

How Symmetry Undid the Particle

A Demonstration of the Incompatibility of Particle
Interpretations and Permutation Invariance

Benjamin C. Jantzen

Carnegie Mellon University
2010

A dissertation submitted in partial fulfillment of the
requirements for the degree of Doctor of Philosophy in
Logic, Computation and Methodology

Committee

Mara Harrell (chair)

John Earman

Jeremy Butterfield

Acknowledgements

Quite a few people have contributed to the work presented here or have facilitated its production in important ways, and I am grateful for their assistance. I would particularly like to thank John Earman, Jeremy Butterfield, Bryan Roberts, Peter Spirtes, and Kristina Jantzen for valuable discussion and criticism. I am especially indebted to Clark Glymour for helping me to clarify the logical structure of my central argument, and to two anonymous referees for *Philosophy of Science* for their insightful critique. I am immensely grateful for the support provided by Mara Harrell. Her guidance was essential to the development of the ideas presented here. Finally, I would like to thank Parisa Farhi for somehow clearing away the impediments in our daily lives to give me the time I needed to complete this dissertation.

Abstract

The idea that the world is made of particles—little discrete, interacting objects that compose the material bodies of everyday experience—is a durable one. Following the advent of quantum theory, the idea was revised but not abandoned. It remains manifest in the explanatory language of physics, chemistry, and molecular biology. Aside from its durability, there is good reason for the scientific realist to embrace the particle interpretation: such a view can account for the prominent epistemic fact that only limited knowledge of a portion of the material universe is needed in order to make reliable predictions about that portion. Thus, particle interpretations can support an abductive argument from the epistemic facts in favor of a realist reading of physical theory. However, any particle interpretation with this property is untenable. The empirical adequacy of modern particle theories requires adoption of a postulate known as permutation invariance (PI)—the claim that interchanging the role of two particles of the same kind in a dynamical state description results in a description of the identical state. It is the central claim of this essay that PI is incompatible with any particle interpretation strong enough to account for the epistemic facts. This incompatibility extends across *all* physical theories.

To frame and motivate the inconsistency argument, I begin by fixing the relevant notion of particle. To single out those accounts of greatest appeal to the realist, I develop the logically weakest particle ontology that entails the epistemic fact that the world is piecewise predictable, an ontology I call ‘minimal atomism’ (MA). The entire series of scientific conceptions of the particle, from Newton’s mechanically interacting corpuscles to the ‘centers of force’ in classical field theories, all comport with MA. As long as PI is left out, even quantum mechanics can be viewed this way. To assess the impact of PI on this picture, I present a framework for rigorously connecting interpretations to physical theories. In particular, I represent MA as a set of formal conditions on the models of physical theories, the mathematical structures taken to represent states of the world. I also formulate PI—originally introduced as a postulate of non-relativistic quantum mechanics—in theory independent terms. With all of these pieces in hand, I am then able to present a proof of the inconsistency of PI and MA.

In the second part of the essay, I survey responses to the inconsistency result open to the scientific realist. The two most plausible approaches involve abandoning particles in one way or another. The first alternative interpretation considered takes the property bearing objects of the world to be regions of space rather than particles. In this view, the properties once attributed to particles in quantum states are attributed instead to one or more regions of space. PI no

longer obtains in this case, at least not as a statement about the permutation symmetry of property bearers. Rather, the new interpretation naturally imposes an analogous constraint on quantum states.

The second major approach to evading the inconsistency result is to dispense with objects altogether. This is the recommendation of so-called ‘Ontic Structural Realism’ (OSR). The central OSR thesis is that structure rather than entities are the basic ontological components of the world. OSR is intended to embrace the ‘miracle’ argument in favor of scientific realism (it would be a miracle if a scientific theory were predictively successful unless it were also approximately true with regard to its description of reality) while avoiding the pessimistic meta-induction (most predictively successful theories along with their associated ontologies have been overturned, so we should expect the same of our current theories). I demonstrate that one principal motivation for OSR based on the under-determination of interpretations in QM is actually dissolved by the incompatibility result. At the same time, I suggest reasons to think that OSR fares no better with respect to the pessimistic meta-induction than traditional realism does. Thus, while PI and MA may be incompatible, object ontologies remain the best option for the realist.

Table of Contents

Chapter 1 Introduction.....	1
1.1 The claim	1
1.2 Particles are unobservable objects	1
1.3 Particles are everywhere.....	2
1.4 Particles are nowhere	2
1.5 A precis	3
Chapter 2 Realism and Particle Ontologies.....	5
2.1 Overview.....	5
2.2 Theories and their interpretations.....	5
2.3 Justifying Scientific Realism	8
2.4 Epistemic Divisibility.....	11
2.5 The Appeal of Particle Ontologies	13
Chapter 3 A minimal particle ontology.....	15
3.1 Overview.....	15
3.2 What is a particle?.....	15
3.3 Minimal atomism.....	17
3.4 Scientific atomism	21
3.5 MA and QM.....	25
3.6 Previous arguments against particles in QM	27
3.7 Conclusion.....	32
Chapter 4 Connecting metaphysics with physics.....	33
4.1 Overview.....	33
4.2 A taxonomy of theories and interpretations.....	33
4.3 Property independence as a formal constraint on theories	36
Chapter 5 Permutation Invariance.....	41
5.1 Overview.....	41
5.2 Formulation and Consequences in QM.....	41
5.3 Empirical support for PI in QM.....	44
5.4 Generalization of PI.....	49
Chapter 6 Proving an inconsistency	52

6.1 More definitions	52
6.2 The general argument	54
6.3 The argument for QM – version 1	55
6.4 The argument for QM – version 2	58
Chapter 7 Objections	61
7.1 Overview.....	61
7.2 Objection: QM tells us when to ignore PI.....	61
7.3 Objection: PI is compatible with classical physics.....	65
7.4 Objection: A metric is the wrong measure of distance.....	67
7.5 Objection: The result is old news.....	70
7.6 Objection: MA is the wrong account of particles.....	71
Chapter 8 A viable object ontology	72
8.1 Overview.....	72
8.2 Privileging the wavefunction.....	73
8.3 Spatial regions as objects	74
8.4 A clarification: single regions and multiple values	82
8.5 SOI and PI.....	83
8.6 SOI and EDiv	84
8.7 SOI and Scientific Practice.....	88
Chapter 9 Against structural realism.....	90
9.1 Overview.....	90
9.2 Two routes to OSR	90
9.3 Undermining Underdetermination	94
9.4 Theory succession, OSR, and SOI.....	96
Chapter 10 Conclusion	101
10.1 Recapitulation	101
10.2 What are macroscopic objects?	102
References.....	106

Chapter 1 Introduction

1.1 The claim

An influential interpretation of modern physics consists of the following assertions:

- (i) material objects are composed of particles;
- (ii) these particles belong to a finite number of types;
- (iii) particles of the same type are indistinguishable in every respect.

These assertions are manifest in the explanatory language of physics, chemistry, and molecular biology. I maintain that they are mutually inconsistent. It is the aim of this essay to clarify, motivate, and ultimately undermine the worldview they espouse, a view in which material objects are made of particles.

1.2 Particles are unobservable objects

Our everyday experience is populated by objects—entities like tables and coffee mugs. These objects possess properties, and the properties any one object possesses through time are independent of what properties most other objects possess. The properties of my coffee mug, for instance, are more or less uninfluenced by the properties of my computer. Many objects like the mug and computer are discrete, but even where divisions between objects are vague—where does the mountain end and the plain begin?—there are many ways of conceptually drawing the line such that whatever entities picked out persist in their properties without regard to one another.

Furthermore, the objects of our experience have smaller objects as parts. I could smash my coffee mug to pieces. I could then take a hammer to each shard to produce a pile of grains. It is no great stretch of the imagination to suppose that similar, even tinier objects can be produced by continuing the demolition process. Specifically, we can imagine the resolution of every material object into a collection of discrete, unobservable objects. These are the particles, the tiny objects that in enormous numbers are supposed to compose the objects of experience. Particles, like tables, bear properties. They persist, they interact, and groups of them do so more or less independently of all the others.

1.3 Particles are everywhere

The physicist Richard Feynman declared that the greatest scientific content one could pack into a single sentence was the claim that “all things are made of atoms — little particles that move around in perpetual motion, attracting each other when they are a little distance apart, but repelling upon being squeezed into one another” (Feynman, Leighton, and Sands 1963, 1.2). As we’ll see in the next chapter, there are good philosophical reasons for the scientific realist to take Feynman’s words at face value. There are also good phenomenological reasons. Many phenomena cry out for a particle interpretation: the integer ratios of chemical reactants, the tracks in bubble chambers, the integer charges on oil droplets, and the striking visualization of ‘atoms’ provided by scanning-tunneling microscopy to name but a few.

Whether or not the particle view is justified, it is everywhere in scientific discourse. Molecular biologists speak as if atoms and molecules are the mechanical constituents of macromolecules. Organic chemists concern themselves with the ‘mechanisms’ of chemical reactions, describing them in terms of the properties of specific atoms within a molecule, the bonds between them, and the manner in which constituent particles like protons and electrons are exchanged. Even in the relativistic quantum theories of physics, notwithstanding caveats about the inapplicability of classical notions, the particle view is ubiquitous. For instance, the computation of scattering amplitudes in quantum field theory—the principal way in which empirical predictions are extracted from the theory—is typically framed in a particle perspective: initial and final state representations are selected with the explicit intention of representing incoming and outgoing spatially isolated particles with well-defined momenta.¹ Though quantum physics has forced the scientific realist to abandon many classical aspects of her conception of particles, the particle view endures in the interpretation of physics.

1.4 Particles are nowhere

In the classical view of particles, all matter is of a kind. Particles, though they are discrete entities, differ only with respect to continuously varying properties such as mass and shape. To be compatible with modern physics, however, we have to admit one seemingly minor modification of this view—all particles belong to one of a finite number of types and every particle of a given type is indistinguishable from every other of that type. Of course, how many types of particle there are depends on the level of analysis; in chemistry the relevant types are the atomic elements while in particle physics they are the quarks, leptons, and bosons of the Standard Model. But at any given level of analysis, the particles of a

¹ See, e.g. (Peskin and Schroeder 1995), Section 4.5.

given type are presumed to be ‘indistinguishable’ in that they satisfy the following condition: interchanging the role of two such particles in a representation of a physical state results in a (possibly distinct) representation of the identical physical state. State representations with this property are said to satisfy *permutation invariance* (PI). If the state representations of our best physical theories capture all there is to say about the observable properties of particles, then particles represented by states obeying PI are exactly indistinguishable with respect to every possible observable property.

This apparently modest modification, however, is profoundly at odds with the particle view. As we will see, PI forces the scientific realist to attribute an identical set of properties to every particle of the same type. This means that such particles cannot bear properties independent of one another, and the central feature that made the particle view such a powerful explanatory apparatus must be abandoned. PI is forced on us by the phenomena and in turn forces us to abandon any substantive account of the world in terms of particles.

1.5 A precis

The first seven chapters of this essay serve to motivate, elucidate, and defend the incompatibility between particles and PI. I begin in the next chapter by showing that arguments similar to those for scientific realism in general favor particle interpretations in particular. In Chapter 3, I provide a rigorous statement of the logically minimal particle interpretation that is supported by the argument of Chapter 2. The interpretations of interest for the remainder of the essay will be those which include this minimal interpretation. In Chapter 4 I provide a framework for connecting interpretations to the theories of mathematical physics. The empirical need for the postulate of PI is taken up in Chapter 5, where it is given a theory-independent formulation. Finally, in Chapter 6 these pieces are combined to demonstrate the incompatibility between PI and even the logically weakest account of particles. It is shown that *any* theory which contains PI cannot be compatible with a particle interpretation. This result is defended from the most serious objections in Chapter 7.

The remainder of the essay is concerned with the consequences of the incompatibility result. In Chapter 8 I construct in some detail an alternative interpretation of non-relativistic quantum mechanics (QM). This interpretation posits spatial volumes as the basic objects out of which material objects are composed, and can be proven to possess the same virtues favored by the realist for particle interpretations. This interpretation is put to work in Chapter 9 in order to dispel the idea that the incompatibility result forces the realist to give up on objects and embrace an ontology of structure. On the other hand, I suggest that the ex-

istence of a viable alternative to particle interpretations does not lend any support to traditional entity realism either. The final chapter of this essay offers a brief consideration of a puzzle facing any realist ontology of unobservable micro-constituents: how should we conceive of the macroscopic objects that populate our experiential world? Whatever the right answer is, it won't involve particles.

Chapter 2 Realism and Particle Ontologies

2.1 Overview

As I stated at the outset, the addition of PI to any physical theory rules out a large class of interpretations, namely those which posit an ontology of particles. I also claimed that the scientific realist ought to be concerned with such a result since particle interpretations are both ubiquitous and appealing. However, I have to this point offered little support for this claim. In Chapter 3, I will argue that particle interpretations have been and remain a principal part of the scientific worldview, appearing as standard interpretations of theories of matter from the 17th century through the present. In each of these cases, however, the appeal of the associated particle ontology is a function of the details of the theory being interpreted. The aim of the present chapter is to give a much more general, theory-independent account of why the scientific realist does or at least ought to find particle ontologies attractive. To make this case, I'll adapt a standard argument in favor of realism in general in order to show that some basic facts about our epistemic access to the world tend to favor particle ontologies in particular. This in turn motivates concern over the consequences of PI that are drawn out in later chapters.

2.2 Theories and their interpretations

Realism about particles is a species of scientific realism. At the expense of blurring together many nuanced positions, one can characterize scientific realism as the claim that what our best scientific theories say about the unobservable entities, processes, or events in the world is approximately true.² To be a little more careful, the realist is not committed to the truth of any particular scientific theory, but rather to the claim that insofar as one has reason to believe that a given theory is scientifically successful—empirically, explanatorily, or in whatever way such things should be judged—then one also has reason to believe that its claims about the world are true. For my purposes, it is not essential to pin

² For similar but more nuanced characterizations of scientific realism, see e.g. (Boyd 1973; Chakravartty 2007; Laudan 1981; van Fraassen 1980).

down precisely what is meant by ‘success’. It is important, however, to draw an often overlooked distinction between a theory and its interpretations.

In the most general sense, I take a *theory* to consist of the following parts:³

- (i) a collection of syntactic objects such as statements in a formal language, a mathematical space (e.g. a phase space), a set of diagrams deemed well-formed (e.g. the structural diagrams of organic chemistry), et cetera;
- (ii) a deductive system which, when provided a subset of the syntactic objects in (i) known as the set of ‘boundary conditions’, specifies a second subset of the syntactic objects in (i). The deductive system might consist of a set of differential equations such as those of Hamiltonian dynamics, or something less formal like a system of rules for determining the oxidation number of atoms in a compound.⁴

Models of the theory are just those syntactic objects (e.g. trajectories in phase space or lists of oxidation numbers) that are possible outputs of the deductive system of the theory when coupled with a set of possible boundary conditions. Finally, we can say that an *interpretation* of a theory is a mapping from models of the theory to descriptions of theoretically possible worlds. These descriptions amount to claims about the way the world is or might be. I’ll often refer to a particular description of a possible world as a ‘state of affairs’.

Every theory requires at least a partial interpretation if it is to entail empirical consequences. That is, every theory must be equipped with a means of extracting claims about observables from its models if it is to have any empirical content. This much the realist and anti-realist can agree on. But it is generally possible to map models of a theory to richer descriptions of the world that contain assertions about more than observable states of affairs. According to the scien-

³ I do not claim that my choice of terminology is standard or representative of the literature. Rather it was chosen to make explicit a number of distinctions that are too often left vague and implicit.

⁴ As a consequence of the way I am defining theories there is no single theory of, say, Hamiltonian mechanics but rather one theory for every new combination of mathematical space and equations of motion. For example, a two-particle system requires the use of a six-dimensional phase space where a one-particle system requires only a three-dimensional phase space. Similarly, the equations of motion contain more terms and derivatives with respect to more variables. Thus, there are distinct versions of the two theoretical components I’ve stipulated in my definition, and therefore two distinct ‘theories’. While this is perhaps linguistically odd, the extra precision is helpful when exploring the compatibility of theories and interpretations.

tific realist, one of these interpretations—at least for a successful theory—is likely to yield true claims about unobservable features of the world.

Of course, the realist does not consider every assignment of states of affairs to models of a theory a viable interpretation. I do not want to attempt a set of necessary and sufficient conditions for viability (I doubt any exist that would adequately capture what have historically been taken as serious interpretations), but I will suggest that the acceptable accounts of unobservables are typically those which explain the relations amongst observables. So for purposes of exploring the arguments of the scientific realist, I'll assume that the viable interpretations of a theory are those for which the claims about the unobservable portions of the world taken by themselves would make the observable statements entailed by that theory probable. For instance, if there really exist chemical atoms of the sort Dalton imagined, then it would have to be the case that chemical reactions involve integer ratios of reagents. Dalton's atomism is a viable interpretation of at least a portion of chemical theory. On the other hand, Thales's supposition that "all is water" does not make the integer ratios found in chemical reactions more or less probable, and so is not a viable interpretation.⁵

Even with the constraints imposed on viable interpretations, it is still the case that theories wear their interpretations loosely. While scientists nearly always have a preferred interpretation in mind, one can nonetheless attach many ontologies to the syntactic components⁶ of any given theory. This is true of qualitative theories like, for instance, the chemical theory of valence bonding. Though there is an obvious reading suggested by the words chosen for use as theoretical terms (e.g. atomic orbital), a chemist is only committed to the predictions made by the theory concerning various chemical outcomes. The predictive claims can be logically separated from the story one tells about the meanings of non-observational theoretical terms. While one might interpret a valence bond as referring to the overlap of two atomic orbitals, one could also interpret a valence bond as referring to a pair of electrons of distinct, atom-specific varieties, or to a region of increased charge density, or to any number of more fanciful entities.

The distinction I've drawn between a theory and its interpretation is much sharper for the mathematical theories of physics than for the semi-phenomenological chemical theories considered above. In this case, theories

⁵ One is also plausibly constrained by concerns about theoretical unity – interpreting chemical reactivity in terms of Aristotelian *potentia* would not cohere well with the established interpretations of other parts of chemistry.

⁶ The logical positivists would read 'syntactic components' as 'theoretical terms' since theories in that view are statements in a formal language.

consist of equations of constraint that pick out mathematical objects—the models—from a well-defined space of possibilities. The models of a theory of mathematical physics are not nearly as suggestive about their intended referents in the world as are the plain-language terms used in stating the valence bond theory. There are in fact many *prima facie* plausible ways of interpreting the models of each such theory in terms of entities or processes in the world. For example, even though the formal structure of QM has been more or less settled since the 1930s, it continues to accrue a bewildering array of interpretations.⁷

Conceptually separating the formal structure of a theory from its interpretations is important, since only by doing so is it coherent to ask about the relative support for competing realist readings of the same formal theory (e.g. QM). The realist wants to know whether a particle interpretation or some more exotic alternative is the best way to read QM. Obviously, this question is without content if a theory and its full interpretation are inseparable. With this distinction in mind it is clear that a realist is concerned primarily with the truth of the interpretations of a theory, not with the truth of theories, whatever that may mean. Thus, one can recast scientific realism as the claim that what the favored interpretations of the best scientific theories tell us about entities and processes is a more or less accurate description of the way the world really is.

2.3 Justifying Scientific Realism

The dominant argument in favor of scientific realism is the infamous ‘miracle’ argument memorably summarized by Hilary Putnam (1979): “Realism is the only philosophy of science that doesn’t make the success of science a miracle.” Typically, the argument is taken to appeal to some version of ‘inference to the best explanation’. I use this controversial term to refer loosely to the class of inductive inferences with the following form:

Inference to the Best Explanation (IBE):

A, B, C, ... are facts.

Hypothesis H is the best available explanation of the facts A, B, C, ...

There is reason to believe H is true.

⁷ For a brief but current survey, see (Dickson 2007), Section 5.5. The reader should note that Dickson’s notion of ‘interpretation’ is a little different from mine. For a dated but apropos review of interpretations in the sense I have used the term, see (Bunge 1956). The following are among the more prominent interpretations: Copenhagen (Faye 2008), Bohmian (Bohm and Hiley 1985), Everett (Geroch 1984), modal (Dickson and Dieks 2009), and transactional (Cramer 1986).

There have been many attempts to make this schema more explicit with the aim of defending such inferences in general and the miracle argument in particular. Alan Musgrave (2007), for instance, adds an additional premise to IBE in order to make inferences of this sort deductively valid:

It is reasonable to believe that the best available explanation of any fact is true.
 A, B, C, ... are facts.
 Hypothesis H is the best available explanation of the facts A, B, C, ...

 It is reasonable to believe that H is true.

Of course, one might argue that Musgrave has merely reduced the problem of justifying IBE to a problem previously unsolved—it is not clear why explanation has anything at all to do with truth.⁸

Another way of unpacking IBE and the no-miracles argument is to appeal to the probability calculus (Magnus and Callender 2004). If we know A, B, C, ... to be true then $\Pr(A \wedge B \wedge C \wedge \dots) = 1$. We can then interpret the second premise of IBE as asserting that $\Pr(A \wedge B \wedge C \wedge \dots | H) \gg 0$ and $\Pr(A \wedge B \wedge C \wedge \dots | \neg H) \ll 1$. It is then supposed to follow that $\Pr(H | A \wedge B \wedge C \wedge \dots) \gg 0$. To make the inference explicit in the case of scientific realism, let $S(x)$ stand for ‘theory x is empirically successful’, and $T(x)$ stand for ‘theory x is true of the world’. We then take as a premise that $\Pr(S(x) | T(x)) \gg 0$ and that $\Pr(S(x) | \neg T(x)) \ll 1$. That is, if a theory is true of the world, it is overwhelmingly probable that it is empirically successful. On the other hand (and rather controversially) we can assume that if a theory is false then it is very unlikely to be successful. From these two premises and the fact that a given theory h is empirically successful (i.e. $\Pr(S(h)) \gg 0$) it ostensibly follows that h is very probably true of the world (i.e. $\Pr(T(h) | S(h)) \gg 0$). Of course, as Magnus and Callender (2004) point out, this argument is fallacious unless we also have a means of fixing the prior probability that a theory is true at some moderately high value. Without knowing the value of $\Pr(T(h))$ we can say nothing about the posterior probability, $\Pr(T(h) | S(h))$.

Even if there exists some reasons for attributing a value to the prior probability such that the argument is no longer invalid, the remaining premises are subject

⁸ While there are many attempts in the literature to formulate just what constitutes an explanatory relation, they share the supposition that ‘explanation’ involves making something (an event, a proposition, an action, etc.) intelligible. While it might be possible to spell out some subject-independent notion of intelligibility, the notion is dangerously close to a psychological one. If scientific ‘truth’ is determined by intelligibility, this would seem to invite any number of anti-realist readings.

to a number of serious objections, particularly worries about the claim that $\Pr(S(x)|\neg T(x)) \ll 1$. There are other ways to account for the empirical success of a theory besides its literal truth. Bas van Fraassen (1980, 40), for instance, argues that empirical success is merely an outcome of Darwinian selection—the theories that have survived are just those whose predictions line up with experience. Van Fraassen’s objection can be strengthened if we take into account the theory/interpretation distinction I emphasized above. For any one theory there exist multiple consistent interpretations. Let $T(y)$ stand for ‘the *interpretation* y is true of the world’. If h is a particular theory and a_i is a compatible interpretation, then our old premises must be revised to read $\Pr(S(h)|T(a_i)) \gg 0$ and $\Pr(S(h)|\neg T(a_i)) \ll 1$. The second premise is now obviously problematic for two reasons: (i) there are many alternate interpretations a_j that might just as well account for theory h ’s empirical success and (ii) it is possible that every available interpretation is false and the theory is nonetheless empirically successful. In either case, it may be that $\Pr(S(h)|T(a_i)) \gg 0$ and simultaneously false that $\Pr(S(h)|\neg T(a_i)) \ll 1$.⁹

The difficulty we’ve encountered is related to the under-determination problem that complicates the use of IBE in theory choice from even a strict empiricist standpoint. In that case, the compatibility of indefinitely many theories with the same evidence means that falsifying any one theory generally fails to pick out a single alternate theory as the best. There is, however, one important difference. With respect to theory choice, it makes no sense to contrast the probability of the evidence given the theory with the probability of the evidence given that *no* theory accounts for it. For any finite amount of evidence there necessarily exists an indefinite number of theories that can reproduce it—it is not possible for all theories to fail empirically. On the other hand, it is conceivable that all interpretations of a given theory—all attempts to map models of a theory onto possible worlds with a fixed ontology—yield false claims about the world. This is just what the anti-realist claims is the case. That means we can sensibly compare the probability of empirical success given that *no* interpretation of the theory is true with the probability of empirical success given the truth of various interpretations.

Because it is coherent to consider the case in which all interpretations are false, we can avoid the suite of problems that follow from multiple interpretations by

⁹ There is another difficulty I’ve been ignoring: we don’t actually know that any given theory is empirically adequate since we only have a finite amount of data. So we have to either take $S(h)$ to mean “as empirically adequate as we’ve had occasion to test” or set $\Pr(S(h))$ to some lower value. If we attempt the latter, it is hard to see what value to use. The choice makes a big difference, which is in part what Magnus and Callender (2004) are getting at.

appealing to the Likelihood Principle (LP).¹⁰ According to LP, evidence e favors hypothesis h_1 over hypothesis h_2 just if $\Pr(e|h_1) > \Pr(e|h_2)$. Thus, while we may not be able to ascertain the probability of the truth of an interpretation, we can nonetheless favor one interpretation over another or over the falsity of all interpretations on the basis of the available evidence. According to LP, the fact that, for a given interpretation a_i , $\Pr(S(h)|T(a_i)) \gg \Pr(S(h)|\forall x \neg T(x))$ means that the empirical success of theory b favors the truth of a_i over the falsity of all compatible interpretations.¹¹ That is, the empirical success of theory b favors realism about each compatible interpretation over anti-realism.

This version of the no miracles argument is much less impressive than Putnam's slogan; it only establishes that at least some interpretations are favored by the evidence relative to the anti-realist stance. It does not single out a globally 'best explanation'. At best, the use of LP only lets one rank available interpretations relative to one another. There is little to guarantee convergence to the true interpretation. Nonetheless, this is a defensible explication of the no-miracles argument.

2.4 Epistemic Divisibility

I am not interested in defending the miracle argument; I leave that to the scientific realist. Rather, I want to bring forward a version of the argument that specifically supports realism about particle ontologies. Whatever force the miracle argument carries for the realist, the same will obtain for this derivative argument. The idea is to show that insofar as the realist has good reason to maintain that some interpretation is true, she has good reason to think that some particle interpretation is true. For this reason, the realist ought to be concerned about whether any particle interpretations are in fact compatible with physical theory.

Like the miracle argument, the argument for particles rests on the observation of empirical success. But rather than consider all of the ways in which a particular scientific theory and an associated interpretation are empirically successful,

¹⁰ I don't want to defend LP as the correct explication of IBE since it actually doesn't say anything about what's true or best, just what's better supported. But I do take it as underwriting the most plausible (i.e. defensible) reconstruction of the no-miracles argument.

¹¹ The assertion that empirical success is much more probable given the truth of a theory (or at least some theory)—versions of which appear in all defenses of the miracle argument—can reasonably be argued to beg the question against the empiricist. If order amongst observables is a brute fact of nature, then the assertion is false. On the other hand, the inference might be supported by appeal to something like a Principle of Sufficient Reason that effectively denies the possibility of irreducible order amongst observables.

let us consider just one simple feature of any moderately successful account of the world: it anticipates the future of some part of the world from a vanishingly small portion of facts about the current state of the world. So for instance, to predict the outcome of a chemical reaction a chemist need only know a few facts about the volume and composition of her reagents; she needn't take account of the density of seawater off the coast of Greenland or the mean temperature of Alpha Centauri, or even of the masses of the beakers holding her reagents. To predict where an artillery shell will land, one need only account for the angle of the gun barrel and the velocity of the shot as it exits the muzzle. It is not necessary to know the details of the shell's manufacture or the properties of the air behind the muzzle. The very possibility of empirical science depends upon this ability to fracture the world into epistemically independent pieces. I'll call the fact that the future of a part of the world can be predicted from limited facts about that part alone 'Epistemic Divisibility' (EDiv).

I want to stress that EDiv is not a locality condition. Locality conditions are frequently invoked in physics, and can be read as either epistemic or ontological assertions of the independence of space-like separated physical systems. A common ontological approach asserts that spatially distant systems are physically independent (or causally independent depending on your predilections). Such an ontological claim further entails an epistemic thesis similar to but not identical with EDiv, namely that in predicting the future of some part of the world, we can ignore all other parts that are sufficiently distant in space (exactly how far will depend on the theory at hand). The epistemic form of any given locality condition is logically stronger than EDiv; if locality holds, so too does EDiv.

However, the converse does not hold. It is possible for the world to be such that it divides into epistemically independent units according to something other than spatial separation. Were I writing in the centuries prior to Copernicus when astrology remained a viable science, it would be plausible to assert that vastly separated physical systems nonetheless exert a profound influence upon one another. In a world where astrology works, it would be implausible to assert a general locality condition since I really might need to know the position of Jupiter relative to the stars in the constellation of Cancer in order to predict the outcome of my chemical experiment. Or consider instead one of Dalton's early theories in which the atoms of a gas are posited to be centers of repulsion and each atom exerts a repulsive force *only on atoms of the same element* (Nash 1950, 19-20). Though Dalton himself did not do so, we might imagine further that each atom exerts a non-negligible force on every other atom of the same type irrespective of distance. But since there would be no interaction between atoms of different elements, the world would still divide into units that are independent in

the sense that EDiv asserts, at least if there were sufficiently many chemical elements. One wouldn't need to know what state the oxygen atoms are in to predict the future of the carbon. The upshot is that a locality constraint is just one way in which EDiv might manifest, but it is not the only way.

Whether a suitably framed locality condition is true of the world is controversial, but the fact of EDiv is not. For a theory to be empirically successful it must make accurate predictions about the future state of some system on the basis of what is necessarily finite, limited information about the system. Newtonian mechanics, for instance, is highly empirically successful for a broad domain of phenomena in that it allows one to predict, for instance, the locations of the planets on the basis of a finite number of past positions. The empirical success of any scientific theory thus entails the fact of EDiv. Since it is not the case that all current or conceivable theories also contain or require a spatial locality condition, let us take EDiv simpliciter (rather than locality) as a necessary condition for the empirical success of any scientific theory.

2.5 The Appeal of Particle Ontologies

It is possible to frame a stripped-down version of the miracle argument for any member of the special class of interpretations that entail the truth of EDiv. This special class of ontologies consists of all those which assert that the world is composed of ontologically independent entities. Roughly, ontological independence is a modal claim: if two entities are ontologically independent then it is possible for either to have existed, occurred, or evolved in time as it has irrespective of the existence, occurrence, or evolution of the other. Such ontologies stand in contrast to radically holistic or monistic ontologies such as the 'blobjectivism' defended by Horgan and Potrč (2002).

To construct a version of the miracle argument, let \mathcal{Q} designate the special class of interpretations that posit ontologically independent entities. Every $q \in \mathcal{Q}$ entails EDiv, and so $\Pr(\text{EDiv}|T(q)) = 1$. On the other hand, if no interpretation in \mathcal{Q} is true of the world, then it would be a 'miracle' if EDiv were true. Thus, it must be the case that $\Pr(\text{EDiv}|\forall_x \neg T(x)) \ll 1$. So for any scientific theory b paired with a compatible interpretation $q \in \mathcal{Q}$ we have the following: $\Pr(\text{EDiv}|T(q)) \gg \Pr(\text{EDiv}|\forall_x \neg T(x))$. Note that I have not appealed to the details of the empirical success of b , only the fact that any such success entails the truth of EDiv. Because the inequality holds for any $q \in \mathcal{Q}$ compatible with b , LP leads us to conclude that all of these interpretations are to be favored over the anti-realist stance. The same argument form used to defend realism about

particular theories—when supplied with only minimal facts about empirical success—gives us a very general argument in favor of those interpretations which posit ontologically independent entities.

We have not yet established the appeal of *particle* ontologies. But it is a trivial extension of the argument to note that what I am calling particle ontologies¹²—interpretations which posit discrete entities (the particles) which bear properties largely independent of one another—sit squarely in \mathcal{Q} . That is, particles just are discrete ontologically independent units, and so particle interpretations are in this special class. Of course, particle ontologies are not the only ontologies which can support EDiv; I'll introduce another possibility in Chapter 8. However, of the available scientific interpretations, particle ontologies make up the majority.

The upshot of this discussion is not that scientific realism is likely to be true. Nor is it that the world really is made of particles. Rather, I am claiming that if one is a scientific realist, then one should be concerned if particle ontologies fail, since this class of ontologies is supported by an argument very closely related to the miracle argument—if one is compelling, so too is the other. Particle interpretations are prime candidates for true scientific claims about unobservable entities. If they fail, realism loses an appealing exemplar.

¹² The notion of 'particle' as I introduced it in Chapter 1 is, as I'll show in Chapter 3, consistent with use of the term in scientific interpretations. One may, of course, insist on a different definition, one that makes particle interpretations no longer members of \mathcal{Q} . But such a notion is no longer supported by the argument from EDiv and is alien to the particle worldview manifest in scientific discourse. When I refer to particle ontologies, I will just mean those that posit discrete, ontologically independent entities.

Chapter 3 A minimal particle ontology

3.1 Overview

In the last chapter I sketched an argument in favor of realism about particles, but did so by appealing to an undeveloped notion of ‘particle’. In this chapter, I spell out precisely what counts as a particle ontology by presenting the logically weakest set of assumptions such an ontology must satisfy. Two considerations come into play when laying out such a minimal set. The first concerns the meaning of the term ‘particle’. The word has been used in many ways, so I stipulate which sense is relevant to my project. Second, not every ontology satisfying that sense of ‘particle’ constitutes a plausible or appealing ontology for the realist. In the previous chapter, we saw that the appeal of particle ontologies derives from the fact that they can account for EDiv and so we will focus our attention on just such ontologies. In the next section, I take up the first of these tasks—specifying the relevant sense of ‘particle’. In Section 3.3, I consider which particle ontologies in this sense can account for EDiv, and use these considerations to motivate the promised minimal ontology. The assumptions of this minimal ontology delineate a very large class of ontologies, each produced by adding one or more elaborations. In Section 3.4 I argue that the various forms of scientific atomism from the 17th century onward fall within this class. In Section 3.5, I show that even QM can be given an interpretation in this class as long as PI is omitted.

3.2 What is a particle?

While the meaning of the term ‘particle’ as it applies to matter has evolved over time, it has retained a stable core amounting to something like ‘a discrete constituent of some material body’. More precisely, a particle is a discrete individual bearing properties and which, along with other particles, may be part of a composite body that also bears properties. When I say that particles are individuals I mean that each stands in definite identity relations with the others (each is identical only with itself and not identical with the others) and that each can be the value of a logical variable. This further entails that particles can be grouped into sets, labeled, and individually attributed properties. I mean nothing more metaphysical than this. Without this meager notion of individuality, we could make no sense of the attribution of properties to particles.

More importantly, to say that particles are *discrete* entities is to say something about the divisibility of material objects. Roughly speaking, to say that a material object is composed of discrete particles is just to say that there is a fixed and determinate set of objects—the particles—into which the material object can be completely decomposed. Classically, we would further stipulate that any finite material object (i.e. finite mass, spatial extent, etc.) is composed of finitely many, spatially separable, indivisible particles. In this view, if we were to keep dividing the object into smaller and smaller pieces, the process would eventually bottom out with a finite set of particles.¹³ But this is much more than we need to account for EDiv and so more than we need to read into the notion of ‘particle’. I will not insist on these further classical stipulations, but no harm is done if the reader thinks in those terms. Of course, without these additional stipulations, all sorts of things not generally associated with particle ontologies are admitted. For instance, insofar as we think of an ideal fluid as an uncountable set of material points, each with a set of properties, then the fluid can be considered to consist of particles in the very general sense given above. Yet, since it would be infinitely divisible, such a fluid would surely not admit a classical particle composition. I am willing to admit such strange cases for the sake of generality, so I will proceed on the assumption that a ‘particle’ is just one part in the determinate set of parts into which a material body can be divided.

I should take care to note that we need not assume there are any mereological simples in order to assert that the material objects are finitely divisible into particles. The statement about divisibility should be understood as relative to a process or set of ‘fundamental’ particles. Each of these particles may in turn be composed of other particles. So in chemistry, for instance, the fundamental particles are the atoms of chemical elements. Insofar as we assume chemical elements to be fundamental, there is a fixed and determinate set of atoms of the chemical elements into which every material body can be divided without any material parts left over. Particle physics on the other hand recognizes the particles of the standard model as the fundamental units into which nothing can be further divided. Whether or not this is true is immaterial for our purposes. All that matters is that positing an ontology of particles means positing a set of components into which material things can be completely resolved at some level of description.

¹³ These assumptions need not be merely stipulated. They follow naturally if we embrace some plausible assumptions about the properties of particles and their aggregates. If, for instance, we insist (i) that particles bear some state-independent properties (like mass) that are represented by an outer measure function, (ii) that each particle has non-zero measure (a finite mass), and (iii) that material objects possess a value of the same property represented by the additive measure, then a universe of countably many objects with finite property values (countably many objects of finite mass) entails countably many particles.

3.3 Minimal atomism

While the decomposability of material bodies is a necessary condition for an ontology to count as a particle ontology, it is insufficient for securing the epistemic facts. So the notion of particle developed in the last section is still too inclusive for my purposes. In Chapter 2, I argued that a slightly stronger notion of particle is needed in order for the proponent of particle realism to garner support from empirical science. Specifically, the realist needs an ontology strong enough to account for EDiv. I will henceforth use the term ‘particle ontology’ to refer just to those ontologies which posit discrete, property-bearing individuals that compose material bodies *and which entail the truth of EDiv*. I maintain that the following set of assumptions is the logically weakest ontology of discrete entities that entails EDiv, and thus the weakest such particle ontology:

minimal atomism (MA):

- (i) Every material object can be decomposed into a collection of discrete objects called ‘particles’ each with a non-empty set of state-independent properties.
- (ii) Every non-empty set of particles possesses at least one state-dependent monadic property.
- (iii) For every set S of particles and for every $\sigma \subset S$, it is the case that for most physically possible conditions, the state-dependent monadic properties of σ are approximately independent of the properties of most subsets of the complement of σ in S over a finite interval of time.

These conditions require clarification on a number of points. To begin, by a ‘state-independent’ property I mean one which remains invariant for a given particle within or across physically possible worlds. Conversely, by state-dependent I mean a property which can assume different values at different times for a particle in a given physically possible world, or different values for the same particle in different physically possible worlds. A monadic property is one borne by a single object. On one view, these might be seen as the non-relational properties of objects, and I will henceforth assume that this is the case. However, one could still be a strict relationist and assert that there are only relational properties without contradicting MA. To spell out MA in strictly relation terms is complex, and I will not do so here. However, the advocate of relationalism is invited to read ‘monadic’ as a shorthand for whatever properties we predicate of single objects, even if these properties are taken to be relational in

some sense. So for instance, I will speak of momentum as if it is a property possessed by a particle. However, in a relationist view of space, this property refers to one or more other objects in relation to which velocity is defined.

Two more terms introduced in MA will be given only a short gloss here, since later chapters are devoted to providing a full account. The first of these is ‘physical possibility’. To say that a state of affairs is physically possible is roughly to say that—for whatever theory and interpretation pair we take to be true of the world—the given state of affairs obtains in the possible world described by one of the theory’s interpreted models. The second is the critical notion of the ‘approximate independence’ of properties. A full account is given in Chapter 4; for now, we can say that the properties of object x are approximately independent of the properties of object y if it is physically possible for x to assume almost any of the properties it might have had if y did not exist and vice versa. For instance, the properties my coffee cup can assume within the limits of the manipulations I can perform in my office would be—so far as I could discern via measurement—indistinguishable from what they would be if my pen were removed from my desk and destroyed.

It remains to show that the conditions of MA are in fact sufficient to entail EDiv. Let S stand for the set of all particles in the universe and let $\sigma \subset S$ stand for the set of particles composing a material object like my coffee cup. Since I have stipulated that σ (my coffee cup) is a material body, we know from MA(i) that σ is non-empty—it contains at least one particle. Furthermore, from MA(ii) we know that σ bears at least one state-dependent property. My coffee cup, for instance, bears the state-dependent properties of momentum, density, and temperature. Now consider the complement of σ in S , denoted σ^c , which contains all of the particles in the universe except for those making up my coffee cup. Every material body other than my coffee cup and its parts—from my desk to distant galaxies—can be identified with some subset of σ^c . From MA(iii), we know that under most physically possible conditions—under most conditions in which we might find the universe—the properties of σ are approximately independent of most subsets of σ^c . That is, the properties of the material object σ are approximately independent of the properties of most other material objects—whether single particles or composites. We can expect the properties of my coffee cup by and large to be independent of the properties of all the atoms in all the other objects in the universe—stars, planets, oceans, atmosphere, etc—with the exception of a small number of objects like my desk with which it strongly interacts. This means that in order to predict the future properties of σ we need only have knowledge of the properties of at most a minority of the subsets of σ^c . That is, to predict the future properties of my coffee cup we need

only know about the properties of at most a minority of the other objects in the universe (a vanishingly small minority in most cases). But this is precisely the content of EDiv. The epistemic facts follow from the assumptions of MA—if MA is true then so is EDiv. I should note that MA does not tell us anything about *which* subsets of σ^C we need to know about. That is the purview of physics. MA is not a stipulation of the correct theory of the physical world, only in effect a guarantee that a tractable physics is possible. Just as EDiv is weaker than an epistemic locality condition, so MA(iii) is weaker than an ontological locality condition. It may be the case that particle independence aligns with spatial separation as in Newtonian physics. But it might also be the case that independence tracks some other property—as in the example of Section 2.4—or no property at all.

In deriving EDiv from MA, I have implicitly invoked a sort of determinism. It is only possible to anticipate the future from a set of facts about the past and present—no matter how complete—if the future is determined by the present state of affairs. This determination need only be in the weak sense that probability distributions over possible observable states of affairs at $t + \Delta t$ are determined by the state at t . At first glance, this assumption might appear problematic. For instance, one might object at this point that even classical mechanics admits various forms of indeterminism that contradict this claim. For instance, because it imposes no speed limit, there are models of classical mechanics that contain “space-invaders”, objects that arrive from spatial infinity in a finite time (Earman 1986). In a different vein, Newton’s Laws when coupled with certain problematic initial conditions present ill-posed initial-value problems. The ‘Norton Dome’—for which a ball can begin rolling away from an unstable equilibrium at any finite time—is in this class (Norton 2008). Regardless of the specific pathology, the point is that, in some models of classical mechanics, even complete knowledge of the state of the world at one time is insufficient for predicting the state of the world at future times. In the space invader case, models in which the invader appears and in which it doesn’t are identical up through the time the invader arrives. For the Norton Dome, it is impossible to predict at time $t=0$ the time at which the ball will commence rolling. This seems to contradict even the mild determinism we need to derive EDiv from MA. As weak as it is, it would thus seem that MA excludes even classical mechanics.¹⁴

To draw this conclusion, however, is to require too much of MA. Recall that EDiv only asserts that, under most conditions, it is possible to reliably predict the future of most isolated bits of the world. To secure this claim, we do not need to require that *every* model of a theory be deterministic or that the future of

¹⁴ This concern was raised by Bryan Roberts in personal communication.

every region in every model be determined by its past. It merely has to be the case that, for any given set of particles, *most* models which contain them are such that these particles are independent of most others and have their future determined by their initial state. Space invaders are only problematic if they invade most regions of space in most models. Norton Domes should trouble us only if most initial conditions produce them. While difficult to formulate precisely, it is implausible to claim that the models of classical mechanics are dominated by indeterminism, even if we leave the set of viable initial conditions unrestricted. However, even if it is true that the space of classical mechanical models with unrestricted initial conditions is saturated with such indeterminism, theories include a specification of the class of viable boundary conditions. That is, theories are not merely naked dynamical equations—a class of models and thus a class of interpretations is only specified when the class of possible boundary conditions is stipulated. The initial conditions that yield Norton Domes are plausibly excluded from the class of physically possible conditions associated with the theory. Classical mechanics understood in this way is sufficiently deterministic to support the inference to EDiv, as are a great many other theories of microphysics.

To head off the worry that MA excludes quantum theories from the outset, I should note that even QM generally possesses the sort of ‘determinism’ required for the above derivation. While specifying an observable state of affairs (i.e. stipulating values of a complete set of commuting observables for a system) is insufficient for fixing the outcomes of measurements made later on the system, it is sufficient for fixing the probability (which can be read as a propensity or disposition) of obtaining each possible measurable outcome. In the quantum domain, EDiv amounts to the claim that we can predict the probability of a measurement outcome for a system on the basis of limited knowledge about that system alone. The relevant probabilities at future times are determined by the state of the world at earlier times, and so the relevant determinism obtains. To put it another way, the quantum state of the world evolves deterministically in the manner assumed in my derivation of EDiv, even if the outcomes of measurements are not fixed by the quantum state. The determinism implicit in moving from MA to EDiv is a harmless background assumption that does no work in ruling in or out particle interpretations for the physical theories of interest, and so I will continue to accept it as a given.

Not only does MA entail EDiv but, as suggested by the name, it is minimal—if we drop any of the assumptions in MA then EDiv fails to follow. Of course, if we drop MA(i) what’s left would not warrant the appellation of ‘particle ontology’, at least not as I’ve explicated the notion of ‘particle’. But it is also the case

that without MA(i) EDiv fails to follow from the remaining propositions. This is because MA(i) guarantees that, if there are at least two objects in the universe, then there are at least two particles in the universe that can bear properties and one can sensibly talk about independence. Without this stipulation, it might be the case that the universe of material objects consists of only a single partless whole¹⁵ and thus MA(iii) would be trivially satisfied even though EDiv does not obtain. This is an exotic possibility but nevertheless one we need to rule out. Given MA(i) and MA(iii), EDiv fails to follow without MA(ii) because it may be the case that the only monadic properties particles possess are state-independent (e.g. mass). If this were the case, then MA(iii) would again be trivially satisfied but there would no longer be any sense in talking about the changing state of some portion of the universe. With only state-independent properties, an object does not change through time—there would be nothing to predict and EDiv would be irrelevant, though not strictly false. When I say an object with only state-independent monadic properties does not change in time I do not intend to deny that its relational properties may change. However, EDiv is an assertion about the features of an object in the world that can be said to change without reference to any other object in the world. These are just the monadic properties of the object in question. Finally, MA(iii) is clearly essential in establishing EDiv. It guarantees an ontological independence of properties amongst the parts of the universe (the existence of which are established by MA(i) and (ii)) that in turn entails the epistemic independence of EDiv. Without MA(iii), propositions MA(i) and (ii) are compatible with EDiv, but are not strong enough to imply the epistemic facts we observe.

3.4 Scientific atomism

If MA looks austere, that is by design—I want to capture the broadest class of theories with the virtue of accounting for EDiv. If MA looks overly metaphysical or distinct from scientific conceptions, that is an impression I wish to dispel. While I haven't the space to present the rich and complex history of particle ideas in its entirety, I can at least point to a couple of representative theories to suggest that MA is consistent with each of the major particle interpretations of microphysics. From the 17th century onward, scientific accounts of atoms or particles have been elaborations of the minimal particle ontology specified by MA.

¹⁵ While perhaps implausible, the view that the universe is the only concrete particular and possess no proper parts has its defenders [e.g. (Horgan and Potrč 2000, 2002)].

Two themes of conceptual development merged in the seventeenth century to produce the first scientific particle theories.¹⁶ On the one hand, an account of chemical change in terms of minimal units with chemical properties—descendants of the Aristotelian natural minima—increasingly looked like a plausible approach to understanding the growing body of data on chemical and physical transformations. On the other hand, the ascendant mechanical philosophy with its emphasis on quantitative primary qualities sought to account for the behavior of matter in terms of inertia, motion, and extension; this would have to include the particulate components of physical bodies insofar as the latter can be viewed as autonomous bodies bearing physical properties of their own.¹⁷ Both approaches are synthesized for the first time in the chemical works of Robert Boyle.¹⁸ In Boyles' view, the world is full of insensible particles that are natural minima in the sense that nature seldom if ever divides them. These tiny particles have their own "determinate bulk and shape" (Boyle 1991, 42). Moreover, they can bind tightly together to make composite particles that in turn comprise even larger units or chunks of bulk matter (Boyle 1949, 30-1). That is, they bind together to form what we might call molecules in modern parlance. The basic particles themselves, like the corpuscles of Descartes or Gassendi, have no chemically relevant qualities. However, the shape and topology of the composite particles (the molecules) constitutes what Boyle calls a "permanent texture" (van Melsen 1952, 103). These permanent textures determine the chemical and sensible qualities of bulk matter which is made up of such composite particles.

Boyle's ontology plainly satisfies MA(i) and MA(ii). It also satisfies MA(iii). The aggregation and disaggregation of particles which accounts for both chemical change (via formation of new molecules) and physical change like evaporation (via change in the motion or relative location of the molecules) is mediated by contact and local motion. This means that the properties of the material objects in any region of the universe depend only upon the properties—such as motion and texture—of the particles within that region and upon those particles whose motion carries them into that region over the time period of interest. In other words, the properties of the particles composing some material body are entirely independent of all other particles with the exception of those which enter the

¹⁶ For an elegant survey of the development and merger of Aristotelian and atomistic lines of thought, see (van Melsen 1952). For an overview of the 17th century transition, see (Chalmers 2009).

¹⁷ See, e.g., (Whyte 1961).

¹⁸ Drawing a sharp line making Boyle the first to provide a theory of physical corpuscles with chemical properties is a bit arbitrary and rather unfair. Boyle borrowed heavily from Sennert, who might plausibly receive credit as the first to combine the Aristotelian natural minima with atomism (Newman 1996; van Melsen 1952). However, since Boyle is usually fingered as providing the first 'mechanical' theory of chemical corpuscles [see e.g. (Levere 2001, 14-5)], I have deferred to precedent for ease of exposition.

region occupied by the body during the interval of interest. Boyle's particles thus satisfy a locality condition that is actually much stronger than MA(iii).

Not much of empirical value derived from Boyle's account. It allowed him to offer plausible qualitative explanations but little in the way of quantitative prediction. A genuinely physical particle theory had to wait until the early nineteenth century when Dalton carried through the synthesis of mechanical philosophy with the chemical theory of minima. Dalton embraces the hypothesis that "...all bodies of sensible magnitude, whether liquid or solid, are constituted of a vast number of extremely small particles, or atoms of matter bound together by a force of attraction..." (Dalton 1808, 143) and thus offers an account consistent with MA(i). While broadly similar to the theory put forth by Boyle, Dalton endowed his particles with very specific and—at least indirectly—*measurable* properties of the sort necessary to do chemistry. In particular, he granted each a weight and put great emphasis on the determination of relative weights. His account thus comports with MA(ii). As with Boyle, Dalton conceived of the union of chemical atoms into molecules to involve simple juxtaposition—they bound together in space without any internal alteration (van Melsen 1952, 139). Likewise, physical changes in bulk matter such as change of phase or crystal formation were to be explained by the spatial rearrangement of atoms, with all interaction mediated by direct contact. Thus, Dalton's theory satisfies MA(iii) in the same way Boyle's did. This view of hard, space-filling atoms interacting via direct contact is carried over into the kinetic theory of gases in the mid-nineteenth century. Dalton's theory is therefore representative of a broad class of physico-chemical accounts one might call corpuscular atomism. The many versions of this corpuscular atomism are, like the accounts of Boyle and Dalton, straightforward instances of MA.

Before leaving the corpuscular view, it is worth emphasizing an aspect of Dalton's theory that persists into modern particle physics. This innovation has to do with the properties of individual atoms and molecules. Dalton argued for the claim that "...*the ultimate particles of all homogeneous bodies are perfectly alike in weight, figure, &c.*" (Dalton 1808, 143) (emphasis in the original). In other words, he argued that the atoms of what we would call the chemical elements must have exactly the same monadic properties. While we no longer view this as exactly correct (there are multiple isotopes of each chemical element), the supposition has carried over into particle physics where it is thought that there are fundamental particles (particles without further internal structure), these particles belong to a finite number of kinds, and every particle of a given kind bears exactly the same state-independent monadic properties. While enormously fruitful in both its chemical and physical contexts, this postulate is exactly what leads to PI and,

ultimately, to the undermining of particle ontologies. The synthesis of chemical minima and mechanical philosophy must ultimately be viewed as a failure.

An alternate conception of material particles developed in parallel beginning in the 18th century, and would ultimately outlive the corpuscular view. In the mid-18th century, Roger Boscovich realized that the corpuscularian view of atoms as hard, undeformable solids which make direct contact in collision leads to insurmountable difficulties with Newton's mechanics. If atoms are solid volumes, then "...in the collision of solid bodies, either there must be compenetration, or the Law of Continuity must be violated by a sudden change of velocity..." (Boscovich 1966, 122). In other words, if atoms are perfectly elastic bodies that fill space and interact by direct contact, then any collision must result in an instantaneous change in velocity, contrary to the differential equations of Newton's mechanics. Boscovich's solution was to reconceive atoms as point-masses surrounded by fields of force. As he puts it:

...matter is unchangeable, and consists of points that are perfectly simple, indivisible, of no extent, & separated from one another; that each of these points has a property of inertia, & in addition a mutual active force depending on the distance in such a way that, if the distance is given, both the magnitude & the direction of this force are given; but if the distance is diminished indefinitely, the force is repulsive, & in fact also increases indefinitely; whilst if the distance is increased, the force will be diminished, vanish, be changed to an attractive force that first of all increases, then decreases, vanishes, is again turned to a repulsive force, & so on many times over; until at greater distances it finally becomes an attractive force that decreases approximately in the inverse ratio of the squares of the distances.
(Boscovich 1966, 121)

If this description reminds the reader of the passage I quoted from Richard Feynman in Chapter 1, this should emphasize the extraordinary prescience of Boscovich. Note that, though it sounds inordinately complicated, the force he describes aligns well with what we understand as a typical interatomic potential (e.g. the Leonard-Jones potential), at least if we neglect the repeated reversal of force direction. Whatever the merits of Boscovich's detailed proposal, his basic view of particles as centers of force became an integral component of atomic theory following the discovery of the electron in the nineteenth century. I'll take Rutherford's model of the atom as the culmination of this classical line of thought. In this model, the fundamental particles are the positively charged and relatively massive nucleus and the negatively charged and relatively massless electrons. Each particle is thus a center of electromagnetic force, and this force

is what binds the electrons in orbits¹⁹ about the nucleus in order to make ‘atoms’, now understood as composite objects.

As with the corpuscular theory, this interpretation of the physics is also in line with MA. It is again the case that material objects are supposed to be divisible into discrete entities, in this case centers of force. Each of these entities possesses the state-independent properties of charge and mass, as well as state-dependent properties such as momentum. So MA(i) and (ii) are again satisfied. MA(iii) is also satisfied, though in this case we really must speak of approximate independence. This is because the fields of force associated with each particle extend indefinitely through space. Particles arbitrarily far from one another nonetheless exert a force and thus influence each other’s properties. However, the forces in the theory drop off with distance (typically as an inverse square). Thus, the properties of the particles in one region of space are almost but not quite what they would have been had the distant particles not existed. Because in most models of the theory particles are distributed through space with significant separation, each material body (each composite of particles) is approximately independent of most other objects (composed of distant particles). In this way MA(iii) obtains.

Of course, Rutherford’s model is not compatible with the laws of electrodynamics, as he was well aware. To pursue this difficulty and the ways in which it spurred the development of a new mechanics would take us too far astray. Skipping over intervening years of confusion and foment, we arrive at last at QM.²⁰

3.5 MA and QM

As long as we leave out PI, interpretations of QM are available that are compatible with MA. In particular, interpretations which attribute properties to particles can be constructed such that those properties satisfy the independence condition of MA(iii). To do so requires some severe departures from classical metaphysics, but this is unsurprising. For instance, if we adjust our notion of ‘bearing a property’ to mean something like possessing a definite probability or propensity for manifesting a property, then we can treat quantum particles as bearing a full set of properties at all times. That is, each particle acts as a hanger

¹⁹ To be historically accurate, this is not quite what Rutherford says. Though he drops hints in his seminal paper on atomic structure that he has in mind the so-called ‘planetary model’, he asks the reader to “[c]onsider an atom which contains a charge $\pm Ne$ at its centre surrounded by a sphere of electrification containing a charge $\mp Ne$ supposed uniformly distributed throughout a sphere of radius R ” (Rutherford 1966, 709).

²⁰ These three major particle interpretations—the corpuscular, the center-of-force, and the quantum—were also pointed out by Whyte (1961, 22-3).

for probability distributions over property values corresponding to every possible observable. This is in contrast to the classical view in which particles were hangers for exact property values. Having adopted this liberal notion of what it means to bear a property, we can then ask whether quantum particles generally exhibit the sort of independence with respect to their properties that MA(iii) requires.²¹

Without PI, this is in fact the case. Non-interacting particles can be in ‘product states’ in which their properties are entirely independent.²² In general, even weakly-interacting particles in separable states will be approximately independent of one another with respect to the distributions over property values they bear. For instance, it is possible to have a particle with a high probability of being found on one side of my desk and another with a similar distribution centered on a point on the other side of my desk and to assign to each particle almost any distribution over its remaining properties that we could have assigned had the other particle not existed. In this way, MA(iii) can be satisfied and quantum particles viewed as the ontologically independent units that account for EDiv. Note that, unlike in the classical case, the independence of particles need not be a function of spatial separation. For instance, we might instead consider two weakly interacting particles that have very definite momenta rather than position. Such particles would be as non-localized as can be—their respective probability distributions over position overlap significantly. There is no sense in which we could say that these particles are spatially separated. Nonetheless, we could assign any distribution over, say, spin in the z-direction to one irrespective of what distribution we give the other. This would mean that it is possible to measure any of a wide range of property values for one particle irrespective of what was measured for the other. In this sense, the properties of one can be independent of the properties of the other despite spatial contiguity.

Of course, for MA(iii) to obtain, it must be the case that such independence holds under most physically possible conditions. It will fail when the particles interact strongly or when states are entangled (as is generally the result after strong interaction). Without a precise way to quantify what fraction of quantum models feature strong interactions or entanglement amongst all or most particles

²¹ An account of ‘unsharp quantum properties’ very similar to notion of properties as distributions over values suggested here is developed in (Busch and Jaeger 2010).

²² Let H^1 be the Hilbert space of states for each particle taken individually. Then the states of the two-particle system live in the outer product space $H^1 \otimes H^1$. If the states $|\psi_i\rangle$ comprise a basis for H^1 , then any state in the outer product space can be expressed in the form $\sum_{ij} c_{ij} |\psi_i\rangle |\psi_j\rangle$. With no condition of permutation invariance imposed, possible states include ‘product states’ of the form $|\psi_i\rangle |\psi_j\rangle$. The joint distributions over property values for two particles in a product state can be factored into independent distributions for each particle separately.

in the universe, I cannot make an airtight case. Nonetheless, it is intuitively plausible that in most models, most particles are independent of one another in the requisite way and MA(iii) is satisfied.

3.6 Previous arguments against particles in QM

I have given only a very coarse sketch of the sort of interpretation of QM that one would have to adopt in order to satisfy MA. It is not my intention to pursue such a view in detail, especially since, as I'll show in Chapter 6, PI precludes any such interpretation. The exercise would be moot. Rather, the point I wish to make is that there is nothing inherent in the structure of QM without the logically independent postulate of PI that precludes a particle interpretation. This claim is perhaps surprising given the number of authors who have already suggested that QM is incompatible with a particle interpretation. But as we'll see below these arguments rely on PI, posit much stronger accounts of 'particle' than MA, or both.

There are myriad arguments in the literature of physics and philosophy in support of the 'non-individuality' of putative quantum particles. The notion of 'non-individuality' is cashed out in nearly as many ways as there are distinct arguments for it. However, we need only concern ourselves with those notions of individuality that must obtain if the entities in an interpretation are to be deemed particles in the sense enshrined in MA. As I mentioned in Sections 3.2 and 3.3, MA assumes that, at any given time, there is a determinate fact of the matter which particle is which and what properties it bears. It assumes that particles stand in identity relations, both synchronic (one can specify which bears which property) and diachronic (a particle at time t is identical with one and only one particle at later time t'). Thus, the notions of 'non-individuality' with which we will be concerned are those which deny that putative quantum particles stand in any such relations.

One argument for the non-individuality of quantum particles rests on the assumption that individuality is determined by spacetime trajectory. In this view, objects are individuals just if they are associated with a unique and continuous trajectory through space and time. This view was most notably advanced by Erwin Schrödinger [see e.g. (Bitbol 2007; Schrödinger 1998)], and is referred to as 'Space-Time Individuality' or STI by French and Krause (2006). Under STI facts about identity reduce to facts about trajectory. If two events involve the same individual, they lie on the same spacetime trajectory. As Schrödinger and others have emphasized, it is in general not possible to attribute a definite spacetime trajectory to putative quantum particles. Thus, since individuality depends on trajectory, would-be quantum particles are not individuals and so cannot be par-

ticles in the sense of MA.²³ As Schrödinger himself put it: “...giving up the path means giving up the particle.”²⁴ While the argument appears valid, it rests on a dubious premise, namely that spacetime trajectories are the only way in which particles might come to possess definite identities. At the very least, this premise is logically much stronger than MA requires. Thus, the argument from STI does not preclude particle interpretations in line with MA.

The remaining arguments for non-individuality (in the sense of non-identity) each invoke PI in one way or another. Therefore, they cannot contradict the claim that QM and MA are compatible in the absence of PI. However, each of the remaining accounts also invokes one or more assumptions of greater logical strength than MA. Making these assumptions explicit serves to emphasize the much greater generality of the result presented here. We can begin with another time-worn argument, this time drawing on features of statistical mechanics [see e.g. (French 2006; French and Krause 2006)]. Very roughly, in classical statistical mechanics, there are four ways in which two sets of properties can be distributed over two particles. That is, there are four distinct states associated with two particles and two different property sets. Call the particles p_1 and p_2 and the property sets M_1 and M_2 . We'll denote the proposition that p_i possesses properties M_j by $M_j(p_i)$. The four possibilities are thus $M_1(p_1) \wedge M_1(p_2)$, $M_1(p_1) \wedge M_2(p_2)$, $M_2(p_1) \wedge M_1(p_2)$, and $M_2(p_1) \wedge M_2(p_2)$. It is taken as a premise that all distinct (if not qualitatively distinguishable) configurations of particles and properties are equiprobable.²⁵ From this specification of the state space and the assumption of equiprobability, the usual Maxwell-Boltzmann statistics follows along with the classical statistical mechanical reduction of thermodynamics. Quantum statistical mechanics—whose departure from classical statistical mechanics is a consequence of PI—poses a problem. If the equiprobability assumption holds, then the observed statistics of putative particles of the same kind entails that we have miscounted distinct states. As we'll see in Chapter 5, there are two classes of particle—bosons and fermions—distinguished by the statistics they obey. For the bosons, $M_1(p_1) \wedge M_2(p_2)$ and $M_2(p_1) \wedge M_1(p_2)$ combined receive the same prob-

²³ A closely related objection to particle ontologies is made in the context of QFT, for which no spatially localized states of a single quantum are possible, at least none with plausible transformation properties under a Lorentz boost. The implicit idea is that particles are necessarily individuated by spacetime location and thus the impossibility of localizing a particle state precludes a particle interpretation of QFT. See (Teller 1995), Chapter 4 for an introductory overview.

²⁴ This quote is from a letter to Henry Margenau dated April 12, 1955. The letter appears in the Archive for History of Quantum Physics (Microfilm 37, Section 9) at the University of Pittsburgh. It is reproduced in (Bitbol 2007, 86).

²⁵ Hans Reichenbach gives a version of the argument recounted here in (Reichenbach 1999). In particular, he makes this assumption explicit on p232. To be fair, Reichenbach—unlike many later authors to advance this argument—motivates the equiprobability claim by drawing out the sorts of ‘causal anomalies’ that alternative approaches would force us to adopt.

ability as either $M_1(p_1) \wedge M_1(p_2)$ or $M_2(p_1) \wedge M_2(p_2)$, and for the fermions they receive a combined probability of 1. This suggests that for either class $M_1(p_1) \wedge M_2(p_2)$ and $M_2(p_1) \wedge M_1(p_2)$ are not in fact distinct states, and thus the putative particles lack identity—there is no fact of the matter as to which possesses which property.

Even if we grant that distinct states are equiprobable—an implausibly strong a priori assumption—one must take care to clarify whether equiprobability pertains to metaphysically possible states or to contingently accessible states. In the strong reading implicitly adopted in the argument for non-individuality, it is assumed that every distinct *possible* state is equiprobable. However, if one assumes that equiprobability applies merely to the physically *accessible* states, then the above is not an argument against the possibility of distinct permuted states. This is because the same statistical weights can be accounted for by denying the accessibility—though not the metaphysical possibility—of distinct permuted states. To explain this rebuttal in detail requires more of the apparatus of QM than I have yet introduced (though see Chapter 5 below). Qualitatively speaking, it can be shown that if a particle begins in any of a particular subset of states corresponding to either bosons or fermions then the dynamics of quantum mechanics guarantees that the future states of the particle will be in the same subset—bosons cannot become fermions or vice versa and neither can come to occupy states outside their prescribed subset. In short, QM guarantees that certain states remain inaccessible to particles that happen to begin in bosonic or fermionic states. Assuming that as a contingent fact all particles began in one of these states, this would account for the observed statistics if accessible states are equiprobable. Thus, in order to conclude that quantum particles really are non-individuals, one must adopt the strong version of equiprobability and insist that we can infer the set of possible states from the observed statistics.²⁶ Such a view takes us beyond MA.

Another argument for non-individuality that invokes PI rests upon the ‘bundle-theory’ account of individuality. In versions of this account, an individual is uniquely picked out by the properties it bears. This notion is captured by Leibniz’s Principle of the Identity of Indiscernibles (PII), which can be stated in a second order language as follows:

$$\text{(PII): } \forall_x \forall_y (\forall_F (F(x) \leftrightarrow F(y)) \rightarrow x = y)$$

²⁶ See (French and Krause 2006), Chapter 4 or (French 2006).

Here, x and y refer to individuals and the predicates F refer to properties of the individuals.²⁷ The argument begins by noting that particles of the same kind subject to PI must have all of their properties in common. Nonetheless, these putative particles are not identical, since we can experimentally count them. This leaves us with two options. If the predicates F refer only to physical properties, then either Leibniz's PII is false and particles are individuated by some non-physical property—what Heinz Post (1963) calls a Transcendental Individuality—or PII is simply irrelevant because quantum particles are not the sorts of things that stand in relations of identity with one another. The former is argued to be distasteful on a variety of philosophical grounds (e.g. even a weak empiricism), and so the latter is favored. Again, this argument rests on PI and so does not contradict the claim that, absent permutation invariance, QM can be given a particle reading. Disregarding this point it is still once again the case that the argument relies on much stronger metaphysical suppositions than MA requires. For instance, in order to argue that quantum particles are non-individuals in this way, we must first accept that genuine individuals obey PII and that there are no non-physical properties such as Transcendental Individuality requires. While perhaps plausible, these claims go beyond what the epistemic facts justify, and leave open the possibility that a more liberal particle ontology in line with MA is still consistent with the physics.

Before considering other arguments against particle interpretations of QM, I would suggest that there is a general difficulty with all arguments for non-individuality. Proponents of such arguments want to claim that, while quantum particles are non-individuals in the sense that they have no identities, there are nonetheless determinate numbers of them in various physical systems. This position is incoherent. As Lavine (1991, 260) puts it, one cannot claim that there are two distinct photons in a box because "...after all, 'distinct' means 'having different identities'..." Thus, statements of cardinality for non-individuals cannot be interpreted via the usual semantics. For a more thorough development of this line of criticism, see (Jantzen 2010).

As I suggested above, there are other arguments against particle interpretations of QM that do not invoke notions of individuality or depend critically on PI. Massimi for instance, argues that no definite properties are attributable to individual particles in joint states that are not 'product states'. For this reason, MA(ii) is false (the particles cannot be said to bear any state-independent properties). It is further the case that, even if we attribute properties to individual

²⁷ There are multiple readings of PII depending on how one restricts the range of F . I won't lay out the possibilities here, but the interested reader should consult (Caulton and Butterfield 2008; French and Redhead 1988; Saunders 2003).

particles when the system is not in a product state, then the state of the system fails to supervene on the states of the individual particles—complete information about the properties of the particles fails to fix the state of the system. For either or both of these reasons, QM cannot be interpreted in terms of particles. This argument, however, rests on strong claims about what it means to bear a property and the relation between particle states and states of particle aggregates. In particular, Massimi assumes the following:

Given any composite system, the quantum states of the subsystems are ontologically separate [that is, the particles can be ascribed properties individually] iff

- 1) each subsystem has definite (though possibly unknown) values for a complete set of compatible observables pertaining to that subsystem alone
- 2) the afore-defined ontologically separate states of the subsystems determine wholly their joint state.

(Massimi 2001, 320-1)

Massimi's first criterion depends on a rather conservative notion of what it is for a particle to bear properties. According to a standard interpretation of QM, "...a system in the state W has a value for the observable F if and only if W assigns probability 1 to one of the possible values of F ..." (Dickson 2007, 285). This principle has come to be known as the 'eigenstate-eigenvalue link'. Even granting the eigenstate-eigenvalue link, Massimi's argument only goes through if we further assert that the only properties a particle can be said to bear are those which it bears with certainty (with probability 1). As I suggested above, there is no a priori reason we need to make this assumption—we can treat distributions over measurable values as perfectly sensible properties of quantum objects. To claim otherwise is to assert a much stronger particle ontology than MA. Furthermore, Massimi's second criteria amounts to the assertion that an object cannot be said to have ontologically independent parts with properties of their own unless the properties of the composite object supervene on those of its parts. Again, while this may be appealing it involves a stronger metaphysical commitment—in this case to a denial of emergent properties—than MA requires. If Massimi's strong assumptions are dropped, both of which go beyond MA, then her rejection of quantum particles does not go through.

There is one more argument to consider against a particle interpretation of QM. This argument actually applies only to relativistic quantum field theory (QFT)²⁸,

²⁸ Actually, the problem of unitarily inequivalent representations, of which this argument is an instance, pertains to any theory with infinite degrees of freedom, including quantum sta-

and thus is tangential to the claim that non-relativistic QM can be given an interpretation in terms of MA. Nonetheless, it is worth having the argument on the table. To begin with, we note that in certain cases, observers in different states of motion will define certain observables differently, in particular the number operator N which gives the total number of particles occupying a particular state. A good example is the Unruh effect. Here, an accelerating observer will, under certain circumstances, see thermal radiation in her rest frame while an inertial observer sees none. Whether one observes ‘particles’ in the world or not thus depends on one’s reference frame. This would seem to cast doubt on the existence of particles as objective entities (like observers) which persist in all reference frames and bear properties. I will not digress into the technical details here. Suffice it to say that, this is a serious objection to particles but one that plays equally well against other realist interpretations of QFT, specifically certain otherwise appealing field ontologies (Baker 2008). I will not attempt to offer a resolution, only suggest that effects of this sort cause trouble for particle interpretations in ways distinct from PI.

3.7 Conclusion

To sum up, the notion of ‘particle’ as motivated by the argument of Chapter 2 is captured by MA, the logically weakest set of propositions constituting a particle ontology. MA is consistent with all of the major scientific particle interpretations, and even QM can be given an interpretation compatible with MA so long as we do not insist on PI. The latter fact is surprising on the face of it. But previous arguments against particle interpretations of QM either invoke PI or put much stronger conditions on particles, properties, or physically possible states than MA does. The question pursued in the remainder of this essay is whether the *weakest* particle ontology that can support EDiv can be maintained for quantum theories or for any successor to QM that respects PI.

tistical mechanics. See (Baker 2008) for an overview, as well as for a concise and much more rigorous statement of the argument considered here.

Chapter 4 Connecting metaphysics with physics

4.1 Overview

In Chapter 2 I distinguished a scientific theory from its models, and the interpretations of these models. In this chapter, my aim is to develop this taxonomy in greater detail for theories with mathematical formulations, such as those of mathematical physics. With this taxonomy, we can then state a formal means of connecting metaphysical interpretations to theories, and for extracting constraints on theories given a particular interpretation. Establishing this formal connection will let us develop with precision one necessary condition that any theory must satisfy in order to be compatible with MA.

4.2 A taxonomy of theories and interpretations

In the context of mathematical physics, a *theory* consists of two parts:

- (i) the specification of a mathematical space such as a phase space or Hilbert space;
- (ii) a system of equations (or ‘laws’) that, when supplemented with boundary conditions pick out a subset of the mathematical space of (i).

To make this more concrete, consider the classical theory of Hamiltonian dynamics. According to my taxonomy, this is actually a large family of theories. Each theory in this family is characterized in part by a phase space (typically $6n$ -dimensional for n particles) pertaining to the physical system of interest. In addition, each theory of Hamiltonian dynamics is equipped with a set of coupled differential equations (Hamilton’s Equations of Motion) that together with a set of boundary conditions (and modulo some technical caveats) determine a unique trajectory in phase space. While in a vague sense Hamilton’s Equations have the same ‘form’ for each theory (for each combination of phase space and equations of motion), they differ both with respect to the specific Hamiltonian used and in the number of conjugate variables.

A *model* of a theory is just a mathematical structure belonging to the mathematical space of the theory and compatible with both the laws of the theory and a set of boundary conditions. A trajectory in phase space is a model of Hamiltonian mechanics. Similarly, a trajectory through rays in a Hilbert space is a model

of time-dependent QM. What I'm calling models of a theory are what physicists often call 'solutions'.²⁹ Aside from a mathematical space and the equations of motion, the determination of a model requires boundary conditions. In Hamiltonian Mechanics, a boundary condition is typically the specification of a single point in phase space taken to represent the initial positions and momenta of the particles. We could think of such a boundary condition as picking out the set of all trajectories in the mathematical space of the theory (phase space) that pass through the specified point. This provides an intuitive way of understanding the role of boundary conditions: *boundary conditions* in general are specifications of a subset of the mathematical space of a theory which is further refined by the laws of that theory. For a fixed theory, the laws of the theory amount to functions from the space of boundary conditions to the space of models. Thus, the set of all physically possible boundary conditions determines the set of all physically possible models.

Interpretations of a theory are sets of property attributions over a set of objects.³⁰ They are supposed to be descriptions of the way the world is, or at least the way part of the world is. Interpretations need not be expressible in first or second order formal languages, but I will treat an interpretation as a collection of statements about the contents of the world. For instance, the interpretation of a trajectory in phase space typically amounts to the assertion of successive locations and momenta of a moving body such as a pendulum bob—it is a collection of claims like “there exists a massive bob with this position and this momentum at time t ,” and so on. The various scientific accounts of particles and their properties discussed in the preceding chapters are also examples of interpretations. To simplify later discussion, it will help to introduce an additional piece of terminology. I'll call that part of an interpretation which describes the properties associated with a single particle a *specification*.

Interpretations are connected to theories by a function on the space of models. This function need not be injective—multiple models might refer to the same state of the world. But for every model there is exactly one description of the way the world is according to the model. To continue our example from Hamiltonian dynamics, each trajectory in phase space—under a typical interpretive function—corresponds to the attribution of a position and momentum to a par-

²⁹ My use of the term 'model' in this case is closely related to the model-theoretic notion from logic and aligns with (at least some) of the ways in which proponents of the 'semantic view' of scientific theories use the term. See, e.g. (Suppe 1977, 1988; van Fraassen 1980).

³⁰ This stipulation rules out many ontologies such as process and event (Simons 2002). We could weaken the notion of interpretation to embrace these possibilities, but it is unnecessary for our purposes since particle ontologies are of the sort captured by this restricted notion of interpretation. Arguably, the vast majority of scientific interpretations fall within this class, no matter how vehemently some philosophers assert that substance-attribute ontologies are dead.

ticle at every moment of time over a given interval. Descriptions that differ with respect to where a particle is at a given time must correspond to distinct trajectories.

For any given theory there are indefinitely many possible interpretations—just as many as there are functions on the domain of models. Despite this freedom in selecting an interpretation, all theories come partially interpreted in the sense that each theory is accompanied by a mapping from its models to a limited set of claims about observable quantities. After all, if a theory didn't tell us what measurable quantities to anticipate it would have no empirical purchase on the world. For instance, any theory of Hamiltonian dynamics will have to come with an interpretation that lets us extract anticipated forces or momentum transfers from its models—I have to at least interpret the phase space trajectory as indicating a certain momentum at a certain time and place so that I can anticipate what will be measured when I stick an instrument there. I refer to these functions from models to observable facts as 'partial' interpretations because it is possible to elaborate any such interpretation—we can create a 'fuller' interpretation by positing a more elaborate state of the world involving unobserved entities, properties, and processes as long as the posited state of the world contains or entails the empirical claims of our partial interpretation. With a partial interpretation given, the space of viable elaborations—the space of full interpretations—is somewhat restricted. In fact, we'll see how a partial interpretation constrains the choice of full interpretation in Chapter 8 below. Despite this constraint, however, full interpretations are still underdetermined by physical theory. For any given physical theory and corresponding partial interpretation, infinitely many full interpretations are possible.

This is not to say, however, that full interpretations are inert or irrelevant. In fact, they serve in two significant capacities in scientific practice: (i) as guides to theory development and replacement and (ii) as guides for building idealizations and simplified models. Providing a detailed account of the first function would take us too far astray, but an example of what I have in mind is provided by the development of Kepler's astronomy. In formulating a theory that computes celestial positions using ellipses for each planet, Kepler was guided by a particular interpretation of Copernicus' work, namely that there really are massive bodies including the Earth that revolve around a central Sun. Practically speaking, it helps to have to have some idea of what your old theory is about in order to develop a new one. Full interpretations do not constitute necessary conditions for scientific change—to hit upon the use of elliptical orbits, it is not necessary to believe that the earth is a planet and that planets actually revolve about the sun. But when confronted with the enormous variety of ways one might adjust

the mathematical system of Copernicus in hopes of better capturing the phenomena, a full interpretation can guide one's choices in a fruitful way.

The second function of full interpretations is much more mundane and easier to illustrate. Typically, physical theories present a set of equations that are too difficult to solve for many or most practical applications. Some sort of approximation is necessary. Sometimes, suitable approximations can be derived from mathematical considerations alone given knowledge that particular (measurable) quantities are small relative to others. This is the case, for instance, when we make the 'small angle' approximation for pendulums. However, most of the time much more elaborate manipulations or modifications of the theory are required, and the partially interpreted theory offers no justification for any such manipulation. Consider, for example, quantum chemistry. In order to determine the shape, binding energy, etc. of a molecule, one must solve the time-dependent Schrödinger equation with a complex Hamiltonian. For even the simplest molecules, this is an intractable problem. A standard approach is to make the 'adiabatic assumption' that the light electrons reconfigure themselves much faster than the heavy nuclei. This motivates solving the electronic problem for static nuclei in order to determine electronic potential energy and configuration as a function of nuclear separation. This 'effective potential' is then used in turn to solve the Schrödinger equation for the nuclei. The point is that the adiabatic assumption is only justified if we think that terms in the Hamiltonian actually refer to bodies that move more or less independent of the nuclei which are in turn massive charged bodies. That is, the adiabatic assumption only makes sense given a *full* interpretation of the quantum mechanical theory in terms of massive, charged particles of the sort MA requires. The full interpretation guides development of tractable approximations to a given physical theory.

4.3 Property independence as a formal constraint on theories

In order to assess the compatibility of PI and particle interpretations, we need a method for ascertaining whether a given theory can sustain full interpretations that satisfy MA. Working backwards, such a method should let us determine what conditions MA places on a theory. In other words, it will let us determine what features a theory must have if its models are to be connected via some function to an interpretation in line with MA. As a first step in connecting MA with specific theories, we will take interpretations to be descriptions of physically possible worlds—to say that a state of affairs is physically possible is to say there is an interpretation of a model of the relevant physical theory that describes that state of affairs. Next, the notion of independence appearing in MA(iii) can be formulated in terms of physical possibility. To say that the monadic properties of one object are entirely independent of the properties of

another is to say that it is physically possible for each object at any one time to possess any of the monadic properties it could have possessed had the other object not existed. By extension, to say that the properties of one object are approximately independent of those of another is to say that it is physically possible at any one time for each object to possess approximately any of the properties it could have possessed had the other not existed. Given that physically possible worlds are described by interpretations, we can restate the condition of property independence this way:

Approximate Property³¹ Independence:

The properties of object A are independent of the properties of object B according to physical theory T just if there exists a non-trivial set of interpretations S corresponding to a set of models of T such that every interpretation $Int \in S$ contains specifications of objects A and B, and there exists a pair of non-trivial sets of interpretations S_1 and S_2 corresponding to sets of models of T' (the one-particle version of T) with the following features: (i) every $Int_1 \in S_1$ contains specifications of all the same objects as each interpretation in S except for object B, (ii) every $Int_2 \in S_2$ contains specifications of all the same objects as each interpretation in S except for object A, (iii) for every pair of interpretations $Int_1 \in S_1$ and $Int_2 \in S_2$ there exists an $Int \in S$ such that the specification of A in Int is approximately the same as it is in Int_1 and the specification of B in Int is approximately the same as it is in Int_2 .

That is admittedly a mouthful, but it is really just the claim that the properties of A are (approximately) independent of those of B if the theory admits models containing both objects that look like combinations of models that contain each object by itself. For example, classical mechanics admits sets of models representing two gas molecules with independent properties. If the molecules interact only very weakly at macroscopic distances (as they do for realistic intermolecular potentials) then there are two-molecule models of the theory that—when interpreted—give descriptions nearly identical to the union of two one-molecule descriptions. Stated in the other direction, it is easy to write down a model interpreted in terms of single molecule of nitrogen flitting about to my left, and another model interpreted as describing a molecule in a volume to my right. This pair of descriptions is (almost) what we would have extracted from a model of the theory that incorporated both molecules. Since this holds for lots of configurations of each single molecule—we can make the left molecule move upward, and the right molecule move downward, we can put them at different

³¹ All properties referred to in this definition are taken to be monadic.

positions, etc. and still find a joint model with approximately the same particle specifications—we can say that the properties of the molecules are independent. Of course, independence will not obtain for all conceivable sets of models of a theory—it is always relative to a set of boundary conditions that demarcate a set of physically possible conditions.

There remain two vague notions in the definition of property independence given above, that of “approximate” property equivalence and that of a “non-trivial” set of properties. Both can be explicated with the same device—a measure of distance or difference between any two specifications. We can do this formally by imposing a metric on the space of all one-particle specifications. A metric on a set A is a function $d: A \times A \rightarrow \mathfrak{R}$ (where \mathfrak{R} is the set of reals) with the following properties:

1. (positive definiteness) $d(x,y) \geq 0$ and $d(x,y) = 0$ iff $x = y$
2. (symmetry) $d(x,y) = d(y,x)$
3. (triangle inequality) $d(x,z) \leq d(x,y) + d(y,z)$

Metric functions are generalizations of the distance function familiar from coordinate geometry, and are natural choices for measuring the distance between specifications. While we would like whatever function we use to establish the degree of difference between object specifications to have the properties of a metric, it turns out that we needn’t be any more specific than that. Whatever metric we choose, the incompatibility result reported in Chapter 6 obtains. Thus, I’ll leave the specification of a particular metric to the reader’s discretion.

Once we have a metric over object specifications, we can define the approximate equivalence of specifications by appealing to a threshold separation, ϵ . That is, we can stipulate some value $\epsilon \geq 0$ such that specifications that are within ϵ of one another according to our metric are considered approximately the same. Obviously there is some arbitrariness in selecting such a threshold. There are some plausible guides we might use. Since we are ultimately concerned with accounting for an epistemic fact (namely EDiv) we might appeal to the degree of difference between two specifications which it is actually possible to measure³². But again, it makes no difference what we choose. No matter what we take the threshold to be, PI will be incompatible with MA. So again, I will not attempt to fix a determinate value.

Once we have a metric and a threshold for establishing approximate specification equivalence, we can also spell out the notion of a non-trivial set of specifications: to say that a set of specifications is non-trivial is just to say that its di-

³² Only part of a specification is observable.

ameter (the greatest distance between two elements in the set according to the metric) is large compared to ϵ . Since any two specifications within ϵ of one another are approximately the same, any set of specifications that is not large compared to ϵ would contain ‘approximately’ one state, and in this sense would be trivial. One can think of the threshold ϵ as setting a natural length scale on the space of specifications. While the value of ϵ is arbitrary, the structure of the set of specifications for this length scale is not.

With the tools laid out above, we can finally formulate a single, rigorous condition of property independence that any theory must satisfy to be compatible with MA. As stated, MA(iii) asserts that property independence obtains (under most conditions) for *any* set of particles S . For simplicity, I will focus on the simplest case in which S contains just two particles. In that case, the condition MA(iii) for particle interpretations can be put in terms of models of the theory as follows:

Two-particle independence:

For a theory T to be compatible with MA (iii), there must exist three sets of models of T —call them α , β , and γ —with the following properties:³³

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

³³ To be more careful, the models in question are of T (models of which are interpreted as describing two particles) and of T' (models of which are interpreted as describing single particles). T and T' involve laws of the same form (e.g. the Schrodinger equation) but different mathematical spaces. I am assuming that there is a determinate method for picking T' given T .

approximately attributes each of these specifications to one particle in the interpretation).³⁴

Note that two-particle independence is a necessary condition for a theory to satisfy MA, but is weaker than MA(iii) because it insists only on the approximate pairwise independence of particles. In Chapter 6, we'll see that even this weak condition cannot be satisfied for any theory that incorporates PI.

To this point, I have pursued a strategy of refinement. By examining the principal argument for realism, one can discern the outlines of a broad class of interpretations with appeal for the realist. In Chapter 3, I refined this class to pin down the logically weakest set of propositions that pick out the particle ontologies. In this chapter, I have shown how to convert interpretations to constraints on the models of a mathematical theory. Specifically, I extracted a formal condition for two-particle independence as a necessary though insufficient condition for a theory to be compatible with MA. In the next chapter, I work in the other direction to generalize the postulate of PI, converting it from an idiosyncratic component of QM to a theory-independent constraint on particle ontologies. Only with such a generalized postulate is it possible to ask whether PI always precludes a particle interpretation.

³⁴ S_α and S_β are subsets of a common set of all possible single-particle specifications. From the interpretation of each model in γ it is possible to extract an ordered pair of specifications, the elements of which are members of the same common set of single-particle specifications.

Chapter 5 Permutation Invariance

5.1 Overview

In Chapter 3, we saw that QM can be given a particle interpretation in the absence of PI. However, it is my central contention that, if PI is imposed on QM or any other theory in which it can be formulated, the resulting theory is incompatible with MA. Accepting PI means rejecting particles. I have dealt with the notion of particle in depth (see Chapter 3) but have yet to state PI very clearly. The present chapter is dedicated to developing an account of that principle and defending its empirical necessity. I'll begin by explicating PI in the context of QM where it began, and sketch out a number of empirical consequences of the principle that are well-confirmed experimentally. I'll then abstract away the details of QM to yield a theory-independent statement of PI using the terminology developed in the preceding chapter.

5.2 Formulation and Consequences in QM

The canonical statement of PI in the literature of QM appears in (Messiah and Greenberg 1964, 250):³⁵

(PI_{QM1}) “Dynamical states represented by vectors which differ only by a permutation of [particles of the same type] cannot be distinguished by any observation at any instant of time.”

The “vectors” to which the authors refer are vectors in the Hilbert space used to represent a quantum system (more below). By itself, PI_{QM1} only tells us something about the connection between permutation operations and observables. Which sets of vectors represent distinct states is fixed by two additional postulates (Hartle and Taylor 1969, 2045):

- (1) Two vectors representing the same state must give the same expectation value with respect to all observables; and
- (2) Two vectors representing distinct states must give different expectation values for at least one observable.

³⁵ Greenberg and Messiah were the first to explicitly endorse the principle I've been calling PI instead of the stronger Symmetrization Postulate (SP) that had been used to resolve the problem of ‘exchange degeneracy’. For more on the historical development of SP and PI, see (French and Krause 2006). For a standard treatment of the problem of exchange degeneracy (which actually turns on an implicit invocation of PI) see (Messiah 1999, 582-5).

From these two postulates, it immediately follows that two vectors represent the same state if and only if they give the same expectation values for all observables. Together with PI_{QM} , this further entails that : (i) any permutation of the parts of a vector that correspond to particles must yield a vector representing the same state, and (ii) observables can only be represented by operators that commute with the operators representing permutation.

What does this mean for the representation of states by vectors in a Hilbert space?³⁶ Suppose that $H^{(1)}$ is the Hilbert space corresponding to the states of a single particle. For n distinguishable particles (for which PI does not apply), the states of the system are given by the vectors of the outer product Hilbert space $H^{(n)} = H^{(1)} \otimes H^{(1)} \otimes \cdots \otimes H^{(1)}$ (the outer product is taken n times). We can write a generic vector in this outer product space in the form $\sum_{\alpha_1, \alpha_2, \dots, \alpha_n} c_{\alpha_1, \alpha_2, \dots, \alpha_n} |u_{\alpha_1}\rangle |u_{\alpha_2}\rangle \cdots |u_{\alpha_n}\rangle$.

The *permutation operators* [see, e.g. (Cohen-Tannoudji, Diu, and Lalo  e 1977), XIV.B] are operators $\hat{P}_{i_1, i_2, \dots, i_n}$ on the outer product space defined by the action:

$$\begin{aligned} \hat{P}_{i_1, i_2, \dots, i_n} & \sum_{\alpha_1, \alpha_2, \dots, \alpha_n} c_{\alpha_1, \alpha_2, \dots, \alpha_n} |u_{\alpha_1}\rangle |u_{\alpha_2}\rangle \cdots |u_{\alpha_n}\rangle \\ & = \sum_{\alpha_1, \alpha_2, \dots, \alpha_n} c_{\alpha_1, \alpha_2, \dots, \alpha_n} |u_{\alpha_{i_1}}\rangle |u_{\alpha_{i_2}}\rangle \cdots |u_{\alpha_{i_n}}\rangle \end{aligned} \quad (5.2.1)$$

Qualitatively, a permutation operator rearranges the kets in each outer product vector $|u_{\alpha_1}\rangle |u_{\alpha_2}\rangle \cdots |u_{\alpha_n}\rangle$. Each permutation operator is unitary and possesses an inverse that is also a permutation operator. They do not commute, and are not generally self-adjoint. The permutation operators form a group. In fact, they constitute a representation of the symmetric group S_n . The outer product space $H^{(n)}$ decomposes into a direct sum of subspaces, each corresponding to an irreducible unitary representation³⁷—or ‘irrep’ for short—of the symmetric group. Any permutation operator acting on any vector in one of those subspaces results in a vector in the same subspace. PI along with conditions (1) and (2) imply that each of these subspaces—the irreps of S_n —represents a single physical state. Put another way, each distinct state of a system of n particles of the same

³⁶ The formalism sketched in this section was first developed in (Messiah and Greenberg 1964; Hartle and Taylor 1969).

³⁷ Informally, a unitary representation of a group G on a Hilbert space H is a mapping f from G to the unitary operators on H that preserves the group structure of G . That is, $f(u * v) = f(u)f(v)$ for any $u, v \in G$. If a subspace $V \subset H$ is invariant under the action of operators representing G , then a mapping from G to the unitary operators restricted to V is a subrepresentation of the group. An irreducible representation of a group is a representation with no trivial subrepresentations (there are no subspaces invariant under the unitary operators representing the group). For an overview, see (Fulton and Harris 1991).

type must be represented by one and only one of the distinct subspaces of $H^{(n)}$ corresponding to an irrep of the symmetric group. Operating on one of the vectors in such a subspace with a permutation operator yields a (possibly distinct) vector in the same subspace.³⁸

As I said, every outer product space $H^{(n)}$ decomposes into a direct sum of irreps of the symmetric group S_n . Some of these are one-dimensional, and consist of either completely symmetric or completely anti-symmetric vectors. These are vectors that are either invariant or change sign under the action of a permutation operator. The former describe the states of so-called ‘bosons’ (e.g. photons) while the latter describe the ‘fermions’ (e.g. electrons). For $n = 2$, these one-dimensional representations exhaust the possibilities. When $n \geq 3$, there are also higher-dimensional subspaces in the decomposition of $H^{(n)}$. In principle, these higher-dimensional subspaces could describe the states of one or another type of particle. The states of such particles—called ‘paraparticles’—would exhibit more complex symmetries and thus, for instance, more complicated statistics when considered in large ensembles. However, only bosons and fermions are believed to exist.

To give a more concrete illustration of the restrictions on states when PI is imposed, consider the case of two particles. If each particle is represented individually by vectors in the Hilbert space $H^{(1)}$, then states of the joint system reside in the outer product space $H^{(2)} = H^{(1)} \otimes H^{(1)}$. There are only two permutation operators on $H^{(2)}$, the identity operator and \hat{P}_{21} and there are only one-dimensional representations of the symmetric group S_2 . These come in two kinds. The first belong to the subspace we’ll call H_S which is spanned by the symmetric vectors of the form $|\psi_S\rangle = c_{ij} (|u_i\rangle|u_j\rangle + |u_j\rangle|u_i\rangle)$ for which $\hat{P}_{21}|\psi_S\rangle = |\psi_S\rangle$. The second sort of representations fall within a subspace H_A spanned by antisymmetric vectors of the form $|\psi_A\rangle = c_{ij} (|u_i\rangle|u_j\rangle - |u_j\rangle|u_i\rangle)$.³⁹ Note that $H^{(2)} = H_S \oplus H_A$. Thus, when PI is imposed in the case of two particles, only two sorts of symmetry under permutation are possible. If the particles are represented by vectors belonging to H_S we call them ‘bosons’ and their states are invariant under the action of the permutation operator. If the particles are represented by states belonging to H_A we call them ‘fermions’—their states change sign under action of the permutation operator.

³⁸ Note that PI restricts the class of observables to operators that commute with the permutation operators.

³⁹ Each vector in H_S or H_A when combined with the identity operator and \hat{P}_{21} constitutes an irrep of S_2 .

5.3 Empirical support for PI in QM

PI has important empirical consequences. It entails effects on the outcome of scattering experiments and the behavior of large aggregates of particles that diverge dramatically from what one would otherwise expect. It also explains the complex structure of atoms. These consequences are easiest to draw out in the case of two particles. By way of contrast, suppose a system of two distinguishable particles is represented by an asymmetric state $|\phi\rangle|\psi\rangle$ in $H^{(2)}$, where $|\phi\rangle$ and $|\psi\rangle$ are each vectors in $H^{(1)}$. Consider a single-particle property represented by the observable $\hat{B}^{(1)}$ on $H^{(1)}$. Suppose that $\hat{B}^{(1)}$ has a non-degenerate spectrum u_i with eigenvectors $|u_i\rangle$. We can then ask for the probability of measuring the joint two-particle system and finding the first particle to have property u_i and the second to have u_j . This is simply given by $\|\langle u_i|\langle u_j|\phi\rangle|\psi\rangle\|^2 = \|\langle u_i|\phi\rangle\|^2 \|\langle u_j|\psi\rangle\|^2$. Now, it may be the case that we do not care or cannot determine using some particular apparatus which of the distinguishable particles bears which property. Then we would want to know the probability that one of the two particles—it doesn't matter which—yields u_i when a measurement is made and the other manifests u_j . Since the states $|u_i\rangle|u_j\rangle$ are orthogonal, this disjunctive probability is given by:

$$\Pr(u_i, u_j) = \|\langle u_i|\phi\rangle\|^2 \|\langle u_j|\psi\rangle\|^2 + \|\langle u_j|\phi\rangle\|^2 \|\langle u_i|\psi\rangle\|^2 \quad (5.3.1)$$

Now suppose that the particles are of the same type such that PI applies. This changes the picture in two ways. First, the vector $|\phi\rangle|\psi\rangle$ cannot represent a state of the system. The particles must occupy either a symmetric (boson) or anti-symmetric (fermion) state of the form $\frac{1}{\sqrt{2}}(|\phi\rangle|\psi\rangle + \epsilon|\psi\rangle|\phi\rangle)$, with $\epsilon = +1$ or $\epsilon = -1$ respectively. Second, it makes no sense to ask about the probability that a specific particle manifests u_i and the other u_j —there is no corresponding observable. Put another way, when PI is imposed, the eigenvectors of observables must be (anti)symmetric. We can only meaningfully ask about the probability of measuring the particles in the state $\frac{1}{\sqrt{2}}(|u_i\rangle|u_j\rangle + \epsilon|u_j\rangle|u_i\rangle)$. This is generally interpreted as the probability of finding one particle (we cannot say which) to have u_i and the second to have u_j . With these adjustments made, we can compute the desired probability for particles of the same type when PI applies:

$$\Pr(u_i, u_j) = \|\langle u_i|\phi\rangle\|^2 \|\langle u_j|\psi\rangle\|^2 + \|\langle u_j|\phi\rangle\|^2 \|\langle u_i|\psi\rangle\|^2 + 2\epsilon \text{Re}(\langle u_i|\phi\rangle\langle u_j|\psi\rangle\langle \psi|u_i\rangle\langle \phi|u_j\rangle) \quad (5.3.2)$$

The third term in Equation (5.3.2) containing ϵ is known as the ‘interference term’, and generically distinguishes distributions derived from (anti)symmetric states from those pertaining to asymmetric product states like $|\phi\rangle|\psi\rangle$.

Interference terms have a significant impact on the probability of scattering outcomes. Consider a process in which two beams of particles are directed at one another, with a detector placed at an angle θ relative to the common axis of the incident beams (see Figure 5.1). One can compute the quantum mechanical differential cross section for such a process as a function of θ .⁴⁰ This differential cross section—which carries units of area per solid angle—is roughly a measure of how likely particles are to scatter into a differential solid angle centered around θ after they collide. If the particles are distinguishable, the quantum result matches the classical analog (Landau 1996, 59):

$$\frac{d\sigma_C}{d\Omega}(\theta) = \frac{\eta^2}{4k^2} \left(\frac{1}{\sin^4(\theta/2)} \right) \quad (5.3.3)$$

Equation (5.3.3) is identical to the differential cross-section for the classical Coulomb scattering of charged particles. In this expression, $k = \sqrt{2\mu E}$ where μ is the reduced mass of a particle pair and E is the energy of each incoming particle. The second constant, η , is called the “Coulomb parameter” (Landau 1996, 58) and depends upon both k and the charges of the colliding species ($\eta = \frac{\mu Z_1 Z_2 e^2}{k}$). If the particles are indistinguishable as far as our experimental apparatus is concerned but are nonetheless of different types, then to predict the measured cross-section we should make Equation (5.3.3) symmetric around $\theta = \pi/2$. That is, we should consider $\frac{d\sigma}{d\Omega}(\theta) = \frac{d\sigma_C}{d\Omega}(\theta) + \frac{d\sigma_C}{d\Omega}(\pi - \theta)$. This is because we cannot tell whether a measured particle was scattered from the beam on the left, or is a recoiling particle from the beam on the right (Landau 1996, 156). Symmetrizing the equation in this way is analogous to constructing the disjunctive probability expressed in Equation (5.3.1). The result is an expression for the quantum mechanical differential cross section when PI does *not* apply:

$$\frac{d\sigma}{d\Omega}(\theta) = \frac{\eta^2}{4k^2} \left(\frac{1}{\sin^4(\theta/2)} + \frac{1}{\cos^4(\theta/2)} \right) \quad (5.3.4)$$

If PI *does* apply, then interference terms appear in the formulas for cross sections. For bosons, the correct expression is (Landau 1996, 159):

$$\frac{d\sigma}{d\Omega}(\theta) = \frac{\eta^2}{4k^2} \left(\frac{1}{\sin^4(\theta/2)} + \frac{1}{\cos^4(\theta/2)} + \frac{8 \cos[2\eta \ln(\tan(\theta/2))]}{\sin^2 \theta} \right) \quad (5.3.5)$$

⁴⁰ See Chapter X, Section 1 of (Messiah 1999) for an accessible overview. See Chapter 10 of (Landau 1996) for derivations of the analogous cross-sections for fermions and bosons. A completely general expression for the quantum cross-section for Coulomb scattering can be found in (Plattner and Sick 1981).

For fermions, the differential cross section is given by (Landau 1996, 160):

$$\frac{d\sigma}{d\Omega}(\theta) = \frac{\eta^2}{4k^2} \left(\frac{1}{\sin^4(\theta/2)} + \frac{1}{\cos^4(\theta/2)} - \frac{4 \cos[2\eta \ln(\tan(\theta/2))]}{\sin^2 \theta} \right) \quad (5.3.6)$$

The third term in parentheses in both Equation (5.3.5) and Equation (5.3.6) is an interference term—positive for bosons and negative for fermions. Each of these three cross sections is plotted in Figure 5.2. The two expressions representing bosons and fermions diverge dramatically from what one would expect if PI did not apply. At $\theta = \pi/2$, the cross section for fermions is 1/2 that for unsymmetrized particles—those described by Equation (5.3.4)—while the cross section for bosons is twice as great. This dramatic difference has been repeatedly confirmed by experiment. Quantum scattering experiments strongly suggest the need for PI.

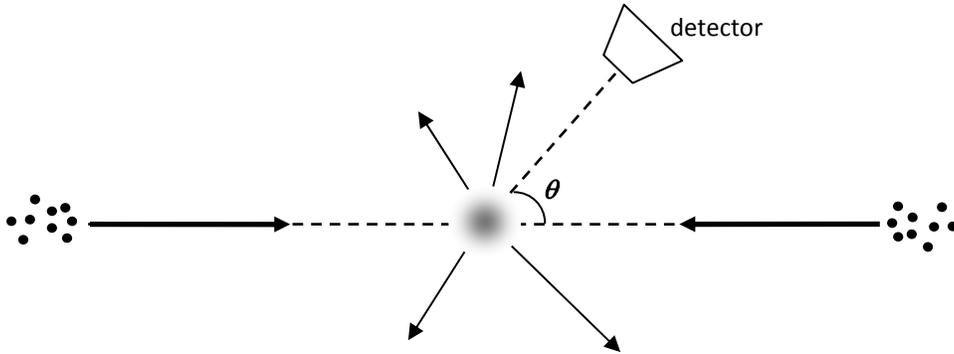


Figure 5.1 A simple scattering experiment in which two beams of particles are directed toward one another along a common axis. A detector is placed at an angle θ relative to this axis, allowing one to count the particles scattered into a solid angle about θ per unit time.

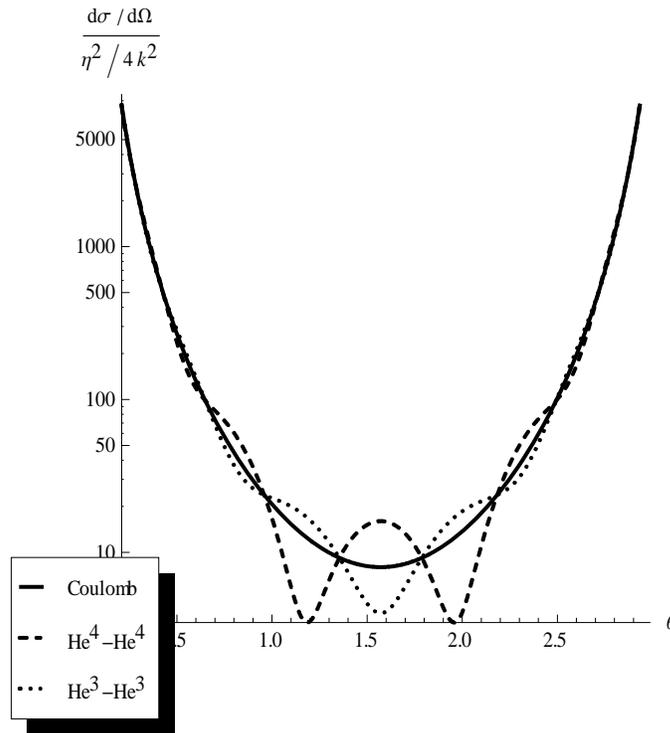


Figure 5.2 Comparison of the differential cross-sections for collisions involving distinguishable He^3 and He^4 nuclei (solid line), pairs of indistinguishable He^4 nuclei (dashed line), and pairs of indistinguishable He^3 nuclei (dotted line). The helium nuclei have a charge of +2; He^3 is a fermion while He^4 is a boson. The solid line is labeled ‘Coulomb’ because it is identical to the classical prediction for Coulomb scattering. All curves plotted are those predicted by QM and do not reflect experimental data.

The imposition of PI also accounts for much of the electronic structure of atoms and molecules. To see this qualitatively, consider the probability of finding two electrons in an atom occupying the same state. Since electrons are fermions ($\epsilon = -1$), Equation (5.3.2) indicates that this probability is 0. This is the basis for the *Pauli Exclusion Principle*: “two [fermions of the same type] cannot be in the same individual state” (Cohen-Tannoudji, Diu, and Lalo  e 1977, 1389). We can read “individual state” in this case as a maximal set of jointly measurable properties. The Exclusion Principle tells us that we cannot find two electrons with the same energy, spin direction, and orbital angular momentum. This explains the hierarchical series of electronic ‘orbitals’ found in atoms. Roughly speaking, individual electron states can be characterized by three numbers:⁴¹ the ‘principal quantum number’ n which corresponds to total energy, the ‘angular momentum number’ l which corresponds to the total state-dependent angular momentum, and the ‘magnetic quantum number’ m which corresponds to the direction of the intrinsic angular momentum of an electron. Since the Exclusion Principle precludes finding two electrons with all three of these numbers in common, the electrons in an atom are forced to distribute over available combinations of quantum numbers. This distribution leads to the rich structure upon which chemistry is predicated. For instance, it guarantees that the alkali metals all have a single electron in the highest energy state, and this in turn can be shown to account for the tendency of atoms in this group to readily ionize and donate an electron in chemical reactions. I haven’t the space to explicate in any detail the central role of the Exclusion Principle—and thus PI—in quantum chemistry. Suffice it to say that the empirical ramifications of PI are legion for atomic and molecular physics.

Finally, I should note that PI carries far-reaching implications for quantum statistical mechanics. Statistical mechanics depends crucially on counting states, in particular the number of states accessible to a system for a given range of values of system properties (e.g. total energy or temperature). PI reduces the number of states accessible to a system. For instance, suppose that $H^{(1)}$ is spanned by two basis states $|0\rangle$ and $|1\rangle$. Consider an operator \hat{A} on $H^{(2)}$ with eigenvectors $|0\rangle|0\rangle$, $|0\rangle|1\rangle$, $|1\rangle|0\rangle$, and $|1\rangle|1\rangle$ with eigenvalues 0, 1, 1, and 2 respectively. If \hat{A} corresponds to some physical property A and we know the system to be in a state with $A = 1$, then without PI there are four states accessible to the system: $|0\rangle|1\rangle$, $|1\rangle|0\rangle$, $\frac{1}{\sqrt{2}}(|0\rangle|1\rangle + |1\rangle|0\rangle)$, and $\frac{1}{\sqrt{2}}(|0\rangle|1\rangle - |1\rangle|0\rangle)$. However, if PI applies then there is only one state accessible to systems made up of bosons [$\frac{1}{\sqrt{2}}(|0\rangle|1\rangle + |1\rangle|0\rangle)$] and only one state accessible to systems of fermions

⁴¹ This material is covered in any introductory chemistry textbook, e.g. (Chang 1994).

$[\frac{1}{\sqrt{2}}(|0\rangle|1\rangle - |1\rangle|0\rangle)]$. In general, bosons and fermions will yield vastly different numbers of possible states compatible with some range of observable values and both yield statistics that are very different from what we would expect if PI did not apply. A treatment of the differences between these two kinds of quantum statistics is beyond the scope of this project. I refer the interested reader to any introductory text on quantum statistical mechanics [e.g. (Pathria 1996)].

Each of the three kinds of empirical consequences considered above—cross sections in collision processes, atomic structure, and divergent statistical behavior—follow from PI, and each is supported by the experimental data. More accurately the absence of any evidence that particles of the same type can occupy states which violate PI weighs heavily in favor of the principle. To the extent that QM is experimentally verified, so too is PI. Thus, if QM is a theory of particles, then it is a theory of particles that respect PI.

5.4 Generalization of PI

Now that we’ve considered PI as it is expressed in QM, we can abstract away those aspects of the principle that refer specifically to the quantum formalism in order to arrive at a general, theory-independent statement. We can begin by combining PI_{QM1} with the two postulates concerning state representations in order to eliminate talk of vectors in a Hilbert space:

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

There are a few things to note about this revised version of PI. To begin with, it eliminates vague talk of permuting particles in favor of permuting parts of state representations—precisely the action of the permutation operators in QM. Second, “physical states” are states of the world, not the mathematical structures used to represent them. Physical states are what *interpretations* describe. Finally, it is clear from this formulation of PI that the principle makes a strong assumption about the manner in which models of the theory (trajectories through sets of vectors in a Hilbert space) are mapped onto interpretations. In general, we need only assume that there is a function taking models of a theory to interpretations. PI assumes further that this interpretation function is a composite of functions which map *parts* of a model to *parts* of an interpretation. Specifically, parts of the quantum model—a mathematical object—are taken to correspond directly to particle specifications in the interpretation of that model. The parts in question for QM are precisely those which are shuffled about by

the permutation operator, namely the vectors lying in the lower dimensional Hilbert space (the $|u_\alpha\rangle$) out of which an outer product space is constructed.

A couple of examples will help to illustrate this idea of piecewise interpretation. Consider the time independent quantum mechanical treatment of the hydrogen atom. The theory in this case consists of an appropriate Hamiltonian and some boundary conditions on wavefunctions. Models of this theory will be vectors in an outer product space built from kets in lower-dimensional spaces which represent the state of an electron and a proton. So models of the theory will contain parts—namely kets like $|\phi\rangle_{\text{electron}}$ and $|\psi\rangle_{\text{proton}}$ —that are mapped individually to single particle specifications in interpretations of the theory. Importantly, it must be the case that the same electron ket maps to the same electron specification across all models and interpretations of this theory. Likewise for the proton kets.

To give another example of the piecewise interpretation implicit in PI, this time from classical mechanics, consider the physics of pendulums. The motion of a single simple pendulum can be characterized by two coordinates: an angular position, and an angular momentum. The 2-D mathematical space constructed from these coordinates is the pendulum's phase space, and models of Hamilton's Equations of Motion are trajectories through this phase space. The interpretation of these models is trivial: one coordinate of each point in the phase space is taken to correspond to the momentum of the bob at a moment of time, and the other coordinate to its position. To construct a classical theory of two pendulums, we can move to a 4-D phase space built from four coordinates, namely the angular positions and momenta of the two bobs. Models are now trajectories in this 4-D phase space. Interpretations of these models contain two specifications, one for each pendulum, and proceed piecewise. In this case, the 'pieces' in question are the 2-D phase spaces from which the 4-D space is built. The properties attributed to each pendulum depend only upon the projection of a trajectory into the 2-D space isomorphic to the phase space of that pendulum considered in isolation. As long as the projected trajectory for a pendulum is the same as the original trajectory in the 2-D phase space for an isolated pendulum (which it won't be if the pendulums interact at all), then the pendulum must be given the same specification in the interpretation of that trajectory. PI assumes that a common set of parts can be extracted from the models of a theory and mapped individually to descriptions of objects one at a time when the model is interpreted.

Our aim in this chapter is to ascertain a clear statement of PI divorced from the details of QM. It is already the case that PI_{QM2} has been purged of any reference

to the specific formalism of QM, but leaves implicit an important constraint PI imposes on the interpretations of a theory. With the terminology introduced in Chapter 4, we can make this presumed relation between parts of models and interpretations explicit:

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

This statement of PI is the theory-independent form of the principle with which we will be concerned for the remainder of this essay.

Chapter 6 Proving an inconsistency

6.1 More definitions

At the end of Chapter 5, I pointed out that PI implicitly refers to parts of the models of a theory. In particular, PI refers to the parts of a model which individually correspond to specifications in interpretations of that model. I will call the part of a model which corresponds to the specification of one particular particle the *role* of that particle in the model.

Some examples might help to clarify the notion. Suppose we are concerned with the motion of a pair of particles in a spring-like potential confined to move in one dimension. If the amplitude of the motion is small enough and the particles are uncoupled, then the relevant theory of Newtonian mechanics is a pair of differential equations $m\ddot{x}_i = -kx_i$, where $i = 1, 2$. Models of the theory are trajectories in a 4-dimensional phase space, and each particle role is a projection of the 4-D trajectory into a 2-D space. Each role is interpreted in terms of the changing position and momentum of a particle through time, with one dimension referring to the momentum and the other to the position of the particle.

As a second example, consider QM. In standard approaches to time-independent QM, models are typically taken to be vectors in an outer-product Hilbert space. Consider vectors of the form $|v_1\rangle = |\phi\rangle|\psi\rangle$, $|v_2\rangle = a_1|\phi\rangle|\psi\rangle + a_2|\chi\rangle|\xi\rangle$, and $|v_3\rangle = \frac{1}{\sqrt{2}}(|\phi\rangle|\psi\rangle + |\psi\rangle|\phi\rangle)$. Vectors of this sort might constitute models of a two-particle theory of QM. A plausible set of particle roles for each of these vectors consists of the parts that get rearranged by the permutation operators defined in Section 5.2. Specifically, they are the sets of component ‘one-particle’ kets corresponding to just one of the subspaces making up the outer-product space weighted by the constant appearing in the term from which the ket is taken. For the vector $|v_1\rangle$, one of the particle roles is $\{|\phi\rangle\}$ and the other is $\{|\psi\rangle\}$, while for $|v_2\rangle$ the roles are $\{a_1|\phi\rangle, a_2|\chi\rangle\}$ and $\{a_1|\psi\rangle, a_2|\xi\rangle\}$. For the third vector, which has the right form to satisfy PI, the roles are identical: $\{\frac{1}{\sqrt{2}}|\phi\rangle, \frac{1}{\sqrt{2}}|\psi\rangle\}$. The interpretation of these roles is typically understood probabilistically. While the first role in $|v_1\rangle$ corresponds to the specification of a particle bearing all the properties of the state $|\phi\rangle$ with probability 1, the first role in $|v_2\rangle$ corresponds to a specification in which a particle is attributed the properties of $|\phi\rangle$ with probability $\|a_1\|^2$ and the properties of $|\chi\rangle$ with probability $\|a_2\|^2$.

Despite what the last example might suggest, the particle roles in a model needn't align with the way we syntactically represent the model—they might instead be extracted from models via more elaborate processes. For example, suppose we represent states with density operators on a Hilbert space rather than vectors. In a standard approach to attributing states to individual particles in QM, each particle is assigned a reduced density operator that is computed by taking a partial trace over the density operator representing the two-particle state. In this case, each of the resulting reduced density operators constitutes a particle role.

Ultimately, the set of mathematical objects which counts as the relevant set of parts of a model—in particular the parts pertinent to PI—is determined by the interpretation scheme one assumes. The particle roles are whichever parts of the model are interpreted as individually corresponding to specifications. In the most general case then, we can say that, given a particular scheme for interpreting models, there is a function from the space of n -particle models to n -tuples of mathematical objects where each object in this n -tuple corresponds to the description of a particle in the interpretation of the model from which the parts were extracted.

In order to talk about PI, I need to make the connection between particle roles and specifications more precise, and to define the action of permutation on both. To do so, suppose we have fixed a method for interpreting models of a theory such that, given any model we can determine how many particles it represents and what the particle roles are. Let A be the set of particle roles for a given model M , and let n be the total number of particles represented by M . Let N be the set of natural numbers $\{1, 2, \dots, n\}$. We can then introduce a surjective function $\mathbf{f} : N \rightarrow A$. The function \mathbf{f} is an indexing function which allows us to recover the roles for each particle. So, for instance, the j^{th} particle role in M is given by $\mathbf{f}(j)$. Note that there needn't be as many distinct roles as there are particles—multiple particles may have the same role in a given model as in the example from QM above.

Corresponding to each model M , there is an interpretation from which we can extract n specifications (descriptions of individual particles). In particular, each particle role corresponds to one specification. So let S be the set of all one particle specifications. Then we can define another surjective function $\mathbf{interp} : A \rightarrow S$ that maps particles roles to their respective specifications in the interpretation of M . Again, more than one role may correspond to the same specification. By nesting functions, we thus have an indexed set of particle specifications corresponding to model M . The j^{th} specification is given by

$\mathbf{interp}(\mathbf{f}(j))$. It will prove helpful to bear in mind the following property of interpretations. Suppose M_1 and M_2 are models which represent the same number of particles, and \mathbf{f}_1 and \mathbf{f}_2 are the functions indexing the particle roles of each model respectively. If the models M_1 and M_2 have the same interpretation, then for all j , $\mathbf{interp}(\mathbf{f}_1(j)) = \mathbf{interp}(\mathbf{f}_2(j))$. The converse is not necessarily true.

Finally, we need to formalize the notion of permuting particle roles. A permutation of a set is a bijective mapping of the set onto itself. Let $\mathbf{P}_i : N \rightarrow N$ be the i^{th} permutation (there are $n!$ of these bijective functions for a set of cardinality n). Then the i^{th} permutation of particle roles for a model is just the indexed set of particle roles for which the j^{th} member is $\mathbf{f}(\mathbf{P}_i(j))$.

6.2 The general argument

Suppose M is a model which is taken to represent a collection of particles all of the same type.⁴² According to PI, if we permute the particle roles in M according to the i^{th} permutation, the result is a model M_i with the identical interpretation. If the j^{th} specification in the interpretation of M is given by $\mathbf{interp}(\mathbf{f}(j))$, then the j^{th} specification in the interpretation of M_i is $\mathbf{interp}(\mathbf{f}(\mathbf{P}_i(j)))$. Now, since M and M_i have the same interpretation, it must be the case that $\forall_j [\mathbf{interp}(\mathbf{f}(j)) = \mathbf{interp}(\mathbf{f}(\mathbf{P}_i(j)))]$. According to PI, this relation holds for every permutation \mathbf{P}_i . That is:

$$\forall_{i,j} [\mathbf{interp}(\mathbf{f}(j)) = \mathbf{interp}(\mathbf{f}(\mathbf{P}_i(j)))] \quad (6.2.1)$$

Formula (6.2.1) is satisfied if and only if $\mathbf{interp}(\mathbf{f}(j)) = \text{constant}$. That is, if PI obtains then all particle specifications must be identical—every particle is attributed the same properties in the interpretation of M .

The fact that all particles represented by any model M of theory T bear the same specification in the interpretation of M poses a problem if the putative particles are to satisfy MA. In Section 4.3, I stated three conditions T must satisfy if it is to admit interpretations involving two particles that satisfy MA(iii). Suppose that the first two of these conditions are satisfied. That is, suppose that there are two sets of one-particle models of T from which we can extract two sets of specifications—call them S_1 and S_2 —each of which is

⁴² The assumption that all particles are of the same type is made for simplicity. The argument goes through—albeit in a more complex form—as long as at least two putative particles are of the same type.

large compared to a given threshold ϵ .⁴³ If the third condition is also satisfied, then for any pair of arbitrary specifications $s_1 \in S_1$ and $s_2 \in S_2$ we must be able to find a two-particle model of T from whose interpretation we can extract a pair of specifications, one of which is within ϵ of s_1 and one of which is within ϵ of s_2 . Now suppose we choose s_1 and s_2 such that the distance between them—denoted $d(s_1, s_2)$ —is greater than 2ϵ . Such a pair must exist because both sets are large compared to the threshold, meaning that the maximum distance between two points in each set is much greater than ϵ . When we turn to the interpretations of two-particle models, PI entails that both specifications are identical—we can only extract specification pairs of the form $\{s^*, s^*\}$. To satisfy MA, we need to find a two-particle model from which we can extract such a pair of specifications for which $d(s^*, s_1) \leq \epsilon$ and $d(s^*, s_2) \leq \epsilon$. That is, we need to find a model such that:

$$d(s^*, s_1) + d(s^*, s_2) \leq 2\epsilon \quad (6.2.2)$$

However, the triangle inequality for metrics guarantees that $d(s_1, s_2) \leq d(s^*, s_1) + d(s^*, s_2)$. By supposition, $2\epsilon < d(s_1, s_2)$, and so

$$2\epsilon < d(s^*, s_1) + d(s^*, s_2) \quad (6.2.3)$$

But this is a contradiction. There does not exist an s^* such that both Equations (6.2.2) and (6.2.3) are satisfied. Thus, the theory *cannot* be compatible with MA. One cannot consistently interpret a theory in terms of particles in the sense of MA if that theory incorporates PI. If PI holds, there are no particles.

6.3 The argument for QM – version 1

In the preceding section, I stated the incompatibility argument in its strongest, most general form. Of course, one can also show that PI and MA are incompatible in the context of a particular theory and with the assumption of a particular metric for comparing particle specifications. A natural theory for which to do this is QM, the first theory to explicitly formulate and endorse PI. In this and the following section, I run the argument using two versions of the quantum formalism and two different metrics. While this is redundant in some respects,

⁴³ More accurately, in order to assess whether the theory T can be interpreted as representing two particles in accord with MA, we need to consider a theory T' which is effectively a one-particle version of T . I am ignoring this complication and assuming that for any scientific theory, e.g. QM, it is unambiguous which set of equations and boundary conditions are single particle theories and which are multiparticle. Cf. fn 33 in Chapter 4.

my intention is to illustrate both the irrelevance of one's choice of metric and the inevitability of the conclusion no matter how liberally one reads the quantum formalism. The reader who is satisfied with the general argument is invited to skip one or both of these illustrations.

To see how the argument works within QM, we first need to be clear on what constitutes the theory and how models of the theory are to be interpreted. In my first approach, I will take the theory of QM to be the dynamical postulates which specify trajectories in a Hilbert space according to either the Schrödinger or Heisenberg equations of motion. For simplicity, I'll ignore time dependence and focus just on time-independent solutions of the equations of motion for suitably specified boundary conditions. In the time-independent case, models are just rays in a Hilbert space.⁴⁴

How should these models be interpreted? What objects are specified by the models and how should we assign properties to them? In the case of a single-particle system, each ray in the Hilbert space is taken to specify one particle in the world along with its properties. If we adopt the conservative eigenvector-eigenvalue link, then our interpretive function would attribute one property to the particle for each of the self-adjoint operators for which the ray representing that particle contains an eigenvector. For instance, if we are considering a stationary state of a particle in a square potential well, then (since it is stationary) its state representation must be an eigenvector of the Hamiltonian. In our interpretation, we would attribute to the particle an energy represented by the corresponding eigenvalue. However, since the state representation cannot also be an eigenvector of the position operator, this interpretive scheme would preclude us from attributing any position property to the particle in our interpretation.

Assuming that particles in an interpretation are only ascribed properties for which the associated ray contains an eigenvector, then there is a one-one correspondence between rays in the Hilbert space and descriptions in the interpretations of one-particle models.⁴⁵ Because models of QM (as I've construed it here) reside in projective Hilbert spaces (the set of rays of a Hilbert space), there is a natural choice of metric for measuring the distance between models. Specifically, the distance between any two rays r_i and r_j in the projective Hilbert space is given by the Fubini-Study metric:⁴⁶

⁴⁴ This is only true if we neglect the possibility of paraparticles.

⁴⁵ This assumption rules out the standard approach of using reduced density operators to represent (mixed) single-particle states. In the next section, I will lift this restriction and employ a different metric.

⁴⁶ See e.g. (Bengtsson and Życzkowski 2006), section 5.3.

$$d_{\text{FS}}(r_i, r_j) = \arccos \left(\sqrt{\frac{\langle \psi | \phi \rangle \langle \phi | \psi \rangle}{\langle \psi | \psi \rangle \langle \phi | \phi \rangle}} \right),$$

where $|\psi\rangle$ is any one of the vectors lying along r_i and $|\phi\rangle$ is any one of the vectors lying along r_j . Since there is a one-one correspondence between one-particle rays and descriptions for the interpretive scheme we've assumed, we can transfer this metric onto the space of particle specifications. That is, the distance between any two specifications may be defined as the distance between the corresponding rays.

If we consider the case of two-dimensional particle states, this metric makes the inconsistency argument easy to visualize. For instance, we might consider just the spin degree of freedom of a spin- $1/2$ particle. In that case, the Fubini-Study metric describes the geometry of the surface of a sphere embedded in Euclidean three-dimensional space (see Figure 6.1).⁴⁷ Each point on the sphere corresponds to a ray in the 2-D Hilbert space (which in turn corresponds to a particle specification), and the Fubini-Study metric gives the distance between two states (or two specifications) along a great circle of the sphere. To visualize the argument on this sphere, pick a threshold distance ϵ small compared to $\sup(d_{\text{FS}}(r_i, r_j))$. Now choose a pair of rays r_1 and r_2 from the space of one-particle models such that $d_{\text{FS}}(r_1, r_2) > 2\epsilon$. According to MA(iii) we must be able to find a two-particle model that—when interpreted—yields two copies of a single-particle description corresponding to a ray r^* that is within ϵ of both r_1 and r_2 . This means that r_1 and r_2 must lie within a disc of radius ϵ centered on r^* . Obviously, if this is the case then it cannot also be true that r_1 and r_2 are more than 2ϵ apart since the diameter of the disc is 2ϵ . Thus we arrive at a contradiction. Given PI and the Fubini-Study metric on the space of quantum models, MA(iii) is not satisfiable.

⁴⁷ Up to a multiplicative constant, the differential form of the Fubini-Study metric is the metric of the 'Bloch sphere'. For a discussion of the Bloch sphere, see, e.g. (Dickson 2007, Sec. 1.3.5; Nielsen and Chuang 2000). For a detailed discussion of the geometric interpretation of the density matrix in arbitrary dimensions (though they do not use the term 'Bloch sphere'), see (Filippov and Man'ko 2008).

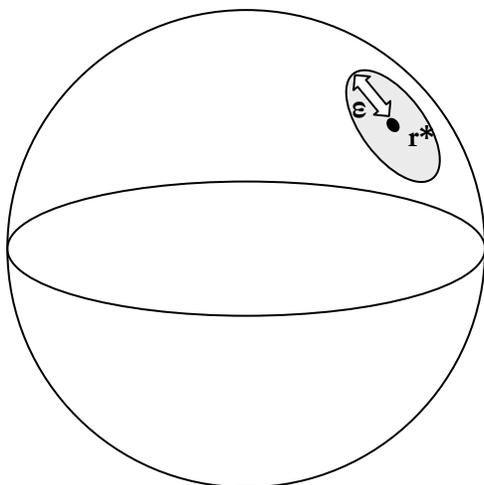


Figure 6.1 The argument visualized for pure states on the Bloch Sphere. The disk indicates a set of specifications within a distance ϵ of the specification \mathbf{r}^* . If QM is to satisfy MA, it must be possible to find two additional specifications within this disk but greater than 2ϵ apart—a clear impossibility.

6.4 The argument for QM - version 2

In the preceding section, I assumed a very conservative view of what constitutes the theory of QM and how its models should be interpreted. In this section, I'll run the argument again, this time taking a more catholic view. I will again ignore time dependence for simplicity, but to be as inclusive as possible, I'll consider the space of state representations in QM to consist of full set of density operators on the relevant Hilbert space. This means that both 'pure' states equivalent to rays in the Hilbert space *and* 'mixed' states—convex combinations of pure states—will count as potential models of the theory. The theory in this view consists of the set of density operators on a Hilbert space and a dynamical equation known as the von Neumann Equation:

$$i\hbar\dot{\rho} = [\rho, \hat{H}]$$

This equation specifies how density operators representing physical states evolve in time or, in the time-independent case, stipulates what states are compatible with a given Hamiltonian.

In this construal of QM, models of the theory are just single density operators. In the case of a single-particle system, each density operator can be interpreted as assigning a full set of properties to a particle in the weak sense of bearing properties that I introduced in Section 3.5. That is, we can drop the eigenvector-eigenvalue link and interpret each density operator as assigning to a particle a set of probability distributions over all the values the particle

might be measured to have for each observable. Equivalently, the density operator assigns to each particle a set of expectation values with respect to all observables, and we treat these as the vague properties of the particle. In this approach, even mixed states—which assign no properties to particles under the eigenvector-eigenvalue link—can still be interpreted in terms of a particle with a determinate set of properties.

Of course, we can no longer use the Fubini-Study metric to assign distances between particle interpretations, since that metric refers to rays in the Hilbert space. Instead, we need a metric on density operators. To once again make the argument easy to visualize, I will take the distance between any two density operators ρ_1 and ρ_2 to be twice the standard ‘trace distance’⁴⁸:

$$d_{\text{TR}}(\rho_i, \rho_j) = \text{Tr} \|\rho_i - \rho_j\|$$

The operation indicated by $\text{Tr} \|\cdot\|$ is the ‘trace-norm’.⁴⁹ Once again, there is a one-one correspondence between models of the theory (density operators) and particle specifications (unique assignments of properties corresponding to sets of expectation values). So once again, we can define a metric on the space of specifications by way of the metric on state representations. In particular, the distance between two specifications is just the trace distance between the two density operators mapped to those specifications in the process of interpretation.

If we consider the case of two-dimensional particle states (e.g. just the spin-degree of freedom of a spin- $\frac{1}{2}$ particle), then this choice of metric again makes the argument easy to visualize. In the two-dimensional case, we can represent density operators as vectors in a three-dimensional Euclidean space. In particular, for each density operator ρ (represented as a 2×2 matrix in the basis of eigenstates of spin in the \hat{z} -direction) there exists a unique vector \vec{r} such that $\|\vec{r}\| \leq 1$ and

$$\rho = \frac{\hat{I} + \vec{r} \cdot \vec{\sigma}}{2},$$

where \hat{I} is the identity operator and $\vec{\sigma}$ is a vector containing the three Pauli matrices ($\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$, $\sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}$, $\sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$). Geometrically, the space of vectors \vec{r} is just the set of points of the Bloch sphere, including those of the interior. That is, for every point inside or on the surface of a Euclidean sphere of radius 1, there is corresponding vector \vec{r} and thus a corresponding density operator ρ . The points on the surface of the sphere (those we considered in the

⁴⁸ See e.g. (Nielsen and Chuang 2000), Section 9.2.1.

⁴⁹ See (Omnès 1994, 245).

previous example) correspond to pure states, while those interior to the sphere correspond to the mixed states. If we map density operators onto points of the Bloch sphere in this way, the trace distance turns out to be equivalent to the Euclidean distance $\|\vec{r}_i - \vec{r}_j\|$.

To visualize the incompatibility result in this framework, choose a threshold distance $\epsilon \ll 1$. Now choose any two points \vec{r}_1 and \vec{r}_2 in the sphere such that $\|\vec{r}_1 - \vec{r}_2\| > 2\epsilon$. If MA(iii) holds for QM, we must be able to find a two-particle model that—when interpreted—yields two copies of a particle specification corresponding to the point \vec{r}^* and such that $\|\vec{r}^* - \vec{r}_1\| \leq \epsilon$ and $\|\vec{r}^* - \vec{r}_2\| \leq \epsilon$. Geometrically, this means we need to find two points that lie within a ball of radius ϵ centered at the point \vec{r}^* but that are nonetheless more than 2ϵ apart (see Figure 6.2). This is clearly impossible. Thus, even in this liberal reading, QM cannot satisfy MA(iii) if PI is imposed.

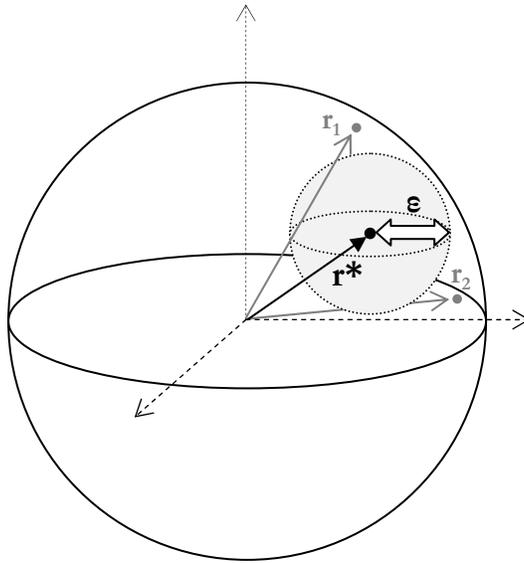


Figure 6.2 The argument visualized for pure and mixed states in the Bloch Sphere. For QM to be compatible with MA, it must be possible to locate two specifications \mathbf{r}_1 and \mathbf{r}_2 that are within a distance ϵ of the specification \mathbf{r}^* (and thus within the shaded sphere) but that are nonetheless a distance of more than 2ϵ apart. This is clearly impossible since it requires one member of the pair to be both within and without the shaded sphere.

Chapter 7 Objections

7.1 Overview

The preceding chapter concluded an argument to the effect that any theory which includes an assertion of PI is necessarily incompatible with MA, and thus cannot be interpreted in terms of particles of the sort that can account for EDiv. There are, of course, a number of ways one might take issue with either the manner in which I've framed the argument or with the conclusion itself. In this chapter, I'll consider what I take to be the five most serious objections.

7.2 Objection: QM tells us when to ignore PI

If PI really imposes an inescapable dependence between the properties of one particle and the properties of every other particle of the same type, then this would seem to make physics intractable. If it were true, how could we successfully predict anything about electrons in the laboratory without knowing about all the electrons in the universe? Since QM is obviously a successful empirical science, it seems that I must be wrong about PI. This is more or less how many introductory texts on QM treat the implications of PI. Physicists often register a concern over property independence, only to promptly dismiss it with a quick calculation. Here, for instance, is how Messiah expresses the worry:

We consider a system of n identical particles. If the particles are electrons, the state of the system will be represented by an antisymmetrical wave function. These are not, however, the only electrons in the universe. To ignore the others, and to treat this system as an entity distinct from the rest, supposes that the dynamical properties of the n electrons are not affected by the presence of the others. The question arises whether such a hypothesis is well founded, or whether the symmetrization postulate, in establishing a certain correlation between these n electrons and the others, renders it invalid.⁵⁰

(Messiah 1999, 600)

Messiah's response is typical of the views expressed in the physics literature:

In practice, the electrons of a system are all inside a certain spatial domain \mathcal{D} , and the dynamical properties in which we are interested all

⁵⁰ By the "symmetrization postulate" (SP), Messiah is referring to a stronger principle than PI. SP is roughly the claim that PI holds *and* all states are represented by single rays in the Hilbert space. We'll meet SP again in the next chapter.

correspond to measurements to be made inside this domain. It turns out to be true that the other electrons may simply be ignored so long as they remain outside \mathcal{D} and so long as their interaction with the electrons of the system remains negligible. This is a general result and applies to bosons as well as fermions.

(Messiah 1999, 600-1)

The claim that electrons outside of the domain of interest can be ignored is backed up by the following calculation (Messiah 1999, 601-3). Consider two separate quantum systems. In the first system, \mathcal{S}_1 , there are two fermions in a state $|\phi\rangle$ with wavefunctions strictly confined to some region of space \mathcal{D} . Denote by $|\chi\rangle$ any of the normalized antisymmetric states whose wavefunctions vanish outside \mathcal{D} , including the state $|\phi\rangle$. The probability that the system \mathcal{S}_1 is found to occupy one of these states $|\chi\rangle$ is just $\|\langle\chi|\phi\rangle\|^2$. Now suppose we consider a second system, \mathcal{S}_2 , made up of $(N - 2)$ additional fermions of the same kind as those in \mathcal{S}_1 and occupying a state $|\Psi\rangle$. Suppose further that the joint wavefunction corresponding to $|\Psi\rangle$ vanishes inside \mathcal{D} , and denote by $|\Theta\rangle$ any ket with this property, and let $|\Theta_i\rangle$ denote members of a set of such vectors that form an orthonormal basis. By supposition $|\Psi\rangle = \sum_i c_i |\Theta_i\rangle$ for some set of $|\Theta_i\rangle$. If the particles were distinguishable (and thus, not fermions), the joint state of the combined system $\mathcal{S}_1 \cup \mathcal{S}_2$ (assuming the electrons of one system do not interact with those of the other) would be represented by $|\phi\rangle|\Psi\rangle$. However, since the fermions are all indistinguishable (they are all of the same kind), the state of the system must be represented by⁵¹

$$|\Phi\rangle = \sqrt{\frac{N!}{2!(N-2)!}} \hat{A}|\phi\rangle|\Psi\rangle,$$

where \hat{A} is the ‘antisymmetrizer’—the projector onto the antisymmetric subspace of the underlying Hilbert space.⁵² In this case, we can now ask for the probability of finding just the \mathcal{S}_1 subsystem “...in any one of the states represented by the orthonormal antisymmetric vectors” (Messiah 1999, 603):

$$|X_i\rangle = \sqrt{\frac{N!}{2!(N-2)!}} \hat{A}|\chi\rangle|\Theta_i\rangle$$

The required probability is given by

$$w = \sum_i \|\langle X_i|\Phi\rangle\|^2 = \|\langle\chi|\phi\rangle\|^2$$

⁵¹ The normalization prefactor differs from the usual $1/N!$ because those permuted vectors which result from swapping kets in $|\phi\rangle$ with kets in $|\Psi\rangle$ are orthogonal to all kets of the form $|\chi\rangle|\Theta\rangle$. See (Messiah 1999, 602).

⁵² See, e.g., (Cohen-Tannoudji, Diu, and Lal e 1977, 1384).

which is exactly the probability obtained above for finding the isolated two-particle subsystem \mathcal{S}_1 in state $|\chi\rangle$. “Thus we can just ignore the existence of the $(N - 2)$ other fermions and still obtain the correct result” (Messiah 1999, 603).

It would appear that we have a straightforward refutation of my thesis. I have argued that, if we take the subsystem of interest to be a single fermion, then it is never the case that the other fermions—whether inside \mathcal{D} or not—can be safely ignored, at least as far as we are interested in the dynamical properties of the lone fermion. But Messiah’s calculation—like the many related calculations appearing in other textbooks—seems to suggest otherwise. The general conclusion is that, so long as the wavefunction attributable to a particle does not overlap the wavefunction of another, we can ignore all other particles. More importantly, there are circumstances under which we can ignore PI and still interpret QM in terms of independent particles.

In one sense, of course, Messiah and his colleagues are exactly right. There is nothing wrong with the mathematics of his derivation. But what he says is also misleading—there is something very wrong with his implied interpretation of the probability w . It is no accident that he considers the properties of subsystems that are themselves defined by a particular set of properties (i.e. having wavefunctions with spatial support in a particular domain), but eschews asking about the properties of any one fermion. The reason is that the only empirically meaningful probabilities we can calculate *cannot make reference to any one particle*. Thus, the properties of any specific particle are not observable properties. As I mentioned in Chapter 5, PI entails that the only physical observables are symmetric operators that commute with the permutation operators. The eigenstates of such operators are (anti)symmetric vectors⁵³. When we summed over the vectors $|X_i\rangle$ in computing w we were effectively projecting the state of the system onto the eigenstates of some symmetric operator. Such an operator when considered as an observable does not make reference to any one particle—it makes an equivalent reference to all. Put another way, the particle roles in the eigenstates of such an observable are identical. Thus, observables never directly attribute properties to individual particles or to specific subsets of particles. This is why the probability w is typically interpreted instead as the probability that a pair of particles is in state $|\chi\rangle$ without reference to *which* particles are in that state. If we restrict ourselves to the observables allowed by PI, then it is meaningless to ask for the dynamical properties of individual particles—we can only ask about the symmetric properties of ‘subsystems’ whose affiliation with particles is indeterminate.

⁵³ Again, more complicated symmetries are possible for systems of more than two particles, but the same conclusions apply.

In his description of the problem and its resolution, Messiah writes as if electrons are particles in the sense of MA. For instance, he talks about a particular set of electrons being excluded from a particular spatial domain and of electrons individually bearing dynamical properties. But the formal expressions computed to resolve the apparent tension between the particle view and PI tell a different tale. In computing w , Messiah computes the probability “...of finding the system in any one of the states represented by the orthonormal antisymmetrical vectors” (Messiah 1999, 603). Though he began by worrying about the independence of *electrons*, he ends up proving facts about *subsystems* instead. Messiah thus does not resolve the question at issue—he does not tell us whether the properties attributable to an individual electron depend upon the presence of others. In fact, he does not tell us how to attribute properties to individual electrons at all, in part because PI says such an assignment is not physically observable.

What happens if we try to individually attribute properties? The general argument of the last section established that no matter how you try to do so—assuming you are consistent across models—it is impossible to satisfy MA. But let’s take a stab at it anyway. A standard approach to assigning individual properties to electrons is to compute a reduced density operator. That is, if we represent the state of the overall system by a density operator $\rho = |\Phi\rangle\langle\Phi|$, then we can attribute a state to the ‘first’ particle by tracing out over the remaining degrees of freedom: $\rho_1 = \text{Tr}^{(2,3,\dots,N)} \rho$. This is equivalent to computing a marginal distribution from the full joint distribution over all property values. It is a property of (anti)symmetric states that for all i and j , $\rho_i = \rho_j$. That is, every reduced density operator is identical to every other. Thus, every fermion in a system is assigned the identical state. If we attempt to assign states to individual electrons, we find that these states and thus the dynamical properties each electron may be said to bear depends upon the presence of other electrons, irrespective of interaction or distance.

Whatever Messiah’s computation shows, it tells us nothing about *particle* properties. There are certainly conditions under which measurements predicted for subsystems are independent—otherwise a quantum physics would be impossible. What’s at issue is the nature of those subsystems—are they composed of particles that are roughly independent of one another and whose independence accounts for the independence of subsystems, or are they something else entirely? I have argued that they must be something else entirely, and the calculations offered by physicists like Messiah do not speak to this claim. Insofar as we are interested in assigning properties to particles, PI forces us to assign identical

properties to all particles of a type, and thus the properties of an individual electron are inextricably linked to every other electron in existence.

7.3 Objection: PI is compatible with classical physics

I have claimed that the incompatibility between PI and MA is general—*any* theory which incorporates the former is incompatible with the latter. It might be objected that this assertion is refuted by the use of reduced configuration spaces⁵⁴ in classical mechanics (CM). The use of certain reduced spaces—as described below—is equivalent to the imposition of PI, and since CM is clearly compatible with a particle interpretation it must be false that PI is always incompatible with MA.⁵⁵

This objection inappropriately attributes properties of CM to the conjunction of CM and PI, which I'll denote CM+PI. The implicit assumption is that every 'classical theory' is compatible with MA. More explicitly, it is simply assumed that, because CM is compatible with MA, CM+PI also must be compatible. However, when CM+PI is examined in detail, it becomes clear that this monstrous form of CM is patently at odds with a particle interpretation. To argue for this claim, we first need to clarify what it means to incorporate PI in CM. Without loss of generality, we can restrict our attention to a pair of particles constrained to move in one dimension. If we ignore time-dependence, models of CM in this case are just points in $\mathbb{R}^2 \setminus D$. (The notation " $\setminus D$ " just means that we remove the points in the 'diagonal' set $D = \{ \langle x, y \rangle \in \mathbb{R}^2 : x = y \}$. We do this because the particles are assumed to be impenetrable.) Each model in $\mathbb{R}^2 \setminus D$ is interpreted in a straightforward manner: one particle is said to have the property of being located at one of the two coordinates of the point, and the other particle is said to be located at a place corresponding to the remaining coordinate.

One way we might try to impose permutation invariance on the theory of CM is by insisting on symmetric probability distributions over points in \mathbb{R}^2 . To remain classical, these distributions should be interpreted either epistemically (we just don't know which model applies) or in terms of ensembles of two-particle systems prepared in identical but coarse-grained states as would be the case in statistical mechanics. The resulting theory is permutation invariant in that the probability of the system occupying a particular state is insensitive to which particle has which properties. This is the notion of permutation invariance Bach

⁵⁴ Or reduced phase spaces, as the case may be.

⁵⁵ This objection was suggested to me both by Chris Smeenk (in discussion) and by an anonymous reviewer for *Philosophy of Science*.

(1997) defends in his monograph on indistinguishable classical particles. But this sort of indistinguishability is insufficient for satisfying PI as it was formulated in Chapter 5. While the distributions over models are permutation invariant, models (which are still points in $\mathbb{R}^2 \setminus D$) are not. If we permute the parts of a model that are taken to refer to a single particle in the interpretation of the model, namely one coordinate of a pair of coordinates, then we attain a distinct model with a distinct interpretation.

Another strategy is to identify the points in the configuration that differ only by a permutation of their coordinates. That is, we can simply declare that points which differ by a permutation are equivalent in the sense that they represent the same state. We can represent this equivalence by taking the quotient of the configuration space $\mathbb{R}^2 \setminus D$ under the action of the permutation operator Π_{21} . The resulting set is denoted $(\mathbb{R}^2 \setminus D) / \Pi_{21}$ and is called a ‘reduced configuration space’.⁵⁶ A version of Hamilton’s laws of motion obtain on this new space, and so we can be said to have a version of CM that respects PI.

Using a reduced configuration space to represent a classical system in this way entails a radical sort of indistinguishability anathema to classical sensibilities and in contrast to the sort of harmless indistinguishability considered by Bach. As Saunders puts it, the use of a reduced phase space “...takes particle indistinguishability all the way down to the microscopic details of individual particle motions, whereas, according to Bach, it ought to concern only statistical descriptions (probability measures)” (Saunders 2006, 197). In fact, Bach asserts that attempting to take indistinguishability to include *every* property of a particle including its trajectory is incoherent. More accurately, he argues that if classical particles are indistinguishable in this sense, they cannot be said to have trajectories. I won’t go through the details of Bach’s argument [see (Bach 1997, 7-8)], but I will point out that it relies on a strong notion of what it is to bear a property: a particle has a trajectory (or position in the time-independent case) only if there is a probability of 1 of measuring the particle to possess this property. While one can take issue with Bach’s particular approach, there is no escaping the fact that if one is to impose PI on CM and retain particles, one will have to be much more liberal in attributing properties to particles—it is indisputably the case that some part of the classical picture must yield if PI is to be added.

As with QM, we can always accommodate PI in CM by loosening our notion of property. But even after making such an adjustment, we still encounter the same problems with property independence as we saw in the quantum case—while we can loosen our notion of ‘property’ enough to allow particles in permutation

⁵⁶ We can construct the reduced phase space from the reduced configuration space.

invariant states to have them, there is just no way to respect PI and avoid attributing the same properties to all particles of a type. The points of $(\mathbb{R}^2 \setminus D) / \Pi_{21}$ are sets of points in \mathbb{R}^2 of the form $\{\langle x, y \rangle, \langle y, x \rangle\}$. The two particle roles in such a pair—the parts which individually correspond to a single particle specification—are identical. Thus, if our interpretation function is indeed a function, the specifications attributed to each of the two particles must also be identical.

How do I know what the roles are? It is implicit in PI—the reason we quotiented out by the action of permutation in the first place was because PI tells us that this action must yield models with identical interpretations. Which points in the original configuration space should be identified as equivalent was determined by an implicit choice of particle role, namely projection onto one dimension or the other of the two-dimensional space. So for instance, the role of particle 1 in the model $\langle u, v \rangle$ is just u . In the quotient space, we can extend this notion of projection, but the result is a pair of identical numbers for either particle. The role of particle 1 in $\{\langle x, y \rangle, \langle y, x \rangle\}$ is $\{x, y\}$. So is the role of particle 2.

One might be inclined to complain at this point that I've made a mistake in understanding the quotient space. One might think the right thing to do is to begin our physics with the reduced space and simply decline to interpret the original space from which it was constructed. This would leave us free to specify some other way of extracting particle roles from models (since PI would no longer fix the choice for us) and thus free to interpret our models as systems of particles with whatever properties we wish. But such a move is not open to us. There are only two possibilities: either the new way of extracting and interpreting particle roles also respects PI such that permuting of the new models will yield identical interpretations, or it won't. In the former case, I can simply run my argument again and show that both particles must get the same specification. In the latter case, the theory is no longer permutation invariant and my thesis does not apply.

The upshot is that it is possible to consistently impose PI on classical theories, provided we are willing to adopt a decidedly non-classical view of what it means to bear a property. However, the resulting theory of PI+CM—unlike CM alone—is strictly incompatible with MA.

7.4 Objection: A metric is the wrong measure of distance

In the argument for incompatibility, the triangle inequality for metrics played an essential role—without this property of metrics my argument does not go through. It might be objected then that I have not sufficiently motivated the use

of a metric to measure the distance between states.⁵⁷ Alternatively, one might claim that some other distance measure—one that fails to respect the triangle inequality—is a more appropriate choice. My response is two-fold: (i) there are good reasons for supposing a full-blown metric to be the most natural choice for the sort of distance measure required and (ii) any alternative measure that allows one to circumvent the incompatibility argument will have some pathological properties that are difficult to justify.

First let's consider the positive reasons for adopting a metric as a measure of distance among particle specifications. Recall that specifications are supposed to be descriptions of the actual physical properties borne by a particle, with no excess formal structure or irrelevant assertions getting in the way—specifications in some sense directly correspond to particles and their properties. It is an implicit premise of physics and an explicit axiom of formal measurement theory that physical quantities—taken one at a time—are representable by the real numbers.⁵⁸ For instance, measurements of mass, the linear components of momentum, kinetic energy, volume, etc. are all represented by real numbers. Furthermore, the standard topology of the real number line is metrizable. Physics generally assume the standard metric, $d(x, y) = |x - y|$, for assessing the difference or distance between two values of a physical quantity like mass. In fact, values of this distance measure itself are often taken to carry ontological implications—the distance between two physical quantities is itself taken to represent a physical quantity. For instance, the difference between two potential energies often represents a quantity of kinetic energy. The point is that when taken one at a time physical quantities—the things directly referred to in specifications—are assumed to be representable by the real numbers equipped with a natural metric structure. This much we inherit from physics itself. So, if our specifications made reference to only a single, empirically accessible physical quantity, it would be uncontroversial to measure the distance between specifications by using a metric over the real numbers representing the value of the physical quantity to which the specification refers. When we consider multiple physical quantities, we have in effect a vector space, a collection of n -tuples of real numbers for which the operation of addition is already defined and for which there is an obvious choice of inner product. A vector space is also metrizable, and thus it seems natural to extend the use of metrics to measure the distance between bundles of physical quantities, each represented by real numbers.

⁵⁷ This objection was suggested to me by Edward Dean and Ruth Poproski in private discussion.

⁵⁸ For an overview, see (Falkenburg 2007), Appendix A. For more thorough treatments of axiomatic measurement theory, see (Narens 2007).

So far, I have only made reference to empirically accessible physical properties. Perhaps the problem is with metaphysical properties for which there is no accepted representation schema. The particle specifications I have been considering need only invoke one such property, namely the identity of each particle. With respect to comparing physical quantities associated with particles however, this seems no more problematic than asserting the identity of measurements in order to talk of sets of measurements. The burden is on the critic at this point to explain why the use of a metric is inappropriate in this case.

If one does not buy the above ‘naturalness’ argument in favor of metrics, then it is still the case that the sort of distance measure required to defeat my argument is profoundly ‘unnatural’. I assume that the first two criteria for a metric are unobjectionable for anything we might accept as a distance measure, namely that it be positive definite and symmetric. If we retain these criteria and look for a distance function that violates the triangle inequality in such a way that my argument fails, the only candidates are distance measures with pathological properties. To be more precise, consider a distance measure $d(x, y)$ on a set X with the following properties:

- (1) $d(x, y) \geq 0$ for all $x, y \in X$ and $d(x, y) = 0 \leftrightarrow x = y$
- (2) $d(x, y) = d(y, x)$ for all $x, y \in X$

Such a measure is often called a ‘semi-metric’ in the literature (Ceder 1961; Wilson 1931). In order to defeat the argument of the preceding chapter, it must also be the case that:

- (3) For all $x, y \in X$ such that $d(x, y) > 2\epsilon$ for any $\epsilon > 0$, there exists a $z \in X$ such that $d(x, z) \leq \epsilon$ and $d(y, z) \leq \epsilon$.

It is easy to show that such a distance function exists. For instance, if we let $X = \mathbb{R}^+$, then the following satisfies (1) – (3):

$$d(x, y) = \begin{cases} 0 & \text{if } x = y \\ \frac{1}{xy} & \text{otherwise} \end{cases}$$

This semi-metric induces a topology⁵⁹ on \mathbb{R}^+ that has some bizarre properties. For one, it is not Hausdorff since for every pair of distinct points x and y , every neighborhood $B_\epsilon(x)$ intersects every neighborhood $B_\delta(y)$. This is because

⁵⁹ It is easy to prove that a basis is provided by the open balls of $d(x, y)$. These are the sets $B_\epsilon(x) = \{y \in S \mid d(x, y) < \epsilon\}$ for which $\epsilon > 0$. Proof sketch: For each $x \in S$, $x \in B_\epsilon(x)$ for all ϵ , so it only remains to show that $B_\gamma(x) \subset \{B_\alpha(y) \cap B_\beta(z)\}$. To do so, suppose $x \in B_\alpha(y)$ and $x \in B_\beta(z)$ for some α, β . This means that $d(x, y) = \frac{1}{xy} < \alpha$ and $d(x, z) = \frac{1}{xz} < \beta$. Introduce a new constant γ such that $\frac{1}{x^2} > \gamma > 0$, and consider the open ball $B_\gamma(x)$. Then for every $p \in B_\gamma(x)$, it’s the case that $d(p, y) = \frac{1}{py} = \left(\frac{x}{y}\right) \left(\frac{1}{xp}\right) < (\alpha x^2) \left(\frac{1}{xp}\right) < (\alpha x^2) \gamma < \alpha$. Similarly, $d(p, z) < \beta$. Thus, by definition $p \in B_\alpha(y)$ and $p \in B_\beta(z)$. Therefore, $B_\gamma(x) \subset \{B_\alpha(y) \cap B_\beta(z)\}$.

property (3) entails that no matter how close to 0 we make ϵ and δ there is some point z such that $d(x, z) < \min(\epsilon, \delta)$ and $d(y, z) < \min(\epsilon, \delta)$. Even more bizarre, for any pair of distinct points, there is a sequence that converges to both.⁶⁰ This last property is intolerable, and will hold for any semi-metric that satisfies property (3). It means that there is a sequence of manipulations we could perform that makes the state of a particle more and more like some other state only to find that, when we have transformed the particle state sufficiently, it is indistinguishable from both the target state and from *any other possible state* no matter how different. This is, at the very least, a counter-intuitive property that militates against dropping the triangle inequality.

In summary, a metric is the obvious choice of distance measure, and is almost forced upon us by the partial interpretations assumed in the practice of physics. The only alternative measures that can do the job are semi-metrics with some odd properties that lead to unacceptable consequences. Thus, the use of a metric is the only viable option for measuring the distance between specifications.

7.5 Objection: The result is old news

Given that a number of authors have previously claimed that quantum ‘particles’ aren’t particles, it might be said that I have merely contributed a fresh argument in support of a foregone conclusion. To some extent, this objection was dealt with in Chapter 3. I am reintroducing the concern here in order to stress two points about the argument made in Chapter 7. First, every one of the previous arguments against quantum particles relies on a much stronger account than MA. Because many of these arguments rest on conflicting accounts of ‘particle’, they do not in fact establish the same conclusion. My argument concerns a very weak account of particle—one that captures all of the significant accounts to be found in physics and philosophy. In fact, the notion of particle captured by MA is so weak as to include within its scope a great many accounts not yet formulated. All of these arguments together exhaust the resources of the realist to account for EDiv in anything remotely resembling an ontology of particles. My

⁶⁰ Recall that a sequence p_1, p_2, \dots converges to x just if, corresponding to every neighborhood of x (in this case, every $B_\epsilon(x), \epsilon > 0$) there is a positive integer N such that $p_n \in B_\epsilon(x)$ for all $n > N$. Proof sketch: Consider the distinct points x and y . Consider the neighborhoods $B_{\delta_1}(x)$ and $B_{\epsilon_1}(y)$. Let the first point in our sequence p_1 be a point within $\min(\delta_1, \epsilon_1)$ of both x and y . The existence of such a point is guaranteed by property (3). Let the next point in the series p_2 be a point within $\min(\delta_2, \epsilon_2)$ of both x and y , where $\delta_2 < \delta_1$ and $\epsilon_2 < \epsilon_1$. We can carry out this procedure indefinitely. The result is a sequence p_1, p_2, \dots with the property that for any neighborhood $B_\delta(x)$ of x or $B_\epsilon(y)$ of y , there is some N such that x_n is within $B_\delta(x)$ and $B_\epsilon(y)$ for all $n \geq N$. In fact, this N is just the first index for which $\delta_N \leq \delta$ and $\epsilon_N \leq \epsilon$ in our construction.

argument is therefore a more general condemnation of particle interpretations than heretofore offered.

Second, my argument is not framed within any specific theory, though it can be cast in such a manner. While previous authors have given reason to think that there are no particles in one sense or another in QM or QFT, none have drawn out the theory-independent consequences of PI. Because my argument applies to *any* theory in which PI can be stated and is presumed true, it precludes particle interpretations for any successors to QM or QFT that retain a version of PI—even quantum gravity should it incorporate the principle. What’s more, insofar as chemistry and molecular biology can be said to ‘reduce’ to quantum physics, they inherit the same difficulty. Chemistry cannot be about atoms if PI holds in the quantum domain. That’s not old news.

7.6 Objection: MA is the wrong account of particles

Finally, one might object that MA fails to capture the correct notion of ‘particle’. This, I suggest is simply beside the point. I am not interested in cataloguing the various ways in which the term is used or in analyzing the intended meaning in each such context. I have given reason to think that MA captures the sense of ‘particle’ as it has been used in every major scientific interpretation to employ that term or various related terms such as ‘atom’. Even if MA fails to capture a few historical notions of particle, it nonetheless captures a very broad class of ontologies that bear special appeal for the scientific realist. This appeal derives from the ability to account for EDiv. Whatever term one prefers to use, if PI precludes all ontologies compatible with MA, then a great many ontologies of interest to scientists and philosophers must be abandoned. Such a result is interesting in its own right.

Chapter 8 A viable object ontology

8.1 Overview

We have seen that particle ontologies are strictly incompatible with theories containing PI. In light of this negative result, the realist must seek an alternate ontology for QM, its plausible successors, and any theory like chemistry that both references particles and ostensibly reduces to QM. That QM needs a new interpretation is hardly a novel claim. It is novel, however, to insist that the most appropriate alternatives may be semi-classical object ontologies rather than radically eliminativist ones. It is my aim in this chapter to sketch one such interpretation that avoids conflict with PI but nonetheless underwrites EDiv.

The interpretation in question—which I’ll call the ‘Spatial Objects Interpretation’ or SOI—serves as a proof of principle. It jettisons particles while nonetheless retaining objects in the sense of things that bear properties, stand in identity relations with one another, and are sufficiently independent of one another to underwrite EDiv. The fact that one can construct a consistent ontology of objects for QM undercuts the motivation for radically revisionist ontologies such as ‘ontic structural realism’ (OSR) which purports to do away with objects altogether. In the next chapter, I’ll take up the positive arguments in favor of OSR, and provide reasons to reject this sort of realist move. Here, I construct an interpretation in terms of classical objects in order to demonstrate that a move to OSR is not necessitated by my incompatibility result.

In outline, SOI asserts that the basic objects of which the universe is composed are regions of space (in the case of non-relativistic QM).⁶¹ That is, they are measurable subsets⁶² of \mathbb{R}^3 . Each such spatial region bears properties in the form of probability distributions over particular, observable property values much as we imagined quantum particles bearing properties in Chapter 3. Some of these properties are peculiar in that they always manifest definite values and are strict-

⁶¹ W. V. O. Quine (1976) fleetingly considered such a possibility before replacing it with an abstract ontology of mathematical objects. In his essay, Quine correctly notes that taking spatial regions as objects as I do in SOI rather than spacetime regions is problematic. However, the generalization of SOI to spacetime regions appears straightforward, and so I have chosen the simpler, non-relativistic setting in which to lay out the approach.

⁶² For our purposes, we can assume they are connected subsets such as most of us intuitively think of when we think of a spatial region (e.g. the computer shaped region in which my laptop can be found or the interior of my desk drawer). However, this restriction is unnecessary.

ly co-occurring—the value of one such special property is always associated with a particular value of the other properties in this special set. So, for instance, if a given volume possess a value of electric charge equal to $-e$ then it must also possess a mass with a definite value of m_e (the mass of the electron). These bundles of co-occurring properties are just the ‘state-independent’ properties associated with particles in the old MA view. I’ll continue to refer to sets of co-occurring property values as state-independent properties.

The principal motivation for considering SOI is to establish the possibility of an object ontology that doesn’t run afoul of PI. As it turns out, SOI is not only compatible with PI, it actually entails a stronger restriction. Before I derive this important feature, however, I want to motivate SOI on grounds independent of our concern with PI.

8.2 Privileging the wavefunction

There is one principal reason to view wavefunctions as privileged representations of quantum systems: every actual measurement of a quantum system involves some finite spatial volume. The spatial degree of freedom is always present, and ignoring this degree of freedom is always a strong idealization. We may talk about systems with only a spin degree of freedom, but to predict measurement outcomes requires one to take into account the wavefunction and the geometry of the measuring apparatus. The idea behind SOI is simply to take this fact seriously—any viable representation of a quantum system must account for the spatial nature of that system.

To get a sense for SOI, let us first consider an analogy with fluid dynamics. There are two basic pictures of fluids: the *Lagrangian* and the *Eulerian* (Panton 1996, Ch 4). In the Lagrangian view, the focus is on individual particles in the fluid. The primary variables in terms of which fluid behavior is described are particle positions \mathbf{r}_i and velocities \mathbf{v}_i which are functions of time t . Fluid flow is conceived in terms of the paths followed by the particles comprising the fluid (see the left panel of Figure 8.1). In the Eulerian view, we focus instead on individual positions in space labeled by coordinates \mathbf{r}_i and chart their changing properties through time. The primary variables by which fluid behavior is described in this view are local properties of spatial locations. Thus, fluid velocity and pressure are functions of both position and time, e.g. $\mathbf{v}(\mathbf{r}_i, t)$ (see the instantaneous velocity field shown on the right of Figure 8.1).

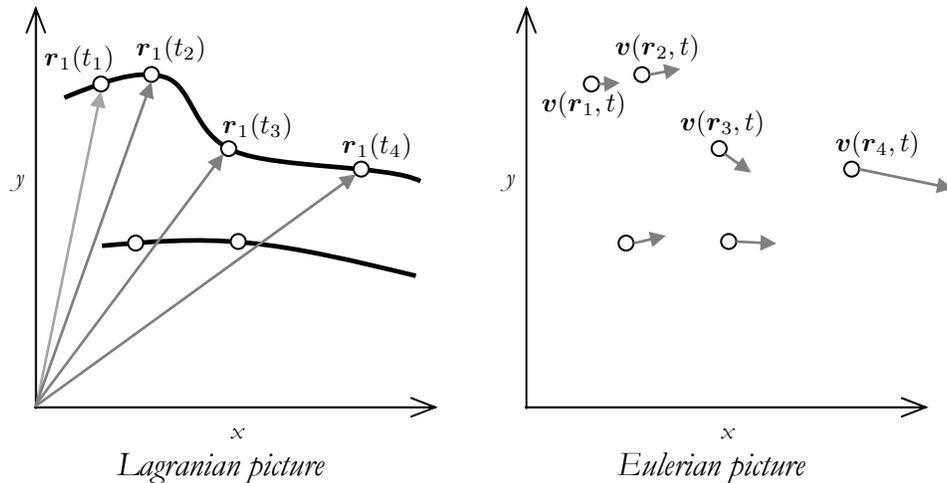


Figure 8.1 Contrasting views of fluid dynamics. The left panel corresponds to the Lagrangian view and shows the changing position of a particle at four different times as it moves through the fluid. The right panel corresponds to the Eulerian view, and depicts the value of the fluid velocity at each of four spatial locations at the same moment of time.

The standard view of QM is akin to the Lagrangian perspective: the basic elements are particles and, though they cannot be attributed well-defined paths in space, kinematics is concerned with the changes in properties of these particles through time. In this view, the wavefunction is just one description of the state of a particle, and position is a property the particle can possess.⁶³

In the interpretation I am proposing, spatial regions are the fundamental elements—what I’ll identify as the objects. Kinematics is described by changing property distributions over each spatial region, much as in the Eulerian approach to fluid dynamics. In SOI, the wavefunction encodes the properties of all regions of space over time. It is akin to the Eulerian specification of fluid pressure and velocity at all points.

8.3 Spatial regions as objects

So far, I have only provided a vague image of SOI. To make it precise, I will state the postulates of QM in terms of the proposed account⁶⁴, demonstrate that these modified postulates coincide with the standard interpretation vis-à-vis observable predictions, and then show how SOI avoids a contradiction with PI.

⁶³ In the standard view, position is a property in the sense that the particle bears a determinate probability distribution over the possible outcomes of position measurements.

⁶⁴ In stating the postulates, I am following the presentation of (Cohen-Tannoudji, Diu, and Lalòe 1977).

First Postulate:

At a fixed time t_0 the joint physical state of every spatial region for a ‘system’ of n sets of state-independent property values is represented by a function $\psi(\mathbf{x}_1, \mathbf{x}_1, \dots, \mathbf{x}_n, t_0)$ belonging to the space \mathcal{L}^2 of square-integrable functions on \mathbb{R}^{3n} . These functions are presumed to be normalized such that $\int dx_1^3 \int dx_2^3 \cdots \int dx_n^3 \psi^*(\mathbf{x}_1, \mathbf{x}_1, \dots, \mathbf{x}_n, t_0) \psi(\mathbf{x}_1, \mathbf{x}_1, \dots, \mathbf{x}_n, t_0) = 1$.

Second Postulate:

Every measurable physical quantity A is represented by an operator \hat{A} acting on $\mathcal{L}^2(\mathbb{R}^3)$. The operator \hat{A} is called an *observable*. State-independent properties are represented by operators proportional to the identity operator.

Third Postulate:

The only possible result of the measurement of a physical quantity A in a particular spatial region \mathcal{V} is one of the eigenvalues of the corresponding observable \hat{A} .

Fourth Postulate:

Case of a discrete non-degenerate spectrum:

When the physical quantities A_1, A_2, \dots, A_n (not necessarily unique) are measured for (possibly overlapping) spatial regions $\mathcal{V}_1, \mathcal{V}_2, \dots, \mathcal{V}_n$ respectively, the probability $\Pr(a_{1,i_1}(\mathcal{V}_1), \dots, a_{n,i_n}(\mathcal{V}_n); t)$ of obtaining the non-degenerate eigenvalues $a_{1,i_1}, a_{2,i_2}, \dots, a_{n,i_n}$ of the corresponding observables $\hat{A}_1, \hat{A}_2, \dots, \hat{A}_n$ at time t is given by:

$$\begin{aligned} \Pr(a_{1,i_1}(\mathcal{V}_1), \dots, a_{n,i_n}(\mathcal{V}_n); t) = \\ \int_{V_1} dx_1^3 \int_{V_1} dx_1'^3 \cdots \int_{V_n} dx_n^3 \int_{V_n} dx_n'^3 a_{1,i_1}^*(\mathbf{x}_1) a_{1,i_1}(\mathbf{x}'_1) \cdots a_{n,i_n}^*(\mathbf{x}_n) a_{n,i_n}(\mathbf{x}'_n) \\ \cdot \psi^*(\mathbf{x}'_1, \mathbf{x}'_2, \dots, \mathbf{x}'_n, t) \psi(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n, t) \end{aligned} \quad (8.3.1)$$

Each function $a_{j,i_j}(\mathbf{x})$ is an eigenfunction of the operator \hat{A}_j with eigenvalue a_{j,i_j} .

Case of a continuous spectrum:

When the physical quantities A_1, A_2, \dots, A_n are measured for spatial regions $\mathcal{V}_1, \mathcal{V}_2, \dots, \mathcal{V}_n$ respectively, the *probability density* $f(a_1(\mathcal{V}_1), \dots, a_n(\mathcal{V}_n); t)$ of obtaining results in the intervals $a_1 + da_1, a_2 + da_2, \dots, a_n + da_n$ for the corresponding observables at time t is

$$\begin{aligned}
f(a_1(\mathcal{V}_1), \dots, a_n(\mathcal{V}_n); t) = & \\
& \int_{V_1} dx_1^3 \int_{V_1} dx_1'^3 \cdots \int_{V_n} dx_n^3 \int_{V_n} dx_n'^3 a_1^*(\mathbf{x}_1) a_1(\mathbf{x}_1') \cdots a_n^*(\mathbf{x}_n) a_n(\mathbf{x}_n') \\
& \cdot \psi^*(\mathbf{x}_1', \mathbf{x}_2', \dots, \mathbf{x}_n', t) \psi(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n, t)
\end{aligned} \tag{8.3.2}$$

One straightforward consequence of the postulates stated in this way concerns the state-independent properties. As I indicated above, these include bundles of properties like mass and charge that co-occur with definite, fixed values. For instance, in some systems whenever a mass of m_e (the mass of an electron) is detected, it is always possible to simultaneously measure a charge $-e$ (the charge of an electron). To say that there is more than one such bundle available to a quantum system is to say that it is possible to simultaneously measure bundles of physical quantities in disjoint volumes, e.g. to measure m_e and $-e$ in one volume and m_p and $+e$ in another. According to the Second Postulate, each of these physical quantities like mass is represented by a scalar multiple of the identity operator \hat{I} in $\mathcal{L}^2(\mathbb{R}^3)$. Every wavefunction is thus an eigenfunction of each of these operators with an eigenvalue corresponding to one of the state-independent property values (e.g. m_e). For simplicity, suppose that mass is the only state-independent property. Then if there are n different masses that can be measured simultaneously, the probability of finding m_1 in $\mathcal{V}_1, \dots, m_n$ in \mathcal{V}_n is given by:

$$\begin{aligned}
\text{Pr}(m_1(\mathcal{V}_1), \dots, m_n(\mathcal{V}_n); t) = & \\
& \int_{\mathcal{V}_1} dx_1^3 \cdots \int_{\mathcal{V}_n} dx_n^3 \psi^*(\mathbf{x}_1, \dots, \mathbf{x}_n, t) \psi(\mathbf{x}_1, \dots, \mathbf{x}_n, t)
\end{aligned} \tag{8.3.3}$$

It is worth stressing a distinction. In the standard view, Equation (8.3.3) is interpreted as the probability that particle 1 with mass m_1 is found in the region \mathcal{V}_1 , and the particle 2 with mass m_2 is found in the region \mathcal{V}_2 , etc. To be more explicit, it represents the probability that the *object* particle 1 has as *properties* a mass of m_1 and a location in the region \mathcal{V}_1 , and so on. According to SOI, this should instead be read as the probability that region \mathcal{V}_1 has a mass m_1 , region \mathcal{V}_2 a mass m_2 , etc. Put another way, the right-hand side (8.3.3) represents the probability that the *objects* $\mathcal{V}_1, \mathcal{V}_2, \dots, \mathcal{V}_n$ are found to possess *properties* m_1, m_2, \dots, m_n respectively.

SOI eliminates an inconsistency in the standard interpretation. In the standard interpretation, how we interpret the expression (8.3.3) depends on whether the particles are of the same type. If they are of n different types, then it represents

the probability of finding particle 1 to have a position in region \mathcal{V}_1 , etc. If the particles are of the same type then PI is in play, and we are led to interpret (8.3.3) via such circumlocutions as “the probability that a particle has a position in region \mathcal{V}_1 , with no fact of the matter as to which it is.” In SOI, the interpretation is uniform. If all of the bundles of state-independent properties are identical, we still read (8.3.3) as representing the probability that each of the regions $\mathcal{V}_1, \mathcal{V}_2, \dots, \mathcal{V}_n$ is found to manifest the indicated property value. It just so happens in this case that $m_1 = m_2 = \dots = m_n$. That is, the property values observed for each region are the same, rather like n coins coming up heads when flipped.

It will help at this point to illustrate the differences between the standard interpretation and SOI by way of a concrete example. Consider the case in which—in the standard interpretation—there is only one particle with a one-dimensional wavefunction given by:

$$\psi(x, t) = \frac{1}{\sqrt{\alpha\hbar(1 + \frac{it}{m\hbar\alpha^2})}\sqrt{\pi}} e^{i(p_0x - p_0^2t/2m)/\hbar} e^{-\frac{(x - p_0t/m)^2}{2\alpha^2\hbar^2(1 + \frac{it}{m\hbar\alpha^2})}} \quad (8.3.4)$$

This is a ‘Gaussian wavepacket’. In the standard interpretation, this wavefunction encodes properties—or rather distributions over property values—pertaining to the particle. If we choose units such that $\hbar = c = 1$ and consider the case when $m = 1$ eV, $p_0 = 1$ eV, $\alpha = 1$ eV⁻¹, then the probability density function for finding the particle between x and $x + dx$ in the standard view is that shown in Figure 8.2. The peak of the probability density corresponds in this case to the expectation value of position which changes linearly with time (specifically, $\langle x \rangle = t$). This describes a particle moving with uniform velocity (and thus constant mean momentum) to the right.

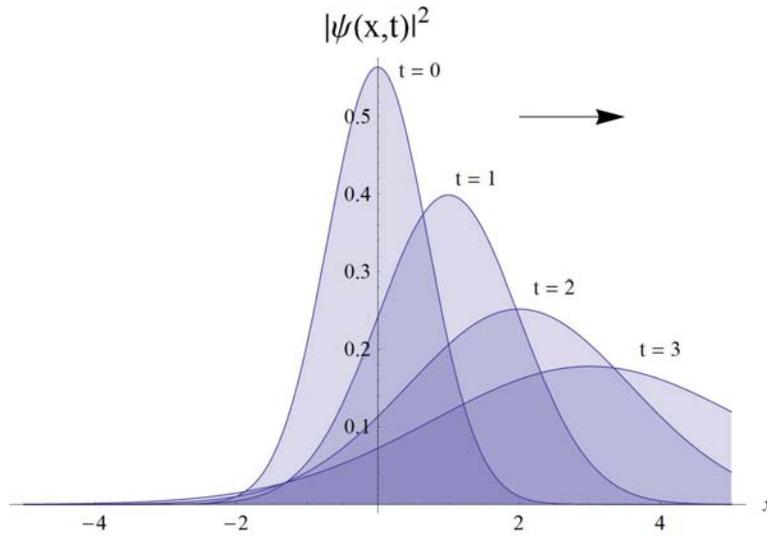


Figure 8.2 A moving particle in the standard interpretation. A plot of the square modulus of wave function for $t = 0, 1, 2,$ and 3 sec. In the standard interpretation, the square modulus is equivalent to the probability density for finding the particle between x and $x + dx$. The sequence shown here represents successive probability densities for a particle moving with constant velocity to the right.

According to SOI, however, there is no particle. There is a mass property that can have a measured value of 0 or m for any given region, and there is a momentum property that can manifest a continuum of values at one instant of time for any given region. In Figure 8.3, the probability of detecting mass m as a function of time is shown for five contiguous spatial regions. The ‘mass property’—really a distribution over the two possible mass property values—is different for each of these regions, and changes in time for each.

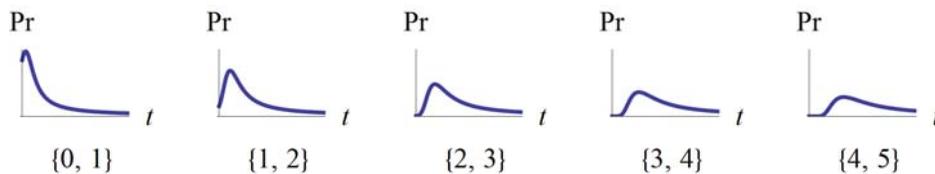


Figure 8.3 The SOI interpretation of the ‘moving’ Gaussian wavepacket. Each frame depicts the probability as a function of time of measuring a mass m within the region indicated below each plot (regions are given as closed intervals on the real line).

Figure 8.4 gives another way of visualizing the mass property for those same contiguous spatial regions. In that plot, the single spatial dimension is shown on the left horizontal axis (labeled ‘ x ’), and time progresses to the right along the other horizontal axis (labeled ‘ t ’). The vertical axis indicates the probability of detecting a mass m in a given region. The five ribbons indicate the respective probabilities as a function of time for each of the five regions. However it is vi-

sualized, the fundamental feature of SOI is that properties are always and only attributed to spatial regions—there are no particles.

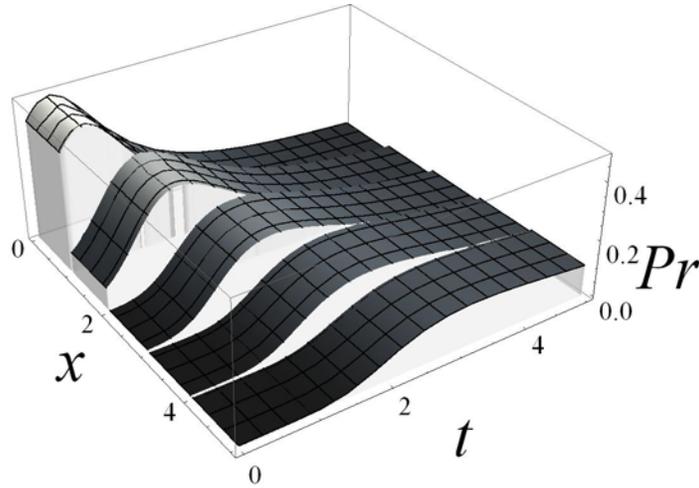


Figure 8.4 An alternate view of the moving wavepacket. In this figure, the x axis represents a one-dimensional space. Each ribbon correspond to one of the spatial regions shown in Figure 8.3. The height of the ribbon indicates the probability of detecting a mass m in the region corresponding to that ribbon. The t axis represents time. Notice that the peak probability occurs later and later for neighboring spatial regions as we move in the positive x -direction.

I intend SOI to be an interpretation of QM, not a modification of the physics. While the postulates above are formally different from those given by, say von Neumann, I claim that SOI leads to exactly the same empirical predictions as the standard interpretation of QM. It differs only as to which formalism for presenting QM is favored, and as to how the parts of that formalism are taken to refer to unobservable things in the world. Rather than attempt a systematic translation of SOI expressions to standard expressions, it will suffice to show that key standard expressions for measurable quantities can be recovered from SOI. To do so, it will help to make use of von Neumann’s formalism via some straightforward identifications. First, we define the *one-place state* corresponding to region \mathcal{V} as follows:⁶⁵

$$|\psi_{\mathcal{V}}\rangle \equiv \int_{\mathcal{V}} dx^3 \psi(\mathbf{x})|\mathbf{x}\rangle \quad (8.3.5)$$

In the standard interpretation, a one-place state is the projection of one-particle state onto the subspace spanned by the position eigenkets with eigenvalues in \mathcal{V} .⁶⁶ The square of modulus of $|\psi_{\mathcal{V}}\rangle$ is given by:

⁶⁵ In the remainder of this section, I suppress the time dependence in all expressions.

⁶⁶ Notice that if the one-place state is given for two contiguous spatial regions \mathcal{V}_1 and \mathcal{V}_2 then the state of the overall region $\mathcal{V}_1 \cup \mathcal{V}_2$ is fully determined. A similar property holds for n -place states. This is in contrast to the spacetime states posited by Wallace and Timpson

$$\begin{aligned}
|\langle \psi_{\mathcal{V}} | \psi_{\mathcal{V}} \rangle|^2 &= \int_{\mathcal{V}} dx'^3 \int_{\mathcal{V}} dx^3 \psi^*(\mathbf{x}') \psi(\mathbf{x}) \langle \mathbf{x}' | \mathbf{x} \rangle \\
&= \int_{\mathcal{V}} dx'^3 \int_{\mathcal{V}} dx^3 \psi^*(\mathbf{x}') \psi(\mathbf{x}) \delta(\mathbf{x}' - \mathbf{x}) \quad (8.3.6) \\
&= \int_{\mathcal{V}} dx^3 \psi^*(\mathbf{x}) \psi(\mathbf{x})
\end{aligned}$$

In the standard interpretation, the right-hand side of Equation (8.3.6) is just the probability of finding the particle in the region \mathcal{V} , while according to SOI it is the probability of finding the region \mathcal{V} to possess mass m . According to both views, this one expression is used to compute the mean mass density of the given spatial region.

Similarly, one can define *n-place states* by projecting a state vector traditionally identified with a many-particle system onto multiple spatial regions:

$$\begin{aligned}
|\psi_{\mathcal{V}_1, \mathcal{V}_2, \dots, \mathcal{V}_n}\rangle &= \\
\int_{\mathcal{V}_1} dx_1^3 \int_{\mathcal{V}_2} dx_2^3 \cdots \int_{\mathcal{V}_n} dx_n^3 \psi(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n) |\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\rangle \quad (8.3.7)
\end{aligned}$$

From Equations (8.3.5) and (8.3.7), one can see that the expectation values derived from the distributions stipulated in the Fourth Postulate are just the standard expectation values when the spatial region of interest is taken to be all of space. So for instance, if the world is described by a one-place state, then the expectation value for a generic operator is given by:⁶⁷

$$\begin{aligned}
\langle \hat{A} \rangle &= \int da \int dx'^3 \int dx^3 a \langle a | \mathbf{x} \rangle \langle \mathbf{x} | \psi \rangle \langle \psi | \mathbf{x}' \rangle \langle \mathbf{x}' | a \rangle \\
&= \int da \langle a | \psi \rangle \langle \psi | a \rangle \\
&= \int da \langle \psi | (a | a \rangle \langle a |) | \psi \rangle \quad (8.3.8) \\
&= \int da \langle \psi | (\hat{A} | a \rangle \langle a |) | \psi \rangle \\
&= \langle \psi | \hat{A} | \psi \rangle
\end{aligned}$$

Finally, the SOI approach can be expanded to incorporate additional degrees of freedom such as spin in the same way they are incorporated into the wavefunc-

(2009). While their proposal is similar in theme to SOI (and unlike SOI, is properly relativistic), their use of reduced density operators means that the state of the union of two space-time regions is not determined by the states of the separate regions.

⁶⁷ This is the discrete case. The continuous case is similar.

tion representation in the standard approach. For instance, we can include a single spin-1/2 property using the standard spinors,

$$\chi \equiv \begin{pmatrix} \alpha \\ \beta \end{pmatrix},$$

where $|\alpha|^2 + |\beta|^2 = 1$. A quantum system is represented by an overall wavefunction of the form $\psi(\mathbf{x})\chi$. If we are working in a basis in which the spin operator \hat{S}_z corresponding to spin along the z-axis is diagonal, then

$$\hat{S}_z = \frac{\hbar}{2} \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$$

This operator has two eigenvectors,

$$\chi^{(+)} = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \text{ and } \begin{pmatrix} 0 \\ 1 \end{pmatrix},$$

with eigenvalues $\pm\hbar/2$.

In effect, this additional degree of freedom is accommodated by working in an expanded vector space. If we let \mathcal{E} stand for the Hilbert space of suitable wavefunctions in $\mathcal{L}^2(\mathbb{R}^3)$ and \mathcal{E}_χ stand for the Hilbert space of spinors, then in order to accommodate spin, we need to work in the outer-product space $\mathcal{E} \otimes \mathcal{E}_\chi$. The Fourth Postulate then needs to be amended as follows:

Amended Fourth Postulate:

Case of a discrete non-degenerate spectrum:

When the physical quantities A_1, A_2, \dots, A_n (not necessarily unique) and spin-quantities Q_1, Q_2, \dots, Q_n are measured for spatial regions $\mathcal{V}_1, \mathcal{V}_2, \dots, \mathcal{V}_n$ respectively, the probability of obtaining the non-degenerate eigenvalues $a_{1,i_1}, a_{2,i_2}, \dots, a_{n,i_n}$ of the corresponding observables $\hat{A}_1, \hat{A}_2, \dots, \hat{A}_n$ and the non-degenerate eigenvalues $q_{1,i_1}, q_{2,i_2}, \dots, q_{n,i_n}$ of the spin observables $\hat{Q}_1, \hat{Q}_2, \dots, \hat{Q}_n$ is

$$\begin{aligned} & \Pr(a_{1,i_1}(\mathcal{V}_1) \wedge q_{1,i_1}(\mathcal{V}_1), \dots, a_{n,i_n}(\mathcal{V}_n) \wedge q_{n,i_n}(\mathcal{V}_n)) = \\ & \int_{\mathcal{V}_1} dx_1^3 \int_{\mathcal{V}_1} dx_1'^3 \cdots \int_{\mathcal{V}_n} dx_n^3 \int_{\mathcal{V}_n} dx_n'^3 a_{1,i_1}^*(\mathbf{x}_1) a_{1,i_1}(\mathbf{x}'_1) \cdots a_{n,i_n}^*(\mathbf{x}_n) a_{n,i_n}(\mathbf{x}'_n) \\ & \quad \cdot (\xi_{1,i_1}^* \chi_1) \cdots (\xi_{n,i_n}^* \chi_n) \psi^*(\mathbf{x}'_1, \dots, \mathbf{x}'_n) \psi(\mathbf{x}_1, \dots, \mathbf{x}_n) \end{aligned} \tag{8.3.9}$$

Each function $a_{j,i_j}(\mathbf{x})$ is an eigenfunction of the operator \hat{A}_j with eigenvalue a_{j,i_j} . Each ξ_{j,i_j} is an eigenvector of the operator \hat{Q}_j with eigenvalue q_{j,i_j} .

From the Amended Fourth Postulate we can see that the expectation value of spin in the \hat{z} -direction for an arbitrary wavefunction $\psi(\mathbf{x})\chi$ is thus

$$\begin{aligned}
\langle \hat{S}_z \rangle &= \int_{\mathcal{V}} dx'^3 \int_{\mathcal{V}} dx^3 \left(\frac{\hbar}{2} \chi^{(+)*} \chi + \frac{\hbar}{2} \chi^{(-)*} \chi \right) \psi^*(\mathbf{x}') \psi(\mathbf{x}) \\
&= \int_{\mathcal{V}} dx^3 \left(\frac{\hbar}{2} \chi^{(+)*} \chi + \frac{\hbar}{2} \chi^{(-)*} \chi \right) \psi^*(\mathbf{x}) \psi(\mathbf{x})
\end{aligned} \tag{8.3.10}$$

It is worth noting that in standard quantum texts, the spatial degree of freedom is often ignored and spin, for instance, is taken to be the only relevant degree of freedom in various toy models. According to SOI, this is equivalent to taking the regions of interest to include all of space. In the limit where \mathcal{V} encompasses all space, Equation (8.3.10) reduces to

$$\langle \hat{S}_z \rangle = \left(\frac{\hbar}{2} \chi^{(+)*} \chi + \frac{\hbar}{2} \chi^{(-)*} \chi \right)$$

which is just the standard expression for expected spin along the z-axis.

8.4 A clarification: single regions and multiple values

There is an important ambiguity to clear up at this stage, one which highlights the inelegance of the mathematical representation of QM I've chosen for formulating SOI, but which does not affect the philosophical point in question. Suppose that mass is the only state-independent property and that just two mass values can manifest simultaneously (i.e. the world is described by a two-place state). According to Equation (8.3.3) the probability that spatial region \mathcal{V}_1 has a mass of m_1 and \mathcal{V}_2 has mass m_2 is given by $\int_{\mathcal{V}_1} dx_1^3 \int_{\mathcal{V}_2} dx_2^3 \psi^*(\mathbf{x}_1, \mathbf{x}_2) \psi(\mathbf{x}_1, \mathbf{x}_2)$. Now suppose that $m_1 = m_2 = m$ and that $\mathcal{V}_1 = \mathcal{V}_2 = \mathcal{V}$. As I stated matters above, we should read the preceding expression as the probability that region \mathcal{V} has a mass of m and a mass of m . This is an ambiguous statement—it might be taken to mean “the mass of \mathcal{V} is m ” or that “the mass of \mathcal{V} is $2m$ ” SOI asserts the latter. The idea is that a single region can manifest integral units of state-dependent properties. Thus a given region can have zero, one, or two units of mass m —the mass of the region might be 0, m , or $2m$. We can extrapolate this interpretation to any physical quantity represented by an operator. If more than one eigenvalue is manifest for a given region, this is to be interpreted in an additive sense—the region bears just one property value equal to the sum of the two repeated values.

I concede that this is not an obvious reading of the two-place wavefunction $\psi(\mathbf{x}_1, \mathbf{x}_2)$. It is, however, a consistent reading. Perhaps a more elegant (and theoretically fruitful) approach would be to assign to each volume a distribution over states in an appropriate Fock space. That is, to each volume one would compute the probability of different sets of ‘occupancy numbers’ where these numbers are understood as representing integral magnitudes for discrete-valued

properties like energy and mass. Developing such a formalism, however, would take us too far astray.

Ultimately, the inelegance of the reading does not affect the philosophical point I want to make. SOI provides a uniform and consistent way of attributing property values to distinct objects. While a formal representation might be constructed which, though isomorphic in some sense to that presented here, is more perspicuous for the purposes of SOI (and perhaps more useful for practical computation), the fact remains that a consistent interpretation of QM in terms of objects can be had despite PI. All that remains to be shown is that SOI is in fact compatible with this postulate.

8.5 SOI and PI

As I suggested above, SOI entails an important consequence in the case of identical bundles of co-occurring properties that is easy to draw out now that we've made these connections with the Hilbert space formalism. If $m_1 = m_2 = \dots m_n$, then it must be the case that

$$\Pr(m_1(\mathcal{V}_1), m_2(\mathcal{V}_2), \dots, m_n(\mathcal{V}_n)) = \Pr(m_{\pi(1)}(\mathcal{V}_1), m_{\pi(2)}(\mathcal{V}_2), \dots, m_{\pi(n)}(\mathcal{V}_n)) \quad (8.4.1)$$

where $\pi : N \mapsto N$ is any permutation on the n indices. From the Second and Fourth Postulate, this means that

$$\begin{aligned} \Pr(m_1(\mathcal{V}_1), \dots, m_n(\mathcal{V}_n)) &= \\ \int_{\mathcal{V}_1} dx_1^3 \int_{\mathcal{V}_1} dx_1'^3 \cdots \int_{\mathcal{V}_n} dx_n^3 \int_{\mathcal{V}_n} dx_n'^3 \psi^*(\mathbf{x}'_1, \dots, \mathbf{x}'_n) \psi(\mathbf{x}_1, \dots, \mathbf{x}_n) &= \\ \Pr(m_{\pi(1)}(\mathcal{V}_1), \dots, m_{\pi(n)}(\mathcal{V}_n)) &= \\ \int_{\mathcal{V}_1} dx_1^3 \int_{\mathcal{V}_1} dx_1'^3 \cdots \int_{\mathcal{V}_n} dx_n^3 \int_{\mathcal{V}_n} dx_n'^3 \psi^*(\mathbf{x}'_{\pi(1)}, \dots, \mathbf{x}'_{\pi(n)}) \psi(\mathbf{x}_{\pi(1)}, \dots, \mathbf{x}_{\pi(n)}) & \end{aligned} \quad (8.4.2)$$

Since the limits of integration are the same, the integrands in both integrals above must be equal. That is,

$$\begin{aligned} \psi^*(\mathbf{x}'_1, \mathbf{x}'_2, \dots, \mathbf{x}'_n) \psi(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n) &= \\ \psi^*(\mathbf{x}'_{\pi(1)}, \mathbf{x}'_{\pi(2)}, \dots, \mathbf{x}'_{\pi(n)}) \psi(\mathbf{x}_{\pi(1)}, \mathbf{x}_{\pi(2)}, \dots, \mathbf{x}_{\pi(n)}) & \end{aligned} \quad (8.4.3)$$

For an arbitrary operator $\hat{O} \equiv \hat{A}_1 \otimes \hat{A}_2 \otimes \cdots \otimes \hat{A}_n$, we have from Equation (8.3.1) and the identity (8.4.3):

$$\begin{aligned}
\langle \hat{O} \rangle &= \sum_{i_1, i_2, \dots, i_n} a_{i_1} a_{i_2} \cdots a_{i_n} \int_{\mathcal{V}_1} dx_1^3 \int_{\mathcal{V}_1} dx_1'^3 \cdots \int_{\mathcal{V}_n} dx_n^3 \int_{\mathcal{V}_n} dx_n'^3 a_{i_1}(\mathbf{x}'_1) \cdots a_{i_n}(\mathbf{x}'_n) \\
&\quad \cdot a_{i_1}^*(\mathbf{x}_1) \cdots a_{i_n}^*(\mathbf{x}_n) \psi^*(\mathbf{x}'_1, \dots, \mathbf{x}'_n) \psi(\mathbf{x}_1, \dots, \mathbf{x}_n) \\
&= \sum_{i_1, i_2, \dots, i_n} a_{i_1} a_{i_2} \cdots a_{i_n} \int_{\mathcal{V}_1} dx_1^3 \int_{\mathcal{V}_1} dx_1'^3 \cdots \int_{\mathcal{V}_n} dx_n^3 \int_{\mathcal{V}_n} dx_n'^3 a_{i_1}(\mathbf{x}'_1) \cdots a_{i_n}(\mathbf{x}'_n) \\
&\quad \cdot a_{i_1}^*(\mathbf{x}_1) \cdots a_{i_n}^*(\mathbf{x}_n) \psi^*(\mathbf{x}'_{\pi(1)}, \dots, \mathbf{x}'_{\pi(n)}) \psi(\mathbf{x}_{\pi(1)}, \dots, \mathbf{x}_{\pi(n)})
\end{aligned} \tag{8.4.4}$$

Equation (8.4.4) holds for any n regions. If we let each \mathcal{V}_i encompass all of space, then Equation (8.4.4) leads to the following:

$$\langle \psi | \hat{O} | \psi \rangle = \langle \psi | \hat{\Pi}^\dagger \hat{O} \hat{\Pi} | \psi \rangle \tag{8.4.5}$$

Here, the operator $\hat{\Pi}$ is the permutation operator which takes $\psi(\mathbf{x}_1, \dots, \mathbf{x}_n)$ to $\psi(\mathbf{x}_{\pi(1)}, \dots, \mathbf{x}_{\pi(n)})$.

At first glance, Equation (8.4.5) is just the first condition for PI. However, in the postulates as stated above, I have implicitly assumed that *any* operator of the form $\hat{O} \equiv \hat{A}_1 \otimes \hat{A}_2 \otimes \cdots \otimes \hat{A}_n$ where each \hat{A}_i is an observable (a self-adjoint operator) on $\mathcal{L}^2(\mathbb{R}^3)$ represents a genuine combination of properties that a collection of spatial regions can manifest. It is thus *not* generally the case that the operator \hat{O} commutes with the permutation operator $\hat{\Pi}$. In order to satisfy Equation (8.4.5) for all permutation operators, the state $|\psi\rangle$ must therefore be invariant up to a phase under permutation. That is, SOI imposes the Symmetrization Postulate (SP) which is stronger than PI.⁶⁸

8.6 SOI and EDiv

To this point, I have shown that SOI retains objects while remaining compatible with PI. I have not yet shown that the resulting interpretation also underwrites EDiv. To do so requires the development of some new conditions. Since SOI does away with particles, the constraint I derived from MA(iii) is inapplicable. What we need is a careful statement of what it means for spatial regions to be largely independent of one another in such a way that knowledge of only a li-

⁶⁸ Actually, to get SP requires the further assumptions that the phase is independent of which permutation operator is in play and that the phase is the same for all states. Even without these extra assumptions, SOI imposes the requirement that permutes lie along the same ray, which is already a stronger constraint than PI.

mitted region is needed to predict the future of a portion of that region. The following condition does the job:

Spatial independence:

For every spatial region \mathcal{V} and for every measurable region $\nu \subset \mathcal{V}$ it is the case that, for most physically possible conditions, the state-dependent properties of ν are approximately independent of the properties of most elements in any given partition of $\nu^c \equiv \mathcal{V} \setminus \nu$ into measurable subsets over a finite interval of time.

Again, we can cash out ‘approximate independence’ in terms of sets of states. To say that one spatial region \mathcal{V} is approximately independent of another region \mathcal{W} , we can appeal to the conditions similar to those introduced in Chapter 4. First, we introduce a metric on the space of descriptions of spatial regions (analogous to the space of particle descriptions). We also introduce a threshold ϵ . For spatial independence to obtain, we then require that there exist three physically possible sets of models of QM—call them α , β , and γ —with the following properties:

- (4) The models in α are interpreted as representing the properties of region \mathcal{V} . Let S_α be the set of all region descriptions extracted from the interpretations corresponding to the models in α . S_α is large with respect to the threshold ϵ .
- (5) The models in β are likewise interpreted as the properties of region \mathcal{W} . Let S_β be the set of all descriptions in the interpretations of the models of β . S_β is also large compared to ϵ .
- (6) The models in γ are interpreted as representing the properties of a spatial region containing as subsets regions \mathcal{V} and \mathcal{W} . For every ordered pair of descriptions in $S_\alpha \times S_\beta$ there exists an interpretation of a model in γ that contains approximately both of these descriptions (that approximately attributes each of these descriptions to the regions \mathcal{V} and \mathcal{W} in the interpretation).

At this point, we face a much more difficult task in proving that SOI satisfies spatial independence than we did in showing that MA (or any other theory) cannot. To obtain a negative result we had only to show that there do not exist any three models satisfying the above conditions. To prove that SOI generally satisfies spatial independence, we first have to stipulate what conditions are physically possible and, for each set of physically possible conditions (e.g. a set

of viable Hamiltonians and some boundary conditions) we would have to work out the set of models which satisfy these conditions. It is not clear that this is a well-posed problem since there appears to be no way to fix a priori the ‘physically possible’ set of conditions. Instead, I offer here only a recipe for constructing such proofs. To do so, I will assume that ‘physically possible’ conditions admit all solutions to the Schrodinger equation subject to the usual conditions of continuity and limits at spatial infinity.

To show that SOI satisfies spatial independence under these conditions, consider the set of possible distributions over single-unit mass properties for a given 1-D region \mathcal{V} , where the probability that \mathcal{V} manifests mass m is denoted q (in the notation used above, $\Pr(m(\mathcal{V})) = q$). I assume that for the metric and threshold chosen these constitute a non-trivial set of states. Choose any $\psi(x)$ such that $\int_{\mathcal{V}} dx \|\psi(x)\|^2 = q$, and $\int_{-\infty}^{+\infty} dx \|\psi(x)\|^2 = 1$, and $\|\psi(x)\|^2 \ll 1$ everywhere in the complement \mathcal{V}^c of \mathcal{V} . We could, for instance, satisfy these conditions by making $\psi(x)$ a tall skinny Gaussian wave-packet with very long but very slender tails. Next, choose a wavefunction $\phi(x)$ that is whatever you like in \mathcal{V}^c , but which vanishes in \mathcal{V} . Obviously, the requirement of continuity imposes some restrictions on $\phi(x)$, but the resulting set of possibilities can easily be made non-trivial. These wavefunctions encode *all* of the distributions pertaining to regions in every possible partition of \mathcal{V}^c . Note that the \mathcal{L}^2 -inner product $(\psi(x), \phi(x)) \approx 0$ (the wavefunctions are almost orthogonal).

All that remains to be shown is that there exists some two-place wavefunction $\Psi(x_1, x_2)$ which entails approximately the same one-place distributions. Let $\Psi(x_1, x_2) = \frac{1}{\sqrt{2}} (\psi(x_1)\phi(x_2) - \psi(x_2)\phi(x_1))$. Consider the region \mathcal{V} and any region $\mathcal{W} \subseteq \mathcal{V}^c$. Then

$$\begin{aligned} \Pr(m(\mathcal{V})) &= \Pr(m(\mathcal{V}), m(\mathcal{W})) + \Pr(m(\mathcal{V}), \neg m(\mathcal{W})) \\ &= \Pr(m(\mathcal{V}), m(\mathcal{W})) + \Pr(m(\mathcal{V}), m(\mathcal{V})) + \Pr(m(\mathcal{V}), m(\mathcal{W}^c \setminus \mathcal{V})) \end{aligned} \tag{8.6.1}$$

In deriving Equation (8.6.1) I have implicitly invoked some basic features of the underlying probability event space used in SOI. In particular, I have assumed that $\Pr(x, \neg m(\mathcal{W})) = \Pr(x, m(\mathcal{W}^c))$ and that, if $\mathcal{V}_1, \mathcal{V}_2, \dots, \mathcal{V}_k$ are disjoint spatial regions, $\Pr(x, \bigcup_i \mathcal{V}_i) = \sum_i \Pr(x, \mathcal{V}_i)$. From (8.6.1) and the Postulates above, we then have

$$\begin{aligned}
\Pr(m(\mathcal{V}), m(\mathcal{W})) &= \int_{\mathcal{V}} dx_1^3 \int_{\mathcal{W}} dx_2^3 \Psi^*(x_1, x_2) \Psi(x_1, x_2) \\
&= \int_{\mathcal{V}} dx_1^3 \int_{\mathcal{W}} dx_2^3 [\psi^*(x_1) \psi(x_1) \phi^*(x_2) \phi(x_2) + \\
&\quad \psi^*(x_2) \psi(x_2) \phi^*(x_1) \phi(x_1) - \psi^*(x_1) \psi(x_2) \phi^*(x_2) \phi(x_1) - \\
&\quad \psi^*(x_2) \psi(x_1) \phi^*(x_1) \phi(x_2)] \\
&= \left(\int_{\mathcal{V}} dx_1^3 \|\psi(x_1)\|^2 \right) \left(\int_{\mathcal{W}} dx_2^3 \|\phi(x_2)\|^2 \right) + \alpha
\end{aligned} \tag{8.6.2}$$

where α is a constant with absolute value much less than unity. Similarly,

$$\Pr(m(\mathcal{V}), m(\mathcal{V})) = \beta \tag{8.6.3}$$

and

$$\Pr(m(\mathcal{V}), m(\mathcal{W}^c \setminus \mathcal{V})) = \left(\int_{\mathcal{V}} dx_1^3 \|\psi(x_1)\|^2 \right) \left(\int_{\mathcal{W}^c \setminus \mathcal{V}} dx_2^3 \|\phi(x_2)\|^2 \right) + \gamma \tag{8.6.4}$$

where β and γ are again constants with vanishing absolute value.

Assembling these pieces, we have from Equations (8.6.2), (8.6.3), and (8.6.4):

$$\begin{aligned}
\Pr(m(\mathcal{V})) &= \left(\int_{\mathcal{V}} dx_1^3 \|\psi(x_1)\|^2 \right) \left(\int_{\mathcal{W}} dx_2^3 \|\phi(x_2)\|^2 \right) \\
&\quad + \left(\int_{\mathcal{V}} dx_1^3 \|\psi(x_1)\|^2 \right) \left(\int_{\mathcal{W}^c \setminus \mathcal{V}} dx_2^3 \|\phi(x_2)\|^2 \right) \\
&\quad + (\alpha + \beta + \gamma) \\
&= \left(\int_{\mathcal{V}} dx_1^3 \|\psi(x_1)\|^2 \right) \left(\int_{\mathcal{W}} dx_2^3 \|\phi(x_2)\|^2 + \int_{\mathcal{W}^c \setminus \mathcal{V}} dx_2^3 \|\phi(x_2)\|^2 \right) \\
&\quad + (\alpha + \beta + \gamma) \\
&= q + (\alpha + \beta + \gamma) \\
&\approx q
\end{aligned} \tag{8.6.5}$$

A similar derivation runs as follows:

$$\begin{aligned}
\Pr(m(\mathcal{W})) &= \Pr(m(\mathcal{W}), m(\mathcal{V})) + \Pr(m(\mathcal{W}), \neg m(\mathcal{V})) \\
&= \Pr(m(\mathcal{W}), m(\mathcal{V})) + \Pr(m(\mathcal{W}), m(\mathcal{W})) + \Pr(m(\mathcal{W}), m(\mathcal{V}^G \setminus \mathcal{W})) \\
&= \left(\int_{\mathcal{V}} dx_1^3 \|\psi(x_1)\|^2 \right) \left(\int_{\mathcal{W}} dx_2^3 \|\phi(x_2)\|^2 \right) + \alpha + \delta \\
&\quad + \left(\int_{\mathcal{V}^G \setminus \mathcal{W}} dx_1^3 \|\psi(x_1)\|^2 \right) \left(\int_{\mathcal{W}} dx_2^3 \|\phi(x_2)\|^2 \right) + \zeta \\
&= \left(\int_{\mathcal{W}} dx_2^3 \|\phi(x_2)\|^2 \right) + (\alpha + \delta + \zeta) \\
&\approx \left(\int_{\mathcal{W}} dx_2^3 \|\phi(x_2)\|^2 \right)
\end{aligned} \tag{8.6.6}$$

Again, the constants δ and ζ are small compared to unity.

Taken together, Equations (8.6.5) and (8.6.6) indicate that the two-place state $\Psi(x_1, x_2)$ does in fact yield specifications for two distinct spatial regions that are approximately the same as any two specifications extracted from the one-place states allowed for those regions considered separately. Thus in the case of a single unit of mass—even in a world containing only ‘particles’ of the same type—a spatial region can be independent of all others in the sense relevant to EDiv.⁶⁹

8.7 SOI and Scientific Practice

SOI was constructed largely to demonstrate that, even after particles have been abandoned, an object ontology which supports EDiv is possible—there is no need to seek structuralist alternatives. Though it was not motivated by any particular scientific concerns, SOI is closely related to at least one interpretation in the chemical literature. ‘Atoms in Molecules’ or AIM is an important competitor of Density Functional Theory in quantum chemistry.⁷⁰ Broadly speaking, AIM treats the electron density function ρ (the square modulus of the many-electron wave-function) as the primary theoretical entity from which all chemically relevant properties can be derived. More specifically, atoms, molecules, and chemical bonds can all be defined in terms of geometric features of the gradient field on ρ . ‘Atoms’, for instance, consist of the union of an attractor (a point at which field lines converge) and its basin (the region of space from which the convergent field lines originate).

⁶⁹ Note that $\Psi(x_1, x_2)$ is exactly the sort of wavefunction Messiah considers when attempting to argue that putative particles can be independent (see Section 7.2 above). Though the discussion in terms of particles was incoherent, Messiah’s claims make sense if we think in terms of spatial regions.

⁷⁰ See (Bader 1990, 1991; Popelier 2000)

AIM—which its proponents refer to as an “*interpretative* theory” (Popelier 2000, xii) (emphasis in original)—and SOI share a great deal in common. To begin with, the wavefunction is given central place. In AIM, it is the means by which ρ is computed. In SOI, it is an encoding of probability distributions corresponding to every spatial region. More strikingly, AIM attributes properties to atoms strictly via volume integrals. That is, ‘atoms’ are spatial regions—denoted by Ω —with properties dictated by integrals over the wavefunction. As Popelier puts it: “*Atomic properties* are then defined as volume integrals over Ω of some integrand which we have to specify for each property” (2000, 46). This is remarkably similar to the view offered by SOI.

Whatever the motivation for its introduction, the convergence of SOI and AIM suggest that such a view is not implausible from a scientific perspective. Furthermore, this similarity suggests a way of accounting for particle talk in chemistry and physics while simultaneously maintaining a viable, particle-free object ontology. In AIM, atoms are *defined* in terms of the properties of spatial regions. As in SOI, it is a spatial region that bears properties, not a particle that bears ‘location’ as one of its properties. Understood in terms of spatial regions that possess particular geometric properties, ‘atoms’ no longer run afoul of PI because there is no sense to be made of permuting their properties. The upshot is that SOI undermines the motivation for ontic structural approaches while providing a means of re-conceptualizing the atoms and other putative particles of chemistry and physics.

Chapter 9 Against structural realism

9.1 Overview

One consequence of the incompatibility of PI with MA is that the scientific realist—if he is to stay a realist—must seek an alternative interpretation for QM and its successors. In the preceding chapter, I sketched one such alternative that treats spatial volumes instead of particles as genuine physical objects. One might instead take the incompatibility to favor a more radical revision of realism, one which disposes of objects altogether. Such is the recommendation of an increasingly vocal group of philosophers who favor ‘Ontic Structural Realism (OSR). It is the aim of this chapter to demonstrate that the fact of incompatibility hinders rather than helps the proponent of OSR, and to suggest ways in which both the incompatibility result and the construction of SOI might be used to argue against realism of either the traditional or OSR variety.

9.2 Two routes to OSR

In Chapter 3, I briefly rehearsed a pair of arguments to the effect that quantum particles are *non-individuals* in the sense that they are numerically distinct entities which do not stand in relations of identity (or nonidentity) with one another. Both of these arguments rest on the consequences of PI. The first is framed in the context of quantum statistical mechanics, and runs something like this:

Argument from statistics:⁷¹

- (S1) All physically possible states accessible to a system are equiprobable with respect to an ensemble of identically prepared systems.
- (S2) Because of PI, quantum states that differ by a permutation of the properties associated with particles of the same type are jointly weighted as a single state.
- (S3) From (S1) and (S2) it follows that states differing only by a permutation of properties amongst particles of the same type are not distinct physical possibilities.
- (S4) If particles were genuinely distinct from one another, then permuting their properties would result in distinct physical possibilities.
- (S5) Quantum particles are numerically distinct.

⁷¹ For a more thorough development of this argument, see e.g. (Belousek 2000; French and Krause 2006, Ch 4; French and Rickles 2003; Reichenbach 1999).

- (S6) From (S3), (S4), and (S5) it follows that particles are non-individuals.

The second argument is closely related to the first. It explicitly assumes Leibniz's Principle of the Identity of Indiscernibles (PII) (see Section 3.5) as the ground for identity relations amongst objects. With this assumption, the argument can be stated as follows:

Argument from PII:⁷²

- (P1) Numerically distinct quantum particles have all of their monadic⁷³ physical properties in common.
- (P2) Because of PI, all particles of the same type have exactly the same monadic properties.
- (P3) From (P2) and the assumption of Leibniz's PII, it follows that the identity of quantum particles cannot be grounded in their physical properties.
- (P4) Metaphysical 'properties' such as haecceity or primitive 'this-ness' are not real.
- (P5) There is nothing to ground relations of identity amongst quantum particles [from (P3) and (P4)].
- (P6) Quantum particles can nonetheless be counted, and are therefore non-individuals.

There are a number of problems with both of these arguments, but critics in the literature have tended to focus on (S1) and (P4). With regard to (S1), it has been pointed out that equiprobability might apply to physically *accessible* states rather than to all physically possible states (see Chapter 3). Because the Hamiltonian is necessarily symmetric, if a particle begins with a state in either the fermionic or bosonic sector of the Hilbert space, it can never come to occupy a state from another sector. In particular, (anti)symmetric states can never evolve into asymmetric states of the sort that would yield distinct states upon permutation. In this way, one can account for the non-classical statistical weighting of quantum states by appealing to the inaccessibility of certain initial states. The reason

⁷² There is a very large literature on PII in QM. For an overview of the debate, see (French and Krause 2006, Sec. 4.2.1; Saunders 2003). For early development of this argument, see (Barnette 1978; Cortes 1976; French 1989, 1989; French and Redhead 1988; Redhead and Teller 1992). For a critique of an argument combining elements of the two I have presented in this section, see (Morganti 2009).

⁷³ There are stronger or weaker versions of PII. The strongest admit only monadic properties into consideration. For a discussion of the various forms of PII (and the related issues of identity and individuality) when relational properties are taken into consideration, see (Caulton and Butterfield 2008; Saunders 2003). For a discussion of relational forms of PII, specifically in quantum mechanics, see (Muller and Saunders 2008).

permutates are all counted as a single state when distributing probability is that particles contingently occupy permutation invariant states, not because asymmetric states are physically impossible.

In stating the argument for non-individuals, I've swept an uncomfortable number of difficulties under the rug, some of which are often overlooked in the literature. For one thing, it is not at all clear that there is a strong analogy between quantum and classical probability distributions over particle states. In the quantum case, when particle roles are permuted in a state representation, the properties assigned to each particle of the same type in the interpretation are not changed. In the classical case, permuting particle roles corresponds to permuting property assignments. It is difficult to read the former in such a way as to make the argument from statistics work. But let us grant for sake of argument that it makes sense to speak of permuting property assignments for quantum particles in (anti)symmetric pure states and that both the argument from statistics and the objection to (S1) are coherent. The upshot is that the physics underdetermines the metaphysics. We cannot settle whether (S1) is true by appealing to physical theory. Likewise, (P4) in the argument from PII is left equally ambiguous by the physics. Steven French has repeatedly asserted that it is in accord with QM to assert either that (P4) is true and thus particles are non-individuals or that (P4) is false and particles are individuals by virtue of some non-physical property like 'this-ness' (Adams 1979).

Whether we consider the argument from statistics or the argument from PII, the physics leaves underdetermined even the broad strokes of the metaphysical view the realist should endorse. This underdetermination is considered by some to weigh heavily against any realist stance that takes quantum objects seriously. Ladyman, for instance, says that "[w]e 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, 419-20).

Ladyman (1998) takes this under-determination as a strong motivation for abandoning an ontology of objects altogether. The alternative he endorses—namely OSR—posits a world made entirely of 'structure'.⁷⁴ Just what structure is supposed to be remains rather vague in the OSR literature. The term is often used as if it refers to the set of relations that obtains in the world (without relations, of course). Rather more vaguely, it might be said that structure is whatever is captured by the equations of mathematical physics without asserting anything

⁷⁴ See (Ladyman 2007; Ladyman et al. 2007).

about the natures (monadic properties, identities, etc.) of any objects ostensibly related by those equations. Whatever structure is, adopting an ontology that posits only structure is supposed to free us from the under-determination that plagues any attempt to provide a particle ontology for QM. By granting ontological primacy to structure, we supposedly obviate the choice between the view of particles as non-individuals or as entities with non-physical haecceities.⁷⁵

The second major motivation for OSR in particular and structuralism in general has to do with the problem of theory succession. OSR has its intellectual roots in an epistemic program that runs back at least as far as Poincaré (1952). Epistemic Structural Realism (ESR) amounts to something like the claim that all we can *know* about the world is structure. This is in contrast to the more radical claim of OSR that all there *exists* in the world is structure. Both ESR and OSR are attractive for the realist in part because they each offer a means of defusing the ‘pessimistic meta-induction’. Larry Laudan (1981) is typically cited as providing the definitive form of this argument, and so we’ll consider Laudan’s formulation. As he presents it, the pessimistic meta-induction invokes a history of ontological discontinuity in scientific interpretations as an objection to a prominent realist argument that in turn is predicated on the cumulative nature of science. The realist argument in question is abductive and runs as follows:

- C1) If earlier theories in a scientific domain are successful and thereby, according to realist principles...approximately true, then scientists should only accept later theories which retain appropriate portions of earlier theories;
 - C2) As a matter of fact, scientists do adopt the strategy of (C1) and manage to produce new, more successful theories in the process;
 - C3) The ‘fact’ that scientists succeed at retaining appropriate parts of earlier theories in more successful successors shows that the earlier theories did genuinely refer and that they were approximately true. And thus, the strategy propounded in (C1) is sound.
- (Laudan 1981, 36)

Laudan argues that C1 requires a continuity of ontology in the sense that the referents of at least some theoretical terms must be preserved across theory change. In order to satisfy (C1), each successor theory—in some suitable limit taken under appropriate boundary conditions and in the region of empirical overlap with the old theory—must contain the theory it succeeded, including its interpretation. That is, the successor theory must—in some suitably restricted fashion—make at least *some* of the same claims about entities, processes, and

⁷⁵ There is a similar argument from diffeomorphism invariance in GR [see e.g. (Ladyman et al. 2007, Sec. 3.2; Pooley 2005; Stachel 2002)].

properties that the original theory made when interpreted. Those claims that are invariant across theory change are precisely those which the realist argues are probably true. However, says Laudan, the history of science suggests that such ontological claims are rarely preserved.

This argument against (C1) constitutes the ‘pessimistic meta-induction’, and the conflict between the history of theory succession and the aforementioned argument for realism is what I take to be the ‘problem of theory change’. The realist needs to account for the apparent lack of continuity in theory change if he is to preserve one of his strongest arguments in favor of realism.

Structuralism—in either epistemic or ontological forms—is supposed to resolve the problem of theory change by making science look far more cumulative. While successive theories may not preserve the entities posited by their predecessors, they supposedly retain the ‘structure’. Of course, to flesh out such a position requires one to take a stand on just what constitutes structure. It is not clear that any convincing account exists that can do what is asked of it. In the last section of this chapter, I’ll offer a few reasons to suspect that no such account will be forthcoming. In the next section, however, I restrict the discussion to the impact of the incompatibility result of Chapter 6 on the argument for OSR that derives from underdetermination in the interpretation of QM.

9.3 Undermining Underdetermination

At first blush, it appears that the general incompatibility between PI and MA weighs in favor of OSR. My argument extends to any theory which incorporates PI, while the original argument from under-determination applied only to QM. I’ve thus ruled out particle interpretations for a broad class of theories rather than one in particular, and this suggests a pervasive problem with object ontologies. The proponent of OSR might take this as further evidence that we ought to quit thinking in terms of objects and simply learn to love structure instead. This picture is badly mistaken. By eliminating particle interpretations, the incompatibility of PI and MA actually dissolves the underdetermination which initially motivated OSR. At the same time, it leaves open the possibility of alternative object ontologies—it is a mistake to take evidence against particles as evidence against objects altogether.

First, let’s consider the support for OSR that ostensibly derives from the underdetermination of ontology by physics in the case of QM. In particular, let’s consider each of the competing ‘packages’ for interpreting the theory. On the one hand, there is the ‘Received View’ which asserts that quantum particles are numerically distinct entities that nonetheless lack identity. This package is simply

incoherent. As I have argued elsewhere (Jantzen 2010), cardinality (and thus numerical distinctness) is inseparable from identity—a collection of entities without identities simply has no determinate cardinality. Thus, rather than represent one of two equally viable particle interpretations, this first package—were it somehow forced upon us by the physics—amounts to a *reductio* against a particle ontology. Of course, the claim that there exist numerically distinct particles is not forced on us by the physics. Any argument one might offer in favor of the claim that a certain number of particles have been counted in some experiment can just as well be taken as evidence for the claim that a certain magnitude has been measured for a discrete property of some spatial volume. This is the SOI interpretation of the data, and SOI is just one example of an interpretation that allows us to retain our well-developed theory of cardinality when speaking about objects. There are likely others. Thus, the premises supporting the first particle package really support an argument against particle ontologies while at the same time suggesting ways to construct alternative object ontologies.

Second, consider the package that takes particles to be individuals but supposes as a brute fact that certain states are inaccessible. The bulk of this essay has been an argument to the effect that this is not a coherent view either. If the putative particles have identities—as indeed they must if they are to be the sort of entities of which MA speaks—then it is still the case that PI prevents them from bearing properties with even approximate independence. Thus, there is no viable particle interpretation for distinct particles either. Without identity, talk of particles is meaningless; grant the particles identities and the account is untenable.

At this point, it might still look as if I've merely strengthened the case for OSR. I've argued that the under-determination which was supposed to push us away from particle accounts does not actually leave the choice of particle ontology undetermined—there simply are no viable particle ontologies. Though the motivation provided by underdetermination has vanished, it looks like an ontology that eschews objects is the best remaining option. But to jump to the OSR conclusion is to move too quickly. This is because the incompatibility of PI and MA does not rule out other object ontologies compatible with the physics. In particular, SOI as it was developed in the previous chapter is proof positive that there are viable object ontologies. Thus, while the traditional realist cannot help himself to an ontology of particles, he still has available at least one ontology that retains objects which bear properties and stand in relations of identity. OSR is therefore not the last ontology standing—its acceptance requires a positive ar-

gument for superior value. For this reason, it is not clear why OSR should be seen as an attractive alternative.

To sum up, the two packages to which the underdetermination argument refers are both inviable. However, these two packages do not exhaust the possibilities. Object ontologies are still available—in particular SOI—and so no motivation for OSR is to be derived from the underdetermined interpretation of QM. Instead, one would need to give reasons why some specific OSR account fares better than SOI for QM. In the absence of such an account, we may as well retain objects, albeit without particles.

9.4 Theory succession, OSR, and SOI

What of the problem of theory change, the second principal motivation for adopting OSR? Here I can only provide some suggestive observations. In particular, there are reasons to suspect that both the incompatibility result and the availability of SOI work against the realist of either stripe. The incompatibility of PI and MA forces the traditional realist toward an alternative such as SOI. Yet the way in which SOI was constructed—though it provides the realist with a way to avoid the incompatibility of PI and MA—suggests that the continuity of ontology required by the traditional realist is too cheaply purchased. The proponent of OSR was right about underdetermination, but pointed to the wrong space of options. The manner in which SOI was constructed suggests the possibility of indefinitely many alternatives that would allow one to retain features of any given ontology. This seriously undermines the claim that the retention of ontology is indicative of truth. But again, this need not push us into the arms of OSR. It looks like any substantive account of OSR will suffer similar problems. Thus, there is reason enough for both the traditional realist and the OSR proponent to be uncomfortable.

To see how neither SOI nor OSR are likely to help the realist with the problem of theory change, let me cast the problem in the terminology of Chapter 4. When we speak of ‘theory change’ we are talking about replacing what I would call a family of theories with another. Any family of theories of mathematical physics (e.g. Hamiltonian mechanics) consists of a class of mathematical spaces (e.g. the class of n -dimensional phase spaces) along with a class of systems of equations (e.g. the class of Hamilton’s equations of motion, with one set of equations for each combination of Hamiltonian and specification of the degrees of freedom in the system). Each theory in a family is associated with a class of models, and each class of models maps to a class of interpretations. For the interpretation of any given model, we can in principle separate the minimal portion which represents only observable features of the world from the remainder

which—according to the realist at least—represents unobservable entities or processes whose existence accounts for the observable facts.⁷⁶ To simplify matters, I will consider theory succession in terms of a pair of representative theories from each family of theories.

If (C1) is to do any work for the realist, it must be the case that, in changing from one theory T_O to its successor T_S , parts of the interpretations of T_O are preserved beyond the minimal claims about observables. Let $I(T_S)$ be the class of interpretations of T_S and let $I(T_O)$ stand for the class of interpretations of T_O . Then there must be some subclass $I_{\text{sub}}(T_S) \subseteq I(T_S)$ with the following property: for most $x \in I(T_O)$ there exists a $y \in I_{\text{sub}}(T_S)$ such that there is an approximate equivalence between x and y . Furthermore, the approximate equivalence must obtain between *full* interpretations, not merely the components of each interpretation that correspond to observables.

A couple of examples from the history of electrodynamics might help to cut through the tangle of terminology. Consider first the phenomenon of refraction. When light passes from one medium to another with a different index of refraction the apparent ray is deflected. The angles relative to the vertical line perpendicular to the interface made by the incoming ray (θ_{in}) and the outgoing ray (θ_{out}) are given by:

$$n_{\text{in}} \sin(\theta_{\text{in}}) = n_{\text{out}} \sin(\theta_{\text{out}}) \quad (9.4.1)$$

The constants n_{in} and n_{out} are the indices of refraction of the medium before and after the interface respectively. In Laplace's corpuscular theory (circa 1800), light constitutes a large number of moving particles which obey some simple equations of motion (Frankel 1976). From these equations one can derive (9.4.1). If we take Equation (9.4.1) to constitute a theory in its own right (in conjunction with an appropriate 4-D configuration space), then the models consist of two ordered pairs $\langle n_{\text{in}}, \theta_{\text{in}} \rangle, \langle n_{\text{out}}, \theta_{\text{out}} \rangle$. According to Laplace, the interpretations of these models involve descriptions of streams of particle with various dynamical properties.

Laplace's theory proved insufficient to account for some puzzling effects of polarization. Consider the transmission of light through two linear polarizers. The intensity of the light exiting the second polarizer, S , is a simple function of the angle between the transmission axes of the two polarizers:

⁷⁶ This is *not* to say that we could strip away the parts of a theory that are interpreted in terms of unobservables. If we could, then the theory would contain a great deal of surplus structure, at least with respect to its capacity for generating empirical predictions. I am only claiming that we can separate the parts of an interpretation that assert things about observables (e.g. propensities for measurement outcomes) from those that assert things about the existence and properties of unobservable entities (e.g. the properties of microscopic particles).

$$S(\theta) = S_{\max} \cos^2(\theta) \quad (9.4.2)$$

The constant S_{\max} is the maximum transmitted intensity. This relation is known as Malus' Law, and was discovered in the early 19th century by Étienne-Louis Malus, who had been working on the problem of double refraction (Frankel 1974). Laplace's dynamical theory lacked the resources to account for this relation. Largely for that reason, we can say it was succeeded by that of Malus himself, who accounted for the new phenomenon by attributing more properties to the corpuscles of light. It is inessential to consider the details here, but Malus' theory retained the Laplacian dynamics used to derive (9.4.1). Thus, we can consider Malus to have provided a new theory of polarization *and* a successor theory of refraction. Now, the models of Malus' theory of refraction (T_O) and Laplace's theory (T_S) are identical. For any given interpretation, there is also substantial overlap. Not only do the interpretations of Malus contain the same assertions concerning observables like the angle of refraction, but they retain all of Laplace's property ascriptions. There are more properties attributed to corpuscles in Malus' account, but the ontology of Laplace has been retained in full. So—even though I have not given an account of approximate equivalence in this case—it is plausible to consider every model of $I(T_O)$ to be approximately the same as a corresponding model in $I(T_S)$. This is the sort of continuity the traditional realist requires.

Consider instead the succession of Malus' theory of polarization (which we'll suppose is restricted to Equation (9.4.2)) and Maxwell's account of the same. Again, we needn't concern ourselves with the details. Suffice it to say that in this case the equations comprising each theory are the same, as are the models. However, in this case, each model for Malus corresponds to claims about the properties of corpuscles while the interpretations of Maxwell's theory contain only descriptions of the state of a disembodied electromagnetic field. Beyond the minimal overlap in assertions about observables, these interpretations are entirely different. This is the sort of succession that—if widespread—is problematic for the realist.

What the construction of SOI suggests is that the alignment of the unobservable portions of the interpretations belonging to a theory and its successor is arbitrary, or at least very flexible. The observable portions of interpretations must coincide. This much the non-realist admits. But there seems to be little constraint on the remainder. SOI was built by fixing the observable portion and then reconfiguring the rest to achieve the desired result—an object ontology free from conflict with PI. The problem in short is that the interpretation of a theory's models in terms of unobservables is underdetermined by the models of

a theory. This problem is similar to, but not identical with the problem of theory under-determination in which the formulation of a theory is underdetermined by any finite set of observations. The point here is that, without both the right sort of relation between the successor theory and that which it replaces *and* strong constraints on plausible interpretations, the argument from the retention of ontology collapses. The concoction of SOI suggests that the right inter-theory relations can be said to obtain but only because the constraints on interpretations are so weak. If I can pick whatever sorts of entities I like in order to populate the interpretations of a scientific theory, then there is no reason to take those entities seriously.

I have ignored quite a few responses open to the realist at this point. For instance, the interpretation of one ‘theory’ as I’ve narrowly understood the term is constrained by other theories as well as various pragmatic concerns such as fecundity. I haven’t the space to pursue these objections. I merely want to suggest that both the incompatibility result and the construction of an alternate interpretation can be construed as weighing against the traditional realist stance.

Would the switch to OSR help? Perhaps, but it doesn’t look promising. If structure is supposed to be something that appears in the interpretations of a theory’s models—something like the set of relations that obtain or the logical form of that set of relations—then it looks as if we face the same dilemma as above. On one horn, we can embrace whatever interpretations were historically adopted but then we would be forced to concede that structure is not preserved across episodes of theory change. Think, for instance, of the different relational structures that obtain in the corpuscular interpretations of Malus’ theory versus the field interpretations of Maxwell’s. On the other horn of the dilemma, we could impose novel interpretations on old theories in order to secure continuity. But again, the worry is that it is too easy to do so. Beyond the relations connecting observables, I am free to reconfigure interpretations as I please. This includes the elements of ‘structure’. So either structure fails to be preserved across theory change or it is preserved but only because we can always find an interpretation that does what we want—interpretations in terms of structure are underdetermined.

To avoid this dilemma, we might take structure to be a property of the class of minimal interpretations—those parts which refer only to observables. Then it is plausible to suppose that structure is preserved across theory change *and* that it is strongly constrained by specification of a theory. But then realism about structure wouldn’t be a very strong sort of realism—it would amount to the claim that there really do exist empirical regularities amongst observable quanti-

ties. Whether we reify those regularities as ‘modal relations that exist in the world’ or not seems to be a distinction without a difference.

Of course, these are only vague and intuitive considerations. One can only pass judgment on specific structuralist proposals. Examining the few proposals that are out there would take us too far astray. I’ll settle for the following speculation: any account of structure that resolves the problem of theory succession collapses into a very weak sort of realism about empirical regularities. OSR is thus no better off at the end of the day than traditional entity realism.

Chapter 10 Conclusion

10.1 Recapitulation

I have tried to convince the reader that anyone predisposed to scientific realism—to the position that our theories provide true descriptions of the unobservable contents of the world—ought to favor particle ontologies. Those arguments which support realism in general tend to favor particle ontologies in particular. In Chapter 2, I defended this claim with the following argument:

Argument for Particles:

- (P1) EDiv is true.
- (P2) Any ontology that entails EDiv is better supported by the fact of EDiv than any ontology—particle or otherwise—that does not.
- (P3) MA is the logically weakest particle ontology and it entails EDiv.
- (P4) Every major scientific interpretation that posits particles (atoms, electrons etc.) has been compatible with MA, including QM (if PI is dropped).
- (P5) The realist should favor particle ontologies compatible with MA.

This argument in favor of particle ontologies suffers all the deficiencies of a likelihood argument and is at best suggestive. But the argument, along with the sheer ubiquity of the particle interpretation in scientific discourse, motivates a consideration of whether such a view is consistent with our best physics. It is not. The bulk of this essay constitutes an argument for the conclusion that any theory containing PI is strictly incompatible with MA or any of its derivative interpretations. This latter incompatibility argument, completed in Chapter 6, can be glossed as follows:

Incompatibility:

- (I1) MA requires particles to bear state-dependent properties approximately independent of one another (Chapter 3).
- (I2) PI requires all particles in the interpretation of a model to possess identical properties (Chapter 6).
- (I3) PI prevents particles from bearing properties even approximately independent of one another (Chapter 6).
- (I4) If PI, then not MA and vice versa.

- (15) PI is empirically necessary (Chapter 5).
- (16) MA is false—there are no particles.

If the incompatibility argument is sound, then the deeply entrenched view of the world described in Chapter 1 is untenable. While laymen and scientists alike speak of atoms as composites of protons, neutrons, and electrons, statements of this sort cannot literally be true. Whatever the world is made of, it isn't particles. Of course, this rejection of particles is predicated on the assumption that any empirically adequate microphysics of the future will retain PI. This may be false—the future may once again offer us a physics compatible with MA. But until such time, the scientific realist must seek alternatives.

Such alternatives exist, at least if we restrict our attention to QM. In Chapter 8 I sketched an interpretation which I called SOI. This interpretation posits spatial volumes as the fundamental objects of the world. SOI avoids conflict with PI and entails EDiv. But it has its shortcomings. For one, it is incompatible with special relativity. It is also formally awkward. Whether these deficiencies can be corrected in a variation of this approach which takes spacetime volumes as objects along the lines of the proposal by Wallace and Timpson (2009) remains to be seen. Either way, it may be worth exploring the consequences of SOI for QM. As I stressed in Chapter 2, interpretations determine the form of approximations and guide theory extension—SOI may inspire some novel physics in both respects.

Finally, we were led to ask whether SOI bolsters the case of the scientific realist or if instead the incompatibility between PI and MA should incline one to modify or reject the realist project. In the preceding chapter, I argued that, while SOI is not necessarily a salve for realism, the incompatibility result ought not to lead one into the arms of OSR. I won't rehearse those arguments again. The upshot is that the two most appealing options are an abandonment of the sort of scientific realism with which we have so far been concerned, or the reconfiguration of our scientific interpretations along lines suggested by SOI. If we choose the latter, then much work remains to be done. In the next section, I'll close this essay by considering a puzzle that challenges any realist interpretation of microphysical theory: given a particular ontology of microphysics, what is the correct ontology for the macroscopic world?

10.2 What are macroscopic objects?

According to the traditional particle interpretation that still permeates the special sciences, macroscopic objects are conglomerates of particles. The particles mak-

ing up an object—such as a stone or a human being—are typically distinguished from particles outside the object by spatial proximity, inter-particle forces of attraction, or some combination thereof. So for instance, the four-centimeter-wide octagonal crystal of fluorite on my shelf is a macroscopic object. In the canonical scientific view, it is composed of atoms of calcium and fluoride. These particles—which bear properties such as charge and position independent of one another—are bound together via strong electromagnetic forces which act only between particles in the crystal and not between particles of the crystal and those outside of it (e.g. the air molecules that surround it). Furthermore, all of the particles of the crystal are contiguous. In the canonical view, a macroscopic object is a coherent assembly of adjacent particles.⁷⁷

The canonical view cannot be right if the incompatibility argument is sound. Our best microphysics is incompatible with a particle interpretation. Therefore, whatever macroscopic objects are, they cannot be conglomerates of particles.⁷⁸ So what might they be? SOI offers an answer for the realist. To see this, it will help to clarify what SOI says about the putative micro-objects of the particles view. According to this interpretation, spatial volumes are the objects out of which the world is made. These volumes possess certain properties with integer magnitudes. More accurately, they possess properties the possible values of which are representable by the integers. In the SOI account, ‘particles’ are not objects in their own right, but may be identified with bundles of co-occurring properties that are instantiated by spatial volumes. So for instance, in the particle view one might say that a fast-moving electron passed through a particular volume of space. According to SOI, this is really saying that a pair of property values for charge density and momentum flux were momentarily highly probable for that volume and the same values were highly probable for adjacent volumes just before and after the time of interest. The view being urged here is roughly that of an ocean wave. Waves are not objects themselves—they lack appropriate identities—but are rather bundles of properties instantiated by different vo-

⁷⁷ I am ignoring the metaphysical possibility of distributed objects. While ‘water’ for instance fails to be a single object under the account I’ve sketched, there are other notions of ‘object’ that would include as an instance all water in the universe no matter how dispersed. What’s more, depending on one’s stance concerning composition, it may be the case that every set of objects, no matter how unrelated, constitute an additional object [see, e.g. (Van Cleve 2008)]. It is not my intention to take a stance on this metaphysical debate. Rather, I intend to offer a plausible reconstruction of the typical notion of a macroscopic object as it appears in scientific discourse.

⁷⁸ The ‘Received View’ of particles in QM (discussed in Chapters 3 and 9) invites the following pseudo-problem. If particles are non-individuals but the macroscopic objects they compose are individuals, why and how does this macro-individuality emerge? There is a straightforward answer suggested by Messiah’s discussion of independent groups of particles (see Chapter 7). But developing the response is irrelevant here since I have argued that the Received View is incoherent.

lumes of seawater at different times. This is how SOI forces us to envision the particles we previously took to be objects.

How does this help us solve the problem of macroscopic objects? The idea is to simply scale up the revised notion of a ‘particle’. SOI commits us to the claim that macroscopic spatial volumes are the genuine objects that bear properties, but as was the case with ‘particles’ it is not to these volumes we refer when casually speaking of, say, crystals of fluorite. Just as with the putative particles, one could instead understand casual references to macro-objects like tables as actually referring to co-occurring bundles of properties that are instantiated at any one time by some spatial volume. In this view, a macro-object is some set of macroscopic properties such as mass, mass distribution (i.e. shape), momentum, etc. For instance, when I refer to my hefty fluorite crystal, I can be understood as referring to a collection of properties that includes an octagonal distribution of mass, a propensity to transmit green light and absorb red, an aggregate mass, etc.

Reducing macroscopic objects to bundles of properties instantiated by spatial volumes requires some care when talking about the parts of a macro-object. By the parts of a macro-object, one might mean either its spatial parts, i.e. the various spatial volumes that instantiate the properties identified with the object at some instant of time, or one might mean some other bundle of co-occurring properties that are associated with the macro-object. That is, one might mean an additional macro-object that is included in or composes the first. So for instance, my crystal has as parts a number of fragments—like the upper pyramid—that are macro-objects in their own right in that they are bundles of co-occurring properties such as a mass (roughly half that of the whole crystal) and density. There is significant overlap between these two notions of ‘part’—the co-occurring properties which constitute a macro-object part are always instantiated by a subset of the spatial volume that instantiates the whole object. Thus, parts that are macro-objects are instantiated by volumes that are parts of the original volume.

This view of macro-objects as co-occurring properties entails one interesting consequence: macro-objects must obey Leibniz’s PII. Because macro-objects in the SOI account are individuated on the basis of their properties—they just are bundles of properties instantiated at different times by different spatial volumes—there cannot be two distinct macro-objects with identical properties. This is not to say that there cannot be two distinct macroscopic spatial volumes with identical properties. It is just to say that, in that case, it is meaningless to ask which of the two is object A and which is object B—property bundles only

possess identity relations like objects insofar as they differ from one another in at least one property.

It is true that this account of macroscopic objects does not comport perfectly with the way we casually talk about them. In our informal language, we speak of macro-objects as if position is a property. If we are to give an account of macro-objects from the perspective of SOI, we must instead view position as a reference to the objects which instantiate the bundle of properties comprising an object, not to a property. But this would seem to be a very mild modification of our intuitive ontology of the immediately observable world, particularly when compared to competitors such as eliminative structural realism. Whether or not this account of macro-objects accords with our common mode of speaking is ultimately irrelevant. The adoption of SOI was motivated by a concern for explaining EDiv, not for rationally reconstructing linguistic practices. This motivation remains even if we don't like the modest linguistic mismatch.

Of course, bending our intuitions to fit the SOI account is only necessary if we seek a realist reading of microphysics. The fact that a coherent account of macro-objects can be given from SOI does not mean that we should be realists about spatial volumes or any other entity posited by an interpretation of QM. In light of the discussion of the preceding chapter, one may simply wish to abandon or restrict the realist project altogether, obviating the problem with macro-objects. Nothing I have said here will settle the question of realism. At best, we can summon the ghost of Descartes and say that perhaps, just perhaps, the world is made of space after all.

References

- Adams, Robert (1979), "Primitive Thisness and Primitive Identity", *The Journal of Philosophy* 76:5-26.
- Bach, Alexander (1997), *Indistinguishable Classical Particles*. Berlin: Springer.
- Bader, Richard F. W. (1990), *Atoms in Molecules: A Quantum Theory*. Oxford: Clarendon Press.
- (1991), "A Quantum Theory of Molecular Structure and Its Applications", *Chemical Reviews* 91 (5):893-928.
- Baker, David (2008), "Against Field Interpretations of Quantum Field Theory", *British Journal for the Philosophy of Science* 60 (3):585-609.
- Barnette, R. L. (1978), "Does Quantum Mechanics Disprove the Principle of the Identity of Indiscernibles?" *Philosophy of Science* 45 (3):466-70.
- Belousek, Darrin (2000), "Statistics, Symmetry, and the Conventionality of Indistinguishability in Quantum Mechanics", *Foundations of Physics* 30 (1):1-34.
- Bengtsson, Ingemar, and Karol Zyczkowski (2006), *Geometry of Quantum States: An Introduction to Quantum Entanglement*. Cambridge, UK: Cambridge University Press.
- Bitbol, Michael (2007), "Schrödinger against Particles and Quantum Jumps", in James Evans and A. S. Thorndike (eds.), *Quantum Mechanics at the Crossroads: New Perspectives from History, Philosophy and Physics*, Berlin: Springer, 81-106.
- Bohm, D., and B. J. Hiley (1985), "Unbroken Quantum Realism, from Microscopic to Macroscopic Levels", *Physical Review Letters* 55 (23):2511-4.
- Boscovich, Roger (1966), "A Theory of Natural Philosophy", in Henry A. Borse and Lloyd Motz (eds.), *The World of the Atom*, New York: Basic Books, Inc., 121-8.
- Boyd, Richard N. (1973), "Realism, Underdetermination, and a Causal Theory of Evidence", *Notis* 7 (1):1-12.
- Boyle, Robert (1949), *The Sceptical Chymist*. London: J. M. Dent & Sons. Original edition, 1661.
- (1991), *Selected Philosophical Papers of Robert Boyle*. Edited by M. A. Stewart. Indianapolis, IN: Hackett Publishing Company, Inc.
- Bunge, Mario (1956), "Survey of the Interpretations of Quantum Mechanics", *American Journal of Physics* 24 (4):272-86.
- Busch, Paul, and Gregg Jaeger (2010), "Unsharp Quantum Reality", *PhilSci Archive*.
- Caulton, Adam, and Jeremy Butterfield (2008), "On Kinds of Indiscernibility in Logic and Metaphysics", *PhilSci Archive*.
- Ceder, Jack G. (1961), "Some Generalizations of Metric Spaces", *Pacific Journal of Mathematics* 11 (1):105-25.
- Chakravartty, Anjan (2007), *A Metaphysics for Scientific Realism: Knowing the Unobservable*. Cambridge: Cambridge University Press.

- Chalmers, Alan (2010), *Atomism from the 17th to the 20th Century* 2009 [cited 7/15 2010]. Available from <http://plato.stanford.edu/archives/fall2008/entries/qm-copenhagen/>.
- Chang, Raymon (1994), *Chemistry*. 5th ed. New York, NY: McGraw-Hill.
- Cohen-Tannoudji, Claude, Bernard Diu, and Franck Laloe (1977), *Quantum Mechanics*. 2 vols. Vol. 1. New York: Wiley.
- (1977), *Quantum Mechanics*. 2 vols. Vol. 2. New York: Wiley.
- Cortes, Alberto (1976), "Leibniz's Principle of the Identity of Indiscernibles: A False Principle", *Philosophy of Science* 43 (4):491-505.
- Cramer, John G. (1986), "The Transactional Interpretation of Quantum Mechanics", *Reviews of Modern Physics* 58 (3):647-87.
- Dalton, John (1808), *A New System of Chemical Philosophy*. 2 vols. Vol. 1. Manchester: S. Russell for R. Bickerstaff.
- Dickson, Michael (2007), "Non-Relativistic Quantum Mechanics", in Jeremy Butterfield and John Earman (eds.), *Philosophy of Physics*, New York: Elsevier, 275-415.
- Dickson, Michael, and Dennis Dieks (2010), *Modal Interpretations of Quantum Mechanics* 2009 [cited 7/14 2010]. Available from <http://plato.stanford.edu/archives/spr2009/entries/qm-modal/>.
- Earman, John (1986), *A Primer on Determinism*. Dordrecht: Reidel.
- Falkenburg, Brigitte (2007), *Particle Metaphysics: A Critical Account of Subatomic Reality*. Berlin: Springer.
- Faye, Jan (2010), *Copenhagen Interpretation of Quantum Mechanics* 2008 [cited 7/14 2010]. Available from <http://plato.stanford.edu/archives/fall2008/entries/qm-copenhagen/>.
- Feynman, Richard P., Robert B. Leighton, and Matthew Sands (1963), *The Feynman Lectures on Physics*. 3 vols. Vol. 1. Reading, MA: Addison-Wesley Publishing Company.
- Filippov, Sergey N., and Vladimir I. Man'ko (2008), "Geometrical Interpretation of the Density Matrix: Mixed and Entangled States", *Journal of Russian Laser Research* 29 (6):564-80.
- Frankel, Eugene (1974), "The Search for a Corpuscular Theory of Double Refraction: Malus, Laplace and the Price Competition of 1808", *Centaurus* 18 (3):223-45.
- (1976), "Corpuscular Optics and the Wave Theory of Light: The Science and Politics of a Revolution in Physics", *Social Studies of Science* 6 (2):141-84.
- French, Steven (1989), "Identity and Individuality in Classical and Quantum Physics", *Australasian Journal of Philosophy* 67 (4):432-46.
- (1989), "Why the Principle of the Identity of Indiscernibles Is Not Contingently True Either", *Synthese* 78:141-66.
- (2010), *Identity and Individuality in Quantum Theory* 2006 [cited 5/15/10 2010]. Available from <http://plato.stanford.edu/archives/spr2006/entries/qt-idind/>.
- French, Steven, and Décio Krause (2006), *Identity in Physics*. Oxford: Clarendon Press.
- French, Steven, and Michael Redhead (1988), "Quantum Physics and the Identity of Indiscernibles", *Brit. J. Phil. Sci.* 39:233-46.

- French, Steven, and Dean Rickles (2003), "Understanding Permutation Symmetry", in K. Brading and E. Castellani (eds.), *Symmetries in Physics: Philosophical Reflections*, New York, NY: Cambridge University Press.
- Fulton, William, and Joe Harris (1991), *Representation Theory: A First Course*. New York: Springer.
- Geroch, Robert (1984), "The Everett Interpretation", *Noûs* 18 (4):617-33.
- Hartle, James B., and John R. Taylor (1969), "Quantum Mechanics of Paraparticles", *Physical Review* 178 (5):2043-51.
- Horgan, Terry, and Matjaž Potrč (2000), "Blobjectivism and Indirect Correspondence", *Facta Philosophica* 2:249-70.
- (2002), "Addressing Questions for Blobjectivism", *Facta Philosophica* 4:311-20.
- Jantzen, Benjamin C. (2010), "No Two Entities without Identity", *Synthese* (forthcoming).
- Ladyman, J. (2007), *Structural Realism* 2007 [cited December 20 2007]. Available from <http://plato.stanford.edu/entries/structural-realism/#OSRSpaPhy>.
- Ladyman, James (1998), "What Is Structural Realism?" *Studies in History and Philosophy of Science* 29 (3):409-24.
- Ladyman, James, Don Ross, David Spurrett, and John Collier (2007), *Everything Must Go: Metaphysics Naturalized*. Oxford: Oxford University Press.
- Landau, Rubin H. (1996), *Quantum Mechanics II: A Second Course in Quantum Theory*. 2nd ed. New York: John Wiley & Sons, Inc.
- Laudan, Larry (1981), "A Confutation of Convergent Realism", *Philosophy of Science* 48 (1):19-49.
- Lavine, Shaughan (1991), "Is Quantum Mechanics an Atomistic Theory?" *Synthese* 89:253-71.
- Levere, Trevor H. (2001), *Transforming Matter: A History of Chemistry from Alchemy to the Buckyball*. Baltimore: The Johns Hopkins University Press.
- Magnus, P. D., and Craig Callender (2004), "Realist Ennui and the Base Rate Fallacy", *Philosophy of Science* 71:320-38.
- Massimi, Michela (2001), "Exclusion Principle and the Identity of Indiscernibles: A Response to Margenau's Argument", *British Journal for the Philosophy of Science* 52 (2):303-30.
- Messiah, A. M. L., and O. W. Greenberg (1964), "Symmetrization Postulate and Its Experimental Foundation", *Physical Review* 136 (1B):248-67.
- Messiah, Albert (1999), *Quantum Mechanics*. Translated by G. H. Temmer. 2 vols. Vol. 1. Mineola, NY: Dover Publications, Inc.
- (1999), *Quantum Mechanics*. Translated by G. H. Temmer. 2 vols. Vol. 2. Mineola, NY: Dover Publications, Inc.
- Morganti, Matteo (2009), "Inherent Properties and Statistics with Individual Particles in Quantum Mechanics", *Studies in History and Philosophy of Modern Physics* 40 (3):223-31.
- Muller, F. A., and Simon Saunders (2008), "Discerning Fermions", *British Journal for the Philosophy of Science* 59:499-548.
- Musgrave, Alan (2007), "The 'Miracle Argument' for Scientific Realism", *The Rutherford Journal: The New Zealand Journal for the History and Philosophy of Science and Technology* 2.

- Narens, Louis (2007), *Introduction to the Theories of Measurement and Meaningfulness and the Use of Symmetry in Science*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Nash, Leonard K. (1950), *The Atomic-Molecular Theory*. Edited by James Bryan Conant, *Harvard Case Histories in Experimental Science*. Cambridge: Harvard University Press.
- Newman, W. (1996), "The Alchemical Sources of Robert Boyle's Chemistry", *Annals of Science* 53:567-85.
- Nielsen, Michael A., and Isaac L. Chuang (2000), *Quantum Computation and Quantum Information*. Cambridge: Cambridge University Press.
- Norton, John D. (2008), "The Dome: An Unexpectedly Simple Failure of Determinism", *Philosophy of Science* 75:786-98.
- Omnès, Roland (1994), *The Interpretation of Quantum Mechanics*. Princeton, NJ: Princeton University Press.
- Panton, Ronald L. (1996), *Incompressible Flow*. 2nd ed. New York: John Wiley & Sons, Inc.
- Pathria, R. K. (1996), *Statistical Mechanics*. 2nd ed. Oxford ; Boston: Butterworth-Heinemann.
- Peskin, Michael E., and Daniel V. Schroeder (1995), *An Introduction to Quantum Field Theory*. Cambridge, MA: Perseus Books.
- Plattner, Gian-Reto, and Ingo Sick (1981), "Coherence, Interference and the Pauli Principle: Coulomb Scattering of Carbon from Carbon", *European Journal of Physics* 2:109-13.
- Poincaré, Henri (1952), *Science and Hypothesis*. New York: Dover Publications, Inc. Original edition, 1905.
- Pooley, Oliver (2005), "Points, Particles, and Structural Realism", in French and Saatsi Rickles (ed.), *The Structural Foundations of Quantum Gravity* Oxford: Oxford University Press.
- Popelier, Paul (2000), *Atoms in Molecules: An Introduction*. Harlow, England: Prentice Hall.
- Post, Heinz (1963), "Individuality and Physics", *The Listener* 70:534-7.
- Putnam, Hilary (1979), *Mathematics, Matter, and Method*. 2nd ed. Cambridge: Cambridge University Press.
- Quine, Willard (1976), "Whither Physical Objects?" in R. S. Cohen, Paul K. Feyerabend and Marx W. Wartofsky (eds.), *Essays in Memory of Imre Lakatos*, Dordrecht: D. Reidel Publishing Company, 497-504.
- Redhead, Michael, and Paul Teller (1992), "Particle Labels and the Theory of Indistinguishable Particles in Quantum Mechanics", *British Journal for the Philosophy of Science* 43:201-18.
- Reichenbach, Hans (1999), *The Direction of Time*. Mineola, NY: Dover
- Rutherford, Ernest (1966), "The Scattering of A and B Particles by Matter and the Structure of the Atom", in Henry A. Boorse and Lloyd Motz (eds.), *The World of the Atom*, New York: Basic Books, Inc., 707-21.
- Saunders, Simon (2003), "Physics and Leibniz's Principles", in Katherine Brading and Elena Castellani (eds.), *Symmetries in Physics*, Cambridge, UK: Cambridge University Press, 289-307.
- (2006), "On the Explanation for Quantum Statistics", *Studies in History and Philosophy of Modern Physics* 37:192-211.

- Schrödinger, Erwin (1998), "What Is an Elementary Particle?" in Elena Castellani (ed.), *Interpreting Bodies: Classical and Quantum Objects in Modern Physics*, Princeton, NJ: Princeton University Press.
- Simons, Peter (2002), "Candidate General Ontologies for Situating Quantum Field Theory", in Meinard Kuhlmann, Holger Lyer and Andrew Wayne (eds.), *Ontological Aspects of Quantum Field Theory*, New Jersey: World Scientific, 33-52.
- Stachel, J. (2002), "‘the Relations between Things’ Versus ‘the Things between Relations’: The Deeper Meaning of the Hole Argument", in D. B. Malament (ed.), *Reading Natural Philosophy. Essays in the History and Philosophy of Science and Mathematics*, Chicago: Open Court, 231-66.
- Suppe, Frederick (1977), *The Structure of Scientific Theories*. 2nd ed. Urbana: University of Illinois Press.
- (1988), *The Semantic Conception of Theories and Scientific Realism*. Urbana: University of Illinois Press.
- Teller, Paul (1995), *An Interpretive Introduction to Quantum Field Theory*. Princeton, NJ: Princeton University Press.
- Van Cleve, James (2008), "The Moon and Sixpence: A Defense of Mereological Universalism", in Theodore Sider, John Hawthorne and Dean W. Zimmerman (eds.), *Contemporary Debates in Metaphysics*, Malden, MA: Blackwell Publishing Ltd, 321-40.
- van Fraassen, Bas C. (1980), *The Scientific Image*. Oxford: Clarendon Press.
- van Melsen, Andrew G. (1952), *From Atomos to Atom: The History of the Concept Atom*. Translated by Henry J. Koren, *Dusquesne Studies: Philosophical Series I*. Pittsburgh, PA: Dusquesne University Press.
- Wallace, David, and Christopher G. Timpson (2009), "Quantum Mechanics on Spacetime I: Spacetime State Realism", *PhilSci Archive*.
- Whyte, Lancelot Law (1961), *Essay on Atomism: From Democritus to 1960*. Middletown, CT: Wesleyan University Press.
- Wilson, Wallace Alvin (1931), "On Semi-Metric Spaces", *American Journal of Mathematics* 53 (2):361-73.