divider, and $\Delta S = 0$ (see Planck, 1927, §234–238). One way to demonstrate this is by use of the extensivity and additivity of the thermodynamic entropy. We have already explained what we mean by "extensive". What we mean by "additive" is that the entropy of a system is equal to the sum of the entropy of its subsystems. By additivity, the initial entropy before the divider is removed is equal to the sum of the entropies of the gases in each chamber. Since the gases in each chamber are of the same kind, the sum of the entropies of two chambers

[13] Gibbs discusses the thermodynamic details of this example more thoroughly in (Gibbs, 1948, p. 165–168). Our presentation differs slightly from his.

is equal to twice the entropy of one chamber. And by extensivity, the final entropy after the divider is removed is just the entropy of a system that is twice the size of one of the chambers. The initial and final entropy is therefore equal and the change in entropy is therefore zero.

Through this example we see that the extensivity is an important feature of the thermodynamic entropy. Now if we are to connect statistical mechanics to thermodynamics, we want a statistical mechanical entropy that we can equate with the thermodynamical entropy, at least for systems in equilibrium.[14] So in order to correctly describe this example of the mixing of gases by use of statistical mechanics, the statistical mechanical entropy must be extensive like the thermodynamical entropy.[15]

Based upon this fact, Gibbs argues that the generic phase, which is insensitive to permutation, is the correct phase for classical statistical mechanics, for it naturally leads to an extensive entropy and explains the necessary $N!$ correction that we must make to equation (3.17) (Gibbs, 1902, ch. 15).[16] The mixing of two gases of the same kind both demonstrates the importance of an extensive statistical mechanical entropy and provides us with a further reason beyond parsimony to view statistical mechanics as permutation symmetric.

3.7.3 Responses to Gibbs

Often Gibbs is taken as presenting a paradox, which often bears his name. Within classical statistical mechanics, the $N!$ correction that we must make to equation (3.17) is supposedly paradoxical. Many hold that we can only

[14] We are not here advocating a reduction of the thermodynamical entropy to the statistical mechanical entropy. It might be that no reduction is possible (see Callender, 1999). But even if one entropy does not strictly reduce to the other, there must be a point of contact between the two in the equilibrium case; for this is necessary in order to ensure some connection between statistical mechanics and thermodynamics.

[15] There is a growing field of nonextensive statistical mechanics that uses a nonextensive entropy to describe phenomena that lie outside the range of equilibrium statistical mechanics (see Tsallis, 2004). Equilibrium statistical mechanics is only a limiting case of this theory. Nonextensive statistical mechanics is generally used to describe quasi-stable states of certain systems and the states of systems that do not have an extensive energy. Since our arguments only concern the relation between equilibrium statistical mechanics and equilibrium thermodynamics, we can safely ignore these cases of nonextensive entropy.

[16] Saunders (2003, p. 302) is a recent proponent of Gibbs's argument.

resolve this paradox by use of quantum theory. That is they think that the $N!$ correction to the partition function is simply a quantum correction that follows from the permutation symmetry of quantum theory (which in turn follows from the fact that quantum particles are not objects).[17] The uncorrected partition function is the correct partition function for classical statistical mechanics. However, since the world is quantum instead of classical, the uncorrected partition function does not lead to the correct expression for the entropy. To arrive at the correct expression for entropy we must consider the high temperature limit of the *quantum* partition function. Since the necessary $N!$ emerges in this limit (see Feynman (1972, p. 97) and Fujita (1991)), we can take the correction as a necessary quantum correction to an incorrect classical theory.

This is a prevalent view that we find in many of the important textbooks on statistical mechanics, such as Tolman (1938), Schrödinger (1948), Huang (1963), Feynman (1972), Kittel and Kroemer (1980), ter Haar (1995), and in numerous other places. Huang summarizes the view well when he writes:

> It is not possible to understand classically why we must divide [the partition function] by $N!$ to obtain the correct counting of states. The reason is inherently quantum mechanical. ...It is something that we must append to classical mechanics in order to get the right answers. (Huang, 1963, p. 154)

The $N!$ correction to the partition function does not follow from the permutation symmetry of classical physics, rather it follows from the permutation symmetry of quantum physics.

David Hestenes objects to this line of argument and defends Gibbs's view that the $N!$ correction indicates that generic phases are the correct phases for classical statistical mechanics. He points out that for Gibbs's argument it "is only necessary to invoke the general requirement that statistical and experimental [that is thermodynamical] entropies agree", and that "this argument requires no appeal whatever to quantum theory" (Hestenes, 1970, p. 841).

[17] French and Rickles (2003) also discuss this popular view.

Hestenes is right: We need not interpret the $N!$ correction as a quantum correction. We can take it as an expression of the permutation symmetry of classical physics without any reference to quantum physics. But we can go further and say that we should not regard the $N!$ correction as a quantum correction.

Even though the world is quantum, we do not have to take classical statistical mechanics simply as a limiting case of quantum statistical mechanics. We can also regard it as a theory that is independent of quantum theory, with its own ontology and interpretation. And this is how we must regard the theory if we are to make any interesting ontological comparisons to quantum theory, for only an independent theory will have a distinct ontology that we can compare to a quantum ontology. When we consider classical physics as an independent (albeit empirically false) theory, we cannot view the $N!$ correction as quantum correction. We need an alternative explanation that only relies upon our interpretation of classical physics. We have such an explanation. The $N!$ correction follows from the permutation symmetry of classical physics.[18]

But even if we accept that the $N!$ correction is not a quantum correction, we might deny that it is an expression of permutation symmetry. The idea

[18]There is also another version of the Gibbs paradox. Instead of questioning why we must correct the classical statistical mechanical entropy to correctly describe the mixing of two gases, it questions why there is a discontinuous change in the thermodynamical entropy for the mixing of two gases. If the gases in each chamber are of the same kind, there is no change in entropy. But if they are of different kinds, there is a change of entropy, and this change is not a function of how different the gases are. What is paradoxical here is why ΔS does not continuously approach zero as the gases in each chamber become qualitatively similar.

For many the two versions of the paradox are related (for example see Lewis (1930) and Post (2000, p. 126–127)). The reason there is a discontinuous change in the thermodynamic entropy is because extensivity leads to a discontinuous change in the statistical mechanical entropy. As we have pointed out, for many the extensivity of the statistical mechanical entropy follows from the permutation symmetry of quantum mechanics, which in turn follows from the fact that quantum particles lack identity. So for some it is ultimately the ontology of quantum particles that resolves this version of the Gibbs paradox.

There is a question as to whether the discontinuity in the statistical mechanical entropy adequately explains this discontinuity in the thermodynamic entropy (see Mosini, 1995). But if the extensive statistical mechanical entropy does resolve this version of the paradox, this solution is not solely the product of quantum theory. For once again we can understand the extensivity of the statistical mechanical entropy in classical physics as well.

is that we can view the $N!$ correction simply as part of the entropy function, lacking any particular physical meaning.[19]

In equation (3.11) above, we neglected the fact that entropy is only defined up to some constant. The correct equation relating entropy to the partition function should read:

$$S = k\,\frac{\partial}{\partial T} T \ln Q \bigg|_{V,N} + C \tag{3.19}$$

where C is some constant. Q can be either the corrected or uncorrected partition function. However, if it is the uncorrected partition function, then, in order to ensure an extensive entropy, C must be a certain function of N. We need not view this as an expression of permutation symmetry. We just take it as part of the definition of the entropy for the uncorrected partition function.

There is nothing incoherent about this position, but it really does not undermine the Gibbsian view. All it shows is that the permutation symmetry of classical physics does not necessarily follow from the $N!$ correction. We can accept this and claim that nonetheless we still have good reason to view the $N!$ correction as an expression of permutation symmetry; for only permutation symmetry offers an explanation as to why the constant must be of the form it is for the uncorrected partition function.

But there is another line of objection to Gibbs's argument. Some claim that the use of the $N!$ correction is simply a convention that does not possess any empirical content. If this is true, then the $N!$ correction is not empirically necessary and obviously does not provide us with any reason to view classical physics as permutation symmetric.

Huggett (1999a, p. 21–22) presents us with an argument along these lines. He claims that “additivity [and by extension extensivity] is not an observable consequence of thermodynamics” (Huggett, 1999a, p. 21). This is because the second law of thermodynamics only defines entropy differences, and not absolute entropies. The absolute entropies for qualitatively identical dis-

[19]This seems to be a claim that Huggett (1999a, p. 20–21) makes and that Gordon (2002, p. 413–414) endorses.

tinct systems can therefore differ by an additive constant, which prevents the function from being additive and extensive.[20]

Now Huggett attributes this argument to van Kampen (1984). But while van Kampen does present an argument against the Gibbsian view outlined above, it is little more complex than the summary given by Huggett leads us to believe.

Let us look at van Kampen's argument in greater detail. Van Kampen takes the definition of thermodynamic entropy, S, to follow from the second law of thermodynamics, which defines entropy as:

$$dS = \frac{dQ}{T} \tag{3.20}$$

where Q is thermodynamical heat. As van Kampen points out, "*the second law defines only entropy differences between states that can be connected by a reversible change*" (van Kampen, 1984, p. 305).[21]

Equation (3.20) only determines the entropy function S up to a constant. This constant is only fixed for states that are connected by a quasi-static (that is to say reversible) process. If the entropy function is extensive, then the constant must be a function of the number of particle N in the system. But as van Kampen writes, "There is no way, however, to compare entropy values belonging to different N, unless one introduces a new kind of process by which N can be varied in a reversible way" (van Kampen, 1984, p. 305).

The argument here follows from the fact that the second law only defines entropy differences for quasi-statistic processes. If there is no change in the number of particles for the process, then the entropy constant can be any function of N without affecting the entropy difference. The entropy constants associated with distinct systems of fixed particle number can differ, therefore, in such a way that the absolute entropy is not extensive.

It is only for reversible processes where the number of particles change that second law requires the entropy constant to be a certain function of

[20] Gordon (2002, p. 413–415) endorses Huggett's position.

[21] The emphasis is van Kampen's. By reversible, van Kampen means that the process is quasi-static. We mention this because there are other notions of reversibility (see Uffink, 2001).

N. It is therefore only for systems of variable particle number that the entropy needs to be extensive. So van Kampen holds that the second law does not require the entropy for systems of fixed particle number to be extensive, but it does require the entropy for systems of variable particle number. Huggett, stating the former, but neglecting the latter, slightly misrepresents van Kampen's position.

But this does not mean that van Kampen agrees with the Gibbsian argument for permutation symmetry. For van Kampen holds that while systems of variable particle number must have an extensive entropy, this extensivity does not reflect any sort of permutation symmetry. He demonstrates this by deriving the density function for the classical grand canonical ensemble, which is the proper statistical ensemble for systems with variable energy and particle number.

Before we present van Kampen's derivation, let us first present an alternative derivation that explicitly assumes permutation symmetry. Consider a large system of fixed particle number described by a canonical ensemble. Let the system have volume V, particle number N, and a Hamiltonian function H. Split the system into two subsystem systems, system 1 and system 2. Let the volume, particle number, and Hamiltonian of system 1 be V_1, N_1, and H_1, and let the volume, particle number, and Hamiltonian of system 2 be V_2, N_2, H_2. It is the case that $N = N_1 + N_2$ and $V = V_1 + V_2$. We further assume that $H = H_1 + H_2$. Let system 1 be much smaller than system 2, so that $N_1 \ll N_2$ and $V_1 \ll V_2$.

We want to know the density function $\rho(p_1, q_1, N_1)$ for N_1 particles in volume V_1 with coordinates (p_1, q_1) (which give the canonical coordinates for all N_1 particles). Assuming that the interactions between the systems 1 and 2 are negligible, it is the case that:

$$\rho(p_1, q_1, N_1) \propto \frac{1}{N_1! N_2!} e^{-H_1(p_1,q_1)/kT} \int_{V_2} e^{-H_2(p_2,q_2)/kT} dp_2 dq_2 \tag{3.21}$$

Once again we have explicitly assumed permutation symmetry, that is we have assumed that permutation of particles does not lead to a distinct physical situation. We therefore divided out the $N_1!$ distinct permutations from system 1 and the $N_2!$ distinct permutations from system 2. We can rewrite expression (3.21) as follows:

$$\rho(p_1, q_1, N_1) = C\frac{1}{N_1!}e^{-H_1(p_1,q_1)/kT}\int_{V_2} e^{-H_2(p_2,q_2)/kT}dp_2dq_2 \tag{3.22}$$

where we have incorporated $N_2!$ into the unspecified proportionality constant. When we take the thermodynamic limit of the whole system, we can eliminate reference to system 2 by introducing the chemical potential μ in the standard way (see Huang, 1963, sec. 8.3), so that we have:

$$\rho(p_1, q_1, N_1) = C\frac{e^{N_1\mu/kT}}{N_1!}e^{(-P_1V_1-H_1(p_1,q_1))/kT} \tag{3.23}$$

where P_1 is the pressure of system 1.

Now van Kampen also derives equation (3.23), but without making any corrections that follow from permutation symmetry (van Kampen, 1984, p. 308–309). He starts by taking the density as follows:

$$\rho(p_1, q_1, N_1) = C\frac{N!}{N_1!(N-N_1)!}e^{-H_1(p_1,q_1)/kT}\int_{V_2} e^{-H_2(p_2,q_2)/kT}dp_2dq_2 \tag{3.24}$$

This version of the density function is not permutation symmetric. It counts a permutation of two particles, whether within a subsystem or between subsystems, as a distinct physical situation. (The combinatorial expression to the right of C expresses the fact that a permutation of particles between subsystems produces a distinct physical situation.) When we take the thermodynamic limit of the whole system, equation (3.24) leads to equation (3.23). We therefore can derive the same density function for the grand canonical ensemble regardless of whether or not we take the system to be permutation symmetric.

Now Gibbs presents his comments about general phases in his discussion of the grand canonical ensemble. There he takes the $N!$ in the density function of the grand canonical ensemble to follow from the use of general phases, which we in turn interpret as permutation symmetry. Van Kampen's claim is that Gibbs's explanation is wrong, and that the $N!$ arises naturally in the derivation of the density function for the grand canonical ensemble without any assumptions of general phases or permutation symmetry. Van Kampen concludes, "*the $N!$ arises from the computation of phase space volume according to the original rules without any additional postulate*, either classical or quantummechanical" (van Kampen, 1984, p. 309).[22]

Van Kampen attributes his conclusion to Ehrenfest and Trkal (1920), who present a similar view to van Kampen's when they write, "*The law of dependence on N can only be satisfactorily settled by utilizing a process in which N changes reversibly and then comparing the ratios of the probability with the corresponding differences of entropy*" (Ehrenfest and Trkal, 1920, p. 163).[23] Ehrenfest and Trkal establish this conclusion, however, by use of a different example. Instead of discussing the grand canonical ensemble of a single gas, they discuss a mixture of gases that are able to chemically interact, where these chemical interactions allow for the particle number of each gas in the mixture to vary. Like van Kampen, they derive expressions with the necessary $N!$'s without making any assumption about general phases or permutation symmetry, thereby removing "any remaining obscurities as regards the occurrence of $N_1!N_2!\ldots$" (Ehrenfest and Trkal, 1920, p. 163).

So according to van Kampen (and Ehrenfest and Trkal) extensivity, and the $N!$ correction it requires, does not lead us to conclude that classical physics is permutation symmetric. The entropy only needs to be extensive for the grand canonical ensemble, and this does not require permutation symmetry.

Van Kampen presents us with a very persuasive argument to which there is no quick response. But there is a response. If we accept a more general statement of the second law thermodynamics, it is not only the entropy for systems of variable particle number that needs to be extensive; the entropy for systems described by the microcanonical and canonical ensembles also needs to be extensive. This requires $N!$ corrections for which permutation symmetry provides the only natural explanation. The details of this response are the subject of the next section.

[22] The emphasis is his.

[23] The emphasis is theirs.

3.8 The Extensivity and Additivity of Entropy

Before we go any further, we should say that van Kampen's argument appears right if we accept his definition of the second law of thermodynamics and entropy. If we accept equation (3.20) as the sole definition of entropy, then only the entropy function of systems of variable particle number need to be extensive, and this extensivity can be derived without recourse to permutation symmetry. Extensivity does not then provide any evidence of permutation symmetry. We can still rely on arguments of parsimony to establish permutation symmetry, but no longer arguments from entropy.

But van Kampen's argument is wrong if we accept a more general definitions of entropy and the second law. As we will see shortly, a more general definition provides us with a thermodynamic entropy that is extensive for all systems, regardless of whether particle number is fixed or not. This requires that the entropy associated with all statistical ensembles is extensive, not only the entropy associated with the grand canonical ensemble. As we have shown above, this requires corrections that are best understood as expressions of permutation symmetry.

3.8.1 The Standard Approach to Entropy

But before we turn to the axiomatic approach, let us review the standard approach to entropy, which is found in nearly every textbook, and which provides the foundation of van Kampen's argument. Fermi (1936) gives a typical presentation of this approach in his widely read introduction to thermodynamics.

Fermi takes the second law (at least initially) to be equivalent to Kelvin and Clausius's principles. These are negative statements about cyclical processes that essentially forbid perpetual motion of the second kind. (A machine in perpetual motion of the second kind transfers heat from a lower to a higher temperature without any other effects.[24]) By use of Kelvin and

[24]Perpetual motion of the second kind is not to be confused with perpetual motion of

Clausius's principles, Fermi proves that no engine is more efficient than a Carnot cycle. By considering the efficiency of a Carnot cycle, he introduces the thermodynamic temperature and shows that it is equivalent to the empirical temperature.

Through further use of Carnot cycles and Kelvin and Clausius's principles, Fermi proves $\sum_{i=1}^{n} Q_i/T \leq 0$ for a system that transfers Q_i heat at T_i temperature from n sources, where Q is positive for heat received by source. The equality holds for quasi-static processes. For a continuous distribution of sources, the summation turns into the integral $\oint dQ/T \leq 0$. Many take this to be the quantitative statement of the second law of thermodynamics. Again the equality only holds for quasi-static processes. Fermi goes on to show that the value of $\int_A^B dQ/T$ for a quasi-static process is independent of the path taken in state-space. It is therefore possible to define the entropy S to be a function of state such that $dS = dQ/T$.

This approach to entropy owes its origin to Clausius (see Uffink, 2001, sec. 6.2). There are several technical problems with this approach. But putting the technical problems aside, the important point of this approach for our purposes is that it introduces entropy through an analysis of Carnot cycles and quasi-static processes. This analysis only uniquely defines entropy differences for quasi-static processes. And since entropy is only defined by the equation $S(B)-S(A) = \int_A^B dQ/T$, van Kampen is right to claim that extensivity is only required for quasi-static processes where the particle number of a system changes.

3.8.2 An Axiomatic Approach to Entropy

There is an alternative tradition, starting with Carathéodory, that views entropy and the second law in a different way than we find in the Clausius approach. This tradition, which we will call the axiomatic approach, attempts to provide a more rigorous mathematical formalism for thermodynamics, and an axiomatic system from which the second law follows. In order

the first kind. A machine in perpetual motion of the first kind is able to do work without consuming heat. Such a machine violates the first law of thermodynamics, which is an expression of the conservation of energy.

to understand how entropy and the second law differ in this tradition, we will examine the work of two of the latest contributers to the tradition, that is the work of Lieb and Yngvason (1999).

We start with some notation and concepts. In thermodynamics we describe a system by use of an equilibrium state-space Γ, where each point X in Γ represents an equilibrium state of the system. For an ideal gas, X is specified by any two of following three thermodynamic quantities: pressure, volume, or temperature. X will depend on other thermodynamic quantities for more complicated systems.

Two operations that we will encounter below are composition and scaling. First composition. We can take the composition of any two wholly distinct system to produce a third system that has the other two as its subsystems. For example we can have two systems sitting on the same workbench at some distance from each other, and consider the composed system that has these two as its subsystems. The two systems may or may not interact with each other; but when each is in equilibrium, the state of the composed system is given by the states of the subsystems. Let Γ_X and Γ_Y be the equilibrium state-spaces associated with two wholly distinct systems, and let X and Y be the state-variables for those systems, where $X \in \Gamma_X$ and $Y \in \Gamma_Y$. We take the Cartesian product $\Gamma_X \times \Gamma_Y$ to be the state-space of the composed system and (X, Y) to be the state-variable for the composed system, where $(X, Y) \in \Gamma_X \times \Gamma_Y$.

Now scaling. We can scale a system in size so that its extensive quantities are multiplied by some factor $c > 0$, but its intensive quantities remain unchanged. Let Γ be the state-space and X be the state-variable of some system. Now consider a duplicate system that is scaled by a factor c. Let Γ^c be the state-space and cX be the state-variable for the scaled duplicate.[25]

With these operations in hand, let us move on to the second law. For Lieb and Yngvason, the second law essentially concerns adiabatic accessibility.

[25]Scaling does have its limits. A scaled duplicate might be too big or too small for thermodynamics to apply. If we consider a scaled duplicate that is too big, we can no longer ignore gravitational effects. If we consider a scaled duplicate that is too small, we can no longer ignore quantum effects. We therefore must limit scaling to a regime where thermodynamics applies.

Take X and Y to be two thermodynamical equilibrium states. We then define adiabatic accessibility as follows:

> A state Y is adiabatically accessible from a state X, in symbols $X \prec Y$, if it is possible to change the state from X to Y by means of an interaction with some device (which may consist of mechanical and electrical parts as well as auxiliary thermodynamic systems) and a weight, in such a way that the device returns to its initial state at the end of the process whereas the weight may have changed its position in a gravitational field. (Lieb and Yngvason, 1999, p. 17)

Adiabatic accessibility should not be confused with an adiabatic process. An adiabatic process is one where no heat is exchanged at any time between the system and the external environment. In this definition of adiabatic accessibility, there can be an exchange of heat between the system and the auxiliary device, so long as the auxiliary device returns to its initial state. However, if two states are adiabatically accessible, there will be an adiabatic process connecting them (Lieb and Yngvason, 1999, p. 23).

Lieb and Yngvason postulate six axioms concerning adiabatic accessibility. They are:

> Reflexivity: $X \prec X$.
>
> Transitivity: $X \prec Y$ and $Y \prec Z$ implies $X \prec Z$.
>
> Consistency: $X \prec X'$ and $Y \prec Y'$ implies $(X, Y) \prec (X', Y')$.
>
> Scaling Invariance: If $X \prec Y$, then $cX \prec cY$ for all $c > 0$.
>
> Splitting and Recombination: For $0 < c < 1$, $X \prec (cX, (1-c)X)$ and $(cX, (1-c)X) \prec X$.
>
> Stability: If $(X, \epsilon Z_0) \prec (Y, \epsilon Z_1)$ holds for a sequence of ϵ's tending to zero and some states Z_0 and Z_1, then $X \prec Y$.
>
> (Lieb and Yngvason, 1999, p. 21)

They also state a further hypothesis, the comparison hypothesis, which states:

> Any two states X and Y in the same state-space are comparable, i. e. $X \prec Y$ or $Y \prec X$. (Lieb and Yngvason, 1999, p. 22)

Lieb and Yngvason are able to derive the comparison hypothesis from a set of further axioms for simple thermodynamic systems and their compositions. Simple systems are systems that have state-spaces with only one energy coordinate.

With all of this in place, we can now present Lieb and Yngvason's formulation of the second law. They state it as follows:

> There is a real-valued function on all states of all systems (including compound systems), called entropy and denoted by S such that:
>
> When X and Y are comparable states then $X \prec Y$ if and only if $S(X) \leq S(Y)$.
>
> If X and Y are states of some (possibly different) system and if (X,Y) denotes the corresponding state in the composition of the two systems, then the entropy is additive and extensive for these states. (Lieb and Yngvason, 1999, p. 19)

They are able to derive this result from the axioms and the comparison hypothesis (along with some further axioms that deal with mixtures and other special cases). Based on the properties of the entropy function, they are able to derive equation (3.20).

The details do not concern us. What does concern us is their presentation of entropy and the second law. For them entropy is an ordering function over states that tells us by use of the second law which equilibrium states are adiabatically accessible from other states. They do derive the standard expressions of entropy and second law that we find in the Clausius approach, but these are simply special cases that follow from their more general expression of the second law.

3.8.3 The Empirical Content of Extensivity and Additivity

With this summary of the axiomatic approach in place, let us see how it affects van Kampen's argument about extensivity. Lieb and Yngvason assume from the very beginning that entropy is both extensive and additive for all systems. They are able to show that there is an unique function (up to an affine transformation) that satisfies these requirements, along with their axioms. This alone might be enough to convince us that the extensivity of entropy is an important feature for all systems, even those of fixed particle number; for it is impressive that there is an unique function in this axiomatic system that is both additive and extensive. And even if this extensivity does not directly express some aspect of nature, it does provide us with a simpler and stronger theory; for if we do not require extensivity, then there is a whole set of entropy functions that satisfy the axioms.[26]

But all might not be convinced by this. Some might still claim that extensivity, even if it does narrow the set of acceptable entropy functions, is just a mathematical nicety with no empirical content for systems of fixed particle number. If this is the case, van Kampen's position remains unaffected.

But we can make a stronger argument for the empirical content of extensivity in the axiomatic approach than by simply pointing out that it narrows the set of entropy functions. We will do this by considering the more general property of additivity. Extensivity follows from additivity given the splitting and recombination axiom stated above, and provided we only scale by rational numbers (which is required if the entropy function is to be bound (Lieb and Yngvason, 1999, p. 20) and reasonable if we assume that all thermodynamic systems are composed of discrete atoms of some sort). Since the splitting and recombination axiom is one of the main axioms of this approach, and appears to be essential to the approach, its empirical content is clear. So we can demonstrate the empirical content of extensivity by demonstrating the empirical content of additivity.

Now if we take a Clausius approach to entropy, and take equation (3.20)

[26] Simon Saunders has expressed something like this view to me in conversation.

as our definition of entropy, then absolute entropy functions associated with different systems of fixed particle number do not need to be additive for the same reason they do not need to be extensive. Equation (3.20) only determines the absolute entropy function up to a constant. That constant only needs to be fixed for a given quasi-static process. The constants of the absolute entropy function can therefore vary from system to system in such a way that the absolute entropy is not additive. This is true even if we hold that entropy differences are additive, so that we can determine the change of entropy for a system by summing the change in entropy for its subsystems.[27]

But in the axiomatic approach, entropy is not just defined as a quantity that changes by a given amount in a given process. As we have already said, entropy is an ordering function that establishes which states are adiabatically accessible from other states. The important feature of entropy in this approach is that it establishes an order not only over states in a single state-space associated with a single system, but also over states in multiple state-spaces associated with multiple systems. That is entropy establishes whether a state in one state-space can be adiabatically accessible from a state, or composition of states, in another state-space. In order to serve this function, entropy must be additive. And it is in this way that additivity possesses empirical content.

We can demonstrate this by considering an example. Consider two wholly distinct thermodynamical systems A and B. Both A and B have their own equilibrium thermodynamical state-space. Let X be the equilibrium state-variable for system A and Y be the equilibrium state-variable for system B. Systems A and B need not be the same sort of thermodynamical system. Their thermodynamical equations of state might be very different.

A and B are initially isolated from each other. But let consider the case

[27]Now van Kampen appears to accept some form of additivity, even though he labels it a convention (van Kampen, 1984, p. 305–306). For it seems that additivity for him is conventional not because it lacks empirical content, but because it is necessary in order to expand the notion of entropy to a broader class of phenomena, such as non-equilibrium systems. Huggett (1999a, p. 21) misrepresents van Kampen's views on this point. This said, there is no reason why van Kampen would hold that absolute entropy is additive, only that entropy differences are, for this is all that is required to expand entropy in a way that is consistent with equation (3.20).

where we bring systems A and B together and allow them to interact in such a way that they form system C, such that system C has its own state-space. Let Z be the state variable for system C. Further let us stipulate that the particle numbers for system A and B are constant, so that if we take equation (3.20) as our definition of entropy, there is no reason that the absolute entropy should be additive.

If the entropy is, however, additive, then we can establish whether $(X, Y) \prec Z$ without knowing the particulars of how system A and system B interact. To see this consider the entropy functions associated with each of the systems. Let $S_A(X)$ be the entropy function for the state-space of system A, $S_B(Y)$ for system B, and $S_C(Z)$ for system C. The entropy functions are of the following form:

$$S_A(X) = f_A(X) + C_A \tag{3.25}$$

$$S_B(Y) = f_B(Y) + C_B \tag{3.26}$$

$$S_C(Z) = f_C(Z) + C_C \tag{3.27}$$

If the entropy is additive, then:

$$S(A, B) = S_A(X) + S_B(Y) \tag{3.28}$$

and the constants C_A, C_B, and C_C will be properly coordinated so that:

$$S_A(X) + S_B(Y) \leq S_C(Z) \tag{3.29}$$

if and only if $(X, Y) \prec Z$.

If the entropy is not additive, then there is no constraint upon the constants C_A, C_B, and C_C. Whether or not expression (3.29) holds therefore depends upon an arbitrary choice of constants, and cannot serve as a necessary and sufficient condition for adiabatic accessibility. We therefore cannot determine whether Z is adiabatically accessible from X and Y.

And so this is the empirical content of additivity: Additivity allows us to determine the adiabatic accessibility between states in different state-spaces.

In the Clausius approach the entropy function is only defined with respect to the processes undergone by a given system. That is it is only associated with a single thermodynamic state-space, and does not encode any information about subsystems. It therefore does not have to be an additive function. In the axiomatic approach, however, the entropy function is an ordering function that applies to more than just the state-space of a single system. It also applies to the state-spaces of that system's subsystems and to the state-spaces of those systems to which it can be a subsystem. The entropy encodes the adiabatic accessibility between states in all of these state-spaces. This requires that it is additive, even for processes where the particle number is fixed.

It is remarkable that we can determine adiabatic accessibility between states in different state-spaces by use of additivity independently of any detailed knowledge of the interaction between the subsystems. For Lieb and Yngvason this is "at the *heart of thermodynamics*" (Lieb and Yngvason, 1999, p. 20).[28] It is clear that, being at the heart of thermodynamics, it is no mere convention. And since additivity has empirical content in this axiomatic approach, so, by extension, does extensivity. For as we have already mentioned, the extensivity of entropy follows from additivity and the axiom of splitting and recombination. Since the thermodynamic entropy needs to additive, even for systems of fixed particle number, it also needs to be extensive, even for systems of fixed particle number.

If we accept this axiomatic approach to entropy, we must disagree with van Kampen, and agree with the Gibbsian. The thermodynamic entropy needs to be extensive, even for processes of fixed particle number. The statistical mechanical entropy needs to be extensive as well. This requires an $N!$ correction to the statistical mechanical entropy. This $N!$ correction follows from the permutation symmetry of classical physics. The conclusion to draw is that classical physics, like quantum physics, is permutation symmetric. Like quantum physics, in classical physics a permutation of two particles of the same kind does not lead to a description of a distinct physical state. Two trajectories in phase space that only differ by a permutation transformation

[28]The emphasis is theirs.

represent the same physical state.

It seems to me that the axiomatic approach is a substantial step forward from the Clausius approach to entropy. But all might not share this view. We can reestablish van Kampen's argument by simply sticking to the Clausius approach. If we do, then we must rely upon arguments of parsimony to establish the permutation symmetry of classical physics.

But whether or not we accept this axiomatic approach, the important point is that we have good reason to view classical physics as permutation symmetric. And because classical physics possesses permutation symmetry, we must deny that the permutation symmetry of particles entails that they are not objects. We must also deny that the difference between classical and quantum statistics is due to a difference in permutation symmetry. While these views are popular, they are wrong. Permutation symmetry does not indicate that quantum particles are not objects like classical particles.

Chapter 4

Quantum Statistics

4.1 The Question of Quantum Statistics

Indistinguishability is often taken to have important consequences for physical statistics. As we have mentioned, it is often held that permutation symmetry implies quantum statistics, and that, therefore, the difference between classical and quantum statistics is due to a difference in permutation symmetry. (In what follows we shall always take "quantum statistics" to refer solely to Bose-Einstein and Fermi-Dirac statistics.) Now our conclusion of the previous chapter has shown this to be incorrect. Classical physics, like quantum physics, is permutation symmetric. This permutation symmetry does not lead to quantum statistics. Permutation symmetry alone, therefore, cannot explain the difference between classical and quantum statistics.

But an important question remains: How is permutation symmetry related to quantum statistics? There obviously is some connection. As we will see shortly, symmetrization, which is a form of permutation symmetry, is responsible for quantum statistics in the quantum formalism. But, as we will discuss in this chapter, the application of permutation symmetry to the quantum formalism does not imply symmetrization or quantum statistics. The application of permutation symmetry to the quantum formalism, however, does imply a departure from classical statistics. As we will see, this does not indicate a difference in permutation symmetry between classical and quantum physics, only a difference in the consequences of permutation symmetry in the two theories.

4.2 The Symmetrization Requirement and Quantum Statistics

In one sense there is no mystery surrounding quantum statistics. Quantum statistics clearly follow from the symmetrization requirement, which is a form of permutation symmetry. The symmetrization requirement requires that all state-vectors are symmetrized upon the permutation of particle labels for particles of the same kind. For a system that consist solely of particles of the same kind, a symmetrized state-vector satisfies the equation:

$$\hat{P}|\Psi\rangle = \pm|\Psi\rangle \tag{4.1}$$

where $|\Psi\rangle$ is a state-vector and $\hat{P}$ is a permutation operator that permutes particle labels in the description of the state-vector. The state-vector is symmetric if the sign is positive and antisymmetric if it is negative. Symmetric state-vectors describe bosons, antisymmetric state-vectors describe fermions. We will use the term "symmetrized" to refer to both symmetric and antisymmetric state-vectors. Symmetric and antisymmetric vectors each form a separate subspace of the tensor product Hilbert space that is invariant (although reducible) under the action of the permutation group.[1] There is, however, no unitary transformation that both commutes with the permutation group and connects a vector in one of the subspaces to any vector outside that subspace.[2]

The quantum Hamiltonian is permutation symmetric—it commutes with

[1] A space is invariant under the action of a group if and only if vectors in the space are only transformed into other vectors in the space by the representation of the group on the space. A space is reducible with respect to a group if it contains a nontrivial subspace that is invariant under the action of the group.

[2] There is another important subspace that is invariant under the action of the permutation group. It consists of parasymmetric states. We will consider this subspace in more detail in the next section.

the permutation group ($[\hat{H}, \hat{P}] = 0$)—if the quantum potential is permutation symmetric. For particles of the same kind, the quantum potential is permutation symmetric for the same reason that the classical potential is permutation symmetric: The potential must supervene upon the dynamical properties of the system. Since particle labels do not represent dynamical properties, a permutation of particle labels cannot lead to a change in the potential function. The quantum Hamiltonian is therefore permutation symmetric. Since it is symmetric, a state-vector in either the symmetric or antisymmetric subspace never evolves into any vector outside that subspace. Symmetric systems will always be symmetric; antisymmetric systems will always be antisymmetric. The symmetrization requirement therefore restricts state-vectors to subspaces of state-space. This is what is responsible for quantum statistics.

We can see this when we look at the quantum canonical ensemble. As with the classical canonical ensemble, the number of particles N and the temperature T of the system are constant, but the energy E of each state in the ensemble varies. In quantum mechanics, we describe the states in the canonical ensemble not by use of a density function defined on phase space, but by use of density matrix ρ defined on Hilbert space. Let $|\Phi_j\rangle$ be a many-body energy eigenfunction of the N-particle system with an associated energy eigenvalue E_j. The density matrix describing the canonical ensemble is then:

$$\rho = \sum_j e^{-E_j/kT} |\Phi_j\rangle\langle\Phi_j| \tag{4.2}$$

where the sum is taken over the set of orthogonal energy eigenfunctions.

The quantum partition function Q is defined as:

$$Q = \sum_j e^{E_j/kT} \tag{4.3}$$

$$= \mathrm{Tr}\rho \tag{4.4}$$

where again the sum is taken over all energy eigenfunctions. The quantum partition function has the same relation to the thermodynamical properties of the system as the classical partition function, which we discussed in the last chapter. The expectation value of an observable $\hat{O}$ for the ensemble is:

$$\langle \hat{O} \rangle = \frac{\mathrm{Tr}(\hat{O}\rho)}{Q} \tag{4.5}$$

Given equation (4.2) it is clear how the symmetrization requirement affects statistics. If we limit $|\Phi_j\rangle$ to symmetric states, then the ensemble described by the density matrix is limited to the symmetric subspace of the Hilbert space. This affects both the partition function and the expectation values of the system, giving us Bose-Einstein statistics. Similarly if we limit $|\Phi_j\rangle$ to antisymmetric states, then we have Fermi-Dirac statistics. If we place no limitation on the states in the density matrix, then the ensemble includes orthogonal energy eigenfunction that span the whole space. This leads to classical statistics.[3]

4.3 Permutation Symmetry in Quantum Theory

Some authors, such as Landau and Lifshitz (1958, p. 204–205) and Blokhinstev (1964, p. 395–396), have claimed that symmetrization, and with it quantum statistics, follows from indistinguishability. They argue that, because of the indistinguishability of quantum particles, a permutation of those particles cannot affect the physical state. This implies that a permutation operator can only affect the physically meaningless phase factor of a state-vector. That is:

$$\hat{P}|\Psi\rangle = \lambda|\Psi\rangle \tag{4.6}$$

where $\hat{P}$ is a permutation operator permuting two given particles and λ is the phase. Now if we apply the same permutation operator twice, we permute the particles, but then permute them back. The state-vector is therefore

[3] If we limit the density matrix to parasymmetric states, then we will have a partition function that gives parastatistics. We will consider the example of parasymmetric states and parastatistics in further detail in the next section.

unchanged so that:

$$\hat{P}^2|\Psi\rangle = |\Psi\rangle \quad (4.7)$$

This implies that $\lambda = \pm 1$ so that:

$$\hat{P}|\Psi\rangle = \pm|\Psi\rangle \quad (4.8)$$

Only symmetric and antisymmetric state-vectors satisfy this equation.

But neither of the two notions of indistinguishability that we have considered imply the symmetrization requirement, and with it quantum statistics. Let us first consider the case of indistinguishability understood in terms of indiscernibility. Several authors, such as de Muynck and van Liempd (1986) and Dieks (1990), have pointed out that indiscernibility (or as de Muynck, van Liempd, and Dieks call it, observational indistinguishability) does not imply symmetrization. At the heart of these arguments is the work of Messiah and Greenberg (1964).[4]

Messiah and Greenberg take the minimum requirement of indistinguishability to hold that "[d]ynamical states represented by vectors which differ only by a permutation of identical particles cannot be distinguished by any observation at any time" (Messiah and Greenberg, 1964, p. 250). They state this formally for a system of particle of the same kind as follows:

$$(\langle\Psi|\hat{P}^\dagger\hat{O}\hat{P}|\Psi\rangle = \langle\Psi|\hat{O}|\Psi\rangle) \quad (4.9)$$

Equation (4.9) is often called the indistinguishability postulate (IP).

IP implies that particles of the same kind in a system will be either weakly discernible or indiscernible (assuming that all the physical properties and relations of quantum particles supervene upon the dynamical equations and the expectation values of observables). But as Messiah and Greenberg note, the symmetrization requirement is only a sufficient, not a necessary condition, for IP.

If we require all observables to be permutation symmetric, so that not all Hermitian operators represent observables, but only those that commute

[4]Messiah and Greenberg's views are further expounded by Hartle and Taylor (1969).

with the operators of the permutation group ($[\hat{O}, \hat{P}] = 0$), then any state-vector satisfies IP, even those that are unsymmetrized. IP therefore allows for statistical ensembles of unsymmetrized states, ensembles that lead to classical statistical mechanical and thermodynamical equations. So indistinguishability understood in terms of indiscernibility does not necessarily lead to symmetrization or quantum statistics.[5]

Neither does indistinguishability understood in terms of permutation symmetry. Consider again the symmetric and antisymmetric subspaces of Hilbert space. These subspaces are invariant under the action of the permutation group, but neither is irreducible for an arbitrary number of particles. The irreducible subspaces of each are rays. A ray is a one dimensional subspace that consists of normalized vectors, where vectors within the same ray only differ by a physically meaningless phase factor. Since they only differ by a phase factor, vectors in the same ray describe the same state. And since these rays are invariant under the action of the permutation group, a permutation transformation will only connect a vector in one ray to another vector in the same ray. That is a permutation will only affect the phase of state-vectors confined to these rays.

The symmetrization requirement restricts state-vectors to these irreducible rays for particles of the same kind. Symmetrization is therefore a form of permutation symmetry. It ensures that two state-vectors connected by a permutation transformation describe the same state; for any two state-vectors connected by a permutation transformation will belong to the same ray.

All this is well and good. But for a Hilbert space that describes three or more particles, there are other invariant subspaces that are irreducible under the action of the permutation group. These irreducible subspaces are of more than one dimension and lie outside the symmetric and antisymmetric subspaces. Messiah and Greenberg (1964) have labeled these subspaces generalized rays. The sum of these generalized rays form the parasymmetric subspace, which is wholly disjoint from both the symmetric and anti-symmetric

[5] Kaplan (1976) and Sarry (1979) both argue that the indiscernibility of quantum particles does ensure the symmetrization requirement, despite Messiah and Greenberg's work. But they do this by imposing further conditions that do not appear warranted. For further criticism of their arguments, see de Muynck and van Liempd (1986) and Dieks (1990).

subspaces.

As Messiah and Greenberg were the first to point out, if we require all observables to be permutation symmetric, then we can also take parasymmetric vectors to represent states without violating permutation symmetry. Since all observables are permutation symmetric, two vectors in the same generalized ray will have the same expectation values for all observables. We can thus take vectors in each generalized ray to represent the same physical state, and claim, in analogy with rays, that two vectors in the same generalized ray only differ by a physically meaningless generalized phase. Since two parasymmetric vectors connected by a permutation transformation belong to the same generalized ray, a permutation transformation does not lead to a distinct state. Parasymmetric vectors are therefore just as permutation symmetric as symmetrized vectors.

Just as with symmetric and antisymmetric vectors, there is no permutation symmetric unitary operator that connects any parasymmetric vector to one that is not parasymmetric. So just as with the symmetric and antisymmetric state-vectors, a parasymmetric vector can never dynamically evolve to a state that is not parasymmetric.

Statistical ensembles that are confined to parasymmetric states are described by different statistics than quantum statistics. When we take a statistical ensemble of parasymmetric states, we arrive at parastatistics, not Bose-Einstein or Fermi-Dirac statistics.

So indistinguishability, understood in terms of permutation symmetry, does not require symmetrization. If we confine state-vectors to invariant subspaces that are irreducible under the action of the permutation group, then state-vector can be permutation symmetric even if they are not symmetrized.

Some might think that, instead of symmetrization, this is actually what is implied by permutation symmetry: the restriction of state-vectors to invariant subspaces that are irreducible under the action of the permutation group. But it is important to note that even this weaker condition is not implied by permutation symmetry.

The minimum that permutation symmetry requires is that two state-descriptions that differ only by a permutation of particles describe the same physical situation. Let us call this form of permutation symmetry the permutation symmetry of state-descriptions. The permutation symmetry of state-descriptions is the most general statement of permutation symmetry. It applies to both classical and quantum formalisms, for a state-description can be a trajectory in phase space or it can be an evolving vector in Hilbert space. This form of permutation symmetry does not require that Hilbert space state-vectors are confined to invariant subspaces that are irreducible under the action of the permutation group.

To see this, consider a quantum system that consists of particles of the same kind. Let us describe the physical situation of the system at some time t by use of a proposition $S(t)$. We here understand a proposition in a very general sense. We take it to be an abstract entity that expresses the content or meaning of a sentence or state-description. Since $S(t)$ is dependent upon our approach to the measurement problem and on what we take to be the relevant properties describing a state and a system, there is no consensus on what specific proposition $S(t)$ is. But luckily what follows does not depend upon the specific nature of the proposition. All that we need to admit is that there is some unique proposition $S(t)$.

Let f be a function that takes us from state-vectors to propositions so that $f(|\Psi(t)\rangle) = S(t)$. f is basically an interpretation of the state-space. It tells us what physical situation is actually associated with a state-description. If two state-descriptions that only differ by a permutation are to describe the same physical situation, then they must express the same proposition and the following must hold:

$$f(e^{-i\hat{H}(t)/\hbar}|\Psi_0\rangle) = f(e^{-i\hat{P}^\dagger\hat{H}(t)\hat{P}/\hbar}\hat{P}|\Psi_0\rangle) \tag{4.10}$$

where $|\Psi_0\rangle$ is the state-vector for the system at time $t = 0$. But since we also maintain that Hamiltonians, both classical and quantum, are invariant under permutation, we can just write:

$$f(e^{-i\hat{H}(t)/\hbar}|\Psi_0\rangle) = f(e^{-i\hat{H}(t)/\hbar}\hat{P}|\Psi_0\rangle) \tag{4.11}$$

This is what the permutation symmetry of state-descriptions requires of a Hilbert space description.

There is an analogous requirement for a classical phase space formalism. For a phase space formalism we take phase-space trajectories ζ to be the state-descriptions. The permutation symmetry of state-description thus requires that:

$$f(\zeta, t) = f(P\zeta, t) \tag{4.12}$$

In quantum theory, if we take both the quantum phase and generalized phase to be physically meaningless, then both symmetrized and parasymmetric vectors satisfy equation (4.11). But unsymmetrized vectors can also satisfy equation (4.11). $|\Psi_0\rangle$ can be an unsymmetrized state-vector and still satisfy equation (4.11), provided that it represents the same physical situation, is associated with the same proposition, as $\hat{P}|\Psi_0\rangle$. The two vectors need not be identical or only differ by a phase. Just as we can take a phase-space trajectory to be permutation symmetric, even though a permutation of particle labels leads to a distinct trajectory, we can also take an unsymmetrized state-vector to be permutation symmetric, even though permutation of particle labels leads to a distinct state-vector. In both cases, this is because the two distinct state-descriptions describe the same physical state.

Now what consequence does this have on statistics? Recall that in the classical case, we had to correct statistical ensembles for redundancies. We only included one state-description within the ensemble for each possible physical situation. The permutation symmetry of state-descriptions requires that we do the same for unsymmetrized quantum ensembles. In the classical case, such a correction only affects the extensivity of the entropy. In the quantum case, it affects more. To see this, consider the quantum canonical ensemble. For the quantum canonical ensemble, we cannot, as we did in the classical case, simply divide the partition function by $N!$ to eliminate redundant states. This follows from the definition of the quantum partition function given in equation (4.4). To correct the partition function, we must remove the redundant states from the sum in equation (4.4). We only want to include one state from all of those connected by a permutation transformation. But when we only take one state, we have the same number of states with the same energy eigenvalues as we do for the symmetric ensemble. The unsymmetrized states are different from the symmetric states, but there is the same number of each for a given energy eigenvalue. The sum in equation (4.4), and with it the partition function, is the same in both cases. The corrected unsymmetrized ensemble is therefore described by the thermodynamic equations of Bose-Einstein statistics.

We must, however, note that the corrected unsymmetrized ensemble and the symmetric ensemble are not equivalent in all respects. While the thermodynamic equations and distribution functions will be the same for both types of ensembles, the set of expectation values for each will differ. Once again the expectation value for some observable $\hat{O}$ for the canonical ensemble is:

$$\langle \hat{O} \rangle = \frac{\mathrm{Tr}(\hat{O}\rho)}{Q} \tag{4.13}$$

Even though Q is the same for both the corrected unsymmetrized ensemble and the symmetric ensemble, $\mathrm{Tr}(\rho\hat{O})$ can differ. Some expectation values are different for symmetric states than they are for unsymmetrized states. We cannot therefore say that the corrected unsymmetrized ensemble leads to quantum statistics, only that it does not lead to classical statistics.

So indistinguishability, understood in terms of permutation symmetry, does not necessarily lead to symmetrization, nor to the more general condition that confines state-vectors to subspaces that are irreducible under the action of the permutation group.[6] Permutation symmetry, however, does prevent classical statistics in the quantum formalism. So it is responsible

[6]Huggett (1999b) purposes an explanation of why state-vectors are confined to the symmetrized and parasymmetric subspaces. He argues that state-vectors are confined to these subspaces because all the symmetry groups in non-relativistic quantum mechanics—the Galilean group and the permutation group—only transform elementary particles by use of irreducible representations. As we have seen above, this condition confines state-vectors to symmetrized and parasymmetric subspaces. Huggett's explanation, however, is not meant to provide an argument for the necessity of any general condition that confines state-vectors to these subspaces. Rather, his stated aim is point out an important connection between the different symmetry groups in non-relativistic quantum mechanics, and thus provide some sort of unified understanding of symmetry in non-relativistic quantum mechanics.

for a change in statistics, even if it is not solely responsible for quantum statistics. We will discuss the implications of this in the section 4.6.

Now before we move on, we should note that all of this is only true of quantum theory as it stands. It still might be the case that if we supplement quantum theory in a certain way, then symmetrization does necessarily follow from the permutation symmetry of state-descriptions. Bacciagaluppi (2003) presents an example of how this can be so. He claims that symmetrization is a necessary consequence of the permutation symmetry of state-descriptions in the de Broglie-Bohm theory. The added structure of particle trajectories only allows for symmetrized states and quantum statistics.

In the de Broglie-Bohm theory we express the state-vector as $|\Psi\rangle = Re^{iS/\hbar}$. This is related to particle trajectories by the guidance equation, which (with a non-zero vector potential) is:

$$m_j \dot{\mathbf{x}}_j - q_j \mathbf{A}(\mathbf{x}_j) = \nabla_j S \tag{4.14}$$

Bacciagaluppi takes the permutation symmetry of the theory to require that "the velocity (average velocity) of particle 1 in a given configuration is equal to that of particle 2 in the configuration with the particles exchanged" (Bacciagaluppi, 2003, p. 5). We formally state this for spinless particles as:

$$\nabla[S(\mathbf{x},\mathbf{y},t) - S(\mathbf{y},\mathbf{x},t)] = 0 \tag{4.15}$$

This is a consequence of the permutation symmetry of state-descriptions. The permutation symmetry of state-descriptions requires that two points in configuration space connected by a permutation transformation represent the same physical state for the system at a given time. And since, according to the de Broglie-Bohm interpretation, a physical state gives the position and velocity of the particles in the system, this requires that equation (4.15) holds.

Bacciagaluppi shows that equation (4.15) implies that:

$$S(\mathbf{x},\mathbf{y},t) = S(\mathbf{y},\mathbf{x},t) + \gamma(\mathrm{mod} 2\pi) \tag{4.16}$$

where γ is equal to 0 or π for particles that move in three or more spatial dimensions. He also shows that equation (4.15) implies that:

$$R(\mathbf{x}, \mathbf{y}, t) = \pm R(\mathbf{y}, \mathbf{x}, t) \tag{4.17}$$

Bacciagaluppi goes on to generalize his results to particles with spin. Since $|\Psi\rangle = Re^{iS/\hbar}$, symmetrization is a necessary consequence of the equation (4.15). Symmetrization therefore follows from permutation symmetry given the extra structure of the particle trajectories of the de Broglie-Bohm interpretation of quantum mechanics. But once again, baring such additions, permutation symmetry does not imply symmetrization or the weaker condition that includes parasymmetric states.

4.4 Permutation Symmetry and the Topological Approach

Let us for the moment indulge in a digression and consider the relation between the permutation symmetry of state-descriptions and the topological approach.[7] In the topological approach we consider a reduced configuration space of a system of particles of the same kind, where we identify points in the full configuration space that differ by a permutation of particle labels, and remove points where particles coincide.[8] Here permutation symmetry is applied to the points of configuration space.

We quantize the system at a fixed time in reduced configuration space by associating a one-dimensional complex Hilbert space h_x with each point in configuration space. h_x is a fiber. The quantum state-vector is defined as a cross-section of the fiber-bundle. "That is, Ψ is assumed to be a single-valued function over the configuration space, whose function value $\Psi(x)$ at the point x is a vector in h_x" (Leinaas and Myrheim, 1977, p. 13). A quan-

[7]For details of the topological approach, see Laidlaw and DeWitt (1971) and Leinaas and Myrheim (1977). For a review of the philosophical implications of the approach, see French and Rickles (2003).

[8]There is some debate as to the legitimacy in this approach of removing points where the particles coincide (see Brown et al., 1999).

tum mechanical state at a given time is a function from configuration space into the complex numbers. As Leinaas and Myrheim (1977) point out in one of the seminal papers of the topological approach, when we quantize a system of particles in this way, only quantum statistics are possible for particles that move in three or more spatial dimensions. Parastatistics and classical statistics are ruled out. Quantum statistics appear to follow from both permutation symmetry and the dimensionality of space.

But the only reason parastatistics are ruled out in this approach is because this approach only assigns a one-dimensional complex Hilbert space to each point in the reduced configuration space. If we associate a complex Hilbert space of more than one dimension with each point, then it is possible to have parastatistics and even more exotic ambistatistics (Imbo et al., 1990). So in the topological approach, permutation symmetry does not inevitably lead to quantum statistics for particles traveling in three or more spatial dimension unless we make the apparently arbitrary stipulation that each fiber is a one-dimensional complex Hilbert space. Any justification of this stipulation appears to beg the question.

But there is another point we can raise concerning the topological approach. The permutation symmetry of state-descriptions does not force us to use a reduced configuration space. It only requires that permuted points represent the same physical situation, not that they are identical. It is perfectly consistent with the permutation symmetry of state-descriptions to take the standard configuration space and quantize that. This leads us to a standard Hilbert space description that includes unsymmetrized vectors. As we have just demonstrated, when we apply the permutation symmetry of state-descriptions to the Hilbert space formalism, we do arrive at a different statistics than classical statistics, but we do not necessarily end up with quantum statistics. So even if we accept that the topological approach leads to quantum statistics, it does so by making a stronger assumption than required by the permutation symmetry of state-descriptions.

4.5 Symmetrization as an Empirical Hypothesis

So where does all of this leave the symmetrization requirement? If it does not follow from the indistinguishability of quantum particles, how are we to understand it? In the absence of any better explanation, it seems best to view the symmetrization requirement as simply a separate empirical hypothesis that is part of our complete quantum theory. While symmetrization is a form of permutation symmetry, the fact that permutation symmetry manifests itself in term of symmetrization, instead of some weaker requirement, is just an empirical feature of quantum physics.

This is similar to a view put forward by French and Redhead (1988), which is further expounded by French (1989a). The idea that they put forward is that we can view symmetrization "as an initial condition in the specification of the situation" (French, 1989a, p. 443). As we have already pointed out, a state in a symmetric or antisymmetric subspace never dynamically evolves to a state outside that subspace since the quantum Hamiltonian is permutation symmetric. We can therefore regard symmetrization "as imposing a restriction on the possible states of the assembly such that certain of them are rendered inaccessible to the particles" (French, 1989a, p. 433).

While this is certainly one way to spell out what it means for symmetrization to be an empirical hypothesis, I prefer not to think of symmetrization simply as a constraint on initial conditions; for this leave the impression that unsymmetrized states are physically possible, even though they are unaccessible, or worse that symmetrization is simply an accidental feature of our actual world (although I do not claim that French or Redhead hold either of these views). Rather I prefer to think of symmetrization as simply an additional postulate of quantum theory, standing alongside other empirical hypotheses like the Schrödinger equation, dictating what is physically possible—that no vector that lies outside a symmetrized subspace of Hilbert space can represent a physically possible state of a system of particles of the same kind.

French and Redhead are careful not to endorse the view that symmetrization is simply an empirical hypothesis. As we mentioned above on page 49, they put it forward in order to show that symmetrization does not necessarily imply that quantum particles lack identity. We do not need to hold that the reduction of states due to permutation symmetry follows from the fact that quantum particles are not objects, for it is possible to hold that the reduction in states simply follows from the initial conditions of the system, which confines the evolution of the system to a specific subspace of Hilbert space.

French and Redhead seem willing to grant that there are better ways to view the symmetrization requirement that do not take it as a "brute fact" (to use French and Rickles (2003) description of the view). But the symmetrization requirement is no more brutish than the other empirical hypotheses of quantum theory. Granted we cannot explain it solely by considering the permutation symmetry or ontology of quantum theory, but we should not feel we need to. There are plenty of feature of quantum mechanics that are only justified by their empirical adequacy, not by any philosophical argument.

Of course if we are to claim that symmetrization is simply an empirical hypothesis of quantum theory, then it has to have clear empirical content. French and Rickles claim that there is no direct evidence of symmetrization or permutation symmetry in general, and raise the possibility that symmetrization is "a kind of 'free-floating' principle" that is not required experimentally (French and Rickles, 2003, p. 231). In order to explain what they mean by direct evidence, they draw upon the work of Kosso (2000).[9] As Kosso puts it, in order to have direct evidence that some physical property or law is left invariant by a symmetry transformation, one "must observe that the specified transformation has taken place, and one must observe that the specified invariant property is in fact the same, before and after" (Kosso, 2000, p. 86). Galileo's ship experiment provides the paradigmatic example of such a direct observation. In the ship experiment we can observe two subsystems (one on the ship and one on the shore) that are related by an active transformation (a Galilean boost). Since the two subsystems have the same internal evolution, we have direct evidence that the dynamical laws governing the internal

[9]For a critical assessment of Kosso's arguments, see Brading and Brown (2004).

evolution of both systems are invariant with respect to the transformation.

There is no analogous experiment that we can conduct to directly observe the validity of symmetrization. This is because there is no way that we can observe that a permutation transformation of two particles of the same kind has taken place. In Galileo's ship experiment, we can observe that the ship is moving with respect to the shore, and therefore establish by a change in relative velocities that the transformation has taken place. But since a permutation of two particles of the same kind leaves all expectation values of the system unchanged, there is no change in any observable property that allows us to establish that the permutation transformation has taken place. We therefore cannot directly observe the permutation symmetry associated with symmetrization.

However, as Kosso points out, even if there is no direct evidence of a symmetry, there can be indirect evidence. We can still observe other consequences of the symmetry, like conservation laws. Of course there are no conservation laws associated with symmetrization, but there is quite a bit of other indirect evidence. The existence of quantum statistics provides us with clear empirical evidence for the symmetry. And since symmetrization affects the set of expectation values for individual pure states, there is also plenty of evidence for symmetrization outside of statistical mechanics. So while symmetrization is not directly observable in the way Kosso outlines, it still has empirical content and clear empirical justification. Once again, it seems best to view the symmetrization requirement as an empirical hypothesis, and not as a requirement that necessarily follows from the indistinguishability of quantum particles.

4.6 The Difference between Classical and Quantum Statistics

We have shown that permutation symmetry is only a necessary, not a sufficient, condition for symmetrization and quantum statistics. But we have also shown that permutation symmetry ensures that quantum theory is not described by classical statistics. So even though permutation symmetry does not entail quantum statistics, we still might think that a difference in permutation symmetry is responsible for the difference in statistics between classical physics and quantum physics. This conclusion appears to be supported not only by the derivations of statistics in Hilbert space, but also by similar derivations in phase space. Now if this conclusion is true, and a difference in statistics is due to a difference in permutation symmetry, then classical physics must lack permutation symmetry, despite what wrote in the last chapter.

In order to evaluate this conclusion, let us look at some of these Hilbert space and phase space derivations in greater detail. Let us start with the Hilbert space derivations. Once again let us consider the case of an ideal gas. Instead, however, of considering the canonical ensemble, as we have throughout, let us consider the microcanonical ensemble, where both the number of particles and total energy of the system are fixed. This will allow for an easier comparison to the standard phase space derivations, which we consider next.

We take a quantum ideal gas to be a collection of N non-interacting particles in a box of infinite potential.[10] The Hamiltonian $\hat{H}$ within the box is:

$$\hat{H} = \frac{-\hbar^2}{2m} \sum_{j=1}^{N} \nabla_j^2 \tag{4.18}$$

For simplicity, we will suppress internal degrees of freedom. For further simplicity, we will assume that the system occupies a large volume, so that the energy spectrum of a single particle in the box is nearly continuous. We therefore can divide the energy spectrum of a single particle into cells, so that in each cell j there are g_j single-particle energy eigenstates with a single-particle energy between ϵ_j and $\epsilon_j + \Delta_j$. Since the particles do not interact, we can determine the statistical mechanical behavior of the N-particle system from its single-particle energy distribution $\{n_j\}$, where $\{n_j\}$ assigns n_j particles to cell j, and where $N = \sum_j n_j$ and $E = \sum_j \epsilon_j n_j$. The statistical mechanical properties of the system are given by the total number G of possible

[10]For the complete details of this calculation see Tolman (1938, ch. 10).

N-particle states of the system that have a single-particle energy distribution $\{n_j\}$. The number G will differ depending upon the symmetrization restrictions we place on the N-particle states of the system.

For an unsymmetrized gas, where there are no symmetrization restrictions on the ensemble, the number of N-particle unsymmetrized states G_{U} that have a distribution $\{n_j\}$ is:

$$G_{\mathrm{U}} = \prod_j \frac{N!}{n_j!} g_j^{n_j} \tag{4.19}$$

For a symmetric gas, where the states of the ensemble are confined to the symmetric subspace of Hilbert space, the number of N-particle symmetric states G_{S} that have a distribution $\{n_j\}$ is:

$$G_{\mathrm{S}} = \prod_j \frac{(n_j + g_j - 1)!}{n_j!(g_j - 1)!} \tag{4.20}$$

And for an antisymmetric gas, where the states of the ensemble are confined to the antisymmetric subspace of Hilbert space, the number of N-particle antisymmetric states G_{A} that have a distribution $\{n_j\}$ is:

$$G_{\mathrm{A}} = \prod_j \frac{g_j!}{n_j!(g_j - n_j)!} \tag{4.21}$$

To be clear, the unsymmetrized ensemble is uncorrected and violates even the most basic requirement of permutation symmetry. Such an ensemble includes all states connected by a permutation transformation. If we correct the ensemble by removing the redundant states connected by a permutation transformation, then the number of states G that have a distribution $\{n_j\}$ is equal to G_{S}, not G_{U}.

Assuming that each state in the microcanonical ensemble is equally probable, G is proportional to the probability that the system is in a state with a distribution $\{n_j\}$. We can therefore calculate the most probable distribution $\{\overline{n}_j\}$ by calculating the maximum value of G. By use of variational calculus and the Stirling approximation, we arrive at the following results. For an unsymmetrized gas, the most probable distribution is the Maxwell-Boltzmann distribution function:

$$\overline{n}_j = g_j e^{-(\epsilon_j - \mu)/kT} \tag{4.22}$$

where μ is the chemical potential, k Boltzmann's constant, and T the temperature. For a symmetric gas, the most probable distribution is the Bose-Einstein distribution function:

$$\overline{n}_j = \frac{g_j}{e^{(\epsilon_j - \mu)/kT} - 1} \tag{4.23}$$

And for an antisymmetric gas, the most probable distribution function is the Fermi-Dirac distribution function:

$$\overline{n}_j = \frac{g_j}{e^{(\epsilon_j - \mu)/kT} + 1} \tag{4.24}$$

A classical distribution function thus follows from a lack of permutation symmetry. Uncorrected ensembles of unsymmetrized states are described by classical statistics. Permutation symmetry is therefore responsible for the difference in statistics in the Hilbert space formalism.

We appear to reach a similar conclusion when we consider the phase space analogs to these quantum microcanonical derivations. While in the Hilbert space derivations, equations (4.19)–(4.21) simply give the number of unsymmetrized and symmetrized states that are consistent with a given distribution, they are analogous to combinatorial equations that give the number of ordered and unordered distributions of particles in phase space cells. To see this, take the single-particle μ-space associated with each particle in a classical ideal gas. Divide the μ-space into cells of volume h^3. Let g_j be the number of cells contained in the volume whose points have a single-particle energy between ϵ_j and $\epsilon_j + \Delta_j$. Further let n_j be the number of particles within the g_j cells. Just as with the quantum microcanonical ensemble, the distribution function of this classical microcanonical ensemble follows from the number of states G associated with a single particle energy distribution $\{n_j\}$.

Just as above, G is given by equations (4.19)–(4.21). However these equation now have a different meaning. For a classical microcanonical ensemble, they do not give the number of unsymmetrized, symmetric, and antisymmetric states that have a distribution $\{n_j\}$, for there are no such states in phase space. Rather, in this phase space derivation, equations (4.19)–(4.21) express the number of complexions or ways we can distribute the particles in the gas among the cells of μ-space so that the single particle energy distribution is $\{n_j\}$. Equation (4.19) gives the number of ordered ways we can distribute the particles, where we arrive at a distinct complexion when the particles are permuted. Equation (4.20) gives the number of unordered ways, where we arrive at the same complexion when the particles are permuted. And equation (4.21) gives the number of unordered ways we can distribute the particles so that there is no more than one in each cell.

As above, equation (4.19) leads to a Maxwell-Boltzmann distribution function, equation (4.20) to a Bose-Einstein distribution function, and equation (4.21) to a Fermi-Dirac distribution function. So in this phase space derivation, we have classical statistics for ordered distributions and quantum statistics for unordered distributions. This difference between ordered and unordered distributions is really a difference in permutation symmetry. Thus permutation symmetry, once again, appears to be responsible for the difference in statistics.

Of course these phase space derivations are related to the derivations of Boltzmann, Planck, Bose, and Einstein, which we have discussed in the last chapter. The phase space derivation of classical statistics is similar Boltzmann's derivation, given in (Boltzmann, 1877) and popularized by Ehrenfest and Ehrenfest (1959, p.26–31). The phase space derivation of quantum statistics is closely related to those given by Planck (1900), Bose (1924), and Einstein (1924). The fact that a difference in permutation symmetry leads to a difference in statistics in these phase space derivations, and the previous Hilbert space derivations, no doubt have lead many to believe that the difference between the statistics of quantum and classical physics follows from a difference in permutation symmetry.

And given the derivations above, it does seem that the view is correct. For in both the classical and quantum derivations, it is a difference in permutation symmetry that leads us from classical to quantum statistics. In the Hilbert space derivation, it is the difference between unsymmetrized and symmetrized states that is responsible for the difference in statistics, while in the phase space derivation, it is the difference between ordered and unordered distributions. Once again, this conclusion appears to conflict with our claim of the previous chapter that classical physics is just as permutation symmetric as quantum physics.

But there is no conflict, as we see when we look again at the phase space derivation of classical statistics.[11] The combinatorial method applied above is of course important in the historical development of statistical mechanics, but its derivation of classical statistics is misleading. If we switch to a more contemporary derivation, it is easy to see this. As we have noted above in our discussion of the classical canonical ensemble, the probability of a distribution is proportional to the volume in phase space occupied by states with that distribution. So instead of considering the number of complexions associated with a distribution, what we need to consider is the volume associated with a distribution.

This is easy enough to calculate for a microcanonical ensemble in phase space. The volume V_Γ associated with a distribution $\{n_j\}$ is given by the following equation:

$$V_\Gamma = \prod_j \frac{N!}{n_j!}\, a_j^{n_j} \tag{4.25}$$

where a_j is the volume in μ-space of single-particle states with an energy between ϵ_j and $\epsilon_j + \Delta_j$ and n_j is the number of particles with a single-particle state within volume a_j. Equation (4.25) is analogous to equation (4.19). They are both of the same form, and both have a maximum value for the Maxwell-Boltzmann distribution function.

Equation (4.25), like equation (4.19), lacks permutation symmetry. There are distinct states in V_Γ that only differ by a permutation transformation. To impose permutation symmetry, we need to divide out the redundant states, as we did for the canonical ensemble in the last chapter. We do that here by dividing the volume V_Γ by $N!$, so that the permutation symmetric expression

[11] The following conclusions follow those presented by Saunders (forthcoming a).

is:

$$V_\Gamma = \prod_j \frac{a_j^{n_j}}{n_j!} \tag{4.26}$$

This equation, like equation (4.25), has a maximum value for the Maxwell-Boltzmann distribution. As with the classical canonical ensemble, the division by $N!$ only affects the expression of the entropy.[12] It leaves the rest of the statistical description unchanged.

So we have a divergence here between the two phase space derivations. In the combinatorial approach, permutation symmetry leads to quantum statistics, but in the phase space volume approach, permutation symmetry still leads to classical statistics. It is clear what is responsible for this difference: it is the use of μ-space cells of definite size in the combinatorial approach.

We can demonstrate this simply by consider the limit where the volume of the cells goes to zero. In this limit $g_j >> n_j$ and equations (4.20) and (4.21) both become:

$$G_S \approx G_A \approx \prod_j \frac{g_j^{n_j}}{n_j!} \tag{4.27}$$

which is analogous to equation (4.26). So in this limit, permutation symmetry does not affect the statistics (except for the extensivity of entropy) within the combinatorial derivation. It is only when the cells are of some definite size that permutation symmetry leads to quantum statistics.

When the cells are of dimension h^3, there is a direct correspondence between unordered distributions among cells in phase space and the number of symmetrized quantum states in Hilbert space and a correspondence between the number of ordered distributions among cells and unsymmetrized states. These correspondences lead to quantum distribution functions in the phase space formalism for unordered distributions, and a difference in the statistics of non-permutation symmetric and permutation symmetric ensembles in the phase space formalism. But when considered solely in the context of a classical formalism, the cells are simply a mathematical fiction that allow for the application of combinatorics. It is only the limit, where the volume of the

[12]See Huang (1963, p. 151–155) for details of how the $N!$ correction leads to an extensive entropy for the classical microcanonical ensemble.

cells goes to zero, that has any physical significance. And in this limit, there is no longer a correspondence between cells and quantum states, and no difference (except for an extensive entropy) in the statistics of non-permutation symmetric and permutation symmetric ensembles. In the proper derivation of classical statistics in phase space, the derivation that does not posit cells of definite size, permutation symmetry does not lead to a difference in statistics.

Planck noted early on the role of the definite cell size in the derivation of his radiation law (see Jammer (1966, p. 53–54) and Kuhn (1978, p. 131)). He took this as an indication of the importance of his constant h, which required some sort of explanation. As he relays in his Nobel lecture, his constant "was completely indispensable for obtaining the correct expression for entropy—since only with its help could the magnitude of the "elementary regions" or "free rooms for action" of the probability, decisive for the assigned probability consideration, be determined" (Plank, 1967). Eventually Planck realized that his constant required a break with classical physics. As he writes, his constant was "elusive and resistant to all efforts to fit it into the framework of classical theory" (Plank, 1967). Bose, in his derivation of Planck's law, also noted that the phase space cells must have a definite size of h^3, and that "[n]othing definite can be said about the method of dividing phase space in this manner" (Bose, 1924). For both Planck and Bose, cells of definite size are essential to their derivations, but they recognize (or came to recognize) that such cells have no classical justification. They are simply quantum additions to the classical theory.

So what can we take from all of this? Well the important point is that permutation symmetry by itself is not responsible for the difference in statistics between classical and quantum physics. It is true that, in the quantum formalism, permutation symmetry does lead to a different statistics from what which we find in classical physics. But this does not indicate a difference in permutation between classical and quantum physics, only a difference in the effects of permutation symmetry in the two theories. Permutation symmetry in the classical formalism only affects the extensivity of the entropy, while permutation symmetry in quantum formalism affects the entire statistics.

Now why the difference? Why does permutation symmetry have a different effect in the two theories? Well it clearly has something to do with the difference between the state-spaces of each theory. We can see this simply by pointing out that there is only a correspondence between the two state-spaces when we divide phase space into cells of definite volume, and that it is only this, essentially quantum, addition that allows for permutation symmetry to affect the statistics in classical physics.

Saunders (forthcoming b) argues that the essential point is that the different state-spaces lead to different probabilistic measures. For phase space we have a continuous measure, but for Hilbert space we have a discrete measure. Permutation symmetry only has an effect on the discrete measure. This is why, when we impose a discrete measure on phase space by dividing into cells of fixed volume, permutation symmetry leads to quantum statistics.

But once again, the fact that permutation symmetry in quantum physics leads to different statistics than in classical physics does not indicate that classical physics lacks permutation symmetry. There is no conflict with our claim of the previous chapter that classical physics possesses permutation symmetry.

Chapter 5

Permutation Symmetry and Anti-Haecceitism

We rejected the idea that permutation symmetry follows from the fact that particles are not objects. But this does not leave permutation symmetry without any sort of philosophical foundation or philosophical implications. For while the permutation symmetry of particles is not due to a lack of particle identity, it is related to an important philosophical thesis—anti-haecceitism.

5.1 Anti-Haecceitism Defined

Let us start with with a review of the philosophical view that is anti-haecceitism. The modern debate between haecceitism and anti-haecceitism has its roots in the papers of Kaplan (1975) and Adams (1979). It is, however, David Lewis that has given the most precise definition of the two positions (Lewis, 1986, sec. 4.4). Lewis defines anti-haecceitism as a supervenience thesis: What a possible world represents *de re* supervenes upon its qualitative character. Haecceitism is the denial of this supervenience thesis.

When we talk about representations *de re* of a possible world, all we are talking about are the facts and properties that a possible world assigns to various objects. A possible world can represent *de re* all sorts of things about all sorts of objects. Some possible worlds represent *de re* that Al Gore wins the 2000 presidential election (let us keep our examples up to date). For example in one possible world there are no hanging chads, and the Democrats win Florida.[1]

There are clear qualitative differences between this possible world and the actual world. One is that this possible world does not have hanging chads, while the actual world does. These two worlds differ in their qualitative character. The qualitative character of a possible world is the pattern of instantiation of qualitative properties (including relations).[2] We can describe, in a possibly infinite description, the qualitative character of a world by using only qualitative properties and bound variables, without using proper names or non-qualitative properties.

The important question for the debate between haecceitism and anti-haecceitism is whether two worlds can differ in what they represent *de re* of Al Gore, or any other object, without differing in their qualitative character. To illustrate the dispute, let us consider another example. Consider a world that consists of a succession of epochs, where the events in each epoch are the same. In such a world the same history of events keeps repeating over and over again. In the seventeenth epoch there is a presidential election where the Democrats lose Florida in a extraordinary close race and in the 137th epoch there is another presidential election that is qualitatively identical to one in the seventeenth epoch, and to one in every other epoch.

Now the question is whether two possible worlds that give the same qualitative description of this reoccuring state of affairs can differ in how they represent *de re* Al Gore. Even though they have the same qualitative character, can one of them represent *de re* that Al Gore lost the election in the the seventeenth epoch, while the other represents *de re* that he lost the election in the 137th epoch? If we accept haecceitism, then the answer is yes; if

[1]For those who might not recall, hanging chads are partially dislodged bits of paper that fill the holes in punch card ballots. Several ballots were voided in the 2000 presidential election because these pieces of paper were not fully dislodged. The margin of victory was so slim in Florida, that, had these ballots not been voided, they could have delivered the election to Al Gore.

[2]As we discussed above on page 12, a qualitative property is one that does not imply or refer to particular objects.

anti-haecceitism, then no.

While we have defined haecceitism above as the rejection of anti-haecceitism, we can refine our definition to account for different strengths of haecceitism (Lewis, 1986, p. 239–243). Moderate haecceitism holds that the qualitative character of a possible world can constrain, if not fully determine, representations *de re*. This is haecceitism coupled with essentialism. How a world represents *de re* Al Gore might not supervene on the qualitative character of that world, but that does not mean that some world can represent *de re* Al Gore as an egg. There are certain essential properties that Al Gore must possess. Extreme haecceitism claims that there are no qualitative constraints on how a world represents *de re*: Al Gore could be an egg.

In our definition of anti-haecceitism we have made use of possible worlds. But this use of possible worlds does not mean that anti-haecceitism implies modal realism. All that is implied is a possible world semantics, not a specific theory of the nature of possible worlds. Possible worlds need not be real concrete objects, but can be sets of abstract propositions or what have you.

We must also keep in mind that anti-haecceitism is not simply the denial of haecceities (Lewis, 1986, p. 225). As we mentioned in chapter 2, an haecceity is a non-qualitative property of being identical to a particular object, such as "identical to Al Gore" or "$= a$". One way to violate anti-haecceitism is by use of haecceities: Two worlds can have the same qualitative character, yet still differ in what they represent *de re*, if they have different instantiations of non-qualitative haecceities. One world can represent *de re* that Al Gore lost the election in the seventeenth instead of 137th epoch because the individual in the seventeenth epoch possesses the haecceity associated with Al Gore.

In chapter 2, we expressed serious doubts about the independent existence of haecceities. If we are to consider haecceities at all, it is only as innocuous gerrymandered properties that are incapable of any interesting metaphysical work, such as individuation or representation *de re*. But even though we have good reason to reject haecceities, anti-haecceitism is still consistent with a metaphysically robust version of haecceities (one where haecceities actually serve some independent role in representations *de re*), so long as two worlds cannot differ in their instantiations of haecceities without differing qualitatively.

And just as anti-haecceitism does not entail the rejection of haecceities, haecceitism does not entail the acceptance of haecceities. It is true that if we accept haecceities, then we have good reason to accept haecceitism; for haecceities can lead to and explain haecceitistic differences.[3] But we can also explain haecceitistic differences without reference to haecceities. As Lewis notes, we can explain haecceitistic differences by use of overlapping possible worlds (although Lewis ultimately rejects overlap):

> Just as two duplicate strings may share a dot though one puts it in the middle and the other puts it at one end so it might happen that two duplicate worlds share an individual though one puts him in the seventeenth epoch and the other puts him in the 137th; and thereby those two worlds could differ in what they represent *de re* about which epoch is his. (Lewis, 1986, p. 228)

Here we do not explain haecceitistic differences by use of haecceities, but rather by use of common parts of overlapping worlds, which might have nothing to do with haecceities.

Both Gordon (2002, p. 409) and Hoefer (1996, p. 15) deny that we can have haecceitistic differences without haecceities. As Gordon puts it, without haecceities "we cannot affirm that haecceitistic differences between worlds are possible, because every difference is a qualitative difference" (Gordon, 2002, p. 409). But this is not so. If we regard cross-world identity simply as numerical identity, then we can take this identity to be primitive in the way we outlined in chapter 2. If we regard cross-world identity in this way, then there is no need to appeal to haecceities. An object is the object it is in two different worlds simply because it is identical to itself in both worlds. There is no need to cite any further facts that constitute the cross-world identity of the object. When we accept such primitive cross-world identity, we can have overlap in the way in which Lewis illustrates, and with it haecceitistic

[3] It is certainly the acceptance of haecceities that leads Adams (1979) to accept haecceitism.

differences, without haecceities.

It is true that one of the main motivations behind haecceitism is the belief that cross-world identity is numerical identity, as opposed to some other relation; for numerical identity need not be constrained in any way by qualitative properties. And if it is not so constrained, then it can lead to haecceitistic differences of the sort that Lewis presents.[4] But, once again, this sort of reasoning need not force us to accept haecceities; for we do not need haecceities to establish numeric cross-world identity. We can take such identity as primitive.

We will consider some further arguments for haecceitism shortly. But let us now consider the philosophical motivations for anti-haecceitism. One of the big ones is Humean supervenience. Humean supervenience is the thesis that everything supervenes on the spatiotemporal arraignment of local natural qualitative properties. This of course implies that representations *de re* supervene on the qualitative character of possible worlds, and therefore implies anti-haecceitism.

Another motivation for anti-haecceitism comes from counterpart theory. Counterpart theory denies that cross-world identity is numerical identity, and therefore denies the main motivation behind haecceitism. Counterpart theory analyzes cross-world identity by use of a counterpart relation, where an object is another object's counterpart if it qualitatively resembles that object in a relevant way. The actual Al Gore is not numerically identical to the possible Al Gore that won the presidential election in the possible world that lacks hanging chads. The possible Al Gore is simply the counterpart of actual Al Gore. The possible Al Gore qualitatively resembles the actual Al Gore in a way that is relevant to our discussion of the 2000 presidential election.[5] Since representations *de re* are determined by qualitative similarity between counterparts, counterpart theory implies anti-haecceitism. Everything that a possible world represents *de re* must supervene upon its

[4]It is a belief in numeric cross-world identity that leads Kaplan (1975, sec. 4) to accept haecceitism.

[5]Lewis gives a detailed presentation and defense of counterpart theory in (Lewis, 1973) and (Lewis, 1986, ch. 4).

qualitative character.[6]

While Humean supervenience and counterpart theory are both popular philosophical views, they are not without critics. Quantum theory itself raises doubts about Humean supervenience. Whether we consider particles or fields, it seems likely that some of the properties of a quantum system cannot be localized to a spatiotemporal region, and that therefore not everything supervenes upon local natural qualitative properties. And while counterpart theory remains one of the more successful theories about representations *de re*, it has come under increasing attack (for example see Merricks (2003)).

But even if we do not count ourselves amongst the supporters of these views, we still might accept anti-haecceitism; for anti-haecceitism does not imply Humean supervenience or counterpart theory. And we still can motivate anti-haecceitism independently of these views. For the key idea behind anti-haecceitism is that the world is essentially a qualitative place. As van Fraassen put it, "*everything*, the world, can be completely described by entirely general propositions" (van Fraassen, 1991, p. 465). Van Fraassen labels this position semantic universalism. Under a suitably interpretation of what general propositions of everything are, semantic universalism is equivalent to anti-haecceitism, a point which van Fraassen (1991, p. 464, fn. 14) himself notes. The central idea behind semantic universalism and anti-haecceitism—that everything can be described qualitatively—is both initially plausible and attractive in its own right. We can therefore motivate anti-haecceitism independently of Humean supervenience and counterpart theory.

This idea that the world is qualitative is also the primary motivation behind the principle of identity of indiscernibles (PII). But we now see that PII is a stronger principle than required by this belief. For while PII implies anti-haecceitism, anti-haecceitism does not imply PII (Lewis, 1986, p. 224). If PII is true for possible worlds, then all distinct worlds, regardless of what they represent *de re*, must differ in their qualitative character. This, of course, implies anti-haecceitism. But anti-haecceitism does not forbid, as PII does,

[6] Lewis does discuss the possibility of a non-qualitative counterpart relation, which would allow for haecceitistic differences. But he rejects such a relation as utterly mysterious (Lewis, 1986, p. 229–230), a view with which I agree.

indiscernible objects within a world or indiscernible possible worlds.

Many balk at the idea of indiscernible possible worlds. They take anti-haecceitism to be a version of PII applied to possible worlds. Anti-haecceitism, however, is consistent with qualitatively indiscernible possible worlds so long as two indiscernible worlds give the same representations *de re*. But some still might deny the existence of distinct possible worlds with the same qualitative character. Such distrust of indiscernible possible worlds is justified if we hold that possible worlds are something like sets of propositions. But if we are modal realists, and possible worlds are real concrete objects, then we might say that there are just as many possible worlds as there are, and some of them might be indiscernible, like other concrete objects.[7] Anti-haecceitism itself remains neutral on this issue, as we shall be in this discussion. But the important point is that, even if we think the world is qualitative, we need not accept PII, only the weaker principle of anti-haecceitism.

5.2 Kripkean Worries

Even if we are attracted to the idea that that we can describe everything in the world qualitatively, some of our intuitions still might suggest that anti-haecceitism is false. Kripke summarizes some of these doubts with a putative counterexample to anti-haecceitism:

> Two ordinary dice (call them die A and die B) are thrown, displaying two numbers face up. For each die, there are six possible results. Hence there are thirty-six possible states of the pair of dice, as far as the numbers shown face-up are concerned, though only one of these states corresponds to the way the dice actually will come out. We all learned in school how to compute the probabilities of various events (assuming equiprobability of the states.) ... Now in doing these school exercises in probability, we were in fact introduced at a tender age to a set of (miniature)

[7]Maudlin (1996) takes the indiscernibility of possible worlds to be a significant defect of modal realism. Lewis (1996), however, does a good job of defending his view on this point.

> 'possible worlds'. The thirty-six possible states of the dice are literally thirty-six 'possible worlds', as long as we (fictively) ignore everything about the world except the two dice and what they show (and ignore the fact that one or both dice might not have existed). Only one of these mini-worlds—the one corresponding to the way the dice in fact come up—is the 'actual world', but the others are of interest when we ask how probable or improbable the actual outcome was (or will be). ... Nor should any school pupil receive high marks for the question 'How do we know, in the state where die A is six and die B is five, whether it is die A or die B which is six? ... The answer is, of course, that the state (die A, 6; die B, 5) is *given* as such (and distinguished from the state (die B, 6; die A, 5)). ... The 'possibilities' simply are not given purely qualitatively (as in: one die, 6, the other, 5). If they had been, there would have been just twenty-one distinct possibilities, not thirty-six. (Kripke, 1980, p. 16–17)

Here the mini-worlds (die A, 6; die B, 5) and (die B, 6; die A, 5) seem to differ in their representations *de re* of the dice without differing in their qualitative character. This haecceitistic difference leads to different statistical assumptions about the behavior of the dice. And so we have a putative counterexample to anti-haecceitism, and a pretty common one at that.

Of course there are striking similarities between Kripke's example and the statistics of classical particles. We can easily substitute particles for dice and single-particle states for the numbers given by the dice. This makes Kripke's example even more important to our study, for it appears to present a very powerful intuition that calls not only anti-haecceitism into question, but also the permutation symmetry of classical physics. For if a permutation of dice leads to a distinct representation *de re*, then, similarly, a permutation of particles should lead to a distinct physical situation. We will come back to the connection between permutation symmetry in physics and anti-haecceitism in the next section. But right now let us focus on the specifics of Kripke's example.

We look once again to Lewis for the dissolution of this counterexample (Lewis, 1986, p. 225–227). As Lewis points out, anti-haecceitism only applies to maximally specific possible worlds. Less than maximally specific possible worlds, such as Kripke's mini-worlds, abstract away from many of the qualitative details of a possible world. When we ignore the full qualitative character of possible worlds and consider mini-worlds, these mini-worlds will differ in what they represent *de re* without explicitly differing in their qualitative properties. But it is only because we are not considering the full qualitative character of the possible worlds that we have this apparent haecceitistic difference. Full descriptions of possible worlds will not differ in their representations *de re* without differing in their qualitative character.

Consider Kripke's dice. When we fill in the complete description of the possible worlds describing the dice, we include the trajectories of the dice (and enough further information to establish a frame of reference for the position of the dice). Die A is then the die that has a continuous trajectory starting from some fixed point, and die B is the die that has a continuous trajectory starting from some other fixed point. (die A, 6; die B, 5) and (die B, 6; die A, 5) give different representations *de re* of the dice. But this does not lead to an haecceitistic difference because the maximally specific possible worlds related to these descriptions differ in their qualitative character: they describe different sets of single-particle trajectories. Of course it does not pay to consider maximally specific possible worlds when calculating probabilities in the classroom or the casino. Kripke's abbreviated description is all that is needed. But these abbreviated descriptions do not imply haecceitism.[8]

In his example, Kripke makes use of the proper names "A" and "B". Let us follow Lewis (1986, p. 222–223) and call the use of proper names in descriptions of possible worlds "Kripkean specification". Kripkean specification is not only a legitimate way to describe possible worlds, it is often the only way to give a finite and manageable description of a possible world. We almost always use Kripkean specification. We describe a possible world where

[8]Lewis (1986, p. 230–235) also discusses other putative examples of haecceitistic differences. He is able to deal with these examples by admitting that not all possibilities are possible worlds.

Al Gore won the 2000 election by using the proper name "Al Gore", and not by giving a purely qualitative description of a world where there is a person qualitatively similar to the actual Al Gore that wins something qualitatively similar to the actual 2000 election.

But even though Kripkean specification uses proper names, it does not conflict with anti-haecceitism. While we have made use of purely qualitative descriptions in order to elucidate the qualitative character of a possible world, we can still use proper names in order to describe possible worlds, so long as we deny that their use can lead to any haecceitistic differences. We deny that (die A, 6; die B, 5) and (die B, 6; die A, 5) specify maximally specific worlds that are qualitatively indiscernible, but that differ in their representations *de re* of the dice. The reason we might have the intuition that these descriptions do present us with worlds that differ haecceitistically is because we often use similar statements to describe less than maximally specific possible worlds. These less than maximally specific possible worlds can differ in there representations *de re* without explicitly differing in their qualitative character. But, as we have argued above, these differences are not truly haecceitistic differences. They simply follow from the fact that we have abstracted away from the full qualitative character of possible worlds.

We can accept anti-haecceitism and use proper names (or singular terms) not only in informal descriptions of possible worlds, but also in a formal modal semantics. Carnap (1950) provides us with a simple example of how we can do this. Carnap considers a set of objects and the set of monadic properties that these objects can possess. He defines two types of descriptions: state descriptions and structure descriptions. State descriptions assign a maximal set of properties to a domain of objects using predicates and singular terms. Structure descriptions are sets of isomorphic state descriptions. Two state descriptions are isomorphic if and only if some permutation of singular terms in one description gives the other. If isomorphic state descriptions all give the same representations *de re*, that is if all representations *de re* are captured by structure descriptions, then the system is anti-haecceitistic, even though structure descriptions are defined by use of state descriptions, and state description include singular terms. This is not peculiar to Carnap's approach. Van Fraassen (1991, ch. 12, sec. 4.2) provides us with another example of a modal semantics that is anti-haecceitistic even though it makes use of singular terms. In van Fraassen's formal semantics, he specifies a domain by use of singular terms. And as he demonstrates, the models of his semantics are anti-haecceitistic as long as they obey a type of permutation symmetry.

These two examples not only demonstrate that we can accept anti-haecceitism and still use proper names or singular term to describe possible worlds, they also demonstrate the connection between permutation symmetry and anti-haecceitism. Both examples use permutation symmetry to eliminate apparent haecceitistic differences.

As Kaplan writes:

> When we construct a model of something, we must distinguish those features of the model which represent features of that which we model, from those features which are intrinsic to the model and play no representational role. The latter are *artifacts of the model.* For example, if we use string to make a model of a polygon, the shape of the model represents a feature of the polygon, and the size of the model may or may not represent a feature of the polygon, but the thickness and three-dimensionality of the string is certainly an artifact of the model. (Kaplan, 1975, p. 722)

Apparent haecceitistic differences are like the thickness of the string. The use of Kripkean specifiers, proper names, and singular terms lead to distinct descriptions that are related by a permutation transformation. These descriptions make apparently haecceitistic distinctions. But any such distinctions are simply artifacts of model that follow from the use of Kripkean specifiers. They do not actually represent haecceitistic differences. And we eliminate such distinctions by use of permutation symmetry.

5.3 Permutation Symmetry and Anti-Haecceitism

It should be clear at this stage that permutation symmetry serves an analogous role in physics. In a similar way as we did with Kripke's dice, we can associate with each state-description of a physical system a maximal possible world, where this world represents *de re* the system evolving as described by the state-description. For example, we can associate with each phase space trajectory describing a system a possible world, where this possible world represents *de re* the particles in the system with position and momentum at a given time as described by the phase space trajectory.

If two state-descriptions represent different physical situations, then each is associated with a possible world that gives a different representation *de re* of the system. Each of the worlds represents *de re* a different possible evolution of the system.

For simplicity, we will take the possible worlds associated with state-descriptions to consist of nothing but the system (and anything else required in order to establish a spacetime background and a coordinate system). And for further simplicity we will ignore the possibility of indiscernible possible worlds.[9]

Given this connection to possible worlds, it is easy to see the connection between the permutation symmetry of state-descriptions and anti-haecceitism. Let us first consider the case where the permutation symmetry of state-descriptions does not apply. In this case two state-descriptions that differ solely by a permutation of particles can represent distinct physical situations. That is they are each associated with worlds that differ in their representations *de re*. We have already pointed out on several occasions that

[9]If we wish, however, we can be more general. We can associate with each state-description a set of possible worlds instead of just one sparse world, where each possible world in the set represents *de re* the same system evolving in the same way. This generalization allows us to accept indiscernible worlds. It also allows us to acknowledge that there are several different possible worlds that represent *de re* the system in the same way. Many worlds will give the same representation *de re* of the system, but differ in their representations *de re* of other objects. For example, two worlds can represent *de re* the same evolution of a gas, but differ in whether Al Gore won the 2000 presidential election. In what follows we can substitute sets of possible worlds for each single sparse world without affecting any of our conclusions.

two state-descriptions that differ solely by a permutation transformation do not differ in their qualitative description of the system. For example two phase space trajectories connected by a permutation transformation do not differ in their qualitative description of single particle trajectories, but only differ in which particles they assign to which trajectories. Because these state-descriptions do not differ in their qualitative description of the system, but do differ in the physical situation that they represent, they are associated with two possible worlds that share the same qualitative character, but differ in their representations *de re* of the system. State-descriptions that violate permutation symmetry are associated with possible worlds that violate anti-haecceitism, that are haecceitistic.

Now consider state-descriptions that satisfy permutation symmetry. Two state-descriptions that differ solely by a permutation transformation represent the same physical situation. They are therefore both associated with the same possible world. There is thus no haecceitistic difference between the possible worlds associated with the state-descriptions. Any two state-descriptions that describe distinct physical situations will differ in their qualitative description of the system. They will thus be associated with possible worlds that differ in their qualitative character. The permutation symmetry of state-descriptions is thus connected to anti-haecceitism. Just as permutation symmetry in modal semantics eliminated haecceitistic redundancies that were artifacts following from the use of singular terms, permutation symmetry in physics removes the haecceitistic redundancies that are artifacts following from the use of particle labels in both the phase space and Hilbert space formalisms.

Permutation symmetry is not related to an ontology where particles are not objects. Rather, it is related to the rejection of haecceitistic differences, which is a common feature of both classical and quantum physics. So instead of grounding a distinction between the ontology of classical and quantum physics, permutation symmetry actually expresses a continuity between the two theories. Both are anti-haecceitistic. Both simply give a qualitative description of the world.

I take this to be essentially the same conclusion that Huggett reaches in his study of haecceitistic differences in classical physics (Huggett, 1999a). In his paper, Huggett demonstrates that classical statistical mechanics is consistent with permutation symmetry (although he, incorrectly in our view, holds back from the conclusion that classical physics is actually permutation symmetric). Classical physics therefore does not make any necessary haecceitistic distinctions. We are in entire agreement with his following conclusion:

> One could easily be led to think that an important difference follows from the different ways in which "identical" particles are handled in the two approaches [that is classical and quantum approaches]. Different statistics point to different metaphysics of individuality in each case: haecceitism in classical mechanics and antihaecceitism in quantum mechanics. My analysis has been directed at showing that there are no heavy metaphysical implications of classical physics, and that therefore anticipated innovations in the notion of an individual in quantum mechanics will not be innovative at all. (Huggett, 1999a, p. 23–24)

Quantum physics satisfies anti-haecceitism, but this does not indicate that quantum particles are somehow strange, for classical physics also satisfies anti-haecceitism.

5.4 Conflicting Claims

At this stage we should address some conflicting claims made by other authors. One author who broadly disagrees with our conclusions is Teller. In his paper (Teller, 2001), he considers the apparent excess of possibilities that follow from the use of particle labels in the quantum Hilbert space formalism. If we take every unsymmetrized orthogonal vector in a Hilbert space description to represent a distinct possibility, then we have an excess of possible cases. For when we assign a distinct probability to each possibility, we do not arrive at the correct statistics for quantum systems. If each orthogonal vector is associated with a distinct possibility, then our statistical ensemble needs to include each vector. Such an ensemble lacks permutation symmetry, which is necessary for quantum statistics.[10]

So in a Hilbert space formalism, there are vectors that are not included in the quantum statistical ensemble. These vectors appear to represent excess possibilities. Teller discusses these excess possibilities in terms of "counterfactual switching". Counterfactual switching is just the permutation of singular terms in some Carnap like state description of a possible world. For Teller, the problem of quantum statistics, as we shall call his problem, asks us to explain away the apparent excess possibilities that follow from counterfactual switching, excess possibilities that lead to the wrong statistics.

Teller claims that the best way to understand these redundancies, and solve the problem of quantum statistics, is by denying "the existence of identity bearing objects" to which particle labels in quantum theory refer (Teller, 2001, p. 377). He calls this the "no-referent option" (Teller, 2001, sec. 7). By denying that particles are objects, he denies that counterfactual switching leads to redundant possibilities; for such switching is not possible when labels do not refer to objects. From this he concludes that "[t]he subject matter of quantum theories . . . must be understood either in terms of a radically conceived "identity free" ontology, or in terms of non-substantial field concepts, or in terms of the multiple instantiation of properties, but not "in" anything" (Teller, 2001, p. 387). Teller's position is diametrically opposed to ours; for none of these approaches treat particles as objects with their own identity.

Teller's argument is analogous to the ontological argument from permutation symmetry, which holds that permutation symmetry follows from the fact that particles are not objects. Teller's argument, however, is cast in more modal terms. While we have already rejected the ontological argument from permutation symmetry, we are now in a position to better see where it goes wrong by examining where Teller's argument goes wrong.

There are redundancies in a Hilbert space formalism, but they do not fol-

[10] Although, as we have pointed out in the previous chapter, permutation symmetry is not sufficient for quantum statistics. We still need the stronger symmetrization requirement, which we have argued is simply an empirical hypothesis.

low from the fact that quantum particles are not objects. They simply follow from anti-haecceitism. As we pointed out above in our discussion of Kripkean specification, we can use such things as particle labels to make haecceitistic distinctions, but these distinctions are not necessarily legitimate. Given anti-haecceitism, these apparently haecceitistic distinctions are simply artifacts of our description that follow from the use of such things as particle labels. We therefore can explain away the excess possibilities of the Hilbert space formalism by use of anti-haecceitism, without any need for a change in ontology. Because the representations *de re* of a possible world supervene upon its qualitative character, two vectors that only differ by permutation must be associated with the same possible world. Because the representations *de re* of a possible world supervene upon its qualitative character, counterfactual switching does not lead to a distinct possible world.

That these redundancies do not follow from a change in ontology is evident from the fact that these redundancies also exist in a classical phase space formalism, where a change in ontology is not a viable explanation. There are particle labels in a phase space formalism just like there are in a Hilbert space formalism. If we associate a distinct possible world with every phase space trajectory, then we have apparent haecceitistic differences. Just as in the quantum case, if these haecceitistic differences are unchecked by permutation symmetry, they lead to incorrect results. In the classical case they lead to the wrong entropy. In the quantum case they lead to the wrong statistics. In both cases we explain these redundancies by use of anti-haecceitism, not by a change in ontology.

Teller, incorrectly in our view, rejects an explanation in terms of anti-haecceitism. We can see this best in his earlier paper (Teller, 1998). In that paper he argues that, if particles are to be objects, they must possess, what he calls, "minimalist haecceities". Minimalist haecceities are not to be confused with the metaphysically robust notion of haecceities that we have considered above. They are something else. Teller explicates the notion with three "tests", which we can take to be necessary conditions. They are as follows:

1. Strict identity: A subject matter comprises things with haec-

ceities just in case the subject matter comprises things to which strict identity applies; that is, just in case there is a fact of the matter for two putatively distinct objects, either that they are distinct or, after all, that they are one and the same thing.

2. Labeling: A subject matter comprises things with haecceities just in case the subject matter comprises things that can be referred to with names directly attaching to the referents; that is just in case these things can be named, or labeled, or referred to with constants where the names, labels or constants each pick out a unique referent, always the same on different occurrences of use, and the names, labels, or constants do not function by relying on properties of their referents.

3. Counterfactual switching: A subject matter comprises things with haecceities just in case the subject matter comprises things which can be counterfactually switched, that is just in case a being A and b being B is a distinct possible case from b being A and a being B, where A and B are complete rosters of, respectively, a's and b's properties in the actual world. (Teller, 1998, p. 121)

The second and third conditions precludes anti-haecceitism.

Teller does not address Lewis's view of anti-haecceitism directly, but he does use his conception of minimalist haecceities to reject van Fraassen's semantic universalism, which, with a few provisos, is equivalent to anti-haecceitism. As he puts it, "if haecceities are admitted [as they must if we are discussing objects] there are more significant propositions than van Fraassen allows" (Teller, 1998, p. 133). Teller claims that van Fraassen's semantic universalism is not true because it does not capture the haecceitistic distinctions that objecthood entails.

It is clear how we should respond to Teller: Minimal haecceities simply beg the question against anti-haecceitism. We need not posit minimalist haecceities in order to make sense of objects that may or may not enter into haecceitistic differences. Only Teller's first condition, strict identity, is necessary for objecthood. We have already discussed in chapter 2 that labeling is not necessary. Even if we cannot label an object, we can still quantify over it and predicate properties to it. And counterfactual switching is not necessary either, unless we presuppose haecceitism.

Teller also has another argument that conflicts with our conclusions. This argument follows from Teller's denial that counterpart theory, which implies anti-haecceitism, can solve the problem of statistics by removing the apparent excess possibilities of the Hilbert space formalism (Teller, 2001, p. 374–376). Teller's argument turns on the claim that we cannot understand the possibilities described by quantum state-descriptions in terms of possible worlds. This is because the use of particle labels requires that we consider the counterparts of specific particles instead of considering possible worlds. As Teller writes:

> Likewise, a counterpart theoretical interpretation of the probabilities for what will happen to Boson 1 and 2 will be in terms of distributions over possibilities involving counterparts of 1 and 2. To interpret in terms of possible worlds not characterized in terms of counterparts for 1 and 2 would be to abandon the interpretation of '1' and '2' as referring labels in discussion of probabilities of what will happen to *them*, or at least to abandon the counterpart theoretic approach to such treatment. (Teller, 2001, p. 375)

If we understand the possibilities described by quantum state-descriptions in terms of the counterparts of specific particles, then, Teller argues, we still have an excess of possibilities. This is because a counterpart of one particle can also be a counterpart of another particle of the same kind. As Teller writes:

> Suppose that in the actual world we have prepared a box with two Bosons, 1 and 2, each with completely indeterminate position, and a measurement apparatus set up to detect the number of particles in the right and in the left side of the box . . . We

> consider a possible world similarly arranged, but in which the measurement apparatus has been triggered and detects one particle in the right and one in the left side of the box in that world. Now, which is the counterpart of the real particle 1 and which the counterpart of the real particle 2? It makes no difference. We can chose arbitrarily. (Teller, 2001, p. 374)

Thus while there is only one possible world, there are two possibilities instead of the one demanded by quantum statistics; for there are two ways we can draw the counterpart relation. So counterpart theory, and with it anti-haecceitism do not appear to solve the problem of statistics.

But it seems to me that this is clearly wrong. The way we explain what possibility a state-description describes is by use of a possible world. Whether or not two state-description describe the same possibility depends upon whether or not they describe the same possible world. And we can interpret the probability measure assigned by a statistical ensemble over state-descriptions as a probability measure over possible worlds. This does not change if we adopt counterpart theory. Even if we adopt counterpart theory, we are still only concerned with the possible worlds described by state-descriptions and not with the counterparts within those worlds of real or actual particles. The use of particle labels in no way indicates that the latter is the case; for the labels within the state-description, whether we interpret them in terms of proper names or bound variables, only refer to particles in the possible world described by the state-description, and not to particles in the actual world (unless the possible world described is the actual world).

Counterpart theory, and with it anti-haecceitism, does solve the Teller's problem of quantum statistics. Once again, we can understand apparent excess possibilities of the Hilbert space (and phase space) formalism as haecceitistic redundancies that follow from the use of particle labels, redundancies that do not express legitimate haecceitistic differences, redundancies that are artifacts of the description. There is no need, and no reason, to think, as Teller does, that these redundancies follow from the fact that quantum particles are not objects.

Another author who broadly disagrees with our conclusions is Gordon.

In his paper (Gordon, 2002), Gordon takes issue with Huggett's conclusion, outlined above (see page 117). Recall that Huggett, like us, does not think that quantum statistics indicates a difference in ontology. This is because there is no difference in the haecceitistic distinctions made by classical and quantum physics. Both are anti-haecceitistic (or as Huggett claims capable of being anti-haecceitistic).

Gordon grants that classical physics does not make any necessary haecceitistic distinctions, but he denies that this provides a continuity between the nature of classical and quantum particles. This is because classical and quantum particles differ in how they are individuated in each theory. We cannot individuate quantum particles, as we can classical particles, by use of spatiotemporal trajectories or PII. If we are to regard quantum particles as objects, then it seems that we need to introduce metaphysically robust haecceities to individuate them. The use of such haecceities again appear to lead to excess possibilities that are inconsistent with quantum statistics. As he puts it:

> The important point to walk away with here is that both the phase-space and distribution-space [Z-space] representations in MB [Maxwell-Boltzmann] statistics presuppose classical individuative criteria for particles. It is precisely these criteria that give *de re* modality a foothold in physical theory, and it is just these criteria that quantum statistics problematizes. (Gordon, 2002, p. 413)

According to Gordon, anti-haecceitism does not solve the problem of quantum statistics.

Given what we have said in chapter 2, this argument in unpersuasive, for it presupposes the necessity of a principle of individuation. We have shown in chapter 2 that a principle of individuation is not necessary in order to regard particles as objects. We have argued that the identity of particles is best viewed as primitive. So we can reject Gordon's worry that the individuation of quantum particles smuggles in haecceitistic distinctions that are problematic for quantum statistics.

5.5 Physical Support for Anti-Haecceitism

But we must acknowledge that Gordon is right in one important respect. Even if physics does not make any explicit haecceitistic distinctions, haecceitism could still be true. There might be haecceitistic distinctions that are not recognized by a physical description. But the fact that physics does not make any haecceitistic distinctions gives us good reason to doubt haecceitism. This is especially true if we are physicalists and believe, roughly, that everything in the actual world supervenes upon the physical character of the world, which is qualitative.[11] This does not imply that what other worlds represent *de re* also supervenes solely upon the qualitative character of those possible worlds. But the onus is now on the metaphysician to explain why these possible worlds are haecceitistic, while the actual world is not. In what relevant way do these possible worlds differ from the actual world? For the supporter of haecceitism and physicalism, there is no easy answer to this question.

But even a supporter of haecceitism who rejects physicalism does not have an easy task. A supporter of haecceitism should be able to provide us with a clear example that can only be understood in terms of haecceitism. Such an example is not necessary, but it should be sought after. Since both classical and quantum physics are anti-haecceitistic, such an example must differ from a classical or quantum example. This does not appear to be a problem in the quantum case. Most philosophical examples, like Kripke's dice, are not explicitly quantum. But examples like Kripke's dice are classical. Because classical physics is anti-haecceitistic, an example of haecceitism must make use of more than just classical objects. Such a task is not obviously impossible (maybe something like consciousness is relevant), but it is obviously difficult.

While physics does not necessarily lead us to accept anti-haecceitism, the fact that both classical and quantum physics lack any haecceitistic distinctions gives us very good reason to hold anti-haecceitism in general. This provides an excellent example of how physics can influence metaphysics.

[11] There is of course an ongoing debate as to the proper definition of physicalism. Two leading candidates are provided by Lewis (1983, p. 361–365) and Jackson (1998, ch. 1).

5.6 Anti-Haecceitism and the Hole Problem

Anti-haecceitism is not only relevant to our understanding of permutation symmetry and quantum statistics. It also has consequences for our understanding of spacetime physics. It is specifically relevant to the so-called hole problem.

Much has been made lately about the connection between permutation symmetry in quantum theory and the hole problem in spacetime physics. There is an active debate as to the proper connection and the proper ontological conclusions that follow. As we will see in this section, anti-haecceitism is really the common feature which unites the two. It leads to similar ontological conclusions in both cases.

Let us first start with a brief introduction of the hole problem. The hole problem, as presented by Earman and Norton (1987), is meant to force a spacetime substantivalist into a corner. According to the argument, if we hold that spacetime exists independently of material objects and their spatiotemporal relations, then we must accept that any generally covariant theory is indeterministic.

The models of a generally covariant theory consist of an ordered $(n+1)$-tuple $< M, O_1, O_2, \ldots, O_N >$, where M is a differentiable manifold, and O_i is a geometrical object that is defined everywhere on M. Since the theory is generally covariant, if $< M, O_1, O_2, \ldots, O_N >$ is a model of the theory, then so is $< M, h^*O_1, h^*O_2, \ldots, h^*O_N >$, where h is a diffeomorphism from M onto M. A corollary of general covariance is that there are distinct models that differ only within an arbitrarily small region, the hole. This is because there exists arbitrarily many different diffeomorphism transformations that are equal to identity outside the hole, but differ from identity within the hole.

Earman and Norton claim that a spacetime substantivalist must regard each of these distinct diffeomorphic models as describing a distinct physical situation, or to use our possible worlds terminology, as equivalent to a distinct possible world (Earman and Norton, 1987, p. 521–522). They hold that a substantivalist must regard the differentiable manifold as representing spacetime and all of the geometrical objects defined on the manifold as representing fields contained in spacetime. If spacetime exists, then each diffeomorphic model represents the same spacetime, but with the contents—the fields—shifted in different ways. The consequence of this is that for a spacetime substantivalist, the theory is indeterministic. Even if we know everything outside the hole, there are several possible situations (corresponding to the different diffeomorphic models) that could exist within the hole. Knowing the state of all the fields outside the hole is not enough to determine a unique model that describes the fields within the hole. This violates nearly every definition of determinism.

It is easy at this stage of our discussion to see that there is analogy between the putative problems of quantum statistics and the hole argument. Both problems point out that the identity and existence of a certain class of objects lead to an apparent excess of possibilities, which are problematic. In the quantum case, when we take quantum particles as objects, permutations lead to distinct vectors that appear to represent distinct possibilities, possibilities that are not accounted for in quantum statistics. In the spacetime case, when we take spacetime points as objects, diffeomorphism transformations lead to distinct models that appear to represent distinct possibilities, possibilities that undermine determinism.

Teller (2001) frames the analogy in terms of counterfactual switching. As he writes:

> The two problems, quantum statistics and the hole argument, have much in common as to how the excess possible cases arise: They both work by what I have called *counterfactual switching*. In both problems we have names—number-labels of quantum particles and number-coordinates of space-time points. In both problems we suppose that there are identity bearing things, the quantum particles or the space-time points, to which these names refer, and that reference is constant across possible cases. Finally in both problems we get descriptions of the problematic possible cases by supposing redistribution of ALL the properties and relations pertaining to one object of reference from that referent to another. (Teller, 2001, p. 371)

Several authors have pointed out that the use of diffeomorphism transformations in the hole argument are similar to permutation transformations. Maudlin (1990, p. 545) points out that if a substantivalist simply regards the set of mathematical points of the manifold, devoid of any topological structure, as representing spacetime, simply permuting the points in the set leads to a distinct mathematical object that the substantivalist must consider as describing a distinct possible world. And both Rynasiewicz (1994) and Liu (1996) argue that the hole argument is analogous to Putnam's paradox, where we consider the permutation of elements in some generic domain. Stachel (2002) builds upon this and formulates a set theoretic version of the hole argument that abstracts away from the differentiability and continuity of the differentiable manifold. In this set theoretic version, diffeomorphism transformations on general covariant models give way to permutation transformations on structured sets.[12]

Now whether or not we can reformulate the hole argument in such a way that we replace diffeomorphism transformations with permutation transformations, the important point is that, in both the quantum and spacetime case, certain transformations lead to an apparent excess of possibilities. In both cases these transformations do not affect the qualitative description of the situation. Because of this, any putative problems that follow from the apparent excess of possibilities disappear in the light of anti-haecceitism.

In the quantum case, permutation transformations only change which particle is in which single-particle state. Since, for particles of the same kind, this does not express a qualitative difference, by anti-haecceitism, which particle is in which state does not represent *de re* anything about the system. We therefore associate states connected by a permutation transformation with the same possible world. (Again we are ignoring, but not rejecting, the possibility of indiscernible worlds.) In the relativistic case, diffeomorphism transformations only change which spacetime points underly which parts of the fields. Since this does not express a qualitative difference, by anti-haecceitism, which spacetime points underly which parts of the fields

[12]See Pooley (forthcoming) for a critical assessment of Stachel's set theoretic hole argument.

does not represent *de re* anything about the system. We therefore associate the set of diffeomorphic models with the same possible world. Taking diffeomorphic models to describe the same physical situation is often called Leibniz equivalence, but we see that Leibniz equivalence follows from the more general thesis of anti-haecceitism.

Once we have accepted Leibniz equivalence, fears of indeterminism melt away. The fields outside the hole do not determine a unique model, but they do determine a set of diffeomorphic models that describe the same physical situation and give same representation *de re* of the system. A choice of a particular model is therefore just like a choice of name: practically important, but physically meaningless.

This anti-haecceitistic solution to the hole problem is consistent with substantivalism. Leibniz equivalence does not lead us to deny that spacetime points are distinct objects any more than permutation symmetry leads us to deny that particles are distinct objects. Earman and Norton disagree and claim that the substantivalist cannot accept Leibniz equivalence. Since a substantivalist holds that the differential manifold represents spacetime, any diffeomorphic transformation that moves the fields on the manifold will represent a distinct physical situation. For in each of these different models different points of the manifold will underly different parts of the fields. But a substantivalist need not accept such haecceitistic differences in order to maintain that spacetime points are distinct objects. Just as we can accept that particles are objects, but deny that two qualitatively identical possible worlds can differ in how they represent *de re* which particular particles are in which single-particle states, we can also accept that spacetime points are objects, but deny that two qualitatively indiscernible possible worlds can differ in how they represent *de re* which particular spacetime points underly particular parts of some field.

We should note that this anti-haecceitistic view is neutral as to whether it is the differentiable manifold that represents spacetime. Leibniz equivalence holds regardless of whether we take the set of points of the manifold, independent of any topological structure, as representing spacetime (as Earman and Norton do), or, going the other way, if we take the differentiable manifold coupled with the metric tensor as representing spacetime (as Maudlin (1990) and Hoefer (1996) do). In either of these cases, diffeomorphism transformations do not alter the qualitative feature of the model, and diffeomorphic models are associated with the same possible world by anti-haecceitism.

Now of course we are not the first to claim that substantivalism is consistent with Leibniz equivalence. There have been several different arguments for that conclusion. But nearly all of them assume or imply anti-haecceitism. One well known exception is Maudlin (1990). Maudlin claims that spacetime points have essential properties, given by the metric tensor, that prevent all diffeomorphic models from representing a legitimate possible situation. There is no indeterminism between diffeomorphic models because only one model among the set of diffeomorphic models correctly expresses the essential properties of spacetime.

It is clear that it is essentialism and not anti-haecceitism that is at work here. While the two solutions are not mutually exclusive, they differ in one important aspect. In an essentialist solution, not all models represent legitimate possible worlds. However, in an anti-haecceitistic solution, we can accept that all models represent legitimate possible worlds; it is just that diffeomorphic models all represent the same possible world.

While Maudlin's solution to the hole problem does not turn on anti-haecceitism, other well known solutions do. One such solution is given by Butterfield (1989). He solves the hole problem by adopting a definition of determinism similar to one given by Lewis (1983, p. 359–361). Lewis takes two objects to be duplicates if and only if they share all of their qualitative intrinsic properties. Using duplicates he defines divergent worlds. Two possible worlds diverge if and only if they are not duplicates but have initial temporal segments that are duplicates. He defines a theory to be deterministic if and only if there are no two divergent worlds that both completely obey the theory.

Butterfield (1989) presents a similar definition of determinism for covariant theories. Butterfield, however, differs from Lewis by defining duplicates spacetime regions without appealing to the notion of qualitative or intrinsic properties. Butterfield takes spacetime regions to be duplicates if and only if they are isomorphic. Two worlds diverge if and only if:

> (1) they both contain regions S, $S^{'}$ of kind **S**; and (2) there is a diffeomorphism $\alpha : S \rightarrow S^{'}$ with $\alpha^{*}(O_i) = O_i^{'}$ on $\alpha(S) = S^{'}$; and (3) there is no global isomorphism $\beta : M \rightarrow M^{'}$ with $\beta^{*}(O_i) = O_i^{'}$ and $\beta(S) = S^{'}$. (Butterfield, 1989, p. 24–25)

A consequence of this definition is that even if diffeomorphic models describe distinct possible worlds, these possible worlds will never diverge, and therefore never violate determinism. The hole problem is solved.

On the face of it, this solution to the hole problem has nothing to do with anti-haecceitism, for it allows that diffeomorphic models describe distinct possible worlds. But anti-haecceitism is related to the definition of determinism used here. The definition of determinism given by Butterfield does not imply counterpart theory, but, as he acknowledges, counterpart theory lends support to this definition (Butterfield, 1989, p. 25). If we accept counterpart theory, then it is quite natural to define divergence in terms of world-bound duplicates. If we reject counterpart theory, and accept that the same spacetime points can exist in more than one possible world, we should define divergence in terms of transworld spacetime points instead of in terms of world-bound duplicates. Such a definition of divergence will not deliver us from the hole problem.

So it is counterpart theory that motivates and justifies Lewis's and Butterfield's definition of determinism. But as we have already pointed out above, counterpart theory implies anti-haecceitism, and with it a solution to the hole problem. The lesson here is that if we accept Lewis's definition of determinism, we should deny from the very beginning that diffeomorphic models represent possible worlds that differ in their representations *de re*. That these possible worlds never diverge becomes trivial.

And once we have accepted anti-haecceitism, we can dismiss the arguments given by Belot (1995) and Melia (1999) against Lewis's definition of determinism and Butterfield's use of it. Both Belot and Melia's arguments essentially consist of various examples that they believe are indeterministic, but which are deterministic by Lewis's definition. One example both consider is a theory that describes the collapse of a column.[13] Consider a perfectly cylindrical column sitting on top of a perfect sphere. Imagine that the column collapses due to a spherical load located at the center of the top of the cylinder. Further imagine that the rest of the universe is empty. Standard treatments of loaded cylinders will hold that, for this symmetric case, the cylinder can buckle in any direction.[14] Such theories are therefore indeterministic, and there should be divergence between two worlds that are initial duplicates. But, according to Lewis's definition, there is no divergence, for in these symmetric worlds there is no qualitative intrinsic difference between a column that collapses one way and column that collapses another way.

This example and the others presented by Belot and Melia postulate haecceitistic differences between possible worlds, a point which Brighouse (1997) also makes. If we accept anti-haecceitism, then there can be no difference in how these symmetrical worlds represent *de re* the column collapsing. The only reason it appears plausible to claim that there should be a difference is because we often describe less than maximally specific possible worlds where the cylinder buckles in one direction instead of another. While, like in Kripke's dice example, these less than maximally specific possible worlds can differ in their representations *de re* without differing in their explicit qualitative features, it is only because they are less than maximally specific, and not because they violate anti-haecceitism. Accepting anti-haecceitism, these counterexamples to Lewis's theory of determinism fail. So we see that anti-haecceitism is not only implied by Butterfield's solution to the hole problem, it is also allows us to defend his position from criticism.

Other well known solutions to the hole problem are more clearly connected to anti-haecceitism. Hoefer (1996) argues that if we reject the primitive thisness of spacetime points, we can accept both substantivalism and Leibniz equivalence.[15] As he argues, when we reject the primitive thisness of space-

[13] Both authors attribute the example to Wilson (1993, p. 215–216).

[14] For simplicity imagine that the buckled cylinder will always have the same shape.

[15] Hoefer actually uses the phrase "primitive identity" instead of "primitive thisness", but he acknowledges that the latter is an appropriate surrogate for the former. I have chosen to substitute the latter since in this thesis I use "primitive identity" to refer to the lack of a principle of individuation, rather than to primitive thisness.

time points, we deny that qualitatively identical models of spacetime points can describe distinct physical situations. But the important issue here is not haecceities but haecceitism. For first of all, Hoefer use anti-haecceitism to support his position on haecceities (Hoefer, 1996, p. 15). And second of all, we can still have haecceitistic distinctions without haecceities as we have pointed out above. Objects can still have primitive cross-world identities without haecceities, and these identities can lead to overlap, and with it haecceitism. So it is not the rejection of haecceities that leads to Leibniz equivalence. It is the rejection of haecceitism.

Brighouse (1994) and Saunders (2003) also claims that a substantivalist can accept Leibniz equivalence. They both use PII to support their views (although Brighouse only considers the application of PII to possible worlds). Since the possible worlds associated with a set of diffeomorphic models possess the same qualitative character, they are identical by PII. Once again it is not really PII, but rather the weaker condition of anti-haecceitism, that is at work here. What is important is not that the worlds are actually identical, but that they have the same representations *de re*.

So while anti-haecceitism does not provide the only solution to the hole problem available to the substantivalist, it is at the heart of many solutions. And it is also the relevant feature that connects permutation symmetry and Leibniz equivalence. Some might think that this analogy suggests that spacetime points, like quantum particles, are not objects since they lead to excess possibilities. But as we have seen, it is the opposite that is true. The excess models, like the excess Hilbert space vectors, are simply artifacts of the description that do not represent any legitimate haecceitistic differences. We account for them by use of Leibniz equivalence, in the spacetime case, and permutation symmetry, in the quantum case. This does not in any way preclude the objects described by the models and vectors from being just that—objects.

Anti-haecceitism is also connected to any other symmetries that do not change the qualitative description of a system, although in a slightly different way. Any two state-descriptions that are connected by such a symmetry will, by anti-haecceitism, be associated with the same possible world. So any apparent difference between the state-descriptions connected by such a symmetry transformation will not actually represent *de re* anything. While the differences might not be haecceitistic differences, as they are in the case of permutation and diffeomorphism symmetries, they will still only be artifacts of the description.

For example, consider the Galilean symmetry of Newtonian mechanics. If we hold that the qualitative properties of Newtonian mechanics are invariant under Galilean transformations, then any two descriptions that are connected by a Galilean transformation give the same representation *de re*, even though they might differ in their assignment of properties like absolute position. Absolute positions are then like the haecceitistic differences that follow from particle labels and diffeomorphic models: they are simply artifacts of the representation that do not represent *de re* anything.

In fact, if we hold, as Saunders (2003) does, that all physical properties are invariant under the exact symmetry transformations of a theory, then any such symmetry transformation will not change the qualitative description. Any differences between state-descriptions connected by such a transformation will be artifacts of the description.[16] But we will not here ask whether all properties of a theory are actually invariant under the exact symmetry transformations of the theory. This is a substantive claim that deserves its own investigation. Let us here just say that this is the case for permutation symmetry and diffeomorphism symmetry for the anti-haecceitistic position that we have outlined above.

Let us recap. We started this chapter by wondering if we could provide a philosophical foundation for permutation symmetry to replace the ontological explanation that we rejected. We have seen that such an explanation is provided by anti-haecceitism, an explanation that allows us to understand the apparent excess possibilities of both classical and quantum physics without requiring any change in ontology. (And as we have just seen, anti-haecceitism also provides a similar explanation of Leibniz equivalence.)

As van Fraassen (1991, p. 435) comments, a theory does not wear its

[16] Ismael and van Fraassen (2003) presents a similar proposal to this one for using symmetry to eliminate "superfluous theoretical structure".

content on its sleeve. And, as Maudlin writes, "the ontological structure of the physical universe does not mirror the ontological structure of the mathematical object representing it" (Maudlin, 1990, p. 545). The haecceitistic differences that follow a literal interpretation of the formalism do not represent *de re* anything. They are simply artifacts of the description that we account for by use of permutation symmetry. Permutation symmetry therefore does express an important feature of physics, but it is not that particles are not objects. It is that the world is qualitative.

Chapter 6

Particles and Fields

6.1 The Relation between Particles and Fields

Our argument up to this point has been that quantum particles are objects despite their indiscernibility and permutation symmetry. But we have not yet addressed the relationship between quantum particles and quantum fields. In quantum theory, wherever we describe things in terms of particles, it seems that we can also describe them in terms of fields. Further, there are situations and regimes where only a field, and not a particle, interpretation seems possible. So it appears that fields provide an equivalent, if not superior, ontology to particles. If this is true, how does it affect our arguments about the individuality of quantum particles? The worry is that a proper understanding of the relationship between quantum particles and quantum fields will undermine our previous arguments, and force us to reject the view that quantum particles are objects. This would revive the thesis that the change from classical to quantum physics is necessarily accompanied by a change in ontology, where particles are no longer objects. In this chapter we will explore the relationship between quantum particles and fields, and clarify in what sense we can take quantum particles as objects.

Let us start by reviewing the connection between quantum particles and fields. We often describe systems of quantum particles by use of the standard tensor product Hilbert space description, where the state-space is a tensor product of N single-particle Hilbert spaces, H_N. As is well known, the direct sum of such tensor product Hilbert spaces define a Fock space, $H_0 \oplus H_1 \oplus H_2 \oplus \cdots \oplus H_N$. We can uniquely map vectors from the tensor product Hilbert space to the Fock space.

Now the Fock space description has a clear field interpretation. Because we can map tensor product Hilbert space states to Fock space states, we can extend this field interpretation to states in the tensor product Hilbert space description, replacing a particle ontology with a field ontology. We can also go the other way and interpret the Fock space states that are equivalent to tensor product Hilbert space states in terms of particles instead of fields.

In quantum field theory, we can decompose a free quantum field into its modes and take the state of the field to be discrete excitations of those modes. We can represent the state in Fock space by use of raising operators acting upon the free field vacuum, where each raising operator represents a discrete excitation into a given mode. Let us call this the occupation number formalism.

We can map every symmetrized state-vector in the tensor product Hilbert space formalism onto an occupation number vector. We can also map all operators and dynamical equations.[1] This allows us to extend our field interpretation of the occupation number formalism to the tensor product Hilbert space formalism, where we replace particles with discrete excitations of modes. It also allows us to interpret states in the occupation number formalism that are equivalent to tensor product Hilbert space states in terms of particles, where we replace discrete excitations of modes with particles.

For example, consider the tensor product Hilbert space state (in position representation):

$$\Psi(\mathbf{r}_1, \mathbf{r}_2) = 2^{-1/2}(u_k(\mathbf{r}_1)u_{k'}(\mathbf{r}_2) + u_{k'}(\mathbf{r}_1)u_k(\mathbf{r}_2)) \tag{6.1}$$

where $u_k(\mathbf{r})$ and $u_{k'}(\mathbf{r})$ are orthogonal single-particle Hilbert space states.

[1] For further details see Robertson (1973) and Huggett (1994).

This tensor product Hilbert space state is equivalent to the Fock space state:

$$|\Psi\rangle = \hat{a}_k^\dagger \hat{a}_{k'}^\dagger |0\rangle \tag{6.2}$$

where $\hat{a}_k^\dagger$ is the appropriate raising operator acting on the free field vacuum $|0\rangle$, an operator that describes a discrete excitation of the field into the kth mode. Because of the equivalence between these two states, we can interpret the tensor product Hilbert space state $\Psi(\mathbf{r}_1, \mathbf{r}_2)$ as describing a field with two excited modes or we can interpret the Fock space state $|\Psi\rangle$ as describing two particles.

A common mistake is to identify a discrete excitation into the kth mode of the field with a particle in a single-particle pure state $u_k(\mathbf{r})$. Because of symmetrization, no single particle in a system of two or more particles of the same kind is in a single-particle pure state, and therefore we cannot identify a single particle with a single excitation.[2] Nonetheless we can still associate discrete excitations with particles, it is just that these excitations are associated with particles in symmetrized states, not single-particle pure states.

So it appears that we can always substitute a field ontology for a particle ontology. The reverse, however, is not always the case; for there are aspects of quantum field theory for which only fields appear to provide an adequate ontology.

Often people cite the lack of localized particles in relativistic quantum field theory as demonstrating that particles provide an inadequate ontology for quantum field theory. Certainly it is true that localization is a problem for relativistic fields.[3] It is true that the standard conception of particles includes some sort of localization. But we are concerned with a weaker notion of particles. We only wish to investigate whether we can regard them as objects at all, even if they are not localizable objects. While we acknowledge that the apparent lack of localization provides several philosophical puzzles,

[2]This point is made by van Fraassen (1991, p. 443).

[3]There is an extensive literature dealing with the problem of locality in relativistic quantum field theory. Saunders (1994) and Halvorson and Clifton (2002) both provide good introductions and references to the field.

we will not consider them here.

Even ignoring problems of localization, there are other aspect of quantum field theory that are problematic for a particle interpretation. States of indeterminate particle number provide one example. In a Fock space formalism, there is a Hermitian operator that gives the total number of particles or mode excitations described by a state. Not all states in Fock space are eigenvectors of this number operator. For example, a vector can be a superposition of two eigenvectors with different particle numbers. If we assume that a system in an indeterminate particle number state does not possess a determinate number of particles, as some interpretations hold, then indeterminate number states are not states of particles.[4]

When we take an object, as we have above, as that which we describe by standard quantifiers and predicates, then we must be able to apply standard set theory to the object, for otherwise the object cannot enter into a fixed domain over which the quantifiers range. Standard set theory obviously does not apply to an indeterminate collection of particles, for a set of such particles does not have a clear extension. We therefore cannot consider such particles as objects in the way we have outlined.

A superposition of eigenvectors of other observables does not lead to the same problem for particles, for a system need not possess a definite value for other properties in order for the particles of the system to be objects. For example, particles do not have to possess definite momentum in order to be objects, and therefore neither does the system. But particles do need to possess a definite existence, and this implies a definite particle number for the system.

While indeterminate particle number states present problems for a particle interpretation, there are no such problems for a field interpretation. Fields have a determinate existence even if the number of field excitations is indeterminate. The number of field excitations become just another property like momentum. In this way a field ontology appears superior to a particle

[4]Not all interpretations hold that the number of particles is actually indeterminate for an indeterminate particle number state. For example, see Teller (1995, p. 32). But claims such as Teller's are controversial.

ontology.

Some might claim that indeterminate particle number states can still describe particles, just not particles that are objects. Rather, such states describe particles that are capable of having an indeterminate existence or indeterminate number. Let us call such particles, as many do, "quanta". The problem with quanta is that we do not appear to have any way of understanding their nature independently of a field interpretation. Quanta appear to be nothing more than discrete excitations of quantum fields by another name: they do not appear to provide us with an alternative ontology to that of fields.

One might object and claim that we can elucidate the nature of quanta independently of fields. One possible way we might understand the nature of quanta independently of a field ontology is by use of alternative set theories, such as quaset or quasi-set theory.[5] Both of these theories are generalizations of set theory. Quaset theory allows for the indeterminate extension of a set. Quasi-set theory allows for the indeterminate identity of elements in a set. While the authors of these theories claim that they capture the unusual nature of non-individual particles, neither of these theories is capable of describing particles of indeterminate number. This is because the generalized sets in both theories still have a cardinality. It is clear that any finite collection of an indeterminate number of particles should not have a determinate cardinality. Another way we might understand quanta is in terms of ontic vagueness. We will discuss this proposal in greater detail in the next chapter. There we will argue that ontic vagueness does not provide us with a way to understand particles of indeterminate existence.

Quanta, therefore, do not appear to present us with a new category of entity, lacking determinate existence, with which we can extend a particle interpretation to indeterminate particle number states. While we can interpret indeterminate particle number states in terms of quanta, such an interpretation is equivalent to a field interpretation. Any indeterminate existence that quanta possess is simply due to the fact that they are related

[5]See Dalla Chiara et al. (1998) for a review of these alternative set theories and their proposed application to quantum particles.

to the properties of the true objects described by the indeterminate particle number states—the fields.

States of indeterminate particle number do not present the only problem for a particle interpretation in quantum field theory. There are also several problems associated with the expectation values of the vacuum state. The vacuum appears to possess a non-zero energy. We might question the reality of the vacuum energy, but if we take quantum field theory seriously, as more than just an effective theory, then we have some good reasons to take the vacuum energy as real.[6] There are also other expectation values that appear to describe vacuum fluctuations. These non-zero expectation values are of observables that do not commute with the particle number operator.

If we take the vacuum to be void of all particles, then we cannot explain these vacuum phenomena in terms of particles. Some claim that we can understand vacuum phenomena in terms of particle fluctuations (see Sciama, 1991). Instead of claiming that the vacuum is the state of a system with no particles, they claim that the vacuum is a state where virtual particles continually come into and out of existence. But there are some serious objections to this view. For example, Saunders (2002) claims that such explanations are semi-classical, and therefore do not present an adequate quantum explanation of vacuum phenomena. And Redhead (1994) argues that these are free field vacuum phenomena, and that virtual particles have no place in our interpretation of free fields.

A particle interpretation of these phenomena, therefore, appears problematic. Yet there does not appear to be any problems for a field interpretation. In a field interpretation, the vacuum is just the ground-state of the field. And of course, ground-states can have non-zero expectation values.

Yet another problem for a particle interpretation is the Unruh effect. In the Unruh effect, an observer that is uniformly accelerating in a Minkowski spacetime observes a thermal background of particles, while an inertial observer observes a vacuum state with no particles. If particles provide the correct objective ontology for quantum field theory, then their existence should

[6] See Saunders (2002) for a detailed analysis of whether the vacuum zero-point energy is real.

not be dependent upon the path of an observer. Again, there does not appear to be such problem for fields. While different observers might measure different states of the field, the existence of the field is not dependent upon the path of the observers.

So here we have some examples of how a field ontology appears superior to a particle ontology.[7] So what are we to make of all of this? Because of this superiority, fields appear a better candidate than particles for the fundamental ontology of quantum theory. Of course all of this could change. Relativistic quantum field theory could turn out to be just an effective field theory. Its fundamental replacement might have a different ontology, perhaps even a particle ontology. But in the absence of such a theory, it is difficult, if not impossible, to claim that particles provide us with a fundamental quantum ontology.

A particle ontology, however, still provides an emergent or effective ontology. By this we mean that a particle ontology is useful in describing phenomena in certain regimes of quantum theory, such as non-relativistic quantum mechanics of fixed particle number or certain non-interacting relativistic cases.[8] Quantum particles supervene upon quantum fields, but they are still real in so far as they provide a useful emergent ontology.

Of course there are plenty of examples of emergent or effective ontologies. The observable objects of our daily existence provide the clearest example. Tables and chairs might ultimately be nothing more than collective excitations of quantum fields, yet they are still indispensable to our understanding of the observable world. And not all emergent ontologies consist of observable objects. Biology provides several examples of unobservable emergent objects. Chromosomes (which are directly unobservable) supervene upon molecules, atoms, and ultimately quantum fields, but they are still useful in describing certain biological phenomena, and in this way are real like molecules, atoms,

[7] A field ontology is not without its own problems. Interacting fields that do not have a Fock space representation and unitarily inequivalent representations present some problems for a field interpretation. A particle interpretation, however, does no better on these matters.

[8] This view about emergent ontology is similar to the one presented by Dennett (1991), who urges us to accept the reality of certain patterns.

and quantum fields. Just like tables, chairs, and chromosomes, quantum particles provide us with another example of an emergent ontology.

We have, however, noted that we can substitute fields for particles, even if we cannot always substitute particles for fields. So even if particles can provide an ontology for phenomena in some limited quantum regime, so can fields. But this fact does not undermine the use of quantum particles or their individuality. Since particles do not provide the fundamental ontology of quantum theory, the ability to substitute fields is not a problem. We adopt a particle ontology, not because particles are the basic constituents of all reality, but simply because particles are useful in describing a certain class of phenomena. So long as particles are useful, it does not matter if there is an alternative description in terms of fields.

And there are plenty of instances where a description in terms of particles is more useful, succinct, and clearer than an alternative description in terms of fields. We have already provided the example of atomic Bose-Einstein condensates (BECs) in chapter 2 (see page 27). A particle ontology provides a simple way of understanding the interactions in an atomic BEC. And atomic phenomena in general seem easier to understand in terms of particles. It is true that many theoretical descriptions of atomic phenomena use effective field theories, but, as was the case for atomic BECs, we generally do not interpret the fields literally. We do not think of a cesium atom as a discrete excitation of some cesium field, even if we use an effective field theory to describe cesium atoms. We might think of a cesium atom as a collective excitation of the fundamental fields of the standard model, but this is not what is described by an effective field description, and it far too difficult to describe cesium atoms in terms of fundamental fields. So while we can substitute a field ontology for a particle ontology, the two are not equally useful.

We should point out that there is analogy here between the effective field theories describing atomic phenomena and use of fields in general relativity. In general relativity, we take a spacetime model to be a collection of fields defined on a manifold. But while fields enter into the formal description of general relativity, as they do in effective quantum field theories, we do not always interpret these fields literally. In general relativity, we often take them to describe planets, suns, and galaxies, just like we take the effective fields of quantum theory to describe atoms or what have you.

And we should also point out that we can also substitute fields for particles in classical physics, like we can in quantum physics, as Redhead (1982, 1988) describes in detail. For we can take a classical point particle to be an excitation of a special sort of field. This field possesses the property of penetrability or impenetrability at every spacetime point. Now instead of talking about a classical particle at a point, we talk about a field that possesses the property of impenetrability at that point. In this way we can substitute a field ontology for classical particle ontology. But while we can substitute classical fields for classical particles, a description in terms of classical fields will not always be as useful as one in terms of classical particles. A classical field of impenetrability does not offer a more concise description of a classical gas, or any other phenomena that we usually describe by use of classical particles. Just like in the quantum case, a particle ontology is not undermined by the fact that we can substitute a field ontology.

The important point that we want to make here is that while a particle ontology is not the most fundamental ontology of quantum theory, and while we can substitute fields for particles, particles still serve as emergent ontology. Now as an emergent ontology there is still a question as to whether particles are objects are not. But our arguments of the previous chapters show that they are.

6.2 The Surplus Structure Argument

Having introduced the Fock space formalism above, we should now consider a very well known argument, an argument that purports to show that Fock space provides a superior description of non-relativistic quantum mechanics (NRQM), a description that suggests that quantum particles are not objects. This argument is Redhead and Teller's surplus structure argument (Redhead and Teller, 1991, 1992; Teller, 1995).

Their argument is this. They start by claiming that if quantum particles are objects, then they must possess "label transcendental individuality," or LTI for short, where LTI is an "attribute transcending principle of individuation, that in virtue of which the individual can bear a label" (Redhead and Teller, 1992, p. 203). The basic idea here is that if a particle is to be an object it must be capable of being labeled, and LTI is that which allows it to be labeled.[9] Since there are no distinct spatiotemporal trajectories that can secure the reference of quantum particle labels, Redhead and Teller take LTI as something, like haecceities or bare particulars, that transcends qualitative properties (Redhead and Teller, 1992, p. 211).

The idea that objects must be capable of being labeled is one that we have already discussed in chapter 5 in relation to some of Teller's other papers. We will revisit our views on this issue shortly. But first let us consider the rest of Redhead and Teller's argument.

Having established, to their own satisfaction, the importance of LTI, they consider the tensor product Hilbert space description of NRQM. They argue that if we take this Hilbert space formalism as describing particles with LTI, then we must take unsymmetrized vectors describing particles of the same kind as representing possible states. As they put it:

> To say that quantum particles are the sorts of things which can be meaningfully thought of as bearing labels is to say that it is meaningful to think of a first particle having one property and a second particle having a second property. This fact is reflected in the labeled tensor product Hilbert state formalism, which, for the two-particle case we are considering, includes vectors such as $|a_1^h\rangle|a_2^t\rangle$. Or to put the point in another way, if the labeled tensor product Hilbert state formalism, complete with its particles labels, is generally meaningful, then in particular the vectors such

[9] Redhead and Teller also believe that LTI is necessary for quantum particles in order to ensure that they persist through time (Redhead and Teller, 1992, p. 203). But since objects need not persist in order to be objects, we can accept that quantum particles are objects, but deny that they persist through time, and therefore deny that they need LTI in order to ensure that they persist through time. This said, what is important about LTI in this surplus structure argument is not that it ensures persistence, but that it ensures labelibility.

> as $|a_1^h\rangle|a_2^t\rangle$ should be physically meaningful. (Redhead and Teller, 1992, p. 208)

Here $|a^h\rangle$ and $|a^t\rangle$ are orthogonal single-particle vectors, and the subscripts above serve as particle labels. The fact that unsymmetrized states of particles of the same kind are meaningful in some way presents an apparent conflict with quantum statistics and the symmetrization requirement. There is, however, a solution. As they write, "The conflict with quantum statistics is resolved by declaring that, though meaningful, such states and properties never occur" (Redhead and Teller, 1992, p. 209).

So quantum particles can be objects, so long as we accept that there are possible unsymmetrized states that never occur. But this is an unpleasant position to take. As Redhead and Teller argue:

> Physicists express an important methodological view with the maxim: 'What is possible happens'. Roughly speaking, this means that if a theory describes a state of affairs, then, insofar as the theory is correct, the state of affairs will be found to occur. Or if not, then at least there should be some specific account of why the phenomenon does not occur. For example, statistical mechanics describes a possible state in which an unheated cup of coffee comes to a boil. We never see such events, but then the theory accounts for this by giving an explicitly account of the event's exceedingly low probability. ... The non-symmetric states of the labeled tensor product Hilbert space formalism provides a flagrant violation of the maxim. Insofar as the formalism is correct, the states which it describes ought to be found in nature. They are not. (Redhead and Teller, 1992, p. 216–217)

Redhead and Teller describe these states, which are possible but never occur, as "surplus structure" (Redhead and Teller, 1992, p. 217). This surplus structure in the tensor product Hilbert space formalism would be acceptable if it were not for the fact that there is an alternative formalism free of this surplus structure, free of states that violate Redhead and Teller's methodological maxim, namely the Fock space formalism. In the Fock space formalism there are no explicit particle labels, and it is quite natural to take particles not to be objects, but rather discrete excitations of the quantum field that are free of LTI. As Redhead and Teller state:

> We have available an alternative, in fact empirically superior formalism, free of both LTI and the undetectable, non-occurring states and properties to which LTI appears to give rise. These facts give us strong methodological grounds for concluding that quantum objects do not have LTI. (Redhead and Teller, 1992, p. 218)

From this they conclude:

> Quantum entities differ from classical objects, not only by failing to have exact trajectories, but by failing to have an attribute transcending identity which would sustain at least conceptual individuation. (Redhead and Teller, 1992, p. 218)

That is quantum particles are not objects.

In the literature, there are two lines of criticism against this surplus structure argument. The first points out that we can express a tensor product Hilbert space description within a Fock space description for non-relativistic symmetrized states of fixed particle number, as we have already discussed. Because of this, we can extend any interpretation of tensor product Hilbert space states to the Fock space states to which they are equivalent. So even if we have methodological reasons for adopting a label free Fock space description, this does not indicate that quantum particles are not objects like they are in the tensor product Hilbert space formalism, for we can extend our interpretation of the latter to a subset of the former.

Huggett and van Fraassen both take this line against the surplus structure argument. As Huggett writes:

> So, I urge, if the QFT Fock representation is equivalent to the many particle formalism, then it too is consistent with individuals, whether we express states with or without labels. (Huggett, 1994, p. 73)

And as van Fraassen writes:

> All models of (elementary, non-relativistic) quantum field theory can be represented by (i.e. are isomorphic to) the sort of Fock space model constructions I have described above. Since the latter are clearly carried out within a 'labelled particle' theory, we have a certain kind of demonstrated equivalence of the particle—and the particleless—picture. (van Fraassen, 1991, p. 448)

Both Huggett and van Fraassen are absolutely right that there is a equivalence between the two formalisms that allows us to interpret both in the same manner. But this does not appear to counter the main point of the surplus structure argument. The argument is not that we should prefer Fock space over the tensor product Hilbert space formalism, and that the Fock space formalism tells us that particles are not objects. Rather, the argument is that if particles are objects, then there are possible unsymmetrized states that do not occur. These states present us with surplus structure that we should excise. We rid ourselves of this surplus structure when we no longer consider particles as objects.

The discussion of Fock space is somewhat unnecessary. It simply provides us with a clear way to understand how the particles are not objects, how they are instead discrete excitations of the fields. We can interpret particles as objects in the Fock space formalism, but this is beside the point for the surplus structure argument. Even if we interpret particles as objects in the Fock space formalism, there is still surplus structure; for we should still take unsymmetrized states as possible, even if they are not explicitly represented in the Fock space formalism.

This brings us quite naturally to the second line of criticism, which is more to the point. This line holds that the surplus structure argument turns on a confusion about possibility. The argument does not properly distinguish between logical possibility, physical possibility, and contingent features of the actual world.

Consider once again Redhead and Teller's methodological maxim: "What is possible happens." Certainly what is physically possible does not always happen. Now Redhead and Teller are speaking with a little bit of rhetorical flair. It is clear that they do not hold this maxim to be absolutely true. But they do demand an explanation of why physically possible events do not occur. In the quotation above, they present the example of the cup of coffee that spontaneously boils. This is a physically possible event that does not occur. Yet there is an explanation for why it does not occur: It does not occur because the statistical probability for such an event is minute. Redhead and Teller demand a similar explanation for why unsymmetrized states do not occur.

We might question whether such explanations are necessary for all events that are physically possible, but not actual. We might just accept without question that some things are physically possible, but do not occur, and that this is no more mysterious than the fact that Al Gore could have been president, but isn't. We can, however, put this question to the side, for there is another issue of greater concern.

Since unsymmetrized states violate the symmetrization requirement, they are not physically possible. Because they are not physically possible, we do not need to offer any explanation of why they do not occur. As Huggett writes, "there is no mystery at all about why non-symmetric states are never realised; they are not within the symmetised Hilbert space that correctly represents the world, and hence do not correspond to physical possibility" (Huggett, 1994, p. 74).[10]

The comparison to a boiling cup of coffee is misleading. In the case of a boiling cup of coffee, we are dealing with a physical possibility, although an improbable one. In the case of unsymmetrized states of particles of the same kind, we are not dealing with a physical possibility. While unsymmetrized states do exist in the tensor product Hilbert space formalism, they are artifacts of the representation that do not represent any physical possibilities for particles of the same kind. To quote Huggett once again, "the fact that we can express non-symmetric states in the same formalism as the symmetric ones hardly makes them possible in any interesting sense" (Huggett, 1994, p. 74–75).

[10]French (1998, p. 104) also presents a similar view.

Since unsymmetrized states do not present us with non-occurring physical possibilities, we should not regard them as surplus structure in this respect. And without surplus structure, Redhead and Teller's argument presents no reason to claim that quantum particles are not objects.

This is a valid objection against their argument as it is stated. But we can reconstrue their argument in such a way that it avoids this objection. The important aspect of Redhead and Teller's argument is that LTI requires that we take the permutation of particles as producing distinct possibilities represented by unsymmetrized vectors. Because of the permutation symmetry of quantum theory, we cannot regard these possibilities as physical possibilities. But we can ask why. Why do unsymmetrized vectors not represent physical possibilities? Why is quantum physics permutation symmetric? On this reconstrual it really does not matter that unsymmetrized vectors are not physically possible, for the question is now why they are not physically possible.

It is clear how Redhead and Teller would answer these questions as we have stated them. Unsymmetrized vectors are not physical possibilities because quantum particles do not possess LTI. Quantum theory is permutation symmetric because quantum particles are not objects. On this reconstrual, Redhead and Teller's argument is an attempt to explain the permutation symmetry of quantum theory by claiming that quantum particles are not objects. We have some reason to think that our reconstruction of their argument is one they would endorse, for in a later paper they claim that the true concern of their argument is "exchange degeneracy" (Teller and Redhead, 2000, p. 954). And of course this argument is very similar to Teller's other views, which we have discussed in chapter 5.

It is important to realize that, even if we accept this argument, it does not provide an explanation of the symmetrization requirement. As we have pointed out in chapter 4, symmetrization is a stronger requirement than permutation symmetry. So even if Redhead and Teller succeed in providing an explanation of permutation symmetry, they do not offer any explanation of why only symmetrized states are physically possible. We must still take this as a further empirical hypothesis.

If we are to take Redhead and Teller's surplus structure argument as an attempt to explain permutation symmetry by claiming that quantum particles are not objects, then their surplus structure argument is really no different than the ontological argument from permutation symmetry. And we have already discussed how this argument fails. Permutation symmetry is not connected to a change in ontology. Rather, it is connected to anti-haecceitism. And we have also already discussed the major flaw in Redhead and Teller's argument: the objecthood or individuality of quantum particles does not require any LTI that leads to haecceitistic differences. Objects need not be labeled. And particle labels in a Hilbert space formalism need not be interpreted as proper names, much less proper names that rigidly refer over different counterfactual contexts.

But we can take something away from Redhead and Teller's argument. There is something to their comments on surplus structure. As they write:

> When we have a second formalism which is free of what are at least apparently surplus elements in a first formalism, and when the second formalism not only has all the correct empirical import of the first but also covers phenomena not covered by the first, we surely have strong grounds for judging the apparently surplus structure of the first formalism to be really surplus. (Redhead and Teller, 1991, p. 56)

If we take surplus structure as artifacts of a description, which do not represent anything in the physical world, then we must agree with their statement. For one way we can tell that such elements truly are artifacts is if there is an adequate alternative formalism where they do not occur.

In this way a Fock space description does demonstrate that unsymmetrized states are simply artifacts of a tensor product Hilbert space description. This relation between the two descriptions is analogous to the relation between a phase space description and a Z-space description of classical dynamics. Since a Z-space description provides an adequate description of classical dynamics without specifying, as a phase-space description does, which particle is traveling along which single-particle trajectory, such information is simply an artifact of the phase space description. This does not necessarily mean that we should prefer either a Fock space or Z-space description. We can use a tensor product Hilbert space or phase space description so long as we realize what elements of those descriptions fail to represent.

This is how we should understand surplus structure.[11] When we understand it as such, it does not provide us with any reason to hold that quantum particles are not objects.

6.3 Symmetrization and Correlations

There is another ontological argument related to fields that deserves our attention. The gist of this argument is that the statistical correlations produced by the symmetrization requirement indicate that a field ontology is the proper ontology for all of quantum theory.

Let us start by pointing out what correlations we have in mind. Statistical mechanical ensembles of symmetrized states encode positive and negative correlations between particles, even if those particles do not interact. Ensembles of symmetric states encode positive statistical correlations, while ensembles of antisymmetric states encode negative correlations. There are no such correlations for ensembles of uncorrected unsymmetrized ensembles.

To illustrate this, consider once again a simple system of two particle of the same kind with two single particle eigenstates $|H\rangle$ (heads) and $|T\rangle$ (tails) of the same energy. There are three symmetric states for this system:

$$|\Psi_1\rangle = |H\rangle|H\rangle \tag{6.3}$$

$$|\Psi_2\rangle = |T\rangle|T\rangle \tag{6.4}$$

[11]Those familiar with the details of Redhead and Teller's argument will recognize that this view of surplus structure is similar to what they call weakly surplus structure, which they take to be uninterpreted elements of a formalism. This is opposed to strongly surplus structure, which they take to be interpreted, but not actual, elements of the formalism (see Teller, 1995, p. 25–26). In their argument they consider unsymmetrized states to be examples of strongly surplus structure when we interpret particles as objects. As we have said above, this is not the case. Unsymmetrized states do present us with strongly surplus structure, but only weakly surplus structure.

$$|\Psi_3\rangle = \frac{1}{\sqrt{2}}(|H\rangle|T\rangle + |T\rangle|H\rangle) \tag{6.5}$$

The microcanonical ensemble of this system assigns an equal probability to each state and is described by the following density matrix (which we have not yet normalized):

$$\rho = |\Psi_1\rangle\langle\Psi_1| + |\Psi_2\rangle\langle\Psi_2| + |\Psi_3\rangle\langle\Psi_3| \tag{6.6}$$

Let $\text{Prob}(H_1)$ be the unconditional probability for this statistical ensemble that particle 1 is heads. Further let $P_H = |H\rangle\langle H|$. Given this:

$$\begin{aligned} \text{Prob}(H_1) &= \langle P_H \otimes I\rangle && (6.7)\\ &= \frac{1}{3}\text{Tr}(\rho(P_H \otimes I)) && (6.8)\\ &= \frac{1}{3}\text{Tr}(\rho_1 P_H) && (6.9)\\ &= \frac{1}{2} && (6.10) \end{aligned}$$

where I is the single-particle identity operator and ρ_1 is the reduced density matrix for particle 1. Now $P_H \otimes I$ is not a permutation symmetric operator, and therefore we might question whether it is actually an observable. But let us for the moment put this objection to the side.

Let $\text{Prob}(H_1|H_2)$ be the conditional probability for this ensemble that particle 1 is heads given that particle 2 is heads. This probability is as follows:

$$\begin{aligned} \text{Prob}(H_1|H_2) &= \frac{\langle P_H \otimes P_H\rangle}{\langle I \otimes P_H\rangle} \\ &= \frac{2}{3} && (6.11) \end{aligned}$$

Given these probabilities, there is a positive statistical correlation between the particles: $\text{Prob}(H_1|H_2) > \text{Prob}(H_1)$.

In the antisymmetric case there is only one state for this simple system:

$$|\Psi\rangle = \frac{1}{\sqrt{2}}(|H\rangle|T\rangle - ||T\rangle|H\rangle) \tag{6.12}$$

For the microcanonical ensemble that consists of this one state, $\text{Prob}(H_1) = 1/2$ and $\text{Prob}(H_1|H_2) = 0$. Since $\text{Prob}(H_1|H_2) < \text{Prob}(H_1)$, there is a negative statistical correlation.

In the unsymmetrized case, there are four states for this simple system:

$$|\Psi_1\rangle = |H\rangle|H\rangle \tag{6.13}$$
$$|\Psi_2\rangle = |T\rangle|T\rangle \tag{6.14}$$
$$|\Psi_3\rangle = |H\rangle|T\rangle \tag{6.15}$$
$$|\Psi_4\rangle = |T\rangle|H\rangle \tag{6.16}$$

For a ensemble consisting of these states, $\text{Prob}(H_1) = 1/2$ and $\text{Prob}(H_1|H_2) = 1/2$. Since $\text{Prob}(H_1|H_2) = \text{Prob}(H_1)$, there is no statistical correlation between the particles.

Of course these statistical correlations are not the only correlations that we find in quantum theory. Ensembles of entangled states will express similar correlations, as we see in descriptions of EPR or Bell-like experiments. But here we are focusing specifically on the correlations expressed by symmetrized statistical mechanical ensembles, correlations that follow from quantum statistics.

The importance of correlations was known early on in the development of quantum statistics. In his second paper on Bose-Einstein statistics, Einstein wrote that these statistics "express indirectly a certain hypothesis on a mutual influence of the molecules which for the time being is of a quite mysterious nature" (Einstein, 1925).[12] (Of course Einstein is here working in a quasi-classical framework, where correlations cause bosonic particles to group together in phase space cells.[13])

[12]The English translation of this quotation is taken from Pais (1982, p. 430).

[13]Costantini et al. (1982) and Costantini and Garibaldi (1989) have presented in detail how such correlations are connected to quantum distribution functions in a quasi-classical

The problem with these statistical correlations, what makes them "mysterious" in the eyes of Einstein and others, is that there is no apparent explanation for them. We cannot explain them by reference to some common cause or local interaction. They exist even for an ideal gas, where there is no interaction between the particles. And they exist for ensembles of states where the particles are widely separated throughout their histories.

This brings us to the ontological argument we wish to consider, which goes as follows. There is no acceptable causal explanation for these statistical correlations between quantum particles. The best way to understand these correlations is to abandon the idea that the quantum system is composed of particles and instead consider a quantum field description of the system. When we take fields as our ontology, we no longer have to worry about how to explain the correlations between particles, because there are no particles to correlate. Instead we talk about the allowable states of the whole field.

Another way of stating the argument is in terms of reduced density matrices. The state of each individual particle is described by its reduced density matrix. The reduced density matrix gives the unconditional probabilities associated with the particle. Now the reduced density matrices of the particles in the system do not always determine the state of the whole system. The unconditional probabilities of the particles therefore do not determine the conditional probabilities of the system. This leads to correlations.

Now the easiest way to explain away these correlations is by dropping any talk of particles and their reduced states. Instead there is just the whole system described by a single state. Conditional and unconditional probabilities do not then express properties associated with particles, but rather express different sorts of measurement on the state of the whole field, the probabilities of which are encoded by the state of the field. There are still correlations of a sort, for the conditional and unconditional probabilities are not equal. But this no longer reflects any correlation between particles, just the different probabilities of different measurements of the same state of the field.

The idea at work here is that there is no reason to use a particle ontology

formalism.

if we cannot understand the state of the system in terms of the states of the individual particles. This just leads to strange correlations between the particles. Better to consider the system as a single field with a single state.

Both Reichenbach (1998) and Dieks (1990) put forward a version of this argument. They both, however, combine it with a version of the ontological argument from permutation symmetry; for they both claim that if we take particles as objects, then there are an excess number of states that prevent us from assigning an equal probability to each state. We have already dismissed the ontological argument from permutation symmetry. Let us now focus on their views about correlations.

Reichenbach's concern is the genidentity or cross-temporal identity of quantum particles. He argues that the statistical correlations between particles lead to causal anomalies since they defy explanation in terms of common cause or interaction. He claims that if we are to describe physical systems without such causal anomalies, then we must deny the genidentity of quantum particles. As he writes:

> Conversely, if we accept the synoptic principle [which holds that we cannot simultaneously specify the position and momentum of quantum particles], its combination with Bose statistics leads to the result that there is no material genidentity for elementary particles, provided we describe the physical world in terms of a normal system, that is, a system without causal anomalies. (Reichenbach, 1998, p. 72)

As we have said, Reichenbach's focus is on the genidentity of quantum particles. Persistence, however, is not really the issue here. The correlations behind Reichenbach's causal anomalies still exist for particles that do not persist through time. Causal anomalies throw doubt not on the persistence of quantum particles, but on quantum particles themselves.

Reichenbach does not actually claim that we must or even should prefer a description free of causal anomalies (Reichenbach, 1998, p 71). He only points out that such anomalies are inevitable for a particle interpretation. Dieks is more direct in his assault on quantum particles. As he writes:

> Quantum statistics is the immediate result of the 'natural', uniform, distribution over these states of quantum field theory. There are no correlations between the states according to this statistical distribution. The impression of a correlation between states is only created if individual particle states are artificially added to the theoretical framework of quantum field theory. (Dieks, 1990, p. 131)

It clear that Dieks views these statistical correlations as the byproduct of an artificial particle ontology, and that the correct field ontology is free of any such confusion.

Now what should we make of this ontological argument? The argument does not establish that it is inconsistent to take particles as the ontology of quantum mechanics. Rather, it points out the perceived deficiencies of doing so. If we take particles as our ontology, then we must accept "mysterious" correlations between the particles that indicate some from of holism, a holism where the state of the system is not determined by the states of the particles. Such mysteries are avoided when we accept fields as the proper ontology.

The obvious response to this argument is to reject the need for an explanation of these correlations. Unlike in classical physics, we cannot always determine the state of a quantum system from the states of its quantum particles. This is one of the important ways in which quantum physics differs from classical physics. To this we must assent. But from this we need not agree that quantum particles present us with a confused ontology. Granted, when we accept quantum particles, we need to accept a minimal form of holism, and with it statistical correlations that do not have any apparent causal explanation. But this lack of causal explanation does not provide a crippling blow to the particle picture, for unexplained correlations pervade quantum theory. Statistical correlations might not have any clear explanation in terms of interactions or common causes, but neither do EPR correlations, as Bell's inequalities show. And these latter correlations exist whether or not we switch to a field ontology.[14] And so the "mysterious" nature of these

[14] This appears to be the core of van Fraassen's response to Reichenbach's argument expressed in (van Fraassen, 1998, p. 88) and (van Fraassen, 1991, p. 372–374).

correlations alone are not enough to undermine the potential use of a particle ontology as an emergent ontology; for the mystery of correlations, if there is a mystery, cloaks all of quantum theory. Since correlations are not limited to the particle picture, correlations do not give us reason to abandon a particle ontology if it is useful as an emergent ontology.

6.4 Underdetermination of Ontology

Hopefully at this point we have established both the possibility and plausibility of a particle ontology where particles are objects. But there is another important issue that we must address—the underdetermination of ontology. For while it is possible to take particles as objects, it still might be the case that it is underdetermined whether the proper ontology is one where particles are objects.

Following Ladyman (1998, p. 419), let us call this type of underdetermination "metaphysical underdetermination". We must distinguish it from other forms of underdetermination. We are not talking about how our choice of a theory is underdetermined by evidence. Rather, we are talking about how our choice of ontology is underdetermined by a given theory.

Metaphysical underdetermination clearly threatens any realist interpretation of a theory. If there is nothing about a theory that suggests one ontology over another, then any choice of ontology appear to be arbitrary or conventional. This will undermine any claim that the chosen ontology is the actual ontology of the world. As Ladyman puts it:

> We need to recognize the failure of our best theories to determine even the most fundamental ontological characteristic of the purported entities they feature. It is an *ersatz* form of realism that recommends belief in the existence of entities that have such ambiguous metaphysical status. (Ladyman, 1998, p. 419-420)

Ladyman and French have proposed several examples of metaphysical underdetermination (Ladyman, 1998; French and Ladyman, 2003b). One of their main examples is the putative underdetermination between individual and non-individual particles. This obviously is the example that is of most interest to our study, and we will examine it in greater detail shortly. But they have other examples. They take it to be metaphysically underdetermined whether the proper ontology of quantum theory is a particle or field ontology. Assuming a field ontology, they also take it as metaphysically underdetermined whether fields are substantive objects in their own right or merely properties of spacetime points. Ladyman has also suggested that in general relativity it is metaphysically underdetermined whether spacetime is relational or substantival (Ladyman, 1998, p. 420).

Ladyman and French have claimed that these examples of metaphysical underdetermination support ontic structural realism—which, they claim, takes structures, and not objects, as the proper ontology of physical theories. The basic idea is that each of these opposing ontologies are just two expressions of a single ontology of structures. As they put it:

> The locus of this metaphysical underdetermination is the notion of an object so one way of avoiding it would be to reconceptualise this notion entirely in structural terms. The metaphysical packages of individuality and non-individuality [of quantum particles] would then be viewed in a similar way to that of particle and field in QFT, namely as two different (metaphysical) representations of the same structure. (French and Ladyman, 2003b, p. 37)

So by taking structure as the proper ontology of physics, we can resolve any metaphysical underdetermination in the examples above.

These examples of metaphysical underdetermination do not command immediate consent. Cao denies that there is any troublesome underdetermination between particles and fields. He points out, as we have above, that "the basic ontology of quantum field theory can only be the quantum fields," and that particles are "objective but not primitive entities" (Cao, 2003, p. 63).[15]

[15] In response to Cao, French and Ladyman (2003a, p. 74) do seem to back down a bit from their claim that there is a troublesome case of metaphysical underdetermination between particles and fields. But they still maintain the correctness of their other examples.

Cao is right on this point. While we can describe non-relativistic states in terms of both particles and fields, this sort of underdetermination does not undermine a realist interpretation of particles; for particles are only introduced as an emergent ontology. They are real in so far as they are successful in this role. And as we have already argued above, particles do succeed in providing such an ontology, even when we acknowledge the possibility of redescribing states in terms of fields. So the underdetermination between particles and fields does not offer any support to structural realism, for it does not undermine the reality of particles.

It is harder to dismiss the example of quantum fields. It is open question as to whether quantum fields are substantial objects in their own right or simply properties of spacetime points. And of course our answer to this questions depends upon whether or not we hold that spacetime points actually exist. If we take quantum fields to be the basic ontology of the physical world, any metaphysical underdetermination between these two views about their nature would undermine a realist interpretation of physics. But this question as to the nature of quantum fields only provides a troublesome example of underdetermination if the two sides are hopelessly deadlocked. It does not seem like this is the case, for we have some reason to believe that the progression of quantum theory and its philosophy will throw some light on the nature of quantum fields, perhaps by replacing them with an even more fundamental ontology.[16] So it seems perfectly acceptable, for the time being at least, to accept the reality of quantum fields, even though their nature is still open to debate.

As for spacetime points, Pooley (forthcoming) denies that there is any metaphysical underdetermination between relationist and substantivalist views of spacetime in general relativity. His argument is that the standard formulation of general relativity is straightforwardly substantivalist. On this point he is very persuasive.

This leaves us with one more example—the metaphysical underdetermi-

[16]Cao also denies that the nature of quantum fields presents a legitimate example of troublesome metaphysical underdetermination, but for slightly different reasons (Cao, 2003, p. 63).

nation between individual and non-individual particles. We can question whether any such underdetermination is problematic. Since particles do not provide us with the fundamental ontology of quantum physics, we need not commit ourselves to one of these two positions in order to consider particles as real. Once again, since particles only provide an emergent ontology, they are real in so far as they provide a useful ontology. If both individual and non-individual particles succeed in this respect, then we can regard both as real, even if our choice between the two is arbitrary.

But this example fails in a more fundamental way, for there is no underdetermination between individual and non-individual particles. Now both Cao (2003, p. 61) and Pooley (forthcoming) agree with this conclusion, but for different reasons than us. They both claim that, while it might be possible to regard particles as individuals, it is not plausible. No doubt they think that the arguments that we have discussed and dismissed above establish this conclusion. By showing that it is not only possible, but reasonable to regard quantum particles as objects, it might seem to some that all we have done is strengthen the claim of metaphysical underdetermination. After all we have overturned the arguments about particle identity that Cao and Pooley rely on to alleviate this underdetermination. The individuality of quantum particles is no longer just a possibility, but a legitimate alternative on a par with the ontology of non-individual quantum particles. The metaphysical underdetermination between individual and non-individual particles is even more entrenched.

But this is not so. Having alleviated any worries about a principle of individuation, permutation symmetry, quantum statistics, and surplus structure, we are left without any reason to doubt that quantum particles are objects, and any clear way to understand how they can fail to be objects. While we might have a formal apparatus like quasi-set theory that is capable of describing things without self-identity, we have no reason to apply it to quantum particles since the identity of such particles is secure. So it is not just that we can regard particles as objects, we should. There is thus no metaphysical underdetermination between individual and non-individual particles. Whenever we can adopt a particle ontology, the particles of that ontology clearly are objects.

While we have rejected these examples of metaphysical underdetermination, we have not demonstrated that ontic structural realism is incorrect. The alleviation of metaphysical underdetermination is only one of the proposed benefits of ontic structural realism. Even if there is no metaphysical underdetermination, ontic structural realism still might be the best way to resolve the apparent conflict between the no miracles argument and the pessimistic meta-induction.[17] But if we do accept structural realism, we must do so in a way that allows us to grant that quantum particles are still objects, at least in some physical regime.

To conclude, a particle ontology does not (currently) provide our most fundamental ontology of quantum theory. It serves merely as an emergent or effective ontology. But this does not undermine the reality or individuality of particles. Nor does the fact that we can in some sense replace a particle ontology with a field ontology. Particles still provide a useful description of certain quantum phenomena. And as our arguments show, these particles are objects.

[17]The no miracles argument holds that the empirical success of a scientific theory is miraculous if the theory does not accurately describe some aspect of reality beyond the empirical phenomena. The pessimistic meta-induction concludes from the fact that all previous scientific theories have turned out to be false that our current scientific theories are also false. See Ladyman (1998) for further details of how ontic structural realism is suppose to reconcile these two arguments.

Chapter 7

Quantum Particles and Ontic Vagueness

7.1 The Question of Ontic Vagueness

All admit that it is often vague whether a given man is bald, but not all agree why. It is safe to say that the most prevalent opinion among philosophers is that vagueness has its origin in our language and concepts, that it is linguistic; although some influential philosophers take vagueness to be a sort of ignorance (see Williamson, 1994). But there is also another theory, which is at the center of a rapidly growing debate. This theory holds that vagueness has its origin in the world, that it is ontic. The idea is that a given statement is vague not simply because of the terms and concepts involved, or because of some sort of ignorance, but because it describes something or some object in the world that actually is vague.

Two of the ways objects might be ontically vague is by having either indeterminate existence or indeterminate identity. In the former case it is a vague matter whether the objects exist. In the latter case it is a vague matter whether the objects are identical to other objects. In both cases the vagueness originates not in the description or knowledge of the situation, but in the nature of the objects themselves.

Following the work of Lowe (which we will shortly consider), many wonder whether quantum theory might provide coherent and plausible examples of ontic vagueness. The idea is that quantum particles of the same kind might be ontically vague by having either an indeterminate existence or indeterminate identity. In this chapter we will examine whether or not this is so. We will start by considering quantum examples of indeterminate existence and then move on to consider quantum examples of indeterminate identity. We will conclude in both cases that, regardless of the plausibility or coherency of ontic vagueness in general, quantum particles do not provide clear examples of ontically vague objects.

This topic is obviously relevant to debates on vagueness, but it is also relevant to our previous arguments. Let us say it is true that in at least some circumstances quantum particles are vague objects. There would then be good reason to doubt our claims that quantum particles are objects in the same way as classical particles, for an ontology of vague objects is very different from the ontology we have argued for, and is a significant departure from the ontology of classical particles. If quantum particles are ontically vague in a way that is different from other classical objects, then our previous arguments are at best incomplete, and at worst wrong. But now let us turn to our topic.

7.2 Indeterminate Particle Number and Indeterminate Existence

States of indeterminate particle number provide one putative quantum example of ontic vagueness. If we take such states to describe particles, then we have an indeterminate number of particles. We can take this to follow from the fact that the particles do not have a determinate existence. That is their existence is vague, and this vagueness follows not from our description or knowledge of the particles, but from the nature of the particles themselves. As in the previous chapter, we will call particles of indeterminate number "quanta". As far as I know, none have taken quanta as an example of indeterminate existence; but if we are to take the possibility seriously that quantum particles are ontically vague, as some no doubt do, then this example deserves some consideration.

This example, however, is easy to dismiss. As we have discussed in the last chapter, our understanding of quanta is parasitic on a field interpretation. We take states of indeterminate particle number to be states of fields and not particles. And once we switch to a field interpretation, there is no longer any indeterminate existence. Fields just exist. So unless we have some plausible way of understanding quanta independently of a field interpretation, quanta cannot provide a plausible example of indeterminate existence.

But perhaps we can turn things around and understand quanta independently of fields in terms of ontic vagueness. We mentioned this possibility in the last chapter. Instead of offering quanta as an illustration of indeterminate existence, we use indeterminate existence to illustrate the nature of quanta.

The main problem with objects of indeterminate existence is that they appear to reside in some incoherent no man's land between being and non-being. Something either exists or it does not. There is no third option.[1] Quanta are not immune from this line of objection. If a particle, or any other object, is to possess any property like mass or charge, or be a part of any system, then that particle must exist; for if it does not exist, then it cannot possess any property or be a part of anything.

Now van Inwagen agrees that there are no objects that "dwell in the twilight between the full daylight of Being and the night of Non-being" (van Inwagen, 1990, p. 277). Yet he holds that some objects might nonetheless have an indeterminate existence. Van Inwagen is led to the topic of indeterminate existence by his views on composition. He does not believe in unrestricted composition. We cannot gather any objects we please into a mereological sum that is itself an object. Only some sets of objects actually compose another object. And he admits that it is sometimes vague whether a set of objects actually does compose something. For example it is sometimes

[1] We are here equating being with existence. Those who accept an Meinongian view of objects separate the two. But even for a Meinongian, there is no indeterminate being or indeterminate existence.

vague whether a collection of bricks compose a pile, or when a collection of cells compose a life. For an object that is a pile or a life, it will sometimes be indeterminate whether the object composed—the pile of bricks or the life—exists. So indeterminate existence does not describe objects between being and non-being, but indeterminate compositions. As he writes:

> If there really were borderline objects [between being and non-being], one might focus one's attention upon one of them and say, "It is neither definitely true nor definitely false of *that* that it exists." And that is nonsense. What there really are, however, are sets such that it is not definitely true and not definitely false of their members that they compose anything. (van Inwagen, 1990, p. 277)

We can question whether van Inwagen's vague compositions manage to avoid the no man's land (or twilight) between being and non-being. If a composition is as much an object as those objects which compose it, then it should have as determinate an existence as those objects. And of course we can question the thesis of restricted composition. If we accept unrestricted composition, then we can explain the vagueness associated with composition as linguistic vagueness instead of ontic vagueness. For if we accept unrestricted composition, then any set of objects compose an object. What is vague is whether that composition is of a given sort. Lewis is an influential proponent of this view. As he writes:

> Restrict quantifiers, not composition. Vague existence, speaking unrestrictedly, is unintelligible; vague existence, speaking restrictedly, is unproblematic. . . . There is a sum, unrestrictedly speaking, but it can perfectly well be a vague matter whether this sum falls within a vaguely restricted domain of quantification. (Lewis, 1986, p. 213)

But even if we put these objections to van Inwagen's position to the side, we cannot understand the indeterminate existence of quanta in the way van Inwagen outlines; for what is indeterminate is the quanta themselves and not their composition. So even if we agree with van Inwagen and hold that compositions can have an indeterminate existence without existing in the no man's land between being and non-being, we cannot extend this conclusion to quanta. If we take quanta as entities with indeterminate existence, then they appear to have no other place to occupy than the no man's land, no other place to bathe but in the twilight between being and non-being. This is incoherent.

So quanta fail to provide a clear example of indeterminate existence. And indeterminate existence, insofar as we can understand it in a coherent manner, fails to offer an illumination of the nature of quanta. We still can only understand quanta in terms of excitations of fields, and this does not present an example of indeterminate existence.

7.3 Lowe's Quantum Example

Now let us move on to discuss putative quantum examples of indeterminate identity. These examples deal with quantum particles of determinate number. They therefore concern the properties of a particle ontology that is independent of a field ontology.

Lowe (1994) was one of the first to introduce quantum examples into the debate of ontic vagueness. We will come to Lowe's quantum example in a bit, but first let us make some preliminary remarks. In (Lowe, 1994) and his subsequent work on the subject (Lowe, 1997, 1998, 1999, 2001), Lowe's target is the argument by Evans (1978) that purports to show the impossibility of indeterminate identity.

Evans's argument consists of the following derivation (Evans, 1978).[2] Let "$\bigtriangledown$" be a sentential operator that expresses indeterminacy. Further let "$\hat{x}[Fx]a$" be an expression of property abstraction that states a has the property of F. Assume that "a" and "b" are singular terms with precise ref-

[2]Two things to note. First, Copeland (1997, p. 515) informs us that the target of Evans's argument was Dummett, who put forward the idea that vagueness might be part of the world. Second, Salmon (1982, p. 243–246) independently presented an argument very similar to Evans's.

erence, at least at a given time. Start by assuming that it is indeterminate whether a is identical with b, so that:

$$\bigtriangledown(a = b) \tag{7.1}$$

Given this, b possesses the property of being indeterminately identical to a:

$$\hat{x}[\bigtriangledown(x = a)]b \tag{7.2}$$

But a is determinately identical to itself:

$$\neg \bigtriangledown (a = a) \tag{7.3}$$

and thus does not possess the property of being indeterminately identical to a:

$$\neg\, \hat{x}[\bigtriangledown(x = a)]a \tag{7.4}$$

But by Leibniz's law, (7.2) and (7.4) imply that a and b are not identical:

$$\neg(a = b) \tag{7.5}$$

which conflicts with our claim that it is indeterminate whether a is identical with b. The apparent conclusion is that we cannot consistently regard identity as indeterminate.

This conclusion is not above suspicion. As a proponent of indeterminate identity, there are several objections we might raise. For example we might deny that ontically vague objects have determinate self-identity and reject (7.3) (we will return to this view below). Or we might deny property abstraction in the case of ontically vague objects. Or we might hold that an alternative logical system is necessary in order to properly describe ontic vagueness. We might adopt a three-valued system where a sentence can be true, false, or indeterminate. Or we might adopt fuzzy logic where there are different degrees of truth. In such systems we cannot reproduce Evans's derivation. Van Inwagen (1990, ch. 18) and Parsons and Woodruff (1995) show that Evans's derivation is invalid for a three-value semantics, and

Copeland (1997) shows that it is invalid in fuzzy logic. So it appears that Evans's argument is not a knockdown argument. It does not by itself demonstrate the incoherency of ontic vagueness. It does, however, demonstrate the cost.

Lowe's stated goal in presenting his quantum example is simply to demonstrate the coherence of indeterminate identity, and thereby show that something is amiss in Evans's argument. Lowe is not claiming that we should regard the objects in his quantum example as ontically vague, or that anything is ontically vague. He is only claiming that there is a coherent interpretation of quantum theory where quantum particles are ontically vague (Lowe, 2001).

But we cannot be satisfied with this modest goal. Since there is already some doubt about Evans's derivation, a questionable quantum example is not going to do any more to demonstrate the coherency of ontic vagueness. What we require is a clear and plausible quantum example of ontic vagueness, one that not only demonstrates the coherency of quantum vagueness, but also provides a reason for thinking that the world—at least the quantum world—is ontically vague.

With this in mind, let us move on to consider Lowe's putative quantum example of ontic vagueness. Lowe presents slightly different versions of the example in his various articles. We will focus on his latest (Lowe, 2001). Lowe asks us to consider the absorption of an electron by a helium atom that already contains one electron. Before the absorption event, the electrons are distinct, and we can uniquely refer to each by definite description. After the absorption event, the electrons become entangled, and this entanglement leads to ontic vagueness. As he writes:

> Let t be a time shortly before the absorption event and let t' be a time shortly after that event. And assume, to avoid unnecessary complications, that the particles under discussion are far removed from any other particles. Then, I want to say, the definite description 'the electron that is outside the helium ion' is a precise designator at t which can be used to fix precisely the reference at t of the name 'a' and, likewise, the definite description

> 'the electron that is inside the helium ion' is a precise designator at t which can be used to fix precisely the reference at t of the name 'a^*'. I also want to say that at the later time t' both electron a and electron a^* still exist. However, at t', I want to say, the names 'a' and 'a^*' are no longer precise designators, because there is then no fact of the matter as to which of the two electrons in the helium atom is a and which is a^*. What I want to deny, then, is that, at t', there is a fact of the matter as to which of the two electrons then contained in the helium atom was formerly outside it and which of them was formerly already inside it. Consequently, I also want to deny that the reason why the names 'a' and 'a^*' are imprecise designators at t' is only that we then have no means of *knowing or deciding* which of those electrons is picked out by the name 'a' and which by the name 'a^*'. In short, I want to deny that the source of the imprecision is merely 'epistemic' or 'semantic' and say instead the it is 'ontic'—that it arises from the absence of a suitable fact of the matter rather than from ignorance or indecision on our part. (Lowe, 2001, p. 241–242)

What is indeterminate for Lowe is the diachronic or cross-temporal identity of the two electrons. It is indeterminate which electron after the absorption event is identical with the electron absorbed. This indeterminacy is not due to our knowledge or any sort of semantic indeterminacy, but is rather due to the entanglement of the two electrons after absorption, which is an objective feature of the world.

Now we can say several things about Lowe's example and argument, but really we only need to mention one thing: it is an incorrect quantum description.[3] As French and Krause (1995) point out, the quantum state before absorption is just as entangled as the state after absorption. The source of the quantum entanglement is the antisymmetrization of the state of the two electrons. The state of the two electrons is antisymmetric at both time t and t'. This will be the case for any example that consists of fermions of the same

[3]For further criticism of Lowe's position see Noonan (1995), Hawley (1998), and Odrowąż-Sypniewska (2001).

kind.

The consequence of this is that if "a" and "a^*" fail to be precise designators at time t' because of entanglement, then they also fail to be precise designators at time t. So if something is ontically vague, it cannot be which of the electrons at time t' is identical with the absorbed electron a at time t, as Lowe claims; for it is already indeterminate at time t which of the electrons is identical to a. And since the singular terms at both times do not have a precise reference, we can claim that any indeterminacy of identity statements follows from the imprecision of reference. The vagueness is therefore linguistic, not ontic. Something is amiss in Lowe's informal example. This does not mean that there clearly is no ontic vagueness, but we need to say more to settle the matter.

7.4 Synchronic Vagueness of Quantum Particles

While Lowe's example does not demonstrate that quantum particles are ontically vague, it does indicate a potential source of vagueness—entanglement. French and Krause (1995, 2003) build upon this position and propose, like Lowe, that entanglement leads to ontic vagueness for particles that are objects. However, unlike Lowe, they believe that entanglement leads to synchronic and not diachronic ontic vagueness. Because particles are constantly entangled, it is identity statements between particles at a given time that are indeterminate. This is because entanglement leads to external relations (or as they call them non-supervenient relations), relations that in turn lead to ontic vagueness.

Recall that an external relation is one that does not supervene upon the intrinsic properties of the relata. In general, particles in entangled states will enter into such external relations. We have already presented an example of this in chapter 2. In that chapter we considered two non-interacting spin-1/2 fermions of the same kind with an antisymmetric spin component of their state-function. Because of the entanglement that followed from the antisymmetrization, these two fermions entered into the relation "has spin component in the opposite direction from", a relation that did not supervene upon the intrinsic properties of the particles and was therefore external.

For French and Krause, external relations lead to ontic vagueness because they mask which particle is which. Identity relations are therefore indeterminate. And this indeterminacy in not a matter of ignorance, but rather an objective feature of the quantum world. As they write:

> One view might be to say that, given that the particles are individuals, their identity is perfectly determinate, only, because of the existence of non-supervenient relations, we cannot *tell* whether electron a is identical to b, or not. However, we cannot *in principle* tell this; assuming that quantum mechanics is correct, we cannot tear away the veil of non-supervenience and get at what is 'really' going on. It is not an epistemic problem but an ontic one. (French and Krause, 1995, p. 22)

We have already noted that the external relations like "has spin component in the opposite direction from" ensure that quantum particles are weakly discernible. This weak discernibility seems to be the relevant feature for French and Krause's argument. Ontic vagueness does not follow from external relations alone, but the fact that the particles that enter into these external relations are only weakly discernible; for it is because the quantum particles are only weakly discernible that we cannot say which is which by use of singular terms.

But relations that follow from quantum entanglement are not the only external relations, nor are quantum particles the only weakly discernible objects. As we noted in chapter 2, spatiotemporal relations are also external. And spatiotemporal relations like "distance d apart measured center to center" allow us to weakly discern between otherwise indiscernible classical objects like Black's two spheres. So if quantum particles are ontically vague because they are weakly discernible, then we also expect other weakly discernible objects like Black's two spheres to be ontically vague; for Black's spheres enter into the same sort of external relations, and we cannot say which sphere is which.

But Black's spheres do not appear to possess indeterminate identity. Black implicitly regards them as identical to themselves and distinct from each other; for this is why they are able to serve as a putative counterexample to the principle of identity of indiscernibles. So unless we want to say that, despite our initial belief to the contrary, classical objects like Black's sphere actually are ontically vague in the same way as entangled quantum particles, we need to explain why there is ontic vagueness in the quantum case and not in the classical case.

Now French and Krause acknowledge that spatiotemporal relations are external, but they claim that there are "differences between spatio-temporal relations and those exemplified by entangled states" (French and Krause, 2003, p. 103). We might be able to hold that these differences explain the difference in ontic vagueness.

French and Krause cite two ways that spatiotemporal relations differ from the external relations of quantum states. First, unlike spatiotemporal relations, external quantum relations are not discriminating, and "thus they are not 'analogous' to spatio-temporal relations" (French and Krause, 2003, p. 103). Here they are making use of Lewis's definition of discriminating relations. Lewis takes a set of relations to be discriminating if and only if "it is at least possible, whether or not it happens at every world where the relations are present, that there be a great many interrelated things, no two of which are exactly alike with respect to their place in the structure of relations" (Lewis, 1986, p. 76). We can restate Lewis's definition of discriminating relations in terms of weak discernibility and indiscernibility. A set of relations are discriminating if and only if the objects that enter into those relations are not always weakly discernible or indiscernible; for if the objects are always weakly discernible or indiscernible, they will always share all the same properties and relations.[4]

Spatiotemporal relations between classical objects are obviously discrim-

[4]It might be that Lewis is more concerned with indiscernibility than weak discernibility, and takes a set of relations to be discriminating just so long as the objects that enter into them are not always indiscernible. But whether or not this is so will make no difference to our conclusions.

inating, so long as we hold classical objects to be impenetrable. Classical objects that are related by spatiotemporal relations can all be absolutely discernible. But external quantum relations do not appear to be discriminating; for quantum particles of the same kind will always be either weakly discernible or indiscernible. So there is an important difference between spatiotemporal relations and external quantum relations.

Second, French and Krause claim that, subject to our views about the nature of spacetime and individuality, spatiotemporal relations and quantum external relations differ in their dependence upon intrinsic properties. As they write:[5]

> If we were to adopt a relationist view of space-time, together with some form of 'bundle' theory of individuality, for example, then we might argue that spatio-temporal relations are still dependent on the intrinsic properties of the relevant objects, since if these were stripped away there would be no objects and without the latter there would be no spatio-temporal relations. However they are not *determined* by such properties and hence spatio-temporal relations can be described as only 'weakly' non-supervenient.

But for quantum particles the case is different:

> Since the properties represented by the superposition are not dependent upon those represented by 'both particles in the same state', then they are not determined by them either and thus they are 'strongly' non-supervenient. (French and Krause, 2003, p. 104)

The difference between weakly and strongly non-supervenient relations seems to be that weakly non-supervenient relations, like spatiotemporal relations, still depend upon intrinsic properties for their existence, even though they do not supervene upon intrinsic properties. Without the intrinsic properties of the objects, spatiotemporal relations would not exist. Strongly non-supervenient relations, like external quantum relations, neither depend upon

[5]At the end of this quotation they cite Cleland (1984).

intrinsic properties for their existence, nor supervene upon intrinsic properties.

But neither of these two differences between spatiotemporal relations and external quantum relations can ground a difference in vagueness. Consider once again discriminating relations. It is true that external quantum relations are not discriminating like spatiotemporal relations. But consider the context of Lewis's discussion of discriminating relations. Lewis is considering what properties a set of relations, like spatiotemporal relations, must have in order to provide a organizational framework for the phenomena in a possible world. He claims that, among other things, they must be discriminating (Lewis, 1986, p. 75–76). It is clear that external quantum relations provide no such framework, and therefore it should be no shock that they are not discriminating.

But this difference in discrimination should not lead to any difference in ontic vagueness. For discrimination does not require that the objects in spatiotemporal relations are never weakly discernible or indiscernible, only that they are not weakly discernible or indiscernible in a good number of possible worlds. Spatiotemporal relations are still discriminating even though there is a possible world containing Black's two spheres, where both spheres possess the same spatiotemporal relations. If we hold that weakly discernible quantum particles are ontically vague, then the question is why Black's two spheres are not also ontically vague. What we want to know is why weak discernibility in a given classical world does not lead to ontic vagueness, while weak discernibility in a given quantum world does. We cannot answer this simply by pointing out that, unlike quantum objects, there are classical objects that are absolutely discernible in other worlds. Discrimination is simply beside the point.

As for the different types of non-supervenience, it is not the case that spatiotemporal relations are dependent upon intrinsic properties in a way that external quantum relations are not. If spatiotemporal relations do not exist without their relata, then neither do external quantum relations. And if classical objects do not exist if we strip away their essential intrinsic properties, then neither do quantum particles. Thus external quantum relations are just as dependent upon the intrinsic properties of their relata as spatiotemporal relations. There is no distinction between 'weakly' and 'strongly' non-supervenient relations; and therefore nothing to ground any supposed difference in ontic vagueness between classical and quantum examples of weakly discernible objects.

There is no relevant difference between the classical and quantum examples of weakly discernible objects. If the identity of quantum particles is indeterminate because of the weak discernibility that follows from entanglement, then the identity of other weakly discernible objects, such as Black's two sphere, should also be indeterminate. And since this does not appear to be the case, we should doubt any claim that entangled quantum particles are ontically vague.

And once we draw the analogy to Black's two spheres, it is easy to claim that the source of any vagueness about the identity of entangled quantum particles is linguistic. For the reason that we cannot say which quantum particle is which is not because the quantum particles possess indeterminate identity, but because we cannot uniquely refer to a specific quantum particle by use of a proper name. As we have already discussed in chapter 2, under both the causal and descriptive theories of reference, proper names cannot uniquely refer to weakly discernible objects since such objects will possess all of the same properties and enter into all of the same relations. So if it is indeterminate whether particle a is identical to particle b, it is because "a" and "b" imprecisely refer. This goes not just for quantum particles, but all weakly discernible objects like Black's spheres.

While much of the emphasis in this discussion has been on weakly discernible objects, what we have said is also true of indiscernible objects, like non-interacting symmetrized bosons. Like weakly discernible objects, we cannot say which indiscernible object is which. And as in the case of weakly discernible objects, this inability to specify which is which does not appear to indicate that indiscernible objects are ontically vague, but rather that we cannot uniquely refer to them by use of a proper name.

Now French and Krause (1995, 2003) also think that there is another potential source of ontic vagueness in quantum theory. If we take particles not to be objects, but rather to be non-individuals that lack self-identity, then quantum particles present a clear example of ontic vagueness and indeterminate identity, for a lack of self-identity is as indeterminate as identity can get. And this ontic vagueness does not run afoul of Evans's derivation above; for as we have already mentioned, if we take ontically vague objects to lack self-identity, then we reject (7.3) in Evans's derivation (a point that French and Krause (2003, p. 110) make). So if indiscernibility, permutation symmetry, quantum statistics, or what have you leads us to claim that quantum particles are non-individuals, then quantum theory does lend support to ontic vagueness.

In agreement with French's views on metaphysical underdetermination, the two authors take both the individual and non-individual views of quantum particles to be equivalent alternatives. But we have already said quite a bit about why quantum particles are objects. Indiscernibility, permutation symmetry, and quantum statistics do not leads us to claim that quantum particles are non-individuals. Quantum particles do possess self-identity. So we cannot just cite the supposed lack of particle self-identity as an example of ontic vagueness.

Now it might be true that vagueness, even apparently linguistic vagueness, has its origin in the world. We really have not said anything against this view. But whether or not vagueness is ontic, quantum particles that are entangled do not provide a plausible example of synchronic ontic vagueness. Any synchronic vagueness associated with the identity of quantum particles has a clear linguistic explanation, an explanation that is more plausible than an ontic explanation in terms of external relations, weak discernibility, or non-individuality. There thus does not appear to be any clear case of synchronic ontic vagueness in quantum theory.

And if quantum particles are ontically vague, then so are other classical objects like Black's two spheres. So any ontic vagueness that might exist does not express any important difference between quantum and classical objects.

7.5 Diachronic Vagueness of Quantum Particles

Perhaps quantum theory provides better examples of diachronic ontic vagueness. Maybe the identity of quantum particles over time is indeterminate. This is after the sort of vagueness that Lowe is interested in.

There are plenty of classical examples of diachronic vagueness. One of the better known is Shoemaker's Brown/Brownson example. In this example we transplant Brown's brain in to Robinson's body, and call the person after the operation Brownson. It is vague whether the individual Brown before the operation is identical to the individual Brownson after the operation. That is there is vagueness about the cross-temporal identity of Brown and Brownson.

Any diachronic vagueness associated with quantum particles will be of a different sort than classical examples like Brown/Brownson. Because of the weak discernibility and indiscernibility of quantum particles of the same kind, we cannot introduce uniquely referring proper names like Brown and Brownson. This is why Lowe's example of quantum diachronic vagueness fails. But this does not mean that there is no vagueness associated with the cross-temporal identity of quantum particles.

Now the nature of any diachronic vagueness, in both quantum and classical examples, depends upon our theory of cross-temporal identity. There are numerous theories of cross-temporal identity. Many introduce unnecessary complications into our discussion. We will therefore only focus on two of the more popular ones: the standard endurantist and perdurantist views. It will be easy for those who are interested to extend what follows to other theories.

Let us start with the endurantist view that takes cross-temporal identity to be numerical identity. We might say that there is diachronic vagueness associated with which particle is which over time. If two quantum particles are of the same kind, they will be either weakly discernible or indiscernible. In both cases there will be no way to express which particle at one time is identical to a particle at another time, for all the properties and relations are the same for both particles at a given time.[6]

But even if we accept this inability to say which particle is which over time as a legitimate case of vagueness, it does not present us with a clear or plausible example of ontic vagueness under an endurantist theory of cross-temporal identity. For since, under this theory, cross-temporal identity is numerical identity, this diachronic vagueness is essentially the same as the synchronic vagueness we have just discussed. And we can offer the same linguistic explanation: Any vagueness simply follows from the fact that we cannot uniquely refer to weakly discernible or indiscernible objects at any time.

Now a perdurance theory of cross-temporal identity offers more hope for ontic vagueness, for here cross-temporal identity is not numerical identity. In the standard perdurance theory, persisting objects are mereological sums of temporal parts that are suitably related either by their qualitative similarity or causal relations. This relation is often called a genidentity relation.

Consider a system of two electrons that are weakly discernible. If we maintain that the two electrons are continuants, and the only continuants, then we run into a potential source of vagueness, for there are more than two mereological sums of temporal parts that we can identify with the continuants. To see this let us simplify the example and stipulate that there are only two times at which the continuants exist, t and t'. At each time there are two weakly discernible temporal parts. The mereological sums that we identify with the two continuants have the following constraints. Each mereological sum only consists of one temporal part at time t and one temporal part at t'. The temporal part at time t of one mereological sum must be weakly discernible from the temporal part at time t of the other mereological sum. The same must also be true for time t'. Because the particles are weakly discernible, there are no further qualitative or causal constraints for this example.

For this example, the problem is that there are two distinct sets of mere-

[6]Of course we can avoid any suggestion of diachronic vagueness simply by denying that quantum particles persist. In this case quantum particles only exist at a given time, and there is a succession of quantum particles over time. This view has the disadvantage of differing from our classical conception of particles as persisting objects.

ological sums that satisfy these constraints and not one. So even though we identify the continuants of this example with mereological sums, there are multiple mereological sums with which to identify each continuant. There are an infinite number for a continuous time spectrum and instantaneous temporal parts. And all of this is also true for indiscernible particles.

This problem is not confined to quantum examples, but also surfaces when we consider other weakly discernible objects. Consider once again Black's two spheres. Let us assume that there is no absolute spatial reference frame. Let us once again simplify the example and assume that there are only two distinct times. For further simplicity, let us assume that there is an absolute temporal reference frame (although we can easily do away with this assumption). Once again there are two distinct continuants, the two spheres. However, if each sphere is identical to a mereological sum of properly related temporal parts, then there are multiple mereological sums with which to identify each continuant.

This is true for Black's spheres for the same reason it is true for the two electrons. Again the only apparent constraints on the mereological sums with which we are to identify the continuants are that they only consist of one temporal part at each time, and that at each time the temporal part in one of the sums is weakly discernible from the other. Again this does not lead to a unique pair of mereological sums. It is, however, easier to picture the situation in the case of Black's two spheres. Take two plain note cards of the same size. Draw two circles of equal size and equal distance on each card. Each card is to represent a given time, and each circle on each card a temporal part of one of the spheres. The different mereological sums are represented by the different ways we can stack the cards so that the circles line up. Since we have assumed an absolute temporal reference frame, one card always has to be on the bottom and one on the top. But since we have assumed that there is no absolute spatial reference frame, there are two different ways we can stack the top card on the bottom, representing the two different pairs of mereological sums. So again we have two distinct sets of mereological sums, but only one set of persisting objects to which they are suppose to be equal.

This problem of weakly discernible and indiscernible continuants is related to Unger's problem of the many, which points out that for ordinary objects there are many distinct mereological sums that appear to have equal claim to being a given object (Unger, 1980). For example consider a single cat.[7] Which mereological sum is identical to that cat? There are several distinct mereological sums that appear to have equal claim to being the cat. For example one sum includes a loose hair that the others do not. Another includes some skin that is nearly shed. And so on. Each of these mereological sums presents us with a legitimate cat, but not one of these sums is clearly more of a cat than the others. And so the problem, once again, is that we have only one cat but several distinct cat-like mereological sums with which to identify the cat. In the same way, we have two persisting quantum particles, but more than two mereological sums with which to identify those particles. Both the cat and the quantum particles present us with a case of vagueness. In both cases it is vague which mereological sum is identical with a given object.

As with vagueness in general, there is no consensus on how to explain this sort of vagueness. Certainly we can claim that this vagueness is ontic, that it is genuinely indeterminate which mereological sum is identical to the cat or persisting particle. Parsons and Woodruff (1995) flesh out such an explanation. But it is safe to say that neither cats nor quantum particles provide examples that demand an ontic explanation; for we can offer a clear linguistic explanation of both examples of vagueness. We can follow Lewis (1999) and claim that this vagueness is a form of semantic indecision. When we talk about a cat, we are only talking about one cat-like mereological sum, we just do not specify which mereological sum. We can offer the same sort of explanation for the quantum example. When we are talking about a persisting quantum particle, we are only talking about one mereological sum of temporal parts, we just do not specify which sum. This semantic indecision follows our inability to uniquely refer to a specific mereological sum.

We are not here advocating this linguistic explanation of quantum vague-

[7]This sort of example originates with Geach (1980, p. 215–216).

ness. Our claim is much more modest. It is simply that this sort of quantum example does not advance the debate of ontic vagueness. For this sort of diachronic vagueness is equivalent to classical examples of vagueness. As with the vagueness in the classical examples, the nature of the quantum vagueness is open to debate. It is not obviously ontic, for there exists a plausible explanation that is linguistic. And if it does turn out to be ontic, then quantum continuants will be ontically vague in the same way as other classical mereological sums.

And this is really our conclusion on quantum vagueness in general. We can say that there is some vagueness that follows from our inability to say which quantum particle is which, either at a given time or over time. This quantum vagueness follows from the fact that quantum particles of the same kind are either weakly discernible or indiscernible. But this vagueness is not clearly ontic. We can offer a linguistic explanation of it. And this quantum vagueness is essentially not any different than the vagueness we find in classical examples. Ontic vagueness is therefore just as questionable in the quantum realm as it is in the classical realm, and the debate on vagueness is no better served by quantum examples than it is by classical examples.

And if further philosophical examination reveals that vagueness is ontic, this ontic vagueness will encompass both quantum and classical vagueness. It will therefore not lead to any important difference between the nature of quantum and classical objects, and therefore not undermine the individuality of quantum particles.

Chapter 8

Conclusion

So now we have come to the end. But before we pass onto the bibliography, let us review our thesis and conclusions. Our thesis is, once again, that quantum particles like classical particles are objects. That quality that we most associate with quantum particles—indistinguishability—in no way undermines the individuality of quantum particles.

As we have mentioned, there are two important ways that we can understand indistinguishability, both of which are related to several arguments about the identity of quantum particles. The first that we have discussed is to understand indistinguishability in terms of indiscernibility. The indiscernibility of quantum particles has consequences for the identity of quantum particles if we accept Leibniz's principle of the identity of indiscernibles (PII), which forbids that objects can be qualitatively indiscernible. PII is not a metaphysically necessary principle, but we have pointed out three main motivations for holding that PII applies to physical theories. There is an epistemological motivation, which holds that we should refrain from postulating physical theories with indiscernible objects since such objects do not add to the empirical content of a theory. There is a metaphysical motivation, which holds that PII provides the only suitable principle of individuation for physical objects. And there is a methodological motivation, which holds that PII provides us with a way to determine the ontology of a physical theory from evident physical properties.

Now we have shown that most quantum particles, including all of those of the actual world, are weakly discernible. If two particles of the same kind are antisymmetrized or interacting, then there are always irreflexive relations that weakly discern between the two. But non-interacting bosons can still be indiscernible. PII therefore still has implications for the ontology and interpretation of quantum theory; for it implies either that these indiscernible particles are not objects or that the quantum description of them is incomplete.

We have, however, pointed out that PII leads to some problematic interpretations of physical phenomena like Bose-Einstein condensates. And we have argued that the motivations supporting PII are too weak to sustain the weighty conclusions that follow from the principle. Indiscernible objects can contribute to the empirical content of a theory. We need not accept PII in order to provide a suitable principle of individuation; for we can take the identity of physical particles as primitive without any confusion or contradiction. And PII does not provide us with a way to determine the ontology of a theory from its evident physical properties; for many of those properties are not evident until we have specified an ontology. In short, there is no reason to shun indiscernible objects by affirming PII. And once we have abandoned PII, the identity of quantum particles is not threatened by their potential indiscernibility. Quantum particles are still objects even if indiscernible.

The second way to understand indistinguishability in terms of permutation symmetry. Once again, many hold the permutation symmetry of quantum particles follows from a lack of particle identity, and that this difference in permutation symmetry is an important way in which quantum physics differs from classical physics. However, when we understand indistinguishability in terms of permutation symmetry, there is no difference between classical and quantum particles; for classical physics is just as permutation symmetric as quantum physics. In classical physics, there are no empirical consequences as to which particles are traveling along which single-particle trajectories. The description of the dynamical evolution of the system need not specify which particle is in which single-particle state; for a permutation transformation does not affect the set of single-particle trajectories that describe the dynamical evolution of the system. The statistical mechanical description of the system also need not specify which particle is in which single-particle state; for statistical ensembles do not need to distinguish between permuted trajectories in order to arrive at the correct statistical mechanical and thermodynamical formulas. And because classical physics does not need to specify which particle is in which single-particle state, we can regard it as permutation symmetric and deny that a permutation of particles leads to a distinct physical situation.

A classical physics that is permutation symmetric is a more parsimonious theory than one that is not; for the permutation symmetric theory does not make any empirically vacuous distinctions as to which particle is in which single-particle state. We can therefore cite parsimony as reason to claim that classical physics is permutation symmetric. But there also is a further reason to regard it as so. As Gibbs first saw, permutation symmetry naturally leads to an extensive entropy function. As we have seen, there are several arguments that deny this. Only one, however, has any merit, the one voiced by Ehrenfest, Trkal, and van Kampen. Their argument holds that only systems of variable particle number must have an extensive entropy, and that the extensivity of the entropy for such systems naturally follows from the relevant derivations without any recourse to permutation symmetry. At the base of their argument is the Clausius approach to entropy and the second law of thermodynamics. As we have argued above, while we must assent to the validity of this argument, we can dispute its main premise. If, instead of following the Clausius approach, we follow a more robust axiomatic approach to entropy and the second law, such as that formulated by Lieb and Yngvason, then the entropy must be extensive for all systems, even those of fixed particle number. We therefore must once again turn to permutation symmetry to explain this extensivity.

So we have two very good reasons to hold that classical physics is permutation symmetric. Classical particles are therefore as indistinguishable, in this permutation symmetric sense, as quantum particles. Thus this sort of indistinguishability cannot ground any distinction between classical and quantum physics, such as a difference in ontology or a difference in statistics.

And we have shown in detail how indistinguishability, understood either in terms of indiscernibility or permutation symmetry, does not entail quantum statistics. Quantum statistics follows from the symmetrization requirement, and neither notion of indistinguishability implies the symmetrization requirement. Symmetrization, and with it quantum statistics, is best viewed as simply an empirical hypothesis of quantum theory instead of as a requirement that follows the indistinguishable nature of quantum particles or a lack of particle identity.

We have, however, noted that permutation symmetry in the quantum formalism rules out classical statistics even though it only provides a necessary condition for quantum statistics. Combinatorial derivations seem to suggest that the same is true in the classical formalism. If this were so, then a difference in permutation symmetry would be responsible for a difference in statistics in both theories, and we could not maintain that classical physics was permutation symmetric. But we have shown how this is not so. The combinatorial derivations only lead to a difference in statistics when we divide phase space into cells of definite volume. This is essentially a quantum correction that allows for a direct correspondence with quantum Hilbert space states. Because of the direct correspondence, a difference in permutation symmetry leads to a difference in statistics in phase space, just as it does in Hilbert space. But the division of phase into cells of definite volume does not have any justification within classical physics. When we remove this stipulation, a difference in permutation symmetry no longer leads to a difference in statistics (with the exception of the entropy) in the phase space formalism. Permutation symmetry does lead to a difference in statistics when coupled with the Hilbert space formalism. But permutation symmetry in the phase space formalism only affects a change in the entropy. The permutation symmetry of classical physics is consistent with classical statistics.

Since classical physics is as permutation symmetric as quantum physics, we cannot understand permutation symmetry as following from a lack of particle identity. But this does not mean that permutation symmetry is without any sort of philosophical explanation. As we have pointed out, the permutation symmetry of both classical and quantum physics is connected to the metaphysical thesis of anti-haecceitism. Permutation symmetry eliminates haecceitistic distinctions, distinctions that are simply artifacts of the formalism that fail to represent anything physical. So permutation symmetry ensures that physics is a qualitative theory, understood in this anti-haecceitistic way.

We have discussed some common objections to this view. Authors like Teller and Gordon claim that haecceitistic distinctions are implied by objecthood, at least in the case of quantum particles. The idea is that the individuation and labeling of objects like quantum particles imply haecceitistic differences. And if we regard quantum particles as objects, these haecceitistic distinctions lead to excess states that are not accounted for in quantum statistical ensembles. But as we have argued, objecthood does not imply haecceitistic distinctions for quantum particles. In order to consider quantum particles as objects, we do not require principles of individuation that lead to haecceitistic distinctions; for instead of introducing haecceities and the like, we can simply take the identity of quantum particles as primitive, which does not imply any haecceitistic distinctions. We further do not need to introduce haecceitistic distinctions in order to ensure that quantum particles can be labeled, for we can regard quantum particles as objects without being able to refer to them by use of rigidly referring proper names. Quantum particles, like classical particles, can be objects and anti-haecceitistic.

Because of this connection between permutation symmetry and anti-haecceitism, permutation symmetry provides a point of continuity between classical and quantum physics. Instead of following from a difference in ontology, permutation symmetry follows from the fact that both classical and quantum physics are only concerned with the qualitative nature of the world. Permutation symmetry thus expresses an important feature that both theories share.

We have also pointed out that there is an analogy between many of these issues surrounding quantum particles and those surrounding spacetime points in spacetime physics. As in quantum mechanics, haecceitistic distinctions lead to excess possibilities in spacetime physics. In quantum mechanics these excess possibilities leads to problems for quantum statistics; in spacetime physics they lead to problems for determinism. Some think that the way to solve both these problems is to deny that both quantum particles and spacetime points are objects. We have rejected that this is the case with respect to quantum particles. We can deny haecceitistic distinctions without undermining the identity of quantum particles. This analogy allows us to extend our conclusion to spacetime points. We can accept that spacetime points are objects without admitting any haecceitistic distinctions that lead to indeterminism.

Now our discussion of anti-haecceitism not only provides us with a philosophical explanation of permutation symmetry, it also provides us with an example of how physics and metaphysics interact. The lack of haecceitistic distinctions in physics lends support to the metaphysical thesis of anti-haecceitism. And this metaphysical thesis provides clarification of the mathematical formalism and the ontology of physics, both quantum physics and spacetime physics. We here have an example of mutual interaction between physics and metaphysics.

So hopefully we have established that the indistinguishability of quantum particles does not undermine their individuality. But as we have mentioned, not all questions about particle individuality center on indistinguishability. We have also discussed the relation between particles and fields. We have noted that quantum particles do not provide the fundamental ontology of quantum theory; for they are incapable of providing an adequate ontology for relativistic quantum field theory. Such things as indeterminate particle number states and vacuum phenomena resist a particle interpretation. Fields, however, appear to provide for a better interpretation of these phenomena and relativistic quantum field theory in general. Further we have pointed out that any state we can interpret in terms of particles, we can also interpret in terms of fields. Fields, not particles, therefore seem to be the fundamental ontology of quantum theory.

But as we have argued, this does not mean that particles are merely properties of fields lacking their own individuality or that they have no place in quantum theory at all. They still serve as an effective or emergent ontology. In certain physical regimes, particles provide a more useful, succinct, and clearer description of phenomena than fields. Atomic phenomena provide a clear example of this. We take atoms to be particles and not discrete excitations of some atomic field. Insofar as particles provide a useful description of phenomena in a certain regime, we can consider them as real objects, even though they ultimately supervene upon more fundamental fields.

And one of the advantages of a particle ontology over a field ontology is that it allows us to consider macroscopic objects as mereological sums of particles. We describe our macroscopic world in terms of objects that are not fields. We talk of tables and chairs, not excitations of table fields and chair fields. Since we want the macroscopic world to emerge from the quantum world, at some point there must be a shift from a field ontology to another ontology. Where that occurs will depend greatly upon our how we explain localization and on our solution to the measurement problem. Quantum particles, however, provide a natural point for that shift in ontology.

When we consider particles instead of fields, where we take particles as objects, we can take macroscopic objects as mereological sums of particles. This is what we commonly do, even without a solution to the measurement problem. We consider tables and chairs as a collection of molecules and we consider molecules as a collection of atoms. Granted, if we are physicalists then tables, chairs, molecules, and atoms better supervene upon fundamental fields, but we still consider all of them as emergent objects that are not fields. And so here we have an example of how particles, when taken as objects, provide an important bridge between the classical macroscopic world and the quantum world.

We have also discussed other influential arguments that purport to show that quantum particles lack identity. There is Redhead and Teller's surplus structure argument. This argument holds that if we take quantum particles to be objects, we have surplus unsymmetrized states that never occur. We can excise this surplus structure if we deny that quantum particles are objects. We have pointed out that this argument is essentially no different than the ontological argument from permutation symmetry, and that it essentially makes the same mistakes. Individuality does not lead to haecceitistic differences, and permutation symmetry does not only apply to particles that lack identity.

Another argument cites the statistical correlations between symmetrized particles. If we take particles to be objects, then there is no explanation for these correlations in terms of interactions or common cause. If, however, we take fields as our ontology, then these correlations simply reflect different measurements on the same state of the field. But we have argued that these correlations do not force us to abandon a particle ontology in favor of a field ontology; for similar unexplained correlations pervade quantum theory, regardless of whether or not we adopt a field ontology. Really the only question of importance as to whether or not we should adopt a particle ontology instead of a field ontology is whether a particle ontology provides a more useful description than a field ontology.

We have also examined a further way in which quantum particles are suppose to differ from classical particles. As we have seen, some claim that quantum particles, unlike classical particles, are ontically vague, and that this ontic vagueness might be due to a lack of identity. While we have admitted that there is a legitimate debate as to whether vagueness is ontic, we have denied that quantum particles are obviously ontically vague. Any vagueness they might enter is no different than the vagueness associated with classical objects, and can be explained in the same way.

So if we have been successful in this thesis, then we have shown that quantum particles are objects in much the same way as classical particles. And it is not just that we can take quantum particles as objects. There is no underdetermination between individual and non-individual particles ontologies. Upon philosophical reflection, it is clear that they are objects like classical particles. No doubt there are important differences between the two. It is not clear that quantum particles possess spatiotemporal trajectories, and systems of quantum possess a minimal form of holism. But quantum particles are objects, and objects in a very conventional sense. They are not ontically vague. They do not belong to some new category of thing. They do not require an alternative logic or set theory. Quantum particles, like classical particles, still possess identity. A quantum particle is identical to itself and numerically distinct from other particles. We can predicate properties to it and quantify over it. We can describe it by use of standard first-order quantificational logic.

We therefore have a very important point of continuity between classical and quantum theory. For there are particles in both theories. And while they possess very different properties in both theories (at least in some interpretations of quantum theory), their basic ontological structure is the same. And the conceptual apparatus we apply to them—logic, set theory, our reasoning about possibility—are the same for both. Any adequate interpretation of quantum theory must respect this important point of continuity. The quantum world is not as different from the classical world as many think.

Bibliography

Adams, R. M. (1979). Primitive thisness and primitive identity. *The Journal of Philosophy*, 76(1):5–26.

Albert, D. Z. (2000). *Time and Chance*. Harvard University Press, Cambridge, Massachusetts.

Armstrong, D. M. (1978). *Nominalism and Realism*, volume 1 of *Universals and Scientific Realism*. Cambridge University Press, Cambridge, England.

Ayer, A. J. (1954). The identity of indiscernibles. In *Philosophical Essays*, chapter 2, pages 26–35. Macmillan & Co. Ltd., London.

Bacciagaluppi, G. (2003). Derivation of the symmetry postulates for identical particles from pilot-wave theories. PITT-PHIL-SCI 993.

Bach, A. (1985). On the quantum properties of indistinguishable classical particles. *Lettere al Nuovo Cimento*, 43(8):383–387.

Bach, A. (1990). Boltzmann's probability distribution of 1877. *Archive for History of Exact Sciences*, 41(1):1–40.

Belot, G. (1995). New work for counterpart theorists: Determinism. *The British Journal for the Philosophy of Science*, 46(2):185–195.

Black, M. (1952). The identity of indiscernibles. *Mind*, 61(242):153–164.

Blokhinstev, D. I. (1964). *Quantum Mechanics*. D. Reidel Publishing Company, Dordrecht.

Boltzmann, L. (1877). Über die Beziehung zwischen dem zweiten Hauptsatz der mechanischen Wärmetheorie und der Wahrscheinlichkeitsrechnung respective den Sätzen über das Wärmegleichgewicht. *Wiener Berichte*, 76:373–435.

Born, M. (1943). *Experiment and Theory in Physics*. Cambridge University Press, London.

Bose, S. N. (1924). Planck's Gesetz und Lichtquantenhypothese. *Zeitschrift für Physik*, 26:178–181.

Brading, K. and Brown, H. R. (2004). Are gauge symmetry transformations observable? *The British Journal for the Philosophy of Science*, 55:645–665.

Brading, K. and Castellani, E., editors (2003). *Symmetries in Physics: Philosophical Reflections*. Cambridge University Press, Cambridge, England.

Brighouse, C. (1994). Spacetime and holes. In *Proceedings of the Biennial Meeting of the Philosophy of Science Association*, volume 1, pages 117–125.

Brighouse, C. (1997). Determinism and modality. *The British Journal for the Philosophy of Science*, 48(4):465–481.

Brown, H. R., Sjöqvist, E., and Bacciagaluppi, G. (1999). Remarks on identical particles in de Broglie-Bohm theory. *Physics Letters A*, 251:229–235.

Butterfield, J. (1989). The hole truth. *The British Journal for the Philosophy of Science*, 40(1):1–28.

Butterfield, J. (1993). Interpretation and identity in quantum theory. *Studies In History and Philosophy of Science*, 24(3):443–476.

Callender, C. (1999). Reducing thermodynamics to statistical mechanics: The case of entropy. *The Journal of Philosophy*, 96:348–373.

Cao, T. Y. (2003). Can we dissolve physical entities into mathematical structures? *Synthese*, 136:57–71.

Carnap, R. (1950). *Logical Foundations of Probability.* Routledge and Kegan Paul, Ltd., London.

Castañeda, H.-N. (1975). Individuation and non-identity: A new look. *American Philosophical Quarterly*, 12(2):131–140.

Castellani, E., editor (1998). *Interpreting Bodies: Classical and Quantum Objects in Modern Physics.* Princeton University Press, Princeton.

Cheung, L. K. C. Quantum statistics, tractarian objects and the existential quantifier.

Cleland, C. (1984). Space: An abstract system of non-supervenient relations. *Philosophical Studies*, 57:248–265.

Copeland, B. J. (1997). Vague identity and fuzzy logic. *The Journal of Philosophy*, 94(10):514–534.

Cornell, E. A. and Wieman, C. E. (1998). The Bose-Einstein condensate. *Scientific American*, 278(3):26–31.

Costantini, D. (1987). Symmetry and the indistinguishability of classical particles. *Physics Letters A*, 123(9):433–436.

Costantini, D., Galavotti, M. C., and Rosa, R. (1982). A rational reconstruction of elementary particles statistics. *Scientia*, 117(1–4):151–159.

Costantini, D. and Garibaldi, U. (1989). Classical and quantum statistics as finite random processes. *Foundations of Physics*, 19(6):743–754.

Dalfovo, F., Giorgini, S., Pitaevskii, L. P., and Stringari, S. (1999). Theory of Bose-Einstein condensation in trapped gases. *Review of Modern Physics*, 71(3):463–512.

Dalla Chiara, M. L., Giuntini, R., and Krause, D. (1998). Quasiset theories for microobjects: A comparison. In Castellani (1998), chapter 8, pages 142–152.

de Muynck, W. M. (1975). Distinguishable- and indistinguishable-particle descriptions of systems of identical paticles. *International Journal of Theoretical Physics*, 14(5):327–346.

de Muynck, W. M. and van Liempd, G. P. (1986). On the relation between indistinguishability of identical particles and (anti)symmetry of the wave function in quantum mechanics. *Synthese*, 67:477–496.

Dennett, D. C. (1991). Real patterns. *The Journal of Philosophy*, 88(1):27–51.

Dieks, D. (1990). Quantum statistics, identical particles and correlations. *Synthese*, 82:127–155.

Dummett, M. (1973). *Frege: Philosophy of Language*. Duckworth, London.

Earman, J. and Norton, J. (1987). What price spacetime substantivalism? the hole story. *The British Journal for the Philosophy of Science*, 38(4):515–525.

Ehrenfest, P. (1911). Welche Züge der Lichtquantenhypothese spielen in der Theorie der Wärmestrahlung eine wesentliche Rolle. *Annalen der Physik*, 36:91–118. Reprinted in Ehrenfest (1959).

Ehrenfest, P. (1959). *Collected Scientific Papers*. North-Holland Publishing Company, Amsterdam.

Ehrenfest, P. and Ehrenfest, T. (1959). *The Conceptual Foundations of the Statistical Approach in Mechanics*. Cornell University Press, Ithaca. Translated by Michael J. Moravcsik.

Ehrenfest, P. and Onnes, H. K. (1914). Simplified deduction of the formula from the theory of combinations which Planck uses as the basis of his radiation theory. *Proceedings of the Amsterdam Academy*, 17:870–873. Reprinted in Ehrenfest (1959).

Ehrenfest, P. and Trkal, V. (1920). Deduction of the dissociation-equilibrium from the theory of quanta and a calculation of the chemical constant based on this. *Proceedings of the Amsterdam Academy*, 23:161–183. Reprinted in Ehrenfest (1959).

Einstein, A. (1905). Über einen die Erzeugung und Verwandlung des Lichtes betreffenden heuristischen Gesichtspunkt. *Annalen der Physik*, 17:132–148.

Einstein, A. (1924). Quantentheorie des einatomigen idealen Gases. *Berliner Berichte*, pages 261–267.

Einstein, A. (1925). Quantentheorie des einatomigen idealen Gases. *Berliner Berichte*, pages 3–14.

Evans, G. (1978). Can there be vague objects? *Analysis*, 38:208.

Fermi, E. (1936). *Thermodynamics*. Dover Publications, New York.

Feynman, R. P. (1972). *Statistical Mechanics: A Set of Lectures*. W. A. Benjamin, Inc., Reading, Massachusetts.

French, S. (1989a). Identity and individuality in classical and quantum physics. *Australasian Journal of Philosophy*, 67(4):432–446.

French, S. (1989b). Why the principle of the identity of indiscernibles is not contingently true either. *Synthese*, 78:141–166.

French, S. (1998). On the withering away of physical objects. In Castellani (1998), chapter 6, pages 93–113.

French, S. and Krause, D. (1995). Vague identity and quantum non-individuality. *Analysis*, 55(1):20–26.

French, S. and Krause, D. (2003). Quantum vagueness. *Erkenntnis*, 59:97–124.

French, S. and Ladyman, J. (2003a). The dissolution of objects: Between Platonism and phenomenalism. *Synthese*, 136:73–77.

French, S. and Ladyman, J. (2003b). Remodelling structural realism: Quantum physics and the metaphysics of structure. *Synthese*, 136:21–56.

French, S. and Redhead, M. (1988). Quantum physics and the identity of indiscernibles. *British Journal for the Philosophy of Science*, 39:233–246.

French, S. and Rickles, D. (2003). Understanding permutation symmetry. In Brading and Castellani (2003), chapter 12, pages 212–238.

Fujita, S. (1991). On the indistinguishability of classical particles. *Foundations of Physics*, 21(4):439–457.

Geach, P. T. (1980). *Reference and Generality*. Cornell University Press, Ithaca, third edition.

Gibbs, J. W. (1902). *Elementary Principles in Statistical Mechanics: Developed with Especial Reference to the Rational Foundation of Thermodynamics*. Charles Scribner's Sons, New York.

Gibbs, J. W. (1948). On the equilibrium of heterogeneous substances. In *The Collected Works of J. Willard Gibbs*, volume 1, chapter 3, pages 55–353. Yale University Press, New Haven.

Gordon, B. L. (2002). Maxwell-Boltzmann statistics and the metaphysics of modality. *Synthese*, 133:393–417.

Halvorson, H. and Clifton, R. (2002). No place for particles in relativistic quantum theories? *Philosophy of Science*, 69:1–28.

Hartle, J. B. and Taylor, J. R. (1969). Quantum mechanics of quantum particles. *Physical Review*, 178(5):2043–2051.

Hawley, K. (1998). Indeterminism and indeterminacy. *Analysis*, 58(2):101–106.

Hestenes, D. (1970). Entropy and indistinguishability. *American Journal of Physics*, 38(7):840–845.

Hilbert, D. and Bernays, P. (1934). *Grundlagen der Mathematik*, volume 1. Springer-Verlag, Berlin.

Hoefer, C. (1996). The metaphysics of space-time substantivalism. *The Journal of Philosophy*, 93(1):5–27.

Huang, K. (1963). *Statistical Mechanics*. John Wiley & Sons, New York.

Huggett, N. (1994). What are quanta, and why does it matter? In *Proceedings of the Biennial Meeting of the Philosophy of Science Association*, volume 2, pages 69–76.

Huggett, N. (1999a). Atomic metaphysics. *The Journal of Philosophy*, 96(1):5–24.

Huggett, N. (1999b). On the significance of permutation symmetry. *The British Journal for the Philosophy of Science*, 50:325–347.

Imbo, T. D., Shah Imbo, C., and Sudarshan, E. C. G. (1990). Identical particles, exotic statistics and braid groups. *Physics Letters B*, 234(1,2):103–107.

Ismael, J. and van Fraassen, B. C. (2003). Symmetry as a guide to superfluous theoretical structure. In Brading and Castellani (2003), chapter 23, pages 371–392.

Jackson, F. (1998). *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Clarendon Press, Oxford.

Jammer, M. (1966). *The Conceptual Development of Quantum Mechanics*. McGraw-Hill Book Company, New York.

Jauch, J. M. (1968). *Foundations of Quantum Mechanics*. Addison-Wesley Publishing Company, Reading, Massachusetts.

Kant, I. (1965). *Critique of Pure Reason*. St. Martin's Press, New York. Translated by Norman Kemp Smith.

Kaplan, D. (1975). How to Russell a Frege-Church. *The Journal of Philosophy*, 27(19):716–729.

Kaplan, I. G. (1976). The exclusion principle and indistinguishability of identical particles in quantum mechanics. *Soviet Physics Uspekhi*, 18(12):988–994.

Ketterle, W. (2002). Nobel lecture: When atoms behave as waves: Bose-Eistein condensation and the atom laser. *Reviews of Modern Physics*, 74:1131–1151.

Kittel, C. and Kroemer, H. (1980). *Thermal Physics*. W. H. Freeman and Company, San Francisco, second edition.

Kosso, P. (2000). The empirical status of symmetries in physics. *The British Journal for the Philosophy of Science*, 51:81–98.

Kripke, S. A. (1980). *Naming and Necessity*. Blackwell, Oxford.

Kuhn, T. S. (1978). *Black-Body Theory and the Quantum Discontinuity: 1894–1912*. Clarendon Press, Oxford.

Ladyman, J. (1998). What is structural realism? *Studies in History and Philosophy of Science*, 29(3):409–424.

Laidlaw, M. G. G. and DeWitt, C. M. (1971). Feynman functional integrals for systems of indistinguishable particles. *Physical Review D*, 3(6):1375–1378.

Landau, L. D. and Lifshitz, E. M. (1958). *Quantum Mechanics: Non-Relativistic Theory*. Pergamon Press, London. Translated by J. B. Sykes and J. S. Bell.

Leibniz, G. W. and Clarke, S. (1956). *The Leibniz-Clarke Correspondence*. Manchester University Press, Manchester.

Leinaas, J. M. and Myrheim, J. (1977). On the theory of identical particles. *Il Nuovo Cimento B*, 37(1):1–23.

Lewis, D. (1973). *Counterfactuals*. Basil Blackwell, Oxford.

Lewis, D. (1983). New work for a theory of universals. *Australasian Journal of Philosophy*, 61(4):343–377.

Lewis, D. (1986). *On the Plurality of Worlds*. Blackwell, Oxford.

Lewis, D. (1996). Maudlin and modal mystery. *Australasian Journal of Philosophy*, 74(4):683–684.

Lewis, D. (1999). Many, but almost one. In *Papers in Metaphysics and Epistemology*, chapter 9, pages 164–182. Cambridge University Press, Cambridge, England.

Lewis, G. N. (1930). The principle of identity and the exclusion of quantum states. *Physical Review*, 36:1144–1153.

Lieb, E. H. and Yngvason, J. (1999). The physics and mathematics of the second law of thermodynamics. *Physics Reports*, 310:1–96. Erratum published in 314:669, 1999.

Liu, C. (1996). Realism and spacetime: Of arguments against metaphysical realism and manifold realism. *Philosophia Naturalis*, 33(2):243–263.

Lorentz, H. A. (1910). Alte und neue Fragen der Physik. *Physikalische Zeitschrift*, 11:1234–1257.

Lowe, E. J. (1994). Vague identity and quantum indeterminacy. *Analysis*, 54(2):110–114.

Lowe, E. J. (1997). Reply to Noonan on vague identity. *57*, (1):88–91.

Lowe, E. J. (1998). *The Possibility of Metaphysics: Substance, Identity, and Time*. Clarendon Press, Oxford.

Lowe, E. J. (1999). Vague identity and quantum indeterminacy: Further reflections. *Analysis*, 59(4):328–330.

Lowe, E. J. (2001). Ontic indeterminacy of identity unscathed. *Analysis*, 63(3):241–45.

Mandl, F. (1971). *Statistical Physics*. John Wiley & Sons, London.

Margenau, H. (1944). The exclusion principle and its philosophical importance. *Philosophy of Science*, 11(4):187–208.

Maudlin, T. (1990). Substances and space-time: What Aristotle would have said to Einstein. *Studies in History and Philosophy of Science*, 21(4):531–561.

Maudlin, T. (1996). On the impossibility of David Lewis' modal realism. *Australasian Journal of Philosophy*, 74(4):669–682.

Melia, J. (1999). Holes, haecceitism and two conceptions of determinism. *The British Journal for the Philosophy of Science*, 50:639–664.

Merricks, T. (2003). The end of counterpart theory. *The Journal of Philosophy*, 100(10):521–549.

Messiah, A. M. L. (1962). *Quantum Mechanics*, volume 2. North-Holland Publishing Company, Amsterdam. Translated by J. Potter.

Messiah, A. M. L. and Greenberg, O. W. (1964). Symmetrization postulate and its experimental foundation. *Physical Review*, 136(1B):248–267.

Mosini, V. (1995). Fundamentalism, antifundamentalism, and Gibbs' paradox. *Studies in History and Philosophy of Modern Physics*, 26(2):151–162.

Natanson, L. (1911). On the statistical theory of radiation. *Belletin de l'Académie des Sciences de Cracovie A*, pages 134–148.

Noonan, H. W. (1995). E. J. Lowe on vague identity and quantum indeterminacy. *Analysis*, 55(1):14–19.

Odrowąż-Sypniewska, J. (2001). Quantum indiscernibility without vague identity. *Analysis*, 61(1):65–69.

Pais, A. (1982). *'Subtle is the Lord ...': The Science and the Life of Albert Einstein*. Clarendon Press, Oxford.

Parsons, T. and Woodruff, P. (1995). Worldly indeterminacy of identity. *Proceedings of the Aristotelian Society*, 95:171–191.

Planck, M. (1900). Zur Theorie des Gesetzes der Energieverteilung im Normalspektrum. *Verhandlungen der Deutschen Physikalischen Gesellschaft*, 2:237–245. An English translation appears in ter Haar (1967).

Planck, M. (1913). *Vorlesungen über die Theorie der Wärmestrahlung*. Barth, Leipzig, second edition.

Planck, M. (1927). *Treatise on Thermodynamics*. Longmans, Green and Co., London, 3rd edition. Translated from the 7th German edition by Alexander Ogg.

Plank, M. (1967). The genesis and present state of development of the quantum theory: Nobel lecture. In *Nobel Lectures: Physics 1901–1921*. Elsevier Publishing Company, Amsterdam.

Pooley, O. Points, particles, and structural realism. In French, S., Rickles, D., and Saatsi, J., editors, *Structural Foundations of Quantum Gravity*. Oxford University Press. Forthcoming.

Post, H. (2000). The impersonal atom. *Physis*, 37(1):107–136.

Quine, W. V. O. (1960). *Word and Object*. MIT Press, Cambridge, Massachusetts.

Quine, W. V. O. (1976). Grades of discriminability. *The Journal of Philosophy*, 73(5):113–116.

Redhead, M. L. G. (1982). Quantum field theory for philosophers. In *Proceedings of the Biennial Meeting of the Philosophy of Science Association*, volume 2, pages 57–99.

Redhead, M. L. G. (1988). A philosopher looks at quantum field theory. In Brown, H. R. and Harré, R., editors, *Philosophical Foundations of Quantum Field Theory*, chapter 1, pages 9–23. Clarendon Press, Oxford.

Redhead, M. L. G. (1994). The vacuum in relativistic quantum field theory. In *Proceedings of the Biennial Meeting of the Philosophy of Science Association*, volume 2, pages 77–87.

Redhead, M. L. G. and Teller, P. (1991). Particles, particle labels and quanta: The toll of unacknowledged metaphysics. *Foundations of Physics*, 21(1):43–62.

Redhead, M. L. G. and Teller, P. (1992). Particle labels and the theory of indistinguishable particles in quantum mechanics. *The British Journal for the Philosophy of Science*, 43:201–218.

Reichenbach, H. (1998). The genidentity of quantum particles. In Castellani (1998), chapter 4, pages 61–72. Excerpted from *The Direction of Time*.

Robertson, B. (1973). Introduction to field operators in quantum mechanics. *American Journal of Physics*, 41:678–690.

Rosenkrantz, G. S. (1979). The pure and the impure. *Logique et Analyse*, 22(88):515–523.

Russell, B. (1956). On the relations of universals and particulars. In Marsh, R. C., editor, *Logic and Knowledge*, pages 105–124. New York.

Rynasiewicz, R. (1994). The lessons of the hole argument. *The British Journal for the Philosophy of Science*, 45(2):407–436.

Salmon, N. (1982). *Reference and Essence*. Basil Blackwell, Oxford.

Sarry, M. F. (1979). Permutation symmetry of wave functions of a system of identical particles. *Soviet Physics JETP*, 50(4):678–680.

Saunders, S. Are quantum particles individuals? Forthcoming a.

Saunders, S. On the explanation for quantum statistics. Forthcoming b.

Saunders, S. (1994). A dissolution of the problem of locality. In *Proceedings of the Biennial Meeting of the Philosophy of Science Association*, volume 2, pages 88–98.

Saunders, S. (2002). Is the zero-point energy real? In Kuhlmann, M., Lyre, H., and Wayne, A., editors, *Ontological Aspects of Quantum Field Theory*, chapter 16, pages 313–343. World Scientific, New Jersey.

Saunders, S. (2003). Physics and Leibniz's principles. In Brading and Castellani (2003), chapter 16, pages 289–307.

Schiff, L. I. (1968). *Quantum Mechanics*. McGraw-Hill Book Company, New York, third edition.

Schrödinger, E. (1948). *Statistical Thermodynamics*. Cambridge University Press, Cambridge, England.

Schrödinger, E. (1998). What is an elementary particle? In Castellani (1998), pages 197–210.

Sciama, D. W. (1991). The physical significance of the vacuum state of a quantum field. In Saunders, S. and Brown, H. R., editors, *The Philosophy of Vacuum*, chapter 6, pages 137–158. Clarendon Press, Oxford.

Stachel, J. (2002). "The Relations Between Things" versus "The Things Between Relations": The deeper meaning of the hole argument. In Malament, D. B., editor, *Reading Natural Philosophy: Essays in the History and Philosophy of Science and Mathematics*, chapter 9, pages 231–266. Open Court, Chicago.

Strawson, P. F. (1959). *Individuals*. Routledge, London.

Swinburne, R. (1995). Thisness. *Australasian Journal of Philosophy*, 73(3):389–400.

Teller, P. (1983). Quantum physics, the identity of indiscernibles, and some unanswered questions. *Philosophy of Science*, 50:309–319.

Teller, P. (1995). *An Interpretive Introduction to Quantum Field Theory.* Princeton University Press, Princeton.

Teller, P. (1998). Quantum mechanics and haecceities. In Castellani (1998), chapter 7, pages 114–141.

Teller, P. (2001). The ins and outs of counterfactual switching. *Noûs*, 35(3):365–393.

Teller, P. and Redhead, M. (2000). Is indistinguishability in quantum mechanics conventional? *Foundations of Physics*, 30(6):951–957.

ter Haar, D. (1967). *The Old Quantum Theory.* Pergamon Press, Oxford.

ter Haar, D. (1995). *Elements of Statistical Mechanics.* Butterworth-Heinemann, Oxford, third edition.

Tiesinga, E., Moerdijk, A. J., Verhaar, B. J., and Stoof, H. T. C. (1995). BEC in ultra-cold Cesium: Collisional constraints. In Griffin, A., Snoke, D. W., and Stringari, S., editors, *Bose-Einstein Condensation*, pages 465–471. Cambridge University Press, Cambridge, England.

Tolman, R. C. (1938). *The Principles of Statistical Mechanics.* Oxford University Press, London.

Tsallis, C. (2004). Nonextensive statistical mechanics: Construction and physical interpretation. In Gell-Mann, M. and Tsallis, C., editors, *Nonextensive Entropy—Interdisciplinary Applications*, Santa Fe Institute Studies in Sciences of Complexity, pages 1–53. Oxford University Press, Oxford.

Uffink, J. (2001). Bluff your way in the second law of thermodynamics. *Studies in History and Philosophy of Modern Physics*, 32(3):305–394. Preprint: cond-mat/0005327.

Unger, P. (1980). The problem of the many. *Midwest Studies in Philosophy*, 5:411–467.

van Fraassen, B. C. (1991). *Quantum Mechanics: An Empiricist View.* Clarendon Press, Oxford.

van Fraassen, B. C. (1998). The problem of indistinguishable particles. In Castellani (1998), chapter 5, pages 73–90.

van Inwagen, P. (1990). *Material Beings.* Cornell University Press, Ithaca.

van Kampen, N. G. (1984). The Gibbs paradox. In Parry, W. E., editor, *Essays in Theoretical Physics in Honour of Dirk ter Haar*, pages 303–312. Pergamon Press Ltd., Oxford.

Weyl, H. (1949). *Philosophy of Mathematics and Natural Science.* Princeton University Press, Princeton.

Weyl, H. (1950). *The Theory of Groups and Quantum Mechanics.* Dover, New York. Translated by H. P. Robertson.

Wiggins, D. (2001). *Sameness and Substance Renewed.* Cambridge University Press, Cambridge, England.

Williamson, T. (1994). *Vagueness.* Routledge, London.

Wilson, M. (1993). There's a hole and a bucket, dear Leibniz. *Midwest Studies in Philosophy*, 18:202–41.

Wittgenstein, L. (1961). *Tractatus Logico-Philosophicus.* Routledge & Kegan Paul, London. Translated by D. F. Pears and B. F. McGuinness.