Issues of Identity and Individuality in Quantum Mechanics

Adam Caulton

Darwin College, University of Cambridge

26th September 2011

This dissertation is submitted for the degree of Doctor of Philosophy.

Abstract

This dissertation is ordered into three Parts. Part I is an investigation into identity, indiscernibility and individuality in logic and metaphysics.

In Chapter 2, I investigate identity and discernibility in classical first-order logic. My aim will will be to define four different ways in which objects can be discerned from one another, and to relate these definitions: (i) to the idea of symmetry; and (ii) to the idea of individuality.

In Chapter 3, the four kinds of discernibility are put to use in defining four rival metaphysical theses about indiscernibility and individuality.

Part II sets up a philosophical framework for the work of Part III.

In Chapter 4, I give an account of the rational reconstruction of concepts, inspired chiefly by Carnap and Haslanger. I also offer an account of the interpretation of physical theories.

In Chapter 5, I turn to the specific problem of finding candidate concepts of particle. I present five desiderata that any putative explication ought to satisfy, in order that the proposed concept is a concept of *particle* at all.

Part III surveys three rival proposals for the concept of particle in quantum mechanics.

In Chapter 6, I define factorism and distinguish it from haecceitism. I then propose an amendment to recent work by Saunders, Muller and Seevinck, which seeks to show that factorist particles are all at least weakly discernible. I then present reasons for rejecting factorism.

In Chapter 7, I investigate and build on recent heterodox proposals by Ghirardi, Marinatto and Weber about the most natural concept of entanglement, and by Zanardi about the idea of a natural decomposition of an assembly.

In Chapter 8, I appraise the first of my two heterodox proposals for the concept of particle, *varietism*. I define varietism, and then compare its performance against the desiderata laid out in Chapter 5. I argue that, despite its many merits, varietism suffers a fatal ambiguity problem.

In Chapter 9, I present the second heterodox proposal: *emergentism*. I argue that emergentism provides the best concept of particle, but that it is does so imperfectly; so there may be no concept of particle to be had in quantum mechanics. If emergentism is true, then particles are (higher-order) properties of the assembly, itself treated as the basic bearer of properties.

Preface

I am grateful to many people for discussions, criticisms and suggestions regarding the claims and arguments of this dissertation. I cannot possibly name all the people who have taught and inspired me, but any attempt at such a list must include Nazim Bouatta, Tim Button, Eric Curiel, Newton da Costa, Foad Dizadji-Bahmani, John Earman, Steven French, Katherine Hawley, Jules Holroyd, Leon Horsten, Nick Huggett, Nick Jones, Jeff Ketland, Brian King, Eleanor Knox, James Ladyman, Øystein Linnebo, Fraser MacBride, Kerry McKenzie, Fred Muller, Sam Nicholson, Tom Pashby, Oliver Pooley, Miklós Rédei, Simon Saunders, Michael Seevinck, Rob Spekkens, Nic Teh, Lee Walters, and Nathan Wildman.

I would also like to thank Margrit Edwards, Charlie Evans, Heather Sanderson and Lesley Lancaster for all their help over the last four years.

Finally, I wish to thank Jeremy Butterfield for his unwavering encouragement, infinite patience, and for being the most supportive supervisor a graduate student could hope for.

This dissertation is the result of my own work and includes nothing which is the outcome of work done in collaboration, except some aspects of Chapters 3 and 4, which have appeared as an article co-authored with Jeremy Butterfield.

This dissertation does not exceed the 80,000 word limit set by the Philosophy Degree Committee.

Contents

Chapter 1

Introduction

What is a quantum particle? What properties and relations do they possess? It is the aim of this dissertation to make some progress in finding out.

There has been discussion over many decades about the treatment of indistinguishable (also known as 'identical') particles.[1] This discussion has been invigorated in the last few years, principally by Saunders' (2003b) revival of the Hilbert-Bernays account of identity (also briefly discussed by Quine) and his, and later Muller's and Seevinck's (Muller and Saunders (2008), Muller and Seevinck (2009)), application of it to quantum theory.

In short, Saunders saw that there is an error in the consensus of the previous philosophical literature. That literature had shown that for any assembly of indistinguishable quantum particles (fermions or bosons), and any state of the assembly (appropriately (anti-) symmetrized), and any two particles in the assembly: the reduced density matrices (reduced states) of the particles (and so all probabilities for single-particle measurements) were equal; and so were appropriate corresponding two-particle conditional probabilities. This result strongly suggests that quantum theory endemically violates the principle of the identity of indiscernibles: for any two particles in the assembly are surely indiscernible.[2]

[1] Castellani (1998) is a valuable collection of classic and contemporary articles. French and Krause (2006) is a thorough recent monograph.

[2] This consensus seems to have been first stated by Margenau (1944, pp. 201-3) and it is endorsed and elaborated by e.g. French and Redhead (1988, p. 241), Butterfield (1993, p. 464),

Saunders' basic insight (2006, p. 7) is that the Hilbert-Bernays account provides ways that two objects can be distinct, which are not captured by these orthodox quantum probabilities: and yet which *are* instantiated by quantum theory. Thus on the Hilbert-Bernays account, two objects can be distinct, even while sharing all their monadic properties and their relations to all other objects. For they can be distinguished by either:

(i) a relation R between them holding one way but not the other (which is called 'being relatively discernible');

(ii) a relation R between them holding in both directions but neither object having R to itself ('being weakly discernible').

Saunders, and later Muller and Saunders (2008), argue that the second case (ii) is instantiated by fermions: the prototype example being the relation $R =$ '... has opposite value for spin (in any spatial direction) to ...' for two spin-$\frac{1}{2}$ fermions in the singlet state. More recently, Muller and Seevinck (2009) have argued that (ii) is instantiated by *all* particles. (I will propose a refinement of these results in Section 6.3.)

Saunders' proposals have led to several developments, including exploring the parallel between quantum particles and spacetime points (e.g. Pooley (2006), Caulton and Butterfield (2011), Muller (forthcoming)). But this is work for another day. Here there will be enough to do focussing solely on quantum theory.

The Hilbert-Bernays account is a reductive account, reducing identity to a conjunction of statements of indiscernibility. So this account is controversial: many philosophers hold that identity is *irreducible* to any sort of qualitative facts, but nevertheless wholly unproblematic because completely understood.[3] This latter view is certainly defensible, perhaps orthodox, once one sets aside issues about

and French (2006, §4). Massimi (2001, pp. 326-7) questions these authors' emphasis on monadic properties, but agrees that quantum mechanics violates the identity of indiscernibles when the quantum state is taken to codify purely *relational* properties.

[3]For example, Lewis (1986, pp. 192-193). More generally: in the philosophy of logic, the Fregean tradition that identity is indefinable, but understood, remains strong—including as a response to the Hilbert-Bernays account: cf. e.g. Ketland (2006, p. 305 and Sections 5, 7).

diachronic identity—as I will. But this is not a problem for me: for I am not committed to the Hilbert-Bernays account (and nor is Saunders' basic insight about weak discernibility in quantum theory)—in fact, I will support an interpretation of quantum mechanics in which particles may be indiscernible. Besides, these controversies do not undercut the rationale for my investigation (in Chapter 2) into the kinds of discernibility; essentially because anyone, whatever their philosophy of identity or their attitude to the identity of indiscernibles, will accept that discernibility is a sufficient condition for diversity (the 'non-identity of discernibles').

It cannot be denied that, with the resources of the Hilbert-Bernays account, two objects previously thought indiscernible may yet be discernible (albeit merely 'relatively' or 'weakly'). Nevertheless, the quantum particles discussed in very nearly all of the philosophical literature (including the recent work by Saunders, Muller and Seevinck) possess a feature which is intuitively *like* indiscernibility. The consensus (shared by Saunders, Muller and Seevinck) is that quantum particles of the same species—electrons, say—possess exactly the same *monadic* properties: they are, as I shall say (in Section 2.3) 'absolutely indiscernible'. The unfortunate upshot of this is that no single one of them can be 'picked out' in language or in thought, without appeal to some sort of 'underlying' or 'non-qualitative' property—namely, an *haecceity*.

But this result is intolerable. In the face of such a result, the correct response can only be that something, somewhere has gone wrong. Thus I claim that a fundamental mis-interpretation of quantum theory pervades the consensus. That is, I claim that the consensus has mis-identified the representational relationships between the quantum formalism and the physical world.

The interpretative proposal that leads to this intolerable result I call *factorism*. Despite its unfamilar name, factorism is a familiar proposal. It says that particles are the physical correlates of the labels of factor Hilbert spaces. I will distinguish this from another doctrine, familiar in philosophy: haecceitism (Section 6.2). I will agree that factorism is right for *distinguishable* systems—i.e. systems for which the Indistinguishability Postulate is not imposed.[4] Since such systems differ in, say,

[4]For a thorough discussion of the Indistinguishability Postulate, see French and Krause (2006, pp. 131-149).

intrinsic charge or mass, the labels can be taken as short for such distinguishing properties. But this construal is not available for indistinguishable systems. Here, the correlates of labels make very poor particles, for the reasons just given; and thus I reject factorism. Indeed, as I will argue (Section 6.4), the factorist's particles are like that familiar abstraction of high-school statistics, 'the average taxpayer', and factorism makes a reification error (what Whitehead (1925) called 'the fallacy of misplaced concreteness') analogous to those who take the average taxpayer as a real person.

I will therefore seek a successor to factorism. That is: I will seek a new account of what a quantum particle is, and of how particles are represented in the formalism of quantum theory. But how does this project proceed? How can one look for a new concept without already knowing what one is looking for?

The answer is that we do, at least in part, know what we are looking for. (Otherwise we could not even have suspected that the consensus was in such error.) We have an intuitive notion of what a particle is, and there are other theories—namely, classical mechanics and quantum field theory—whose interpretations provide, in each of their own domains, a more precise and complete account. With an eye on our cloudy, pre-theoretic notion, and on "neighbouring" theories' accounts, we may search the quantum formalism for objects more worthy of the name 'particle'. It is the work of Chapters 4 and 5 to set up the philosophical framework that will make these rough ideas precise.

I will discuss two heterodox interpretative proposals for what these natural candidates are. I call them *varietism* and *emergentism*. They both propose that the real particles are the things of which the factorist's "particles" are the average, just as the average taxpayer encodes statistical information about a population of real people.

Furthermore, both proposals make a break from the current debate in philosophy of physics about whether quantum particles are discernible. We will see that particles in the senses of varietism and emergentism *will* often be absolutely discernible, *contra* the received consensus. But, regrettably, the *weak* discernibility results that in all states any two particles are weakly discernible (mentioned above) will *not* carry over either to varietism or emergentism. Therefore, for both varietism and emergentism, bosons and paraparticles may fail to be discernible at all.

An important difference between varietism and emergentism is whether particles exist in all states of the assembly. According to varietism, there are particles always, i.e. in all states of the assembly. But according to emergentism, there are particles only *sometimes*, i.e. only in some states of the assembly. This prompts a revision in our understanding of the relationship between an assembly and its particles. I investigate these matters in Chapter 9.

To explain and assess my two anti-factorist proposals, I need first to present (Chapter 7) some little-noted subtleties of entanglement for indistinguishable systems. Here I follow an analysis by Ghirardi, Marinatto and Weber (2002). The leading idea is that the definition of entanglement must not appeal to factor-space labels, since it must respect the requirement that all physical quantities be symmetric (i.e. permutation invariant). This can be done, while meshing with the familiar definition for distinguishable systems (namely, that a non-entangled state is a separable, or product state). One defines a state of the assembly to be non-entangled iff the *symmetrization* of some 1-dimensional projector on a factor space has probability 1. Then the main result (Section 7.1) is, roughly speaking, that a state is non-entangled iff it is the appropriate symmetry projection of a product state (taken as including, for bosons, the case of a product of identical factors). (Here, 'appropriate symmetry projection' means, for fermions and bosons, the familiar (anti)-symmetrization; and for paraparticles, projection onto their subspace.) This implies, in particular, that fermionic states do *not* always count as entangled; and bosonic states need not be product states to count as non-entangled.

Later, in Section 7.2, I extend Ghirardi and Marinatto's work by developing a means of *individuating*—i.e., picking out from other constituents of the assembly—a system, or collection of systems, *qualitatively*—i.e., in way that appeals not to factor-space labels, but to single-system states. This development dovetails nicely with recent work by Zanardi (2001). I then provide (Section 7.2.3) a recipe for calculating reduced density operators for qualitatively individuated systems.

With Chapter 7's results in hand, I can state and assess my two anti-factorist proposals: starting (in Chapter 8) with *varietism*. Roughly speaking, varietism proposes that particles are those objects whose statistics are given by pure states whenever (and only whenever) the assembly's state is non-entangled in Ghirardi and Marinatto's sense. The proposal has many merits (reviewed in Section 8.2) in terms of satisfying Chapter 5's desiderata for the concept of particle. However, this proposal founders on an ineliminable ambiguity (Section 8.3), for fermionic or paraparticle states, in the specification of what these objects could be.

This ambiguity is troubling in three ways.

1. It is not a matter of philosophical argument or controvertible interpretation, but follows from the results reviewed in Chapter 7.

2. It is an ambiguity, not between some handful of alternatives, but between continuum-many: for a pair of fermions it is parameterized by points on the Riemann sphere, i.e. by an extended complex number $z \in \mathbb{C} \cup \{\infty\}$.

3. The ambiguity cannot be assimilated to the philosophically familiar cases, such as reference to macroscopic bodies with vague spatial boundaries (Geach's (1980) and Unger's (1980) 'problem of the many'). Such cases are less troubling because the rival disambiguations almost coincide, and we know the parts whose shared containment underpins the near-total overlap of the rival disambiguations. Neither of these features carry over for varietism's disambiguated rival particles.

This leaves my third proposal, *emergentism* (Chapter 9). This has two leading ideas, one philosophical and one physical. I believe the physical idea reflects the practice of physics; the philosophical idea is more idiosyncratic.

The physical idea arises from my admission that factorism is correct for distinguishable systems. This prompts one to require that indistinguishable particles should also be 'label-able'. That is: which properties of the assembly count as particles is governed by which properties can act as labels. In general, these properties must be *qualitative*, i.e. they correspond to states belonging to some basis in the single-particle Hilbert space. Specifically *which* single-particle states they are will, in general, depend on the state of the assembly; but they *must* be states in which no two separate degrees of freedom are entangled. In short: there is a preferred basis problem which emergentism seeks to solve in the same way that nowadays proponents of the Everett interpretation of quantum mechanics seek to solve the traditional preferred basis problem: by appeal to the particular dynamical features of the situation at hand (cf. e.g. Wallace (2003, 2010)).

The philosophical idea is that the basic objects of predication are *not* particles, but something else: perhaps spatial regions or the entire assembly itself. Emergentism then saves particles from being an additional ontological extravagance by identifying them with higher-order properties of the basic object or objects, whether spatial regions or the assembly. This proposal:

(i) accommodates the ambiguity about fermions and paraparticles which besets varietism; but also

(ii) violates the desideratum that particles be the building blocks of assemblies, in the sense (cf. Chapter 5) that particles' properties and relations determine all properties of the assembly: for the assembly may be in a state in which there are no particles.

I conclude that emergentism is the winner in a poor field. There are two main regrets. First, there seems to be no way of preserving the recent results about weak discernibility (by Saunders, Muller, Seevinck, and here in Section 6.3) which recently revived the debate about the discernibility of quantum particles. And second, there seems to be no way to think of assemblies as being "built out of" the emergentist's particles. In this way, the emergentist sees the assembly as many see the field in quantum field theory. Indeed: under emergentism, the assembly of elementary quantum mechanics appears as nothing but the quantum field in a particular limit of (what the textbooks call) conserved total particle number. This leads to two surprises. First, particles under emergentism are *also* closer to *classical* particles than the consensus has it. And second, if the arguments in the following Chapters are right, then, in *non-relativistic elementary quantum mechanics*, we already have reason enough to adopt an basic ontology of fields, rather than particles.

I emphasize at the outset that my focus is orthodox quantum mechanics. More specifically, I set aside:

(i) heterodox cousins such as pilot-wave theory, with its distinctive (and attractively straightforward!) meaning of 'particle', and dynamical reduction models, as developed by Ghirardi and others (reviewed by Bassi and Ghirardi (2003));

(ii) some programmatic responses in defence of varietism that would involve changing the formalism (details in Section 8.3); and

(iii) thorough investigations of quantum field theory (QFT), although I will briefly mention QFT in several places, particularly in Chapters 8 and 9.

I also give notice that I will not shed light on the measurement problem (as the exclusion (i) above hints). What I *will* do is combine considerations about logic, metaphysics, philosophical methodology and philosophy of language with subtleties about entanglement and individuation—odd bedfellows, you may say!—to state and assess three proposals of what is a quantum particle.

1.1 Prospectus

Here I give a short summary of the content of each Chapter.

This thesis is ordered into three Parts. Part I is a rather self-contained investigation into identity and indiscernibility in logic and metaphysics.

In Chapter 2, I investigate identity and discernibility in classical first-order logic. My aim will will be to define four different ways in which objects can be discerned from one another, and to relate these definitions: (i) to the idea of symmetry; and (ii) to the idea of individuality.

In Chapter 3, the kinds of discernibility defined in the previous Chapter are put to use in defining four rival metaphysical theses about identity and individuality. These theses are linked to various positions in the extant literature on identity and indiscernibility, and are compared by their commitments as to what, for each them, is possible. The Chapter concludes with a discussion of a heterodox semantics that is more congenial to three of the four rival metaphysical theses.

Part II sets up a philosophical framework for the work of Part III. It gives an account of what I am doing by asking the question: What is the best concept of particle for quantum mechanics?

In Chapter 4 I give an account of the rational reconstruction of concepts, inspired chiefly by Carnap (1950) and Haslanger (2006). I also propose a way of understanding the interpretation of physical theories. The idea of a representation relation between mathematical and physical realms explains how a theory's mathematical formalism is afforded physical content. This unifies the two projects of theory interpretation and rational reconstruction.

In Chapter 5, I turn to the specific problem of finding candidate concepts of particle. I present five desiderata that any putative concept ought to satisfy, in order that the concept is a concept of *particle* at all. The role of these desiderata is then demonstrated for single-system Hilbert spaces.

Part III surveys the three rival proposals for the concept of particle in quantum mechanics.

In Chapter 6, I define factorism and distinguish it from haecceitism. I then propose an amendment to recent work by Saunders, Muller and Seevinck, which seeks to show that factorist particles are all at least weakly discernible. I then present reasons for rejecting factorism.

In Chapter 7, I turn to more formal matters, and investigate recent heterodox proposals by Ghirardi, Marinatto and Weber about the most natural concept of entanglement. I link this work to Zanardi's proposed conditions for natural decompositions of an assembly's Hilbert space. I build on this work to develop a means of 'qualitatively individuating' a system, and propose a recipe for calculating expectation values and reduced density operators for such systems.

In Chapter 8, I appraise the first of my two heterodox proposals for the concept of particle, *varietism*. I define varietism, and then compare its performance against the desiderata laid out in Chapter 5. I argue that, despite its many merits, varietism suffers a fatal flaw.

In Chapter 9, I present the second of the two heterodox proposals for the concept of particle, *emergentism*. I argue that emergentist provides the best concept of particle, but that it is does so imperfectly; so there may be no concept of particle to be had in quantum mechanics. If emergentism is true, then particles are (higher-order) properties of the assembly, itself treated as the basic bearer of properties.

Part I

Identity in Logic and Metaphysics

Chapter 2

Identity and indiscernibility in logic

The main aim of this Chapter is to define four kinds of discernibility, inspired by Quine (1976a) and Saunders (2003b), and investigate their formal interrelations. These kinds of discernibility, and the formal results regarding them, are important in their own right, and they will be useful in later Chapters. In Section 2.1, I begin by laying down some stipulations about the philosophical terms I will use in this Chapter and thereafter. Then in Section 2.2, I briefly discuss identity in classical first-order predicate logic. In Section 2.3, I define the four kinds of discernibility, and investigate further what I call *individuality.* Finally (Section 2.4), I present some interesting formal results which link certain kinds of discernibility to permutations on the domain of quantification.

2.1 Stipulations about jargon

Here I make some stipulations about philosophical terms. I think that all of them are natural and innocuous, though the last one, about 'individual', is a bit idiosyncratic.

'Object', 'identity', 'discernibility' — I will use 'object' for the broad idea, in the tradition of Frege and Quine, of a potential referent of a singular term, or

value of a variable. The negation of the identity relation on objects I will call (indifferently): 'non-identity', 'distinctness', 'diversity'; (and hence use cognate words like 'distinct', 'diverse'). When I have in mind that a formula applies to one of two objects but not the other, I will say that they are 'discerned', or that the formula 'discerns' them. I will also use 'discernment' and 'discernibility': these are synonyms; (though the former usefully avoids connotations of modality, the latter often sounds better). Their negation I will call 'indiscernibility'.

'Individual' — Following a recent tradition started by Muller & Saunders (2008), I will also use 'individual' for a narrower notion than 'object', viz. an object that is discerned from others by a strong, and traditional, form of difference—which I will call 'absolute discernibility'. Anticipating the following section, this usage may be illustrated by the fact that a haecceitist (in my terms) would demand that all objects be individuals, in virtue of each possessing its own unique property.

It is worth distinguishing cases according to whether the discernment is by an arbitrary language or an 'ideal' one. That is: since the Hilbert-Bernays account will be cast in a formal first-order language, and such languages can differ as to their non-logical vocabulary (primitive predicates), our discussion will sometimes be relative to a choice of such vocabulary. So, for example, an object that fails to be an individual by the lights of an impoverished language may yet be an individual in an ideal language adequate for expressing all facts (especially all facts about identity and diversity). Therefore, one might use the term 'L-individual' for any object which is absolutely discerned from all others using the linguistic resources of the language L. The un-prefixed term 'individual' may then be reserved for the case of ideal language. However, I will stick to the simple term 'individual', since it will always be obvious which language is under consideration. In Chapter 3, I will use 'individual' in the strictly correct sense just proposed, since there I envisage a language which *is* taken to be adequate for expressing all facts.

'Haecceitism' — Though the details will not be needed until Chapter 3, I should say what we will mean by 'haecceitism' (a venerable doctrine going back to Kaplan (1975) if not Duns Scotus!). The core meaning is advocacy of haecceities, i.e. non-qualitative thisness properties: almost always associated with the claim that every object has its own haecceity. But this core meaning is itself ambiguous,

and authors differ about the implications and connotations of 'haecceity'—about how 'thick' a notion is advocated. Some discussion is therefore in order.

2.1.1 Haecceitism

At first sight, there are (at least) three salient ways to construe haecceitism.[1] For each, I give a description and one or more proponents. All three will refer to possible worlds, and will use the language of Lewis's (1986) metaphysics (though it will not require a commitment to his form of modal realism, or to any form of modal realism for that matter). The first is the weakest and the second and third are equivalent; I favour the second and third (and prefer the formulation in the third).

1. *Haecceitistic differences.* Following Lewis (1986, p. 221), we may take haecceitism to be about the way possible worlds represent the modal properties of objects. It is the denial of the following supervenience thesis: a world's representation of *de re* possibilities (that is, possibilities pertaining to particular objects) supervenes on the qualitative mosaic, i.e. the pattern of instantiation of qualitative properties and relations, within that world.[2] Thus a haecceitist allows that two worlds, exactly alike in their qualitative features, may still disagree as to which object partakes in which property or relation.[3]

   This version of haecceitism makes no further claims as to what exactly is left out of the purely qualitative representations, and so it is the weakest of the three haecceitisms. But since two qualitatively identical worlds may disagree about what they represent *de re*, they must differ in some non-qualitative

[1] I owe much of the discussion here to my conversations with Jeremy Butterfield and Fraser MacBride.

[2] This prompts the question, 'What is a *qualitative* property or relation?' That is not a question I will here attempt to answer.

[3] Note that this is a stronger position than one that just allows duplicate worlds, that is, several worlds exactly alike in their qualitative features. The existence of duplicate worlds does not entail haecceitism in Lewis's sense, since, for each individual, every world in a class of mutual duplicates may represent the same dossier of *de re* possibilities. Lewis (1986, p. 224) himself refrains from committing either to duplicate worlds or a "principle of identity of indiscernibles" applied to worlds; though the question seems moot for anyone who does not take possible worlds to be concretely existing entities.

ways: ways which somehow represent (actual) objects in a way that does not rely on how things are qualitatively (whether accidentally or essentially) with those objects. There are two natural candidates for these representatives: the objects themselves (or some abstract surrogate for them), divorced from their qualitative clothing; or some non-qualitative properties which suitably track the objects across worlds. These two candidates prompt the second and third kinds of haecceitism (which, we argue below, are in fact equivalent). I see no sensible alternative to these two candidates for the missing representatives—though perhaps the difference could be taken as a primitive relation between worlds. But haecceitism in my first sense does not entail either of the haecceitisms below, though they each entail haecceitism in my first sense of the acceptance of haecceitistic differences.

2. *Combinatorial independence.* Lewis's definition was inspired by a definition by Kaplan (1975, pp. 722-3); but Kaplan's is stronger. It is phrased explicitly in terms of trans-world identification.[4] But I believe there is a version of haecceitism, clearer than Kaplan's, defined in terms of combinatorial possibility; a version which is still stronger than Lewis's, but does not commit one to claims about non-spatiotemporal overlap between worlds or trans-world mereological sums.

   According to this version of haecceitism, objects partake independently—that is, independently of each other and of the qualitative properties and relations—in the exhaustive recombinations which generate the full space of possible worlds. For example: with a domain of N objects there are: 2^N-many possible property extensions (each of them distinct); 2^{N^2}-many distinct binary relation extensions (each of them distinct); so the number of distinct worlds[5] containing N-objects, k monadic properties and l binary relations (and no other relations), is 2^{kN+lN^2}.

[4]Kaplan: haecceitism is 'the doctrine that holds that it does make sense to ask—without reference to common attributes and behavior—whether this is the same individual in another possible world, that individuals can be extended in logical space (i.e., through possible worlds) in much the way we commonly regard them as being extended in physical space and time, and that a common "thisness" may underlie extreme dissimilarity or distinct thisnesses may underlie great resemblance.'

[5]Or distinct equivalence classes of world-duplicates!

Combinatorial independence would appear to favour the doctrine that objects "endure" identically through possible worlds, since it seems sensible to identify property extensions with sets of objects, and the *same* objects are added to or subtracted from the extensions in the generation of new worlds. But trans-world "perdurance"—the doctrine that trans-world "identity" (in fact a misnomer, according to the doctrine) is a relation holding between parts of the same trans-world continuant—can be accommodated without difficulty. (It is the commitment to there being a unique, objective trans-world continuation relation, and therefore something for a rigid designator to get a grip on, which distinguishes the perdurantist from those, like Lewis, who favour modal talk of *world-bound* objects: cf. Lewis (1986, pp. 218-20).)

3. *Haecceitistic properties.* Perhaps the most obvious form of haecceitism—and perhaps *prima facie* the least attractive—is the acceptance of haecceitistic properties. According to this view, for every object there is a property uniquely associated with it, which that object and no other necessarily possesses, and which is not necessarily co-extensive with any (perhaps complex) qualitative property. An imprecise (and even quasi-religious) reading of these properties is as "inner essences" or "souls" (hence the view's unpopularity). Perhaps a more precise, and less controversial, reading is that each thing possesses some property which "makes" it that thing and no other: a property "in virtue of which" it is that thing (cf. Adams (1981, p. 13)). Whether or not that is in fact more precise, this reading implies a notion of ontological primacy—that the property comes in some sense "before" the thing—about which I remain silent. But besides, I disavow this implication. I intend our third version of haecceitism to be no more controversial than combinatorial independence; in fact I take this version of haecceitism to be *equivalent* to combinatorial independence.

   When I say that, according to this view, there is a haecceitistic property corresponding to each object, by 'object' is not meant 'world-bound object', which would entail a profligate multiplication of properties. Rather, these haecceitistic properties are envisaged as an alternative means to securing trans-world continuation (understood according to either endurantism or perdurantism). The motivation is as follows.

Throughout this thesis, I will be a quidditist, meaning that I take for granted the trans-world identity of properties and relations. (I remain agnostic as to whether this trans-world identity is genuine identity, i.e. endurance of qualities; or whether there is instead a similarity relation applying to qualities between worlds (either second-order, applying directly to the qualities themselves; or else first-order, applying to the qualities indirectly, via the objects that instantiate them), i.e. perdurance.) So for the sake of mere uniformity, haecceitism may be accommodated into my quidditistic framework by letting a trans-world monadic property do duty for each trans-world object. Haecceitism and anti-haecceitism alike can then be discussed neutrally in terms of world-bound objects and trans-world properties and relations, the difference between them reconstrued in terms of whether or not there are primitive monadic properties which allow one to simulate rigid designation.[6]

This position may be characterized syntactically as one that demands that, in a language adequate for expressing all facts, the primitive vocabulary contains a 1-place predicate '$N_a x$' for each envisaged trans-world object a.[7] That a monadic property *can* do duty for a trans-world object—or, to rephrase in terms of the object-language, that a monadic predicate can do duty for a rigid designator—without any loss (or gain!) in expressive adequacy, is well known (cf. Quine (1960, §38). Given a haecceitistic predicate '$N_a x$', one can introduce the corresponding rigid designator by definition: $a := \imath x.N_a x$; and conversely, given a rigid designator 'a' and the

[6]Anti-quidditists (like Black 2000) can still use our framework, by populating the domain of objects with properties and relations, now treated as the value of first-order variables, and using the new predicate or predicates 'has', so that 'Fa' becomes 'a has F' (cf. Lewis 1970b). Anti-quidditism may then amount to what I later (Chapter 3) call 'anti-haecceitism', but applied to these hypostatized qualities. However, this purported anti-quidditism must be "quidditist" about the 'has' relation(s), a situation that clearly cannot be remedied by hypostatizing again, on pain of initiating a Bradley-like regress. There is no space to pursue the issue here. But we endorse the view of Lewis (2002) that 'has' is a 'non-relational tie', and speculate that, if it is treated *like* a relation in the logic, then it is better off treated as one whose identity across worlds is not in question—that is: treated "quidditistically".

[7]To simplify, I assume that a haecceitist is a haecceitist about every object. This leaves some (uninteresting) logical space between haecceitism and what we later characterise as anti-haecceitism (see the end of Section 3.1).

identity predicate '=', one can likewise introduce the corresponding haecceitistic predicate: $\forall x(N_a x \equiv x = a)$. We therefore urge the view that the difference between the two stronger versions of haecceitism is *merely notational*, meaning that the addition of the haecceitistic predicate '$N_a x$' to one's primitive vocabulary ontologically commits one to no more, and no less, than the addition of the name 'a', together with all the machinery of rigid designation.[8]

From now on, I will take haecceitism in a sense stronger than the first, Lewisian version. The second and third, stronger versions are equivalent, but I favour the notational trappings of the third version, i.e. the acceptance of haecceitistic properties. Thus when I later ban names from the object-language (in Section 2.2.1), this ought to be seen *not* as a substantial restriction against the haecceitist, but merely a convenient narrowing of notational options for the sake of a more unified presentation. We simply require the haecceitist to express her position though the adoption of haecceitistic predicates, though all are at liberty to—indeed, all *should*—read each instance of the haecceitist's '$N_a x$' as '$x = a$'.

A concern may remain: how can ontological commitment to a *property* be equivalent to ontological commitment to a (trans-world) *object*? There is no mystery, once we lay down some principles for what ontological commitment to properties involves. I take it that ontological commitment to a collection of properties, relations and objects entails a commitment to all the logical constructions thereof, because instantiation of the logical constructions can be defined away without residue in terms of instantiation and the existence of the original collection. (For example, for any object, the complex predicate '$Fx \wedge Gx$' is satisfied by that object

[8] By saying that I intend the acceptance of haecceitistic properties to be equivalent to combinatorial independence, I do not intend to rule out as counting as a (strong) version of the acceptance of haecceitistic properties the view described above, viz. that the haecceities are in some sense ontologically "prior" to the objects, and serve to "ground their identity" (whatever that may mean). That view would *entail* combinatorial independence between objects and other objects, and between objects and the qualitative properties and relations (though not, of course, between objects and their haecceities), and would equally well be served by rigid designators as by the addition of haecceitistic predicates to the primitive vocabulary. The point is that the second and third versions of haecceitism match completely in logical strength, so they are equally consistent with more metaphysically ambitious views which seek to "ground" trans-world identity in non-qualitative properties.

just in case both 'Fx' and 'Gx' are satisfied by it.) So ontological commitment to logical constructions is no *further* commitment at all. Now, ontological commitment to certain objects is revealed clearly enough: one need only peer into the domain of quantification to see if they are there. Ontological commitment to complex properties and relations is equally straightforward, being a matter of commitment to their components. But what about ontological commitment to the properties and relations taken as *primitive*, i.e. not as logical constructions—what does that involve? Well, since ontological commitment to objects is clear enough, let us use that: let us say that ontological commitment to the primitive properties and relations is a commitment to their being *instantiated* by some object or objects.

With these principles laid down, we can now prove that commitment to the trans-world continuant a and commitment to the haecceitistic property, *being* a, entail each other. *Left to right:* We can take two routes. First route: Commitment to any object at all (i.e. a non-empty domain) entails commitment to the identity relation, which is instantiated by everything. So commitment to the trans-world continuant a entails commitment to the identity relation, since $a = a$. But the property *being* a is just a logical construction out of the identity relation and the trans-world continuant a, so commitment to the trans-world continuant a entails commitment to the property *being* a. Second route: Commitment to the trans-world continuant a entails commitment to the property *being* a being instantiated. But that is just to say that commitment to the trans-world continuant a entails commitment to the property *being* a. *Right to left:* Commitment to the property *being* a entails either: (i) a commitment to its being instantiated (if taken as primitive); or (ii) a commitment to the entities of which it is a logical construction (if taken as a logical construction). If (i), then we are committed to something's being a, that is, the existence of a. If (ii), and 'being a' is understood properly as containing a rigid designator, not as an abbreviated definite description à la Russell, then we are committed to its logical components, i.e. the identity relation, and a itself. *QED.*

To sum up: the acceptance of haecceitistic differences (the first version of haecceitism, above) need not commit one to either combinatorially independent trans-world objects (our second version), or to non-qualitative properties that could do duty for them (our third version), though the two latter doctrines are perhaps the most natural way of securing haecceitistic differences, and may themselves be considered as notational variants of each other. I stipulate that we mean by 'haecceitism' one of the stronger versions, and for reasons purely to do with uniformity of presentation, I stipulate that the advocacy of this stronger version of haecceitism be expressed through the acceptance of haecceitistic properties. This version of haecceitism will be further developed, along with three other metaphysical theses, in Chapter 3.

2.2 A logical perspective on identity

So to sum up, our aim is to use the Hilbert-Bernays account as a spring-board so as to give a precise "logical geography" of discernibility. This logical geography will be in terms of the syntax of a formal first-order language. But we will also relate our definitions to the idea of permutations on the domain of quantification, and to the idea of these permutations being symmetries. These relations seem not to have been much studied in the recent philosophical literature about the Hilbert-Bernays account; and we will see that they turn out to be subtle—some natural conjectures are false.[9]

2.2.1 The Hilbert-Bernays account

What I will call the Hilbert-Bernays account of the identity of objects, treated as (the values of variables) in a first-order language, goes as follows; (cf. Hilbert and Bernays (1934, §5) : who in fact do not endorse it!), Quine (1970, pp. 61-64) and Saunders (2003a, p. 5)). The idea is that there being only finitely many primitive predicates enables us to capture the idea of identity in a single axiom. In fact, the axiom is a biconditional in which identity is equivalent to a long conjunction of statements that predicates are co-instantiated. The conjunction exhausts, in

[9]We thank N. da Costa and J. Ketland for alerting us to their results in this area, which we (culpably!) had missed. Further references below.

a natural sense, the predicates constructible in the language; and it caters for quantification in predicates' argument-places other than the two occupied by x and y.

In detail, we proceed as follows. Suppose that F^1_i is the ith 1-place predicate, G^2_j is the jth 2-place predicate, and H^3_k is the kth 3-place predicate; (we will not need to specify the ranges of i, j, k). Suppose that the language has no names, or function symbols, so that predicates are the only kind of non-logical vocabulary. Then the biconditional will take the following form:

$$\begin{aligned}\forall x\forall y\Big\{ x = y \;\equiv\; \Big[\;& \ldots \wedge (F^1_i x \equiv F^1_i y) \wedge \ldots \\ & \ldots \wedge \forall z\big((G^2_j xz \equiv G^2_j yz) \wedge (G^2_j zx \equiv G^2_j zy)\big) \wedge \ldots \\ & \ldots \wedge \forall z\forall w\Big((H^3_k xzw \equiv H^3_k yzw) \wedge (H^3_k zxw \equiv H^3_k zyw) \\ & \qquad \wedge (H^3_k zwx \equiv H^3_k zwy)\Big) \wedge \ldots \Big]\Big\}\end{aligned} \tag{HB}$$

(For primitive predicates, I will usually omit the brackets and commas often used to indicate argument-places.) Note that for each two-place predicate, there are two biconditionals to include on the right-hand side; and similarly for a three-place predicate. The general rule is: n biconditionals for an n-place predicate.[10]

This definition of the Hilbert-Bernays account prompts three comments.

1. *Envisaging a rich enough language*:— The main comment is the obvious one: since the right hand side (in square brackets) of (HB) defines an equivalence relation—which from now on I will call 'indiscernibility' (or for emphasis: 'indiscernibility by the primitive vocabulary')—discussion is bound to turn on the issue whether this relation is truly identity of the objects in the domain. Someone who advocates (HB) is envisaging a vocabulary rich enough

[10]We do not need to include explicitly on the right-hand side clauses with repeated instances of x or y, such as $(G^2_j xx \equiv G^2_j yy)$, since these clauses are implied by the conjunction of other relevant biconditionals. For example, from $\forall z(G^2_j xz \equiv G^2_j yz)$ we have, in particular, that $(G^2_j xx \equiv G^2_j yx)$. And from $\forall z(G^2_j zx \equiv G^2_j zy)$ we have, in particular, that $(G^2_j yx \equiv G^2_j yy)$. It follows that $(G^2_j xx \equiv G^2_j yy)$. A similar chain of arguments applies for an arbitrary n-place predicate.

(or: a domain of objects that is varied enough, by the lights of the vocabulary chosen) to discern any two distinct objects, and thereby force the equivalence relation given by the right hand side of (HB) to be identity. This comment can be developed in six ways.

(i) Indiscernibility has the formal properties of identity expressible in first-order logic, i.e. being an equivalence relation and substitutivity. (Cf. Equation (2.1) in (i) of comment 3 below; and for details, Ketland (2009, Lemma 12)).

(ii) Let us call an interpretation of the language, comprising a domain D and various subsets of D, D^2 etc. a 'structure'; (since in the literature on identity, 'interpretation' also often means 'philosophical interpretation'). Then: if in a given structure, the identity relation is first-order definable, then it is defined by indiscernibility, i.e. the right hand side of (HB) (Ketland (2009, Theorem 16)).

(iii) On the other hand: assuming '$=$' is to be interpreted as the identity relation (cf. (i) in 3 below), there will in general be structures in which the leftward implication of (HB) fails; i.e. structures with at least two objects indiscernible from each other. We say 'in general' since from the view-point of pure logic, '$=$' might itself be one of the 2-place predicates G^2_j: in which case, the leftward implication trivially holds.

(iv) The last sentence leads to the wider question whether the interpretation given to some of the predicates F, G, H etc. somehow presupposes identity, so that the Hilbert-Bernays account's reduction of identity is, philosophically speaking, a charade. This question will sometimes crop up below (e.g. for the first time, in footnote 12); but I will not need to address it systematically. Here I just note as an example of the question, the theory of pure sets. It has only one primitive predicate, '$\in$', and the axiom (HB) for it logically entails the axiom of extensionality. But one might well say that this only gives a genuine reduction of identity if one can understand the intended interpretation of '$\in$' without a prior understanding of '$=$'.

(v) There is the obvious wider question, why I restrict my discussion to first-order languages. Here my main reply is twofold: (a) first-order languages are favoured by the incompleteness of higher-order logics, and I am anyway sympathetic to the view that first-order logic suffices for the formalisation of physical theories (cf. Boolos & Jeffrey (1974, p. 197); Lewis (1970b, p. 429)); and (b) the Hilbert-Bernays account is thus restricted—and though, as I emphasized, I do not endorse it, it forms a good spring-board for discussing kinds of indiscernibility. There are also some basic points about identity in second-order logic, which I should register at the outset. Many discussions (especially textbooks: e.g. van Dalen (1994, p. 151); Boolos & Jeffrey (1974, p. 280) take the principle of the identity of indiscernibles to be expressed by the second-order formula $\big(\forall P(Px \equiv Py) \supset x = y\big)$: which is a theorem of (any deductive system for) second-order logic with a sufficiently liberal comprehension scheme. But even if one is content with second-order logic, this result does not diminish the interest of the Hilbert-Bernays account (or of classifications of kinds of discernibility based on it). For the second-order formula is a theorem simply because the values of the predicate variable include singleton sets of elements of the domain (cf. Ketland (2006, p. 313)). And allowing such singleton sets as properties of course leads back both to haecceitism, discussed in Section 2.1, and to (iv)'s question of whether understanding the primitive predicates requires a prior understanding of identity. In any case, I will discuss the principle of the identity of indiscernibles from a philosophical perspective in Chapter 3.

(vi) Finally, there is the question what the proponent of the Hilbert-Bernays account is to do when faced with a language with infinitely many primitive predicates. The right hand side of (HB) is constrained to be finitely long, so in the case of infinitely many primitives it is prevented from capturing all the ways two objects can differ. Here I see two roads open to the Hilbert-Bernays advocate: (a) through parametrization, infinitely many n-adic primitives, say $R_1xy, R_2xy, \ldots$ may be subsumed under a single, new $(n+1)$-adic primitive, say $Rxyz$, where the extra argument place is intended to vary over an index set for the previous primitives; and (b) one may resort to infinitary logic, which allows for infinitely long formulas—suggesting, as regards (HB), in philosophical terms, a supervenience of identity rather than a reduction of it. (I will consider infinitary logic just once below, towards the end of Section 2.4.3.)

To sum up these six remarks: I see the Hilbert-Bernays account as intending a reduction of identity facts to qualitative facts—as proposing that there are no indiscernible pairs of objects. This theme will recur in what follows. Indeed, the next two comments relate to the choice of language.

2. *Banning names*:— From now on, it will be clearest to require the language to have no individual constants, nor function symbols, so that the non-logical vocabulary contains only predicates. But this will not affect my arguments: they would carry over intact if constants and function symbols were allowed.

   As I see it, only the haecceitist is likely to object to this apparent limitation in expressive power. But here, Section 2.1's discussion of haecceitism comes into its own. For even with no names, the haecceitist has to hand her thisness predicates $N_a x, N_b x$ etc., with which to refer to objects by definite descriptions. Thus I propose, following Quine (1960, §§37-39), that we, and in particular the haecceitist, replace proper names by 1-place predicates; (each with an accompanying uniqueness axiom; and with the predicates then shoe-horned into the syntactic form of singular terms, by invoking Russell's theory of descriptions).[11] And as emphasised in Section 2.1, $N_a x$ etc. are to be thinly construed: the predicate $N_a x$ commits one to nothing beyond what the predicate $a = x$ commits one to. Thus the presence of these predicates in the non-logical vocabulary means that the haecceitist should have no qualms about endorsing the Hilbert-Bernays account—albeit in letter, rather than in spirit.[12]

[11]Then any sentence $S =: \Phi(a)$ containing the name a, where the 1-place formula $\Phi(x)$ contains no occurrence of a, may be replaced by the materially equivalent sentence $\forall x\,(N_a x \supset \Phi(x))$, which contains no occurrence of the name a. Note also that Saunders (2006, p. 53) limits his inquiry to languages without names; but no Quinean trick is invoked. Robinson (2000, p. 163) calls such languages 'suitable'.

[12]For the spirit of the Hilbert-Bernays account is a *reduction* of identity facts to putatively

A point of terminology: Though I ban constants from the formal language, I will still use a and b as names in the *meta*-language (i.e. the language in which I write!) for the one (or two!) objects, with whose identity or diversity we are concerned. I will also always use '=' in the meta-language to mean identity!

3. *'=' as a logical constant?*:— It is common in the philosophy of logic to distinguish two approaches by which formal languages and logics treat identity:

   (i) '=' is a logical constant in the sense that it is required, by the definition of the semantics, to be interpreted, in any domain of quantification, as the identity relation. Then for any formula (open sentence) $\Phi(x)$ with one free variable x, the formula

   $$\forall x \forall y\, (x = y \supset (\Phi(x) \equiv \Phi(y))) \tag{2.1}$$

   is a logical truth (i.e. is true in every structure). Thus on this approach, the rightward implication in (HB) is a logical truth.

   But the leftward implication in (HB) is not. For as discussed in comment 1, the language may not be discerning enough. More precisely: there are structures in which (no matter how rich the language!) the leftward implication fails.[13]

   (ii) '=' is treated like any other 2-place predicate, so that its properties flow entirely from the theory with which we are concerned, in particular its axioms if it is an axiomatized theory. Of course, we expect our theory to impose on '=' such properties as being an equivalence relation

qualitative facts, about the co-instantiation of properties. Indeed: according to the approach in which '=' is not a logical constant (cf. comment 3(ii)), acceptance of (HB) entails that the language with the equality symbol '=' is a definitional extension of the language without it. However, representing each haecceity by a one-place primitive predicate, accompanied by a uniqueness axiom, assumes the concept of identity through the use of '=' in the axiom: making this reduction of identity, philosophically speaking, a charade. By contrast, if only genuinely qualitative properties are expressed by the non-logical vocabulary, the philosophical reduction of identity facts to qualitative facts can succeed—provided, of course, that the language is rich enough or the domain varied enough.

[13] Such is Wiggins' (2001, pp. 184-185) criticism of the Hilbert-Bernays account as formulated by Quine (1960, (1970).

(cf. Ketland (2009, Lemmas 8-10)). We can of course go further towards capturing the intuitive idea of identity by imposing every instance of the schema, eq. 2.1. So on this approach, the Hilbert-Bernays account can be regarded as a proposed finitary alternative to imposing eq. 2.1, a proposal whose plausibility depends on the language being rich enough. Of course, independently of the language being rich: imposing (HB) in a language with finitely many primitives implies as a theorem the truth of eq. 2.1.[14]

To sum up this Section: the conjunction on the right hand side of (HB) makes vivid how, on the Hilbert-Bernays account, objects can be distinct for different reasons, according to which conjunct fails to hold. In Section 2.3, I will introduce a taxonomy of these kinds. In fact, this taxonomy will distinguish different ways in which a single conjunct can fail to hold. The tenor of that discussion will be mostly syntactic. So I will complement it by first discussing identity, and the Hilbert-Bernays account, in terms of permutations on the domain of quantification.

2.2.2 Permutations on domains

I will now discuss how permutations on a domain of objects can be used to express qualitative similarities and differences between the objects. More precisely: I will define what it is for a permutation to be a *symmetry* of an interpretation of the language, and relate this notion to the Hilbert-Bernays account.

Definition of a symmetry

Let D be a domain of quantification, in which the predicates F^1_i, G^2_j, H^3_k etc. get interpreted. So writing 'ext' for 'extension', $\text{ext}(F^1_i)$ is, for each i, a subset of D; and $\text{ext}(G^2_j)$ is, for each j, a subset of $D \times D = D^2$; and $\text{ext}(H^3_k)$ is, for each k, a subset of D^3 etc. For the resulting interpretation of the language, i.e. D together

[14] Hilbert and Bernays (1934, p. 186) show that '=' as defined by (HB) is (up to co-extension!) the *only* (non-logical) two-place predicate to imply reflexivity and every instance of the schema, eq. 2.1. The argument is reproduced in Quine (1970, pp. 62-63).

with these assigned extensions, we write $\mathcal{I}$. I shall also call such an interpretation, a 'structure'.

Now let π be a permutation of D.[15] I now define a symmetry as a permutation that "preserves all properties and relations". So I say that π is a *symmetry* (aka *automorphism*) of $\mathcal{I}$ iff all the extensions of all the predicates are invariant under π. That is, using $o_1, o_2, ...$ as meta-linguistic variables: π is a symmetry iff:

$$\begin{array}{rrcl} \forall o_1, o_2, o_3 \in D, \forall i,j,k: & o_1 \in \mathrm{ext}(F_i^1) & \text{iff} & \pi(o_1) \in \mathrm{ext}(F_i^1); \text{ and} \\ & \langle o_1, o_2 \rangle \in \mathrm{ext}(G_j^2) & \text{iff} & \langle \pi(o_1), \pi(o_2) \rangle \in \mathrm{ext}(G_j^2); \text{ and} \\ & \langle o_1, o_2, o_3 \rangle \in \mathrm{ext}(H_k^3) & \text{iff} & \langle \pi(o_1), \pi(o_2), \pi(o_3) \rangle \in \mathrm{ext}(H_k^3); \end{array} \tag{2.2}$$

and similarly for predicates with four or more argument-places.[16]

Relation to the Hilbert-Bernays account

Let me now compare this definition with the Hilbert-Bernays account. For the moment, I make just two comments, (1) and (2) below. They dispose of natural conjectures, about symmetries leaving invariant the indiscernibility equivalence classes. That is, the conjectures are false: the first conjecture fails because of the somewhat subtle notions of weak and relative discernibility, which will be central later; and the second is technically a special case of the first.[17] In Section 2.4, I will discuss how the conjectures can be mended: roughly speaking, we need to replace indiscernibility by a weaker and "less subtle" notion, called 'absolute indiscernibility'.

[15] I note *en passant* that since a permutation is a bijection, the definition of permutation involves the use of the '=' symbol; so that which functions are considered to be permutations is subject to one's treatment of identity. But no worries: as I noted in comment 2 of Section 2.2.1, this definition is cast in the meta-language!

[16] So a haecceitist (cf. comment 2 in Section 2.2.1) will take only the identity map as a symmetry—unless they stipulate that haecceitistic properties are exempt from the definition of symmetry.

[17] My two comments agree with Ketland's results and examples (2006, Theorem (iii) and example in footnote 17) or (2009, Theorem 35(a) and example). But I will not spell out the differences in jargon or examples, except to report that Ketland calls a structure 'Quinian' iff the leftward implication of (HB) holds in it, i.e. if the identity relation is first-order definable, and so (cf. comment 1(ii) of Section 2.2.1) defined by indiscernibility, i.e. by the right hand side of (HB).

(1): *All equivalence classes invariant?*:— Suppose that we adopt the relativization to an arbitrary language: so we do not require the language to be rich enough to force indiscernibility to be identity. Then one might conjecture that, for each choice of language, a permutation is a symmetry iff it leaves invariant (also known as: *fixes*) each indiscernibility equivalence class. That is: π is a symmetry of the structure $\mathcal{I}$ iff each member of each equivalence class $\subset \mathrm{dom}(\mathcal{I})$ is sent by π to a member of that same equivalence class.

In fact, this conjecture is false. The condition, leaving invariant the indiscernibility equivalence classes, is stronger than being a symmetry. I first prove the true implication, and then give a counterexample to the converse.

So suppose π leaves invariant each indiscernibility class; and let $\langle o_1, o_2, ..., o_n\rangle$ be in the extension $\mathrm{ext}(J^n)$ of some n-place predicate J^n. Since π leaves invariant the indiscernibility class of o_1, it follows that $\langle \pi(o_1), o_2, ..., o_n\rangle$ is also in $\mathrm{ext}(J^n)$. (For if not, (HB)'s corresponding conjunct, i.e. the conjunct for the first argument-place of the predicate J^n, would discern o_1 and $\pi(o_1)$). From this, it follows similarly that since π fixes the indiscernibility class of o_2, $\langle \pi(o_1), \pi(o_2), ..., o_n\rangle$ is also in $\mathrm{ext}(J^n)$. And so on: after n steps, we conclude that $\langle \pi(o_1), \pi(o_2), ..., \pi(o_n)\rangle$ is in $\mathrm{ext}(J^n)$. Therefore π is a symmetry.

Philosophical remark: one way of thinking of the Hilbert-Bernays account trivializes this theorem. That is: according to comment 1 of Section 2.2.1, the proponent of this account envisages that the indiscernibility classes are singletons. So only the identity map leaves them all invariant, and trivially, it is a symmetry. (Compare the discussion of haecceitism in footnotes 12 and 16.)

Counterexample to the converse: Consider the following structure, whose domain comprises four objects, which we label a to d. The primitive non-logical vocabulary consists of just the 2-place relation symbol R. R is interpreted as having the extension

$$\mathrm{ext}(R) = \{\langle a,b\rangle, \langle a,c\rangle, \langle a,d\rangle, \langle b,a\rangle, \langle b,c\rangle, \langle b,d\rangle\}\ .$$

Now let '=' be defined by the Hilbert-Bernays axiom (HB). From this the reader can check that '=' has the extension

$$\text{ext}(=) = \{\langle a,a\rangle, \langle b,b\rangle, \langle c,c\rangle, \langle d,d\rangle, \langle d,c\rangle, \langle c,d\rangle\} \ ;$$

i.e. c and d are indiscernible. So '=' is *not* interpreted as identity (cf. the relativization of '=' to the language, discussed in 3(ii) of §2.2.1). The relation '=' carves the domain into three equivalence classes: $[a] = \{a\}, [b] = \{b\}$ and $[c] = [d] = \{c, d\}$.

Now consider the permutation π, whose only effect is to interchange the objects a and b; i.e. $\pi := \begin{pmatrix} abcd \\ bacd \end{pmatrix} \equiv (ab)$. This permutation is a symmetry, since it preserves the extension of R according to the requirement (2.2); yet it does *not* leave invariant the equivalence classes $[a]$ and $[b]$ (Fig. 2.1).[18]

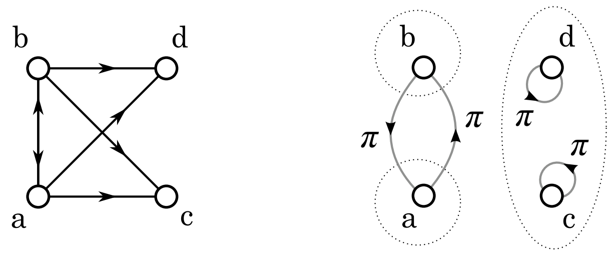

Figure 2.1: A counterexample to the claim that symmetries leave invariant the indiscernibility classes. In this structure, drawn on the left, the permutation π which swaps a and b is a symmetry; but a and b form their own separate equivalence classes under '=', as shown on the right.

(2): *Only the trivial symmetry?*:— Suppose now that in $\mathcal{I}$, the predicate '=' *is* interpreted as identity; and that (HB) holds in $\mathcal{I}$. So we are supposing that the objects are various enough, the language rich enough, that indiscernibility in $\mathcal{I}$ is identity. Or in other words: the indiscernibility classes are singletons. On these suppositions, one might conjecture that that the only symmetry is the trivial one, i.e. the identity map $id : D \to D$. (Such structures are often called 'rigid' (Hodges

[18]Two remarks. (1): Agreed, this counterexample could be simplified. A structure with just a and b, with $\text{ext}(R) = \{\langle a,b\rangle, \langle b,a\rangle\}$ has a, b discernible, but the swap $a \mapsto b, b \mapsto a$ is a symmetry. But this example will also be used later. (2) Accordingly, in the counterexample, $(ab)(cd)$ would work equally well: i.e. it also is a symmetry that does not leave invariant $[a]$ and $[b]$. Looking ahead: Theorem 1 in Section 2.4.1 will imply that $\{a, b\}$ and $\{c, d\}$ are each subsets of absolute indiscernibility classes; in fact, each is an absolute indiscernibility class.

(1997, p. 94)).)

In fact, this is false. Comment (1) has just shown that there are symmetries that do not leave invariant the indiscernibility classes. Our present suppositions have now collapsed these indiscernibility classes into singletons. So there will be symmetries which do not leave invariant the singletons, i.e. are not the identity map on D. To be explicit, consider the structure, and its indiscernibility classes, drawn in Fig. 2.2.

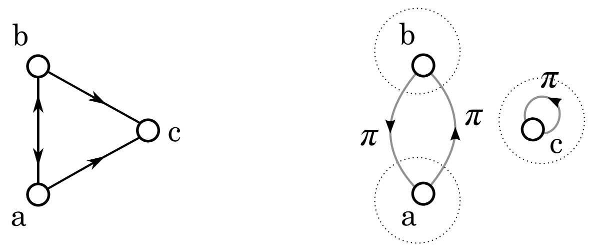

Figure 2.2: A counterexample to the claim that if the indiscernibility classes are singletons, the only symmetry is the identity map. In this structure, the indiscernibles of Figure 2.1 have been identified. a and b are still distinct (they are discernible) but are swapped by the symmetry π.

2.3 Four kinds of discernment

I turn to defining the different ways in which two distinct objects can be discerned in a structure. These kinds of discernment will be developed (indirectly) in terms of which conjuncts on the RHS of the Hilbert-Bernays axiom (HB) are false in the structure.[19]

[19]Other authors, notably Muller and Saunders 2008), consider the discernment of two objects *by a theory* (say T), so that the RHS of (HB), applied to the two objects in question, is a theorem of T; whereas our concern is the discernment of two objects *in a structure*. My focus on structures is necessary: given our ban on names, the sentence expressing the satisfaction of the RHS of (HB) by the two objects cannot even be written!

2.3.1 Three preliminary comments

(1) *Precursors*: Broadly speaking, my four kinds of discernment follow the discussions by Quine (1960, 1970, 1976a) and Saunders (2003a, 2003b, 2006). Quine (1960, p. 230) endorses the Hilbert-Bernays account of identity and then distinguishes what he calls absolute and relative discernibility. His absolute discernibility will correspond to (the disjunction of) my first two kinds—and I will follow him by calling this disjunction 'absolute'. Besides, his relative discernibility will correspond to (the disjunction of) my third and fourth kinds. But I will follow Saunders ((2003a, p. 5); (2003b, pp. 19-20); (2006, p. 5)) by reserving 'relative' for the third kind, and using 'weak' for the fourth. (So for me 'non-absolute' will mean 'relative or weak'.)[20]

(2) *Suggestive labels*: I will label these kinds with words like 'intrinsic' which are vivid, but also connote metaphysical doctrines and controversies (e.g. Lewis 1986, pp. 59-63). I disavow the connotations: the official meaning is as defined, and so is relative to the interpretation of the non-logical vocabulary.

(3): *Two pairs yield four kinds*: The four kinds of discernment arise from two pairs. We begin by distinguishing between a formula with one free variable (labelled **1**) and a formula with two free variables (labelled **2**).[21] Each of these cases is then broken down into two subcases (labelled **a** and **b**) yielding four cases in all: labelled **1(a)** to **2(b)**. We will also give the four cases mnemonic labels: e.g. **1(a)** will also be called **(Int)** for 'intrinsic'.

The intuitive idea that distinguishes sub-cases will be different for **1** and **2**. For **1**, the idea is to distinguish whether discernibility depends on a relation to another object; while for **2**, the idea is to distinguish whether discernibility depends on an asymmetric relation. Both these ideas are semantic, and even a bit vague. But the

[20] Quine (1976a, p. 113) defines what he calls grades of discriminability, which is a spectrum of strength. Saunders (2006, pp. 19-20) agrees that there is such a spectrum of strength, although in his (2003a, p. 5) he makes the three categories 'absolute', 'relative' and 'weak' mutually exclusive. I say 'kinds' not 'categories' or 'grades' to avoid the connotation of mutual exclusion or a spectrum of strength.

[21] We thank Leon Horsten for the observation that the notion of discernibility may be parameterized to other objects (so that, e.g., we might say that a is discernible from b relative to $c, d, \ldots$), which would involve formulas with more than two variables. The idea seems to us workable, but we will not pursue it here.

definitions of the sub-cases will be syntactic, and precise—and we will therefore remark that they do not completely match the intuitive idea.

As announced in comment 2 of Section 2.2.1, a and b will be names in the meta-language (in which I am writing) for the one or two objects with which we are concerned.

2.3.2 The four kinds defined

1(a) (Int) *1-place formulas with no bound variables, which apply to only one of the two objects a and b.* This of course covers the case of primitive 1-place predicates, e.g. Fx: so that for example, $a \in \text{ext}(F)$ but $b \notin \text{ext}(F)$, or vice versa. But we also intend this case to cover 1-place formulas arising by slotting into a polyadic formula more than one occurrence of a single free variable, while the polyadic formula nevertheless does *not* contain *any* bound variables.

The intent here is to exclude formulas which quantify over objects other than the two we are concerned with. So this case will include formulas such as: Rxx and $Fx \wedge Hxxx$; (R, H primitive 2-place and 3-place predicates, respectively; not abbreviations of more complex open sentences). But it will exclude formulas such as $\forall z(Fz \supset Rzx)$, which contain bound variables.[22]

I will say that two objects that do not share some monadic formula in this sense are discerned *intrinsically*, since their distinctness does not rely on any relation either object holds to any other. An everyday example, taking 'is spherical' as a primitive 1-place predicate, is given by a ball and a die. Another example, with Rxy the primitive 2-place predicate 'loves' (so that Rxx is the 1-place predicate 'loves his- or herself') is *Narcissus* $\in \text{ext}(R**)$,

[22]Recall my ban on individual constants ((2) in Section 2.2.1). If we had instead allowed them, this sub-case **1(a)** would be defined so as to also exclude all formulas containing any constant, including a and b. The exclusion of formulas such as Rcx, which refer to a third object, is obviously desirable, given the intuitive idea of discernment by intrinsic properties. However, the exclusion of formulas involving only the constants a and-or b may be more puzzling. My rationale is that, for any formula of the type Rax, Rbx, etc. which is responsible for discerning two objects, there will be an alternative formula (either Rxx or Rxy) which we would instead credit for the discernment, and which falls under one of the three other kinds.

but (alas) *Echo* $\notin$ ext$(R**)$.

I shall say that any pair of objects discerned other than by **1(a)**—i.e. discerned by **1(b)**, or by **2(a)**, or by **2(b)** below—are *extrinsically* discerned.

1(b) **(Ext)** *1-place formulas with bound variables, which apply to only one of the two objects a and b.* That is: this kind contains polyadic formulas that *do* contain bound variables. So it contains formulas such as $\forall z(Fz \supset Rzx)$. And a and b are discerned by such formulas if: for example, $a \in \text{ext}(\forall z(Fz \supset Rz*))$ but $b \notin \text{ext}(\forall z(Fz \supset Rz*))$. I will say that they have been discerned *externally.* An example, taking F = 'is a man', R = 'admires' is: *Cleopatra* $\in \text{ext}(\forall z(Fz \supset Rz*))$ but *Caesar* $\notin \text{ext}(\forall z(Fz \supset Rz*))$. (Recall that I reserve the term *extrinsic* to cover all three kinds **1(b), 2(a), 2(b)**: so external discernment is more specific than extrinsic.)

The intent is that in this kind of discernment, diversity follows from the relations the two objects a and b have to *other* objects. However, as I said in (3) of Section 2.3.1, the precise syntactical definition cannot be expected to match exactly the intuitive idea. And indeed, there are examples of external discernment where the relevant value of the bound variable in the discerning formula is in fact a or b, even though this is invisible from the syntactic perspective. (In the example just given, the universal quantifier in $\forall z(Fz \supset Rzx)$ quantifies over a domain that includes Caesar himself.)[23]

I will say that two objects discerned by a formula either of kind **(Int)** or of kind **(Ext)** are *absolutely discerned.* Note that a pair of objects could be both intrinsically and externally discerned. But since **(Ext)** is intuitively a "weaker" form of discernment, I shall sometimes say that a pair of objects that are externally, but not intrinsically, discerned, are *merely* externally discerned.

Interlude: Individuality and absolute discernment. I will also say that an object that is absolutely discerned from *all* other objects is an *individual* or *has individu-*

[23]As to my ban on individual constants: if we had instead allowed them, this sub-case **1(b)** would be defined so as to also exclude formulas containing a and b, but to allow other constants c, d etc., so as to capture the idea of discernment by relations to other objects. But as in the case of bound variables, it could turn out that the "third" object picked out is in fact a or b. Cf. footnote 22.

ality. Note that if an object is an individual, some or all of the other objects might themselves fail to be individuals (cf. figure 2.3).

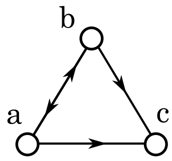

Figure 2.3: An object's being an individual requires its being absolutely discerned from *all* others, but not their being absolutely discerned from anything else. Here, c is absolutely discerned from both a and b, e.g. by the formula $\forall z(Rzx \supset Rxz)$, while a and b are themselves non-individuals.

Being an individual is tantamount to being the bearer of a uniquely instantiated definite description: where 'tantamount to' indicates a qualification. The idea is: given an individual, we take *seriatim* the formulas that absolutely discern it from the other objects in the structure, and conjoin them and so construct a definite description that is instantiated only by the given individual. The qualification is that in an infinite domain, there could be infinitely many ways that a given individual was absolutely discerned from all the various others: think of how a finite vocabulary supports arbitrarily long formulas, and so denumerably many of them. Thus in an infinite domain the above "*seriatim*" procedure might yield an infinite conjunction—preventing a finitely long uniquely instantiated definite description.[24]

I will examine absolute discernibility in Section 2.4. For the moment, I return to our four kinds of discernibility: i.e. to presenting the last two kinds. *End of Interlude.*

2(a) (Rel) *Formulas with two free variables, which are satisfied by the two objects a and b in one order, but not the other.* For example, for the formulas Rxy and $\exists zHxzy$, we have: $\langle a, b\rangle \in \mathrm{ext}(R)$, but $\langle b, a\rangle \notin \mathrm{ext}(R)$; and $\langle a, b\rangle \in$

[24]A terminological note: Saunders (2003b, p. 10) says that an object that is the bearer of a uniquely instantiated definite description is 'referentially determinate', and Quine (1976a, p. 113) calls such an object 'specifiable'. So, modulo my qualification about infinite domains, these terms correspond to my (and Muller & Saunders' (2008)) use of 'individual'. Cf. also comment 1(vi) in Section 2.2.1.

ext($\exists zH * z\bullet$), but $\langle b,a\rangle \notin$ ext($\exists zH * z\bullet$). Here the diversity of a and b is an extrinsic matter (both intuitively, and according to my definition of 'extrinsic', which is discernment by any means other than **(Int)**), since it follows from their relation to each other. But it is not a matter of a relation to any third object. Following Quine (1960, p. 230), I will say that objects so discerned are *relatively* discerned.

And as above, I will say that objects that are relatively discerned but neither intrinsically nor externally discerned, are *merely* relatively discerned. Merely relatively discerned objects are never individuals in my sense (viz. absolutely discerned from all other objects).[25]

2(b) (Weak) *Formulas with two free variables, which are satisfied by the two objects a and b taken in either order, but not by either object taken twice.* For example, for the formulas Rxy and $\exists zHxzzy$, we have: $\langle a,a\rangle, \langle b,b\rangle \notin$ ext(R), but $\langle a,b\rangle, \langle b,a\rangle \in$ ext(R); and $\langle a,a\rangle, \langle b,b\rangle \notin$ ext($\exists zH * z\bullet$), but $\langle a,b\rangle, \langle b,a\rangle \in$ ext($\exists zH * z\bullet$). (We say 'but not by *either* object taken twice' to prevent a and b being intrinsically discerned.) Again, the diversity of a and b is extrinsic, but does not depend on a third object; rather diversity follows from their pattern of instantiation of the relation R. I call objects so discerned *weakly* discerned. And I will say that objects that are weakly discerned but neither intrinsically nor externally nor relatively discerned (i.e. fall outside **(Int), (Ext), (Rel)** above), are *merely* weakly discerned.

Objects which are discerned *merely* weakly are not individuals, in my sense (since they are not absolutely discerned). Max Black's famous example of two spheres a mile apart (1952, p. 156) is an example of two such objects. For the two spheres bear the relation 'is a mile away from', one to another; but not each to itself. The irony is that Black, apparently unaware of weak discernibility, proposes his duplicate spheres as a putative example of two ob-

[25] Note that this kind of discernment does not require the discerning formula, for example Gxy, to be asymmetric for all its instances; i.e., we do not require $\forall x\forall y(Gxy \supset \neg Gyx)$. In this I agree with Saunders (2003a, p. 5) and Quine (1960, p. 230). My rationale is that intuitively, this kind of discernment does not require anything about the global pattern of instantiation of the relation concerned. The same remark applies to **(Weak)** below, where now I differ from Saunders (2003a, p. 5), who demands that the discerning relation be irreflexive. Nevertheless, I adopt Saunders'/Quine's word 'weak'.

jects that are qualitatively indiscernible (and therefore as a counterexample to the principle of the identity of indiscernibles).[26]

I will also say that two objects that are not discerned by any of our four kinds (i.e. by no 1-place or 2-place formula whatsoever) are *indiscernible*. For emphasis, I will sometimes call such a pair *utterly indiscernible*. In particular, I will say 'utter indiscernibility' when contrasting this case with the failure of only one (or two or three) of our four kinds of discernibility. Of course, utter indiscernibles are only accepted by someone who denies the Hilbert-Bernays account.

2.4 Absolute indiscernibility: some results

In Section 2.2.2, we saw that a permutation leaving invariant the indiscernibility classes must be a symmetry; then we gave a counterexample to the converse statement, and to a related conjecture that if indiscernibility is identity, there is only the trivial symmetry. But now that I have defined absolute discernibility (viz. as the disjunction, **(Int)** or **(Ext)**), we can ask about the corresponding claims that use instead the absolute concept. That is the task of this Section. (But its results are hardly needed for the discussions and results in later Sections.)

I will prove that with the absolute concept, Section 2.2.2's converse statement is "resurrected", i.e. a symmetry leaves invariant the absolute indiscernibility classes (Section 2.4.1). Then I will give some illustrations, including a counterexample to the converse of *this* statement (Section 2.4.2). Then I will show that for a *finite* domain of quantification, absolute indiscernibility of two objects is equivalent to the existence of a symmetry mapping one object to the other (Section 2.4.3).[27]

[26] That is, assuming that space is not closed. In a closed universe, an object may be a non-zero distance from itself, so the relation 'is one mile away from' is not irreflexive, and cannot be used to discern. French's (2006, §4) and Hawley's (2009, p. 109) charge of circularity against Saunders (2003a) enters here: it seems that the irreflexivity of 'is one mile away from' relies on the prior guarantee that the two spheres are indeed distinct; but their distinctness is supposed, in turn, to be grounded by that very relation being irreflexive. The openness or closedness of space would decide the matter, of course, but that too seems to stand or fall with the irreflexivity or otherwise of distance relations—between spatial or spacetime points, if not material objects. Cf. also comment 1(iv) in Section 2.2.1.

[27] Absolute discernibility and individuality are closely related to *definability* in a formal lan-

But first, beware of an ambiguity of English. For relations of indiscernibility, we have a choice of two usages. Should we use 'absolute indiscernibility' for just 'not absolutely discernible' (which will therefore include pairs of objects that are discernible, albeit by other means than absolutely)? Or should we use 'absolute indiscernibility' for some kind (species) of indiscernibility—as, indeed, the English adjective 'absolute' connotes? (And if so, which kind should we mean?)[28]

I stipulate that I mean the *former*. Then: since absolute discernibility is a kind of (implies) discernibility, we have, by contraposition: indiscernibility implies absolute indiscernibility (in my usage). Since both indiscernibility and absolute indiscernibility are equivalence relations, this implies that the absolute indiscernibility classes are unions of the indiscernibility classes; cf. Figure 2.4. With this definition, Section 2.4.1's theorem will be: a symmetry leaves invariant the absolute indiscernibility classes.

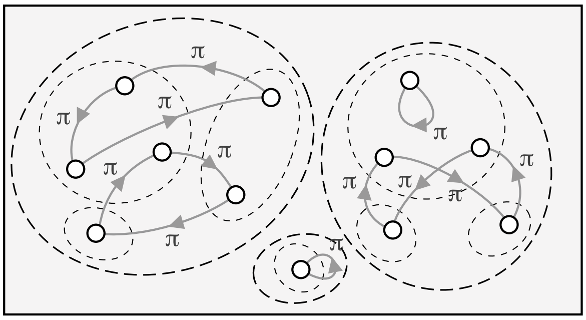

Figure 2.4: Preview to Section 2.4.1's theorem. A generic symmetry π acting on a domain must preserve the absolute indiscernibility classes (thick broken lines), but may break the indiscernibility classes (thin broken lines). Note also that the object at centre-bottom, alone in its absolute indiscernibility class and therefore an individual, must be sent to itself under π.

Besides, for later use, I make the corresponding stipulation about the phrases 'intrinsic indiscernibility', 'external indiscernibility' etc. That is: by 'intrinsic

guage; (for example, an object that is definable is an individual in the sense of Section 2.3.2's Interlude). The interplay between definability (and related notions) and invariance under symmetries is given a sophisticated treatment by da Costa & Rodrigues (2007), who consider higher-than-first-order structures. Some of their results have close affinities with our two theorems; in particular their theorems 7.3-7.7. As in footnote 9, I thank N. da Costa.

[28]This sort of ambiguity is of course not specific to discernment: it is common enough: should we read 'recalcitrant immobility' as 'not-(recalcitrant mobility)' or as 'recalcitrant not-mobility'?

indiscernibility' and 'intrinsically indiscernible', I will mean 'not-(intrinsic discernibility)' and 'not-(intrinsically discernible)', respectively; and so on for other phrases.

2.4.1 Invariance of absolute indiscernibility classes

Recall that in Section 2.2.2, we saw that leaving invariant (fixing) the indiscernibility classes was sufficient, but not necessary, for being a symmetry. That is: Section 2.2.2's counterexample showed that being a symmetry is not sufficient for leaving invariant the indiscernibility classes. This situation prompts the question, what being a symmetry *is* sufficient for. More precisely: is there a natural way to weaken Section 2.2.2's sufficient condition for being a symmetry—viz. indiscernibility invariance—into being instead a necessary condition? In other words: one might conjecture that leaving invariant some supersets of the indiscernibility classes yields a necessary condition of being a symmetry.

In fact, my concept of absolute indiscernibility is the natural weakening. (N.B. The ban on names is essential to its being a weakening: allowing names in a discerning formula makes (the natural redefinition of) absolute discernment *equivalent* to weak discernment.) That is: being a symmetry implies leaving invariant the absolute indiscernibility classes. Cf. Figure 2.4. I will first prove this, and then give a counterexample to the converse statement: it will be similar to the counterexample used in Section 2.2.2 against that Section's converse statement.

Theorem 1: For any structure (i.e. interpretation of a first-order language): if a permutation is a symmetry, then it leaves invariant the absolute indiscernibility classes.

Proof: I will prove the contrapositive: I assume that there is some element a of the domain which is absolutely discernible from its image $b := \pi(a)$ under the permutation π, and I prove that π is not a symmetry. (Remember that 'a' and 'b' are names in the metalanguage only; we stick to our ascetic object-language demands set down in comment 2 of Section 2.2.1.[29]) So our assumption is that

[29] I am extremely grateful to Leon Horsten for making me aware of the problems with a previous version of this proof, in which names were reintroduced into the object language.

for some object a in the domain, with $b := \pi(a)$, there is some formula $\Phi(x)$ with one free variable for which $a \in \mathrm{ext}(\Phi)$ while $b \notin \mathrm{ext}(\Phi)$, or vice versa ($a \notin \mathrm{ext}(\Phi)$ while $b \in \mathrm{ext}(\Phi)$):

The proof proceeds by induction on the logical complexity of the absolutely discerning formula Φ. From the assumption that $a \in \mathrm{ext}(\Phi)$ iff $b \notin \mathrm{ext}(\Phi)$, I will show that, whatever the main connective or quantifier used in the last stage of the stage-by-stage construction of Φ, there is some logically simpler open formula, perhaps with more than one (n, say) free variable, $\Psi(x_1, \ldots x_n)$, and some objects $o_1, \ldots o_n \in D$ (not necessarily including a, and not necessarily n in number, since maybe $o_i = o_j$ for some $i \neq j$) such that we have: $\langle o_1, \ldots o_n\rangle \in \mathrm{ext}(\Psi)$ iff $\langle \pi(o_1), \ldots \pi(o_n)\rangle \notin \mathrm{ext}(\Psi)$, where $\pi(o_i)$ is the image of o_i under the permutation π. That is, we continue to break Ψ down to logically simpler formulas until we obtain some *atomic* formula whose differential satisfaction by some sequence of objects and the sequence of their images under the permutation π directly contradicts π's being a symmetry.

The proof begins by setting $\Psi := \Phi(x)$ (so to start with, the adicity n of our formula equals 1 and our objects o_i comprise only a). We then reiterate the procedure until we reach an atomic formula. Thus:—

Step one. We have that $\langle o_1, \ldots o_n\rangle \in \mathrm{ext}(\Psi)$ iff $\langle \pi(o_1), \ldots \pi(o_n)\rangle \notin \mathrm{ext}(\Psi)$. (Remember that to start with, we set $n = 1$ and $o_1 = a$. And the 'iff' means only material equivalence.)

Step two. Proceed by cases:

- If $\Psi(x_1, \ldots x_n) = \neg\xi(x_1, \ldots x_n)$, then we have
  $\langle o_1, \ldots o_n\rangle \in \mathrm{ext}(\neg\xi)$ iff $\langle \pi(o_1), \ldots \pi(o_n)\rangle \notin \mathrm{ext}(\neg\xi)$; that is,
  $\langle o_1, \ldots o_n\rangle \in \mathrm{ext}(\xi)$ iff $\langle \pi(o_1), \ldots \pi(o_n)\rangle \notin \mathrm{ext}(\xi)$;
  so ξ is our new, simpler Ψ.

- If Ψ is a conjunction, then we can write
  $\Psi(x_1, \ldots x_n) = \big(\xi(x_{i(1)}, \ldots x_{i(l)}) \ \wedge\ \eta(x_{j(1)}, \ldots x_{j(m)})\big)$,
  where $l, m \leqslant n$ and $l + m \geqslant n$, and $i : \{1, 2, \ldots l\} \to \{1, 2, \ldots n\}$ and $j : \{1, 2, \ldots m\} \to \{1, 2, \ldots n\}$ are injective maps. First of all, we recognise that

$\langle o_1, \ldots o_n \rangle \in \text{ext}(\Psi) \quad \text{iff} \quad \langle o_{i(1)}, \ldots o_{i(l)} \rangle \in \text{ext}(\xi) \text{ and } \langle o_{j(1)}, \ldots o_{j(m)} \rangle \in \text{ext}(\eta)$.
Then, given step one, namely
$\langle o_1, \ldots o_n \rangle \in \text{ext}(\Psi)$ iff $\langle \pi(o_1), \ldots \pi(o_n) \rangle \notin \text{ext}(\Psi)$,
this is equivalent to
$\langle o_{i(1)}, \ldots o_{i(l)} \rangle \in \text{ext}(\xi)$ and $\langle o_{j(1)}, \ldots o_{j(m)} \rangle \in \text{ext}(\eta)$ iff
$\langle \pi\left(o_{i(1)}\right), \ldots \pi\left(o_{i(l)}\right) \rangle \notin \text{ext}(\xi)$ or $\langle \pi\left(o_{j(1)}\right), \ldots \pi\left(o_{j(m)}\right) \rangle \notin \text{ext}(\eta)$.
That is:
$\langle o_{i(1)}, \ldots o_{i(l)} \rangle \in \text{ext}(\xi)$ iff $\langle \pi\left(o_{i(1)}\right), \ldots \pi\left(o_{i(l)}\right) \rangle \notin \text{ext}(\xi)$, or
$\langle o_{j(1)}, \ldots o_{j(m)} \rangle \in \text{ext}(\eta)$ iff $\langle \pi\left(o_{j(1)}\right), \ldots \pi\left(o_{j(m)}\right) \rangle \notin \text{ext}(\eta)$.
So either the formula $\xi(x_1, \ldots x_l)$ or $\eta(x_1, \ldots x_m)$ is our new formula Ψ; with adicity l, respectively m, replacing the adicity n; and the objects $o_{i(1)}, \ldots o_{i(l)}$, respectively $o_{j(1)}, \ldots o_{j(m)}$ replacing the objects $o_1, \ldots o_n$. (This is an inclusive 'or': if either formula suffices, imagine that only one is chosen to continue the inductive procedure. Heuristic remarks: (i) It is only in this clause that the process can reduce the number of variables occurring in Ψ, and hence the number n of objects under consideration. (ii) The next three cases can be dropped in the usual way, if we suppose the language to use just $\neg, \wedge$ as primitive connectives.)

- If $\Psi = \big(\xi \vee \eta\big)$, then continue with $\Psi = \neg\big(\neg\xi \wedge \neg\eta\big)$.
- If $\Psi = \big(\xi \supset \eta\big)$, then continue with $\Psi = \neg\big(\xi \wedge \neg\eta\big)$.
- If $\Psi = \big(\xi \equiv \eta\big)$, then continue with $\Psi = \Big(\neg\big(\xi \wedge \neg\eta\big) \ \wedge\ \neg\big(\eta \wedge \neg\xi\big)\Big)$.
- If $\Psi(x_1, \ldots x_n) = \exists z \xi(z, x_1, \ldots x_n)$, then we have, using '$*$' to mark the n-component argument-place,
  $\langle o_1, \ldots o_n \rangle \in \text{ext}(\exists z \xi(z, *))$ iff $\langle \pi(o_1), \ldots \pi(o_n) \rangle \notin \text{ext}(\exists z \xi(z, *))$.
  So we have
  $\langle o_1, \ldots o_n \rangle \in \text{ext}(\exists z \xi(z, *))$ and $\langle \pi(o_1), \ldots \pi(o_n) \rangle \in \text{ext}(\neg\exists z \xi(z, *))$, or
  $\langle o_1, \ldots o_n \rangle \in \text{ext}(\neg\exists z \xi(z, *))$ and $\langle \pi(o_1), \ldots \pi(o_n) \rangle \in \text{ext}(\exists z \xi(z, *))$.
  That is:
  $\langle o_1, \ldots o_n \rangle \in \text{ext}(\exists z \xi(z, *))$ and $\langle \pi(o_1), \ldots \pi(o_n) \rangle \in \text{ext}(\forall z \neg\xi(z, *))$, or
  $\langle o_1, \ldots o_n \rangle \in \text{ext}(\forall z \neg\xi(z, *))$ and $\langle \pi(o_1), \ldots \pi(o_n) \rangle \in \text{ext}(\exists z \xi(z, *))$.

- The first disjunct entails that there is some object in D—call it c—for which
  $\langle c, o_1, \ldots o_n\rangle \in \text{ext}(\xi)$ and $\langle \pi(c), \pi(o_1), \ldots \pi(o_n)\rangle \in \text{ext}(\neg\xi)$, i.e.
  $\langle c, o_1, \ldots o_n\rangle \in \text{ext}(\xi)$ and $\langle \pi(c), \pi(o_1), \ldots \pi(o_n)\rangle \notin \text{ext}(\xi)$.
  (The second conjunct holds for $\pi(c)$, since it holds for *all* objects in D.)
- The second disjunct entails that there is some object in D—call it d—for which
  $\langle \pi^{-1}(d), o_1, \ldots o_n\rangle \in \text{ext}(\neg\xi)$ and $\langle d, \pi(o_1), \ldots \pi(o_n)\rangle \in \text{ext}(\xi)$, i.e.
  $\langle \pi^{-1}(d), o_1, \ldots o_n\rangle \notin \text{ext}(\xi)$ and $\langle d, \pi(o_1), \ldots \pi(o_n)\rangle \in \text{ext}(\xi)$.
  (The first conjunct holds for $\pi^{-1}(d)$, since it holds for *all* objects in D.)

But we can give $\pi^{-1}(d)$ the name c; so that we can recombine the disjuncts and conclude that, for some object c in D, $\langle c, o_1, \ldots o_n\rangle \in \text{ext}(\xi)$ iff $\langle \pi(c), \pi(o_1), \ldots \pi(o_n)\rangle \notin \text{ext}(\xi)$.
So the formula $\xi(x_1, \ldots x_{n+1})$ is our new Ψ, $n + 1$ is our new adicity, and $c, o_1, \ldots o_n$ are our new objects. (Heuristic remark: It is only in this clause that the process can increase, by one, the adicity of Ψ, and hence the number of objects under consideration.)

- If $\Psi(x_1, \ldots x_n) = \forall z\xi(z, x_1, \ldots x_n)$, then continue with $\Psi(x_1, \ldots x_n) = \neg\exists z\neg\xi(z, x_1, \ldots x_n)$. (Heuristic remark: This case can be dropped in the usual way, if we suppose the language to use just $\exists$ as the primitive quantifier.)

- If Ψ is an atomic formula, then:
  - either $\Psi = F_i^1$ for some primitive 1-place predicate F_i^1, in which case:
  $$o_1 \ \text{ext}(F_i^1) \quad \text{iff} \quad \pi(o_1) \notin \text{ext}(F_i^1);$$
  - or $\Psi = G_j^2$ for some primitive 2-place predicate G_j^2, in which case:
  $$\langle o_1, o_2\rangle \in \text{ext}(G_j^2) \quad \text{iff} \quad \langle \pi(o_1), \pi(o_2)\rangle \notin \text{ext}(G_j^2)$$
  (I emphasize that this case includes the 2-place predicate G_j^2 being '=', i.e. equality);
  - and so on for any 3- or higher-place predicates.

Each case directly contradicts the original assumption that π is a symmetry (cf. Equation (2.2)).

End of proof

Corollary 1: An individual is sent to itself by any symmetry.
Proof: Section 2.3.2 defined an individual as an object that is absolutely discerned from every other object. So its absolute indiscernibility class is its singleton set. QED.

This implies, as a special case, the "resurrection" of Section 2.2.2's second conjecture. That is, we have

Corollary 2: If all objects are individuals, the only symmetry is the identity map.

2.4.2 Illustrations and a counterexample

I will illustrate Theorem 1 and Corollary 2, with examples based on those in Section 2.2.2. Roughly speaking, these examples will show how Section 2.2.2's counterexamples to its two conjectures are "defeated" once we consider absolute indiscernibility instead of utter indiscernibility. Then I will give a counterexample to the converse of Theorem 1.

Theorem 1 illustrated:— In Section 2.2.2's counterexample (1), a and b are absolutely indiscernible. Thus Figure 2.5 illustrates the theorem.

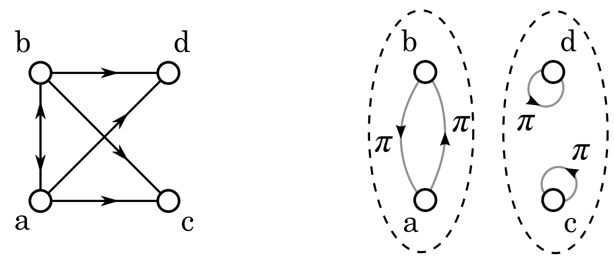

Figure 2.5: Illustration of Theorem 1, viz. that a symmetry leaves invariant the equivalence classes for the relation 'is absolutely indiscernible from'. Here, the symmetry π from Fig. 2.1 leaves invariant the absolute indiscernibility classes, shown on the right.

Corollary 2 illustrated:— I similarly illustrate Corollary 2 by modifying Section 2.2.2's second counterexample, i.e. counterexample (2) (Figure 2.2) to Section 2.2.2's second conjecture. The rough idea is to identify absolute indiscernibles; rather than just *utter* indiscernibles (as is required by (HB)). But beware: identifying absolute indiscernibles that are not utter indiscernibles will lead to a contradiction. Figure 2.2 (and also Figure 2.5) is a case in point: it makes true Rab and $\neg Raa$, so that if one identifies a and b, one is committed to the contradiction between Raa and $\neg Raa$. But by increasing slightly the extension of R, turning absolute indiscernibles into utter indiscernibles, we can give an illustration of Corollary 2, based on Figure 2.2, which avoids contradiction. Namely, we require that Raa and Rbb; this makes a and b utterly indiscernible, not merely absolutely indiscernible. Then we identify a and b, yielding Figure 2.6.

Figure 2.6: Illustration of Corollary 2. When 'is not absolutely discernible from' is taken as identity, the only symmetry for each structure is the identity map.

Against the Theorem's converse:— I turn to showing that Theorem 1's converse does not hold: there are structures for which there are permutations which preserve the absolute indiscernibility classes, yet which are not symmetries. Consider the structure in Figure 2.7.[30] In this structure the relation R has the extension $\text{ext}(R) = \{\langle a, b\rangle, \langle b, a\rangle, \langle a, c\rangle, \langle b, d\rangle\}$. So as in Fig 2.1 (i.e. the counterexample in (1) of Section 2.2.2), a, b are weakly discernible, and c, d are indiscernible. But a and b are absolutely indiscernible. (Proof using Theorem 1: the permutation $(ab)(cd)$ is a symmetry, so $\{a, b\}$ and $\{c, d\}$ must each be (subsets of) absolute indiscernibility classes.) Then the familiar permutation π, which just swaps a and b, clearly preserves the absolute indiscernibility classes. Yet π is not a symmetry, since e.g. $\langle a, c\rangle \in \text{ext}(R)$, but $\langle \pi(a), \pi(c)\rangle = \langle b, c\rangle \notin \text{ext}(R)$.

Figure 2.7 illustrates the general reason why the class of symmetries is a subset of the class of permutations that leave invariant the absolute indiscernibility classes. Namely: for a permutation to be a symmetry, it is not enough that it map

[30] I thank Tim Button for convincing me of this, and for giving this counterexample.

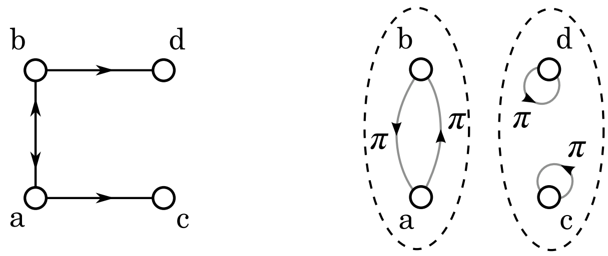

Figure 2.7: A counterexample to the claim that if the absolute indiscernibility classes are left invariant by the permutation π, then π is a symmetry.

each object to one absolutely indiscernible from it; it must, so to speak, drag all the related objects along with it. For example, in Figure 2.7, it is not enough to swap a and b; the objects "connected" to them, namely c and d respectively, must be swapped too. (In more complex structures, we would then have to investigate the objects "connected" to these secondary objects, and so on).

To sum up: this counterexample, together with Theorem 1 and the results of Section 2.2.2, place symmetries on a spectrum of logical strength, between two varieties of permutations defined using our notions of utter indiscernibility and absolute indiscernibility. That is: for a given structure, we have:

$$\begin{array}{c}\pi \text{ leaves invariant}\\ \text{the indiscernibility}\\ \text{classes}\end{array} \xRightarrow{\text{(cf. §2.2.2)}} \pi \text{ is a symmetry} \xRightarrow{\text{(cf. Th. 1)}} \begin{array}{c}\pi \text{ leaves invariant}\\ \text{the absolute}\\ \text{indiscernibility}\\ \text{classes}\end{array}$$

2.4.3 Finite domains: absolute indiscernibility and the existence of symmetries

For structures with a finite domain of objects, there is a partial converse to Section 2.4.1's Theorem 1: viz. that if a and b are absolutely indiscernible, then there is a symmetry that sends a to b. To prove this, we will temporarily expand the language to contain a name for each object. I will also use the Carnapian idea of a *state-description* of a structure (Carnap (1950, p. 71)). This is the conjunction of all the true atomic sentences, together with the negations of all the false ones. But for our purposes, the state-description should also include the conjunction of all the true statements of non-identity between the objects in the domain. This will ensure that a map that I will need to define in terms of the state-description is a bijection (and thereby a symmetry).

I see no philosophical or dialectical weakness in the proof's adverting to these non-identity statements. But I agree that the reason it is legitimate to include them is different, according to whether you adopt the Hilbert-Bernays account of identity or not. Thus the opponent to the Hilbert-Bernays account will include in the state-description all the non-identity sentences holding between any two objects in the domain; for the state-description is to be a *complete* description of the structure, so these non-identity facts should be included. For the proponent of the Hilbert-Bernays account, on the other hand, facts about identity and non-identity are entailed by the qualitative facts, in accordance with (HB). So a description of a structure (in particular, a Carnapian state-description) can be complete, i.e. express all the facts, without explicitly including all the true non-identity sentences. But it is also harmless to include them as conjuncts in the state-description.[31]

I will state the Theorem as a logical equivalence, although one implication (the leftward one) is just a restatement of Section 2.4.1's Theorem 1, and so does not need the assumption of a finite domain.

Theorem 2: In any finite structure, for any two objects x and y:
x and y are absolutely indiscernible $\Longleftrightarrow$ there is some symmetry π such that $\pi(x) = y$.

Proof: *Leftward*: This direction is an instance of Section 2.4.1's Theorem 1, that symmetries leave invariant the absolute indiscernibility classes.

Rightward: Consider an arbitrary finite structure with n distinct objects, $o_1, o_2, \ldots o_i, \ldots o_n$ in its domain D, and any two absolutely indiscernible objects in that domain, o_1, o_2 (so we set $x = o_1, y = o_2$). We temporarily expand the

[31]Harmless, that is, provided the HB-advocate is clear-headed. Recall from comment 1 of Section 2.2.1, that the proponent of the Hilbert-Bernays account assumes that the language is rich enough, or that the domain is varied enough, for each object to be discerned in some way (maybe: relatively or weakly) from every other. Thus one could also argue that this assumption involves no loss of generality: for if it does not hold for a structure, then the indiscernible objects are to be identified; or else—on pain of contradiction for the HB-advocate—the primitive vocabulary is to be expanded so as to discern them, and the proof is then run again with a structure of discerned objects.

language to include a name $\hat{o}_i$ for each object o_i.[32] Then:

1. Construct the state-description $\mathbf{S}$ of the structure. Thus for example, if the language has just one primitive 1-place predicate F and one primitive 2-place predicate R, define:

$$\mathbf{S} := \bigwedge_{i,j} S_{ij} \ \wedge \ \bigwedge_{i} S_i \ \wedge \ \bigwedge_{i<j} \hat{o}_i \neq \hat{o}_j \ \wedge \ \forall z_0 \left(\bigvee_i z_0 = \hat{o}_i \right)$$

$$\text{where} \qquad S_{ij} := \begin{cases} R\hat{o}_i\hat{o}_j & \text{if } \langle o_i, o_j \rangle \in \text{ext}(R) \\ \neg R\hat{o}_i\hat{o}_j & \text{if } \langle o_i, o_j \rangle \notin \text{ext}(R) \end{cases}$$

$$\text{and} \qquad S_i := \begin{cases} F\hat{o}_i & \text{if } o_i \in \text{ext}(F) \\ \neg F\hat{o}_i & \text{if } o_i \notin \text{ext}(F) \end{cases} .$$

2. Define the n-place formula ς from $\mathbf{S}$ by replacing all instances of each name with an instance of a free variable $z_1, ..., z_n$. Writing $\frac{\hat{o}_1}{z_1}$ for the substitution of $\hat{o}_1$ by z_1 etc., we define:

$$\varsigma := \varsigma(z_1, z_2, \ldots z_n) := \mathbf{S}\left(\frac{\hat{o}_1}{z_1}, \frac{\hat{o}_2}{z_2}, \frac{\hat{o}_3}{z_3}, \ldots \frac{\hat{o}_n}{z_n}\right).$$

   Note that $\mathbf{S}$ is $\varsigma(\hat{o}_1, \hat{o}_2, \hat{o}_3 \ldots \hat{o}_n)$.

3. From $\varsigma(z_1, \ldots z_n)$, we define n one-place formulas, the ith being the "finest description" of the ith object in D, i.e. o_i, by existentially quantifying over all but the ith variable, which is itself replaced by another variable, say x:

$$\Sigma_i(x) := \exists z_1 \ldots \exists z_{i-1} \exists z_{i+1} \ldots \exists z_n \ \varsigma(z_1, \ldots z_{i-1}, x, z_{i+1}, \ldots z_n).$$

4. The structure clearly makes true $\Sigma_1(\hat{o}_1)$, since it is entailed by $\mathbf{S}$. But by our assumption, o_1 is absolutely indiscernible from o_2. This means that o_1

[32] Some readers—especially those worrying about our sudden use of names in the object-language—may like to convince themselves that one achieves the same results if, within each absolute indiscernibility class, names are permuted among their denotations. Therefore the proof is invariant under permutations of non-individuals.

and o_2 satisfy the same one-place formulas. So the structure also makes true $\Sigma_1(\hat{o}_2)$. So **S** also entails $\Sigma_1(\hat{o}_2)$.

5. We apply existential instantiation to every existential quantifier in $\Sigma_1(\hat{o}_2)$. Formally, this introduces $n-1$ *new* names, say $\alpha_1, ..., \alpha_{n-1}$. But we know (from the last conjunct in **S**) that there are at most n objects in D; so each of the α_k must name one of the $o_1, ..., o_n$. And the non-identity conjuncts in **S** also entail that any pair of names among the $\alpha_1, \ldots \alpha_{n-1}$ denote distinct objects, all of which are distinct from o_2. Thus we infer that **S** entails the sentence
$$\varsigma(\hat{o}_2, \hat{o}_{\alpha(1)}, \hat{o}_{\alpha(2)} \ldots \hat{o}_{\alpha(n-1)})$$
where the bijective map $\alpha : \{1, 2, 3, \ldots n-1\} \to \{1, 3, 4, \ldots n\}$ is defined by the fact that for each $k = 1, \ldots n-1$ the new name α_k refers to $o_{\alpha(k)}$.

6. We now construct a map π on D using the sentences **S** and $\varsigma(\hat{o}_2, \hat{o}_{\alpha(1)}, \hat{o}_{\alpha(2)} \ldots \hat{o}_{\alpha(n-1)})$ as follows:
$$\begin{array}{cccccccc} \mathbf{S} & \equiv & \varsigma(& \hat{o}_1, & \hat{o}_2, & \hat{o}_3, & \ldots & \hat{o}_n &) \\ & & & \downarrow \pi & \downarrow \pi & \downarrow \pi & \cdots & \downarrow \pi & \\ & & \varsigma(& \hat{o}_2, & \hat{o}_{\alpha(1)}, & \hat{o}_{\alpha(2)}, & \ldots & \hat{o}_{\alpha(n-1)} &) \end{array}$$
That is: π maps o_1 to o_2, and o_2 to $o_{\alpha(1)}$, and o_3 to $o_{\alpha(2)}$, and so on for all the objects in D.

7. Then π is a bijection, since all the $o_2, o_{\alpha(1)}, \ldots o_{\alpha(n-1)}$ are distinct. And D is the range of π. So π is a permutation. But by construction π is also a symmetry. This is because: (i) it induces a map from the state-description **S** to another state-description $\varsigma(\hat{o}_2, \hat{o}_{\alpha(1)}, \hat{o}_{\alpha(2)}, \ldots \hat{o}_{\alpha(n-1)})$; (ii) the latter sentence is entailed by the former and, since both are state-descriptions, both sentences have maximal logical strength; so (iii) the sentences are materially equivalent; and (iv) this material equivalence, together with their maximal logical strength, entails that the extensions of all primitive predicates are

preserved under π. Thus π is the symmetry sought.

End of proof

I now sketch two examples showing the need for finiteness in the statement of Theorem 2. The first uses a countable domain and is sufficient on its own to prove the need for finiteness in Theorem 2; the second, which involves an uncountably infinite domain, is unnecessary, but illustrative.

The countable example is familiar (e.g. Boolos & Jeffrey (1974, p. 191)). I use the fact that first-order arithmetic, in the language $\mathcal{L}_{\mathcal{A}}$ with primitive symbols $s, +, \times$,[33] is not $\aleph_0$-categorical, by finding two elementarily equivalent but non-isomorphic structures, which we then use to create a single structure with an absolutely indiscernible pair not related by any symmetry. Take the standard model of arithmetic $\mathfrak{N} := \langle \mathbb{N}, 0, s^{\mathfrak{N}}, +^{\mathfrak{N}}, \times^{\mathfrak{N}} \rangle$ and a non-standard model $\mathfrak{N}^* := \langle \mathbb{N}^*, \emptyset, s^{\mathfrak{N}^*}, +^{\mathfrak{N}^*}, \times^{\mathfrak{N}^*} \rangle$, where $\mathbb{N}^*$ consists of an initial segment, whose first element is an object $\emptyset$ and which is isomorphic to $\mathbb{N}$, followed by countably infinitely many $\mathbb{Z}$-chains, densely ordered without a greatest or least $\mathbb{Z}$-chain. (So $\mathbb{N}^*$ has the same order type as $\mathbb{N} + \mathbb{Z} \cdot \mathbb{Q}$.)

We then create a new structure $\mathfrak{M} = \langle \mathcal{N}, 0, \emptyset, s^{\mathfrak{M}}, +^{\mathfrak{M}}, \times^{\mathfrak{M}} \rangle := \langle \mathbb{N} \cup \mathbb{N}^*, 0, \emptyset, s^{\mathfrak{N}} \cup s^{\mathfrak{N}^*}, +^{\mathfrak{N}} \cup +^{\mathfrak{N}^*}, \times^{\mathfrak{N}} \cup \times^{\mathfrak{N}^*} \rangle$; (let us assume that the domains $\mathbb{N}$ and $\mathbb{N}^*$ are disjoint, so that $\emptyset \neq 0$). Now, the structures $\mathfrak{N}$ and $\mathfrak{N}^*$ are elementarily equivalent, so any one-place formula in $\mathcal{L}_{\mathcal{A}}$ which is true of 0 in $\mathfrak{N}$ is also true of $\emptyset$ in $\mathfrak{N}^*$, and *vice versa*. So then any one-place formula in $\mathcal{L}_{\mathcal{A}}$ which is true of 0 in $\mathfrak{M}$ will also be true of $\emptyset$ in $\mathfrak{M}$, and *vice versa*.[34] In that case, 0 and $\emptyset$ are absolutely indiscernible in $\mathfrak{M}$. But $\mathfrak{N}$ and $\mathfrak{N}^*$ are not isomorphic; therefore $\mathfrak{M}$ has no symmetries which map 0 to $\emptyset$ or *vice versa*. So we have two objects which are absolutely indiscernible but which are not related by any symmetry.[35]

[33]Remember (comment 2 in Section 2.2.1) that I ban names from primitive vocabularies, so I take $\mathcal{L}_{\mathcal{A}}$ not to contain $\mathbf{0}$, the name assigned to 0 (the number zero) in the standard model of arithmetic, as a primitive. The standard results which I use here are not affected, for we can *introduce* $\mathbf{0}$ by description, in terms of the other primitive vocabulary, i.e. $\mathbf{0} := \imath x\ \neg\exists y\ x = s(y)$.

[34]That is of course not to say that the *same propositions* will be true of 0 in $\mathfrak{M}$ as in $\mathfrak{N}$ (and similarly for $\emptyset$ and $\mathfrak{N}^*$). In particular, the description proposed in fn. 33 as a definition of $\mathbf{0}$, namely $\imath x\ \neg\exists y\ x = s(y)$, is *false* of 0 in $\mathfrak{M}$, because the uniqueness claim fails; but of course the description is also false of $\emptyset$ in $\mathfrak{M}$.

[35]If we expand $\mathcal{L}_{\mathcal{A}}$ to the *infinitary* language $\mathcal{L}_{\mathcal{A}}(\omega_1, \omega)$ (which allows countably infinite con-

Now I provide an example of an uncountably infinite structure in which two objects a_0 and b_0 are absolutely indiscernible, but no symmetry maps one to the other. In this example, there is just one 2-place relation R. We require that a_0 bears R to each of denumerably many distinct $a_1, a_2, \ldots a_i, \ldots$; and each of these a_i bears R only to itself.[36] We also require that b_0 bears R to each of *continuum many* distinct $b_1, b_2, \ldots b_j, \ldots$; and each of these b_j bears R only to itself. Thus we have:

$$\mathrm{ext}(R) = \{\langle a_0, a_1\rangle, \langle a_1, a_1\rangle, \langle a_0, a_2\rangle, \langle a_2, a_2\rangle, \ldots \langle b_0, b_1\rangle, \langle b_1, b_1\rangle, \langle b_0, b_2\rangle, \langle b_2, b_2\rangle, \ldots\}\,.$$

with

$$\mathrm{card}(\{x : \langle a_0, x\rangle \in \mathrm{ext}(R)\}) = \aleph_0\ ; \qquad \mathrm{card}(\{x : \langle b_0, x\rangle \in \mathrm{ext}(R)\}) = 2^{\aleph_0}.$$

So a_0 and b_0 are each like the centre of a wheel, with relations to the a_i, b_j respectively like their wheel's spokes.

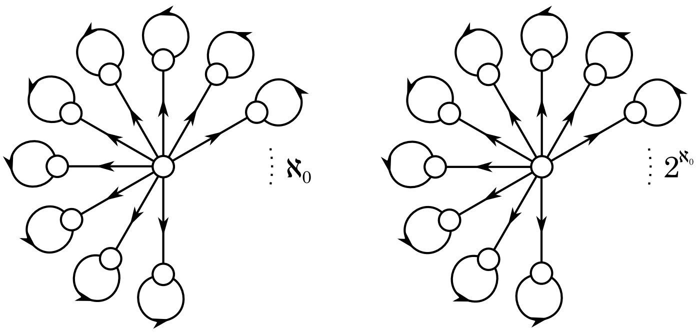

Figure 2.8: A structure with an uncountable domain in which two objects are absolutely indiscernible, yet no symmetry relates them.

These different cardinalities imply that there is no symmetry π such that

junction and disjunction, but only finitary quantification), 0 and $\emptyset$ become absolutely discernible, since we now have the linguistic resources to distinguish between the order types of $\mathbb{N}$ and $\mathbb{N}^*$. Specifically, if we define the relation $x < y := \exists t(x + s(t) = y)$, and recursively define the function $s^0(x) := x$; $s^{n+1}(x) := s\,(s^n(x))$, then 0 and $\emptyset$ are absolutely discerned by the formula $\Phi(x) := \forall y\left(x < y \supset \bigvee_{n\in\omega} y = s^n(x)\right)$, since $0 \in \mathrm{ext}_{\mathfrak{M}}(\Phi)$ but $\emptyset \notin \mathrm{ext}_{\mathfrak{M}}(\Phi)$. Intuitively, the formula $\Phi(x)$ says that every element greater than x lies only finitely far away from x.

[36] So the a_is are weakly discerned from each other by $\neg Rxy$; and so will be distinct even for a proponent of the Hilbert-Bernays account.

$\pi(a_0) = b_0$ or $\pi(b_0) = a_0$. (Recall the discussion after Figure 2.7, at the end of Section 2.4.2.) But on the other hand, the fact that the "only difference" between a_0 and b_0 is the order of the infinity of the objects to which they are related means that they are absolutely indiscernible; for the Löwenheim-Skolem theorem implies that no first-order $\Phi(x)$ can express this difference. More precisely: the Löwenheim-Skolem theorem states that no set of formulas of a first-order language has only models each with denumerably many objects. From this it follows that no first-order $\Phi(x)$ can express the fact that a_0 (and not b_0) has only denumerably many relata. And this implies that no first-order $\Phi(x)$ can discern a_0 and b_0; so they are absolutely indiscernible.

So much by way of developing various results about absolute discernibility. I now turn to metaphysical matters.

Chapter 3

Four metaphysical theses

In this Chapter, I state and discuss four rival metaphysical theses, framed in terms of Section 2.3.2's four kinds. I will first define the theses and make some comments about them (Section 3.1). Then I describe how they disagree in distinguishing possibilities, with one thesis distinguishing two possibilities where another thesis sees only one—thereby condemning the first as proposing a distinction without a difference, or as trading in chimeras (Section 3.2). I will not be committed to any of our four theses—my aim is merely to describe! In Section 3.3, I describe a salient recarving of the four theses into two umbrella positions. One of these positions, which I call *structuralism*, is of particular interest, since it prompts formal procedures that are to be found in a variety of modern physical theories (which is discussed in Caulton and Butterfield (2011)).

It will be obvious that several other metaphysical theses could be formulated as readily as my four. But for my purposes, these are the salient ones.

In this Chapter, I shift focus from syntax and logic to semantics and ontology. Accordingly, I imagine that all parties to the debate about identity and indiscernibility envisage a language that is adequate for expressing all facts, in particular facts of identity and diversity (Adams (1979, p. 7), Lewis (1986, §4.4)); and if they adopt or even just consider the Hilbert-Bernays account, they will imagine deploying it for such a language. (Agreed, the parties will differ about what this language, in order to be adequate, must contain. In particular, I noted in comment 2 of Section 2.2.1 that the haecceitist will require the language to express haecceities with predicates $N_a x, N_b x$ etc.)

I will make three further comments about this shift of focus. They concern, respectively: how uncontentious I will be; the jargon I will use; and my ignoring some modal issues.

(1): In (2) of Section 2.3.1, I officially disavowed any connotations that words like ‘intrinsic’ or ‘external’ might have, while allowing ourselves to use them as mnemonic labels. Now that we are entering on metaphysics, this disavowal is muted, in so far as we I envisage an adequate language.

(2): Given the discussion of jargon at the end of Section 2.1, I now adopt the term ‘individual’ in its fully proper sense, as defined in Section 2.3.2, as an object that is absolutely discerned from all others, relative to the envisaged “ideal” language which is adequate for expressing all facts. Thus ‘individual’ can be used by all parties: though they may well differ about the vocabulary of this language, and also about which objects (if any!) are individuals.

(3): I emphasise that I will not pursue all the issues about modality that are raised by our theses and our comments on them. I will mention only briefly and in passing some implications for the identity of objects across structures (our formal representatives of “possible worlds”); I will not consider modal languages and quantification over structures, and I will steer clear of the issue of essential properties. Of course much of the literature *is* thus concerned. But I have plenty to do, while setting aside these issues: this self-imposed limitation will be prominent in some comments in Section 3.1, and in Section 3.2.

3.1 The four theses

The four theses fall into two pairs. Broadly speaking, each thesis of the first pair is part of a reductive account of identity; while the third and fourth theses are explicitly non-reductive about identity. (We will shortly see another, equally good, way to sort the theses into pairs.)

For each thesis, I will mention an author or two who endorses it, or who is sympathetic to it. But beware: some authors formulate their positions in terms of whether facts of identity and diversity are reduced to or grounded in other facts (qualitative or not). But 'reduced' and 'grounded' are vague, and we will for simplicity instead say 'implied by', etc.

The first pair are both versions of the Principle of the Identity of Indiscernibles (PII), and both conform to the Hilbert-Bernays account of identity. The first is strong; so I write SPII. It vetoes not only objects that are indiscernible one from another, but also objects that are merely relatively or merely weakly discerned. So it amounts to a substantial demand on the primitive vocabulary, that it be rich enough to discern any pair of objects absolutely. (One might say: 'discern any pair of objects in any structure'. But as discussed in (3) just above, I will not emphasise this sort of universal quantification over structures.) So, according to this thesis, all objects are individuals. This thesis seems to be what Hacking (1975) had in mind, given his treatment of his examples. Saunders (2003b, p. 17) ascribes a strengthening of the view to Leibniz, in view of Leibniz's thesis that relational properties are always reducible to the intrinsic properties of the relata; (cf. comment (2) at the end of this Chapter).[1] Adams (1979, p. 11) says of what appears to be SPII, that it is "a most interesting thesis, but much more than needs to be claimed in holding that reality must be purely qualitative".

The second is a weak version of PII: written WPII. It understands 'indiscernibility' as covering all of our four kinds, i.e. as what I called at the end of Section 2.3.2, 'utter indiscernibility'. So this weak PII corresponds to the leftward implication in (HB), that indiscernibility implies identity. Since this implication is the contentious half of (HB), this weak PII is in effect a metaphysical statement of the Hilbert-Bernays account. Like the first thesis, this amounts to a demand on the primitive vocabulary, albeit a weaker one: viz. that it be rich enough to discern any pair of objects, by one or other of our four kinds. In the words of Adams (1979, p. 10), the demand on the primitive vocabulary is that each "thisness" be 'analyzable into, equivalent with, or even identical with, purely qualitative properties or suchnesses.' So, according to this thesis, there can be objects that are not

[1]The fact that absolute discernibility does not imply intrinsic discernibility, and therefore that a "Leibnizian" SPII is a strengthening of SPII, was first pointed out by Ladyman (2007a).

individuals: objects that are only relatively or weakly discerned from at least one other object. To put it as a slogan: there can be diversity without individuality. My central examples of authors who advocate this thesis are Quine (1960, 1970, 1976) and Saunders (2003a, 2003b, 2006). Other authors who are, or who have been, sympathetic are Robinson (2000), Ladyman (2005) and Button (2006).[2]

On the other hand, the third and fourth theses conflict with the Hilbert-Bernays account, in spirit if not in letter. The third thesis is my version of Haecceitism, introduced in Section 2.1.1. As I noted in footnote 12 of Chapter 2, the existence of haecceities contradicts the spirit of the Hilbert-Bernays account.[3] For the latter seeks a reduction of identity to qualitative facts; whereas the haecceitist denies that identity is in all cases reducible to qualitative facts—some pairs of objects differ just by their non-qualitative thisnesses. I shall take my version of haecceitism to hold that any pair of objects differ by their thisnesses, and thus are absolutely discerned—so that every object is an individual. (This does not imply that the discerning work is being done by properties rather than objects: recall our earlier remarks in Section 2.1.1.) Authors who advocate this thesis include Kaplan (1975, pp. 722-3; cf. footnote 4) and Adams (1979, p. 13). Cf. also Wittgenstein's *Tractatus*, remarks 2.013 and 4.27.[4]

I note *en passant* that the first and third theses—the strong version of PII, and Haecceitism—are perhaps the two main traditional positions in the philosophy of

[2]Muller (forthcoming) offers an alternative definition of PII, in which a theory satisfies PII iff $T \vdash \forall x \forall y\,(\mathrm{AutInd}(x, y) \supset x = y)$, where '$\mathrm{AutInd}(x, y)$' is satisfied just in case x and y fail to be even weakly discernible by properties and relations whose extensions are invariant under all automorphisms (what I call symmetries). The idea is to impose a restriction on the properties that are permitted to discern. (Note that it excludes haecceitistic properties, but fails to exclude the use of the identity relation itself.) It is an alternative way to focus on the purely qualitative properties and relations, other than our strategy of banning names. Therefore it is akin to my WPII, except that Muller's PII applies to theories rather than classes of structures.

[3]Agreed, my tactic of banning names, and instead using predicates $N_a x$ etc., enabled the haecceitist to agree with the letter of the Hilbert-Bernays account.

[4]Remark 2.013: 'Each thing is, as it were, in a space of possible states of affairs. This space I can imagine empty, but I cannot imagine the thing without the space.' I interpret this as a commitment to the type of combinatorial independence described in paragraph 2 of Section 2.1.1. Also witness remark 4.27: 'For n states of affairs, there are $K_n = \sum_{\nu=0}^{n} \begin{pmatrix} n \\ \nu \end{pmatrix}$ $[= 2^n]$ possibilities of existence and non-existence. Of these states of affairs any combination can exist and the remainder not exist.' The commitment to haecceitism may not be obvious from this remark, but compare my discussion of a specific example in Section 3.2.1.

identity. In terms of my usage of the term 'individual', they agree that all objects are individuals. But they disagree about whether qualitative indiscernibility is sufficient for being one and the same object (or equivalently, for them: one and the same individual).

Finally, the fourth thesis accepts that a pair of objects can be merely relatively, or merely weakly, discerned—but goes on to deny (HB)'s leftward implication. That is, it denies that indiscernibles must be identical: there are pairs of objects that are utterly indiscernible from one another. In other words: there are utter indiscernibles. But unlike Haecceitism, this thesis does not permit non-qualitative properties which could absolutely discern qualitatively indiscernible objects; so not all objects are individuals. So unlike Haecceitism, this thesis is a denial of the Hilbert-Bernays account in letter as well as in spirit. Here my ban on names comes in particularly useful: according to the thesis under discussion, '$x = y$' is a legitimate primitive formula of the language, representing a particular relation (namely identity); but '$x = a$', which I demand be written as '$\forall y(N_a y \equiv y = x)$', is illegitimate: it represents *no* property at all, since there are no haecceitistic properties.[5]

So far as I can tell, this fourth thesis has only recently been formulated, in part in response to Saunders' advocacy of WPII. Ladyman (2007a, p. 37) calls this thesis 'contextual ungrounded identity'. But I will label it QII, standing for 'Qualitative Individuality with Indiscernibles'; (so here, the 'II' does *not* stand for 'identity of indiscernibles'!). Other formulations, close or identical to my QII, are in: Esfeld (2004), Pooley (2006, ms.), Esfeld and Lam (2006), Ketland (2006), Ladyman (2007a) and Leitgeb & Ladyman (2008).[6] So this fourth thesis, QII, is:

[5] In Robinson's (2000, pp. 163, 173) terminology, the appropriate language for a proponent of QII is *suitable*, since it vetoes thisness predicates, but not *extra-suitable*, since it takes the equality symbol as a primitive.

[6] I also think the quasi-set theory developed by Krause and co-authors (e.g. Krause (1992), Dalla Chiara, Giuntini & Krause (1998), French & Krause (2006, Ch. 7), Krause and French (2007) and da Costa & Krause (2007)) is a kindred position: 'kindred' since it allows objects to be permuted without engendering any change (it is therefore apt to describe it as a theory of non-individuals); but only kindred, i.e. not the same position, since it vetoes talking of identity as a relation among such objects, i.e. '$x = y$' is not a wff if x and y are taken to range over so-called "m-atoms".

(i) like haecceitism, and unlike the two versions of PII, in that it disagrees with the Hilbert-Bernays account (and also it denies any other reduction of identity to qualitative facts); but also

(ii) like the weak version of PII, and unlike both the strong version and Haecceitism, in that it allows there to be objects that are not individuals.

Note that while (i) corresponds to pairing the first two theses, in contrast to the last two (as I announced at the start of this Section), (ii) corresponds to pairing the first and third theses (both demand individuals, in our sense), in contrast to the second and fourth theses (which both allow non-individuals, in my sense).

In my opinion, these two ways of pairing the theses are equally natural. Thus we can consider the four theses as illustrating the four possible combinations of answers to two equally natural Yes-No questions. Namely, the questions:

(a) Is indiscernibility sufficient for identity? (Or equivalently, the contrapositive: is diversity (non-identity) sufficient for discernibility? Here 'discernibility' is to mean 'qualitative discernibility', i.e. it excludes appeal to haecceities.)

(b) Is every object an individual (in our sense of being absolutely discerned from every other object)?

Thus the four metaphysical theses can be placed in Table 3.1.

		Is qualitative indiscernibility sufficient for identity?	
		Yes	*No*
Is every object	*Yes*	**SPII**	**Haecceitism**
an individual?	*No*	**WPII**	**QII**

Table 3.1: Two questions and four metaphysical theses.

To sum up, here is the official statement of the four rival metaphysical theses.

SPII : *The Strong Principle of the Identity of Indiscernibles*

Any two objects are absolutely discernible, in Section 2.3.2's sense, by qualitative properties or relations. Therefore every object is an individual. The possibility of objects discernible merely relatively or merely weakly is denied.

WPII : *The Weak Principle of the Identity of Indiscernibles*

Any two objects are discernible, through some 1- or 2-place formula involving only qualitative properties and relations, according to the Hilbert-Bernays axiom. The gap between discernibility and absolute discernibility allows for objects which are non-individuals; such objects are either merely relatively or merely weakly discernible.

Haecceitism :

Any object a has a haecceity $N_a x$, a property with no other instances. (But we construe each of these properties "thinly": specifically, as identical with the property expressed by '$x = a$'.) This means that the diversity of objects is primitive in the sense that two objects need not be at all qualitatively discernible. And every object is an individual. Individuality is also primitive in the sense that an object's absolute discernibility from all others may rely solely on the haecceities, and need not involve qualitative properties or relations.

QII : *Qualitative Individuality with Indiscernibles*

The diversity of objects is primitive in the sense that two objects can be qualitatively utterly indiscernible. Individuality, by contrast, is *not* primitive: each object's individuality requires absolute discernment from all others by purely qualitative properties or relations. The gap between diversity and individuality allows non-individual objects. So a non-individual is either merely relatively discernible, or merely weakly discernible, or else utterly indiscernible, from at least one other object.

Finally, I emphasise that the four theses are not the only occupants of their respective quadrants of logical space, as carved out by the two questions in the Table. Examples are as follows.

1. We saw already in Section 2.1 that there are different versions of haecceitism: how thickly should haecceities be construed?

2. SPII affords another example. One can say Yes to both questions in different ways than our SPII. One obvious way is by requiring every primitive predicate of the envisaged adequate language to be 1-place.[7] And my previous discussion has shown two other less obvious ways:

   (a) One could require that every object be specifiable in Quine's (1976a) sense (cf. the Interlude, and footnote 24, in Section 2.3.2);

   (b) One could allow only those structures that have only the trivial symmetry, i.e. for which the identity map is the only symmetry. Corollary 2 of Theorem 1 (Section 2.4.1) means that in infinite domains this requirement is weaker than all objects being individuals (i.e. my SPII). For it is only in finite domains that they are equivalent (cf. Theorem 2 in Section 2.4.3).

3. WPII affords another example. That is: one can say Yes to 'Is qualitative indiscernibility sufficient for identity?, and No to 'Is every object an individual?', other than by endorsing WPII. For one could allow relative discernibles, but forbid structures containing mere weak discernibles.

4. Finally, QII affords another example. One can say No to both questions in different ways than my QII: for example, by saying that only some objects have haecceities—this again secures the No answers, that not all objects are individuals, and that there are utter indiscernibles.

3.2 The theses' verdicts about what is possible

I turn to illustrating exactly what structures each of our four theses permit, for the simple case of exactly two objects a and b and one (qualitative) relation R.

[7]Such a language is envisaged by Saunders (2003b, p. 17) for a modern treatment of Leibniz's metaphysics. See also footnote 10.

(This case will be complicated enough!) This exercise will be governed by three main rules.

1. All the state-descriptions will have common elements which we will *not* show explicitly. Specifically, every state-description contains conjuncts which constrain the cardinality of the domain to be two: $(x \neq y \;\wedge\; \forall z(z = x \vee z = y))$, with x and y bound by existential quantification. For WPII and SPII, the '=' symbol is of course governed by the Hilbert-Bernays axiom (HB); for those metaphysical theses, (HB) should be considered as an additional implicit conjunct in the state-descriptions written below.

2. As usual, the haecceitist is asked to add two haecceitistic predicates, 'N_a' and 'N_b', to the primitive vocabulary, one for each object, and associated uniqueness conditions (again implicit in the state-descriptions below).

3. I characterise the structures syntactically; that is, by the logically strongest sentences for which the structures are models. For finite structures, this is equivalent to characterising structures up to isomorphism. So when we "count the structures" allowed by each metaphysical thesis, we are in fact counting the *equivalence classes* of mutually isomorphic structures. Sometimes I will say 'structure' when I mean 'equivalence classes of structures'—when counting it is useful only to mean the latter!

3.2.1 Haecceitism

I begin with Haecceitism, which will turn out to be the most "generous" metaphysical position, in the sense of allowing more possibilities. The inclusion of the haecceitistic predicates in the primitive vocabulary allows us always to discern the two objects in any structure. Consider the state-description sentences

$$\exists x \exists y (N_a x \wedge N_b y \wedge \pm Rxx \wedge \pm Rxy \wedge \pm Ryx \wedge \pm Ryy)$$

where $\pm Rxy$ designates either Rxy or its negation.[8]

[8]Here we see how our ban on names (comment 2 in Section 2.2.1) has not lost the haecceitist any expressive power. This expression is equivalent, in a language with names, to the template

A decision, for each of the R-formulas, whether it is asserted or denied, together with the haecceitistic uniqueness conditions, suffices to completely describe a structure (up to isomorphism). There are 4 R-formulas, so there are $2^4 = 16$ distinct (equivalence classes of) structures.

The structures may be partially ordered, according to which R-formulas are asserted in their descriptions. That is: we can partially order, by 'is a subset of', the power-set of the set of R-formulas; this induces a partial order on the structures, by associating each structure with the subset of R-formulas that are true in it. For the Haecceitist case, this partial order produces a lattice of structures. (A *lattice* is a partially ordered set, in which any two elements have a greatest lower bound, and a least upper bound.) It is shown in Figure 3.1.

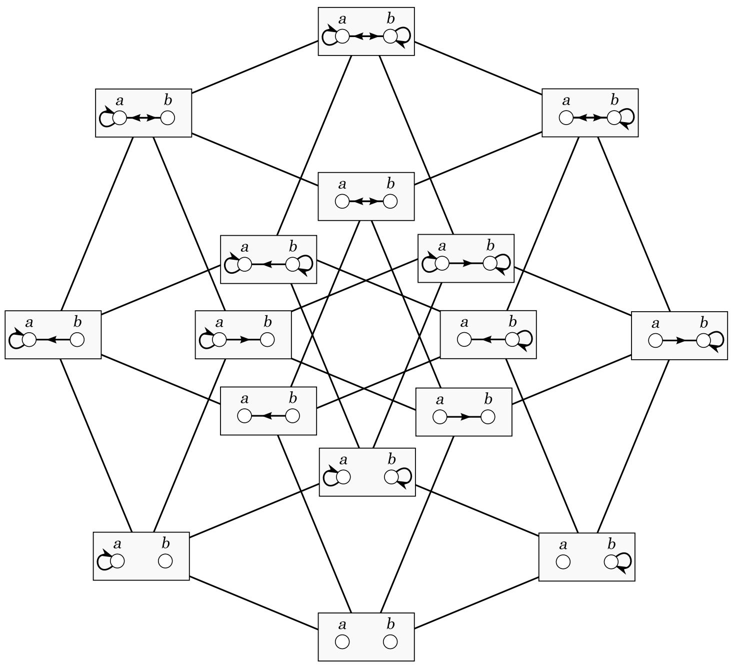

Figure 3.1: The lattice of distinct, permitted two-object structures, according to the haecceitist.

In Figure 3.1, lines connecting structures indicate a single alteration in the assertion/denial of an R-formula. The Figure also merits two further comments:

sentence ($\pm Raa \wedge \pm Rab \wedge \pm Rba \wedge \pm Rbb$).

1. It adopts an unusual convention about how to display the partial order. It is usual to draw a lattice by putting all the elements that are immediately greater than the least element together on a horizontal row (often called 'rank'), above the least element; and all the elements that are immediately greater than any of them, on a next higher rank; and so on. Figure 3.1 does *not* do that. Instead, I have chosen to display, using the four different *directions* of the various upward lines (e.g. the four different lines emanating from the least element), the four different truth-value flips that one can make, so as to make more R-formulas true. (An example in tribute to Hitchcock and Cary Grant: the direction North by North-West corresponds to affirming the formula Rba.)

2. An assignment of a truth-value, 0 or 1, to each of the four R-formulas (a "fourfold decision"), can be thought of as a vertex of the 4-dimensional unit hypercube $[0,1]^4 \subset \mathbb{R}^4$. There are $2^4 = 16$ such vertices, and Figure 3.1 is a parallel projection of the hypercube onto the plane of the paper.

Note how the use of haecceitistic predicates distinguishes structures that would otherwise be equivalent. For example, the (equivalence class of) two-object structures picked out by the sentence

$$\exists x \exists y (\neg Rxx \wedge Rxy \wedge \neg Ryx \wedge \neg Ryy)$$

splits, for the haecceitist, into two equivalence classes: those specified by

$$\exists x \exists y (N_a x \wedge N_b y \wedge \neg Rxx \wedge Rxy \wedge \neg Ryx \wedge \neg Ryy)$$

and those specified by

$$\exists x \exists y (N_a x \wedge N_b y \wedge \neg Rxx \wedge \neg Rxy \wedge Ryx \wedge \neg Ryy) \ .$$

3.2.2 QII

The proponent of QII has no recourse to haecceitistic predicates, but still allows the distinctness of the two objects to be primitive in the sense of allowing utter indiscernibles. Her state-description sentences for describing the two-object structures are

$$\exists x \exists y (x \neq y \wedge \pm Rxx \wedge \pm Rxy \wedge \pm Ryx \wedge \pm Ryy) \ .$$

The lack of any way to *absolutely* discern one object from the other, except by the pattern of instantiation of R, means that the number of distinguishable structures is less than the *prima facie* number, 16. Any choice of the combination of R-formulas to assert that renders the two objects absolutely discernible is double-counted, since we obtain an equivalent sentence by the interchange of the bound variables x and y. Since there are two structures in which the objects are utterly indiscernible, and two in which they are merely weakly discernible, and none in which they are merely relatively discernible, we conclude that $16 - 2 - 2 - 0 = 12$ of our 16 *prima facie* allowable structures are double-counted. So there are in fact only $2 + 2 + \frac{12}{2} = 10$ distinct structures. The partially ordered set (unlike for Haecceitism, it is not a lattice) of these structures is shown in Figure 3.2. (Note: in accordance with our discussion, this Figure lacks names for the objects.) Thus the effect of eliminating haecceitistic differences is to "fold over" the haecceitist's lattice, given in Figure 3.1, along its central vertical axis.

The result is no longer a lattice, since it is no longer the case that any two structures have a unique join and meet. This seemingly unremarkable technical result underlies a significant metaphysical one: namely that there is no unique transworld identity relation. We can see this as follows. Take, for example, the two (equivalence classes of) structures α and β that lie in the middle row of Figure 3.2. What is the supremum of these two structures, i.e. the structure in which all the relations that hold in either of these two structure hold, and no more?

Without haecceitistic predicates, the answer depends on how else we decide to cross-identify the objects. We cannot, as with Haecceitism, simply take the union of ext(R) in each of the two structures (i.e. the union of two sets of ordered pairs) to obtain a unique, new ext(R), since for QII (and WPII and SPII, below) it is not a set of ordered pairs but an (isomorphism) *equivalence class* of such sets that represents a single world—and the unions of two isomorphism equivalence

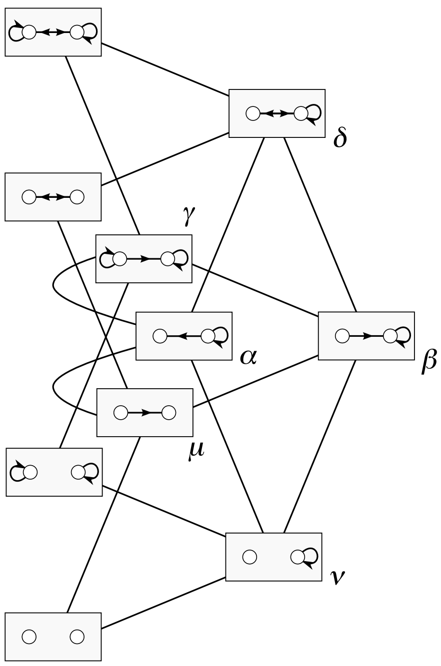

Figure 3.2: The poset of distinct, permitted two-object structures, according to the proponent of QII.

classes does not yield an isomorphism equivalence class. So we must employ an alternative strategy: to form the union *first*, for two structures *haecceitistically conceived*, and *then* to take equivalence classes under isomorphism. But which pair of haecceitistically-conceived structures do we choose? Each of α and β is an equivalence class containing two structures, so there are $2 \times 2 = 4$ pairs to choose from, which, after taking equivalence classes, yields two candidates for the supremum. They are γ and δ in Figure 3.2.

In other words: we must decide on some QII-acceptable way to cross-identify the objects between structures. In our example, we could either take our transworld individuals to be identified by the formulas $\pm Rxx$, (yielding as supremum the equivalence class δ), or by the formulas $\pm\exists y(y \neq x \mathbin{\&} Rxy)$ (yielding as supremum the equivalence class γ).

So under QII the partial order is not a lattice. But using formulas to cross-identify objects between structures brings out four further differences from the situation under Haecceitism: differences which will also be shared by WPII and

SPII, below.

1. There is no *uniquely salient* option as to which qualitative properties and relations ground the transworld identity relation.

2. One may need to use individuating formulas that are *disjunctive* in the primitive vocabulary to cross-identify objects between structures whose equivalence classes of properties- and relations-in-extension are disjoint. (Cf. e.g. structures μ and ν in Figure 3.2. There are two rival candidate pairs for trans-structural objects, and it is hard to say which cross-identification criteria are most natural.)

3. Even allowing disjunctive individuating formulas, there will not always be cross-identifying formulas, since at least one of the structures may contain *non-individuals* (cf. e.g. any of the four unlabelled (i.e. leftmost) structures in Figure 3.2 with any other in the poset). (In such a case, cross-identification is anyway unnecessary to yield a unique supremum and infimum.)

4. Cross-identification of an object between structures by the same formula fails to be transitive, since the formula may fail to be uniquely instantiated in every structure. Therefore we shouldn't, strictly speaking, speak of transworld *identity* at all; and talk of counterparts à la Lewis (1986, Ch. 4) seems far more appropriate.[9]

[9]On a related note, I thank Leon Horsten for drawing my attention to the fact, shown by Quine (1960) and Føllesdal (1968), that the use of definite descriptions to individuate objects in quantified modal logics leads to a collapse of modal distinctions, if we require that any two definite descriptions which are actually co-instantiated are necessarily co-instantiated. This latter requirement (which Quine (1960, p. 198) calls a 'disastrous assumption') is a last-gasp solution to the problem of referential opacity for modal operators. I suggest an alternative, seemingly more radical, escape: to veto as ill-formed any sentence in which modal operators are put in the scope of a quantifier. This apparently onerous restriction means that first-order variables are no longer rigid designators that come for free. I welcome this result: for Haecceitists this is no restriction at all, since for transworld identification there is always the recourse to haecceitistic predicates; and for everyone else transworld identification is provided for, if at all, by qualitative matters.

3.2.3 WPII

We saw that the QII proponent complains that the Haecceitist double-counts some structures. The proponent of PII agrees with this accusation against Haecceitism, but in turn complains that both the Haecceitist and the QII proponent "over-count" some structures by counting them at all, i.e. by accepting them as possible—or at least as possible *two*-object structures.

This equivocation between the elimination of two-object structures, and the re-description of them as one-object structures corresponds to the two ways of understanding the Hilbert-Bernays axiom (HB), discussed in comment 3 of Section 2.2.1. On one hand, if we take '=' as a logical constant, then we take the specification of the cardinality of the domain as logically independent of the specification of the pattern of instantiation of the qualitative properties and relations, and therefore as possibly in contradiction with it, taken in conjunction with (HB). In such a case we take (HB) as eliminating these problematic structures. On the other hand, if we take '=' as *defined* by (HB), then we have no business fixing the cardinality of the domain independently of the pattern of instantiation: the problematic structures just mentioned are then construed as clumsily described one-object structures. For definiteness, let us talk from now on of elimination rather than re-description.

The proponent of PII has at her disposal for the specification of structures only the state-description sentences

$$\exists x \exists y (\pm Rxx \wedge \pm Rxy \wedge \pm Ryx \wedge \pm Ryy) \ .$$

But, as just discussed, the specifications obtained from these sentences are hostage to the Hilbert-Bernays axiom (HB); consequently not all of them will describe structures of as many as two objects. Those structures for which the PII proponent is committed to identify the objects are excluded as *genuine* two-object structures. Which structures these are depends on which version (strength) of the principle of identity of indiscernibles (PII) is endorsed.

The most liberal PII proponent, the proponent of WPII (Weak PII), rules out only those structures in which there are any *utterly* indiscernible objects. There

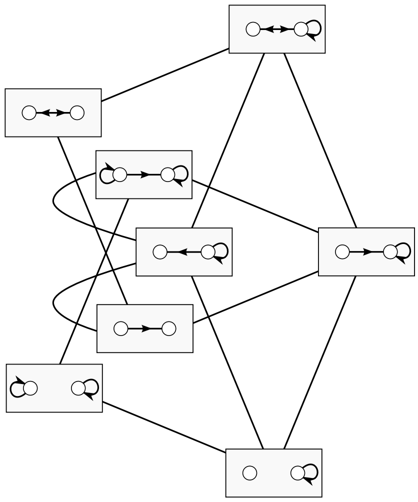

Figure 3.3: The poset of distinct, permitted two-object structures, according to the proponent of WPII.

are only two such structures in our example: viz. corresponding to all four of the R-formulas being asserted, or all four being denied. This reduces the number of distinct, permitted structures from QII's ten to eight; see Figure 3.3.

3.2.4 SPII

The partially ordered set of structures permitted by WPII, Figure 3.3, is further diminished if we endorse SPII: i.e. the strengthening of PII to eliminate structures that contain absolute indiscernibles. Figure 3.3 has two such structures: those in which the objects are merely weakly discernible. By eliminating these two structures, we are left with only six distinct structures; see Figure 3.4.[10] Note that in the move from WPII to SPII, the elimination of structures *cannot*, as in the move from QII to WPII, be instead understood as re-description. For, even if we refrain from independently specifying the cardinality of the domain, a state-description on its own may contradict the requirement that any two objects be

[10] Returning to the 'Leibnizian' proponent of SPII: if we follow Saunders' (2003b, p. 17) idea that the 'Leibnizian' requires all objects to be be *intrinsically* discernible, this would reduce the number of permitted structures still further to *four*. These are the structures in which only one of the two objects satisfy Rxx. I will discuss this view, which I call *qualitative intrinsicalism*, in Section 3.3.

absolutely discernible. This fact underscores the logical strength of SPII.

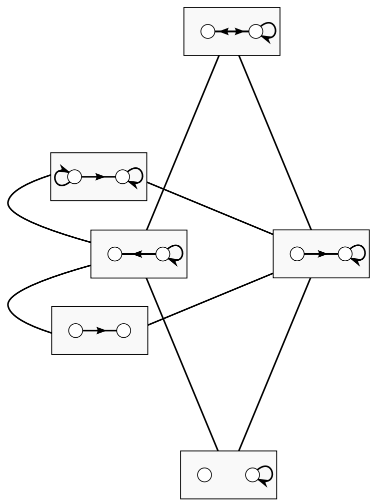

Figure 3.4: The poset of distinct, permitted two-object structures, according to the proponent of SPII.

3.2.5 A glance at the classification for structures with three objects

Finally, we glance at structures with three objects. The existence of a third object means that we here get the "first", i.e. most elementary, illustrations of external discernibility in the *intuitive* sense that requires a third object. We also see here the first cases of mere relative discernibility. There is no space for details; I just report that, as in the preceding Subsections, one can count and classify the structures according to each of their symmetries. Out of the 512 structures accepted by the Haecceitist, 420 have no non-trivial symmetries (and therefore lie in isomorphism equivalence classes with cardinality 6), 84 have only pair-wise swaps as non-trivial symmetries (with 3 structures per isomorphism equivalence class), 4 have only cyclic permutations as non-trivial symmetries (with 2 structures per isomorphism equivalence class), and there are 4 totally symmetric structures (for which isomorphism equivalence classes are singletons). 46 of the structures contain indiscernible objects, 42 of which contain only two indiscernibles and 4 of which contain three.

From this information, we may compute the following totals:

Metaphysic	*Number of structures*	*Breakdown of numbers*
Haecceitism	512	$= 420 + 42 + 42 + 4 + 2 + 2$
QII	104	$= \frac{420}{6} + \frac{42+42}{3} + \frac{4}{2} + 2 + 2$
WPII	88	$= \frac{420}{6} + \frac{42}{3} + \frac{4}{2} + 2$
SPII	70	$= \frac{420}{6}$

3.3 Structuralism and intrinsicalism

The four metaphysical theses just discussed have various similarities and differences, but one distinction in the logical space is of particular interest: I believe it meshes with much of the recent logico-philosophical literature falling under the buzz-words 'structuralism' and 'structural realism',[11] and also with modern philosophical discussions of physical theories—in particular, general relativity and non-field-theoretic quantum mechanics.[12]

This distinction is about *individuality* (a main topic of Section 2.4): specifically, whether the individuality of an object relies on, or is associated with, *qualitative* and (typically) *relational*, differences to other objects; or whether instead an object's individuality relies on, or is always associated with, purely *intrinsic* differences (defined in Section 2.3.2; which may be qualitative *or* solely haecceitistic). The position we wish to call *structuralism* holds that: either there may be objects which are not individuals; or at least, if every object in every possible world is an individual, then it is not in all cases due to differences which are purely intrinsic. So the denial of structuralism holds that every object, in every world, is an individual due to purely intrinsic differences. We call this position (for want of a better word) *intrinsicalism*.

Why are these positions philosophically salient? That is: what unites the various, more specific views that fall under the umbrellas 'structuralist' or 'intrin-

[11] See e.g. MacBride (2005), Button (2006), Ladyman (2005, 2007a), Leitgeb & Ladyman (2008) and Ketland (2006, 2009).

[12] See e.g. Saunders (2003a, 2003b, 2006) Esfeld & Lam (2006), Pooley (2006), Ladyman (2007a), Muller & Saunders (2008) and Muller & Seevinck (2009). Jeremy Butterfield and I add to this discussion in Caulton & Butterfield (2011).

sicalist'? To this I have two answers: one historical, one philosophical. Firstly, intrinsicalism encompasses what I believe are two central views which have dominated, up until recently, philosophical debate about identity and individuality; namely, haecceitism and a strengthening of SPII, in which all objects are taken to possess a unique, though maybe very complex, individuating essence; (the 'bundle theory' of objects being one such view: cf. Adams (1979, p. 7). On the other hand, structuralism encompasses a variety of alternatives to intrinsicalism which have recently been articulated.[13]

Secondly, a common feature of structuralist, as against intrinsicalist, views is that they prompt a certain interpretation of, or even a revision of, traditional formalisms in both semantics and physical theories. To end this Chapter and part, I will briefly discuss this (Section 3.3.2). But first (Section 3.3.1), I relate these two new positions to the four metaphysical theses of Section 3.1.

3.3.1 Relation to the four metaphysical theses

Structuralism, as I define it, includes some of our four metaphysical theses (roughly two "and a half" of them), which may be linearly ordered on a spectrum, according to their strictness, i.e. what they rule out. (This is already suggested by the possible world diagrams of Section 3.2, in which we pass from haecceitism to SPII by a successive vetoing of worlds.)

At the weak end of this spectrum we find QII, for which the numerical diversity of objects is accepted as primitive, just as it is by the haecceitist. QII is structuralist both because it allows non-individuals and because it allows external discernibility as a basis for individuality.

Along the way to the strong end of the spectrum, we find versions of structuralism that ground diversity, as well as individuality, in the differential instantiation of qualitative properties and relations. In doing so, we find a commitment to the

[13]Hintikka & Hintikka (1983) present a similar dichotomy to our structuralism vs. intrinsicalism (which they call the "structural" and "referential system", respectively), and bemoan the bias towards intrinsicalism in the historical development of formal logic. I agree that the historical dominance of intrinsicalism is unfortunate and unjustified; and note with interest their psychological evidence (from studies with children) that the evidence may be sex-linked.

Hilbert-Bernays account of identity; so here we find WPII. WPII is structuralist both in allowing for non-individuals (though not indiscernible non-individuals) and in allowing for external discernibility.

At the strong end of the spectrum, we find that our previous metaphysical thesis, SPII, has been split. For SPII demands that every object be an individual. This does not yet commit one to intrinsicalism: for intrinsicalism requires also that the individuating properties be intrinsic. We may therefore distinguish between an SPII which *allows* for objects individuated by merely external formulas—that is, a structuralist SPII—and an SPII which does *not* allow for such objects—an intrinsicalist SPII. To repeat the main idea: to deserve the label 'structuralism' (in my usage), the individuality of objects must in some cases involve their absolute discernment by *extrinsic*, qualitative properties, not by *intrinsic*, qualitative properties.[14]

Finally, I note that haecceitism is a form of intrinsicalism, since haecceities are intrinsic individuating properties *par excellence*. Thus my taxonomy is summarised in Figure 3.5, below.

3.3.2 The semantics of the structuralist

As we have seen, grounding individuality qualitatively has consequences for the specification of possible states of affairs. Thus a structuralist cannot make sense of what Lewis (1986, p. 221) called 'haecceitistic differences': differences to do

[14]I should point out that even intrinsicalist SPII has a whiff of structuralism about it. Consider, for example, the three-object structure $\mathfrak{A} := \langle\{a, b, c\}, F^{\mathfrak{A}}, G^{\mathfrak{A}}, H^{\mathfrak{A}}\rangle$, with three primitive monadic predicates with extensions $F^{\mathfrak{A}} = \{a, b\}, G^{\mathfrak{A}} = \{b, c\}, H^{\mathfrak{A}} = \{a, c\}$. This structure is permitted under intrinsicalist SPII, the objects being individuated by the monadic predicates $F \wedge G, G \wedge H, F \wedge H$. The whiff of structuralism is obvious when we consider the combinatorial theory of logical possibility, which takes the structure $\mathfrak{A}$ to indicate the (logical) possibility of other structures, some of which have a, b and c allocated to the extensions of F, G and H in ways that do not secure their qualitative individuality—one such case is $a, b, c \in \mathrm{ext}(F \wedge \neg G \wedge \neg H)$. It should not come as a surprise that SPII (whether intrinsicalist or structuralist) restricts naïve, combinatorially-conceived possibility, since the admissibility under SPII of affirming or denying a given elementary proposition depends on global features of the structure description (namely, whether certain other elementary propositions are affirmed or denied). Note that, in contrast, QII and, perhaps, WPII do *not* restrict possibility in this way. For, as explained in Section 3.2.3, QII and (given the reductionist attitude to identity) WPII may be taken as *redescribing* rather than *vetoing* the structures permitted under unrestricted Haecceitism.

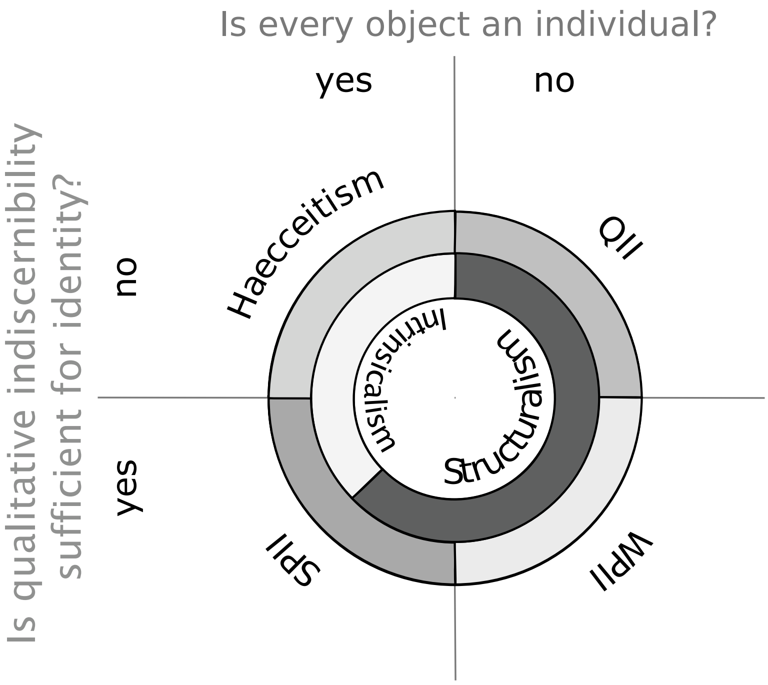

Figure 3.5: Four metaphysical theses and two salient positions.

with which object occupies which role in the mosaic of qualitative matters of fact. Since an object's individuality is provided, if it all, by the role that it plays in the structure of which it is a part, there simply is nothing else to grasp it with, that could be used to give sense to, for example, the object swapping roles with another. A structuralist simply does not agree with the haecceitist that a permutation of objects "underneath" the mosaic makes any sense.

But here the structuralist runs up against an aspect of formal semantics (model theory) which threatens to make her position hard to state. Model theory allows without demur that an object may be individuated independently of its qualitative role in a structure.[15] Or think of Lewis's (1986, p. 145) 'Lagadonian' languages, in which objects are their own names. In philosophical terms, this trick ensures that haecceities are always at hand to give sense to a permutation of "bare" objects underneath the qualitative properties and relations. The same point applies in physical theories. Permutations of the objects that the theory treats (e.g. particles

[15] This point is often emphasised in the philosophy of modality; a good early example is Kaplan (1966).

in classical mechanics or quantum mechanics) can be implemented on the theory's state space (e.g. a classical phase space, or quantum Hilbert space, for several particles); and typically, a state is changed by such a permutation, i.e. its image is a different state, even if the mosaic of qualitative properties and relations is unchanged.

In response, the structuralist can adopt either of two approaches, which I now briefly describe; (more details are in Caulton & Butterfield (2011)). The first involves a certain interpretation of the usual formalism (of the formal semantics or physical theory in question). The second involves revising the formalism.

First, the structuralist can accommodate the orthodox practice in model theory and physical theory—that permutations of "bare" objects make sense—by an interpretative move. That is, she can admit that such permutations typically induce a different formal representation of a state of affairs: a different model in the sense of model theory, or a different state in the physical theory's state-space. But she then expresses her position by saying that such a permutation does not produce a representation of a *distinct* state of affairs.

On this view, it is misleading to say that model-theoretic structures, or states in the state-space, represent possible states of affairs; for the representation relation is not one-to-one, but many-to-one. Rather, a single state of affairs is represented by the *equivalence class* of mutually isomorphic structures (states), the isomorphism being given by a permutation of the 'bare' objects in the structure's domain (a permutation of the physical theory's objects, e.g. particles). This ascent from structures to equivalence classes suits a theory which seems "not to care" how its objects are arranged under the mosaic of properties and relations, so long as the same qualitative pattern is instantiated; such theories support a structuralist intepretation.

The second approach proposes to keep the original structuralist prohibition against permutations making sense, by revising the formalism. The idea is to replace each equivalence class of isomorphic representations by a less structured item, so that permutations do not make even a *formal* difference, or else cannot be made sense of. It is well established, for both model theory and physical theories. (See Belot (2001, p. 59) for a philosophical introduction.) Thus in elementary model theory, it is straightforward to show that if we quotient a structure by the indiscernibility relation, then: (i) the resulting structure is elementarily equivalent to the given one; and (ii) in it, identity is first-order definable, namely by indiscernibility—i.e. the leftward implication of (HB) holds (Ketland (2009, Theorems 18, 20)). In physical theories (both classical and quantum), one considers the quotient of the state space (often called a *reduced state-space*) by various equivalence relations (often going by the name 'gauge-equivalence')—not just by isomorphisms induced by permutations of the underlying objects (particles). For more details, from a philosophical perspective, cf. e.g. Belot (2003, especially §5) and Butterfield (2006a, §§2.3, 7).

To close, I shall only note the situation for quantum mechanics, with isomorphisms induced by permutations of the underlying quantum particles. Here also one can quotient: in effect, particles of the various possible symmetry types (including paraparticles) can be represented in a "reduced Hilbert space" (in general, of lower dimension than the original, "naïve" one) in which each particle permutation is represented either trivially or not at all. (We develop this topic in Caulton & Butterfield (2011).) Here I note only that this second approach (unlike quasi-set theory, cf. footnote 6) retains the orthodox ideas and methods of model theory and classical logic, merely limiting their application to the favoured, i.e. quotient, structures or states.

This concludes my purely logical and metaphysical study of identity and discernibility. I now turn to quantum mechanics and its interpretation. The concepts and results forged in the foregoing two Chapters will be implemented in the following Chapters to answer the question: What, in quantum mechanics, are particles? But before I tackle that question specifically (in Part III), I need first (in Part II) to lay down a philosophical framework which will make explicit what I take questions of this sort to be asking, and how best to answer them.

Part II

Representing particles

Chapter 4

Concepts and representation

In this Chapter, I discuss philosophical preliminaries about the nature of our enquiry (Section 4.1) and the contrast between the mathematical and physical realms, and how one represents the other (Section 4.2). This discussion was not necessary for Part I, since those results may be taken as applicable for both mathematical and physical ontologies. However, for the purposes of interpreting a physical theory, it is crucial to distinguish between one's mathematical ontology and one's physical ontology. For, as I will argue, interpretation involves settling upon the right relation between the two.

4.1 Concepts: analysis, explication and reform

In analytic philosophy, there is a strong tradition of conceptual enquiry; and therefore also, considerable debate about what such enquiry should involve and could achieve. I will briefly locate our project in relation to this tradition and debate.

I will sketch three views of conceptual enquiry. I will soon say more about what I take a concept to be, but as a first approximation we can take a concept to be a division, i.e. a binary distinction, among some relevant domain of objects: it gathers a subset of the domain as being similar to each other in some respect. I briefly discuss the first two views (Section 4.1.1); but I will emphasize the third view, due to Haslanger (2006), which I endorse (Section 4.1.2).

4.1.1 The Moorean and Carnapian views

The first view of conceptual enquiry is the most traditional. It is often called 'conceptual analysis', but I will call it *Moorean*, after G.E. Moore. (But there are countless recent distinguished practitioners e.g. Chisholm; and on reading middle-period Plato, one might equally call it 'Socratic'!) One takes a concept which seems to be central to our thought and language, and-or philosophically important, and-or problematic: e.g. causation, perception, freedom, or duty. One then endeavours, by considering the way the concept is used (especially how it is expressed in our language), to give an definition of the concept (or term). The definition is to be true to how the concept or term is in fact used. So it is an analytic proposition in the usual (admittedly rough) sense of being true 'in virtue of the meanings of words'; and in particular, it must face no counter-examples. It is also meant to reveal the connections between the concept or word and other related ones: and to thereby give us a clearer view of how a whole group of our concepts work, and of the sector of reality we use them to describe.

Of course, this view has several variants. For example, there is the view that careful attention to the *minutiae* of linguistic usage can be revealing (a view popular in Oxford in the 1950s and '60s; cf. in particular Austin (1961)); or that the philosophical problem is not ignorance of an analytic proposition or propositions, which is to be overcome by formulating definition(s) which face no counter-examples, but is rather a sort of 'intellectual headache', which will be 'cured' by carefully surveying the uses—and *not* by trying to formulate a general definition, nor any other sort of general doctrine (a view associated with Wittgenstein (1953)).

But these variants share two main ideas, which I take as characteristic of the Moorean view. Namely:

(i) Conceptual analysis is an 'armchair' i.e. non-empirical activity, closely tied to the examination of language, and 'intuitions' about what we would say in various imagined scenarios.

(ii) Conceptual analysis is descriptive, not reformative. It 'leaves everything as it is', i.e. it reports how we use our words and concepts (either aiming to give, or ducking out of, generalizations, according to the variant), and does not try to reform our use, not even to the extent of regimenting it.

The second view of conceptual enquiry, I call *Carnapian*, after Carnap (e.g. 1950). Carnap rejects both of (i) and (ii), above. To take (ii) first: he is willing to reform usage, for the sake of a better, in particular more precise, doctrine. To signal this difference, he says he aims to give an *explication*, not analysis, of the concept in question. And 'explication' harbours a systematizing ambition: he aims, not so much for a single concept's explication, or a handful of them, as for a theory, tying together several philosophically central or problematic concepts. (The main example in his own work was his theory of probability.)

As to (i): Carnap's explication projects are to find the best language to fit a particular scientific purpose. Sometimes (but not always), that means we must look to our best scientific theories for guidance in finding superior replacements for our old concepts. Nevertheless, the received concept—the concept in need of explication—is conceived as available to armchair reflection.

The Moorean and Carnapian views have had a large influence. One obvious case is 'Canberra Plan metaphysics' (cf. e.g. Braddon-Mitchell and Nola (2009)) and its inspiration, the writings of David Lewis: for example in his analytical functionalism about mind (with its appeal to the 'platitudes of folk psychology'; cf. e.g. Lewis (1966)), and his analysis of causation (e.g. Lewis (1973)). I will also be close to the general Carnapian idea, when endorsing the third view of conceptual enquiry, viz. Haslanger's; cf. Section 4.1.2. But I should first here register that both Moore's and Carnap's ideas can be, and have been, doubted and denied.

In the first place, one can doubt the quality of the evidence supplied by intuitions about what we would say. Do academic philosophers' intuitions match what would in fact be said by people in the street, i.e. by the 'folk' whose concepts the philosopher has set herself to analyse or explicate? Besides, the recent 'experimental philosophy' movement (Knobe and Nichols (2008) provide a recent collection of work) has gathered a lot of evidence that what people say is so various as to cast doubt on the existence of a communal, or common-sense, conceptual scheme that could be articulated, even if questionairres supplement armchair reflection.

In particular, what people say about some imagined scenario is often very sensitive to exactly how the scenario is described, or how the question about it is asked to them, or even their mood, or the topic of their preceding conversation. Such sensitivity, and consequent variety, certainly suggests there is no stable body of shared opinions or 'intuitions' of the type that philosophers in the Moorean tradition usually invoke in support of their analyses.[1,2]

On the other hand, these considerations may be thought to be all grist to the Carnapian mill: Carnap agrees that our received concepts are untrustworthy or unsuited to scientific ends (Demopoulos (2007))—this is why we must engage in projects of explication, with an eye on our scientific theories. However, explication is philosophical, as opposed to scientific, work, and requires a great deal of reconstruction of science as we find it. It may be argued that this process of 'rational reconstruction' relies as much on intuitive judgements as the Moorean projects.

But I will not here try to answer these doubts or denials, even for our restricted topic, the concept of a particle in quantum mechanics. I in fact think that, for our topic at least, Haslanger's view—her methodology of conceptual enquiry—is sound; and that an empirical enquiry into how 'particle' and cognate or similar words are used by quantum physicists would vindicate what I argue in Chapter 5 and later. I admit that these claims need to be defended. But I will not take space in this Chapter to do so: sufficient unto the day is the work thereof!

4.1.2 Haslanger's scheme

Haslanger's view (2006) is a variant of the Carnapian view. Like Carnap, she accepts that one can, even should, reform concepts and doctrines (rejecting (ii) above); and, also like Carnap, she agrees that reform must pay attention to exter-

[1] Ladyman and Ross (2007, pp. 10-15) also argue against the unbridled use of intuition in metaphysics, both on the grounds mentioned above, and because it fails to pay heed to the deliverances of modern science.

[2] One might also cast doubt on the armchair conception of philosophical enquiry, by showing it to be a contingent historical development. Thus Kusch (1995) describes how in Germany in the period 1870-1930, the fact that psychology grew apart from philosophy, defining itself as an experimental subject needing funding for laboratories etc. prodded philosophy into adopting a contrasting self-definition, i.e. as armchair enquiry.

nal, even empirical, matters (rejecting (i) above). But she develops this idea by distinguishing not just two concepts, the received concept (Carnap's *explicandum*) and a proposed improvement (Carnap's *explicatum*): but instead, three concepts.

(i): A trichotomy: Her trichotomy is based on considerations in philosophy of language, about how a speaker, or a whole linguistic community, can be mistaken, even self-deceiving, about what concept they operate with. In speaking, they may avow that they use a certain concept; while the rest of their behaviour suggests, or even shows, that they in fact operate with another concept. Often, and especially in cases of self-deception, this other behaviour is the person's actions: as the old adage has it, actions speak louder than words. So we need to distinguish the *avowed* concept (Haslanger calls it the 'manifest' concept) from the *operative* concept—even before raising the question of improving on one or both of them. The proposed replacement concept—a close cousin of Carnap's *explicatum*—is what Haslanger calls the *target* concept. This, then, is her trio of notions. And for her, the aim of conceptual reform is always to bring both the avowed and operative concepts into coincidence with the target concept.

However, I think that her distinction between the avowed and operative concepts is not controversial. Many authors use some such distinction; though of course they use different labels. Two examples from the (voluminous) literature in philosophy of mind, on self-deception (so now setting aside morality and politics) are Mellor (1977) and Dennett (1981, Chapter 16). Both take belief to be shown primarily by how one acts (corresponding to Haslanger's notion of the operative concept, and to 'actions speak louder than words'); while what one says reflects what one believes that one believes (corresponding to Haslanger's notion of the avowed concept).

On the other hand, in expounding the details of Haslanger's scheme, and in applying it in Chapter 5 and beyond to the concept of particle, it will help to keep in mind an uncontroversial scientific example of her trichotomy in action: that is, an example of conceptual reform, or transition to a target concept, being prompted by a perceived disparity between the avowed and operative concepts. (Haslanger's interest is in the reform of morally and politically charged concepts, such as 'race' and 'parent'.) I take as my example (from many that could be chosen) the scrutiny of, and changes in, the concept of *chemical element*, that arose from the discovery of isotopes.

(ii): An example: chemical elements: Thus the existence of isotopes entails that defining elements (or just individuating them) in terms of their atomic weight will yield a different classification scheme—a different periodic table—from defining them in terms of their atomic number. The historical summary (at the risk of gross over-simplification!) is that:

(a) before isotopes—or, indeed, electrons and protons—were discovered, the basic property guiding the (various) classifications of the elements was, due to Mendeleev, atomic weight (Scerri and Worrall (2001, pp. 438-9)); but

(b) since chemical behaviour is correlated with the electronic configuration, and therefore also with atomic number, rather than atomic weight (and chemists of course focus on chemical behaviour), the discovery of isotopes prompted chemists (after Moseley) to define elements in terms of atomic number (Scerri and Worrall (2001, p. 439)).

So using Haslanger's labels, we can say:

(a′) Before the discovery, the avowed concept of element tied it to atomic weight, while the operative concept tied it to atomic number, and so to chemical behaviour—or perhaps, tied it to some vague combination of weight and number, and perhaps other properties.

(b′) Given the strict association of atomic number with chemical behaviour, the target concept of element was tied to atomic number alone.

(c′) After the discovery of isotopes, chemists recognized the disparity in (a′) and the truth of (b′); so they self-consciously adopted atomic number as their new avowed and operative concepts, so that all three now coincided.

(iii): Subjective or objective?: The contrasting examples of race or parent (which are Haslanger's), and chemical element, raise another broad issue: whether a concept—equivalently: the distinction it draws in its domain of objects, or the

similarity it encodes—is subjective or objective. I deliberately choose these vague words, since I will not need to choose between the various possible precise meanings. Broadly speaking, 'subjective' means: relative to, or causally influenced by, any of various, individual or collective, biological or historical or socio-cultural factors. So 'objective' means, broadly, not being thus relative or causally influenced.

But these glosses leave plenty of scope for debate about the details. For example, suppose we say a concept is objective: the concept 'carves nature at the joints', as the saying goes. Is that to mean that the concept is independent of the theory in which concept is embedded? Or even independent of broader or more general topics: such as (a) theories that are kindred to the given one, and (b) the constitution of the human mind or community that formulated the theory? And the same questions can of course be asked about the domain of objects, in which our concept draws a distinction: how independent, if at all, is the domain's definition from the given theory, or from the broader topic of the enquiring human mind or community?

I will (fortunately!) not need to decide these questions; though I would endorse the common-sense view that the concepts of race and parent are 'more subjective' than that of chemical element or our concern, that of particle.[3] But I need to raise them since they of course bear on the idea of conceptual reform, and especially that of the target concept. For the idea is that a concept different from (but no doubt kindred to) both the avowed and operative concepts is more appropriate or successful than they are for articulating the division/distinction/similarity at issue in the relevant domain of objects. Often, especially in scientific or uncontroversial contexts, the suggestion will be that the target concept is "more objective", or more objective in some desired regard, than the avowed or operative concepts. But we need not always think of conceptual reform in this objectivist way. In particular, if one thinks it too objectivist to think of a concept as encoding a division/distinction/similarity, one might express conceptual reform as: the target

[3]Selections from a vast literature include, from philosophy of science, Jardine (1986), and Hacking (1999, Chapter 3); and from metaphysics, Lewis (1983) and Taylor (1993). For example: Lewis advocates a very strongly objective conception of similarity that is independent of theory, mind, society and indeed all contingency; Taylor replies that one can enjoy most of the benefits of Lewis' conception with a much less daring, in particular theory-relative, conception of similarity.

concept being more appropriate or successful for the goals associated with the use of the term. In any case, whether or not we think of conceptual reform objectivistically: adopting the target concept is meant to be an improvement. Thus Haslanger calls the project of determining the target concept *ameloriative* (Haslanger (2006, pp. 95-97)).

(iv): Norms: Mention of improvement raises the topic of normativity, i.e. the involvement of norms in determining the target concept. (Of course, such norms might not be subjective: only a strong, and probably false, subjectivism would say that they must be. So this topic cuts across the topic of subjectivity vs. objectivity.) I admit that there are such norms, but *what* these norms are will vary from one case to another—so I do not try to spell them out, except to make explicit two general considerations.

The first consideration is that the target concept should be in some sense 'natural'. It is hard to pin down this requirement in general terms,[4] since naturalness need not be a simple matter of "carving nature at the joints"; nor need it be a simple matter of suiting a particular anthropocentric end (Haslanger (2006, p. 109-11)); cf. also comment (iii), above. However, I can be more specific for the case of interpreting a physical theory. *For a theory's formalism comes with its own standard of naturalness.* The idea is to rule out gruesome interpretations of the theory's formalism, and ad hoc physical posits whose only role is to support over-stretched interpretations.[5]

The second consideration is that the target concept is subject to the usual canons of interpreting linguistic agents; particularly some form of the principle of charity (Davidson (1970)). So we ought not to interpret uses of the operative concept as too grossly in error, when compared with the proposed target concept. (If particles turned out to be the sorts of things that never existed in laboratories, for example, then that would make nonsense of a lot of what experimental physicists say!) Of course, what counts as error, and to what degree error is to be tolerated,

[4]Belnap (1993, p. 117) offers the phrase, attributed to Alan Ross Anderson, that an explication of a word, or its associated target concept, is '"a good thing to mean" by the word'.

[5]This claim is intended to be neutral between Lewis's (1983) dyed-in-the-wool realist commitment to natural properties and relations, and Taylor's (1993) more humble 'T-cosy' properties. Cf. footnote 3.

are questions best answered only in a specific context. However, I note here that: (i) considerations of naturalness and interpretative charity may well come into conflict; and (ii) the idea that moving to the target concept is an improvement (ameloriative, in Haslanger's jargon) depends on *some* significant degree of error being tolerable, for otherwise we could hope to learn nothing from such a project.

I also stress that it is not only for the target concept that normativity enters. It also enters in the determination of the operative concept. For the operative concept has a good claim to be the current linguistic meaning of the term concerned; and this also involves norms of interpretation, such as the principle of charity. Indeed, here charity must be applied rather more stringently than in the case of the target concept.

(v): A unique correct conceptual reform? No!: I will not claim to *prove* my view, i.e. to show to all that emergentism, our third target concept, is the best particle concept for the quantum mechanics of indistinguishable systems. Here, there is both a general point, and a specific one. The general point is that most philosophers, in particular analytic metaphysicians, agree that ontology is controversial, since the criteria for assessing proposals—and therefore what counts as relevant evidence—are not universally agreed.[6] Undoubtedly there are cases—let us call them the *good cases*—in which it is uncontroversial which concept is most apt to explicate an old one. (The identification of genes, the unit of inheritance, with more or less vaguely defined segments of a DNA strand may count as a good case, as does my example of 'element', above.) There are also cases—let us call them the *bad cases*—in which the general consensus is that the old concept is better left buried, because there is no target concept to go to. (Here surely belong phlogiston and the *élan vital.*) In between we have the *hard cases*, in which it can only be a matter of stipulation whether we take the received concept to be explicated or eliminated. (Perhap's Lewis's (1995) account of qualia falls into this category.) The majority of cases in science may well be hard.

The specific point is that, as we will see in Chapters 8 and 9, even for our circumscribed topic, viz. how quantum mechanics treats indistinguishable systems,

[6]Discussions include Sider (2001, Introduction) on the ontology of how objects persist through time; and Oliver (1996) on the ontology of properties and relations.

there are no clear winners: ours is one of the hard cases. So in fact there are several rival ontological pictures with various merits and de-merits, and it is a matter of judgment how to weigh them. But my judgment will be that emergentism is best.

Nor do I claim that there is, or should be, some single best meaning for the word 'particle': I allow that like many other physical concepts, e.g. least action and entropy, the concept of particle is a bit vague and flexible—and should no doubt be kept like that, both for convenience and for heuristic purposes.[7]

(vi): Neologisms: Finally, I should mention the topic of *linguistic reform*, i.e. whether we should mint a new word for a new target concept. Note first that projects of conceptual reform of course vary in many ways: for example, in how much the avowed and operative concepts diverge from the target concept; and in how unique is the target concept. Such considerations influence the decision whether to preserve the original term, or to make a clean break by introducing a new term. Good cases surely prompt the retention of the old vocabulary, while bad cases prompt linguistic reform. In hard cases, Haslanger distinguishes between what she calls 'constructivists', who favour retaining the old term, and 'eliminativists' (also called: 'error theorists') who favour replacing it with a new term.

As to our own topic, the concept of a particle: I have announced that I will argue that nothing in quantum mechanics satisfies all the desiderata. In Haslanger's jargon, nothing satisfies all the strands of the avowed, or indeed operative, concept. So any target concept must be substantially different from the avowed or operative one: and the next Chapters are devoted to hunting for the best such target concept. I will be a constructivist, not an eliminativist, about the term 'particle'—that is, I will preserve the term for the modified concept (Haslanger (2006, pp. 90-1). Indeed, even if we thought it best to change to a different term, we would have to admit that 'particle' is so entrenched that it would be a mug's game to try and abolish it: witness the way that neologisms like 'wavicle' that were invented so as to convey the un-particle-like, and even weird, behaviour of

[7]And I will accept that the jargon that ties 'particle' or 'quantum particle' to the factor Hilbert spaces, and their labels, is *so* entrenched as to be unchangeable. I do not expect the reader, even if convinced by this work, to never again say 'particle' or similar for the physical correlate of a label of a factor Hilbert space: I admit that I often do this myself!

so-called quantum particles, have never caught on.[8]

4.2 Physics, mathematics and language

In this Section, I make some specific comments about my conception of a physical theory and its interpretation. This will clarify how I intend to apply the idea of conceptual reform to the specific case of 'particle' in quantum mechanics. First (Section 4.2.1), I separate a theory's representational apparatus (mathematics) from its ultimate subject matter (physics). Then (Section 4.2.2) I borrow—and specialise—some important notions from intensional semantics to clarify my terms, particularly 'concept'. Then I turn (Section 4.2.3) to the representation relation between a theory's mathematical formalism and (putatively) physical entities, and conclude with some brief remarks about physical properties (Section 4.2.4).

4.2.1 Two realms: the physical and the mathematical

One of my central claims below will be that a certain element of mathematical formalism (a factor Hilbert space label, or, more properly, the position of a factor Hilbert space in a tensor product) does not represent what it is ordinarily taken to represent. I will also make some positive suggestions as to what pieces of mathematics represent what pieces of the physical world. To frame this discussion, I need first to make some orienting and clarifying—but hopefully uncontroversial—remarks about representation, and the realms that lie either side of the represen-

[8]In apolitical contexts (such as my project here), it is hard to disagree with Carnap that the choice between constructivism and eliminativism can only be a convention; though in more suasive/political contexts, other considerations may come into play, such as the desire to preserve a way of speaking. In this connection, Haslanger argues that constructivism is often preferable to eliminativism, since it allows us to give breathing space to "framework concepts" (roughly: concepts whose employment is endemic and central), whereas elimination requires a "wholesale adoption of a new conceptual scheme" (Haslanger (2006, p. 115)). But I reply that there can be no substantive difference here: on *both* approaches one is adopting a new conceptual scheme; the question is simply whether we make that plain by the adoption of new vocabulary, or whether we engineer the appearance of continuity by preserving the old vocabulary. There may indeed be *political* or *suasive* reasons to favour one approach over another, but *semantic* matters would be the same no matter which approach is taken.

tation relation: the physical and the mathematical.

I begin by endorsing the principle that, as Lewis (1992, p. 218), borrowing from Bigelow (1988, pp. 132-133, 158-159), puts it, 'truth supervenes on being': or, as Dummett (1991, p. 328) says, 'The principle that, if a statement is true, there must be something in virtue of which it is true, is a regulative principle that can hardly be gainsaid.' But I also maintain that this principle is not a substantive metaphysical commitment, for example to truthmakers. Rather, I see it, like Dummett, as a piece of philosophical methodological commitment, a 'regulative principle' that is a precondition for the discussion of substantive claims, not itself one of those claims.

The principle allows us to define a *realm* linguistically, i.e. as the subject matter of a particular *sphere of discourse*. I assume that our linguistic practice may usefully be divided into spheres of discourse. Then a realm is the minimal collection of objects, properties, and whatever else, whose configurations suffice to determine the truth or falsity—if anything does—of the sentences belonging to a single sphere of discourse.[9]

Thus the physical realm determines the truth or falsity (if anything does) of the sentences of a physical theory. It may contain quarks, electromagnetic fields and quasars; it may also contain polymers, people, and beautiful symphonies; it almost certainly contains Stern-Gerlach apparatuses, laboratory benches, and bubble chambers. It may even contain (as advocated by Lewis 1986) *possible* (i.e. non-actual) quarks, people and bubble chambers.

Similarly, the mathematical realm determines the truth or falsity (if anything does) of the sentences of mathematics. It contains integers, topological spaces, groups, sets and categories. Agreed, this list may double-count: for example, perhaps the mathematical realm contains *only* sets or *only* categories. But that is not a question I will need to address.[10]

[9]For a detailed account along this line, see Lewis (1988).

[10]Nor will I address the precise nature of the determination, by ontology, of truth or falsity. In the case of physics, the prospects are still good that supervenience can do the job. However, in the case of mathematics (along with many other spheres of discourse, such as ethics and aesthetics), supervenience notoriously appears useless.

I also make no commitment to realms being distinct, or to their overlapping or not overlapping. On the contrary, it seems, for example, that physical discourse shades off gradually and vaguely into discourse about everyday objects. In that case the physical realm and the realm of everyday objects will no doubt overlap; and they may well coincide. (As various authors, e.g. Healey (1978) and Ladyman & Ross (2007, p. 44), point out, physics is essentially "imperialist", since it is about whatever *anything* is made of, or at least anything concrete.) Similarly, I am not committed to the physical and mathematical realms being distinct. For all I know, the mathematical realm may be physical, as would be the case if some version of nominalist structuralism were true (e.g. Field (1980), though the term 'nominalist structuralism' is Horsten's (2007)). Or the physical realm may be mathematical, as would be the case if 'Blanket Pythagoreanism (Quine (1969, p. 59)) were true.

What is important for us is the *role* that mathematical and physical entities play: specifically, the fact that mathematics is used to represent physical facts and possibilities—perhaps indispensably. A popular, and tempting, image is of mathematical structures as being something like abstract ready-mades, to be plucked off the Platonic shelf according to their suitability as representative surrogates for the patterns in the concrete, physical world. This image makes a lot of commitments that I wish to shirk, but two less controversial features of it that should be preserved in understanding my project here are:

(i) the role of mathematical entities as *representations* (and the physical entities as *represented*); and

(ii) the *stipulative* nature of the way that mathematical elements are picked out or developed so as to serve those representational ends.

4.2.2 Language

One sphere of discourse that I have so far left out, but which is fundamental to my project, is language itself. There are items that we (probably) only need to refer to for the study of language, or allied topics such as theories of meaning. These include linguistic items (such as words and sentences) and meanings; possibly also things like universals and possible worlds. But the subject matter of language also includes everything else that can be talked, written, perhaps even thought, about. For all these things are necessary to determine the truth or falsity of all sentences about language and meaning. So the linguistic realm includes everything.

Let me be more specific about what a language is. I mean, roughly, the various systems of linguistic expressions—also called 'syntactic types' or 'signs'—and the various rules for their concatenation. Words, terms, phrases and sentences—conceived as types not tokens—are all signs. For examples, the numeral '2' is a sign, which refers to the number 2 (as it has just done!), and the noun-phrase 'the numeral "2"' is also a sign, which itself refers (in English) to a sign, viz. the numeral '2'. Signs are syntactic types in the sense that the numeral '2' is a single, repeatable object—normally taken to be abstract—of which the preceding instances (e.g. the very ink marks on the paper you are holding, or the collections of coloured pixels on the screen you are looking at) are repeated tokens.

This notion of language is entirely syntactic: a specification of a language in my sense does not include a specification of the referents or meanings of the signs taken to be meaningful; i.e. it does not include an *interpretation*. Interpretations ascribe to each sign of the language an item which codifies its contribution to the truth-conditions of the sentences in which it occurs. Following Carnap (1956) and Lewis (1970a), I call these items *intensions*.[11] Interpretations will also ascribe, to a privileged class of signs—notably, sentences and noun phrases—an *extension*. Extensions constitute the existential commitments of the sentences of a language.

Intensions, being functions, are undoubtedly mathematical items, but I remain agnostic about the ontology of possible worlds—i.e. about whether they are non-actual, concrete things or actual, abstract things—and therefore I am agnostic about what these mathematical items are made *from*. Extensions may be physical, or mathematical; and, if there *is* anything else, then extensions may be those things

[11]In the (degenerate) case of a sentence, the intension codifies its own truth-conditions, and may therefore be identified with a function from indices (each a specification of a possible world and a context of utterance) to truth-values. A sentence's extension is its truth-value; so, its intension is a function from indices to possible extensions. This general scheme is repeated for noun-phrases, whose extensions are objects, and whose intensions are therefore functions from indices to possible objects. For a sketch of this type of semantics, see Carnap (1956) and Lewis (1970a).

too.

Intensions are a suitable precisification of what I mean by *concept*. Thus, when I talk about the concept of particle throughout this paper, I can be taken to be talking about some candidate intension (or some class of candidate intensions) for the term 'particle'. Purists may hesitate to identify concepts with intensions, taking them to belong to essentially separate ontological categories. But no matter: my use of the term 'concept' is as adequately explained by the one-to-one *association* of intensions with concepts. I admit that the identification has only the practical value of tidiness; so I can remain equivocal on the matter.

4.2.3 Representation

In Section 4.2.1, I separated (if only functionally) the mathematical and the physical realms by putting each at one side of a representation relation. Thus mathematical items represent physical items, or the circumstances of physical items: mathematical quantities represent physical quantities; equations, i.e. relations between mathematical quantities, represent relations between physical quantities.

But representation is too widespread a relation to ground categorically the mathematics/physics distinction, because it is a relation which can hold in 'every direction': between items both of which are in either of the two realms, or between items both of which are in one realm. Aside from the representation involved in a physical theory, in which we use mathematical items to represent physical items (e.g. when we provide equations of motion for a simple harmonic oscillator), we also use physical items to represent mathematical items (e.g. when we write on a blackboard in maths class). Equally, we use mathematical items to represent other mathematical items (von Neumann ordinals represent the integers), and physical items to represent other physical items (an orrery represents the planets of our solar system).

Nevertheless, the interpretation of physical theories concentrates exclusively on the representation of the physical by the mathematical. We have our mathematical formalism; the task is to interpret it. Central to this task is the identification of elements within the mathematical formalism which are genuinely representative of something physical, or at least hypothetically physical; in other words: the identification of representative, as opposed to *surplus*, or *redundant*, structure. A theory cannot be true or false *of the physical realm* without such representatives. (A trivial caveat: it can be false if its mathematical formulation is inconsistent—that much we can know prior to interpretation.)

But can't anything represent anything else? In which case: how is the interpretative project well-defined? If all we demand is the theory's truth, isn't success too easy? (Cf. Putnam (1977, 1980).) The project *is* well-defined, given suitable concessions. There are two possibilities to consider, concerning our familiarity with the physical items being represented:

1. On the one hand, we may have *semantic access* to a given physical item that is independent of, or at least detachable from, the theory in question. That is, we may be able to refer to that item without recourse to the theory being interpreted. (This access may, of course, be due to the representational machinery of another theory that has already been interpreted.) In this case, we may specify the representation relation by specifying the relata: the mathematics of the new theory and the antecedently grasped physical item. (Ideally, ascertaining a theory's *experimental* claims involves setting down the representational relation in this way. Of course, this is far from a simple affair!—cf. Chang (1996).)

2. On the other hand, there may be no independent semantic access to a given physical item, or at least we may *know* of no independent semantic access to the item. (In fact it may be a familiar entity, unfamiliarly described.) In this case, the laying down of the representation relation and the specification of the represented physical item is the *same* act: the physical item is implicitly defined as that which corresponds, one-to-one, to some specified representative mathematical structure. This "structural" definition is, for all we know, the very best we can hope for (Poincaré (1905, p. 160-2)).

   Here I connect with a recent wealth of literature that fall under the general heading of *Structural Realism*.[12] There is not space here to engage with this

[12] See e.g. Worrall (1989), Ladyman (1998, 2002, 2007b), Worrall and Zahar (2001), French

literature, except to mention the link, and the suggestion (first made, about physical objects, by Ladyman (1998, pp. 420-2)) that structural characterisations of physical entities may be not just all we *can* know, but all there *is* to know.

The important point for me is that stipulating the representation relation for (apparently) unfamiliar physical entities is far from being a trivial exercise. It involves postulating a physical ontology (whether structural or not) for the representational mathematics to correspond to—in Lewis's (1983, 1984) jargon, it involves postulating an elite collection of natural properties and relations. But this does *not* entail any realism about that ontology—since the postulation of the physical ontology is *hypothetical*, dependent on the theory being true. And I need make no commitment to the truth of a theory in order to interpret it.[13]

I now connect my discussion of representation with my discussion of language, in Secton 4.2.2 above. The connection is this: in a physical theory we use mathematical language to refer to mathematical items, whose representation relation to the physical realm is specified as just described, by us, the interpreters. A physical theory may therefore be usefully divided into two parts: (i) its *formalism*, which consists of its language, in the sense of Section 4.2.2, and the mathematical entities to which the expressions of that language *immediately* refer; and (ii) the *representation relation* which holds between these mathematical entities and the physical world and-or its hypothetical surrogates. In this picture, mathematical language refers to the physical world by proxy, or by a 'zig-zag': first, by referring to the mathematical entities in its formalism; and second, to the physical realm, in virtue of the representation relation.

Typically, a theory's formalism will be interpreted to represent *modal* features of the physical world; and here I connect with general intensional semantics. For this modal structure provides the raw materials to construct intensions, such as the concept of particle.

and Ladyman (2003), Psillos (2006), Ladyman and Ross (2007).

[13] In this way, I retain my agnosticism between Lewis's natural properties and Taylor's "T-cosy" properties, as mentioned in footnote 3.

I will presently consider an explicit example. But first, I will list in general terms the interpretative commitments that I will for granted, both in my example and for the rest of the thesis.

1. As already indicated, I take myself to be giving a *realist* interpretation of the quantum formalism. That means I take vectors of the Hilbert space to represent physical *states*: which means, roughly speaking, temporal slices of the physical world, i.e. the way things are at a particular time in three-dimensional space. That means that I propose to take superpositions "seriously", i.e. as a feature of the physical world. "The way things are" is also supposed to capture other degrees of freedom, such as spin, or another "internal space". I do *not* interpret vectors of the Hilbert space as representing e.g. dispositions for various measurement outcomes (except insofar as these supervene on the occurrent features of the physical state *already* represented by the vector), as in the Copenhagen interpretation (Wheeler/Zurek); or our—or anyone else's—subjective or ideally rational state of information, except in so far as this may be gleaned from the physical facts.

2. Luckily for us, since I am interested here only in the quantum mechanics of systems presumed to be microscopic, I do not need to commit ourselves to any account about whether superpositions amplify to macroscopic bodies or conscious observers. So I will not need to address the measurement problem. I will, however, make use of probabilities (especially in my technical discussions in Chapter 7), so I assume that any solution of the measurement problem will vindicate the Born rule.

3. I take the quantum formalism at face value; I will not add to it. Therefore I put Bohmian intepretations of quantum mechanics to one side, for the most part; although I will occasionally mention implications for the Bohmian, only if these are interesting.

4. I take Hermitian operators to represent physical quantities, in the usual way. For instance: if $\hat{Q}$ represents the generalized position associated with the co-ordinate x, so that $(\hat{Q}\psi)(x, y, \ldots) = x\psi(x, y, \ldots)$, then $\hat{P}$ defined by $(\hat{P}\psi)(x, y, \ldots) = -i\hbar\partial_x\psi(x, y, \ldots)$ represents that system's momentum

along co-ordinate x. But I need only take the implication to go one way: i.e. an operator represents a physical quantity *only if* it is Hermitian.

Let me now turn to our explicit example. In quantum mechanics the sentence '$\langle \hat{Q} \rangle = q$' ostensibly asserts a relation between two mathematical entities: the Hermitian operator $\hat{Q}$ and the real number q. Namely: that its trace (when multiplied by a certain operator that represents the system's state—which is here left unspecified) is q. But suppose that we have specified that $\hat{Q}$ represents the distance of a certain particle, say, from the North-West corner of the laboratory floor, projected along the skirting board of the East-facing wall; and that q is a distance measured in metres. We also apply the Born rule, which specifies that for any quantity $\hat{A}$ and any state ρ, the mean value of $\hat{A}$ is given by $\langle \hat{A} \rangle_\rho = \mathrm{Tr}(\rho \hat{A})$. And we also stipulate that if no state is specified, context dictates that we take as ρ the actual present state of the system. Thanks to these stipulations about representation, '$\langle \hat{Q} \rangle = q$' says that the particle's present mean distance from the North-West corner of the lab floor, projected along the skirting board of the East-facing wall, is q metres. Thus we have (partially) specified an intension[14] for the sentence whose extension—i.e. whose truth-value—is determined by the way things happen to be (which may be established, at least in principle, and using other, admittedly substantive, assumptions, by an appropriate measurement).

Returning to our specific project: the concept of particle, as applied in quantum mechanics, will be an intension, which is fixed by specifying a function from quantum mechanical states to physical objects (namely, particles!). Since quantum mechanical states are represented mathematically, by vectors and density operators, what we seek is a general rule which tells us, for any such vector or density

[14] There are two subtleties which result from my proposed interpretation that I have suppressed here. (1) I have substantially coarse-grained the intension: an intension is a function from *indices* to extensions, but what I have just called a "intension" is a function from *states* (and contexts, somewhat impoverished) to extensions. But no worries: a state—which we can take as a possible world's timeslice—defines the set of worlds which have that state as a temporal slice; and we may treat values of the intension as constant over all variations of the context, except time. (2) I have restricted the intension: the "intension" is only defined for worlds which obey quantum mechanics (and which contain the particle, and the lab, in question). Whether this is a worry will depend on one's interpretative ambitions: surely it is no problem to leave the intension undefined if quantum mechanics were false or the particle did not exist. But perhaps we would prefer a physical interpretation that did not require the existence of our lab.

operator, what that vector or density operator says about the particles: e.g. what states they are in, how many there are, etc. It will obviously be a partial function, since it is restricted to quantum mechanical states, and so the corresponding intension will be partial. It is also clear that the identification of the "correct" function—i.e. the best target concept, in the sense of Section 4.1.2—will rely on a lot of other interpretative work having already been done. But that, I submit, is all right: my aim here is to tweak realist interpretations of the quantum formalism, not to create a new one *ab initio*.

4.2.4 A note about physical properties

As I have said, the task of theory interpretation is to find the best representation relation between the mathematical realm and the physical realm. It is therefore important that physical entities, and the mathematical entities which may or may not represent them, are kept at least conceptually distinct. However, there is one collection of physical entities—the physical properties—which seem to make trouble for this separation. The purpose of this Subsection is to clear up the trouble.

The problem is that the entities often deployed to do duty for properties—including physical properties—are mathematical objects. So, for example, some authors (e.g. Carnap 1956) identify properties with functions from possible worlds (or state-descriptions) to sets; others (e.g. Lewis 2002) identify them with sets (or set-like objects,[15] like classes), of any kind; and others still (e.g. Lewis 1970a) identify them with functions from intensions to intensions.

I acknowledge the benefits of each identification and I agree that for any property there is some intension, or some function, or some set (or set-like object), that may play its role. Therefore I agree that intensions, functions from intensions to intensions, and more generally sets, can indeed *do duty* for properties for various purposes.

[15]The property of being a set that does not belong to itself cannot be identified with a set! However, it may be identified with a proper class, which is for most intents and purposes "set-like".

How then do I save all properties from being mathematical by default? There are two main responses.

1. The most obvious move would be to acknowledge the benefits of the above identifications but ultimately resist them. There is, for example, an alternative, venerable tradition of identifying properties (at least some properties) with concrete universals (Murdoch (1970), Armstrong (1978)). Thus physical properties are essentially non-mathematical, although mathematical items (like those mentioned above) may represent them, for various purposes. This move is tantamount to breaking my agnosticism (in Section 4.2.1) about the overlap between the physical and mathematical realms.

2. Alternatively, I may accept any one of the identifications above, and admit that physical properties so construed are in some sense an abstraction from physical entities. However, even if I succumb to any of these identifications, I must still give an account of what properties are abstractions *from.* Standard moves are to take them as abstractions from primitive resemblance relations holding between either particulars or tropes, or as abstractions from a primitive distinction made between sets themselves, into natural and unnatural (Lewis (1983, p. 14; 1986, pp. 64-9)). In this case, it is the resemblance relations *themselves*, or the natural/unnatural distinctions, that are represented by the properties which have been identified with functions or sets. And the difference between the former and the latter is sufficient to make the work of theory interpretation non-trivial.

To sum up: physical properties do not make trouble for my account of theory interpretation, even if they are identified with mathematical objects. This concludes the set up of my general framework for conceptual reform. I now turn to my specific project: the concept of particle.

Chapter 5

Particles: what's in a name?

In this Chapter I discuss the concept of a particle in general terms (Section 5.1). This will build on the discussion in Chapter 4, especially Sections 4.1.2 and 4.2.3. Included is a brief discussion of a famous proposal of Wigner's (Section 5.1.2's interlude). Since I am aiming here to give an account of the *general* concept of particle, I will treat classical and quantum mechanics equally. Finally, I specialize to a single elementary quantum mechanical system, (Section 5.2). This will conclude (at last!) all the pre-requisites needed for my analysis, in Part III, of assemblies of quantum systems, especially indistinguishable systems.

But first I must make a crucial distinction. This distinction is between the operative concept of particle, as it is applied without reference to a particular physical theory (let us call it the *general* operative concept of particle), and the operative concept of particle-as-applied-in-quantum-mechanics (let us call it the *local* operative concept).

The distinction is crucial because, as I will argue in Chapter 6, the use of the term 'particle' in quantum mechanics—specifically in the context of indistinguishable systems—is in crisis: the local operative concept *conflicts* with the general operative concept. Indeed, this observation will form the basis of my criticism of factorism, since factorism holds that particles are the physical correlates of factor Hilbert space labels; and that, I contend, is true only for the local operative concept.

How can this be? Isn't the local operative concept just the restriction of the general concept to quantum mechanical worlds? No: the use of the term 'particle' outside of quantum mechanics suggests extensions to the quantum mechanical realm which are *subverted* by the current use of the term within quantum mechanics. This is the spirit of my objection against factorism in Chapter 6, and something which varietism and emergentism both seek to avoid. Thus the general operative concept is what we need to provide desiderata for any putative target concept, in order that the latter may be seen as an explication, specific to quantum mechanics, of the former (cf. (iv) in Section 4.1.2). Let us now turn to outlining the general operative concept.

5.1 Five desiderata for the concept of particle

Here I give a list of five constraints on, or desiderata for, any putative target concept of particle. They are strands, or components, of the operative concept of particle, which I have gleaned from the word's use by physicists, philosophers and philosophers of physics. (Admittedly, some of my evidence is introspective, in Moorean fashion (Section 4.1), so I may be accused here of giving some aspects of the avowed concept, in addition to the general operative concept.) I will devote a Subsubsection to each desideratum (Sections 5.1.1 to 5.1.5). They are like conjuncts that, taken together, constitute a functional definition of 'particle', i.e. an *intension* to associate with the term 'particle'.

This intension constrains our search for a target concept of particle as applied in quantum mechanics in the following way: as much as possible, the target concept ought to coincide with the restriction of this intention to quantum mechanical worlds. The proviso that it coincide 'as much as possible' is meant to echo the considerations from Section 4.1.2 above, that it may turn out (indeed, it *will* turn out) that *no* intension that is natural in terms of the quantum mechanical formalism will perfectly coincide with the restriction of this intension to quantum mechanical worlds. And again: there may be no concept which satisfies, to a sufficient degree, requirements of interpretative charity to physicists and philosophers of physics.

I do not claim that my list of five desiderata is the uniquely best list one could come up with. One might wish to lay down other constraints, or make some of my five more precise. For example, I would like to draw attention to the vast physics literature on solitons, in which a particle is taken to be essentially localised.[1] I make no such claim, but it is a natural strengthening of my second desideratum (Section 5.1.2), namely that particles have location as a degree of freedom.

Nor do I claim that the five desiderata have equal weight. In Part III, I will favour some of the constraints at the expense of others; that is, I give the constraints unequal weight.

5.1.1 Being physical

Our first constraint is that a particle is *physical*, in the sense discussed in Section 4.2.1. This I consider a compulsory (i.e. an indefeasible) component or strand, of the concept of a particle. The role of this desideratum is to exclude candidate target concepts which identify particles with what seem to be merely mathematical artefacts of the formalism of quantum mechanics. (This will be the basis of my criticism of factorism in Chapter 6.)

But I do not demand that particles turn out to be *particulars*, in the sense of items of which properties are predicated, but which cannot themselves be predicated of anything. That is, I allow that particles may be properties of something else. This allowance will be welcome for three reasons.

1. It will be welcome for any view that takes spacetime points, and perhaps regions (perhaps conceived as fusions of points), as the only particulars. For such a view needs to take whatever is the material content of the universe—particles, if such there be—as certain (perhaps higher-order) properties of spacetime points or regions. Indeed one version of emergentism, which I will consider in Chapter 9, proposes precisely this ontology.[2]

[1] For a physicists' introduction, see Manton and Sutcliffe (2004). For a philosophers' introduction, see Teh (forthcoming).

[2] This ontology appears to be popular; supporters include Carnap (e.g. 1934, §82), Quine (e.g. 1977) and Field (1980).

2. This allowance does *not* prevent particles being *objects* in the 'thin' sense (cf. Section 2.1) of Frege and Quine, i.e. potential values for first-order variables. In other words: whatever particles might be (whether particulars or properties), they could not fail to possess self-identity, and so to be objects in this thin sense. Thus I agree with some authors, such as (perhaps surprisingly!) Quine (1970, p. 28) and Lewis (1970b, p. 429), that a property ('attribute') can fall in the domain of (first-order) quantification. But I see no reason to follow the Quinean orthodoxy that existential commitment *must* be revealed by the contents of that domain alone. I reserve the right to use predicates *alone* to refer to properties. These declarations mesh with the first of the two responses to the problem of mathematically representing physical properties, as discussed in Section 4.2.4 (but neither do I rule out the second response).

3. This allowance has the advantage that it easily accommodates a denial one may want to make—indeed, *I will* want to make, when defending emergentism, in Chapter 9. Thus one may want to deny that particles exist in all quantum mechanical worlds. That is, maybe there are states in an assembly's Hilbert space for which it would simply be false that there are any particles. That is unpuzzling if particles are properties: for of course properties can be a feature of some states of an object but not others. Agreed: this situation would naturally lead us to ask *of* what, then, is an assembly an assembly? I will return to this puzzle in Chapter 9.

The requirement that particles be physical is compulsory, but how do we tell whether or not it has been satisfied? In essence this is a specific instance of the more general question facing theory interpretation, namely: How do we separate a theory's representative content from its purely mathematical artefacts? There is clearly no straightforward or complete answer to this question, but the task is considerably easier if we have at least a partial idea of what we are looking for. Thus, in the case of particles, we can appeal to the other four of our five desiderata. In this way the five desiderata are mutually supporting. It is to the other four desiderata that I now turn.

5.1.2 Being located

A particle is the sort of thing that is *located*, at least in some approximate sense. I will not need to be precise about what in general would count as 'located in an approximate sense'. It will be enough for me that classical mechanical particles are located, and quantum particles have position probabilities, i.e. probabilities to be found in various regions, were a position measurement made.

In more detail: It goes without saying for classical point particles, or the infinitesimal material points (also called: elements) in a classical continuous medium, are located. In quantum theory, the wave mechanics of a single 'particle' makes us familiar with the position representation: i.e. with (i) the state being represented as a $\mathbb{C}$-valued function on physical space, like a complex cousin of a classical scalar field, and so (ii) the position-operator $\hat{Q}$ being represented as multiplication—i.e. $(\hat{Q}\psi)(x) = x\psi(x)$—and $|\psi(x)|^2$ being a probability density for results of position measurements.

I require particles to have location as a degree of freedom, but as I mentioned in the preamble to this Chapter, I will not require them to be spatially local*ized*. I admit that this requirement has played a large role in the discussions of solitons and relativistic quantum theories.[3]

I register that one might argue that my desideratum that particles have a location makes redundant my previous desideratum that particles be physical. For, while many may disagree with the common sense view of Aristotle's time that everything (physical) that exists must exist *somewhere* (Morison (2002, pp. 18-19)), surely it is true that anything with a location is physical.

I am sympathetic to this claim—though there cases (such as phonons and centres of mass) that are by no means clear-cut! But there is a justification for listing both desiderata separately, which can be expressed as a dilemma. The dilemma concerns the identification of elements of the quantum formalism as representing genuine particles (i.e. genuinely physical entities), or as mere mathematical arte-

[3]Cf. Malament (1996), Fleming and Butterfield (1999), Halvorson (2001) and Halvorson and Clifton (2002) for the peculiarities of relativistic localization. I note that Fraser (2008) does *not* take localizability as a necessary condition for particlehood in her treatment. However, Fraser does lay down a necessary condition for particlehood which I have ignored, namely countability.

facts. If you do not take having a location as sufficient for being physical, then clearly both desiderata must be imposed. But if you take having a location to be sufficient for being physical, then you will still face the problem of distinguishing between mathematical elements which genuinely represent physical entities with location, and elements which only *seem* to: as we shall see, many putatively physical objects can be associated with the appropriate mathematical state space. In this case, we need to know *first* whether there *is* a genuinely physical object being represented, before we can say that it has a location. Therefore it is still valuable to list 'being physical' as a separate desideratum.

My inspiration for requiring particles to have location is a famous proposal of Wigner's for the definition of a concept related to that of particle, namely an *elementary system*. Let me briefly discuss the Wigner's proposal, and its relation to our project.

Interlude: Wigner's proposal
If one were to ask a practicing theoretical physicist today, 'What is a particle?', the answer one is very likely to receive is 'an irreducible representation of the Poincaré group', proposed by Wigner (1939). Does this not immediately provide us with our target concept of particle? No, but let me explain.

Wigner's proposal is an answer to a slightly different question. It was proposed as a definition of 'elementary system', and I will not require that particles be elementary (i.e. part-less) in this sense. (See also my discussion of compositionality below, in Section 5.1.4.) Nevertheless, Wigner's proposal provides a useful precisification of our requirement that particles be located.

For the uninitiated, Wigner's proposal is extremely puzzling: how can an elementary system, supposedly a physical object that bounces around in familiar, three-dimensional space, be something as mathematical as an irreducible representation of a group? The puzzle is merely a result of abbreviation: fully expanded, Wigner's proposal is that (the mathematical representation of) an elementary system's *state space supports* an irreducible representation of the relevant spacetime symmetry group. The Poincaré group is most often mentioned explicitly in this context, since this is the group of transformations between inertial frames in special relativity.

Thus Wigner's proposal says what an elementary system *is* by saying what it *can do*: i.e., what physical states are available to it, given the kinematical restrictions incumbent on it due to the spacetime symmetries governing the physical laws. It is inspired by the idea that an elementary system exists in space and has no internal structure; so the only things that it can do within the confines of the physical laws is to be translated, rotated and boosted with respect to the spacetime background. Therefore its state space is generated by the spacetime transformations which are symmetries of the theory, and contains no further structure; so it is irreducible. And there is no problem of Wigner's proposal being mathematical in nature, since we *expect* that the modal profile of an elementary system be *represented* by something mathematical, like a state space.

Particles need not be elementary systems, but some may be. Common to both is their having location as a degree of freedom: this determines that the relevant state space be a representation of the spacetime symmetry group. We may borrow this aspect of the concept of elementary system for our purposes, but drop the requirement of partlessness that constrains the state spaces to be *irreducible* representations.

Therefore my desideratum that particles have location can be cashed out in the following terms: any target concept must pick out as particles entities whose modal profile is represented in the formalism by a state space that supports a representation (reducible or irreducible) of the spacetime symmetry group.

The idea that particles need not be elementary systems is supported by the idea of *dressing* in interacting quantum field theory.[4] In that theory, it is a common technique to simplify a physical problem by treating a "cloud" of what one first construes as distinct systems—for example a "bare" electron surrounded by local vacuum fluctuations—as a single system. Non-elementary particles also include systems which are themselves thought of as made out of particles, such as protons and neutrons. I return to the issue of dressing in Section 5.2.2.

The requirement that the formalism contain a representation of the spacetime symmetry group for each particle is perhaps also more familiar than it may

[4]For a philosophers' introduction, see Mattuck (1992, Chs. 0-2).

first sound. For example, for my special concern with composite systems: the tensor product structure in quantum mechanics and the Cartesian product structure in classical mechanics—both used to represent assemblies of *distinguishable* particles—fulfil this requirement straightforwardly. Important for us, however, is the fact that less familiar structures may also fulfil this requirement. For example (though there is no space here to pursue the matter in detail), the reduced configuration space still contains subspaces which support representations of the Galilei group. This is important because in the case of interest for us—namely, quantum mechanics for *indistinguishable* systems—the assembly's state space, which is a symmetric *sector* of a tensor product Hilbert space, has no obvious tensor product structure. I will return to this matter in Section 7.2.
End of Interlude.

To conclude this Section, I stipulate that the desideratum that particles—or any nearby surrogate for them—have a location be utterly non-negotiable.

5.1.3 Persisting over time

Particles do not just have a location. They are treated as spatio*temporal*, that is, persisting, entities. In classical mechanics, a point-particle has a *trajectory*: which we can model as a map from an interval of time (some interval of real numbers) into spacetime. This idea is generalised in quantum mechanics (in the Schrödinger-picture position representation) to a map from a time to spatial wavefunctions.

The idea of particles as persisting entities may seem non-negotiable, but to a philosopher it immediately suggests a nearby concept for which persistence over time is *denied*: namely, *particle-stages*, which exist momentarily, i.e. at one time only.

Here we meet the philosophical controversy over how to understand persistence over time. According to *endurantists* (e.g. Wiggins (1980), Haslanger (1989), Merricks (1999)), an object at any two times is identical: it is 'wholly present' at all and only those times at which it exists. According to *perdurantists* (e.g. Quine (1976b), Lewis (1986, p. 210), Sider (2001), Hawley (2001)), a persisting object is composed of temporal parts, any two of which from different times are distinct. I will endorse perdurantism, but there is no room here to defend my endorsement. Particle-stages are thus minimally persisting temporal parts of a particle.

I also endorse Lewis's view (1986, pp. 218-20) that, like world-bound objects and their trans-world composites, object-stages are in a sense primary to their trans-temporal composites. The sense in which they are primary is not ontological but practical: both objects and object-stages exist "on an equal footing"; but while it will always be *useful* to talk of object-stages when it is useful to talk of objects, there are conceivable contexts in which it is not useful to talk of objects, but still useful to talk of object-stages (Reichenbach (1958, §43); Butterfield (2006b, §4)).

Object talk breaks down in favour of object-stage talk when the world is unkind enough to provide no uniquely appropriate, non-overlapping trans-temporal mereological sums of object-stages. Cases like these include unmanageable fission and fusion of trans-temporal objects (Parfit 1971) but are not limited to these examples. We will see, in Section 7.3 and Chapter 8, that quantum mechanical worlds *are* unkind in this way, at least to the varietists' and emergentists' particles. So here, the corresponding particle-stages may provide the best target concept.

To sum up (and looking ahead a bit): (i) by endorsing perdurantism and temporal parts, we allow ourselves to take advantage of otherwise unfortunate circumstances in which talk of persisting objects is problematic, but talk of their temporal parts is not; and (ii) I forewarn the reader that this unfortunate circumstance faces us upon the rejection of factorism.

5.1.4 Composing assemblies

In both everyday life and science, we all believe (or at least avow!—cf. (i) in Section 4.1.2) that countless familiar objects—like the solid objects that Austin (1962, p. 8) jokingly called 'moderate-sized specimens of dry goods'—are made of particles. This yields our fourth desideratum for any target concept of particle: we hope for a reasonable sense of 'being made of' according to which this claim comes out true. I call this requirement 'compositionality' or 'being compositional'. In this subsubsection, I will sharpen this vague idea of something's 'being made of' some other things. I will end by commenting on how particles being compositional is compatible with: (i) holism of the kind much discussed in quantum philosophy; and (ii) particles *not* being fundamental.

In philosophy, the most discussed and widely accepted notion of composition is mereology.[5] The mereological sense of an X 'being made of' some Ys is perhaps the most direct sense possible: namely, identity. On this understanding, X is made out of the Ys in virtue of X *being* the Ys. Each Y is part-identical to X by its being a part of X. (In general, two objects are part-identical iff they share a common part, but one need not be a part of the other.) The notion of mereological composition as part-identity has recently been attacked (e.g. Yi 1999). But I will accept it in this paper: in fact it will be useful to me in my discussion of varietism in Chapter 8.

Although I accept mereological composition as a sense of 'being made of' that is legitimate as a strand of the operative concept of particle, it is not the only legitimate sense. It may turn out, for example, that the route from particles to the familiar objects of everyday experience—Austin's moderate-sized specimens of dry goods—is rather less direct than simple part-identity. Two examples immediately come to mind:

(i) Macroscopic objects are often thought to have vague boundaries, but mereological fusions have boundaries that are as sharp as their parts.[6]

(ii) Applying mereology to a collection of entities technically requires those entities to comprise the basic ontology; but we may not wish to include particles *or* macroscopic objects in our basic ontology; i.e. we may not want to *quantify* over them. For example, a nominalist about properties may still wish to treat either particles or macroscopic objects (or both) as properties. This nominalist will have to construe talk of particles or macroscopic objects as disguised talk about whatever concrete particulars *are* taken to exist. Then it is hard to see how mereological fusions of particles could either make sense, or else produce the correct candidates for macroscopic objects.

[5]For an introduction to mereological composition, see the classic paper by Leonard & Goodman (1940).

[6]This is, of course, Unger's (1980) and Geach's (1980) problem of the many. I address a close cousin of this problem in Chapter 8.

So I turn to a natural alternative sense of 'being made of': namely, supervenience. The rough idea is to stipulate that X be made of some Ys iff X's properties supervene on Y's properties and relations. Some clarification is in order: I make five comments.

1. Note that my proposal is only reasonable if the Ys in question *exhaust* X, in the sense that there can be nothing which composes X which is not already one of, or itself composed of, the Ys. For otherwise supervenience can easily fail in cases of composition. Take a house, for example: it is composed of its bricks, amongst other things, but the way things are with the house could easily change without anything changing for the bricks—we might, for example, replace all the windows.

2. The bi-conditional above is trivially satisfied if we range sufficiently liberally over the properties of X and the Ys: e.g. if it is taken as a property of X that $X = X$ and the Ys are thus-and-so. Therefore I exclude relational properties of X which make explicit reference to the Ys, and demand that the Ys' properties and relations make no explicit reference to X. (X's properties may be *reducible* to the Ys' properties and relations, but that is another thing.) Another way of putting this is that Xs properties on one hand, and the Ys' properties and relations on the other, should belong to different *levels*. Let us call the way things are with X the *high-level* facts, and the way things are with the Ys the *low-level* facts.

3. Supervenience is covariance over worlds, and we require covariance over all quantum mechanical "worlds"; i.e. all states in the assembly's Hilbert space. This is the most liberal quantification over worlds possible while still remaining within the quantum mechanical framework. This is important, since supervenience of a sort unsuitable for our purposes can often arise by restricting the range of worlds one quantifies over. For example, if we restrict to only the dynamically possible worlds, then in deterministic worlds (in virtue of the dynamical laws), the properties instantiated at distinct Cauchy surfaces will supervene on one another—yet we do not want to say that distinct Cauchy surfaces compose one other!

4. A superficial objection to my suggested precisification of 'being made of' is that, as e.g. Horgan (1993, pp. 560-1) has pointed out, supervenience does not exclude distinctness of the properties concerned: it is not a necessary truth—still less an analytic truth—that covariance precludes distinctness. I reply that the relation between supervenience and identity is a matter of one's conceptual scheme,[7] and that anyone who denied identity in instances of supervenience (where the quantification over worlds is sufficiently liberal) should also deny that composition is a relation of part-identity. Either way, composition may be cashed out in terms of supervenience.

5. Finally, I note that my supervenience understanding of composition is a logical weakening of the mereological understanding, if we assume that the properties of a mereological fusion supervene on the properties of and relations between the mereological proper parts. This assumption is not an axiom in, e.g. Leonard and Goodman's (1940) calculus, but it is a feature of Lewis's modal framework.[8]

To conclude the discussion of compositionality, I will make three comments: about how particles being compositional is compatible with: (i) holism of the kind much discussed in quantum philosophy; (ii) particles *not* being fundamental; and (iii) multiple realizability.

(i) Our notion of compositionality makes no restrictions on what the lower-level facts mention, so long as they *are* purely *lower*-level facts: in particular they

[7]Lewis's (e.g. 1983) implementation of Hume's dictum, that there are no necessary connections between distinct existences, commits him to the conclusion that supervenience implies at least part-identity.

[8]This seemingly innocuous assumption vetoes a peculiar form of property dualism. A person differs from a zombie, according to substance dualism, because the person has proper parts (namely, a soul) that the zombie doesn't, even though they may be physically identical. A person differs from a zombie, according to one form of property dualism, because the person's proper parts bear non-physical properties, or relations to each other, that the zombie's parts don't, even though, once again, they may be physically identical. What the assumption above rules out is an alternative form of property dualism in which a person may differ from a zombie only in the possession of a higher-level property (consciousness, say), despite being composed of the same (physical and non-physical) proper parts, and despite those parts possessing the same (physical and non-physical) properties and relations. Suffice it to say that I find this doctrine bewildering, and I doubt that anyone has ever believed it.

can be facts about *relations between* lower-level entities, as well as facts about the lower-level entities' intrinsic properties. This bears on so-called 'quantum holism': that for many states of an assembly—namely, entangled states—the state does not supervene on the intrinsic states of the assembly's constituents (Teller (1986, pp. 78-80)). This failure of supervenience has been taken as evidence against the constituents' existence "prior to" the assembly (Massimi (2001), Hawley (2009)). (As we shall see in Chapter 9, I am sympathetic to this position, but *not* because of the phenomenon of entanglement.)

But a more conservative response would be to take it—*ceteris paribus*—merely as evidence for the existence of irreducible *relations* between the particles which go to make up the assembly (Teller (1986, p. 80)). In fact, this response is also available to those who stick to the narrow notion of composition as mereological composition: for the properties of a mereological sum typically do not supervene on the intrinsic properties of its parts. Therefore, our requirement of compositionality does not exclude particles from being the constituents of an assembly, despite the fact that their states typically fail to subvene the state of the assembly.

I should point out that this debate over holism in quantum mechanics assumes the interpretative position that I deny for indistinguishable systems—namely, factorism. But Teller's point still stands (especially for distinguishable systems, for which I endorse factorism). However, as we will see in Chapter 9, we are, after all, forced to except a holism about quantum assemblies that *cannot* be overcome by countenancing irreducible relations between constituents. It is ironic that the states which prompt this holistic retreat are characterised, not by being entangled, but rather by *failing* to exhibit a form of entanglement (cf. Section 8.3). To anticipate: the culprit that compels holism is not entanglement, but the indistinguishability postulate.

(ii) The idea that particles are fundamental is in a way "complementary" to the idea that macroscopic objects are made out of them. Compositionality looks upward, as it were, in its claim that entities at higher levels are made of particles: fundamentality looks downward, in its claim that particles are not themselves made out of anything. But, as I mentioned in the Interlude in Section 5.1.2, above, I do not take it as necessary that particles be ontologically fundamental.

(iii) Finally, I note that the requirement that particles compose macroscopic objects does not entail that they *necessarily* do so. That is: being composed out of particles is consistent with being composed out of particles only *contingently*. This is not just a philosophical side-remark: it will be important later, in Chapter 9, where I entertain the idea that there are states of the assembly in which there are no particles, and yet there is still the possibility for stable macroscopic structures to emerge.

5.1.5 Being applicable across different theories

If possible, we should hope to arrive at a single concept of particle which vindicates its use in a variety of physical theories. The ideal is that there be a single intension, defined here by the five desiderata, whose restriction to the worlds of any theory coincide with that theory's most natural candidates for the particle concept. It would be a shame to have to retreat to a collection of homonyms, each one indexed to a particular theory, and all connected to each other by some sort of 'family resemblance'.

Prima facie, there is reason to be optimistic. *Functional* definitions offer the promise of being general enough to be applicable in a variety of situations (in our case, a variety of theories), and yet specific enough to pick out, in each of those situations, a salient entity or category of entities (e.g. Lewis 1972). Giving a functional definition of 'particle' falls significantly short of providing a *reduction* of particle talk, independent of any particular theory. On the contrary: different theories already provide what for each of them are eligible candidates for reference; our functional definition then simply picks the candidates (if any) which satisfy the definition, i.e. which (if any) of them are worthy of the name 'particle'.[9]

It cannot be a *component*, or strand, of the concept of particle that it have inter-theoretic applicability: a concept, being an intension, is supposed to determine a unique extension in a variety of possible worlds and contexts, and extensions cannot themselves have 'inter-theoretic applicability'. Rather, inter-theoretic applicability

[9]We can, if we so wish, then investigate further to provide a complete account of what particles in classical mechanics are like, or what particles in quantum field theory are like; but none of this will compromise the fact that these will be investigations into the nature of *particles*.

is a desideratum for the concept itself, a constraint which guides the evaluation of putative target concepts. Nevertheless, the constraint, like the three above, is inspired by the actual use of the term 'particle', since it is certainly applied in a way that is independent of any particular theory.

The obvious fleshing out of inter-theoretic applicability is as follows: *if two theories are connected by a relation of partial reduction in some limit, then the particles of the reduced theory ought to coincide with the particles of the reducing theory in that limit.* Thus, for example, if classical particles have determinate trajectories, then particles in quantum mechanics ought to be such that, in the appropriate classical limit, they too have determinate trajectories. (It is precisely this consideration, of course, that led Schrödinger (1926) to identify Gaussian wave packets in quadratic potentials as the quantum analogues of classical particles; cf. Landsman (2007, p. 429).) In this way, inter-theoretic applicability can be seen as a generalisation of the correspondence principle.

It may suggested that, with inter-theoretic applicability as a desideratum, we need only apply the other desiderata for *one* theory T, and let inter-theoretic applicability do the rest for any other theory T^* in a relation of partial reduction to T. Certain entities in T^* would then inherit particle-hood purely in virtue of their acting like the particles of T in the limit in which T^* and T nearly coincide. Away from this limit, T^*'s entities could be as unlike particles as you please, and yet they would still count as particles.

But I am more strict about our use of the term 'particle': for it seems more natural to say of the above case, that T^*'s entities are particles *only* in the limit in which T^* nearly coincides with T. Therefore, I take the requirement of inter-theoretic applicability as an *additional* desideratum on putative target concepts, and *not* an excuse to relax the other desiderata. I must also treat it as a *negotiable* desideratum amongst the other constraints: for we may be forced to be constructivists (in the sense discussed in comment (vi) of Section 4.1.2, above) about the concept of particle for a variety of theories, and we cannot expect different theories to pull in the same direction.

5.2 The desiderata for a single quantum system

To sum up the discussion in Section 5.1, and to introduce my main topic, quantum mechanics, I report here how our five desiderata are largely satisfied by the elementary wave mechanics of a single system. This Section has two parts. In the first part (Section 5.2.1) I state the usual surprises of elementary wave mechanics and argue that, despite these, the Hilbert space of a single system represents a particle according to our five desiderata. (This trival case, in which the indistinguishability postulate does not apply, is the only one for which I endorse factorism.) However, an interesting result will be that, even in this simple case, there will be *different* natural interpretative moves to make; these will correspond to my later three proposals, factorism, varietism and emergentism. The second part (Section 5.2.2) discusses natural factorisations of Hilbert spaces, and the idea of *dressing*.

5.2.1 Satisfying the desiderata for "single-system" Hilbert spaces

(i): Being physical. The idea that the vectors—or, more properly, the rays—of the single system Hilbert space represent *physical* states is uncontroversial; the question is whether we may say that the physical states are possessed by a particle. Thus we are led immediately to consider the other four desiderata.

(ii): Being located. If we consider no other degrees of freedom, then the Hilbert space for a single system is $L^2(\mathbb{R}^3)$ and is equipped with an algebra of quantities, which includes the position and momentum operators $\hat{Q}, \hat{P}$. We can then define a family of unitary operators $U(g), g \in G$ which comprise a unitary representation of the group G of transformations which preserve the background spacetime. The space generated by these transformations is the state space of a particle, as required by the Interlude in Section 5.1.2.[10]

This simple situation becomes interestingly complicated if the "single-system" Hilbert space incorporates more degrees of freedom—e.g. spin—so that the Hilbert space is $L^2(\mathbb{R}^3)\otimes\mathbb{C}^{2s+1}$, for some non-negative half-integer s. In this case, different

[10]For details where G is the Galilei group, see e.g. Jordan (1969, Ch. 7).

interpretative strategies part ways; (so here we glimpse the issues to come in Part III). Some may continue to consider the full Hilbert space as representing the states of a single particle; indeed this is the orthodoxy. However, this interpretative move is not compulsory. Some interpretations may instead consider the Hilbert space to have states which represent *more than one* particle at once, since the extra degree of freedom gives "room" for more than one representation of the spacetime symmetry group. For example, with each projector E_j onto a ray generated by a unit vector $|j\rangle$ of the spin Hilbert space $\mathbb{C}^{2s+1}$, we may associate a distinct representation $U(G) \otimes E_j$ of the Galilei group G.[11] So any state in which the spin and spatial degrees of freedom are entangled, i.e.

$$|\psi\rangle = \sum_j a_j |\phi_j\rangle \otimes |j\rangle \tag{5.1}$$

(where $a_j \neq 0$ for $j \in \{j_1, j_2\}$, for some $j_1 \neq j_2$), will yield different orbits under action by each $U(G) \otimes E_j$ for $j \in \{j_1, j_2\}$. I.e., each E_j selects out the $L^2(\mathbb{R}) \otimes |j\rangle\langle j|$ copy of $L^2(\mathbb{R})$.

As we shall see later in Chapters 8 and 9, this interpretative proposal goes hand in hand with using states, as opposed to Hilbert space labels, to individuate particles—which is the hallmark of any interpretative strategy which I will (in Part III) call *anti-factorist*. I emphasize that under this interpretation, individuation by Hilbert space labels is rejected even in the case of a single system, where there is only *one* Hilbert space label.

It may well be asked of proponents of these anti-factorist interpretative strategies: If states in which the degrees of freedom are entangled represent *more than one* particle, in what sense, then, are we still talking about a "single-system" Hilbert space? The answer is that, while these states represent more than one particle, each particle in one such state belongs to its own non-entangled "branch" of the superposition. Thus for *all* states, it is still true that each non-entangled *branch* contains only one particle. Equivalently, it is true both for the orthodox interpretation, and the heterodox ones here outlined, that every state of the Hilbert space is an eigenstate of the 'total particle number' operator, familiar from

[11]Note that these representations are not unitary, since the projectors E_j have no inverse.

quantum field theory, with eigenvalue 1.

If the anti-factorist countenances the representation of more than one particle in some states, then the question arises how particles are to be cross-identified between states and between branches of a superposition. I postpone this important discussion until Section 8.3.3, where I will have had the chance to more thoroughly investigate anti-factorist interpretations.

(iii): Persisting over time. Proponents of factorism, who consider the "single system" Hilbert space to represent a single particle—indeed, the *same* particle for each state—can straightforwardly account for trans-temporal persistence. To see this, note that in the position representation, one may conceive the particle as a $\mathbb{C}$-valued field, spread out over space. Since the states of the Hilbert space are the configurations of this field, it is unproblematic to consider the field itself as persisting over time. I note here that, for similar reasons, trans-temporal identification is always unproblematic for the entire assembly, and that the case here is special in that the assembly comprises only a single particle.

Proponents of what I have called the anti-factorist interpretation may account for persistence equally easily, so long as the Hilbert space represents only the spatial degree of freedom. If other degrees of freedom are included, then once again the story becomes complicated, since the number of particles represented in each state may not even be conserved over time—although note that the 'total particle number' operator is a conserved quantity, since the number of particles *in each branch* is conserved. Thus issues of trans-temporal identification for anti-factorists overlap with issues of cross-identification between states and branches of superpositions. Once again I postpone further discussion, until Section 8.2.3.

(iv): Composing assemblies. Composition is vacuously satisfied in this case, since for each state—or, at least, for each branch of every state—the assembly comprises only one particle.

(v): Being applicable across different theories. The various results pertaining to the quantum/classical boundary for a single system are a vast subject, and so deserve a separate treatment in their own right. Here I will merely refer to the review by Landsman (2007) relating to both the $\hbar \to 0$ and $N \to \infty$ limits. In

particular, the role of coherent states in achieving the right coincidence between the classical and quantum theories may suggest a strengthening of our desiderata in Section 5.1—for example, to the demand that particles be local*ized*. But I will not pursue that line here. Recall (from Section 5.1.5) that I only require that the quantum particle have an approximately classical trajectory *for those states* (in fact they are coherent states) for which we can define a classical limit. The non-classical behaviour of the quantum particle outside of that limit will not preclude its deserving the name 'particle', so long as our five desiderata are satisfied.

5.2.2 Natural decompositions and dressing

Section 5.2.1's invocation of $L^2(\mathbb{R}^3)$ set aside an important aspect of our topic, to which I now turn. This will pick up on Section 5.1.4's mention of *dressing*.

Recall that a Hilbert space can be factorized in many different ways! More precisely, a Hilbert space of non-prime dimension can in general be factorized (with factors of dimension greater than one!) in many different ways. And a Hilbert space of denumerable dimension can be factorized, with factors of denumerable dimension, in infinitely many different ways. Besides, these different ways are, in general, physically significant. For example: For the second statement, recall the isomorphisms, essentially due to Fourier analysis:

$$L^2(\mathbb{R}) \cong L^2(\mathbb{R}^2) \cong L^2(\mathbb{R}) \otimes L^2(\mathbb{R}) \tag{5.2}$$

where the final tensor product can be realized in infinitely many ways, by rotating axes in $\mathbb{R}^2$. That is: it can be realized as either:

1. the square-integrable functions on the x-axis, tensored with the square-integrable functions on the y-axis; or as
2. the square-integrable functions on the $y = x$ line tensored with the square-integrable functions on $y = -x$ line; or as
3. the square-integrable functions on the line $y = ax$ tensored with the square-integrable functions on the line $y = -\frac{1}{a}x$, where $0 < a < \infty$;

and so on. So quite apart from whatever subtleties beset the use of 'particle' for indistinguishable systems—and such subtleties will be centre-stage later—we here see the question: What counts as the, or at least, *a* physically natural factorization of a quantum Hilbert space?

And we see that this is a question that arises even for what is usually called a 'single particle', or for a pair of distinguishable such particles. Fortunately, I will not need to address this question in full generality. But I should here register three main ingredients of its answer: ingredients which I will sometimes mention, in connection not only with factorism but also with my two other proposals, varietism and emergentism.

(i) I should no doubt allow that the answer is not uniform; as indeed, the vague phrase 'physically natural' suggests. Thus the answer could vary according to the problem at hand. It might depend upon some salient physical quantity, pre-eminently of course the Hamiltonian. It might even depend on human interests: on what one is trying to calculate. These sorts of consideration are familiar in the concept of *dressing*. Thus one can sometimes simplify a problem by 'dressing' the system, i.e. one associates to the system degrees of freedom that one initially considers external to it. Formally, this amounts to revising how one factorizes the total Hilbert space. To anticipate: the non-uniqueness of natural factorizations, for Hilbert spaces of assemblies of indistinguishable systems, will eventually lead me to endorse emergentism, in Chapter 9.

(ii) The example of dressing also brings in another main ingredient in the answer: the idea of locality or—a better word—proximity. That is: if one adds degrees of freedom, so as to define a bigger sub-system (in the formalism: augmenting a factor space), one will normally add degrees of freedom that are spatially localized at or beside the originally given sub-system.

(iii) The quantum information community *has* addressed the above question in general terms, often emphasising how the physically natural factorization depends on what interactions and measurements are operationally accessible (cf. e.g. Zanardi (2001), Zanardi *et al* (2004)). This work will inform some significant new developments in Section 7.2, in which we seek natural factorizations of *subspaces* of an assembly's Hilbert space, given the assumption of the indistinguishability postulate.

This concludes our consideration of the "single-particle" case in quantum mechanics, and with it the outline of my general philosophical framework. We have constructed a general functional definition of 'particle' which we will use to identify, within the quantum formalism, the most natural bearers of that name. In the next and final Part, I will investigate three rival interpretative proposals, each of which stakes its own claim as to what, exactly, quantum particles are.

Part III

What is a quantum particle?

Chapter 6

Against factorism

In Section 6.1, I define my first interpretative proposal, factorism. In Section 6.2, I distinguish it from haecceitism: Section 6.2.1 defines haecceitism, and Section 6.2.2 establishes the distinction between these 'isms'. Section 6.3 discusses the recent debate on the fate in quantum mechanics of the principle of the identity of indiscernibles. In that Section, I make some novel amendments to the proposals made by Muller and Saunders (2008) and Muller and Seevinck (2009), and argue that the factorist should endorse the amended proposals. Then in Section 6.4, I reject factorism for indistinguishable (but not distinguishable) systems, by appealing to the framework of Chapter 4, and to Chapter 5's desiderata for the concept of particle. The leading idea will be that the factorist's particles are not physical—they are a statistical construction, like the average taxpayer.

6.1 Factorism defined

Factorism says: Particles are the physical correlates of the labels of factor Hilbert spaces. This view is orthodox: indeed, not just orthodox, but well-nigh universal. It is deeply entrenched in the way we all speak and think, and learn, about quantum mechanics for more than one system. To explain this, and how the view is nonetheless deniable, it will be clearest to begin by considering, first, a single quantum system, and then *distinguishable* systems.

For a single quantum system, we recall Section 5.2: there, we accepted the 'surprises' of the quantum mechanics of one system (such as superpositions, states being wave-functions, etc.), and agreed nevertheless to call the system a 'particle', since our five desiderata (Section 5.1) were adequately satisfied.

Now consider an assembly of *distinguishable* quantum systems. This pair is represented using a tensor product Hilbert space $\bigotimes_i \mathcal{H}_i$. (We can take the systems to be distinguished by properties not represented in the Hilbert space, for example mass or charge.) Each factor Hilbert space $\mathcal{H}_i$ represents the space of pure states for each particle. The full space of states—including the mixed states—for each particle is then represented by $\mathcal{D}(\mathcal{H}_i)$, the space of density operators defined on $\mathcal{H}_i$. Each of these spaces is subject to the conclusions reached for single systems in Section 5.2.1, despite the surprises associated with non-product states. Factorism takes the orthodox position there outlined—namely that each Hilbert space represents the possible states for a particle, so that each Hilbert space label may taken to represent its corresponding particle.

I concur. I am happy to take this step: I agree that distinguishable particles are the physical correlates of the labels of factor Hilbert spaces, in the usual tensor-product formalism.

But factorism goes beyond this agreement. It says that the same goes for *indistinguishable* systems: that also for an assembly of indistinguishable systems, particles are the physical correlates of the labels of factor Hilbert spaces. Or in other words: *although* such an assembly is described by the symmetric or antisymmetric subspace of the tensor product (according as the systems are bosons or fermions)—or for paraparticles: by a similar sort of subspace—this does *not* disrupt the factor spaces' labels referring to particles. Thus when one treats an assembly using the symmetric or antisymmetric subspace (or an appropriate paraparticle subspace) of $\bigotimes_i^N \mathcal{H}_i$, factorism says that there are N particles, one for each factor space, and the ith particle's states are represented by the density operators in $\mathcal{D}(\mathcal{H}_i)$.

I agree that factorism is well-nigh universal, in the way we all speak and think about indistinguishable systems. We all tend to call the subscripts i 'particle labels', and to say 'an assembly of N indistinguishable particles'; and so on. And

I agree that this way of speaking is:

(i) convenient, not least because it meshes with what we all say about distinguishable systems—which I myself endorse; and

(ii) so entrenched that it would be difficult, if not impossible, to try and change it if one disagreed with it (cf. fn. 7 in Chapter 4).

Nevertheless, I will disagree with it, in Section 6.4.

I end this Section with a clarifying remark about my denial of factorism. Recall that, according to my philosophical framework (see Section 4.2.3), a theory comprises a formalism and an interpretation. The advantage of this is that we may separate questions of physical significance from straightforward issues of reference. That point has particular application here, since I do not deny that the mathematical expressions '$\mathcal{H}_i$', '$\mathcal{D}(\mathcal{H}_i)$', etc. have referents. All should agree that they *do* have referents: namely, certain mathematical objects. The point of controversy is whether these mathematical objects *represent* anything, or anything straightforwardly; i.e., whether they have *physical* correlates. Factorism says that they do: they represent the states of particles. I disagree. Thus my criticism of factorism is not so much that its particles don't exist, but rather that they are nothing but mathematical constructs. They cannot, therefore, really *be* particles, since particles are physical.

6.2 Factorism and haecceitism

In this Section, I distinguish factorism from another doctrine, *haecceitism*, which is more familiar in philosophy (both in metaphysics and philosophy of quantum theory). I first introduce haecceitism (Section 6.2.1); and then contrast it with factorism (Section 6.2.2). The leading idea of this contrast will be: factorism is a proposal about what particles are (and thereby about what (some) physical objects are); while haecceitism is a doctrine about how the identity or 'which-is-which-ness' of objects, in particular of particles, contributes to the individuation of states. So factorism, and later Sections' other proposals for what particles are (varietism and emergentism), are in a sense *prior* to the question of haecceitism, in that they specify the physical objects about which one can then assess doctrines about 'which-is-which-ness'; although haecceitism, as usually defined, presumes factorism. Thus haecceitism will come up again in my later discussions of varietism and emergentism.

6.2.1 Haecceitism defined

Haecceitism and its denial, anti-haecceitism, are instances of a general issue, which relates to my discussion of representation in Section 4.2.3. In philosophers' jargon, the issue is whether a distinction is real, as against 'merely verbal', 'spurious' or a 'distinction without a difference'. In the jargon of my philosophical framework, the issue is whether there are *redundancies* in the representation of physical possibilities by the mathematical objects in our theory's formalism. That is, whether the representation relation between mathematical states (vectors or rays of the Hilbert space) and physical states (instantaneous "possible worlds") is many-to-one. In physics there is also the jargon of 'gauge': the issue in these terms is whether certain quantities are 'gauge-invariant' or 'physical', as against 'gauge-non-invariant', or 'gauge', or 'redundant'.

Having stated the general issue, we only need to specify the distinction with which haecceitism is concerned. It concerns the action of permutations on states. One envisages, in either classical or quantum mechanics, an assembly of N indistinguishable systems. The symmetric group S_N acts on these N systems; and there is a natural induced action on the space of states. In classical mechanics, the space of states is the N-fold Cartesian product of the single-system phase space; and in quantum mechanics, it is the suitably symmetrized N-fold tensor product $\mathcal{S}_\mu(\bigotimes_i^N \mathcal{H}_i)$, where the parameter μ registers whether we project onto the symmetric (boson) or anti-symmetric (fermion) sector, or else one of the sectors of 'mixed' symmetry (paraparticles). The induced action of S_N is of course permutations of system labels (which are usually called, in factorist jargon, 'particle labels'!). We can write, for a state s, be it classical or quantum: $s \mapsto \pi(s), \pi \in S_N$.

There may be some states that are wholly symmetric in the sense that their orbit under this action is a singleton set, i.e. contains only the state in question: $s = \pi(s), \forall \pi \in S_N$. But the generic state will have a non-singleton orbit. (This is so in both classical and quantum mechanics—see below.) So the question arises whether all the elements of the orbit represent the same physical state of affairs. Here of course we invoke the distinction between mathematical states s and the physical states (sometimes dubbed, for emphasis: physical states of affairs) that they represent.

Thus I define *anti-haecceitism* as always answering 'Yes' to this question. This answer implements the intuitive idea of treating the underlying identity of each system, the 'which-is-which-ness' of the systems, as merely gauge or verbal.

On the other hand, I define *haecceitism* as saying that distinct mathematical states in an orbit represent distinct physical states. Intuitively, this implements the idea that the underlying individuality of the systems is physical or real, but this further claim is not required—cf. Section 6.2.2.

To link up with Part I: Note that my new definition of 'anti-haecceitism' is logically stronger than the one I settled on in Section 2.1.1. My new definition is the denial of haecceitistic differences (thereby coinciding with Lewis's (1986, p. 221) definition) rather than the denial of combinatorial independence; and, as we saw, combinatorial independence entails haecceitistic differences, but not *vice versa*. My new definition of haecceitism is therefore correspondingly weaker. However, our new anti-haecceitism is still *structuralist* in the sense of Section 3.3; and our new haecceitism is *intrinsicalist* if one also assumes factorism, so that Hilbert space labels are taken to represent an underlying individuality to the systems.

Haecceitism, in our new sense, can be assessed in classical or in quantum mechanics. Finite-dimensional classical mechanics (e.g. of N point particles) is, so far as I know, almost always formulated haecceitistically, i.e. so as to distinguish states differing by a permutation of indistinguishable particles.[1] And in infinite-dimensional classical mechanics (i.e., the mechanics of continuous media—fluids or solids), one must be a haecceitist (cf. Butterfield (2011, pp. 358-61)).

[1]Belot (2001, pp. 56-61) considers the anti-haecceitistic alternative.

In quantum mechanics, it is conventional to impose the Symmetrization Postulate (SP): i.e. to require that particles all be bosons or all be fermions. This trivializes the assessment of haecceitism: only if we consider paraparticles (i.e. multi-dimensional irreducible representations of the symmetric group S_N) do haecceitism and anti-haecceitism get clearly distinguished (Caulton and Butterfield 2011). Thus a haecceitist says that if a vector (or ray) lying in such a multi-dimensional irreducible representation is the image under a permutation operator of another, then they represent different physical states. But the anti-haecceitist says that all vectors (rays) in the given irreducible representation represent the same physical state.

I am an anti-haecceitist about quantum mechanics (and of course haecceitists about classical mechanics). My rationale is the usual, orthodox one; (at least: usual and orthodox since the work of Messiah and Greenberg (1964)). Namely: the indistinguishability of quantum systems implies that any physical (gauge-invariant!) quantity Q must be symmetric i.e. permutation-invariant, as follows.

More precisely, indistinguishability implies, in the first instance, that quantum expectation values are permutation-invariant, i.e. for any state-vector ϕ, and any permutation operator P representing $\pi \in S_N$: $\langle P\phi|Q|P\phi\rangle = \langle\phi|Q|\phi\rangle$. But this implies $P^\dagger QP = Q$, i.e. $[P, Q] = 0$: Q is permutation-invariant, also known as: symmetric (French and Krause (2006, pp. 142-3)). This is called the *Indistinguishability Postulate*, (IP), by Messiah and Greenberg (1964).

It follows (by Schur's Lemma) that any two vectors (rays) in a given irreducible representation of S_N give the same expectation value to any symmetric (i.e. permutation-invariant) quantity. But I identify a physical state with an assignment of expectation values to all physical quantities (cf. comments 2 and 3 in Section 4.2.3). So I conclude that any two such vectors (rays) represent the same physical state. In short: I infer that anti-haecceitism is true.

I submit that almost all physicists and philosophers would agree with my anti-haecceitism, and would agree with my rationale for it, based on (IP); and would also agree that the contrast here, between classical haecceitism and quantum anti-haecceitism, is the root of all the debate about 'quantum identity'. In particular, anti-haecceitism prompts one to re-evaluate traditional philosophical doctrines about identity and indiscernibility that were first formulated with the particles of classical mechanics (or more generally, the objects of everyday life or of classical physics) in mind. In the literature, the main focus has been on one such doctrine, the principle of the identity of indiscernibles (PII); and in Section 6.3, I will review the debate over this—it is a debate that assumes factorism.

Let me sum up. In the literature about identity and indiscernibility in quantum mechanics, factorism is well-nigh universally assumed, while haecceitism is widely rejected. This is itself enough to strongly suggest that factorism and haecceitism are distinct, and thereby accomplish my present goal of distinguishing them. If they were the same, how could so many wise authors have thought one could endorse one and deny the other? But in the next Subsection I go into fuller detail about their differences.

6.2.2 Factorism and haecceitism distinguished

The spirit of haecceitism is that the underlying identity, or 'which-is-which-ness' of indistinguishable objects is a real, or physical, or non-gauge, matter. But in Section 6.2.1, I defined haecceitism more narrowly and precisely, in terms of permutations of system-labels (in factorist jargon: particle-labels), in classical or quantum mechanics, yielding a different physical state. It will be important here to keep distinct the vague spirit of haecceitism and its precise definition in Section 6.2.1, since it turns out that the precise definition expresses the spirit of haecceitism only if factorism is assumed.

Consider: if the precise definition of haecceitism, in terms of system permutations, is to be true to the spirit of haecceitism, the system labels must represent the objects whose which-is-which-ness the haecceitist wants to defend as real or physical. Besides, these objects must be physical, not mathematical, in the philosophical sense discussed in Section 4.2.1: for I am not concerned with haecceitism, or its denial, about mathematical objects.

Now recall from Section 6.1 that factorism also is precisely the doctrine that the system-labels represent physical objects that deserve the name 'particles'. Thus we see that haecceitism as defined in Section 6.2.1 implies—indeed one might say, presupposes—factorism.

But apart from this inference, haecceitism and factorism are independent. There are two points here. First: there is the gap just noted between the spirit of haecceitism and Section 6.2.1's definition: so by adopting a different conception of what the particles are, we can construct a different precisification of haecceitism that does not entail factorism. As I mentioned at the start of Section 6.2, varietism and emergentism will adopt such different conceptions, and so prompt a renewed assessment of haecceitism. Furthermore, anti-factorists can separate out Section 6.2.1's doctrine from the spirit of haecceitism altogether, and ask the more neutrally worded question, whether mathematical states differing only by a permutation of system labels (*whatever* they represent, if anything) represent the same physical state. This will no longer be a question about haecceitism, but considerations of descriptive redundancy will still apply, and will favour the physical identification of permuted mathematical states.

Second: factorism does not imply haecceitism, even in the narrow sense of Section 6.2.1's definition. I noted there that almost all philosophers of quantum theory (and physicists, if they care to consider the issue!) are factorist anti-haecceitists. So here we should beware of a *non sequitur*: from the formalism distinguishing two factor Hilbert spaces, and their referring to particles (as factorism asserts), to taking the distinct mathematical states related by a permutation (s and $\pi(s)$ in Section 6.2.1's notation) to represent distinct physical states. Remember that, anyway, the issue only arises for paraparticle states, since for bosons and fermions even the *mathematical* states related by a permutation are identical.

So a factorist need not be a haecceitist. And, while haecceitism as usually defined entails factorism, it does so in virtue of presuming it. In fact, what you will accept as a precise definition of haecceitism depends on your interpretative position about what the particles are. So there may be haecceitists who are not factorists. I will address haecceitism again, in my treatments of varietism and emergentism; but first I turn to a familiar topic within factorist quantum mechanics: the discernibility of quantum particles.

6.3 Factorism and discernibility

Here I review the recent philosophical debate about whether quantum particles can be discerned. In that debate, all hands assume the factorist concept of a particle—which I will deny. But despite this, I need briefly to review the debate: for the ideas are needed both in Section 6.4's rejection of factorism, and in my later discussion of varietism and emergentism.

Throughout this Section, for the sake of charity and clarity, I will acquiesce with the factorist orthodoxy, and use the term 'particle' to mean the physical correlate of a factor Hilbert space.

6.3.1 The old orthodoxy

In short, the situation is this. Until about eight years ago, the orthodoxy was that quantum particles (in the factorist sense) cannot be discerned, so that Leibniz's principle of the identity of indiscernibles fails. But then Saunders, later joined by Muller and Seevinck, argued that they can be discerned: the springboard for these arguments is the idea of *weak discernibility*, as articulated in the Hilbert-Bernays account of identity (see Part I).

To explain this in more detail, it will be clearest to start with an earlier tradition, from the founding fathers of quantum mechanics. This says that Pauli's exclusion principle for fermions—or better: symmetrization for bosons and anti-symmetrization for fermions—means that:

(a) bosons can be in the same state; but

(b) fermions cannot be; so that

(c) PII holds for fermions but not bosons.

(For an expression of these three views, see e.g. Weyl (1928, p. 241).) In fact, these claims can and should be questioned. Under scrutiny, each of (a) to (c) fail (at least for factorist particles), and it seems that PII is pandemically *false* in quantum mechanics.

For first: according to the factorist, two bosons *or* fermions that are constituents of an assembly are absolutely indiscernible (in the sense of Section 2.3). For any assembly of indistinguishable quantum systems (whether fermions or bosons), and any state of the assembly (appropriately (anti-) symmetrized), and any two particles in the assembly: the two constituents' probabilities for all single-particle quantities are equal; and so are appropriate corresponding two-particle conditional probabilities, including probabilities using conditions about a third constituent.

This can also be expressed in terms of the reduced density operators of the constituent particles. According to the usual procedure of yielding the reduced density operator of a particle by tracing out the states for all the other particles in the assembly, we obtain the result that for all (anti-) symmetrized states of the assembly, one obtains *equal* reduced density operators for every particle.[2]

Thus not only can fermions 'be in the same state', just as much as bosons can be—*pace* the usual slogan form of the exclusion principle. Also, a pair of indistinguishable particles of either species *must* be in the same state. This result suggests that PII is pandemically *false* in quantum theory.[3]

6.3.2 The new orthodoxy?

Such was the orthodoxy until eight years ago. But this orthodoxy has *also* been recently shown to fail! For (building on the Hilbert-Bernays account of identity) there are ways of distinguishing particles that outstrip the orthodox notion of a quantum state for a particle (and the associated probabilities, including conditional probabilities)—and yet which *are* supported by quantum theory.

For as the Hilbert-Bernays account teaches us, two objects can be discerned, even while sharing all their monadic properties and their relations to all other

[2]A selected bibliography for this result runs as follows: Margenau (1944), French & Redhead (1988), Butterfield (1993), Huggett (1999, 2003), Massimi (2001), French & Krause (2006, pp. 150-73).

[3]I note parenthetically that these results can also be shown to hold for paraparticles, so long as one follows Messiah and Greenberg's (1964) recommendation of working with 'generalised rays' (i.e. multi-dimensional subspaces) instead of one-dimensional rays. See Huggett (2003).

objects (cf. Section 2.3)—and even though any relation that they hold to one another is held symmetrically. That is, they can be discerned *weakly*. Thus if, for some relation R and two objects a and b, we have that Rab and Rba, then a and b must be distinct if either Raa or Rbb (or both) fails.

The remaining task is to provide such a relation in the context of quantum mechanics. This task was undertaken in its most general form for fermions by Muller and Saunders (2008), and for all particles by Muller and Seevinck (2009).[4] I will now (in Sections 6.3.3 and 6.3.4) briefly appraise the results in these two papers. I will conclude that Saunders, Muller and Seevinck were largely correct, but that their proofs make incorrect assumptions, on their own terms, about which aspects of the quantum formalism represent genuine physical structure. I will propose a friendly amendment to the Saunders-Muller-Seevinck results in Section 6.3.5, and secure the fact that the factorist's particles are always merely weakly discernible, whether they be bosons, fermions or paraparticles.

6.3.3 Muller and Saunders on discernment

Here I briefly present the main result contained in Muller and Saunders (2008). First I follow these authors in establishing three important distinctions in the way that particles may be discerned.

1. *Absolute vs. relative vs. weak discernment.* The first distinction relates to the logical form of the predicates used to discern the particles (cf. Section 2.3). As we have seen, all fermions and bosons are absolutely indiscernible; they are also relatively indiscernible. Thus our only hope is to discern them *weakly.*

2. *Mathematical vs. physical discernment.* Of course, it is crucial that the properties and relations used to discern the particles be *physical*, in the sense of Section 4.2.3: we cannot appeal to elements of the theory's mathematical formalism which have no representational function. Thus, for example, we

[4]This work built upon an original suggestion by Saunders (2003b), which took inspiration from the fact that two particles in the spin singlet state may be said to have opposite spin (or to have vanishing combined total spin) without picking a basis.

cannot discern two particles in an assembly merely by appealing to the fact that the Hilbert space for that assembly is a tensor product of two copies of a factor Hilbert space. For all we know, this representative structure may be redundant; there may in fact only be one particle. So we must instead appeal to quantities in the formalism which genuinely represent *physical* quantities. Like Muller, Saunders and Seevinck, I call this sort of legitimate discernment 'physical discernment'. I call instances of spurious discernment 'mathematical discernment'—Muller and Saunders instead use the phrase 'lexicon discernment'; but I distinguish between mathematical objects (like Hilbert spaces) and mathematical language, and so my jargon meshes better with my general philosophical framework.

3. *Categorical vs. probabilistic discernment.* The final distinction relates to the assumptions required to secure the discernment. Muller and Saunders call an instance of discernment 'categorical' just in case it requires no appeal to the Born rule, and 'probabilistic' otherwise. The main advantage of categorical, as opposed to probabilistic, discernment is that by by-passing probabilistic notions its validity need not wait on any solution to the quantum measurement problem. However, this advantage is in my view only slight, since surely any solution to the measurement problem must secure at least an approximate vindication of the Born rule.

We are now in a position to state the main Muller-Saunders result:

(SMS1) Fermions are categorically, weakly, physically discernible.

Reconstruction of proof (cf. Muller and Saunders (2008, p. 536): We consider an assembly of only two fermions, so our Hilbert space is $\mathcal{A}(\mathcal{H}\otimes\mathcal{H})$; the result is easily extendible for more than two particles (cf. Muller and Saunders (2008, p. 534)). Select some complete set of projection operators $\{E_i\}, \sum_i E_i = \mathbb{1}$, for the single-particle Hilbert space $\mathcal{H}$ and define $P_{ij} := E_i - E_j$. Then define $P_{ij}^{(1)} := P_{ij} \otimes \mathbb{1}$ and $P_{ij}^{(2)} := \mathbb{1} \otimes P_{ij}$, where the superscripts are *labels* for particle 1 and 2. We then define the following relation:

$$R_t(x, y) \quad \text{iff} \quad \sum_{i,j} P_{ij}^{(x)} P_{ij}^{(y)} \rho = t\rho, \tag{6.1}$$

where $t \in \mathbb{R}$, ρ is the density operator representing the state of the assembly, and the indices i, j sum over the projectors E_i.

First we prove that 1 and 2 are categorically and weakly discerned by R_t for some value of t. To see that the discernment is categorical, it can be shown (cf. Muller and Saunders (2008, p. 533)) that, with $\dim(\mathcal{H}) \geqslant 2$, for every state $|\Psi\rangle \in \mathcal{A}(\mathcal{H} \otimes \mathcal{H})$,

$$\sum_{i,j} P_{ij}^{(1)} P_{ij}^{(2)} |\Psi\rangle = \sum_{i,j} P_{ij}^{(2)} P_{ij}^{(1)} |\Psi\rangle = -2|\Psi\rangle \tag{6.2}$$

and

$$\sum_{i,j} \left(P_{ij}^{(1)}\right)^2 |\Psi\rangle = \sum_{i,j} \left(P_{ij}^{(2)}\right)^2 |\Psi\rangle = 2(d-1)|\Psi\rangle \ , \tag{6.3}$$

where $d = \dim(\mathcal{H})$. Thus every state of the assembly is an eigenstate of the operators used in the definition of R_t; and so we do not need to assume the Born rule. R_t therefore promises to provide categorical discernment.

To see that R_t discerns the particles weakly for some t, note that $R_t(1,1)$ and $R_t(2,2)$ iff $t = 2(d-1)$, whereas $R_t(1,2)$ and $R_t(2,1)$ only if $t = -2$.[5] So the relations $R_{2(d-1)}$ and R_{-2} both serve to weakly discern particles 1 and 2.

Finally, it remains to be shown that R_t is a *physical* relation. I turn to Muller and Saunder's criteria (2008, pp. 527-8):

> (Req1) *Physical meaning.* All properties and relations should be transparently defined in terms of physical states and operators that correspond to physical magnitudes, as in [the weak projection postulate],[6] in order for the properties and relations to be physically meaningful.
>
> (Req2) *Permutation invariance.* Any property of one particle is a property of any other; relations should be permutation-invariant, so binary relations are symmetric and either reflexive or irreflexive.

[5] Remember that '1' and '2' serve as particle labels in the expressions '$R_t(1,2)$', etc.

[6] The weak projection postulate is effectively Einstein, Podolsky and Rosen's (1935) reality condition that the assembly's being in an eigenstate of any self-adjoint operator Q with eigenvalue q is a sufficient condition for the assembly's possessing the property corresponding to the quantity's Q having value q. This is an interpretative principle, which, like Muller and Saunders (2008) and Muller and Seevinck (2009), I take for granted.

(Req2) is clearly true of R_t. (Req1) is also true of R_t, provided that: (i) the projectors E_i are physically meaningful; and (ii) the physical meaningfulness of operators is preserved under mathematical operations; for our purposes these must include: arithmetical operations, i.e. addition and multiplication; and tensor multiplication with the identity. (Note: Muller and Saunders take (i) (along with (Req2)) to be sufficient to establish that R_t is a physical relation (2008, pp. 534-5). However, it is clear that (ii) is also required.) $\square$

Commentary. I take no issue with Muller and Saunders' claim that their relations R_t provide categorical and weak discernment. However, I question whether the relations R_t may properly be considered physical. I take no issue with the idea that projectors *per se* be physically meaningful (like Muller and Saunders, I agree that these can be considered to represent specific experimental questions with a yes/no answer); but R_t is defined in terms of *non-symmetric* projectors $E_i \otimes \mathbb{1}$, etc. Yet, being anti-haecceitists, I take it as compulsory—that is, as a necessary condition for representing a physical quantity—that the quantities obey the Indistinguishability Postulate.

This brings us to my criticism of (Req2). My criticism has two components. First: it misapplies the correct idea that physical quantities must be symmetric. By requiring only that the relations defined *from* the quantum mechanical quantities be symmetric, (Req2) fails to rule out quantum mechanical quantities which may themselves be non-symmetric. To take a simple illustration of this point: 'x is particle 1 and y is particle 2' clearly fails to be a physical relation, both in the proper sense, and in terms of (Req2). But the relation 'x is particle 1 and y is particle 2, or x is particle 2 and y is particle 1' is equally unphysical, yet it satisfies (Req2).

It may be replied: this is where (Req1) comes in. But this brings us to the second component of my criticism of (Req2): it is redundant. For it is anyway necessary for a quantity to be symmetric to satisfy (Req1), since any non-symmetric quantity contravenes IP, and therefore cannot represent a 'physical magnitude'. Indeed: since (Req1) already demands that the quantities be physical, why do we need another requirement at all?

It may be objected on behalf of Muller and Saunders that, while the quantities $P_{ij}^{(1)}$ and $P_{ij}^{(2)}$ indeed fail to be symmetric, the quantities defined in terms of them—namely, the $\sum_{i,j} P_{ij}^{(x)} P_{ij}^{(y)}$—*are* symmetric. This is indeed true: $\sum_{i,j} \left(P_{ij}^{(1)}\right)^2 = \sum_{i,j} \left(P_{ij}^{(2)}\right)^2 = 2(d-1)\mathbb{1} \otimes \mathbb{1}$ and $\sum_{i,j} P_{ij}^{(1)} P_{ij}^{(2)} = \sum_{i,j} P_{ij}^{(2)} P_{ij}^{(1)} = 2(\sum_i E_i \otimes E_i - \mathbb{1} \otimes \mathbb{1})$, where $\mathbb{1}$ is the identity on $\mathcal{H}$. (Note that the restrictions of both quantities to the anti-symmetric sector, $\mathcal{A}(\mathcal{H} \otimes \mathcal{H})$, are multiples of the identity on that sector.) But I see no force in the objection: the physical significance of these quantities was supposed to rest on their being constructions out of quantities like $E_i \otimes \mathbb{1}$; yet it is precisely these quantities which run afoul of IP.

Without any convincing account of the physical significance of the building blocks of the $\sum_{i,j} P_{ij}^{(x)} P_{ij}^{(y)}$, these quantities must be assessed for their physical significance on their own terms. But, since they are all multiples of the identity on the assembly's state space, this significance is trivial: they all correspond to experimental questions which yield the same answer on every physical state.

This triviality is a problem for Muller and Saunders, since it blocks the R_t from being physical relations. If we now attempt to redefine the R_t in a way that avoids misleading reference to the chimerically physical quantities $P_{ij}^{(x)}$ we obtain:

$$R_t(x,y) \quad \text{iff} \quad (x = y \text{ and } 2(d-1)\rho = t\rho) \text{ or } (x \neq y \text{ and } (-2)\rho = t\rho) \tag{6.4}$$

This is equivalent to:

$$R_t(x,y) \quad \text{iff} \quad (x = y \text{ and } t = 2(d-1)) \text{ or } (x \neq y \text{ and } t = -2). \tag{6.5}$$

So long as we have a definition of the R_t in terms of quantities that seems (i.e. from the point of view of the syntax) to treat the $x = y$ and $x \neq y$ cases equally, the fact that a *different* quantity (i.e. a different multiple of the identity) underlies each of these two cases is tolerable. (In just the same way, Rxy and Rxx are strictly speaking different predicates—since one refers to a relation while the other refers to a monadic property—yet it is normal to treat any instance of Rxx as a (special) instance of Rxy. Indeed: weak discernment *relies* on this being legitimate.) But since the quantities $\sum_{i,j} P_{ij}^{(x)} P_{ij}^{(y)}$ must be taken at face value—that is, as nothing but multiples of the identity—we must adopt definition (6.5) over definition (6.1), and definition (6.5) is hopelessly gerrymandered and unphysical. Thus Muller and Saunders' proof that any two fermions are *physically* discernible does not go through.

In Section 6.3.5, I propose an alternative relation which will discern fermions physically and weakly, though not categorically. But first let me address the main results in Muller and Seevinck (2009).

6.3.4 Muller and Seevinck on discernment

Muller and Seevinck use a similar framework to Muller and Saunders (2008): specifically, they carry over the three distinctions between kinds of discernment presented above, and the two requirements for physical significance, (Req1) and (Req2).[7] There are two main results to discuss: the first concerns spinless particles with infinite-dimensional Hilbert spaces; the second concerns spinning systems with finite-dimensional Hilbert spaces.

I begin with their Theorem 1. (Note that I rephrase their Theorems; cf. Muller and Seevinck (2009, pp. 189).)

(SMS2) In an assembly with Hilbert space $\bigotimes^N L^2(\mathbb{R}^3)$ and the associated algebra of quantities $\mathcal{B}(\bigotimes^N L^2(\mathbb{R}^3))$, any two particles are categorically, weakly, physically discernible.

Reconstruction of proof (cf. Muller and Seevinck (2009, p. 189)): Again, for simplicity's sake, I restrict attention to the case of two particles ($N = 2$). Let Q be the position operator for a single particle in some dimension (say x), and let be P be the momentum operator in that same dimension. (So Q and P are (partially) defined on $L^2(\mathbb{R}^3)$; and I shall not go into detail about the partialness of the domains of definition, which are adequately discussed by Muller and Seevinck.)

[7] Muller and Seevinck (2009, pp. 185-6) entertain adding a third requirement, to the effect that discernment by a relation is 'authentic' only if it is irreducible to monadic properties, and discernment by a monadic property is 'authentic' only if it is irreducible to relations. They reject this extra requirement, as do I; but my reasons are different. For me, physical meaning (embodied in (Req1)) is all one could, and should, reasonably ask for—so long as that is taken to entail the requirement that IP is satisfied; cf. my commentary of Muller and Saunders' proof in Section 6.3.3.

Now define $Q^{(1)} := Q \otimes \mathbb{1}$ and $Q^{(2)} := \mathbb{1} \otimes Q$, and $P^{(1)} := P \otimes \mathbb{1}$ and $P^{(2)} := \mathbb{1} \otimes P$, where $\mathbb{1}$ is the identity on $L^2(\mathbb{R}^3)$.

We may now define a relation C as follows:

$$C(x, y) \quad \text{iff} \quad [P^{(x)}, Q^{(y)}]\rho = c\rho, \text{ for some } c \neq 0 \ , \tag{6.6}$$

where ρ is the density operator representing the state of the assembly. Now for every state we have $C(1, 1)$ and $C(2, 2)$, since $[P^{(1)}, Q^{(1)}] = [P^{(2)}, Q^{(2)}] = -i\hbar\mathbb{1} \otimes \mathbb{1}$. And we also have $\neg C(1, 2)$ and $\neg C(2, 1)$, since $[P^{(1)}, Q^{(2)}] = [P^{(2)}, Q^{(1)}] = 0$. Thus C weakly discerns particles 1 and 2. This discernment is categorical, since C holds or not categorically, i.e. without probabilistic assumptions. And the discernment is physical, since C meets (Req1) and (Req2). $\square$

Commentary. First of all I note that the condition in (SMS2), that each particles' state space be $L^2(\mathbb{R}^3)$, I grant, of course, since I take it to be a compulsory requirement that particles be located (cf. Sections 5.1.2 and 5.2.1). Second: since the discernment is categorical, it is no restriction that the full (i.e. un-symmetrized) Hilbert space is used in the proof: the proof carries over for all restrictions to symmetry sectors.

As in Section 6.3.3, again I take no issue with the claim that the discernment is weak or that it is categorical, but I deny that it is physical. The reason is the same as for Muller and Saunders (2008): namely, the proof uses unphysical quantities. (Thus I deny that (Req1) is satisfied.) Again we demand not just that the discerning relation be symmetric, but also that it be defined using only physical—*a fortiori*, only symmetric—quantities. And $Q^{(x)}$ and $P^{(x)}$, despite their tantalising intuitive interpretation, do not count as physical quantities.

I now turn to Muller and Seevinck's second main Theorem:

(SMS3) In an assembly with a finite-dimensional Hilbert space $\bigotimes^N \mathbb{C}^{2s+1}$, where $s \in \{\frac{1}{2}, 1, \frac{3}{2}, \ldots\}$ and the associated algebra of quantities $\mathcal{B}(\bigotimes^N \mathbb{C}^{2s+1})$, any two particles are categorically, weakly, physically discernible using only their spin degrees of freedom.

Reconstruction of proof (cf. Muller and Seevinck (2009, p. 193-7)): Again I restrict attention to the case of two particles ($N = 2$). Let $\mathbf{S} = \sigma_x\mathbf{i} + \sigma_y\mathbf{j} + \sigma_k\mathbf{k}$

be the quantity representing a single particle's spin (so $\mathbf{S}$ acts on $\mathbb{C}^{2s+1}$). Then we define $\mathbf{S}_1 := \mathbf{S} \otimes \mathbb{1}$ and $\mathbf{S}_2 := \mathbb{1} \otimes \mathbf{S}$, and the relation T as follows:

$$T(x, y) \quad \text{iff} \quad \text{for all } \rho \in \mathcal{D}(\mathbb{C}^{2s+1} \otimes \mathbb{C}^{2s+1}),\ |(\mathbf{S}_x + \mathbf{S}_y)|^2 \rho = 4s(s+1)\hbar^2 \rho. \tag{6.7}$$

I note that $|\mathbf{S}|^2 = s(s+1)\hbar^2 \mathbb{1}$; this entails that $|2\mathbf{S}_1|^2 = |2\mathbf{S}_2|^2 = 4s(s+1)\hbar^2 \mathbb{1} \otimes \mathbb{1}$; so $T(1,1)$ and $T(2,2)$ both hold. Meanwhile, $|(\mathbf{S}_1 + \mathbf{S}_2)|^2 = |\mathbf{S}|^2 \otimes \mathbb{1} + \mathbb{1} \otimes |\mathbf{S}|^2 + 2\mathbf{S} \otimes \mathbf{S} = 2s(s+1)\hbar^2 \mathbb{1} \otimes \mathbb{1} + 2\mathbf{S} \otimes \mathbf{S}$. But the eigenvalues of $|(\mathbf{S}_1 + \mathbf{S}_2)|^2$ never exceed $(2s)(2s+1)\hbar^2 < 4s(s+1)\hbar^2$, so $\neg T(1,2)$ and $\neg T(2,1)$ both hold. This discernment is clearly weak. It is categorical, since it relies on no probabilistic assumptions, and it is physical, since T satisfies (Req1) and (Req2). $\Box$

Commentary. I note that, in order to put the physical significance of T on firmer ground, Muller and Seevinck extend the EPR reality condition (cf. footnote 6) to a necessary and sufficient condition, which they call the 'strong property postulate'. According to this postulate, the assembly possesses the property corresponding to the quantity's Q having value q *if and only if* the assembly's state is an eigenstate of the self-adjoint operator Q, with eigenvalue q. This strengthening is required to establish that the assembly does not possess combined total spin $\sqrt{4s(s+1)}\hbar$ when it is not in an eigenstate of the total spin operator.

Freedom from this stronger reality condition can be bought at the price of a concession to settle for probabilistic rather than categorical discernment. For we may define the new relation T':

$$T'(x, y) \quad \text{iff} \quad Tr\left(\rho |(\mathbf{S}_x + \mathbf{S}_y)|^2\right) = 4s(s+1)\hbar^2. \tag{6.8}$$

It is clear that T' discerns iff the "de-modalized" version of T discerns. But the definition of T' involves a commitment to the Born rule, so T''s discernment is probabilistic. This trade-off between the strong reality condition and the Born rule will also be a feature of my proposed discerning relations in the following Section.

My previous objection, which I levelled against (SMS1) and (SMS2), *appears* to be valid here too. For, even though $|(\mathbf{S}_1 + \mathbf{S}_2)|^2$ and $|2\mathbf{S}_1|^2 = |2\mathbf{S}_2|^2$ are symmetric, once again their building blocks ($\mathbf{S}_1$ and $\mathbf{S}_2$) are not, and (it may be argued) it is only when defined in terms of these components that T is not a gerrymandered relation.

However, our usual objection does *not* hold in this case. On the contrary, it seems reasonable to take $T(x, y)$ as a natural physical relation, even though its explicit mathematical form depends on whether $x = y$ or $x \neq y$. To see this, it should be enough that T can be parsed in English as the relation: 'the combined total spin of x and y has the magnitude $\sqrt{4s(s+1)}\hbar$ in all states'. Combined total spin is a symmetric quantity, and it has obvious physical significance. Therefore I do not take issue with the discerning relation being physical.

I have two further objections in this case: one mild, the other more serious. The mild problem is that the relation T is different in a significant way from the previous relations R_t and C. While R_t and C both applied to a given state of the assembly, the definition of T involves quantification over all states of the assembly. It is therefore a *modal* relation. But appeal to modal relations in this context is problematic, since it threatens to trivialise the search for a discerning relation for every state. It would turn out that PII is necessarily true if it is possibly true: a result that is at best controversial.

(Note, incidentally, that we cannot quite criticise the use of modal relations on the grounds that it assumes haecceitism. The natural thing to do for a factorist is to use the Hilbert space labels to cross-identify, and this seems to have a whiff of haecceitism about it. However, the factorist strategy need not commit one to haecceitism, since the quantification over states may be restricted to the (anti-)symmetric sectors, in which all states are permutation-invariant.)

This mild problem is easily addressed. We simply drop the quantification over states in the definition of T. If we do this, then the (unquantified) right-hand side of the definition (6.7) is still satisfied iff $x = y$, for all states ρ. We thereby drop the modal involvement. Thus we define a new relation, to be parsed as 'the combined total spin of x and y has the magnitude $\sqrt{4s(s+1)}\hbar$'. The discernment remains categorical, since no probabilistic assumptions have been made.

The serious problem is that (SMS3) is clearly only applicable to assemblies whose constituent particles have non-zero spin. This might seem to be only a mild omission, since as a matter of fact no zero-spin particles actually exist (except those that are composed of particles with non-zero spin, and might therefore be discerned by their internal structure). However, it would be nice to establish the discernibility of quantum particles for *all* values of spin, whether or not they happen to be realised.

To sum up: the same problem beleaguers the first two results (SMS1) and (SMS2), which aim to demonstrate the discernibility of (respectively) fermions and any particles with spatial degrees of freedom. The problem is that they both appeal to quantities which, in virtue of contravening IP, are non-physical. The third result, (SMS3), avoids this problem (modulo dropping some unnecessary modal involvement). However, it does not apply to particles with zero-spin. I now turn to my proposal for discerning any species of particle, for any value of spin.

6.3.5 A better way to discern factorism's particles

Muller and Saunders Theorem 3 (pp. 539-40) contains the germ of a better way to secure discernment; i.e. a way free of the criticisms discussed in Sections 6.3.3 and 6.3.4. This Section develops the germ. I proceed in stages. First I outline the basic idea, and propose a relation which weakly and physically discerns two particles in any two-particle assembly, using statistical variance. Then I investigate discernment for heterodox state spaces, in which particles may have definite position, and give a relation that will weakly and physically discern there too. Finally, I propose a relation that weakly and physically discerns any two particles in an assembly of *any* number of particles.

Stage A: The basic idea.

My basic idea is that particles may be discerned by taking advantage of anti-correlations between single-particle states. In the case of fermions, this is 'easy' because of Pauli exclusion: in *any* basis the occupation number for any single-particle state never exceeds one. In the case of the other particles, it is more tricky, due to the fact that states for non-fermionic particles may have as terms product states with equal factors. In these states, two or more particles are fully correlated, so there does not seem to be any quantum property *or* relation which would discern them. The solution is to change the basis to one in which anti-correlations appear with non-zero amplitude; the quantity associated with this new basis can then form the basis of a discerning relation.

Thus my strategy is discernment through anti-correlations, and the finding of anti-correlations through dispersion. For any state in which two particles are fully correlated, there will be dispersion in *some other* basis; in particular, the dispersion will involve branches with non-zero-amplitudes in which the particles are anti-correlated.

Stage B: The variance operator.

For simplicity, I focus exclusively on the two particle case. We may take the assembly Hilbert space to be $L^2(\mathbb{R}^3) \otimes L^2(\mathbb{R}^3)$, but my results still carry over if we restrict to a symmetry sector, or add additional (e.g. spin) degrees of freedom. Anti-correlations between single-particle states in an eigenbasis for some single-particle quantity A may be indicated by means of the following 'standard deviation' operator:

$$\Delta_A := \frac{1}{2}\left(A \otimes \mathbb{1} - \mathbb{1} \otimes A\right) \ . \tag{6.9}$$

Actually, I will use its square Δ_A^2, the 'variance' operator, which, like Δ_A, is self-adjoint (since A is). Unlike Δ_A, Δ_A^2 *is* a symmetric operator, so it is in line with the Indistinguishability Postulate (IP), and is therefore eligible to represent a physical quantity.

I also introduce the symmetrized quantity $\overline{A}$, which may be viewed as a mean of A taken over the two particles:

$$\overline{A} := \frac{1}{2}\left(A \otimes \mathbb{1} + \mathbb{1} \otimes A\right) \ . \tag{6.10}$$

Note that the over-line does *not* indicate an expectation value: $\overline{A}$ is an operator.

By similarly defining $\overline{A^2} = \frac{1}{2}\left(A^2 \otimes \mathbb{1} + \mathbb{1} \otimes A^2\right)$ we can express the variance operator more suggestively:

$$\begin{aligned}
\Delta_A^2 &= \frac{1}{4}\left(A\otimes\mathbb{1} - \mathbb{1}\otimes A\right)^2 \\
&= \frac{1}{4}\left(A^2\otimes\mathbb{1} + \mathbb{1}\otimes A^2 - 2A\otimes A\right) \\
&= \frac{1}{2}\left(\overline{A^2} - A\otimes A\right)
\end{aligned} \tag{6.11}$$

and

$$\begin{aligned}
\Delta_A^2 &= \frac{1}{4}\left(A^2\otimes\mathbb{1} + \mathbb{1}\otimes A^2 - 2A\otimes A\right) \\
&= \frac{1}{2}\left(A^2\otimes\mathbb{1} + \mathbb{1}\otimes A^2\right) - \frac{1}{4}\left(A^2\otimes\mathbb{1} + 2A\otimes A + \mathbb{1}\otimes A^2\right) \\
&= \overline{A^2} - \overline{A}^2 \ .
\end{aligned} \tag{6.12}$$

It is the latter equation (6.12) which justifies the term 'variance' for Δ_A^2. But note again that it is not the (c-numbered) statistical variance of A over a given wavefunction; it is the variance of the operator A over the two particles: Δ_A^2 is itself still an operator. The former equation (6.11) makes it most clear that Δ_A^2 measures the anti-correlation between each of the two particles' A-eigenstates. In particular, for any state all of whose terms are product states with equal factors in the A-basis:

$$|\Psi\rangle = \sum_k c_k |\phi_k\rangle \otimes |\phi_k\rangle \ , \tag{6.13}$$

where

$$A|\phi_k\rangle = a_k|\phi_k\rangle \ , \tag{6.14}$$

the variance has eigenvalue zero:

$$\begin{aligned}
\Delta_A^2|\Psi\rangle &= \frac{1}{4}\sum_k c_k\left(A^2\otimes\mathbb{1} + \mathbb{1}\otimes A^2 - 2A\otimes A\right)|\phi_k\rangle\otimes|\phi_k\rangle \\
&= \frac{1}{4}\sum_k c_k\left(a_k^2 + a_k^2 - 2a_k^2\right)|\phi_k\rangle\otimes|\phi_k\rangle \\
&= 0 \ .
\end{aligned} \tag{6.15}$$

In general, however, a state with anti-correlations will not be an eigenstate of Δ_A^2. For a generic state-vector

$$|\Phi\rangle = \sum_{ij} c_{ij} |\phi_i\rangle \otimes |\phi_j\rangle \tag{6.16}$$

we have

$$\begin{aligned}
\Delta_A^2|\Phi\rangle &= \frac{1}{4}\sum_{ij} c_{ij} \left(A^2 \otimes \mathbb{1} + \mathbb{1} \otimes A^2 - 2A \otimes A\right) |\phi_i\rangle \otimes |\phi_j\rangle \\
&= \frac{1}{4}\sum_{ij} c_{ij} \left(a_i^2 + a_j^2 - 2a_i a_j\right) |\phi_i\rangle \otimes |\phi_j\rangle \\
&= \frac{1}{4}\sum_{ij} c_{ij} \left(a_i - a_j\right)^2 |\phi_i\rangle \otimes |\phi_j\rangle
\end{aligned} \tag{6.17}$$

so that

$$\begin{aligned}
\left\langle \Delta_A^2 \right\rangle &:= \langle \Phi|\Delta_A^2|\Phi\rangle \\
&= \frac{1}{4}\sum_{ijkl} c_{kl}^* c_{ij} \left(a_i - a_j\right)^2 \langle \phi_k|\phi_i\rangle \langle \phi_l|\phi_j\rangle \\
&= \frac{1}{4}\sum_{ij} |c_{ij}|^2 \left(a_i - a_j\right)^2 \quad .
\end{aligned} \tag{6.18}$$

If we assume that A is non-degenerate ($a_i = a_j$ implies $i = j$), then it is clear from (6.18) that there is a positive contribution to the value of $\left\langle \Delta_A^2 \right\rangle$ from *every* anti-correlation that has a non-zero amplitude.

Stage C: Variance provides a discerning relation.

If a two-particle state has anti-correlations in a single-particle quantity A, we can build a symmetric, irreflexive relation which discerns them. The main idea is: if the expectation of the variance operator is non-zero, then this can be expressed as a relation between the two particles which neither particle bears to itself.

Following Muller & Saunders (2008) and Muller & Seevinck (2009), we define the operators

$$A^{(1)} := A \otimes \mathbb{1} \ ; \qquad\qquad A^{(2)} := \mathbb{1} \otimes A \ . \tag{6.19}$$

These quantities, being non-symmetric, are unphysical, but they can be used to define physical quantities: note, for example, that $\Delta_A \equiv \frac{1}{2}\left(A^{(1)} - A^{(2)}\right)$ and $\overline{A} \equiv \frac{1}{2}\left(A^{(1)} + A^{(2)}\right)$. We then define the relation R as follows:

$$R(A, x, y) \quad \text{iff} \quad \frac{1}{4}\left(A^{(x)} - A^{(y)}\right)^2 \rho \neq 0\ . \tag{6.20}$$

In English: $R(A, x, y)$ holds for the state ρ if and only if ρ is not an eigenstate of the absolute difference between x's and y's operator A, with eigenvalue zero. Here the variable A ranges over single-particle quantities, x and y range over the particles 1 and 2, and t ranges over the reals. This definition implies that $R(A, 1, 2)$ iff $R(A, 2, 1)$, iff $\Delta_A^2 \rho \neq 0$. And $\neg R(A, 1, 1)$ and $\neg R(A, 2, 2)$. So $R(A, x, y)$ is symmetric and irreflexive for each A. If Δ_A^2 does not annihilate ρ, then we have $R(A, 1, 2)$ and $R(A, 2, 1)$; so in this case $R(A, x, y)$ weakly discerns particles 1 and 2. Moreover, the discernment is categorical.

The question remains whether this discernment is physical. I claim that it is, on the assumption of factorism, since the quantity $\frac{1}{4}\left(A^{(x)} - A^{(y)}\right)^2$, which is symmetric, can be understood as a measure of anti-correlations between x and y for the single-particle quantity A—i.e., a measure of difference between x's and y's values for A. Thus it is no surprise that $\frac{1}{4}\left(A^{(x)} - A^{(y)}\right)^2 = 0$ for $x = y$; for no object can take a value for any quantity that is different from itself. So long as the single-particle operator A has physical significance, so does $\frac{1}{4}\left(A^{(x)} - A^{(y)}\right)^2 = 0$. I emphasize that the physical meaning of $\frac{1}{4}\left(A^{(x)} - A^{(y)}\right)^2 = 0$ should not be thought of as depending on $A^{(x)}$'s having physical meaning.

There is an important analogy here with relative distance. The relative distance between particle x and particle y need not be thought of as *deriving its meaning from* the absolute positions of x and y, even though the mathematical formalism of our theory may indeed allow us to define the relative distance in terms of these absolute positions. We need not take these mathematical definitions as representative of any physical fact, since we are not forced to admit that an element of the theory's formalism which has a physical correlate also has physical correlates for all of its mathematical building blocks. This is because these mathematical building blocks may contain redundant structure which is not transmitted to all of their by-products. Such is the case of relative distance. And in fact, relative distance is more than an analogy: for (squared) relative distance is an *instance* of Δ^2_A, if we set $A = \mathbf{Q}$, the single-particle position operator.

Note that an additional assumption is required to transmit physical significance from $\frac{1}{4}\left(A^{(x)} - A^{(y)}\right)^2 = 0$ to $R(A, x, y)$: we need to assume Muller and Seevinck's 'strong property postulate'. Recall that this states that any physical quantity of the assembly takes a certain value if and only if the assembly is in the appropriate eigenstate for that physical quantity's corresponding operator. What is important here is the 'only if' component of the biconditional: this enables us to say that the difference in x's and y's values for A is non-zero just in case the assembly is not in the eigenstate with eigenvalue zero—including when the assembly is not in an eigenstate at all.

I summarise the foregoing discussion in the following Lemmas:

Lemma 1 For all two-particle assemblies, and all single-particle quantities A, the relation $R(A, x, y)$ has physical significance if A does, on the assumption of the strong property postulate.

Lemma 2 For each state ρ of an assembly of two particles, and each single-particle quantity A, the relation $R(A, x, y)$ discerns particles 1 and 2 weakly, categorically and physically if and only if $\Delta^2_A\rho \neq 0$, on the assumption of the strong property postulate.

Proofs: See above. $\square$

As with (SMS3), in the previous Section, we can instead forego the strong property postulate and instead take advantage of the Born rule, thereby settling for probabilistic discernment. We then define the relation R' as follows:

$$R'(A, x, y) \quad \text{iff} \quad \frac{1}{4}\mathrm{Tr}\left[\rho\left(A^{(x)} - A^{(y)}\right)^2\right] \neq 0 \ . \qquad (6.21)$$

Similar considerations to those above entail that $R'(A, 1, 2)$ iff $R'(A, 2, 1)$, iff $\langle\Delta^2_A\rangle \neq 0$. And $\neg R(A, 1, 1)$ and $\neg R(A, 2, 2)$. So $R(A, x, y)$ weakly discerns particles 1 and 2 just in case $\langle\Delta^2_A\rangle \neq 0$. Thus:

Lemma 3 For all single-particle quantities A, the relation $R'(A, x, y)$ has physical significance if A does, on the assumption of the Born rule.

Lemma 4 For each state ρ of the assembly, and each single-particle quantity A, the relation $R'(A, x, y)$ discerns particles 1 and 2 weakly, probabilistically and physically, if and only if $\langle \Delta_A^2 \rangle \neq 0$ for that state.

Proofs: See above. $\square$

Stage D: Discernment for all two-particle states.

So far we have seen that two particles in a state with non-zero variance in some single-particle quantity A—i.e. two particles which are anti-correlated in A—may be discerned. To guarantee discernment in all two-particle states it remains to be shown that, for *any* such state, there *will be* some single-particle quantity whose eigenbasis has anti-correlations. In fact I prove a stronger result: namely that there is some single-particle quantity which discerns the two particles in all states of the assembly. Moreover, this quantity is familiar: it is position; and since I require all particles to have a location (cf. Section 5.1.2), it will be a quantity that will always be available to discern.

Theorem 1 For each state ρ of an assembly of two particles, the relation $R(\mathbf{Q}, x, y)$ discerns particles 1 and 2 weakly, categorically and physically; where $\mathbf{Q}$ is the single-particle position operator; on the assumption of the strong property postulate.

Proof: We assume the strong property postulate. From Lemma 2, we know that $R(\mathbf{Q}, x, y)$ discerns particles 1 and 2 weakly, categorically and physically, in the state ρ if and only if $\Delta^2_{\mathbf{Q}}\rho \neq 0$. Let us first consider only pure states, and later generalise to all states.

Pure states. Since we are working in the position representation, we use wavefunctions rather than state-vectors or density operators. The most general form for a wavefunction for the assembly is

$$\Psi(\mathbf{x}, \mathbf{y}) = \sum_{ij} c_{ij} \phi_i(\mathbf{x}) \phi_j(\mathbf{y}) \ , \tag{6.22}$$

where the ϕ_i are an orthonormal basis for $L^2(\mathbb{R}^3)$. (We assume zero spin, but the proof is trivially extended for any non-zero value for spin.) Now

$$\begin{aligned}(\Delta^2_{\mathbf{Q}}\Psi)(\mathbf{x},\mathbf{y}) &= \sum_{ij} c_{ij}\left(x^2\phi_i(\mathbf{x})\phi_j(\mathbf{y}) + \phi_i(\mathbf{x})y^2\phi_j(\mathbf{y}) - 2\mathbf{x}\phi_i(\mathbf{x}).\mathbf{y}\phi_j(\mathbf{y})\right)\\ &= \left(\sum_{ij} c_{ij}\phi_i(\mathbf{x})\phi_j(\mathbf{y})\right)(\mathbf{x}-\mathbf{y})^2\\ &= \Psi(\mathbf{x},\mathbf{y})(\mathbf{x}-\mathbf{y})^2 \qquad (6.23)\end{aligned}$$

(cf. Equation (6.17)). This is the zero function only if $\Psi(\mathbf{x},\mathbf{y}) = 0$ whenever $\mathbf{x} \neq \mathbf{y}$. But then it cannot be represented in $L^2(\mathbb{R}^3) \otimes L^2(\mathbb{R}^3)$, since it is not a function. (We essentially appeal to the fact that no wavefunction is infinitely peaked at the diagonal points of the configuration space. The necessary Ψ can be written as a measure: $\Psi(\mathbf{x},\mathbf{y}) = f(\mathbf{x})\delta^{(3)}(\mathbf{x}-\mathbf{y})$, for some function $f \in L^2(\mathbb{R}^3)$. I return to this point in Theorem 3, below.) Therefore we conclude that $(\Delta^2_{\mathbf{Q}}\Psi)(\mathbf{x},\mathbf{y}) \neq 0$. It follows that $\Delta^2_{\mathbf{Q}}|\Psi\rangle\langle\Psi| \neq 0$.

Mixed states. We extend to density operators by taking convex combinations of (not necessarily othogonal) projectors. We have that

$$\Delta^2_{\mathbf{Q}}\left(\sum_i p_i|\Psi_i\rangle\langle\Psi_i|\right) = \sum_i p_i\Delta^2_{\mathbf{Q}}|\Psi_i\rangle\langle\Psi_i| \neq 0 \qquad (6.24)$$

since both the p_i and the spectrum of $\Delta^2_{\mathbf{Q}}$ are positive.

From Lemma 2, we conclude that $R(\mathbf{Q},x,y)$ discerns particles 1 and 2 weakly. The discernment is categorical since we made no probabilistic assumptions. Finally, the discernment is physical, as follows from Lemma 1, the strong property postulate, and the fact that $\mathbf{Q}$ is physical. □

We can now also prove

Theorem 2 For each state ρ of an assembly of two particles, the relation $R'(\mathbf{Q},x,y)$ discerns particles 1 and 2 weakly, probabilistically and physically; where $\mathbf{Q}$ is the single-particle position operator; on the assumption of the Born rule.

Proof. We assume the Born rule. Then for any state ρ we have (cf. Equations

6.23, 6.24):

$$
\begin{aligned}
\langle \Delta^2_{\mathbf{Q}} \rangle &= \mathrm{Tr}\left(\rho \Delta^2_{\mathbf{Q}}\right) \\
&= \sum_i p_i \langle \Psi_i | \Delta^2_{\mathbf{Q}} | \Psi_i \rangle \\
&= \sum_i p_i \int \mathrm{d}^3\mathbf{x} \int \mathrm{d}^3\mathbf{y} \; |\Psi_i(\mathbf{x}, \mathbf{y})|^2 (\mathbf{x} - \mathbf{y})^2 \; ,
\end{aligned}
\tag{6.25}
$$

which is positive, i.e. non-zero (cf. Equation (6.18)). From Lemma 4, $R'(\mathbf{Q}, x, y)$ therefore discerns weakly. The discernment is probabilistic, since we assumed the Born rule. Finally, the discernment is physical, as follows from Lemma 3, the Born rule, and the fact that $\mathbf{Q}$ is physical. □

It may be objected against the proofs of our two Theorems that we rely too heavily on a technical feature of the assembly's Hilbert space, namely that it contains no states which exhibit no spread in $(\mathbf{x} - \mathbf{y})^2$. Effectively, unfavourable cases for discernment have been ruled out of the assembly's Hilbert space *a priori*. But this objection is easily dealt with.

Theorem 3 If we permit two-particle states to be represented by measures as well as by functions, then for all such states, *either* $R(\mathbf{Q}, x, y)$ *or* $R(\mathbf{P}, x, y)$ discerns particles 1 and 2 weakly, categorically and physically; where $\mathbf{Q}$ is the single-particle position operator, $\mathbf{P}$ is the single-particle momentum operator; on the assumption of the strong projection postulate.

Proof. The guiding idea is that any state will exhibit spread in either relative position or relative momentum, so no state is annihilated by both $\Delta^2_{\mathbf{Q}}$ and $\Delta^2_{\mathbf{P}}$.

We now allow measures, as well as functions, to represent states of the assembly. Recall from the proof of Theorem 1 that $(\Delta^2_{\mathbf{Q}}\Psi)(\mathbf{x}, \mathbf{y}) = 0$ only if $\Psi(\mathbf{x}, \mathbf{y}) = 0$ whenever $\mathbf{x} \neq \mathbf{y}$. In this case $\Psi(\mathbf{x}, \mathbf{y}) = f(\mathbf{x})\delta^{(3)}(\mathbf{x} - \mathbf{y})$, for some measure $f(\mathbf{x})$. Note at this point that the two particles *cannot be fermions*, since $\Psi(\mathbf{x}, \mathbf{y}) - \Psi(\mathbf{y}, \mathbf{x}) = 0$. We now move to the momentum basis by performing a Fourier transform on Ψ:

$$\begin{aligned}
\overline{\Psi}(\mathbf{k},\mathbf{l}) &= \int \mathrm{d}^3\mathbf{x} \int \mathrm{d}^3\mathbf{y}\ \Psi(\mathbf{x},\mathbf{y}) e^{-i\mathbf{k}.\mathbf{x}} e^{-i\mathbf{l}.\mathbf{y}} \\
&= \int \mathrm{d}^3\mathbf{x} \int \mathrm{d}^3\mathbf{y}\ f(\mathbf{x})\delta^{(3)}(\mathbf{x}-\mathbf{y}) e^{-i(\mathbf{k}.\mathbf{x}+\mathbf{l}.\mathbf{y})} \\
&= \int \mathrm{d}^3\mathbf{x}\ f(\mathbf{x}) e^{-i(\mathbf{k}+\mathbf{l}).\mathbf{x}} \\
&= \overline{f}(\mathbf{k}+\mathbf{l}). \qquad (6.26)
\end{aligned}$$

This yields

$$(\Delta^2_{\mathbf{P}}\overline{\Psi})(\mathbf{k},\mathbf{l}) = (\mathbf{k}-\mathbf{l})^2 \overline{f}(\mathbf{k}+\mathbf{l}) \ , \qquad (6.27)$$

which is the zero function only if $\overline{f}(\mathbf{k}+\mathbf{l}) = 0$ whenever $\mathbf{k} \neq \mathbf{l}$. But we can only satisfy this requirement if $\overline{f}$ is the zero function. But in that case $\Psi(\mathbf{x},\mathbf{y})$ is the zero function, and so does not represent a state. So if $(\Delta^2_{\mathbf{Q}}\Psi)(\mathbf{x},\mathbf{y})$ is the zero function, then $(\Delta^2_{\mathbf{P}}\overline{\Psi})(\mathbf{k},\mathbf{l})$ can't be. This result is easily extended to mixed states.

With this result and Lemma 2 we conclude that either $R(\mathbf{Q},x,y)$ or $R(\mathbf{P},x,y)$ (or both) discerns particles 1 and 2 weakly. The discernment is categorical, since we made no probabilistic assumptions. Finally, the discernment is physical, given Lemma 1, the strong property postulate, and the fact that both $\mathbf{Q}$ and $\mathbf{P}$ are physical. □

It only remains to state our

Theorem 4 If we permit two-particle states to be represented by measures as well as by functions, then for all such states, *either* $R(\mathbf{Q},x,y)$ *or* $R(\mathbf{P},x,y)$ discerns particles 1 and 2 weakly, probabilistically and physically; where $\mathbf{Q}$ is the single-particle position operator, $\mathbf{P}$ is the single-particle momentum operator; on the assumption of the Born rule.

Proof: Left to the reader. □

So we have established the weak discernibility of indistinguishable particles in any two-particle assembly. But my results are restricted to the two particle case. Therefore, I now turn to the many-particle case, and present Theorems for assemblies of any number of particles.

Stage E: Discernment for all many-particle states.

I begin by defining a generalized N-particle variance operator for each single-particle quantity A, for any $N \geqslant 2$. For any single-particle quantity A we define

$$\begin{aligned}
\left(\Delta_A^{(N)}\right)^2 &:= \overline{A^2} - \overline{A}^2 \\
&= \frac{1}{N}\sum_i^N \mathbb{1} \otimes \cdots A_i^2 \otimes \cdots \mathbb{1} - \left(\frac{1}{N}\sum_i^N \mathbb{1} \otimes \cdots A_i \otimes \cdots \mathbb{1}\right)^2 \\
&= \frac{1}{N}\sum_i^N \mathbb{1} \otimes \cdots A_i^2 \otimes \cdots \mathbb{1} \\
&\quad -\frac{1}{N^2}\left(\sum_i^N \mathbb{1} \otimes \cdots A_i^2 \otimes \cdots \mathbb{1}\right. \\
&\qquad\qquad \left. + 2\sum_{i<j}^N \mathbb{1} \otimes \cdots A_i \otimes \cdots A_j \otimes \cdots \mathbb{1}\right) \\
&= \frac{1}{N^2}\left((N-1)\sum_i^N \mathbb{1} \otimes \cdots A_i^2 \otimes \cdots \mathbb{1}\right. \\
&\qquad\qquad \left. - 2\sum_{i<j}^N \mathbb{1} \otimes \cdots A_i \otimes \cdots A_j \otimes \cdots \mathbb{1}\right) \\
&= \frac{1}{N^2}\sum_{i<j}^N \left(\mathbb{1} \otimes \cdots A_i \otimes \cdots \mathbb{1} \;-\; \mathbb{1} \otimes \cdots A_j \otimes \cdots \mathbb{1}\right)^2 \; . \qquad (6.28)
\end{aligned}$$

Note that $\left(\Delta_A^{(2)}\right)^2 = \Delta_A^2$; cf. Equations (6.11) and (6.12).

Again, my strategy is to discern by setting $A = \mathbf{Q}$, the single-particle position operator. If we act on any wavefunction Ψ in $\bigotimes^N L^2(\mathbb{R}^3)$ with $\left(\Delta_{\mathbf{Q}}^{(N)}\right)^2$ we obtain, using (6.28),

$$\left(\left(\Delta_{\mathbf{Q}}^{(N)}\right)^2 \Psi\right)(\mathbf{x}_1, \ldots \mathbf{x}_N) = \frac{1}{N^2}\left(\sum_{i<j}^N (\mathbf{x}_i - \mathbf{x}_j)^2\right)\Psi(\mathbf{x}_1, \ldots \mathbf{x}_N) \; . \qquad (6.29)$$

Now it is clear from Equation (6.29) that we cannot proceed in the general case exactly as we did in the two-particle case. That is: we cannot discern two particles—*a* and b, say—by relying on the variance operator's annihilating the wavefunction. For the vanishing of the right-hand side of Equation (6.29) is not a necessary condition for a and b's having vanishing relative distance: this relative distance may be zero, and yet there may still be non-zero contributions from the other particles.

However, we need only make mild adjustments to our previous strategy. My idea is to look at regions of the configuration space for which $(\mathbf{x}_i - \mathbf{x}_j)^2$ is constant, except for when i or j equals a or b. We then independently vary $\mathbf{x}_a$ and $\mathbf{x}_b$. If the wavefunction is non-zero for $\mathbf{x}_a \neq \mathbf{x}_b$, then we find variation in the right-hand side of Equation (6.29) which can only be attributed to a and b's having non-vanishing relative distance, i.e. to their being discernible.

We first define a new dyadic relation between particles:

$$D^{(N)}(x,y) \quad \text{iff}$$
$$\left(\left(\Delta_{\mathbf{Q}}^{(N)}\right)^2 \Psi\right)(\mathbf{x}_1, \ldots \mathbf{x}_N) \neq \frac{1}{N^2} \sum_{\substack{i<j; \\ \langle i,j\rangle \neq \langle x,y\rangle; \\ \langle i,j\rangle \neq \langle y,x\rangle}}^{N} (\mathbf{x}_i - \mathbf{x}_j)^2 \, \Psi(\mathbf{x}_1, \ldots \mathbf{x}_N) \,. \tag{6.30}$$

Note that $D^{(2)}(x,y)$ iff $R(\mathbf{Q},x,y)$; so $D^{(2)}$ is a physical relation. Is $D^{(N)}$ a physical relation for *all* N? First we note that the N-particle variance operator for position, $\left(\Delta_{\mathbf{Q}}^{(N)}\right)^2$, is a physical quantity, as is evident from its definition (6.28). Now we need to make physical sense of the condition in the definition of $D^{(N)}$ (Equation (6.30)).

Recall that $R(A,x,y)$'s defining condition is to the effect that the wavefunction is not an eigenstate of the variance operator for some quantity (with eigenvalue zero); with the strong property postulate, this entails that the assembly does not have the corresponding physical property (namely, zero variance in that quantity). Therefore, there can be no doubt that R's defining condition is physically meaningful (so long as the strong property postulate is valid). However, in the case of $D^{(N)}$, the condition is *not* that Ψ not be an eigenstate; the condition is rather that Ψ not be sent to some specific function by the N-particle variance operator for position. The strong property postulate is therefore no help in giving $D^{(N)}$'s defining condition physical significance.

We regretfully settle for probabilistic discernment. $D^{(N)}$'s defining condition makes perfect physical sense if we assume the Born rule, since then the condition could be interpreted as the N-particle variance operator for position having an expectation value not equal to the value specified in the right-hand side of Equation (6.30). We can make this more explicit by defining another relation:

$$D'^{(N)}(x,y) \quad \text{iff}$$
$$\left\langle \left(\Delta_{\mathbf{Q}}^{(N)}\right)^2 \right\rangle \neq \frac{1}{N^2}\int \mathrm{d}^3\mathbf{x}_1 \cdots \int \mathrm{d}^3\mathbf{x}_N \sum_{\substack{i<j;\\ \langle i,j\rangle \neq \langle x,y\rangle;\\ \langle i,j\rangle \neq \langle y,x\rangle}}^{N} (\mathbf{x}_i - \mathbf{x}_j)^2 \, |\Psi(\mathbf{x}_1, \ldots \mathbf{x}_N)|^2 \,. \tag{6.31}$$

We may now prove

Thereom 5 For each state ρ of an assembly of N particles, the relation $D'^{(N)}(x,y)$ discerns any two distinct particles x and y weakly, probabilistically and physically, on the assumption of the Born rule.

Proof. We prove this only for pure states and zero spin; the extension to mixed states and non-zero spin will be obvious, given our proof of Theorem 1.

It is clear that $\neg D'^{(N)}(x,x)$ for all x, since when $x = y$ the right-hand side of Equation (6.31) corresponds to the definition of the left-hand side (cf. Equation (6.29)), and therefore must be equal to it. Thus $D'^{(N)}$ is irreflexive. To show that $D'^{(N)}$ discerns any two particles weakly, we need to prove that $D'^{(N)}(x,y)$ holds whenever $x \neq y$.

This we do by *reductio*: assume that there are two particles a and b ($a \neq b$) for which $\neg D'^{(N)}(a,b)$. Then we must have, by subtracting the right-hand side of Equation (6.31) from its left-hand side:

$$\frac{1}{N^2}\int \mathrm{d}^3\mathbf{x}_1 \cdots \int \mathrm{d}^3\mathbf{x}_N \, (\mathbf{x}_a - \mathbf{x}_b)^2 \, |\Psi(\mathbf{x}_1, \ldots \mathbf{x}_N)|^2 = 0 \,. \tag{6.32}$$

This holds only if $\Psi(\mathbf{x}_1, \ldots \mathbf{x}_N) = 0$ whenever $\mathbf{x}_a \neq \mathbf{x}_b$. So $\Psi(\mathbf{x}_1, \ldots \mathbf{x}_N) =$

$f(\mathbf{x}_1, \ldots \mathbf{x}_a, \ldots \mathbf{x}_{b-1}, \mathbf{x}_{b+1}, \ldots \mathbf{x}_N)\delta^{(3)}(\mathbf{x}_a - \mathbf{x}_b)$, where f is some $3(N-1)$-place function. But then Ψ is not a function, so it is not a state in $\bigotimes^N L^2(\mathbb{R}^3)$. Thus $D'^{(N)}(a,b)$, and $D'^{(N)}$ weakly discerns any two particles in the assembly.

The definition of $D'^{(N)}$ involves taking an expectation value, so it discerns probabilistically. Finally, the foregoing discussion establishes that $D'^{(N)}$ is a physical relation. □

Finally, for completeness, I present our final Theorem of this Section:

Theorem 6 If we permit states of an assembly of N particles to be represented by measures as well as by functions, then for all such states, *either* $D'(x,y)$ *or* its momentum analogue discerns any two particles x and y weakly, probabilistically and physically, on the assumption of the Born rule.

Proof: Left to the reader. The method is to carry over to the N-particle case the way in which proofs of Theorems 3 and 4 developed Theorems 1 and 2. □

Let me now sum up the results of Section 6.3. On the assumption of factorism, a strong version of the principle of the identity of indiscernibles (PII) fails for all particles. This version of PII permits discernment of two objects only by monadic properties, or relations to other objects not in the pair. However, a weaker (and still non-trivial) version of PII is available, which allows particles to be discerned *weakly*, i.e. by symmetric and irreflexive relations. This version of PII holds for all pairs of particles: fermions, bosons and paraparticles.

Previous attempts to establish this general result by Muller and Saunders (2008) and Muller and Seevinck (2009) have been seen to fail, due to the use of mathematical machinery which could be given no physical interpretation. However, physical relations exist which secure the result. They derive their physical significance from the single-particle position operator (and, if needed, the single-particle momentum operator). In the case of two-particle assemblies, this discernment may be categorical—that is, independent of all probabilistic assumptions—, but we must assume the strong property postulate. In the general case, for any number of particles, the discernment is probabilistic, in that we must assume the Born rule.

I now turn to my criticism of factorism, the interpretative doctrine on which all of the foregoing study of discernibility rests. We will see later (in Chapters 8 and 9) how far this Section's results may be preserved in an anti-factorist setting.

6.4 The defects of factorism

In this Section I argue that factorism is false, since its target concept of particle rests on a false interpretative assumption: namely, that, under IP, factor Hilbert space labels have physical correlates. My argument therefore seeks to undermine this assumption.

I proceed in three stages. First (Section 6.4.1), I argue that the celebrated absolute indiscernibility of factorist particles is, on its own, grounds for suspicion that the representative elements of the quantum formalism have been misidentified. Second (Section 6.4.2), I show that, under factorism, we fail to recover classical particles in the classical limit; thereby contravening our desideratum of inter-theoretic applicability (cf. Section 5.1.5). Third (Section 6.4.3), I combine the first two stages to argue that the factorist's particles are not physical, but rather a mere artefact of the formalism; thereby contravening our (compulsory!) desideratum that particles be physical (cf. Section 5.1.1). I strengthen the argument by appeal to an analogous error in statistics: namely, the reification of statistical constructs, such as the average taxpayer. I conclude the Section, and the Chapter, with a possible escape route for haecceitistic factorists (Section 6.4.4).

All three stages of my argument appeal to the same phenomenon exhibited by factorist particles, just explored in Section 6.3:

> In any state of the assembly, every particle is in the same quantum state as any other of the same species: viz. an improper mixture of *every* state instantiated by particles of that species.[8] In the philosophical jargon: in any state, the particles of a species are absolutely indiscernible from each other: they are therefore *non-individuals* (cf. Section 2.3's

[8]The only exception, of course, is bosonic states in which *all* of the bosons of the same species take the same single-particle state, and so the assembly state is separable.

Interlude).

6.4.1 Losing individuality

The strangeness of the result just mentioned is often not fully acknowledged.[9] Here I argue that it is so strange as to be unbelievable. The considerations I adduce split into philosophical and physical, but they are complementary to one another.

From the philosophical angle: the result that all particles of the same mass, spin and charge are absolutely indiscernible one from another, and are therefore all non-individuals,[10] causes an enormous problem for reference where there should be none. For example, we cannot, under factorism, refer to *the* electron that is in state $|\phi\rangle$, *even if* there is a probability of one that the state $|\phi\rangle$ is instantiated (i.e. even if the assembly is in an eigenstate of the number operator associated with $|\phi\rangle$ with eigenvalue 1). According to factorism, there is *no* electron that is in state $|\phi\rangle$ with probability 1; rather *every* electron has the same non-zero probability of being in that state. This blocks us from engaging in the most intuitive way of speaking when, for example, chemists describe the progression of atomic number in the periodic table as corresponding to the addition of an electron (and nucleons) in a stationary state of definite energy.

From a physical angle: the factorist relies on a decomposition of the assembly's Hilbert space that is not natural, once the indistinguishability postulate (IP) is imposed. (We will return to the idea of 'natural decompositions in Section 7.2.) The factorist's decomposition relies on imagining the Hilbert space of the assembly—which is a *subspace* of a tensor product space, being the range of a projection onto a single symmetry type—as still embedded in the original tensor product space (Earman (ms.)). Yet the restriction on quantities effected by IP makes inaccessible the other regions of this Hilbert space—so it is improper to consider the assembly's state space to be given by the original, unsymmetrized tensor product. (I consider a reply to this objection in Section 6.4.4.)

[9]Notable exceptions are Pooley (2006, p. 116), Dieks and Lubberdink (2008), Earman (ms.) and James Ladyman and Thomas Bigaj (both from personal correspondence).

[10]Here I set aside haecceitism, which *will* afford each factorist particle individuality, albeit an empirically inaccessible one. I return to haecceitism in Section 6.4.4.

Thus there is something suspicious about the factorist's decomposition of the assembly's Hilbert space. Perhaps the most adverse result of this is factorism's account of the classical and QFT limit; a topic to which I now turn.

6.4.2 Losing limits

One consequence of the non-individuality of factorist particles is that they cannot become classical particles in an appropriate limit. This phenomenon is well discussed by Dieks and Lubberdink (2011), but to summarise: In the classical limit, factorist particles do not acquire the approximate trajectories we associate with classicality, since the former must remain in statistically mixed states all the way to the classical limit. Factorist particles cannot tend, in any limit, to become distinguished one from another in space—like classical particles—by zero or at least negligible overlap, since each factorist particles always possesses the "entire spatial profile" of the assembly.

A similar point can be made about quantum field theory: factorist particles do not tend to the behaviour of QFT-quanta if we consider the limit in which the total particle number in conserved. This is because QFT-quanta—which are associated with creation and annihiliation operators $a^\dagger(\phi), a(\phi)$, for some state ϕ in the single-particle Hilbert space—always occupy pure states. Indeed: it may be shown (though I will not here) that QFT-quanta behave just like classical particles in an appropriate classical limit for QFT. Thus the factorist particles appear to be the odd ones out. (This vindicates our claim, at the start of Chapter 5, that the "local" concept of particle as applied in quantum mechanics conflicts with the "general" concept, which is informed by both classical mechanics and QFT.)

Furthermore, to anticipate the results of the next Chapter: the unnaturalness of the factorist's envisaged decomposition of the assembly's Hilbert space means that traditional accounts of entanglement are misleading when IP is in place. For the traditional definition of entanglement—namely, non-separability—makes sense only if the subsystems are taken to correspond to factor Hilbert space labels. We will explore these issues much more deeply soon, but I emphasise here that a further charge against factorism may be that it fails to distinguish "genuine" entanglement—i.e. entanglement which may manifest in the algebra of admissible quantities under IP—from mere non-separability, which is nothing but a *formal* feature of the way states are represented in quantum mechanics.[11]

Thus factorism contravenes our desideratum that the target concept of particle for quantum mechanics be applicable across its neighbouring theories. Yet these neighbouring theories themselves—i.e., classical mechanics and QFT—seem to agree better with each other than either does with quantum mechanics under a factorist understanding. Clearly, something has gone wrong.

6.4.3 Analogy with the average taxpayer

The factorist's error may be better understood by analogy with a well known fallacy: the reification of statistical constructs, such as the average taxpayer. Both of these commit what Whitehead (1925) called the 'fallacy of misplaced concreteness'. In my jargon: both mistakenly attributive representational power to elements of the theory's formalism that have no straightfoward physical correlate.[12]

For consider a factorist particle. Necessarily, its state is statistically mixed iff the assembly's state is *heterogeneous*—i.e. iff the assembly has non-zero probabilities for two or more distinct single-particle states being instantiated. Even in the classical limit, a factorist particle's state is a statistical mixture of several definite trajectories.

In this way, a factorist particle behaves much like the average taxpayer (for example). For any heterogeneity in the taxpaying population is necessarily correlated with a statistical spread in the corresponding attributes of the average taxpayer.[13] For example, if 10% of the taxpaying population are left-handed, then the average taxpayer is 10% left-handed. Attributing properties to an object named 'the

[11]Earman (ms.) calls this latter phenomenon "pixie dust entanglement".

[12]Dieks and Lubberdink's (2011) criticism is much the same: they say that the factor Hilbert space labels have "a merely formal significance".

[13]I emphasise that by 'average' I mean 'mean' not 'median'. It seems convincing to me that there *is* such a person as the *median* taxpayer—i.e. that this person is a physical entity. But of course, 'the median taxpayer' is a non-rigid designator (like 'the prime minister'), so the particular physical entity which is the median taxpayer is not invariant over time or possible worlds.

average taxpayer' may well be more vivid, but we must be careful not to mistake a means of representation for what is being represented. If the 'average taxpayer' exists at all, he/she (?) is not a physical thing. So too, I claim, for the factorist particles. Thus factorism contravenes a compulsory desideratum for the concept of particle: namely, that particles be physical. Rather, they are merely aspects of the mathematical formalism used to represent the physical world.

My claim prompts two immediate questions, one might say objections, which arise from my comparison of factorist particles with the average taxpayer. I answer each as they arise.

(i) In elementary statistics, there is, so to speak, just one average taxpayer. That is: each population defines a unique 'average element', that by definition possesses, for each quantity, the population's mean value. But in the quantum mechanical treatment of an assembly of N quantum particles—as we usually say!—there are N factor Hilbert spaces, and so N factorist particles. So my view faces the question: why are there N such factors, not just one?

    The answer is that the average taxpayer is constructed only to represent the *monadic properties* of the population, and only one is required to serve this purpose. In contrast, N factorist particles are necessary to represent the full collection of properties *and relations*, which together are taken to subvene the properties of the "N-particle" assembly.

(ii) In the case of the average taxpayer, it is clear that we are dealing with a construct, since it is obvious what physical things it is a construction *of*. However, in the case of factorist particles, I have not (yet) given any clue as to what physical things of which *they* are to be considered the constructions. What are these physical things?

    This objection is fair, and the rest of this thesis is dedicated to answer it. I will therefore say nothing now, except that it is a familiar phenomenon in interpretative philosophy of physics to correctly suspect a theory's formalism to contain unrepresentative elements before one has a firm grasp on the *represented* ontology. Such is the case with factorism.

6.4.4 A haecceitistic response

There is one final objection I wish to consider. It can be levelled only by a haecceitistic factorist, and runs as follows. My comparison between factorist particles and the average taxpayer is flawed because the shared feature which is supposed to arouse suspicion about their physical reality—i.e., the possession of a statistically mixed state whenever the population's state is heterogeneous—is an *essential* feature of the average taxpayer but a merely *accidental* feature of factorist particles. It is accidental in the latter case, since it relies on the indistinguishability postulate holding, and it is contingent whether IP is true.

This objection must come from a haecceitistic factorist, since it is only for her that IP may plausibly be construed as contingent. As we saw (Section 6.2.1), anti-haecceitism about factorist particles is *defined* as the view that any given physical state is represented equally well up to a permutation of factor Hilbert space labels. Therefore IP is compulsory for an anti-haecceitistic factorist. A haecceitistic factorist, on the other hand, may view IP as contingent, since she denies that distinct mathematical states related by a permutation represent the same physical state. For her, the restriction of the assembly's Hilbert space to a particular symmetry sector may be viewed, not as a representational necessity, but rather as a consequence of the conjunction of IP (as a contingent fact) with particular initial conditions (cf. French and Krause (2006, p. 148-9)).

This haecceitistic response is a good one. And were not haecceitism doomed for other reasons (cf. Caulton and Butterfield 2011), then haecceitistic factorism may, after all, have offered a rival to emergentism, which I endorse in Chapter 9. However, it is perhaps worth noting that, because (as I will argue in Chapter 9) emergentism vitiates the weak discernibility results of Section 6.3, neither position can use those results to discern their particles.

This concludes my discussion of factorism. I now turn to answering the question left at the end of Section 6.4.3, namely: Of *what* are the factorist particles statistical constructs? These leads us to my two rival target concepts of particle, varietism and emergentism. But first we need some subtleties about the notion of entanglement, as applied to indistinguishable systems, that have already been alluded to. This is the topic of the next Chapter.

Chapter 7

Entanglement and individuation for anti-factorists

The abandonment of factorism—the association of factor Hilbert spaces with the state spaces of particles, even under the imposition of the Indistinguishability Postulate—prompts a reappraisal of the use of some central concepts in quantum mechanics. Chief among these is entanglement: for entanglement is commonly defined as an assembly's state being non-separable (e.g. Bengtsson and Życzkowski (2006, p. 336)); and the physical significance of this mathematical feature is obviously tied to the association of particles with factor Hilbert spaces.

This Chapter is part report, part novel research; both parts are somewhat technical, and more mathematical than interpretative. The Chapter lays the ground for interpretative work in Chapters 8 and 9. It is organised in the following way: Section 7.1 advertises some recent technical results regarding the appropriate definition of entanglement in the context of the Indistinguishability Postulate. These results are due to Ghirardi, Marinatto and Weber (Ghirardi, Marinatto and Weber (2002), Ghirardi and Marinatto (2003, 2004, 2005)). They suggest a new definition of entanglement that is more attuned to the algebra of *permutation-invariant* operators than the previous notion of non-separability. In Section 7.2, I turn to the related idea of *individuating* (picking out, uniquely referring to) quantum systems. I present results which allow one to individuate systems by appealing not to factor Hilbert space labels but to single-system *states*; I call this *qualitative* individuation. This work is novel, but is inspired by the recent work of Zanardi and co-authors (Zanardi (2001), Zanardi *et al* (2004)). I show that qualitatively individuated systems possess an algebra of operators which has a natural tensor product structure, and I give a general prescription for calculating the reduced density operators for qualitatively individuated systems. Finally, Section 7.3 briefly investigates individuation over time.

7.1 Subtleties of entanglement

Ghirardi, Marinatto and Weber (2002) and Ghirardi and Marinatto (2003, 2004, 2005) show that the usual definition of entanglement—i.e., non-separability—can be naturally adapted to quantum systems governed by the Indistinguishability Postulate. Under their adapted definition, which I call *GM-entanglement*, the "entanglement" exhibited by fermions and bosons (and in fact also paraparticles, which they do not consider) is far closer to the usual entanglement exhibited by distinguishable systems. In particular, under their definition assemblies of fermions are *not* always GM-entangled, even though they must (due to anti-symmetrization) occupy non-separable states; and bosons need not occupy product states with identical factors (the only separable state available to an assembly of bosons) in order to be non-GM-entangled.

It is the purpose of this Section to present and explain the notion of GM-entanglement. First, I will discuss the familiar notion of entanglement for distinguishable systems (Section 7.1.1), before moving on to GM-entanglement for indistinguishable systems (Section 7.1.2).

Without further ado, let me lay down some convenient jargon. Following the literature, and previous Chapters, I shall call systems whose quantities are restricted by the Indistinguishability Postulate (IP) 'indistinguishable', and systems for which IP is not imposed 'distinguishable', even though (as we shall see in Section 8.2.6) both sorts of systems may be discernible—even absolutely discernible. And to clearly separate Ghirardi *et al*'s proposed definition of entanglement from the usual definition, I will continue to use the entrenched term 'entanglement' to mean the familiar mathematical idea of non-separability, and so introduce the term 'GM-entanglement' for Ghirardi *et al*'s concept. I will also universally use the word 'consituent' or 'system', instead of 'particle', since the results below do not depend on fulfilling the desiderata of Chapter 5.

In what follows, I will restrict myself to assemblies of two systems. However, the extension of the results to $N > 2$, and also to paraparticles, will in most cases be obvious; I will highlight any subtleties as they arise.

7.1.1 Entanglement for two distinguishable systems

In the case of two distinguishable systems (S_1 and S_2, say), the assembly's Hilbert space is simply $\mathcal{H}_1 \otimes \mathcal{H}_2$, where $\mathcal{H}_1$ is the Hilbert space for S_1, etc. (Thus factorism rules for distinguishable systems.) IP is not imposed, so we do not concentrate on the symmetric or anti-symmetric subspaces.

Entanglement for distinguishable systems is then defined in terms of non-separability. However, to show more clearly the naturalness of Ghirardi, Marinatto and Weber's extension of the concept of entanglement to the case of indistinguishable systems, and to follow their presentation more closely, we define:

> The system S_1, constituent of the assembly $S = S_1 + S_2$, described by the pure density operator ρ, is *non-entangled* with subsystem S_2 iff there exists a projection operator P onto a one-dimensional subspace of $\mathcal{H}_1$ such that:
> $$\mathrm{Tr}[\rho \; P \otimes \mathbb{1}] = 1. \tag{7.1}$$
> where $\mathbb{1}$ is the identity on $\mathcal{H}_2$.

This is equivalent to each of the following familiar conditions (assuming pure ρ):

1. The reduced density operator $\rho_1 = \mathrm{Tr}^{(2)}(\rho)$ of subsystem S_1 is a projection operator onto a one-dimensional subspace of $\mathcal{H}_1$ (i.e. S_1's state is pure);
2. Writing $\rho = |\psi\rangle\langle\psi|$: the state vector $|\psi\rangle$ is factorizable (i.e. separable); i.e., there exist a state $|\phi\rangle \in \mathcal{H}_1$ and a state $|\chi\rangle \in \mathcal{H}_2$ such that $|\psi\rangle = |\phi\rangle \otimes |\chi\rangle$.

The *support* of a density operator ρ is defined as the smallest ("logically strongest") projector P that ρ makes certain: $\mathrm{Tr}(\rho P) = 1$ and there is no $Q < P$ with $\mathrm{Tr}(\rho Q) = 1$. The support of ρ is the projector onto the range $\mathrm{ran}(\rho)$. So: for a system in the state ρ, a quantity A whose spectral decomposition contains the projector onto $\mathrm{ran}(\rho)$ has (with certainty) the corresponding value. But no refinement of A in the corresponding part of its spectrum does so. I again adopt the usual eigenvalue-eigenstate link (cf. Section 4.2.3). Then *only* such a quantity A, or a function of it, has a definite value. (We may also say that a quantity B for which the projector onto $\mathrm{ran}(\rho)$ is a sum of spectral projectors has an 'unsharp' value in the corresponding range.)

So: S_1 is entangled iff the reduced density operator $\rho_1 = Tr^{(2)}(\rho)$ of subsystem S_1 has a multi-dimensional support, $P_{\mathcal{M}_1}$ say, projecting onto a multi-dimensional subspace $\mathcal{M}_1 \subset \mathcal{H}_1$. With the eigenvalue-eigenstate link: this is so iff no one-dimensional projector has a value.

The extreme case ('total' or 'maximal' entanglement) is where $\mathcal{M}_1 = \mathcal{H}_1$, i.e. $P_{\mathcal{M}_1} = \mathbb{1}$. Then there is no self-adjoint operator on $\mathcal{H}_1$ for which one can claim with certainty that the outcome of its measurement will belong to any proper subset of its spectrum.

A phrase which is vivid, and is adopted by Ghirardi and Marinatto (2004, p. 2), is 'complete set of properties': a system with density operator ρ has a *complete set of properties* iff ρ is a one-dimensional projector, i.e. $\mathrm{ran}(\rho)$ is one-dimensional; and similarly, they say a system 'has the complete set of properties identified by ρ' etc. To avoid connotations of completions of quantum mechanics (hidden variable theories etc.), we will prefer to say: the system (or its state) is *maximally specific*; and similarly, I will say the system is *maximally specific à la* ρ. And similarly, if $\rho = |\psi\rangle\langle\psi|$, then I will say that the system is maximally specific *à la* $|\psi\rangle$.

Note that for two distinguishable systems, one system is maximally specific iff the other is. We will see in Section 7.1.2 that this is *not* true for indistinguishable bosons!

Example 1 (Ghirardi and Marinatto (2004, p. 3)): Suppose that an e^-e^+ system is described by the state vector (with obvious notation):

$$|\psi\rangle = \frac{1}{\sqrt{2}}\left(|\uparrow\rangle_{e^-} \otimes |\downarrow\rangle_{e^+} - |\downarrow\rangle_{e^-} \otimes |\uparrow\rangle_{e^+}\right) \otimes |R\rangle_{e^-} \otimes |L\rangle_{e^+}, \tag{7.2}$$

where $|R\rangle$ and $|L\rangle$ are two orthonormal states, whose coordinate representations have compact disjoint supports at the spatial regions Right and Left, respectively. (Note the use of the tensor product to bind both states of different systems and states of the same system associated with different degrees of freedom.) The reduced density operator describing the electron e^- acts on $\mathcal{H}_{e^-} = \mathbb{C}^2 \otimes L^2(\mathbb{R}^3)$ and has the following form:

$$\rho_{e^-} = \frac{1}{2}\left(|\uparrow\rangle\langle\uparrow| + |\downarrow\rangle\langle\downarrow|\right) \otimes |R\rangle\langle R| = \frac{1}{2}\mathbb{1} \otimes |R\rangle\langle R|. \tag{7.3}$$

Although there is no value of the spin along any direction, the electron is, with certainty, inside the right region R; (and an analogous statement can be made concerning the positron, i.e. that it is inside L). In short: some 'definite properties' are possessed because the range of the density operator (7.3) is a proper subspace of $\mathcal{H}_{e^-}$, i.e. the two-dimensional subspace spanned by $|\uparrow\rangle \otimes |R\rangle$ and $|\downarrow\rangle \otimes |R\rangle$.

Example 2: We now consider instead, for the e^-e^+ system:

$$|\psi\rangle = \frac{1}{\sqrt{2}}\left(|\uparrow\rangle_{e^-} \otimes |\downarrow\rangle_{e^+} - |\downarrow\rangle_{e^-} \otimes |\uparrow\rangle_{e^+}\right) \otimes \left(\sum_i c_i |\phi_i\rangle_{e^-} \otimes |\chi_i\rangle_{e^+}\right) \tag{7.4}$$

where $\forall i, c_i \neq 0$, and $\{|\phi_i\rangle\}$ and $\{|\chi_i\rangle\}$ are two complete orthonormal sets of the Hilbert spaces $L^2(\mathbb{R}^3)$ associated to the spatial degrees of freedom of the electron and positron, respectively. The reduced density operator for the electron is:

$$\rho_{e^-} = \mathrm{Tr}^{(2)}\left(|\psi\rangle\langle\psi|\right) = \frac{1}{2}\mathbb{1} \otimes \sum_i |c_i|^2 |\phi_i\rangle\langle\phi_i|. \tag{7.5}$$

In Equation (7.5), $\mathbb{1}$ is the identity operator on the spin space of the electron. Since the range of ρ_{e^-} is now the whole Hilbert space of the first particle, all we can say with certainty about the measurement of any self-adjoint operator is that the outcome will be in its spectrum.

Finally, we recall the familiar fact that non-entanglement corresponds to factorization of all probabilities for all joint measurements: An assembly's pure state $|\psi\rangle$ is non-entangled (i.e. its constituent systems are not entangled with one another) iff the following equation holds for any pair of quantities A of $\mathcal{H}_1$ and B of $\mathcal{H}_2$ such that $|\psi\rangle$ belongs to the domains of $A \otimes \mathbb{1}$ and $\mathbb{1} \otimes B$:

$$\langle\psi|A \otimes B|\psi\rangle = \langle\psi|A \otimes \mathbb{1}|\psi\rangle\langle\psi|\mathbb{1} \otimes B|\psi\rangle. \tag{7.6}$$

Since $|\psi\rangle$ is non-entangled iff $|\psi\rangle = |\varsigma\rangle \otimes |\chi\rangle$ for some $|\varsigma\rangle \in \mathcal{H}_1, |\chi\rangle \in \mathcal{H}_2$, the above expectation value may be expressed as

$$\langle\psi|A \otimes B|\psi\rangle = \langle\varsigma|A|\varsigma\rangle\langle\chi|B|\chi\rangle. \tag{7.7}$$

We will return to this last feature of non-entangled states for distinguishable systems much later (in Stage E of Section 7.1.2 and Section 7.2). Before that, we turn to the definition of GM-entanglement.

7.1.2 GM-entanglement for two indistinguishable systems

I now explore the consequences of Ghirardi and Marinatto's adaptation of the previous definition (Equation (7.1) above) of non-entanglement. The main results will be:

(i) Since IP must be obeyed for indistinguishable systems, the basic idea of non-entanglement—viz. some 1-dimensional single-system projector being certain—needs to be made precise using *symmetric* quantities. This revision of the definition will involve a projector representing, *for a factorist*, the idea that *at least one* of the systems is in the state associated with that projector.

(ii) GM-entanglement for indistinguishable systems shares many of the attractive features of entanglement for distinguishable systems. In particular: (a) there is a sense in which an assembly's state is non-GM-entangled iff it supervenes on the states of its constituents; and (b) analogues of Bell's Theorem (1964) and Gisin's Theorem (1991) hold for the GM-entanglement of two-system assemblies. (This latter result will be proved much later, once we have introduced the idea of *qualitative individuation*, in Section 7.2.)

(iii) The relation between maximal specificity and non-entanglement is weaker for bosons and paraparticles than for fermions or for distinguishable systems. To be more precise: the state of two distinguishable systems is non-entangled iff *either* system is maximally specific, iff *both* systems are maximally specific. This is because for distinguishable systems and fermions we have an equivalence: one system is maximally specific iff the other is. We will see that this is *not* true for bosons or paraparticles.[1] Since we will want non-GM-entanglement to be a symmetric relation between the systems—amongst other things, this will allow us to treat non-entanglement as a property of *states of the assembly*—non-GM-entanglement is defined in terms of *both* systems of the assembly being maximally specific.

(iv) Non-GM-entangled states of fermions and paraparticles possess a feature which is not shared by either non-entangled distinguishable systems, nor non-GM-entangled bosons. This feature is that, for any non-GM-entangled state of the assembly, decomposition into maximally specific systems is *not unique*. This result will be especially important for Chapter 8, below.

This Section is divided into five Stages. In Stage A, I define in one go, both for bosons and fermions: first non-entanglement; second, at least one system being maximally specific. In Stage B, I show that at least one system being maximally specific is equivalent to the composite's state-vector being obtained by symmetrizing or anti-symmetrizing a factorized state. In Stage C, I specialize this theorem to fermions. In Stage D, I specialize this theorem to bosons and paraparticles. In Stage E, I turn to entanglement and correlations—connecting with the examples at the end of Section 7.1.1. I emphasise that these results are not novel: they are all borrowed from Ghirardi and Marinatto (2003).

Stage A: Non-GM-entanglement, and being maximally specific

[1] The exception arises for multiply occupied single-system states, so it does not apply to fermions.

To represent an assembly of two indistinguishable systems, we usually begin with the Hilbert space $\mathcal{H} \otimes \mathcal{H}$, where $\mathcal{H}$ is the Hilbert space for one such system. As is familiar (cf. e.g. French and Krause (2006, Ch. 4)), the Indistinguishability Postulate induces a decomposition of the assembly's Hilbert space into symmetry sectors—in this case the symmetric subspace $\mathcal{S}(\mathcal{H} \otimes \mathcal{H})$ (bosons) or the anti-symmetric subspace $\mathcal{A}(\mathcal{H} \otimes \mathcal{H})$ (fermions).

Inspired by Ghirardi, Marinatto and Weber (2002), I define:

> The indistinguishable constituents of a two-system assembly are *non-GM-entangled* iff *both* systems are maximally specific.

(The assembly is then defined as GM-entangled iff it is not non-GM-entangled.) And we define:

> Given an assembly of two indistinguishable systems, one of the systems is *maximally specific* iff there exists a one-dimensional projector P, defined on $\mathcal{H}$, such that:
> $$\mathrm{Tr}(\rho E) = 1 \tag{7.8}$$
> where
> $$E := P \otimes \mathbb{1} + \mathbb{1} \otimes P - P \otimes P. \tag{7.9}$$

Indeed, extending my terminology from Section 7.1.1, we may say that one of the systems is *maximally specific à la P*.

E is invariant under action by the symmetric group S_2, and it is a projection operator: $E^2 = E$. Furthermore $\mathrm{Tr}(\rho E)$ is the probability of finding *at least* one of the two systems in the state onto which the one-dimensional operator P projects. For E can also be written as

$$E = (\mathbb{1} - P) \otimes \mathbb{1} + \mathbb{1} \otimes (\mathbb{1} - P) + P \otimes P \ ; \tag{7.10}$$

so this definition of 'one of the systems is maximally specific' is really a definition of 'at least one system is maximally specific'.

Stage B: 'At least one being maximally specific': a general theorem

The definition just given is equivalent to the assembly's state taking a certain form:

Theorem (cf. Ghirardi and Marinatto (2003, Theorems 4.2 & 4.3)): At least one of the systems in a two-system assembly is maximally specific iff the assembly's state is obtained by symmetrizing or anti-symmetrizing a separable (i.e. product) state.

Sketch of Proof:
Right to left (easy half): If $|\psi\rangle$ is obtained by symmetrizing or anti-symmetrizing a factorized state of two indistinguishable constituents, then:

$$|\psi\rangle = N\left(|\phi\rangle \otimes |\chi\rangle \pm |\chi\rangle \otimes |\phi\rangle\right). \tag{7.11}$$

By expressing the state $|\chi\rangle$ as

$$|\chi\rangle = \alpha|\phi\rangle + \beta|\phi^{\perp}\rangle, \qquad \langle\phi|\phi^{\perp}\rangle = 0 \tag{7.12}$$

and choosing $P = |\phi\rangle\langle\phi|$, one gets immediately

$$\mathrm{Tr}(\rho E) \equiv \langle\psi|E|\psi\rangle = \frac{2(1 \pm |\alpha|^2)}{2(1 \pm |\alpha|^2)} = 1. \tag{7.13}$$

Left to Right (hard half): If one chooses a complete orthonormal set of single-system states whose first element $|\phi_0\rangle := |\phi\rangle$ spans the range of P, writing

$$|\psi\rangle = \sum_{ij} c_{ij}|\phi_i\rangle \otimes |\phi_j\rangle, \qquad \sum_{ij} |c_{ij}|^2 = 1, \tag{7.14}$$

and, using the explicit expression for E in terms of P, one obtains:

$$E|\psi\rangle = |\phi_0\rangle \otimes \left(\sum_{j\neq 0} c_{0j}|\phi_j\rangle\right) + \left(\sum_{j\neq 0} c_{j0}|\phi_j\rangle\right) \otimes |\phi_0\rangle + c_{00}|\phi_0\rangle \otimes |\phi_0\rangle. \tag{7.15}$$

Imposing condition (7.8) implies that $E|\psi\rangle = |\psi\rangle$ (since E is a projector), which implies that $c_{ij} = 0$ for $i, j \neq 0$. Taking into account that for indistinguishable systems $c_{0j} = \pm c_{j0}$, the normalization condition of the state $|\psi\rangle$ becomes

$$|c_{00}|^2 + 2\sum_{j\neq 0} |c_{0j}|^2 = 1. \tag{7.16}$$

We have then shown that:

$$|\psi\rangle = |\phi_0\rangle \otimes \left(\sum_{j\neq 0} c_{0j}|\phi_j\rangle\right) + \left(\sum_{j\neq 0} c_{j0}|\phi_j\rangle\right) \otimes |\phi_0\rangle + c_{00}|\phi_0\rangle \otimes |\phi_0\rangle. \tag{7.17}$$

In the case of fermions $c_{00} = 0$. Then, introducing a normalized vector $|\xi\rangle := \sqrt{2}\sum_{j\neq 0} c_{0j}|\phi_j\rangle$ we obtain

$$|\psi\rangle = \frac{1}{\sqrt{2}}\left(|\phi_0\rangle \otimes |\xi\rangle - |\xi\rangle \otimes |\phi_0\rangle\right), \tag{7.18}$$

where $\langle\phi_0|\xi\rangle = 0$. For bosons, defining the following normalized vector

$$|\theta\rangle = \sqrt{\frac{4}{2-|c_{00}|^2}}\left(\sum_{j\neq 0} c_{0j}|\phi_j\rangle + \frac{c_{00}}{2}|\phi_0\rangle\right), \tag{7.19}$$

the two-particle state vector (7.17) becomes

$$|\psi\rangle = \sqrt{\frac{2-|c_{00}|^2}{4}}\left(|\phi_0\rangle \otimes |\theta\rangle + |\theta\rangle \otimes |\phi_0\rangle\right). \tag{7.20}$$

Note that in this case the states $|\phi_0\rangle$ and $|\theta\rangle$ are orthogonal iff $c_{00} = 0$, in which case $|\theta\rangle = |\xi\rangle$. $\square$

I now deal separately with the case of fermions and bosons. I will denote the appropriate restrictions of the operator E of equation (7.9) as E_f and E_b in the two cases, respectively.

Stage C: 'At least one being maximally specific': the theorem for fermions

Since $P \otimes P = 0$ on the space of anti-symmetric states $\mathcal{A}(\mathcal{H} \otimes \mathcal{H})$, one can drop such a term in all previous formulae. Accordingly, $E_f = P \otimes \mathbb{1} + \mathbb{1} \otimes P$.

In accordance with the definition of maximal specificity in Stage A: due to the orthogonality of $|\phi_0\rangle$ and $|\xi\rangle$, for the state in equation (7.18), we conclude not only that there is one fermion that is maximally specific *à la* the state $|\phi_0\rangle$ (in Ghirardi and Marinatto's jargon: one fermion possessing the complete set of properties identified by the state $|\phi_0\rangle$), but also that *the other* fermion is maximally specific *à la* $|\xi\rangle$.

So according to Stage A's definition of non-entanglement, we have proved:

> The fermions of a two-system assembly, whose (pure) state is given by $|\psi\rangle$, are non-GM-entangled iff $|\psi\rangle$ is obtained by anti-symmetrizing a separable state.

We may now extend the definition of GM-entanglement to apply to the states of an assembly: a two-fermion assembly's *state* may then be described as GM-entangled iff its constituent fermions are GM-entangled.

We can say more. If in expression (7.18), we write $P = |\phi_0\rangle\langle\phi_0|$ and $Q = |\xi\rangle\langle\xi|$ and we define the operators $E_f := P \otimes \mathbb{1} + \mathbb{1} \otimes P$ and $F_f := Q \otimes \mathbb{1} + \mathbb{1} \otimes Q$, we see that

$$\left.\begin{array}{rcl} \mathrm{Tr}\,(E_f|\psi\rangle\langle\psi|) & = & 1, \\ \\ \mathrm{Tr}\,(F_f|\psi\rangle\langle\psi|) & = & 1, \end{array}\right\} \tag{7.21}$$

Moreover, $E_f F_f = F_f E_f = P \otimes Q + Q \otimes P$ (since $P \perp Q$) is a one-dimensional projector on the Hilbert space of the assembly, and it satisfies:

$$\langle\psi|E_f F_f|\psi\rangle \equiv \langle\psi|\,(P \otimes Q + Q \otimes P)\,|\psi\rangle = 1. \tag{7.22}$$

It is this identity that allows us to say, not only that one fermion is maximally specific *à la* P and one fermion is maximally specific *à la* Q, but that one fermion is maximally specific *à la* P and *the other* fermion is maximally specific *à la* Q. This crucial difference will be important in the next stage, where we consider bosons and paraparticles, since for those symmetry types the two formulations come apart.

I note here an important feature of fermionic states, which will be crucial later, in Chapter 8. There is an arbitrariness, up to two dimensions, about the states *à la* which the two fermions are maximally specific. For suppose the state $|\psi\rangle$ is (7.18), and consider the two-dimensional subspace of $\mathcal{H}$ spanned by the single-system states $|\phi_0\rangle$ and $|\xi\rangle$: clearly, if one chooses *any* two other orthogonal single-constituent states $|\kappa\rangle$ and $|\lambda\rangle$ spanning the same subspace, then $|\psi\rangle$ can also be written (up to an overall phase factor) as:

$$|\psi\rangle = \frac{1}{\sqrt{2}} \left(|\kappa\rangle \otimes |\lambda\rangle - |\lambda\rangle \otimes |\kappa\rangle \right) . \tag{7.23}$$

Thus the two fermions are *also* maximally specific *à la* $|\kappa\rangle$ and $|\lambda\rangle$. Since this arbitrariness involves fixing an orthonormal basis in $\mathbb{C}^2$, it may be parameterized by $\mathbb{CP}^1 \cong \mathbb{C} \cup \{\infty\}$, the Riemann sphere. In fact, each orthonormal basis is represented by a *pair* of antipodal points $\{z, \frac{1}{z}\}$ on the Riemann sphere, since we can permute the two basis vectors without change.

A vivid visual metaphor of this basis freedom is provided by the fact that a pair of antipodal points on the Riemann sphere specifies a line through its centre, and a plane orthogonal to that line that divides the sphere into two hemispheres. Thus the arbitrariness in the single-system states *à la* which two non-GM-entangled fermions are maximally specific corresponds to the continuum-many ways one may divide a sphere into two hemispheres; cf. Figure 7.1.

In the case of three or more fermions, the above results all apply. Specifically: as is true for distinguishable systems, one fermion in the assembly is maximally specific iff they all are, iff the assembly's state is non-GM-entangled, iff the state is obtained by anti-symmetrizing a separable state.

Furthermore, the "preferred basis problem" applying to two-fermion assemblies is only exacerbated by the presence of more fermions. To see this, it is enough to notice that total anti-symmetrization is constituted by all possible pair-wise anti-symmetrizations. For example, in the non-GM-entangled $N = 3$ state $|\Psi\rangle$,

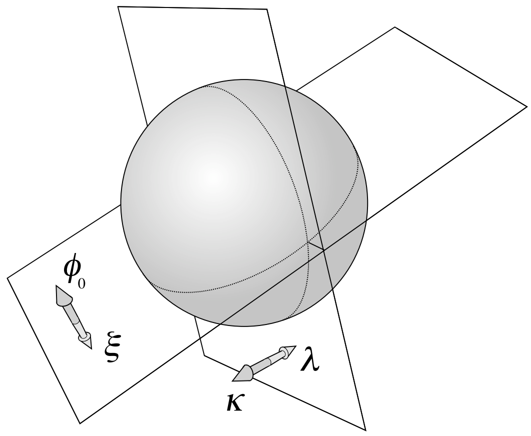

Figure 7.1: Two ways to halve the Riemann sphere, and two pairs of orthogonal states, *à la* which two non-entangled fermions may be maximally specific.

where we define $s(i) := (i+1) \bmod 3$ (so $s^3(i) \equiv i$),

$$
\begin{aligned}
|\Psi\rangle &= \frac{1}{\sqrt{6}} \begin{vmatrix} |\alpha\rangle_1 & |\beta\rangle_1 & |\gamma\rangle_1 \\ \\ |\alpha\rangle_2 & |\beta\rangle_2 & |\gamma\rangle_2 \\ \\ |\alpha\rangle_3 & |\beta\rangle_3 & |\gamma\rangle_3 \end{vmatrix} \\
&\equiv \frac{1}{\sqrt{6}} \sum_{i=1}^{3} |\alpha\rangle_i \otimes \Big(|\beta\rangle_{s(i)} \otimes |\gamma\rangle_{s^2(i)} - |\gamma\rangle_{s(i)} \otimes |\beta\rangle_{s^2(i)} \Big) \\
&\equiv \frac{1}{\sqrt{6}} \sum_{i=1}^{3} |\beta\rangle_i \otimes \Big(|\gamma\rangle_{s(i)} \otimes |\alpha\rangle_{s^2(i)} - |\alpha\rangle_{s(i)} \otimes |\gamma\rangle_{s^2(i)} \Big) \\
&\equiv \frac{1}{\sqrt{6}} \sum_{i=1}^{3} |\gamma\rangle_i \otimes \Big(|\alpha\rangle_{s(i)} \otimes |\beta\rangle_{s^2(i)} - |\beta\rangle_{s(i)} \otimes |\alpha\rangle_{s^2(i)} \Big) \qquad (7.24)
\end{aligned}
$$

(where we have added labels simply to make the ordering in the tensor product clear), all of the bracketed two-fermion states are subject to the usual basis arbitrariness (each in a different two-dimensional subspace of $\mathcal{H}$).

Generally, any non-GM-entangled state of N fermions:

$$\frac{1}{\sqrt{N!}} \begin{vmatrix} |\phi_1\rangle_1 & |\phi_2\rangle_1 & \cdots & |\phi_N\rangle_1 \\ |\phi_1\rangle_2 & |\phi_2\rangle_2 & \cdots & |\phi_N\rangle_2 \\ \vdots & \vdots & \ddots & \vdots \\ |\phi_1\rangle_N & |\phi_2\rangle_N & \cdots & |\phi_N\rangle_N \end{vmatrix} \tag{7.25}$$

(where $\{|\phi_1\rangle, |\phi_2\rangle, \ldots |\phi_N\rangle\}$ is an orthonormal basis of the single-system Hilbert space) is the *unique* totally anti-symmetric state in the Hilbert space $\bigotimes^N \mathfrak{h}$, where $\mathfrak{h} := \mathrm{span}(\{|\phi_1\rangle, |\phi_2\rangle, \ldots |\phi_N\rangle\}) \cong \mathbb{C}^N$. Therefore the same state (7.25) is picked out when the single-system Hilbert space $\mathfrak{h}$ is decomposed into *any other* orthonormal basis. Therefore the basis arbitrariness for any non-GM-entangled state of a N-fermion assembly corresponds to the arbitrariness in selecting a basis in $\mathbb{C}^N$. Each basis—and therefore each set of N single-particle states *à la* which N systems may be said to be maximally specific—corresponds to a point in the manifold

$$\left(\mathbb{CP}^{N-1} \times \mathbb{CP}^{N-2} \times \cdots \times \mathbb{CP}^1\right)/S_N \tag{7.26}$$

(where we quotient by the natural group action of S_N, since a permutation of basis vectors does not change the basis). This manifold has $(N-1)!2^{N-1}$ real dimensions!

Finally, since in the state of an assembly of any type of paraparticles, the states of at least two systems are pair-wise anti-symmetrized (Tung (1985, Ch. 5)), this basis arbitrariness applies as much to paraparticles as to fermions, though parameterizing this arbitrariness is a more subtle matter.

Stage D: 'At least one being maximally specific': the theorem for bosons

The broad similarity to fermions is clear, especially from Equations (7.11) and (7.20). As for fermions, the requirement that one of the two bosons is maximally specific implies that the state is obtained by symmetrizing a factorized state. However, there are some remarkable differences from the fermion case. For bosons, three cases are possible, according to the single-system states that are the factors of the separable state:

1. $|\theta\rangle \propto |\phi_0\rangle$, i.e. $|c_{00}| = 1$. Then the state is $|\psi\rangle = |\phi_0\rangle \otimes |\phi_0\rangle$ and one can infer that there are two bosons each maximally specific *à la* (each with the complete set of properties associated to) $P = |\phi_0\rangle\langle\phi_0|$. It may checked that for this state $\langle\psi|P \otimes P|\psi\rangle = 1$.

2. $\langle\theta|\phi_0\rangle = 0$, i.e. $c_{00} = 0$. One can then consider the operators E_b and F_b, defined as I did for fermions in Stage C, and their product $E_bF_b = F_bE_b = P \otimes Q + Q \otimes P$. Then exactly the same argument as for fermions implies that one of the two bosons is maximally specific *à la* P and the other of the two bosons is maximally specific *à la* Q. That is, $\langle\psi|E_bF_b|\psi\rangle = \langle\psi|(P \otimes Q + Q \otimes P)|\psi\rangle = 1$.

3. Finally, it can happen that $0 < |\langle\theta|\phi_0\rangle| < 1$. In this case, it is true that there is a boson maximally specific *à la* P—i.e., $\langle\psi|E_b|\psi\rangle = 1$—and it is true that there is a boson maximally specific *à la* Q—i.e., $\langle\psi|F_b|\psi\rangle = 1$. But it is *not true* that *one* is maximally specific *à la* P and *the other* is maximally specific *à la* Q. (Note that in this case $E_bF_b \neq F_bE_b$, and neither are equal to $P \otimes Q + Q \otimes P$.) For there is a non vanishing probability of finding both particles in the same state, since $\langle\psi|P \otimes P|\psi\rangle = \langle\psi|Q \otimes Q|\psi\rangle = |c_{00}|^2 > 0$.[2]

According to our our definition of non-GM-entanglement, both bosons are non-GM-entangled for the first two cases. But in the last case we cannot say that *both* bosons are non-GM-entangled, even though we may say that one system is maximally specific *à la* one projector, and one system is maximally specific à la another, distinct (but not orthogonal) projector. The worry, of course, is that we are counting contributions from the same system each time, so we must resist the plural article 'both'. So we count any state of this third type as GM-entangled, being a superposition of states of the first and second types. To sum up:

[2]A plausible candidate projector to associate with the proposition, 'One system is maximally specific *à la* P and the other is maximally specific *à la* Q' is $G := P \otimes Q + Q \otimes P - (P \otimes P + Q \otimes Q - PQ \otimes PQ)$, where we must take advantage of the fact that $\langle\psi|PQ \otimes PQ|\psi\rangle = \langle\psi|QP \otimes QP|\psi\rangle$ for $|\psi\rangle$ in (7.20). Then $\langle\psi|G|\psi\rangle = \frac{2-4|c_{00}|^2+3|c_{00}|^4}{2-|c_{00}|^2}$, which is less than 1 iff $0 < |\langle\theta|\phi_0\rangle| < 1$.

> The bosons of a two-system assembly, whose (pure) state is given by $|\psi\rangle$, are non-GM-entangled iff either: (i) $|\psi\rangle$ is obtained by symmetrizing a factorized product of two orthogonal states; or (ii) $|\psi\rangle$ is a product state of identical factors.

We may then again extend the definition to apply to states of the assembly in the obvious way: the assembly's state is non-GM-entangled iff both systems are.

Note that in the boson case, for any non-GM-entangled state, the two states *à la* which the two systems are maximally specific are uniquely determined—contrary to what happens for fermions (or paraparticles). That is: there are no other orthogonal states $|\kappa\rangle$ and $|\lambda\rangle$, differing from $|\phi_0\rangle$ and $|\theta\rangle$, such that one can write $|\psi\rangle$ in the form

$$|\psi\rangle = \frac{1}{\sqrt{2}}\left(|\kappa\rangle \otimes |\lambda\rangle + |\lambda\rangle \otimes |\kappa\rangle\right). \tag{7.27}$$

Finally, the foregoing results carry over to the case of three or more bosons. The results and definitions also apply to paraparticles, since for any assembly of paraparticles of any single type, at least two constituent systems may occupy the same state (Tung (1985, Ch. 5)). However, we must, of course, replace the phrase 'symmetrized state' in the above with the phrase 'state with symmetry type μ', where μ is the appropriate paraparticle type.

To sum up the previous two stages, for non-GM-entanglement, in our sense that *all* constituents are maximally specific, it must be the case that:

(i) the state for the assembly is obtained, by the appropriate symmetry projection, from a separable state in $\bigotimes^N \mathcal{H}$; and

(ii) the factors of the separable state in question must be orthogonal in the fermion case, and they can be either orthogonal or equal in the boson or paraparticle case.

Stage E: Entanglement and correlations

I turn at last to entanglement! GM-entangled states can very well occur; and here we connect with the examples at the end of Section 7.1.1. Thus consider the following state of two spin-1/2 constituents:

$$|\psi\rangle = \frac{1}{\sqrt{2}}\left(|\uparrow\rangle \otimes |\downarrow\rangle - |\downarrow\rangle \otimes |\uparrow\rangle\right) \otimes |\omega_{12}\rangle \tag{7.28}$$

with $|\omega_{12}\rangle$ a symmetric state of $L^2(\mathbb{R}^3) \otimes L^2(\mathbb{R}^3)$. State (7.28) cannot be written as a symmetrized product of two orthogonal states, and, consequently no constituent is maximally specific (in Ghirardi-argot: possesses any complete set of (internal and spatial) properties).

This sort of example also returns us to the topic of non-entanglement corresponding to factorization of probabilities for joint measurements. To connect with the Bell's theorem literature, let us consider case in which the two constituents are in different spatial regions. Let us then consider two indistinguishable constituents with space and internal degrees of freedom and let us denote as $\mathcal{H}_{sp}$ and $\mathcal{H}_{int}$ the corresponding single-system Hilbert spaces. The Hilbert space for the whole system is the appropriate symmetric or antisymmetric subspace, $\mathcal{S}(\mathcal{H}_{int}^{(1)} \otimes \mathcal{H}_{sp}^{(1)} \otimes \mathcal{H}_{int}^{(2)} \otimes \mathcal{H}_{sp}^{(2)})$ or $\mathcal{A}(\mathcal{H}_{int}^{(1)} \otimes \mathcal{H}_{sp}^{(1)} \otimes \mathcal{H}_{int}^{(2)} \otimes \mathcal{H}_{sp}^{(2)})$, respectively. Let us also assume that the pure state associated to the composite system is obtained by (anti-)symmetrizing a factorized state of the two particles corresponding to their having different spatial locations. To be explicit, we start from a state:

$$|\psi_{fact}\rangle = |\varsigma, R\rangle \otimes |\chi, L\rangle, \tag{7.29}$$

where e.g. $|\varsigma, R\rangle$ is an abbreviation for the single-system state $|\varsigma\rangle \otimes |R\rangle$, and $|\varsigma\rangle$ and $|\chi\rangle$ are two arbitrary states of the internal space of a single systems and $|R\rangle$ and $|L\rangle$ are two orthogonal states whose spatial supports are compact, disjoint and far away from each other. This situation is the one of interest for all experiments about the non-local features of quantum states. From the state (7.29) we pass now to the properly (anti-)symmetrized state:

$$|\psi\rangle = \frac{1}{\sqrt{2}}\left(|\varsigma, R\rangle \otimes |\chi, L\rangle \pm |\chi, L\rangle \otimes |\varsigma, R\rangle\right). \tag{7.30}$$

We already know that for the operators

$$\begin{aligned} E &= P \otimes \mathbb{1} + \mathbb{1} \otimes P - P \otimes P, \quad P = |\varsigma, R\rangle\langle\varsigma, R| \\ F &= Q \otimes \mathbb{1} + \mathbb{1} \otimes Q - Q \otimes Q, \quad Q = |\chi, L\rangle\langle\chi, L| \end{aligned} \tag{7.31}$$

the following equations hold:

$$\mathrm{Tr}\,(E|\psi\rangle\langle\psi|) = 1, \qquad \mathrm{Tr}\,(F|\psi\rangle\langle\psi|) = 1, \tag{7.32}$$

which imply that the constituents are maximally specific *à la* projectors P and Q.

However, here we are interested in the correlations between the outcomes of measurement processes on the constituents—this will offer a taster for the results in Section 7.2. For any two operators $A, B \in \mathcal{B}(\mathcal{H}_{int})$, we construct the operators

$$\left.\begin{aligned} A_R &:= (A \otimes |R\rangle\langle R|)_1 \otimes \mathbb{1}_2 + \mathbb{1}_1 \otimes (A \otimes |R\rangle\langle R|)_2 \\ B_L &:= (B \otimes |L\rangle\langle L|)_1 \otimes \mathbb{1}_2 + \mathbb{1}_1 \otimes (B \otimes |L\rangle\langle L|)_2 \end{aligned}\right\} \tag{7.33}$$

(where I use labels for clarity, and $\mathbb{1}$ is the identity on $\mathcal{H}_{int} \otimes \mathcal{H}_{sp}$). Now it may be checked that A_R and B_L are symmetric—that is, they obey IP—*for any* $A, B \in \mathcal{B}(\mathcal{H}_{int})$. Therefore, any Hermitian A, B satisfy the necessary condition for representing a physical quantity under IP. In fact, I wish to interpret A_R as an operation on any system whose spatial wavefunction's support overlaps R (and similarly for B_L and the region L).

It may then be checked that

$$\langle\psi|A_R B_L|\psi\rangle = \langle\varsigma|A|\varsigma\rangle\langle\chi|B|\chi\rangle\ . \tag{7.34}$$

Equation (7.34) shows that the probabilities referring to the internal degrees of freedom factorize, just as in the case of two distinguishable constituents; cf. Equations (7.6) and (7.7). The same conclusion does not hold when the state is a genuinely entangled state, such as:

$$|\psi'\rangle = \frac{1}{2}\left(|\varsigma\rangle_1 \otimes |\chi\rangle_2 - |\chi\rangle_1 \otimes |\varsigma\rangle_2\right) \otimes \left(|R\rangle_1 \otimes |L\rangle_2 \pm |L\rangle_1 \otimes |R\rangle_2\right), \tag{7.35}$$

for which

$$\langle\psi'|A_R B_L|\psi'\rangle = \frac{1}{2}\langle\varsigma|A|\varsigma\rangle\langle\chi|B|\chi\rangle + \frac{1}{2}\langle\chi|A|\chi\rangle\langle\varsigma|B|\varsigma\rangle - \mathrm{Re}\left[\langle\varsigma|A|\chi\rangle\langle\chi|B|\varsigma\rangle\right] \ , \tag{7.36}$$

in which quantum interference is clearly manifested. (Compare the expectation value $\langle\psi''|A\otimes B|\psi''\rangle$ for the entangled state of two *distinguishable* systems, with only internal degrees of freedom: $|\psi''\rangle = \frac{1}{\sqrt{2}}\left(|\varsigma\rangle_1\otimes|\chi\rangle_2 - |\chi\rangle_1\otimes|\varsigma\rangle_2\right)$.)

The upshot should be obvious: from the point of view of the correlations, and consequently of the implications concerning nonlocality, the non-GM-entangled states of two indistinguishable systems have some of the (rather nice) features as the non-entangled states of two distinguishable systems.

This prompts the suggestion (made in (ii)(a) at this beginning of this Section) that non-GM-entangled states supervene on the states of constituent, maximally specific systems—just as non-entangled states supervene on the states of constituent, distinguishable systems. Indeed, we see that this true, *so long as* the symmetry type of the assembly is determined by the states of the constitutent systems.[3] For, given a symmetry type, any collection of N single-system states determines at most *one* non-GM-entangled state, of that symmetry type, for the N-system assembly for which there are N systems maximally specific *à la* one of the N single-system states. I return to this in Section 8.2.4.

This concludes my reconstruction of the Ghirardi, Marinatto and Weber results. I now turn to a novel idea, which is inspired by these results: qualitative individuation of quantum systems.

7.2 Qualitatively individuating quantum systems

A central idea that may be taken from Ghirardi *et al*'s idea of a maximally specific system is picking out a system according to the state it occupies; for any maximally specific system is maximally specific *à la* some state. Let me use the term *individuation* for this act of picking out a object, or collection of objects,

[3]The spin-statistics theorem would provide the necessary connection; but I emphasise that, strictly speaking, this theorem lies outside the realm of elementary quantum mechanics.

according to some property that it may have; and let me call the property in question the *individuation criterion*. Thus an object need not be an *individual* (cf. Section 2.3's Interlude) in order to be individuated in the present sense. But it will be *uniquely* picked out only if it *is* an individual. In general, an object is individuated by some individuation criterion iff every other object in its absolute indisernibility equivalence class is likewise individuated (cf. Section 2.4). That is: any individuation criterion succeeds in uniquely picking out, not single objects, but absolute indiscernibility classes of objects.

Let us further say that an object, or class of objects, is *qualitatively* individuated iff its individuation criterion is a qualitative property. In quantum mechanics, I submit, qualitative properties are represented by projectors. (Factorist individuation—i.e. individuation according to factor Hilbert space labels—may be considered non-qualitative individuation.) Thus maximally specific systems, since they are individuated using (one-dimensional) projectors, are qualitatively individuated systems. And since the individuating projectors are one-dimensional, the corresponding qualitative property is logically strong, and maximally rich.

But qualitatively individuated systems need not be maximally specific. I will propose individuating systems using *multi*-dimensional projectors. This is the guiding idea of the remaining Sections of this Chapter. In this Section, I will first introduce the main ideas behind qualitative individuation, and link it to natural decompositions of the assembly's Hilbert space (Section 7.2.1). Then (Section 7.2.2) I will propose two reasonable ways to deal with failure of individuation. Finally (Section 7.2.3), I will give a general prescription for calculating the reduced density operator for a qualitatively individuated system.

7.2.1 Qualitative individuation and natural decompositions

One may cash out the idea of a constituent of an assembly in terms of natural decompositions of the assembly's Hilbert space. In the case of distinguishable systems—for which I endorse factorism—the natural decomposition is given *ab initio*: it is the decomposition into the factor Hilbert spaces corresponding to each distinguishable system. Even in this case, the systems are qualitatively individuated—the difference is that the individuation criteria are state-invariant ('intrinsic') properties, so they are not represented in the formalism by projectors.

In the case of indistinguishable systems, since I deny factorism, we must work in the opposite direction: i.e., we are given the assembly's Hilbert space and we must search for its natural decompositions. I will argue here that *subspaces* of the assembly's Hilbert space may be naturally decomposed into spaces which represent the possibilities for qualitatively individuated systems.

What counts as a "natural decomposition"? The answer, suggested by Zanardi (2001, p. 1), is provided by the algebra of quantities defined for the assembly. To be more specific (Zanardi (2001, p. 3)):

> When is it legitimate to consider a pair of observable algebras as describing a bipartite quantum system? Suppose that $\mathcal{A}_1$ and $\mathcal{A}_2$ are two *commuting* $*$-subalgebras of $\mathcal{A} := \mathrm{End}(\mathcal{H})$ such that the subalgebra $\mathcal{A}_1 \vee \mathcal{A}_2$ they generate, i.e., the minimal $*$-subalgebra containing both $\mathcal{A}_1$ and $\mathcal{A}_2$, amounts to the whole $\mathcal{A}$, and moreover one has the (noncanonical) algebra isomorphism,
>
> $$\mathcal{A}_1 \vee \mathcal{A}_2 \cong \mathcal{A}_1 \otimes \mathcal{A}_2 \tag{7.37}$$
>
> The standard, *genuinely* bipartite, situation is of course $\mathcal{H} = \mathcal{H}_1 \otimes \mathcal{H}_2, \mathcal{A}_1 = \mathrm{End}(\mathcal{H}_1) \otimes \mathbb{1}$, and $\mathcal{A}_2 = \mathbb{1} \otimes \mathrm{End}(\mathcal{H}_2)$. If $\mathcal{A}_i' := \{X \mid [X, \mathcal{A}_1] = 0\}$ denotes the *commutant* of $\mathcal{A}_1$, in this case one has $\mathcal{A}_i' = \mathcal{A}_2$.

Thus Zanardi's proposal is to work by analogy with the distinguishable case: we look for commuting subalgebras whose tensor product is isomorphic to the entire algebra for the assembly's Hilbert space. The requirements of commutativity and isomorphism with the entire algebra are tantamount to two of Zanardi *et al*'s (2004, p. 1) three necessary and jointly sufficient conditions for a natural decomposition: what he calls *subsystem independence* and *completeness*, respectively. The remaining requirement, *local accessibility*, is that the subalgebras be "controllable". For us, this is tantamount to satsifying IP. Since we will only consider subalgebras of the algebra of symmetric quantities, we can take this requirement to be fulfilled by default.

However, unlike for distinguishable systems, in the case of indistinguishable systems, prospects seem dim for finding natural decompositions of the assembly's entire Hilbert space. For the Hilbert space can even have a *prime* number of dimensions. (E.g., the Hilbert space for a pair of two-level bosons has three dimensions.)

My basic idea in this Section is that we may instead look for natural decompositions of *subspaces* of the assembly's Hilbert space. The collection of constituents corresponding to these decompositions must then be interpreted as co-existing *only* in those states belonging to the given subspace. But this is not objectionable *per se*. Agreed: in the case of distinguishable systems—and even for haecceitistic factorists—there are means of individuating systems which will suffice for all states. But if one is not a haecceitist, why should one demand or expect this across the board?

Now that we have limited our search for natural decompositions to subspaces of the assembly's Hilbert space, I will show that qualitatively individuated systems provide the natural decompositions being sought.

So let us consider what algebra of operators we may associate with a qualitatively individuated system. For simplicity I will concentrate on the two-system case. Recall that qualitative individuation is individuation by projectors. So suppose that our two individuation criteria (one for each of the two systems) are represented by the projectors E_α, E_β, each of which acts on the single-system Hilbert space. I require that $E_\alpha \perp E_\beta$, i.e. $E_\alpha E_\beta = E_\beta E_\alpha = 0$, so that none of the two systems is individuated by the other's criterion. (The importance of this condition will soon become clear.) Call the system individuated by E_α the α-system, and the system individuated by E_β, the β-system.

Now consider the subspace of the assembly's Hilbert space

$$\mathcal{M}_\lambda(\alpha, \beta) := \mathrm{ran}(E_\alpha \otimes E_\beta + E_\beta \otimes E_\alpha) \ , \tag{7.38}$$

where $\lambda \in \{s, a\}$ indicates whether the subspace lies in the symmetric or anti-symmetric sector; i.e., whether the assembly consists of bosons or fermions. (Again, this is an assembly of two systems, so there are no paraparticle states.)

In every state in this subspace, and in only these states, every term of the state has a single-system state in the range of E_α and a single-system state in the range of E_β. In the case that E_α and E_β are both one-dimensional, this subspace is one-dimensional, and is spanned by the unique non-GM-entangled state for the symmetry type λ in which one system is maximally specific *à la* E_α and the other is maximally specific *à la* E_β. In general I will say that, for all and only states whose support lies solely in $\mathcal{M}_\lambda(\alpha, \beta)$, the individuation is *successful*, or that it *succeeds*. The condition of success is equivalent to

$$\mathrm{Tr}\left[\rho(E_\alpha \otimes E_\beta + E_\beta \otimes E_\alpha)\right] = 1 \ . \tag{7.39}$$

I now claim that, for both bosonic and fermionic two-system assemblies, the subspace $\mathcal{M}_\lambda(\alpha, \beta)$ may be naturally decomposed into two spaces: one corresponding to the α-system and one corresponding to the β-system.

To prove this, we need to fulfil Zanardi's requirements of completeness and subsystem independence. That is, we need to find two commuting subalgebras $\mathcal{A}_\alpha$ and $\mathcal{A}_\beta$ (one for the α-system and one for the β-system) whose tensor product is isomorphic to the algebra of symmetric operators on $\mathcal{M}_\lambda(\alpha, \beta)$. Since $\mathcal{M}_\lambda(\alpha, \beta)$ is either symmetric or anti-symmetric, the latter is the full set of bounded operators on $\mathcal{M}_\lambda(\alpha, \beta)$, i.e. $\mathcal{B}(\mathcal{M}_\lambda(\alpha, \beta))$.

First of all, we can limit our search for $\mathcal{A}_\alpha$ and $\mathcal{A}_\beta$ to subalgebras of $\mathcal{B}(\mathcal{H})$, where $\mathcal{H}$ is the single-system Hilbert space. This is because our individuation criteria E_α, E_β are single-system projectors, so we expect the Hilbert spaces associated with the α-system and β-system to be no larger than $\mathcal{H}$.

Second, we ought to demand that every operator $A \in \mathcal{A}_\alpha$ commute with the individuation criterion E_α: i.e., $\forall A \in \mathcal{A}_\alpha, [A, E_\alpha] = 0$. (And similarly for the β-system.) The reason is so we do not lose track of the α-system by operating on it with operators from its own algebra. As we shall see, this condition is crucial in order to secure the required algebraic structure for $\mathcal{A}_\alpha$ (and $\mathcal{A}_\beta$).

We may now use Schur's Lemma to establish that, when E_α is not trivial (i.e. not the identity), the representation, on the single-system Hilbert space, of the algebra of operators fulfilling both of these conditions, i.e. $\mathfrak{A} := \{A \in \mathcal{B}(\mathcal{H}) \mid [A, E_\alpha] = 0\}$ must be reducible. I.e., $\mathfrak{A} = \mathfrak{A}_\alpha \oplus \mathfrak{A}'$, where $\mathcal{H} = \mathcal{M} \oplus \mathcal{M}_\perp$, and $\mathcal{M} = \mathrm{ran}(E_\alpha)$ supports a representation of $\mathfrak{A}_\alpha$ and $\mathcal{M}_\perp = \mathrm{ran}(\mathbb{1} - E_\alpha)$ supports a representation of $\mathfrak{A}'$. Therefore, $E_\alpha \mathfrak{A} E_\alpha = \mathfrak{A}_\alpha \oplus \mathbf{0}$. Furthermore, since E_α is the identity on $\mathcal{M}$, $\mathfrak{A}_\alpha$'s representation on $\mathcal{M}$ is *irreducible*; consequently $\mathfrak{A}_\alpha = \mathcal{B}(\mathcal{M})$, the algebra of all bounded linear operators on $\mathcal{M}$. Similar results hold for the β-system.

Let us make the identifications $\mathcal{A}_\alpha = \mathfrak{A}_\alpha \equiv \mathcal{B}(\mathrm{ran}(E_\alpha))$ and $\mathcal{A}_\beta = \mathcal{B}(\mathrm{ran}(E_\beta))$. Then $\mathcal{A}_\alpha$ and $\mathcal{A}_\beta$ commute, since $E_\alpha \perp E_\beta$, and so the representations of $\mathcal{A}_\alpha$ and $\mathcal{A}_\beta$ on $\mathcal{H}$ are disjoint. This satisfies Zanardi's first condition for the decomposition being natural.

It remains to be shown that $\mathcal{A}_\alpha \otimes \mathcal{A}_\beta$ and $\mathcal{B}(\mathcal{M}_\lambda(\alpha, \beta))$ are isomorphic. For this I define a linear map π_λ. It acts on the assembly Hilbert space, its dual space, the algebra of operators on the assembly Hilbert space, and (consequently) matrix elements of such operators. It is defined as follows.

For all $|\phi\rangle \in \mathrm{ran}(E_\alpha), |\chi\rangle \in \mathrm{ran}(E_\beta)$; and all $|\Psi\rangle, |\Phi\rangle \in \mathrm{ran}(E_\alpha) \otimes \mathrm{ran}(E_\beta)$; and all $A \in \mathcal{A}_\alpha, B \in \mathcal{A}_\beta$; and all $P, Q \in \mathcal{A}_\alpha \otimes \mathcal{A}_\beta$; and all $a, b \in \mathbb{C}$:

$$|\phi\rangle \otimes |\chi\rangle \quad \overset{\pi_\lambda}{\longmapsto} \quad \frac{1}{\sqrt{2}} \left(|\phi\rangle \otimes |\chi\rangle \pm |\chi\rangle \otimes |\phi\rangle \right) \tag{7.40}$$

$$\pi_\lambda \left(a|\Psi\rangle + b|\Phi\rangle \right) \quad := \quad a\pi_\lambda(|\Psi\rangle) + b\pi_\lambda(|\Phi\rangle) \tag{7.41}$$

(and similarly for the dual space, and where the '$\pm$' in (7.40) corresponds to whether λ is symmetric or anti-symmetric); and

$$A \otimes B \quad \overset{\pi_\lambda}{\longmapsto} \quad A \otimes B + B \otimes A \tag{7.42}$$

$$\pi_\lambda(aP + bQ) \quad := \quad a\pi_\lambda(P) + b\pi_\lambda(Q) \tag{7.43}$$

$$\pi_\lambda(\langle\Psi|Q|\Phi\rangle) \quad := \quad \pi_\lambda(\langle\Psi|)\pi_\lambda(Q)\pi_\lambda(|\Phi\rangle) \ . \tag{7.44}$$

Then we have, for example, for any $A, C \in \mathcal{A}_\alpha$ and $B, D \in \mathcal{A}_\beta$,

$$\begin{aligned}
\pi_\lambda(A \otimes B)\pi_\lambda(C \otimes D) &:= (A \otimes B + B \otimes A)\,(C \otimes D + D \otimes C) \\
&= AC \otimes BD + BD \otimes AC + AD \otimes BC + BC \otimes AD \\
&= AC \otimes BD + BD \otimes AC \\
&= \pi_\lambda(AC \otimes BD), \qquad (7.45)
\end{aligned}$$

where we use the fact that $AD = BC = 0$, since the pairs A, D and B, C have disjoint representations. Note also that the ranges of C of D are the domains of A and B, respectively. This all relies on our original stipulation that a qualitatively individuated system's algebra commute with its individuation criterion. It follows from all this that $\pi_\lambda(\mathrm{ran}(E_\alpha) \otimes \mathrm{ran}(E_\beta)) = \mathcal{M}_\lambda(\alpha, \beta)$ and $\pi_\lambda(\mathcal{A}_\alpha \otimes \mathcal{A}_\beta) = \mathcal{B}(\mathcal{M}_\lambda(\alpha, \beta))$.[4]

To see that π_λ is an isomorphism, note that it is one-to-one, and that it preserves the matrix elements of all operators in $\mathcal{A}_\alpha \otimes \mathcal{A}_\beta$. Since for all $|\phi_i\rangle, |\phi_k\rangle \in \mathrm{ran}(E_\alpha)$; and all $|\chi_j\rangle, |\chi_l\rangle \in \mathrm{ran}(E_\beta)$; and all $A \in \mathcal{A}_\alpha$ and all $B \in \mathcal{A}_\beta$:

$$\begin{aligned}
\pi_\lambda\left(\langle\phi_i|A|\phi_k\rangle\langle\chi_j|B|\chi_l\rangle\right) &\equiv \pi_\lambda\left(\langle\phi_i| \otimes \langle\chi_j|A \otimes B|\phi_k\rangle \otimes |\chi_l\rangle\right) \\
&= \pi_\lambda(\langle\phi_i| \otimes \langle\chi_j|)\pi_\lambda(A \otimes B)\pi_\lambda(|\phi_k\rangle \otimes |\chi_l\rangle) \\
&= \frac{1}{2}\langle\phi_i|A|\phi_k\rangle\langle\chi_j|B|\chi_l\rangle + \frac{1}{2}\langle\phi_i|B|\phi_k\rangle\langle\chi_j|A|\chi_l\rangle \\
&\quad + \frac{1}{2}\langle\chi_j|A|\phi_k\rangle\langle\phi_i|B|\chi_l\rangle + \frac{1}{2}\langle\chi_j|B|\phi_k\rangle\langle\phi_i|A|\chi_l\rangle \\
&\quad + \frac{1}{2}\langle\phi_i|A|\chi_l\rangle\langle\chi_j|B|\phi_k\rangle + \frac{1}{2}\langle\phi_i|B|\chi_l\rangle\langle\chi_j|A|\phi_k\rangle \\
&\quad + \frac{1}{2}\langle\chi_j|A|\chi_l\rangle\langle\phi_i|B|\phi_k\rangle + \frac{1}{2}\langle\chi_j|B|\chi_l\rangle\langle\phi_i|A|\phi_k\rangle \\
&= \langle\phi_i|A|\phi_k\rangle\langle\chi_j|B|\chi_l\rangle \qquad (7.46)
\end{aligned}$$

(since $A|\chi_j\rangle = A|\chi_l\rangle = B|\phi_i\rangle = B|\phi_k\rangle = 0$). The linearity of π_λ covers all linear combinations of the above, and so all states in $\mathrm{ran}(E_\alpha) \otimes \mathrm{ran}(E_\beta)$ and all operators in $\mathcal{A}_\alpha \otimes \mathcal{A}_\beta$.

Thus the algebra $\mathcal{B}(\mathcal{M}_\lambda(\alpha, \beta))$ has a natural decomposition into commuting

[4] For greater clarity, one can imagine each $A \in \mathcal{A}_\alpha$ flanked on both sides by the projector E_α, and each $B \in \mathcal{A}_\beta$ flanked on both sides by E_β. This is harmless, since $A \equiv E_\alpha A E_\alpha$ and $B \equiv E_\beta B E_\beta$. The above results then manifestly follow from $E_\alpha \perp E_\beta$ and $[A, E_\alpha] = [B, E_\beta] = 0$.

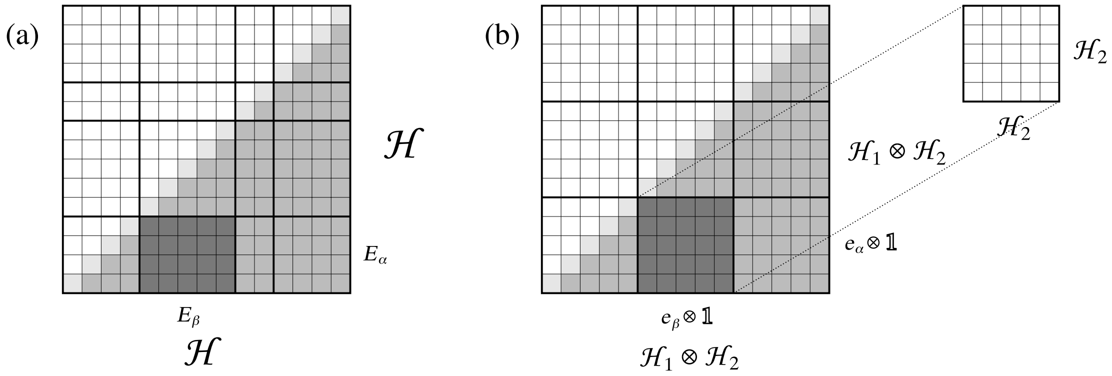

Figure 7.2: (a) The (anti-) symmetric projection of the tensor product of two Hilbert spaces may be decomposed into spaces which exhibit a tensor product structure. (Light grey squares indicate condensed states, which remain under symmetrization but not anti-symmetrization.) (b) If the two Hilbert spaces are decomposed into eigensubspaces of only one degree of freedom, then the "off-diagonal" elements of the decomposition serve as irreps for the full algebra of operators for the other degrees of freedom.

single-system algebras $\mathcal{A}_\alpha$ and $\mathcal{A}_\beta$, corresponding to the systems qualitatively individuated by E_α and E_β, respectively, where $E_\alpha \perp E_\beta$. The result is easily generalised to assemblies of more than two systems (and therefore also to paraparticles).

The above results apply to *any* subspace $\mathcal{B}(\mathcal{M}_\lambda(\alpha, \beta))$ of the assembly Hilbert space, as defined in (7.38), so long as $E_\alpha \perp E_\beta$. I call such subspaces *off-diagonal*, since they contain no even partially condensed states—i.e., states in which the same single-system state is multiply occupied. We may now decompose the *entire* assembly Hilbert space into diagonal and off-diagonal subspaces, and give natural decompositions for each of the off-diagonal subspaces. Each off-diagonal subspace is associated with its own pair of qualitatively individuated systems, and thus behaves, in its own right, like a Hilbert space for an assembly of distinguishable systems; cf. Figure 7.2(a).

In more detail: We may decompose the single-system Hilbert space $\mathcal{H}$ using a complete family of projectors $\{E_i\}, \sum_i E_i = \mathbb{1}$:

$$\mathcal{H} = \left(\sum_i E_i\right)\mathcal{H} = \bigoplus_i E_i(\mathcal{H}) =: \bigoplus_i \mathfrak{h}_i \tag{7.47}$$

Then, with $\mathcal{S}_\lambda$ the appropriate symmetry projector (boson, fermion, etc.), the assembly Hilbert space is

$$\begin{aligned}
\mathcal{S}_\lambda(\mathcal{H}\otimes\mathcal{H}) &= \mathcal{S}_\lambda\left[\left(\bigoplus_i \mathfrak{h}_i\right)\otimes\left(\bigoplus_i \mathfrak{h}_i\right)\right] && (7.48)\\
&= \mathcal{S}_\lambda\left[\bigoplus_i (\mathfrak{h}_i\otimes\mathfrak{h}_i)\oplus\bigoplus_{i<j}[(\mathfrak{h}_i\otimes\mathfrak{h}_j)\oplus(\mathfrak{h}_j\otimes\mathfrak{h}_i)]\right] && (7.49)\\
&= \bigoplus_i \mathcal{S}_\lambda(\mathfrak{h}_i\otimes\mathfrak{h}_i)\oplus\bigoplus_{i<j}\mathcal{S}_\lambda[(\mathfrak{h}_i\otimes\mathfrak{h}_j)\oplus(\mathfrak{h}_j\otimes\mathfrak{h}_i)] && (7.50)\\
&= \bigoplus_i \mathcal{S}_\lambda(\mathfrak{h}_i\otimes\mathfrak{h}_i)\oplus\bigoplus_{i<j}\pi_\lambda(\mathfrak{h}_i\otimes\mathfrak{h}_j) && (7.51)
\end{aligned}$$

where π_λ is as defined above. Each off-diagonal subspace $\pi_\lambda(\mathfrak{h}_i\otimes\mathfrak{h}_j)$ is isomorphic to the un-symmetrized assembly Hilbert space $\mathfrak{h}_i\otimes\mathfrak{h}_j$, and so the former inhererits all of the features of the latter. One such example is that familiar results, such as Bell's Theorem (1964) and Gisin's Theorem (1991), now apply, in the appropriate subspaces, for assemblies of qualitatively individuated systems.[5] This fulfils my promise, made in comment (ii)(b) at the start of Section 7.1.2.

A class of instances of qualitative individuation that is of particular interest arises when the single-system Hilbert space $\mathcal{H}$ represents more than one degree of freedom. In this case, if the individuation criteria E_α, E_β apply to less than the full degrees of freedom, then the full algebra of linear bounded operators on the remaining degrees of freedom is available to the qualitatively individuated systems.

For simplicity, suppose that $\mathcal{H}$ represents two degrees of freedom; i.e., $\mathcal{H} = \mathcal{H}_1\otimes\mathcal{H}_2$. Now let us choose the individuation criteria $E_\alpha = e_\alpha\otimes\mathbb{1}, E_\beta = e_\beta\otimes\mathbb{1}$,

[5] Provided, of course, that $\dim(\mathfrak{h}_i), \dim(\mathfrak{h}_j) \geqslant 2$. Note also that there is no analogue of Gisin's Theorem for paraparticles, since paraparticle states only arise for three or more systems, and Gisin's Theorem cannot be extended beyond the two-system case. (See Żukowski *et al* (2002) for more details).

where e_α and e_β act on $\mathcal{H}_1$ and $\mathbb{1}$ is the identity on $\mathcal{H}_2$. From the results above, it follows that $\mathcal{A}_\alpha = \mathcal{B}(\text{ran}(E_\alpha)) = \mathcal{B}(\text{ran}(e_\alpha)) \otimes \mathcal{B}(\mathcal{H}_2)$ and $\mathcal{A}_\beta = \mathcal{B}(\text{ran}(E_\beta)) = \mathcal{B}(\text{ran}(e_\beta)) \otimes \mathcal{B}(\mathcal{H}_2)$. Thus the full algebra $\mathcal{B}(\mathcal{H}_2)$ is available to both qualitatively individuated systems.[6] (Cf. Figure 7.2(b).)

To conclude this Subsection, I will say something briefly about Huggett and Imbo's (2009, pp. 313) recent claim that it is not necessary to impose the Indistinguishability Postulate (IP) on systems with identical intrinsic (i.e. state-independent) properties. This is because, they claim, systems may be distinguished according to their 'trajectories' (i.e. single-system states). If they are correct, this would entail that factorism is, after all, a viable interpretative position for such systems—so long as we understand factor Hilbert space labels as representing these trajectories (just as, in the case of distinguishable systems, we use factor Hilbert space labels to represent distinct state-*independent* properties of the systems).

The results of this Subsection show that Huggett and Imbo are partly correct. In my jargon: they are right that an un-symmetrised Hilbert space is an equally adequate (since isomorphic) means to represent an assembly of qualitatively individuated systems—*for those states in which the individuation criteria succeed*; and that therefore there is no *practical* need to impose IP, or, therefore, to repudiate factorism *when representing those states and those states alone*. But they are wrong to claim that it is not necessary to impose IP to represent *all of the available states* for systems with identical intrinsic properties. For the isomorphism result above, on which Huggett and Imbo's claim depends, holds *only for the appropriate off-diagonal subspace*. As soon as the assembly's state has components that lie outside of this subspace, the isomorphism breaks down.

I must emphasise too that the breakdown of isomorphism outside of the relevant subspace does not just mean that, for states outside this subspace, the two formalisms yield conflicting empirical claims; empirical claims that confirm IP. Rather, the quasi-factorist formalism ceases to make physical *sense* for states outside of the relevant subspace. For, outside of this subspace, the systems no

[6] This case corresponds to Huggett and Imbo's (2009, pp. 315-6) 'approximately distinguishable' systems. It also corresponds the the example I presented at the end of Stage E of Section 7.1.2, for which $e_\alpha = |R\rangle\langle R|$ and $e_\beta = |L\rangle\langle L|$.

longer occupy the states upon which their individuation—and therefore the entire quasi-factorist formalism—was based.

Huggett and Imbo mistakenly suppose that imposing IP prevents one from qualitatively individuating systems. (As Huggett and Imbo (2009, p. 315) put it: 'IP $\Rightarrow$ trajectory indistinguishability'.) But that assumes what I deny: namely, factorism. Without factorism, we can agree with Huggett and Imbo that systems may be qualitatively individuated, *without* contravening IP. Moreover: without factorism but with IP, we may represent *all* of the states available to the assembly, without fear that our representational apparatus will break down.

7.2.2 Russellian vs. Strawsonian approaches to individuation

All of the results of the previous Subsection apply only to 'the relevant' subspace of the assembly Hilbert space. This is the subspace for which the individuation criteria for the systems succeeds; i.e. for which the projector

$$\mathcal{E}(\alpha, \beta) := E_\alpha \otimes E_\beta + E_\beta \otimes E_\alpha \tag{7.52}$$

has expectation value 1. What about states for which individuation does not succeed? The question is important, since we want a procedure for calculating expectation values of quantities which belong to the joint algebra of the qualitatively individuated quantities; and we want that procedure to be as general as possible.

The way one proceeds depends on one's stance toward reference failure for individuation criteria. I see two equally acceptable routes, which may be associated (perhaps tenuously) with the classic debate over reference failure for definite descriptions. With a little poetic licence, I call the two routes *Russellian* and *Strawsonian.*

The Russellian route (cf. Russell 1905) takes the claim of success of the individuation criteria E_α and E_β to be an implicit tag-along claim *in addition to* any explicit claim which implements those criteria. Thus the expectation value for any $A \in \mathcal{B}(\mathcal{M}_\lambda(\alpha, \beta))$ is given by $\mathrm{Tr}(\mathcal{E}(\alpha, \beta)\rho\mathcal{E}(\alpha, \beta)A)$, which uses the usual quantum mechanical specifications for a conjunction. But A commutes with $\mathcal{E}(\alpha,\beta)$, since it is a sum of products of single-system quantities, each of which commutes with E_α and E_β. So we may simplify to $\mathrm{Tr}(\rho\mathcal{E}(\alpha,\beta)A)$.

The Strawsonian route (cf. Strawson 1950) instead takes the joint success of the individuation criteria E_α and E_β as a *presupposition* of any claim which uses that strategy. Therefore any expectation values calculated under the presupposition of the success of $\mathcal{E}(\alpha,\beta)$ must be renormalized by conditionalizing on that success. This is done using the usual Lüder rule

$$\rho \quad \mapsto \quad \rho_{\alpha\beta} := \frac{\mathcal{E}(\alpha,\beta)\rho\mathcal{E}(\alpha,\beta)}{\mathrm{Tr}(\rho\mathcal{E}(\alpha,\beta))} \ . \tag{7.53}$$

The expectation value of any quantity $A \in \mathcal{B}(\mathcal{M}_\lambda(\alpha,\beta))$ is then given simply by $\mathrm{Tr}(\rho_{\alpha\beta}A)$.

Note that conditionalization requires that $\mathrm{Tr}(\rho\mathcal{E}(\alpha,\beta)) > 0$, which means that the state must have *some* terms for which individuation succeeds. The fact that $\mathrm{Tr}(\rho_{\alpha\beta}A)$ is undefined when $\mathrm{Tr}(\rho\mathcal{E}(\alpha,\beta)) = 0$ meshes rather nicely with Strawson's famous claim that statements containing failed definite descriptions do not possess a truth value.

It will have been noted that the difference between the Russellian and Strawsonian routes for expectation values lies only in the multiplicative factor $\frac{1}{\mathrm{Tr}(\rho\mathcal{E}(\alpha,\beta))}$. A point in favour of the Strawsonian approach is that the identity operator indexed to the two qualitatively individuated systems, $\pi_\lambda(\mathbb{1}\otimes\mathbb{1})$, has expectation value 1 for all normalizable states, while under the Russellian route the identity's expectation value is equal to $\mathcal{E}(\alpha,\beta)$'s expectation value. A point in favour of the Russellian approach is that expectation values may be defined for *all* states. On this approach, if the assembly has *no* terms for which the individuation succeeds, then expectation value for every $A \in \mathcal{B}(\mathcal{M}_\lambda(\alpha,\beta))$ is zero.

7.2.3 Qualitatively individuated systems on their own

In this Subsection, I turn away from the problem of *completely* decomposing an assembly into natural constituent systems, and turn instead to the problem of picking out a *single* constituent system from the assembly. I seek a means to calculate expectation values for quantities associated with a single qualitatively individuated system, whose individuation criterion we may choose.

The way I will proceed is inspired in part by the Strawsonian approach to individuation in Section 7.2.2. The main idea, there and here, is to *conditionalize* upon the success of the individuation. As usual, I work, for the sake of simplicity, in the $N = 2$ case (unless otherwise stated); the generalization to $N > 2$ will be obvious.

We begin with a chosen individuation criterion, a projector E_α, which acts on the single-system Hilbert space. Then it may be checked that the operator

$$n_\alpha := E_\alpha \otimes \mathbb{1} + \mathbb{1} \otimes E_\alpha \tag{7.54}$$

is a number operator for the two-system assembly's Hilbert space. That is, it "counts" the number of systems which are picked out by E_α.

I now define a linear map π_α from the single-system algebra $\mathcal{B}(\mathcal{H})$ into a particular subalgebra of the assembly's algebra. This subalgebra will be the operators which are associated with the α-system. I define

$$\pi_\alpha(A) := E_\alpha A E_\alpha \otimes \mathbb{1} + \mathbb{1} \otimes E_\alpha A E_\alpha \ . \tag{7.55}$$

(Note that, if $A = E_\alpha A E_\alpha$, then π_α is just the *symmetrizer* for A.)

I now claim that the expectation value for any single-system quantity A, associated with the α-system is

$$\langle A \rangle_\alpha := \frac{\langle \pi_\alpha(A) \rangle}{\langle n_\alpha \rangle} \ . \tag{7.56}$$

I will establish this claim by considering a few examples.

1. The state of the assembly $|\psi\rangle = \frac{1}{\sqrt{2}}(|\alpha\rangle \otimes |\beta\rangle \pm |\beta\rangle \otimes |\alpha\rangle)$, where $E_\alpha|\alpha\rangle = |\alpha\rangle$ and $E_\alpha|\beta\rangle = 0$, and $Q|\alpha\rangle = q|\alpha\rangle$. Then $\langle n_\alpha \rangle = 1$ and $\langle \pi_\alpha(Q) \rangle = q$; so $\langle Q \rangle_\alpha = q$. That is, the system individuated by E_α takes as its expectation for Q the value q, associated with the state $|\alpha\rangle$, for which individuation succeeds

(i.e., the state that it is in the range of E_α). (Indeed, the α-system is in an eigenstate for Q, since $\langle Q^2 \rangle_\alpha = q^2$.)

2. $|\psi\rangle = c_1 \frac{1}{\sqrt{2}} (|\alpha_1\rangle \otimes |\beta_1\rangle \pm |\beta_1\rangle \otimes |\alpha_1\rangle) + c_2 \frac{1}{\sqrt{2}} (|\alpha_2\rangle \otimes |\beta_2\rangle \pm |\beta_2\rangle \otimes |\alpha_2\rangle)$, where $|c_1|^2 + |c_2|^2 = 1$; and for all $i = 1, 2$: $E_\alpha |\alpha_i\rangle = |\alpha_i\rangle$ and $E_\alpha |\beta_i\rangle = 0$, and $Q|\alpha_i\rangle = q_i |\alpha_i\rangle$. Then $\langle n_\alpha \rangle = 1$ and $\langle \pi_\alpha(Q) \rangle = |c_1|^2 q_1 + |c_2|^2 q_2$; so $\langle Q \rangle_\alpha = |c_1|^2 q_1 + |c_2|^2 q_2$. That is, the system individuated by E_α takes as its expectation for Q the average for all single-system states $|\alpha_i\rangle$, for which the individuation succeeds. The weights for this average are given by the relative amplitudes of the non-GM-entangled terms.

3. $|\psi\rangle = |\alpha\rangle \otimes |\alpha\rangle$, for $|\alpha\rangle$ as above. Then $\langle n_\alpha \rangle = 2$ and $\langle \pi_\alpha(Q) \rangle = 2q$; so $\langle Q \rangle_\alpha = q$. In this case, E_α individuates *two* systems, and the expectation (indeed, eigenvalue) for Q for both of them is q.

4. $|\psi\rangle = \frac{1}{\sqrt{2}} (|\alpha_1\rangle \otimes |\alpha_2\rangle \pm |\alpha_2\rangle \otimes |\alpha_1\rangle)$, for $|\alpha_1\rangle, |\alpha_2\rangle$ as above. Then $\langle n_\alpha \rangle = 2$ and $\langle \pi_\alpha(Q) \rangle = q_1 + q_2$; so $\langle Q \rangle_\alpha = \frac{1}{2}(q_1 + q_2)$. In this case, E_α again individuates two systems, one whose expectation value for Q is q_1, and one whose value is q_2; thus we take the average. However, the weights for this average are *not* given, as above, by relative amplitudes for non-GM-entangled terms; (the entire state is non-GM-entangled). Rather, they are given by the relative *frequency*, in a single non-GM-entangled state, of each single-system state for which individuation succeeds.

5. $|\psi\rangle = c_1 \mathcal{S}_\lambda (|\alpha_1\rangle \otimes |\alpha_2\rangle \otimes |\beta_1\rangle) + c_2 \mathcal{S}_\lambda (|\alpha_3\rangle \otimes |\beta_1\rangle \otimes |\beta_2\rangle)$, where $\mathcal{S}_\lambda$ is the (anti-) symmetrizer on the assembly Hilbert space, and the single-system states are defined as before. (So $N = 3$; and for simplicity we set aside paraparticles.) Then $\langle n_\alpha \rangle = 2|c_1|^2 + |c_2|^2$ and $\langle \pi_\alpha(Q) \rangle = |c_1|^2(q_1 + q_2) + |c_2|^2 q_3$; so $\langle Q \rangle_\alpha = \frac{|c_1|^2(q_1+q_2)+|c_2|^2 q_3}{2|c_1|^2+|c_2|^2}$. In this case, the weights for the average are determined *jointly* by relative amplitudes *and* relative frequencies. If $|c_1| = |c_2|$, then $\langle Q \rangle_\alpha = \frac{1}{3}(q_1 + q_2 + q_3)$; thus each of the three states in the range of E_α are afforded equal weight, whether or not they belong to the same non-GM-entangled term.

Thus my claim—that the expectation value of any single-system quantity A, for the system qualitatively individuated by E_α, is given by (7.56)—yields the right results, at least for cases 1 to 4. Case 5 seems to me less clear cut, since one might favour a different way to calculate statistical weights from the relative amplitudes and relative frequencies. However, I submit, there are no clear intuitions to rely on in this case, and I can see no objection to the way given by (7.56).

It may have been noticed that π_α is *not* an isomorphism between the single-system algebra $\mathcal{B}(\mathcal{H})$—or indeed any subalgebra thereof—and the range of π_α. It is not even a homomorphism. For it may be checked that $\pi_\alpha(AB) \neq \pi_\alpha(A)\pi_\alpha(B)$ does not hold, even for all those $A, B \in \mathcal{B}(\mathcal{H})$ that commute with E_α.

This is not an objection to (7.56), and should come as no surprise. For there are states of the assembly in which E_α fails to individuate a *unique* system (cf. examples 3-5, above). In these states, we should not expect that $\pi_\alpha(AB) = \pi_\alpha(A)\pi_\alpha(B)$. To perform the operation B, followed by A, on a given α-system (corresponding to $\pi_\alpha(AB)$) relies on a *re-identification* of that system (and that system alone) *after* we have operated with B. But the individuation criterion E_α cannot be guaranteed to pick out that very same system, if more than one system is picked out by E_α.[7] On the other hand, it may be checked that, for any two states $|\psi\rangle$ of the assembly that are eigenstates of n_α with eigenvalue 1—i.e., for all states in which exactly *one* system is individuated by E_α—we have $\langle\psi|\pi_\alpha(AB)|\psi\rangle = \langle\psi|\pi_\alpha(A)\pi_\alpha(B)|\psi\rangle$, as expected.

Thus we have a recipe for calculating the expectation value of any single-system quantity for a qualitatively individuated system or systems. It remains for me to give a general prescription for calculating the reduced density operator for such a system. We require that the reduced density operator ρ_α satisfy the condition that, for all $A \in \mathcal{B}(\mathcal{H})$: $\mathrm{Tr}(\rho_\alpha A) = \langle A\rangle_\alpha$, as given in (7.56). We know from Gleason's Theorem that a unique such operator exists.

As usual, I work by analogy with the case of "distinguishable" systems. The usual prescription for the reduced density operator of a constituent system, say

[7] It is worth emphasising that the possibility of multiple α-systems does not arise *only* for bosons, since E_α need not be a one-dimensional projector.

the kth, of the assembly is (with ρ the state of the assembly)

$$\rho_k := \mathrm{Tr}_k\left(\rho\right), \tag{7.57}$$

where Tr_k denotes a partial trace over all but the kth factor Hilbert space. Now this prescription is obviously no use to anti-factorists; but an equivalent formulation to (7.57) exists that will be of far more use. First we choose a complete orthobasis $\{|\phi_i\rangle\}$ for the single-system Hilbert space $\mathcal{H}$. Then

$$\rho_k := \sum_{i,j} \mathrm{Tr}\left(\rho|\phi_j\rangle\langle\phi_i|_k\right)|\phi_i\rangle\langle\phi_j| \tag{7.58}$$

where

$$|\phi_j\rangle\langle\phi_i|_k := \bigotimes^{k-1} \mathbb{1} \otimes |\phi_j\rangle\langle\phi_i| \otimes \bigotimes^{N-k} \mathbb{1} \tag{7.59}$$

and we now perform a *full* trace on the assembly Hilbert space.

We may adapt (7.58) for anti-factorist, qualitatively individuated systems in the following way. First, we replace each operator $|\phi_j\rangle\langle\phi_i|_k$, which is indexed to a factor Hilbert space, with $\pi_\alpha(|\phi_j\rangle\langle\phi_i|)$, as given in (7.55). And second, we "conditionalize" by dividing by $\langle n_\alpha\rangle = \mathrm{Tr}(\rho n_\alpha)$. Thus

$$\rho_\alpha = \frac{1}{\langle n_\alpha\rangle}\sum_{i,j} \mathrm{Tr}\left[\rho\ \pi_\alpha(|\phi_j\rangle\langle\phi_i|)\right]|\phi_i\rangle\langle\phi_j| \tag{7.60}$$

Written out in full, and for any N, we have

$$\rho_\alpha = \frac{\displaystyle\sum_{i,j}|\phi_i\rangle\langle\phi_j|\ \mathrm{Tr}\left[\rho\left(\sum_{k=1}^{N}\bigotimes^{k-1}\mathbb{1}\otimes E_\alpha|\phi_j\rangle\langle\phi_i|E_\alpha\otimes\bigotimes^{N-k}\mathbb{1}\right)\right]}{\mathrm{Tr}\left[\rho\left(\displaystyle\sum_{k=1}^{N}\bigotimes^{k-1}\mathbb{1}\otimes E_\alpha\otimes\bigotimes^{N-k}\mathbb{1}\right)\right]} \tag{7.61}$$

It may then be shown (as required) that for any $A \in \mathcal{B}(\mathcal{H})$, $\mathrm{Tr}(\rho_\alpha A) = \langle A\rangle_\alpha$.

For this, let $\{|\xi_i\rangle\}$ be a complete eigenbasis for A, where $A|\xi_i\rangle = a_i|\xi_i\rangle$. Then

$$\begin{aligned}
\mathrm{Tr}(\rho_\alpha A) &= \frac{1}{\langle n_\alpha\rangle}\sum_{i,j,k}\mathrm{Tr}\left[\rho\ \pi_\alpha(|\xi_j\rangle\langle\xi_i|)\right]\langle\xi_k|\xi_i\rangle\langle\xi_j|A|\xi_k\rangle && (7.62)\\
&= \frac{1}{\langle n_\alpha\rangle}\sum_{i,j,k}a_k\mathrm{Tr}\left[\rho\ \pi_\alpha(|\xi_j\rangle\langle\xi_i|)\right]\delta_{ki}\delta_{jk} && (7.63)\\
&= \frac{1}{\langle n_\alpha\rangle}\sum_{k}\mathrm{Tr}\left[\rho\ a_k\ \pi_\alpha(|\xi_k\rangle\langle\xi_k|)\right] && (7.64)\\
&= \frac{1}{\langle n_\alpha\rangle}\mathrm{Tr}\left[\rho\ \pi_\alpha\left(\sum_k a_k\ |\xi_k\rangle\langle\xi_k|\right)\right] && (7.65)\\
&= \frac{1}{\langle n_\alpha\rangle}\mathrm{Tr}\left(\rho\ \pi_\alpha\left(A\right)\right) && (7.66)\\
&=: \langle A\rangle_\alpha\ . && (7.67)
\end{aligned}$$

Remember that ρ_α as given in (7.60) and (7.61) is the *average* state of any system individuated by E_α. So long as the state ρ is an eigenstate of n_α with eigenvalue 1—or even a superposition of $n_\alpha = 0$ and $n_\alpha = 1$ eigenstates—then ρ_α yields the state of *the* α-system. However, if ρ contains eigenstates with $n_\alpha > 1$, then the interpretation of ρ_α as the state of *the* α-system can no longer be sustained, since in those terms we effectively average *all* systems picked out by E_α.

For this reason, I note as an aside that it seems appropriate to favour something like a Lewisian *counterpart theory* for claims involving qualitatively individuated systems, over any doctrine of "trans-state" identity. In Lewis's theory, it is perfectly in order for an object to have multiple counterparts in some possible worlds, just as there may be multiple α-systems in some states. Additionally, the obvious freedom in the choice of the individuating criterion E_α meshes well with other flexibilities in Lewis's counterpart relation (cf. Lewis 1968, pp. 115-6). This suggestion warrants a much more in-depth treatment, but there is no space to do that here.

To conclude this Section, I note two important limiting cases of Equation (7.60). The first is when we are maximally discriminating in our individuation; i.e., where

E_α is a one-dimensional projector. Let $|\alpha\rangle$ be the state for which $E_\alpha|\alpha\rangle = |\alpha\rangle$. Then, so long as $\langle n_\alpha\rangle > 0$, $\rho_\alpha = |\alpha\rangle\langle\alpha|$, which is to be expected.

The second limiting case lies at the other extreme, in which we individuate with maximum *in*discriminateness, i.e. with $E_{\boldsymbol{\alpha}} = \mathbb{1}$. In this case $\langle n_\alpha\rangle = N$ and $\pi_\alpha(A) = \sum_{k=1}^{N} \bigotimes^{k-1} \mathbb{1} \otimes A \otimes \bigotimes^{N-k} \mathbb{1}$; so

$$
\begin{aligned}
\rho_\alpha &= \frac{1}{N}\sum_{i,j} |\phi_i\rangle\langle\phi_j| \ \mathrm{Tr}\left[\rho\left(\sum_{k=1}^{N}\bigotimes^{k-1}\mathbb{1}\otimes|\phi_j\rangle\langle\phi_i|\otimes\bigotimes^{N-k}\mathbb{1}\right)\right] && (7.68)\\
&= \frac{1}{N}\sum_{k=1}^{N}\sum_{i,j} |\phi_i\rangle\langle\phi_j| \ \mathrm{Tr}\left[\rho\left(\bigotimes^{k-1}\mathbb{1}\otimes|\phi_j\rangle\langle\phi_i|\otimes\bigotimes^{N-k}\mathbb{1}\right)\right] && (7.69)\\
&= \frac{1}{N}\sum_{k=1}^{N}\sum_{i,j} |\phi_i\rangle\langle\phi_j| \ \mathrm{Tr}\left[\rho\left(|\phi_j\rangle\langle\phi_i|_k\right)\right] \quad \text{(from (7.59))} && (7.70)\\
&= \frac{1}{N}\sum_{k=1}^{N}\rho_k \qquad \text{(from (7.58))}. && (7.71)
\end{aligned}
$$

So with maximum indiscriminateness ρ_α is the "average" of the standard reduced density operators obtained by partial tracing. In the case of indistinguishable systems (i.e. when IP is imposed), we of course have $\rho_1 = \rho_2 = \ldots = \rho_k =: \overline{\rho}$, in which case $\rho_\alpha = \overline{\rho}$. This vindicates my claim in Section 6.4.3 that, in the context of indistinguishable systems, standard reduced density operators obtained by partial tracing codify only the state of the average system, and not the state of any particular system.

7.3 Qualitative individuation over time

In this Section I take a brief look at individuation criteria which evolve *over time*. The investigation here will be all too brief, but will hopefully give a flavour of the direction of future investigation.

Let $\mathcal{E} := E_\alpha\otimes E_\beta + E_\beta\otimes E_\alpha$ represent our individuation criteria for two particles at time $t = 0$. The obvious way to turn this into evolving individuation criteria is to use the usual Heisenberg prescription for time-dependent quantities. Thus, if $U(t) = e^{-it\frac{H}{\hbar}}$ is the evolution operator for the assembly, with Hamiltonian H, then we may define time-dependent operator

$$\mathcal{E}(t) = U(t)\mathcal{E}U^\dagger(t) \tag{7.72}$$

The expectation value of $\mathcal{E}(t)$ is a constant of the motion:

$$\langle\mathcal{E}(t)\rangle_t = \mathrm{Tr}(U(t)\rho U^\dagger(t)U(t)\mathcal{E}U^\dagger(t)) = \mathrm{Tr}(\rho\mathcal{E}) = \langle\mathcal{E}\rangle_0, \tag{7.73}$$

so if the individuation criteria succeed at $t = 0$, then the dynamics preserves this successfulness over time.

This proposal may be seen as analogous to individuation procedures in the classical mechanics of point particles, where we quotient by the symmetric group (Belot 2001). Working in the reduced phase space, by evolving any equivalence class of system points along the Hamiltonian flow we achieve natural trans-temporal identifications for the point particles by demanding continuous trajectories for each particle.

However, returning to the quantum case, we cannot guarantee, unlike in the classical case, that the time-evolute of $\mathcal{E}$ has the right features to count as a pair of individuation criteria for the two particles. For this it is necessary and sufficient that

$$U(t)\left(E_\alpha \otimes E_\beta + E_\beta \otimes E_\alpha\right)U^\dagger(t) = E_\alpha(t) \otimes E_\beta(t) + E_\beta(t) \otimes E_\alpha(t), \tag{7.74}$$

where, at any time t, $E_\alpha(t) \perp E_\beta(t)$. That is, it is necessary and sufficient that $\mathcal{E}$'s evolution may be expressed in terms of a piecemeal evolution of the single-system projectors E_α and E_β.

Too see that condition (7.74) does not hold generally, one need only consider an evolution that takes a heterogeneous state of two bosons at $t = T_i$ to a product state at $t = T_f$. In this case, at time T_i, $E_\alpha(T_i) \perp E_\beta(T_i)$; but by $t = T_f$, we have $E_\alpha(T_f) = E_\beta(T_f)$, which contradicts the requirement of Section 7.2.1 that individuation criteria for distinct systems be orthogonal. I know of no general result, which gives conditions on either the evolution $U(t)$ or the assembly's state, for the satisfaction of (7.74).

However, it is easily seen that condition (7.74) is satisfied if the dynamical evolution is factorizable; i.e. if $U(t) = W(t) \otimes W(t)$, where $W(t)$ is a continuous one-parameter family of unitaries on the single-system Hilbert space $\mathcal{H}$. Under this evolution, the purity of single-system states is preserved over time.

But factorizable evolutions present a problem. For there may be no *unique* pair of time-dependent individuation criteria $E_\alpha(t), E_\beta(t)$ for which condition (7.74) is satisfied. This result is as bad as there being no such pair, if our goal is to find uniquely natural trans-temporal identity conditions. (I return to this point in Section 8.2.3.)

The fact that uniqueness is not guaranteed for factorizable evolutions is illustrated by the following simple example. Consider the singlet state for two spin-$\frac{1}{2}$ fermions:

$$|\psi\rangle = \frac{1}{\sqrt{2}}\left(|{\uparrow}\rangle \otimes |{\downarrow}\rangle - |{\downarrow}\rangle \otimes |{\uparrow}\rangle\right) \ , \tag{7.75}$$

and the trivial evolution $U(t) = e^{i\gamma(t)}$, where $\gamma : \mathbb{R} \to \mathbb{R}$ is any real-valued function of the time for which $\gamma(0) = 0$, so that $|\psi(t)\rangle = e^{i\gamma(t)}|\psi\rangle$. And suppose that at time $t = 0$ we individuate two systems using the projectors $|\uparrow\rangle\langle\uparrow|$ and $|\downarrow\rangle\langle\downarrow|$. (Remember from Stage C of Section 7.1.2 that non-GM-entangled fermionic states suffer a basis arbitrariness, so that any pair of orthogonal projectors associated with the same two-dimensional subspace spanned by $|{\uparrow}\rangle$ and $|{\downarrow}\rangle$ would have been equally good individuators.) We now seek time-dependent individuation criteria for these two systems. Condition (7.74) suggests the time-dependent projectors $W(t)|\uparrow\rangle\langle\uparrow|W^\dagger(t)$ and $W(t)|\downarrow\rangle\langle\downarrow|W^\dagger(t)$, where $U(t) = W(t) \otimes W(t)$ for some unitary $W(t)$ on $\mathcal{H}$. *But $W(t)$ is under-determined by this requirement.*

We make use of the group isomorphism $U(2) \otimes U(2) \cong U(3) \oplus U(1)$. Let $w(t)$ be *any* continuous one-parameter family of unitary 2×2 matrices

$$w(t) = \begin{pmatrix} \alpha(t)e^{i\phi(t)} & \beta(t)e^{i\phi(t)} \\ -\beta^*(t) & \alpha^*(t) \end{pmatrix} \tag{7.76}$$

where $|\alpha(t)|^2 + |\beta(t)|^2$. Then

$$w(t){\otimes}w(t) = \begin{pmatrix} \alpha^2(t)e^{2i\phi(t)} & \alpha(t)\beta(t)e^{2i\phi(t)} & \alpha(t)\beta(t)e^{2i\phi(t)} & \beta^2(t)e^{2i\phi(t)} \\ -\alpha(t)\beta^*(t)e^{i\phi(t)} & |\alpha(t)|^2 e^{i\phi(t)} & -|\beta(t)|^2 e^{i\phi(t)} & \alpha^*(t)\beta(t)e^{i\phi(t)} \\ -\alpha(t)\beta^*(t)e^{i\phi(t)} & -|\beta(t)|^2 e^{i\phi(t)} & |\alpha(t)|^2 e^{i\phi(t)} & \alpha^*(t)\beta(t)e^{i\phi(t)} \\ [\beta^*(t)]^2 & -\alpha^*(t)\beta^*(t) & -\alpha^*(t)\beta^*(t) & [\alpha^*(t)]^2 \end{pmatrix} \tag{7.77}$$

With a suitable change of basis this becomes

$$w(t) \otimes w(t) = \begin{pmatrix} \alpha^2(t)e^{2i\phi(t)} & \alpha(t)\beta(t)e^{2i\phi(t)} & \beta^2(t)e^{2i\phi(t)} & 0 \\ -\alpha(t)\beta^*(t)e^{i\phi(t)} & (|\alpha(t)|^2 - |\beta(t)|^2)e^{i\phi(t)} & \alpha^*(t)\beta(t)e^{i\phi(t)} & 0 \\ [\beta^*(t)]^2 & -\alpha^*(t)\beta^*(t) & [\alpha^*(t)]^2 & 0 \\ 0 & 0 & 0 & e^{i\phi(t)} \end{pmatrix} \tag{7.78}$$

where the change of basis is such that $w(t) \otimes w(t)$ is decomposed into symmetric components (the 3×3 matrix at top-left) and anti-symmetric components (the c-number at bottom-right).

Therefore, the restriction of $w(t){\otimes}w(t)$ to the anti-symmetric sector $\mathcal{A}(\mathbb{C}^2{\otimes}\mathbb{C}^2)$, spanned by $|\psi\rangle$, is $e^{i\phi(t)}$. So if we make the identification $W(t) = w(t)$, we have $U(t) = e^{i\phi(t)}$ on $\mathcal{A}(\mathbb{C}^2 \otimes \mathbb{C}^2)$. Thus $W(t)$ may be *any* unitary 2×2 matrix of the form given in Equation (7.76), subject only to the requirement that $\phi(t) = \gamma(t)$.

The upshot is that the trans-temporal identity conditions for constituent systems in the trajectory $|\psi(t)\rangle$ are subject to somewhat the same problems as face the parts of a rotating sphere of uniform, continuous matter (cf. e.g. Zimmerman (1998), Butterfield (2006b)). In both cases no uniquely natural trans-temporal identity conditions (which, in the case of the rotating sphere, would determine its angular velocity) appears to be available.

To sum up: Factorizable evolutions give favourable conditions under which trans-temporal individuation criteria for constituent systems may be defined; but the conditions are *too* favourable: the criteria may not be unique. At the other end of the spectrum, non-factorizable evolutions, which do not preserve non-GM-entanglement, may not even allow time-dependent individuation criteria to be defined. For there is no guarantee that the condition (7.74) can be satisfied.

This concludes our tour of the technicalia associated with a general anti-factorist approach to permutation-invariant quantum mechanics. The foregoing results prompt a new understanding of particles in quantum mechanics, but we are far from a complete picture. We now turn our attention to the attempt to put some metaphysical meat on these mathematical bones.

Chapter 8

Against varietism

My claim that factorism makes an interpretative error similar to the reification of the average person prompts the question: If the factorists' 'particles' are statistical constructs, *of what* are they statistical constructs? With the results of Section 7 in hand, we are now in a position to make positive attempts at an answer. In the first Section of this Chapter, I will outline the first such attempt, a doctrine I call *varietism*, because it attempts to "transfer" the qualitative variety in an assembly's state to the monadic properties of its constituent particles. Then (Section 8.2), I will assess the degree to which varietism satisfies the desiderata for the concept of particle outlined in Section 5.1, and argue that it fares well. However, varietism suffers from a problem that may be fatal: I outline this problem in Section 8.3.

8.1 Varietism defined

Recall that I endorse factorism in the case of distinguishable systems. So, as usual, I proceed by analogy with that case. (Of course, factorism's associating particles with factor Hilbert spaces was a strategy that proceeded by analogy with the case of distinguishable systems! But extending a different feature to the case of indistinguishable particles will give us a different result.)

For an assembly of distinguishable systems, a state is non-entangled iff it is separable. And, by definition, a separable state is one in which the constituent systems are in pure states; in Section 7.1.1, we called such systems *maximally specific*.

This general rule, which applies to any distinguishable quantum systems whatever, can be adapted to the specific case of an assembly of distinguishable *particles* by ensuring that the assembly's Hilbert space, and associated algebra of operators, is of the right kind. A particle is then just a consituent system of an assembly of "the right kind". "The right kind" is determined by the desiderata of Section 5.1: the Hilbert space must be decomposable into single-systems Hilbert spaces, each of which supports a representation of the spacetime symmetry group (usually the Galilei group). Given the discussion in Section 5.2.1, an assembly of N distinguishable particles will have a Hilbert space $\mathcal{H} = \bigotimes^N (L^2(\mathbb{R}^3) \otimes \mathcal{H}_{int})$, where $\mathcal{H}_{int}$ allows for more (internal) degrees of freedom (we may set $\mathcal{H}_{int} = \mathbb{C}$).

Returning to non-entanglement, we may say that, in the case of distinguishable particles, the assembly's state is non-entangled iff the constituent particles are maximally specific. We now have a template functional definition of 'particle' that we can apply to the case of indistinguishable particles, bearing in mind Ghirardi and Marinatto's heterodox definition of entanglement for indistinguishable systems. So the varietist says that *particles are those systems that are maximally specific just in case the state of an assembly of the right kind is non-GM-entangled.*

In this case, an assembly is "of the right kind" iff its Hilbert space is a *symmetry sector* of the corresponding Hilbert space for an assembly of distinguishable particles. That is, iff $\mathcal{H} = \mathcal{S}_\mu \left[\bigotimes^N (L^2(\mathbb{R}^3) \otimes \mathcal{H}_{int})\right]$, where $\mathcal{S}_\mu$ is a projector onto the symmetry sector associated with the irreducible representation μ of S_N.

Let us investigate this functional definition a bit further: we say that 'particles are those systems that . . . ', but what exactly is a 'system' for varietists? Recall that being maximally specific is that same as being successfully qualitatively individuated by a one-dimensional projector. (The fact that we appeal, even in the distinguishable case, to *qualitative* individuation is precisely what allows us naturally to extend the functional definition to the indistinguishable case.) But factorist particles may also be individuated non-qualitatively, by the appropriate factor Hilbert space label. This permits us to latch on to (i.e. individuate) a particle in a non-GM-entangled state of the assembly by appealing to whichever single-particle state *á la* which it is maximally specific (i.e. whichever single-particle state makes its qualitative individuation successful), and *then go on to re-identify* that particle in *other* states of the assembly: other states in which it may occupy a different single-particle state altogether—pure or mixed. For the factorist, non-qualitative individuation is king: we might say that, for a factorist, being associated with a particular Hilbert space label is an *essential* property of a distinguishable particle; whichever single-particle state it may occupy is merely *accidental*.

In the indistinguishable case, the varietist rejects factorism, so she has no non-qualitative means of individuating systems. Therefore she can only appeal to single-particle states to cross-identify systems between different states of the assembly. In metaphysicians' jargon, the varietist must adopt a qualitative essentialism.

Here a multitude of possible routes present themselves for the varietist. I will now investigate these routes. I begin by restricting ourselves to non-GM-entangled states (Sections 8.1.1 and 8.1.2); I then turn to GM-entangled states in Section 8.1.3. I will conclude with a more specific commitment to what, for a varietist, particles are.

8.1.1 What varietism says about non-GM-entangled fermions

Let us take the simplest example, i.e. a non-GM-entangled state of an assembly of two fermions, belonging to the Hilbert space $\mathcal{A}(\mathcal{H} \otimes \mathcal{H})$:

$$|\psi\rangle = \frac{1}{\sqrt{2}} \left(|\phi_i\rangle \otimes |\phi_j\rangle - |\phi_j\rangle \otimes |\phi_i\rangle \right). \tag{8.1}$$

where $\langle \phi_i | \phi_j \rangle = 0$. (It will be obvious how to generalize these considerations for assemblies of three or more particles.) $|\psi\rangle$ is non-GM-entangled, since we have $\mathrm{Tr}(E_i |\psi\rangle\langle\psi|) = 1$, where $E_i := |\phi_i\rangle\langle\phi_i| \otimes \mathbb{1} + \mathbb{1} \otimes |\phi_i\rangle\langle\phi_i|$, and $\mathrm{Tr}(E_j|\psi\rangle\langle\psi|) = 1$ for E_j defined similarly; thus we have two maximally specific systems.

But which systems are they? Here the varietist stays silent. Since the varietist can only individuate qualitatively, and since she has already individuated using the most specific criteria possible (i.e. one-dimensional projectors), there is *nothing* *more* for her to say. In words:

> '$|\psi\rangle$ is a non-entangled state of two fermions, one of which is in the pure state $|\phi_i\rangle$, and the other of which is in the pure state $|\phi_j\rangle$.'

This is what we would say, in the distinguishable case (where we are all factorists), about the product state $|\phi_i\rangle \otimes |\phi_j\rangle$ (except we would not describe these particles as fermions), with the notable difference that we can also *add* to this description which particle occupies which state.[1] The varietist can make no such addition: the permutation-invariant information is *all* the information.

Note that $|\psi\rangle$ can be obtained from the product state by anti-symmetrization: recall (cf. Section 7.1.2) that *any* non-GM-entangled state can be obtained from a product state by the appropriate symmetry projection. We may therefore begin our characterization of varietism by laying down as a general rule:

(V0) The varietist description of a non-GM-entangled fermionic state $|\Psi\rangle$ reads like the factorist description of the *permutation-invariant* information of the corresponding product state from which $|\Psi\rangle$ is obtained, upon anti-symmetrization.

An immediate problem arises for fermions. As we saw in Stage C of Section 7.1.2, any non-GM-entangled fermionic state is obtainable by anti-symmetrization from a *variety* of different product states. There are as many pairs of varietist particles for the state $|\psi\rangle$ as there are ways to halve a sphere: the particles in $|\phi_i\rangle$ and $|\phi_j\rangle$ are only one such pair.

This problem is significant, and I will argue later that it may be fatal for varietism. But for now we bracket the problem, and continue our exposition. There are two reasons for this: (i) despite this problem, there is something intuitively appealing about varietism, so it is worth developing a full account of it for its own sake; and (ii) the problem may yet be overcome, in which case a full account of varietism is all the more valuable. I return to this "preferred basis problem" for fermions in Section 8.3; in the meantime we return to the exposition.

[1] Remember that this does not make our factorist a haecceitist, for the particle labels may represent distinct intrinsic qualitative properties of the particles, such as mass or charge.

What would the factorist say about $|\psi\rangle$ in (8.1), above? The answer depends on the subspecies of factorist. A haecceitistic factorist describes $|\psi\rangle$ in the following way:

> 'Set aside Ghirardi's understanding of "entanglement"! $|\psi\rangle$ is an entangled state, with two non-entangled terms. In one of the terms, particle 1 is in the pure state $|\phi_i\rangle$ and particle 2 is in the pure state $|\phi_j\rangle$; these states are transposed in the other term. The two terms have a relative amplitude of -1; this means that the two particles are fermions.'

There is something here for the varietist to agree with: namely that $|\psi\rangle$ is an eigenstate of having one fermion in $|\phi_i\rangle$ and one in $|\phi_j\rangle$. But the varietist takes this as an *exhaustive* characterisation of $|\psi\rangle$, while for the haecceitistic factorist there is (in principle) more to ask: in particular, *which* particle is in which state.[2]

This makes the varietist sound like an anti-haecceitist: indeed, varietism *is* anti-haecceitistic in the sense of Section 6.2.1, i.e. in taking the group action of S_N on the assembly's Hilbert space to represent no physical change, i.e. in taking the permutation-invariant information to be exhaustive. But that is anti-haecceitism about the factorist's particles: it is not the sort of "anti-haecceitism" that matters most to a varietist. Recall, from Section 6.2.2, that the question of haecceitism *properly so-called* comes after the question, 'What are the objects?' The varietist's particles do not correspond to factor Hilbert space labels, so she would not consider permutations of factor Hilbert space labels (i.e. the representations of S_N) to correspond to a genuine swapping of particles among the single-particle states.

Hence it is a separate question whether possibilities for the *varietist's* particles supervene on the qualitative character of physical states. But the formalism as it stands compels the answer Yes, i.e. anti-haecceitism for the varietist's particles too. A specification of single-particle pure states, together with a specification of symmetry type, suffices to determine a unique (non-GM-entangled) state for

[2]The haecceitistic factorist would have to admit that, in practice, no measurement could be performed which would answer this question; otherwise the projection postulate would come into conflict with the Indistinguishability Postulate. The haecceitist, of course, takes the Indistinguishability Postulate to be a *contingent* fact.

the assembly (cf. Stage E of Section 7.1.2);[3] therefore there are never two distinct non-GM-entangled states which do not differ with regard to qualitative character. (I will return to this point below.) So, barring any claim that the formalism is incomplete (which I will consider in Section 8.2.6), anti-haecceitism in the sense appropriate to varietism is forced upon the varietist.

As we have said, varietists are *also* anti-haecceitistic about the factorist's particles, but they are of course not anti-haecceitistic factorists. Admittedly, varietists and anti-haecceitistic factorists alike will be uncomfortable about talking, as the haecceitistic factorist does, of the two separable terms of $|\psi\rangle$ as if each were a possible pure state of the assembly. (For an anti-haecceitist factorist, heterogeneous separable states represent *mixed* states of the assembly.) Thus here is what the anti-haecceitist factorist says about the state $|\psi\rangle$:

> 'Set aside Ghirardi's understanding of "entanglement"! $|\psi\rangle$ is an entangled state, since its particles are not in pure states. Both particles are in the same mixed state $\frac{1}{2}|\phi_i\rangle\langle\phi_i| + \frac{1}{2}|\phi_j\rangle\langle\phi_j|$; but don't take those "$\frac{1}{2}$"s as probabilities for the assembly occupying the corresponding product states, like the haecceitist does. These states are not available to the assembly—in fact their mathematical representations don't really make sense to me. The minus sign between the two terms in $|\psi\rangle$ tells us that the particles in question are fermions; but again, don't think of $|\psi\rangle$ as a *superposition* of physically possible product states.'

(Note how the anti-haecceitistic factorist has managed to describe the state $|\psi\rangle$ in a way that relies on associating particles with factor Hilbert spaces, yet *without* mentioning Hilbert space labels. This trick relies on the fact that both particles are in the *same* mixed state.) The contrast with what the varietist would say is stark. For the varietist, the particles *are* in pure states; therefore $|\psi\rangle$ is *not* GM-entangled.

[3] Strictly speaking, that is only true for fermionic and bosonic states. For paraparticle states one would also have to specify certain relations between the particles; but these are *qualitative* relations (i.e., they are permutation-invariant), so the spirit of my claim here holds for paraparticles too.

8.1.2 What varietism says about non-GM-entangled bosons

Section 8.1.1's account carries over *mutatis mutandis*, i.e. with judicious substitutions of 'boson' for 'fermion', etc., for *heterogeneous* bosonic states, i.e. states such as

$$|\psi'\rangle = \frac{1}{\sqrt{2}}\left(|\phi_i\rangle \otimes |\phi_j\rangle + |\phi_j\rangle \otimes |\phi_i\rangle\right) \tag{8.2}$$

where $\langle\phi_i|\phi_j\rangle = 0$. But bosons may also exist in homogeneous product states, such as

$$|\psi''\rangle = |\phi_i\rangle \otimes |\phi_i\rangle. \tag{8.3}$$

What might the varietist say about these? The case is interesting, since, despite there being two maximally specific systems, the systems in question are maximally specific *à la* the *same* one-dimensional projector. Therefore, it cannot be said that the systems have been successfully qualitatively individuated.

Nevertheless, we may say that both systems are maximally specific without actually having to individuate them. It is enough that $(|\phi_i\rangle\langle\phi_i| \otimes |\phi_i\rangle\langle\phi_i|)\,|\psi''\rangle = 2|\psi''\rangle$, i.e. that $|\psi''\rangle$ is an eigenstate of there being exactly 2 particles in state $|\phi_i\rangle$. That is, it is enough so long as we do not impose any form of the identity of indiscernibles (cf. Chapter 3) on the particles (we will come back to this point in Section 8.2).

Homogeneous product states offer a rare opportunity for consensus between factorists, of both haecceitistic and anti-haecceitistic persuasion, and varietists. All three would describe $|\psi''\rangle$ as a pure state of two particles in which both particles were in the pure state $|\phi_i\rangle$. The consensus is no surprise: in these states (and *only* these states) factor Hilbert space labels—which the factorist uses to individuate particles—align perfectly with single-particle states—which the varietist uses to individuate particles.

The only remaining case to consider for bosons are states such as

$$|\psi'''\rangle = \frac{1}{\sqrt{2(1+|\langle\phi_i|\theta_j\rangle|^2)}}\left(|\phi_i\rangle \otimes |\theta_j\rangle + |\theta_j\rangle \otimes |\phi_i\rangle\right) \tag{8.4}$$

where $0 < |\langle\phi_i|\theta_j\rangle| < 1$. But recall from Stage D of Section 7.1.2 that this state is in fact GM-entangled, being (perhaps) more perspicuously written as

$$|\psi'''\rangle = \frac{2\langle\phi_i|\theta_j\rangle|\phi_i\rangle \otimes |\phi_i\rangle + \sum_{k\neq i}\langle\phi_k|\theta_j\rangle\left(|\phi_i\rangle \otimes |\phi_k\rangle + |\phi_k\rangle \otimes |\phi_i\rangle\right)}{\sqrt{2(1+|\langle\phi_i|\theta_j\rangle|^2)}} \ . \tag{8.5}$$

We are therefore led to consider varietism's account for entangled states. But first I sum up Sections 8.1.1 and 8.1.2 by laying down an extension of (V0) to apply to *all* non-GM-entangled states:

(V1) The varietist description of a non-GM-entangled state $|\Psi\rangle$ reads like *some* factorist description of the permutation-invariant information of the corresponding product state from which $|\Psi\rangle$ is obtained, *by the appropriate symmetry projection.*

The modifications have been italicized: I say 'some factorist description', to bracket the basis arbitrariness problem for fermions and paraparticles; and we now incorporate all symmetry types (including paraparticles) by referring to the 'appropriate symmetry projection'. This concludes the preliminary exposition of varietism for non-GM-entangled states.

8.1.3 What varietism says about GM-entangled states

The functional definition of 'particle' given at the beginning of Section 8 applied only to non-entangled states. For a factorist, this is easily extended to entangled states by appealing to non-qualitative individuation with factor Hilbert space labels. The anti-factorists—of which varietists are a subspecies—can make no such appeal in the case of GM-entangled states.

The available routes ahead for the varietist may be placed into two broad options. (One of these options further bifurcates, as we shall see.) The first option is to be cagey, so I call it 'cagey varietism'. Cagey varietists avoid the problem of GM-entangled states by claiming that particles exist *only* in non-GM-entangled states of the assembly. That would be to admit that the concept—the intension—of particle failed to pick out an extension in some (indeed most) quantum mechanical states. The second option is to be heroic. Heroic varietists maintain that particles exist in *every* state of the assembly, including the GM-entangled ones. I consider the two options in turn.[4]

Cagey varietism. The problem with denying the existence of particles for GM-entangled states is that the cagey varietist is then faced with the question of what *does* exists when the assembly's state is GM-entangled. Presumably the assembly continues to exist, yet without its supposedly constituent particles. So a cagey varietist believes an assembly can exist without its particles. She must therefore give up on the idea that particles always compose the assembly. This contradicts the strong version of our compositionality desideratum (cf. Section 5.1.4), since according to that desideratum an assembly must be the mereological sum of its particles. It also contradicts the weak version of our compositionality desideratum, since (trivially) *any* two distinct GM-entangled states differ without there being any corresponding difference in the states of the particles (of which there are none). (However, supervenience still holds between the *non*-GM-entangled states of the assembly and the states of the particles.)

But matters are worse for the cagey varietist. If the assembly always exists, and may exist even though its particles don't, why do we need particles at all? The assembly's state alone is enough to make any statement about it true or false—nevermind the particles!

The cagey varietist has a response. There is a natural way, she argues, to admit that the assembly's state suffices to make any statement about it true or false *without* threatening the particles with redundancy. For the particles could themselves be features—i.e., properties—of the assembly's non-GM-entangled states. They are not extra idle objects, but rather ontological free-riders. That way, talk of particles just *is* convenient talk about the assembly when it is in a non-GM-entangled state. The cagey varietist need not give up on particles altogether, but she must give up on them as *objects*.

The view that particles are not objects but properties is part of the view I call *emergentism*. I discuss this view below in Chapter 9, so we will say no more about it here. It is enough to note that cagey varietism collapses into a particular version

[4]Of course, these options are not exhaustive: one may maintain that particles exist in some but not all GM-entangled states. This option is unmotivated and arbitrary, so I exclude it.

of it. Henceforth, 'varietism' will always mean the heroic kind.

Heroic varietism. The second option is to claim that particles exist for all states of the assembly. We might therefore hope to retain the familiar principle that the assembly is composed of particles, at least in the weak sense of Section 5.1.4. But once again we find that our way ahead is not determined. *Prima facie*, there are (at least) two natural ways to proceed.

1. *Weaken.* Inspired by the technical results in Sections 7.2, the varietist may consider weakening the individuation criteria in order to successfully individuate across several non-GM-entangled terms. In this case a particle need no longer be a maximally specific system, but a system which is specific *à la* some (possibly multi-dimensional) single-system projector P, i.e.

$$\mathrm{Tr}(\rho E) = 1, \tag{8.6}$$

where ρ is the state of the assembly and $E := P \otimes \mathbb{1} + \mathbb{1} \otimes P - P \otimes P$.

2. *Relativize.* An alternative suggestion is to retain maximum specificity, but take advantage of the fact that any GM-entangled state is a superposition of non-GM-entangled states. Thus the varietist may consider treating any two non-GM-entangled terms as representing *distinct* collections of particles, related by being superposed. Under this proposal, particles are "branch-bound" entities, where by "branch" we mean a non-GM-entangled state. So under this suggestion, a particle is a system that is maximally specific on at least *one* of the branches of the assembly's state. This is equivalent to

$$\mathrm{Tr}(F|\Xi\rangle\langle\Xi|) = 1, \qquad \mathrm{Tr}(\rho|\Xi\rangle\langle\Xi|) > 0. \tag{8.7}$$

for some non-GM-entangled state $|\Xi\rangle$, where ρ is the (pure) state of the assembly and $F := Q \otimes \mathbb{1} + \mathbb{1} \otimes Q - Q \otimes Q$, for some one-dimensional projector Q.

An example may help to illustrate these two suggestions. Consider the following state for a two-fermion assembly:

$$|\psi\rangle = \alpha\frac{1}{\sqrt{2}}\left(|\phi_1\rangle \otimes |\phi_2\rangle - |\phi_2\rangle \otimes |\phi_1\rangle\right) + \beta\frac{1}{\sqrt{2}}\left(|\phi_3\rangle \otimes |\phi_4\rangle - |\phi_4\rangle \otimes |\phi_3\rangle\right) \quad (8.8)$$

where $|\alpha|^2 + |\beta|^2 = 1$. I assume that the states $|\phi_i\rangle$ belong to a single-system Hilbert space that is appropriate for a particle interpretation, in accordance with Section 5.2.1. What, according to our two new species of (heroic) varietist, are the constituent particles in this state?

Individuating particles under Weaken

The proponent of *Weaken* needs to find single-system projectors which satisfy Equation (8.6). I set aside for now the basis arbitrariness problem (cf. Stage C in Section 7.1.2) for fermion states by considering only projectors in the $\{|\phi_i\rangle\}$ basis. Still there are many options. They are:

$$\left.\begin{array}{rcl}
P_{13} &=& |\phi_1\rangle\langle\phi_1| + |\phi_3\rangle\langle\phi_3| \\
P_{14} &=& |\phi_1\rangle\langle\phi_1| + |\phi_4\rangle\langle\phi_4| \\
P_{23} &=& |\phi_2\rangle\langle\phi_2| + |\phi_3\rangle\langle\phi_3| \\
P_{24} &=& |\phi_2\rangle\langle\phi_2| + |\phi_4\rangle\langle\phi_4| \\
P_{123} &=& |\phi_1\rangle\langle\phi_1| + |\phi_2\rangle\langle\phi_2| + |\phi_3\rangle\langle\phi_3| \\
P_{124} &=& |\phi_1\rangle\langle\phi_1| + |\phi_2\rangle\langle\phi_2| + |\phi_4\rangle\langle\phi_4| \\
P_{134} &=& |\phi_1\rangle\langle\phi_1| + |\phi_3\rangle\langle\phi_3| + |\phi_4\rangle\langle\phi_4| \\
P_{234} &=& |\phi_2\rangle\langle\phi_2| + |\phi_3\rangle\langle\phi_3| + |\phi_4\rangle\langle\phi_4| \\
P_{1234} &=& |\phi_1\rangle\langle\phi_1| + |\phi_2\rangle\langle\phi_2| + |\phi_3\rangle\langle\phi_3| + |\phi_4\rangle\langle\phi_4|
\end{array}\right\} \quad (8.9)$$

and any other projector Q such that $Q > P_\lambda$ for any P_λ listed above. (It may be checked that Equation (8.6) holds for all such projectors.) Must the advocate of *Weaken* therefore say that the state $|\psi\rangle$ in Equation (8.8) above contains *at least* 9 (and potentially infinitely many—depending on the dimension of the single-particle Hilbert space) particles?

She must, but be careful not to misunderstand her. For it is normally an unspoken rule (except perhaps for analytic metaphysicians!) that when counting a collection of objects, the objects in question are taken to be wholly distinct. If this unspoken rule is explictly relaxed, and the objects in question may overlap, the question, 'How many are there?' can have a surprisingly large answer—consider the 'How many triangles are there?' puzzles in old IQ tests.[5] So the proponent of *Weaken* sanguinely admits that, yes, at least 9 particles are described by $|\psi\rangle$, but many of them overlap many of the others.

Which overlap which? This may be answered by calculating probabilities for being in the states corresponding to the projectors listed in (8.9), for each particle qualitatively individuated[6] using those same projectors. We use the results of Section 7.2.3 (in particular Equation (7.56)). It follows that for the particle individuated by P_λ and the single-particle state P_μ we have

$$p(\mu|\lambda) := p(P_\mu|P_\lambda) = \langle P_\mu\rangle_\lambda = \frac{\langle P_\lambda P_\mu P_\lambda \otimes \mathbb{1} + \mathbb{1} \otimes P_\lambda P_\mu P_\lambda\rangle}{\langle P_\lambda \otimes \mathbb{1} + \mathbb{1} \otimes P_\lambda\rangle} \tag{8.10}$$

The resulting probabilities for the state $|\psi\rangle$ in (8.8) are shown in Table 8.1.

These probabilities may be interpreted as a measure of degree of overlap.[7] For example, $p(13|24) = p(24|13) = 0$, so the particle that is specific *à la* P_{13} is wholly distinct from the particle that is specific *à la* P_{24}; and similarly for P_{14} and P_{23}. (We may use the definite article in all these cases, since the denominator in Equation (8.10) is equal to 1 for all these projectors.) Meanwhile, $p(123|13) = 1$ but $p(13|123) = \frac{1}{1+|\alpha|^2} < 1$, so the particle that is specific *à la* P_{13} is a proper part of the sum of particles (note the plural!—the denominator in Equation (8.10) is more than 1 for $P_\lambda = P_{123}$) that are specific *à la* P_{123}. And every particle or sum of particles that is/are specific *à la* any one of the projectors in (8.9) is a part of the sum of the particles that are specific *à la* P_{1234}, since $p(\mu|1234) = 1$ for all P_μ in (8.9). There are exactly *two* particles that are specific *à la* P_{1234} (the denominator of Equation (8.10) is equal to 2 for $P_\lambda = P_{1234}$); we may identify their sum with the assembly itself.

[5] A more concrete example: How many pairs of socks are there for 4 similar socks? The answer is ${}^4_2C = 6$, not $\frac{4}{2} = 2$.

[6] Recall from Section 7.2.3 that more than one system may satisfy the same individuation criterion. In this case, as usual, we average over all such systems.

[7] This measure satisfies the necessary conditions for giving the *fraction* of overlap, so that $p(\mu|\lambda)$ gives the fraction of the sum of particles specific *à la* P_λ that are also specific *à la* P_μ. It may checked that e.g. $p(23|1234) = p(23|123)p(123|1234) + p(23|234)p(234|1234)$. So, for example, $p(13|1234) = \frac{1}{2}$ may then be interpreted as the particle specific *à la* P_{13} overlapping exactly half of the two particles specific *à la* P_{1234}.

Single-particle state

Individuation criterion	P_{13}	P_{14}	P_{23}	P_{24}	P_{123}	P_{124}	P_{134}	P_{234}	P_{1234}
P_{13}	1	$\lvert\alpha\rvert^2$	$\lvert\beta\rvert^2$	0	1	$\lvert\alpha\rvert^2$	1	$\lvert\beta\rvert^2$	1
P_{14}	$\lvert\alpha\rvert^2$	1	0	$\lvert\beta\rvert^2$	$\lvert\alpha\rvert^2$	1	1	$\lvert\beta\rvert^2$	1
P_{23}	$\lvert\beta\rvert^2$	0	1	$\lvert\alpha\rvert^2$	1	$\lvert\alpha\rvert^2$	$\lvert\beta\rvert^2$	1	1
P_{24}	0	$\lvert\beta\rvert^2$	$\lvert\alpha\rvert^2$	1	$\lvert\alpha\rvert^2$	1	$\lvert\beta\rvert^2$	1	1
P_{123}	$\frac{1}{1+\lvert\alpha\rvert^2}$	$\frac{\lvert\alpha\rvert^2}{1+\lvert\alpha\rvert^2}$	$\frac{1}{1+\lvert\alpha\rvert^2}$	$\frac{\lvert\alpha\rvert^2}{1+\lvert\alpha\rvert^2}$	1	$\frac{2\lvert\alpha\rvert^2}{1+\lvert\alpha\rvert^2}$	$\frac{1}{1+\lvert\alpha\rvert^2}$	$\frac{1}{1+\lvert\alpha\rvert^2}$	1
P_{124}	$\frac{\lvert\alpha\rvert^2}{1+\lvert\alpha\rvert^2}$	$\frac{1}{1+\lvert\alpha\rvert^2}$	$\frac{\lvert\alpha\rvert^2}{1+\lvert\alpha\rvert^2}$	$\frac{1}{1+\lvert\alpha\rvert^2}$	$\frac{2\lvert\alpha\rvert^2}{1+\lvert\alpha\rvert^2}$	1	$\frac{1}{1+\lvert\alpha\rvert^2}$	$\frac{1}{1+\lvert\alpha\rvert^2}$	1
P_{134}	$\frac{1}{1+\lvert\beta\rvert^2}$	$\frac{1}{1+\lvert\beta\rvert^2}$	$\frac{\lvert\beta\rvert^2}{1+\lvert\beta\rvert^2}$	$\frac{\lvert\beta\rvert^2}{1+\lvert\beta\rvert^2}$	$\frac{1}{1+\lvert\beta\rvert^2}$	$\frac{1}{1+\lvert\beta\rvert^2}$	1	$\frac{2\lvert\beta\rvert^2}{1+\lvert\beta\rvert^2}$	1
P_{234}	$\frac{\lvert\beta\rvert^2}{1+\lvert\beta\rvert^2}$	$\frac{\lvert\beta\rvert^2}{1+\lvert\beta\rvert^2}$	$\frac{1}{1+\lvert\beta\rvert^2}$	$\frac{1}{1+\lvert\beta\rvert^2}$	$\frac{1}{1+\lvert\beta\rvert^2}$	$\frac{1}{1+\lvert\beta\rvert^2}$	$\frac{2\lvert\beta\rvert^2}{1+\lvert\beta\rvert^2}$	1	1
P_{1234}	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1+\lvert\alpha\rvert^2}{2}$	$\frac{1+\lvert\alpha\rvert^2}{2}$	$\frac{1+\lvert\beta\rvert^2}{2}$	$\frac{1+\lvert\beta\rvert^2}{2}$	1

Table 8.1: Single-particle probabilities for the various single-particle states for the various qualitatively "individuated" particles in state $|\psi\rangle$ in (8.8), associated with the projectors in (8.9). N.B. $|\alpha|^2 + |\beta|^2 = 1$.

If we count only strictly non-overlapping particles, then we recover the reassuring result that state $|\psi\rangle$ in (8.8) contains two particles. The particles must not overlap, and their sum must be identical to the assembly. So we require either two projectors P_κ and P_λ from (8.9) such that $p(\kappa|\lambda) = p(\lambda|\kappa) = 0$ and $\langle P_\kappa \otimes P_\lambda + P_\lambda \otimes P_\kappa\rangle = 1$; or else just one projector P_κ from (8.9) such that $\langle P_\kappa \otimes P_\kappa\rangle = 1$. There are three ways to satisfy this requirement: we may use the pair P_{13} and P_{24} (there is one particle specific *à la* each of this pair), or the pair P_{14} and P_{23} (similarly), or the single projector P_{1234} (of which there are two corresponding particles). Doesn't this mean now that there are six wholly distinct particles and not two? No: since it may be checked that the sum of any one of the three pairs is identical to the sum of any other.[8]

Thus the advocate of *Weaken* recovers the results we expect by appealing to mereology. The weakening from 'maximally specific' to merely 'specific' has allowed constituent particles to overlap, and be parts of, one another. What does the advocate of *Relativize* have to say about state $|\psi\rangle$ in (8.8)? What, for her, are the constituent particles?

Individuating particles under Relativize

The proponent of *Relativize* must express $|\psi\rangle$ as a superposition of non-GM-entangled "branches", all of which satisfy Equation (8.7). The constituent particles are then the maximally specific systems on each branch. But how are the branches determined? It turns out that, while it is a basis-independent matter whether or not a state is GM-entangled, it is *not* determined which non-GM-entangled states superpose to yield a given GM-entanged state. This is a kind of basis arbitrariness that affects systems of *all* symmetry types, not just fermions and paraparticles. I address this problem below. For now, we will work in the $\{|\phi_i\rangle\}$ product basis to give a flavour of the account given by the proponent of *Relativize*.

The state $|\psi\rangle$ in (8.8) is expressed as the superposition of two non-GM-entangled states in the $\{|\phi_i\rangle\}$ product basis:

$$\frac{1}{\sqrt{2}}\left(|\phi_1\rangle \otimes |\phi_2\rangle - |\phi_2\rangle \otimes |\phi_1\rangle\right) \quad \text{and} \quad \frac{1}{\sqrt{2}}\left(|\phi_3\rangle \otimes |\phi_4\rangle - |\phi_4\rangle \otimes |\phi_3\rangle\right). \tag{8.11}$$

Each non-GM-entangled state in (8.11)—each a branch of $|\psi\rangle$ in (8.8)—is identified with a *separate* collection of particles, each of which is maximally specific *à la* some single-particle state, just as in Section 5.2.1 each branch was associated with a different particle. In my example we have two branches. One consists of the pair of particles maximally specific *à la* $P_1 = |\phi_1\rangle\langle\phi_1|$ and $P_2 = |\phi_2\rangle\langle\phi_2|$; the other consists of the pair of particles maximally specific *à la* $P_3 = |\phi_3\rangle\langle\phi_3|$ and $P_4 = |\phi_4\rangle\langle\phi_4|$. (It may be checked that Equation (8.7) is satisfied for these branches and these single-particle projectors.)

[8] I.e., $p([P_{13}\otimes P_{24}+P_{24}\otimes P_{13}]|[P_{14}\otimes P_{23}+P_{23}\otimes P_{14}]) = p([P_{14}\otimes P_{23}+P_{23}\otimes P_{14}]|P_{1234}\otimes P_{1234}) = p(P_{1234} \otimes P_{1234}|[P_{13} \otimes P_{24} + P_{24} \otimes P_{13}]) = 1$.

The superposition of these two branches in (8.8) is then understood as co-existing (in a world, at a time) pairs of particles. The relative amplitude $\frac{\beta}{\alpha}$ between the two branches may be interpreted as a relation—itself a determinate of a single dyadic determinable—holding between the two pairs of branch-bound particles. The assembly itself may be identified with the sum of these two superposed pairs, so related.

As peculiar as it sounds, this suggestion is akin to the by-now-familiar account of macroscopic objects suggested by many proponents of the Everettian response to the measurement problem (Wallace (2003), Butterfield (2001)). Macroscopic objects, according to this account, are high-level patterns described by the universal wavefunction. However, these patterns are instantiated only in some branches of the universal quantum state and not others; one should therefore not expect to be able to find the same macroscopic objects in each branch. We may even say that macroscopic objects *exist* only in some branches and not others, so long as that is not taken to imply any sort of semantic indeterminacy of the existence claim. On the contrary: for the Everettian, existence in a branch entails existence *simpliciter* (just as, say, existence in Leicester entails existence *simpliciter*). And so as for macroscopic objects under the Everettian's suggestion, so too for the varietist's particles, under my suggestion.

(Note, however, that the varietist need not be an Everettian: so far I have only considered states of *microscopic* assemblies, and have said nothing about the microscopic/macroscopic boundary or the measurement process. Nevertheless, it cannot be denied that the varietist's account holds promise for a simple Everettian story for how particles compose macroscopic objects. Perhaps mereology—which is composition in the strong sense in Section 5.1.4—will do after all.)

It is also worth emphasising at this point that the ontological picture recommended by the proponent of *Relativize* is easily extended to accommodate states which are superpositions of *different numbers* of particles. Such states do not arise in elementary quantum mechanics, but are typical (indeed, characteristic!) of the theory of quantum fields, and are taken by some philosophers to preclude the possibility of a particle interpretation.[9]

[9]Of course, in QFT unsharp particle number is not the only problem thought to face particle

However, according to the proponent of *Relativize*, particles in different branches are strictly distinct, so there is no need for each branch to contain the same number of particles. Any (Fock space) state of the quantum field could still be understood in terms of (branch-bound) particles possessing certain properties and relations, so long as we include also relations that encode relative amplitudes between branches. Therefore, unsharp particle number is no special problem: indeed the effective restriction, in elementary quantum mechanics, to a particular summand of the Fock space, now has no special ontological significance. (Of course, the restriction is perfectly natural, practically speaking, if the interactions are such as to constrain the assembly's state to a particular Fock space summand, as indeed they do for low energies.)

Let us now pursue the problem raised at the outset for the proponent of *Relativize*, namely: How do we determine, for a given GM-entangled states, which non-GM-entangled states are to be its branches? To see the problem, let us consider a slightly different state:

$$\begin{aligned} |\psi'_\pm\rangle \;&=\; \alpha\frac{1}{\sqrt{2}}\left(|\phi_1\rangle \otimes |\phi_2\rangle \pm |\phi_2\rangle \otimes |\phi_1\rangle\right) + \beta\frac{1}{\sqrt{2}}\left(|\phi_2\rangle \otimes |\phi_3\rangle \pm |\phi_3\rangle \otimes |\phi_2\rangle\right) \\ &\quad + \gamma\frac{1}{\sqrt{2}}\left(|\phi_3\rangle \otimes |\phi_4\rangle \pm |\phi_4\rangle \otimes |\phi_3\rangle\right) \; . \end{aligned} \tag{8.12}$$

(We consider both the bosonic and fermionic version, to show that the problem is not peculiar to any single symmetry type.) Now we define two new single-particle states

$$\left.\begin{aligned} |\chi_1^\pm\rangle &:= \tfrac{1}{\sqrt{|\alpha|^2+|\beta|^2}}\left(\alpha|\phi_1\rangle \pm \beta|\phi_3\rangle\right); \\ |\chi_3^\pm\rangle &:= \tfrac{1}{\sqrt{|\alpha|^2+|\beta|^2}}\left(\beta^*|\phi_1\rangle \mp \alpha^*|\phi_3\rangle\right) \end{aligned}\right\} \tag{8.13}$$

Note that $\langle\chi_1^\pm|\phi_2\rangle = \langle\chi_3^\pm|\phi_2\rangle = \langle\chi_1^\pm|\phi_4\rangle = \langle\chi_3^\pm|\phi_4\rangle = 0$. The state $|\psi'_\pm\rangle$ may now

interpretations. Other problems facing any proponent of particles, include the Unruh effect and difficulties of localization. For a survey of these problems, see Baker (2009).

be re-expressed as

$$\begin{aligned}|\psi'_\pm\rangle &= \alpha'_\pm \frac{1}{\sqrt{2}}\left(|\phi_2\rangle \otimes |\chi_1^\pm\rangle \pm |\chi_1^\pm\rangle \otimes |\phi_2\rangle\right) + \beta' \frac{1}{\sqrt{2}}\left(|\chi_1^\pm\rangle \otimes |\phi_4\rangle \pm |\phi_4\rangle \otimes |\chi_1^\pm\rangle\right) \\ &\quad + \gamma' \frac{1}{\sqrt{2}}\left(|\phi_4\rangle \otimes |\chi_3^\pm\rangle \pm |\chi_3^\pm\rangle \otimes |\phi_4\rangle\right) \qquad (8.14)\end{aligned}$$

where

$$\left.\begin{array}{rcl}\alpha'_\pm &:=& \pm\sqrt{|\alpha|^2 + |\beta|^2} \\ \beta' &:=& \frac{\alpha\gamma}{\sqrt{|\alpha|^2+|\beta|^2}} \\ \gamma' &:=& \frac{\beta^*\gamma}{\sqrt{|\alpha|^2+|\beta|^2}}\end{array}\right\} \qquad (8.15)$$

It can be seen from (8.12) and (8.14) that the state $|\psi'_\pm\rangle$ is of the same form in the product basis induced by the single-particle basis $\{|\phi_1\rangle, |\phi_2\rangle, |\phi_3\rangle, |\phi_4\rangle\}$ as in the product basis induced by $\{|\phi_2\rangle, |\chi_1^\pm\rangle, |\phi_4\rangle, |\chi_2^\pm\rangle\}$. Therefore it is hard to see what consideration could favour the former basis without equally favouring the latter, and *vice versa*.

But although we may not be able to solve the problem, we can convert into another that I have already acknowledged. For the arbitrariness in basis for the *fermionic* GM-entangled state $|\psi'_-\rangle$ would be overcome if the arbitrariness in basis for the constituent non-GM-entangled states, first discussed in Stage C of Section 7.1.2 were overcome. A breaking of the under-determination in the latter case induces a breaking of the under-determination of the former. The proponent of *Relativize* may therefore delegate the solving of *her* basis arbitrariness problem, for *fermionic* GM-entangled states, to *any* varietist, who must solve the basis arbitrariness problem for fermionic non-GM-entangled states.

Besides, maybe the proponent of *Relativize* will shirk their responsibility even for *bosonic* GM-entangled states. Recall that bosonic non-GM-entangled states do not suffer a basis arbitrariness problem. But if they did, it seems plausible that any solution for fermionic states would be exportable to bosonic states: after all, in the former problem we seek a privileged basis for the *single*-particle Hilbert space. This privileged single-particle basis, if such there be, would induce a privileged product basis for states of any symmetry type. Thus I allow the proponent of *Relativize* to see her basis arbitrariness problem for *any* GM-entangled state as not especially fatal to her project. The real problem is the basis arbitrariness for *non*-GM-entangled states, and that faces every varietist.

Weaken or Relativize?

We have seen what each of the two kinds of heroic varietism has to say about GM-entangled states. So which particular species of heroic varietism should we prefer: one that has particles as possibly less than maximally specific, but as constituents always of an entire state; or one that has particles as branch-bound systems, but as always maximally specific?

The satisfying answer is that we need not choose. For there is no good reason to take *Weaken* and *Relativize* as mutually exclusive options. On the contrary: the mereological machinery that both endorse can be used to establish their concordance. The proposal for unification is simple: The (maximally specific, branch-bound) particles according to *Relativize* are *parts* of the (typically not maximally specific, trans-branch) particles according to *Weaken*.

For an illustration of this proposal, consider again the state $|\psi\rangle$ in (8.8). The constituent particles of this state under *Relativize* are four, and each is individuated by the projector $P_i := |\phi_i\rangle\langle\phi_i|$, where $i = \{1, 2, 3, 4\}$. The constituent particles under *Weaken* are individuated by the projectors shown in (8.9). But *just as*, under *Weaken*, we were encouraged to think of the particle specific *à la* P_{13} (say) as a *proper part* of the particle specific *à la* P_{123}, why not think of the particle maximally specific—in the relevant branch—*à la* P_1 as a *proper part* of the particle specific *à la* P_{13}?

This identification has advantages beyond reconciling *Weaken* and *Relativize*: it serves to explain the probabilities in Table 8.1. For example, in that Table $p(14|13) = |\alpha|^2$: this can be understood as the particle specific *à la* P_{14} overlapping the particle specific *à la* P_{13}, where the overlap is identical to the particle maximally specific *à la* P_1, which exists on a branch with amplitude α. Conversely, $p(24|13) = 0$, so the particle specific *à la* P_{24} has no common part with the particle specific *à la* P_{13}: i.e. there is no particle maximally specific on some branch that is a part of both. Overlap is part-identity, and part-identity is identity of parts. The particles according to *Relativize* provide the parts whose identity grounds the overlap of the particles according to *Weaken*.

Thus I return to an observation I made in Section 7.2.3, that we are led to something like a quantum counterpart theory (Lewis (1968; 1986, Ch. 4)) in which the basic entities are not world-bound, but rather branch-bound particles. These particles are maximally specific on their branch; that is, each may be qualitatively "individuated" on that branch with a one-dimensional projector that acts on the single-particle Hilbert space. With single-particle projectors—which may be multi-dimensional—we define counterpart relations, which select an (integer) number of particles (possibly zero) in each non-GM-entangled branch.[10]

As with familiar counterpart relations, the projector may fail to individuate a *unique* particle in each branch. In the example above, this occurs for projectors $P_{123}, P_{124}, P_{134}, P_{234}$ and P_{1234} (and more besides). In these cases the counterpart relation will not do as a surrogate for a "trans-branch" *identity* relation. However, by arbitrarily selecting *one* of the selected branch-bound particles from each branch, we define a "trans-branch individual", consisting of at most one particle from each branch. These trans-branch individuals typically overlap one another, and a single counterpart relation (a single projector) may define several such. If a trans-branch individual has a branch-bound particle in every branch, we may call it *ubiquitous*. Ubiquitous trans-branch individuals are precisely what were called 'particles' by the original proponent of *Weaken* (cf. Figure 8.1). Branch-bound particles must, of course, occupy pure states—the states *à la* which they are maximally specific. Trans-branch individuals, on the other hand, may occupy mixed states, so long as they are individuated by a multi-dimensional projector.

But which of the two better deserve the term 'particle': the branch-bound kind or the trans-branch kind? There is no sensible answer to this question and none is needed. Better to let context decide. If we are talking about constitution, we are likely to want to talk about the non-overlapping, maximally specific, branch-bound objects which are the parts of all the others. If we are talking about modality, or if we are qualitatively individuating (perhaps for the purposes of calculating an

[10]Note that, while maximally specific particles are branch-bound, even a one-dimensional projector may succeed in selecting some particle in *several* branches. A single branch-bound particle is uniquely selected only by giving maximally specific individuation criteria for *every* particle on its branch, using the individuation methods in Section 7.2.1.

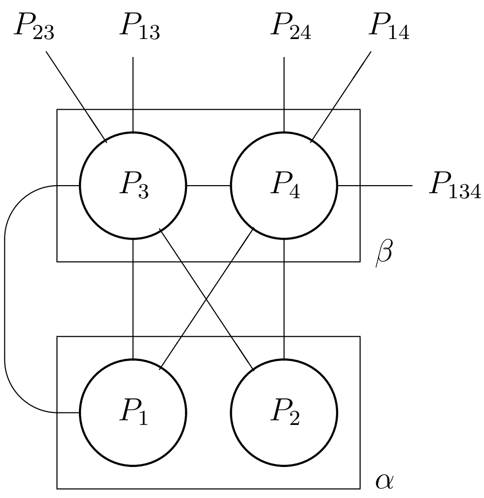

Figure 8.1: Branch-bound particles and some counterpart relations for the state $|\psi\rangle$ in (8.8). The counterpart relation induced by P_{134} defines *two* overlapping trans-branch individuals, which are the trans-branch individuals uniquely defined by P_{13} and P_{14}. All of these trans-branch individuals are ubiquitous.

expectation value), we are likely to want to talk about the (possibly overlapping) objects that have parts in more than one branch, and more than one state.

This concludes my outline of varietism. I sum up with a final explicit statement of what, for any state $|\Psi\rangle$ of the assembly, the varietist's particles are:

(V2) Any state $|\Psi\rangle$ is a sum of terms, each representing a non-GM-entangled "branch". (*Which* sum of terms is determined by whatever solution we find for the basis arbitrariness problem.) Each branch is composed exhaustively of (i.e. is the mereological sum of) wholly distinct maximally specific systems (which therefore must be in pure states); these are the most basic particles. Most generally, a particle is any mereological sum of basic particles, where at most one constituent basic particle belongs to each branch (a "trans-branch" individual). Any two branches are related by a relative amplitude. The full state $|\Psi\rangle$ is the mereological sum of these branches, so related.

I now turn to the desiderata for the concept of particle laid out in Section 5.1, and argue that the varietist's particles satisfy these desiderata to an acceptable degree. In the following Section (Section 8.3) I finally address the basis arbitrariness problem which threatens varietism's viability.

8.2 The merits of varietism

In this Section, I will address to what degree varietism satisfies Section 5.1's desiderata for the concept of particle. I will argue that varietism provides a tolerable target concept for particle—so long as we may assume the basis arbitrariness problem solved—a challenge I postpone until Section 8.3. In a final Section (8.2.6), we address the question whether varietist particles are discernible. But first I address, in order, the various desiderata for the concept of particle.

8.2.1 Varietist particles are physical

Factorism erroneously affords physical existence to a statistical construct, and varietism is the most natural proposal for what the factorist's particles are statistical constructs of. We saw in Section 7.2.3 that we recover the factorist's prescription for calculating the reduced density operator of a constituent particle by setting the qualitative individuation criterion as broad as possible (cf. Equation (7.71)). Under a varietist interpretation, this is tantamount to laying down a counterpart relation which will select *every* branch-bound particle in *every* branch; the resulting density operator is therefore a statistical average over the states of *all* of these branch-bound particles. That covers changeable properties. The unchangeable properties—what are often called 'intrinsic' or 'kinematical' properties—such as mass, spin and charge, are the same for any one of a collection of varietist particles of the same species as for the factorist particles which are statistical constructs of them.[11]

It might be objected that the varietist attributes more unchangeable properties to her particles than the factorist does to his, since it is with single-particle projectors that the varietist individuates her particles in the first place. Although these

[11]There is of course one notable exception; namely the unchangeable property of existing in the same mixed state as all other particles of the same species. This is, strictly speaking, a state-invariant property for an anti-haecceitistic factorist.

projectors represent quantities that are changeable for a factorist particle—insofar as their expectation values vary from state to state—they are unchangeable for a varietist particle—insofar as they constitute essential properties for that particle.

This claim has a more than passing resemblance to an erroneous potential criticism of counterpart theory (cf. Lewis (1986, pp. 9-13)). This potential criticism is that, since under counterpart theory every object exists only on one world, all objects have *all* of their properties essentially. The correct response is to point out that the modal properties of a world-bound object—say real-world Humphrey—are *represented* by the occurrent properties of certain objects in *other* worlds—other-world Humphreys, or Humphrey counterparts—which need not share the same properties as our original, real-world Humphrey. Thus world-bound objects have modal properties 'vicariously' (Lewis (1986, p. 10)). This response need only be slightly modified to suit the varietist. The slight modification registers the fact that particles are not only world-bound; they are branch-bound too. So a varietist particle may not only have modal properties vicariously, but may have *occurrent* (but other-branchly) properties vicariously too. They are represented by branch-bound particles that exist in other branches of the same state.

The claim that the varietist's particles are physical must be judged on whether they behave in a way befitting of physical entities. This judgement must be informed, in part, by their satisfaction of certain of our other strands of meaning for the concept of particle: namely, compositionality and inter-theoretic applicability. For, if the varietist's particles are good candidates for the constituents (in the broad sense; cf. Section 5.1.4) of familiar, macroscopic, physical objects, then they are good candidates for being physical entities. And if the varietist's particles look like classical particles in the classical limit, then we have a defeasible reason to extend our belief in the physicality of classical particles to the varietist's particles, even outside the classical limit. (Remember that it was here that the factorist proposal failed.) We must therefore look to these other desiderata, as we shall do below, in Sections 8.2.4 and 8.2.5.

Identity over time does not speak for or against the varietist's particles being physical, since plenty of physical things (like table-stages or events) do not exist over time. Locationality does not help either: if you don't believe that locationality entails physicality, then it is clearly no help; and if you believe that locationality *does* entail physicality, then presumably there would be no way to convince you that something was locational *before* you were already convinced that it was physical. All agree that there are plenty of unphysical things (e.g. centres of mass) which at least *seem* to be locational, or are treated as such.

Therefore, let me say no more about physicality here, and turn to the other strands of meaning.

8.2.2 Varietist particles are locational

That the varietist's particles are locational is clear (given that they are physical). For either one of the following two claims will always hold:

1. *Location is used to qualitatively individuate the particles.* This need not mean individuating by a precise location, as in e.g. 'The particle occupying the location with co-ordinates (x, y, z)'. (Just as well, given that Hilbert spaces do not contain eigenstates of position!) Rather, it means that particles are individuated by spatial wavefunctions, upon which there is no such restriction. (Since spatial wavefunctions determine momentum-space wavefunctions and *vice versa*, we count it as an instance of locational individuation even when the wavefunctions used yield sharply peaked expectation values for momentum and not position.)

   However, we demanded in Section 5.1.2 that locationality be cashed out in the following way: the state space for any particle must support a representation of the spacetime symmetry group. But particles that are individuated by projectors whose support is less than all of co-ordinate space simply cannot be found outside of that region of support. So the state space for such a particle do *not* support a representation of the spacetime symmetry group because it is not closed under action by that group.

   Despite this, it remains true that individuating criteria may appeal to any state in a state-space that *does* support such a representation. And because of this, varietist particles that are individuated by spatial wavefunctions may be said to fulfil the spirit of the locationality desideratum. That is to say, they satisfy the desideratum well enough.

2. *Location is not used to individuate the particles.* From Section 7.2.1 we know that this case requires the single-particle Hilbert space to contain at least one internal degree of freedom. So suppose we individuate with the projector $E_\sigma := e_\sigma \otimes \mathbb{1}_{sp}$, where e_σ acts on the factor single-particle Hilbert space corresponding to the internal degree of freedom, and $\mathbb{1}_{sp}$ is the identity on the spatial degree of freedom. By applying Equation (7.55), we find that the operator representing the position of the particle specific *à la* e_σ is

$$\pi_\sigma(\mathbf{Q}) = \sum_{k=1}^{N} \left[\bigotimes^{k-1} \mathbb{1} \otimes (e_\sigma \otimes \mathbf{Q}) \otimes \bigotimes^{N-k} \mathbb{1} \right] \tag{8.16}$$

where $\mathbf{Q}$ is the single-particle operator representing position, and $\mathbb{1}$ is the identity on the full Hilbert space.

In this case the particle individuated by e_σ *does* have a state space which supports a representation of the full spacetime symmetry group. The crucial fact is that, since e_σ only acts on the internal degree of freedom, the individuation criterion E_σ commutes with every generator of the group action on the single-particle Hilbert space.

8.2.3 Varietist particles do not (always) persist over time

The varietist countenances trans-branch individuals—objects that have as parts at most one basic particle per branch. And if one is not averse to countenancing objects at other times than the present, or objects in other worlds than the actual, then there is no reason to restrict these to actual, present branches. Thus, depending on one's ontological commitment to other times and other possible worlds, a trans-branch individual may also be a trans-temporal and trans-world individual.

According to varietism, an actual, present trans-branch individual deserves the name 'particle' as much as the branch-bound particles that are its parts. Therefore the varietist may countenance particles that exist over time.

However, as pointed out in Section 5.1.3, the commitment to trans-temporal individuals comes to more than a commitment to arbitrary mereological sums of objects existing at different times. The commitment is not to the existence of such entities—which is all too easy, given that one is happy to countenance non-present objects at all—but to the *naturalness* of such entities (cf. also Lewis (1986, p. 213)).

The naturalness of trans-temporal particles is threatened by the preliminary considerations in Section 7.3. There we found that for many possible quantum histories, either no *uniquely* natural trans-temporal individuation strategies exist, or else no trans-temporal individuation strategy exists at all. In fact matters are worse for the varietist, since even a single individuation strategy may (under-) determine several trans-branch individuals (cf. Section 8.1.3).

If we rule that any trans-temporal individual must be defined by a *natural* trans-temporal individuation strategy to earn the name 'particle', then the quantum world admits too many pathologies for it to be true that whenever there are branch-bound particles, then there are also particles of the trans-temporal kind. This means that varietism fails to satisfy the trans-temporal desideratum. But trans-temporal persistence was a non-compulsory constraint: trans-temporal individuation problems are not enough to prevent varietism's proposed target concept of particle being viable; they simply suggest that, if varietism is right, we are often better off to talk about particles of the branch-bound, therefore time-bound, kind.

8.2.4 Varietist particles compose assemblies

According to the recommended ontological picture in Section 8.1.3, the state of the assembly *always* supervenes on the properties and relations of the constituent particles, so long as we also include the relations which encode relative amplitudes between non-GM-entangled branches—these must be interpreted as multi-grade relations between the particles themselves. Therefore weak compositionality—i.e., supervenience of the assembly's properties on the particles' properties and relations—is satisfied. In fact, *strong* compositionality—i.e., mereological compositionality—is satisfied, since in any state of the assembly, the assembly is a mereological sum of suitably related branch-bound particles.

As we saw in Stage E of Section 7.1.2, there is an interesting supervenience result for varietism that is analogous to a more familiar supervenience result for factorism. The result for factorism is that an assembly's properties supervene on its constituent particles' properties *alone* iff the assembly is non-entangled. The corresponding result for varietism is that an assembly's properties supervene on its constituent particles' properties *and symmetry type* iff the assembly is non-GM-entangled.[12]

There can be no doubt that varietist particles satisfy the compositionality desideratum. Whether varietist particles might even compose macroscopic objects mereologically is an interesting question, but one that may be decoupled from our interests here. For the answer depends not on anything specific to quantum mechanics, but rather on whether macroscopic objects may be said to have precise characteristics.

8.2.5 Varietist particles have inter-theoretic applicability

Recall from Section 5.1.5 that the inter-theoretic applicability of a QM-local particle concept is a matter of ontological continuity in the limits of successful partial reduction of the other theories in question. 'Ontological continuity' is used here rather elastically: it is enough for me that the relevant objects in each theory *behave alike* in the appropriate limit. Of course, to make sense of 'behaving alike' I need a language to describe this behaviour that stands astride the different theories. And I need 'behaving alike' to come to more than just 'satisfies the other desiderata for any target concept of particle', since that much is already guaranteed. What I want to know is whether, in some reasonable sense, the varietist's particles become classical particles in the classical limit, and whether they become

[12] The supervenience result does not hold for paraparticles, since for any specific paraparticle type, distinct assembly states (that is, distinct generalized rays) exist which yield identical occupation numbers for single-particle states. We may regain supervenience by allowing relations between constituent particles into the picture which encode, for any two particles, whether their states are symmetrized or anti-symmetrized in the assembly's state. (This information determines a unique standard Young tableau, and the standard Young tableaux are in one-to-one correspondence with states of all assemblies of all symmetry types; cf. Tung (1985, Ch. 5).).

Fock space quanta in a limit of a QFT, of conserved total particle number.

Making 'behaving alike' precise in a way that is sufficiently general is—thankfully!—not a task I need undertake here.[13] I have a specific cases to consider, so let me appeal to the details of those cases. I consider classical mechanics and QFT in turn.

Varietist particles in the classical limit

Here I will only consider the $\hbar \to 0$ limit, and my discussion will be somewhat elementary.[14] This limit is typically modelled by a sequence of coherent states that are ever-narrowing Gaussians on the system's configuration and momentum spaces (Landsman (2007, §5)). These Gaussians, parameterized by values for $\hbar$, are used to calculate expectation values for the various quantities. In the $\hbar \to 0$ limit, we effectively obtain a Dirac delta function centred at a point in the system's phase space, which acts as a surrogate for the classical state associated with that point. That is: $\lim_{\hbar\to 0}\left(\langle \mathbf{Q}^2\rangle_\hbar - \langle \mathbf{Q}\rangle_\hbar^2\right) = 0$ and $\lim_{\hbar\to 0}\left(\langle \mathbf{P}^2\rangle_\hbar - \langle \mathbf{P}\rangle_\hbar^2\right) = 0$.

The case we are concerned with is more complicated, for two reasons. The first complication is that we are dealing with systems that are themselves assemblies of constituent systems. For distinguishable systems, this is easily handled by defining the appropriate coherent states as products of single-constituent coherent states. Thus in the classical limit we take for granted that the assembly's state is not entangled. However, the second complication is that we are dealing with *indistinguishable* systems, and the resulting superselection rule means that we are not at liberty to consider the behaviour of arbitrary products of coherent states.

This second complication is easily overcome by simply acting on any given product state with a projector of the appropriate symmetry type. Thus the assembly's wavefunction will be a superposition of multivariate Gaussians (i.e. products of single-constituent Gaussians), each centred at an image, under action by S_N, of some point in the configuration space. (The wavefunction has the same character in the momentum representation too.) No information is lost—so long as we

[13] Optimism that the right sense of 'behaving alike' can be made sufficiently precise for general purposes surely underpins Ladyman's (1998, 2002) ontic structural realism.

[14] An indispensable and thorough introduction to the classical limit in quantum mechanics is given by Landsman (2007).

also know the symmetry type of the state—if we instead represent the state as a Gaussian centred at a *single* point in the assembly's *reduced* configuration space, formed by quotienting by S_N. Given the results in Stages C and D of Section 7.1.2, this means that in the classical limit we take for granted that the assembly's state is non-GM-entangled.

That quantum states in the classical limit are taken to be non-GM-entangled is the key to the varietist's success in establishing ontological continuity between her particles and classical particles. For, as we have seen in Sections 8.1.1 and 8.1.2, in any non-GM-entangled state of the assembly, the constituent varietist particles are all in pure states. In the approach to the classical limit, these pure states are ever-narrowing Gaussians in both co-ordinate and momentum representations, so that in the classical limit itself, the varietist particles may be attributed a definite location and momentum.

In more detail: Recall from Sections 7.2.1 that qualitative individuation of quantum systems will be successful so long as the quantum state has its support restricted to some off-diagonal block of the reduced configuration space (RCS) of the assembly. For example, consider two particles on the real line. Then each off-diagonal block of the RCS is defined by a pair of intervals $\langle \Delta_1, \Delta_2 \rangle$, such that Δ_1 and Δ_2 are connected open regions of $\mathbb{R}$, and for every $x_1 \in \Delta_1$ and every $x_2 \in \Delta_2$, $x_1 < x_2$. Now for any state $\Psi(x, y)$, if

$$\int_{\Delta_1} \mathrm{d}x \int_{\Delta_2} \mathrm{d}y \; |\Psi(x,y)|^2 \; + \; \int_{\Delta_2} \mathrm{d}x \int_{\Delta_1} \mathrm{d}y \; |\Psi(x,y)|^2 = 1 \tag{8.17}$$

then

$$\langle E_{\Delta_1} \otimes E_{\Delta_2} + E_{\Delta_2} \otimes E_{\Delta_1} \rangle = 1 \tag{8.18}$$

where

$$(E_{\Delta_1} f)(x) := \begin{cases} f(x), & \text{for } x \in \Delta_1 \\ 0, & \text{for } x \notin \Delta_1 \end{cases} \; ; \quad (E_{\Delta_2} f)(x) := \begin{cases} f(x), & \text{for } x \in \Delta_2 \\ 0, & \text{for } x \notin \Delta_2 \end{cases} \tag{8.19}$$

Therefore, if condition (8.17) holds, we may say that there is one system specific *à la* E_{Δ_1} and one specific *à la* E_{Δ_2}. And since $E_{\Delta_1} \perp E_{\Delta_2}$, we may further say that these systems are distinct. If, further, the state $\Psi(x, y)$ is non-GM-entangled—as it will be when considering the classical limit—then each of these individuation criteria succeeds in picking out a unique branch-bound particle. One will occupy a state whose support lies in Δ_1; the other's state will have its support in Δ_2. Moreover, these two particles exhaustively constitute the assembly itself.

When considering the classical limit, we restrict attention to states $\Psi^{(\hbar)}_{(\mathbf{q},\mathbf{p})}(x, y)$, where $\Psi^{(\hbar)}_{(\mathbf{q},\mathbf{p})}(x, y)$ is a Gaussian centred at $\mathbf{q} = (q_1, q_2)$ (where $q_1 \leqslant q_2$) in the RCS, and its Fourier transform is a Gaussian centred at $\mathbf{p} = (p_1, p_2)$ (where $p_1 \leqslant p_2$) in the reduced momentum space. Now let us set $\Delta^{(\epsilon)}_1 := (q_1 - \epsilon, q_1 + \epsilon), \Delta^{(\epsilon)}_2 \equiv (q_2 - \epsilon, q_2 + \epsilon)$ for some $\epsilon > 0$. Then for any $\epsilon > 0$, Equation (8.17) holds for $\Psi(x, y) \equiv \Psi^{(\hbar)}_{(\mathbf{q},\mathbf{p})}(x, y)$ for some value of $\hbar > 0$, and all smaller values. Thus, given the comments in the previous paragraph, for any $\epsilon > 0$, at some point along the way to the classical ($\hbar \to 0$) limit, *if* $q_2 - q_1 \geqslant 2\epsilon$, *then* the projectors defined from $\Delta^{(\epsilon)}_1$ and $\Delta^{(\epsilon)}_2$ as in Equation (8.19) will each succeed in individuating a branch-bound particle, whose spatial wavefunctions are centred at q_1 and q_2 respectively, and whose momentum wavefunctions are centred at p_1 and p_2 respectively, and which together exhaustively compose the assembly.

Therefore, by selecting the correct ϵ we can, for *any* state $\Psi^{(\hbar)}_{(\mathbf{q},\mathbf{p})}(x, y)$, such that $q_2 > q_1$, individuate branch-bound particles for some value $\hbar > 0$ and thereafter, along the way toward the classical limit. These branch-bound particles will have increasingly definite locations and momenta as $\hbar \to 0$. Thus they are perfect candidates for being the temporal parts of classical particles (with matching locations and momenta) in that limit.

The remaining wrinkle is the case where $q_1 = q_2 =: q$. In this case we can use one "individuation criterion" defined, as in Equation (8.19), from the interval $\Delta^{(\epsilon)} := (q - \epsilon, q + \epsilon)$. There is no ambiguity which branch-bound particle is picked out in this case: both are. Therefore we may additionally individuate using momentum. Provided that the assembly wavefunction is centred at a point that attributes different momenta of the two particles, then we can run a similar individuation campaign to that in the previous two paragraphs, will equal success.

If, on the other hand, the assembly's wavefunction is centred at a point which attributes the *same* location *and* momentum to both particles, then no two projectors will serve to individuate one without also picking out the other. But this is no objection to identifying the indiscernible branch-bound particles with their classical counterparts; for their shared location and momentum is definite in the classical limit, and in the classical case too the two particle-stages are indiscernible. (Whether such a situation is metaphysically possible is not a question the varietist has to answer.)

Thus the varietist recovers classical particle-stages in the $\hbar \to 0$ limit. Unlike the factorist's particles, in this limit the varietist's particles may be attributed a definite position and momentum. But we can do more.

One crucial characteristic of classical particles is that they have definite *trajectories*. To recover genuine classical *particles* (and not just particle-stages) in the classical limit, the varietist must provide uniquely natural trans-temporal individuation criteria for her particles. But along the way to the classical limit, uniquely natural criteria exist. We need only make our original intervals $\Delta_1^{(\epsilon)}$ and $\Delta_2^{(\epsilon)}$ time-dependent. So define $\Delta_1^{(\epsilon)}(t) := (q_1(t) - \epsilon, q_1(t) + \epsilon)$ and $\Delta_2^{(\epsilon)}(t) := (q_2(t) - \epsilon, q_2(t) + \epsilon)$, where $(q_1(t), q_2(t))$ is the location of the system point in the RCS at time t, according to the *classical* dynamics. Then for any $\epsilon > 0$, at each time t, so long as $q_1(t) < q_2(t)$, we succeed in individuating two branch-bound particles somewhere along the way to the classical limit (and thereafter).

Trans-temporal individuals may then be constructed out of these individuated branch-bound particles with the requirement that any two branch-bound particles that exist at "near" times compose the same trans-temporal individual only if their wavefunctions significantly overlap. (This requirement can be made precise.) Again, collisions create a wrinkle: for if two branch-bound particles at a time *themselves* have significantly overlapping wavefunctions, then it will be indeterminate which ought to belong to which trans-temporal individual.

But again, this is no objection to the varietist's claim to have recovered persisting classical particles in the $\hbar \to 0$ limit. For collisions provide as much of a problem for the trans-temporal individuation of classical particles. Thus it is enough for the varietist's case to have recovered a unique pair of trans-temporal individuals for any segment of history for which the two individuals do not collide.

Varietist particles in QFT's limit of conserved particle number

Cross-theoretic ontological identifications are far more straightforward between quantum mechanics and QFT, since the state space of any quantum mechanical theory is a *subspace* of the state space for some quantum field theory. In QFT, the field's state is restricted to one such subspace (in the sense that these subspaces are selected) when the dynamics preserve particle number.

In this case the varietist's proposal is simple: *branch-bound particles are Fock space quanta*. We have already seen, in Section 8.1.3, that varietism is very naturally extended to states with variable particle number. However, it is a further—and perhaps more surprising fact—that varietist particles have been familiar to us all along, albeit in a different theory.

Recall what we know about Fock space quanta. We can glean all we need to know about a quantum from the operators with which it is associated—these are the creation and annihilation operators and the operators constructed from them. Starting from the unique vacuum state $|0\rangle$, it is familiar (cf. e.g. Maggiore (2005, p. 84)) that the mathematical state

$$a^{\dagger}_{\mathbf{k},s}|0\rangle \tag{8.20}$$

represents (up to a multiplicative factor) a physical state in which there is one quantum, whose state comprises a definite momentum with wavevector $\mathbf{k}$ and internal state (spin, polarisation, etc., if any) s. Let us represent this state by $|\mathbf{k}, s\rangle$. Then, by combining typical formalisms from QFT and elementary quantum mechanics, we may write

$$a^{\dagger}_{\mathbf{k},s}|0\rangle \propto |\mathbf{k}, s\rangle. \tag{8.21}$$

Operations on (8.21) with more creation operators yield multiple-particle states with an ever-increasing number of particles. For example, if $(\mathbf{k}, s) \neq (\mathbf{l}, r)$, then

$$a^{\dagger}_{\mathbf{l},r}a^{\dagger}_{\mathbf{k},s}|0\rangle \propto \frac{1}{\sqrt{2}} \left(|\mathbf{k}, s\rangle \otimes |\mathbf{l}, r\rangle \pm |\mathbf{l}, r\rangle \otimes |\mathbf{k}, s\rangle\right). \tag{8.22}$$

where '$\pm$' is positive for bosons and negative for fermions. The space of all states generated from finite operations on the vacuum state $|0\rangle$ by the creation operators

$a^\dagger_{\mathbf{p},r}$ is the Fock space for the quantum field.

The operators $a^\dagger_{\mathbf{p},r}$ and $a_{\mathbf{p},r}$, where $\mathbf{p} \in \mathbb{R}$ and r is an internal state index, are the creation and annihilation operators for momentum quanta, which (by definition) satisfy the (anti-) commutation relations

$$\left[a_{\mathbf{p},r}, a^\dagger_{\mathbf{q},s}\right]_\pm = (2\pi)^3\delta^{(3)}(\mathbf{p}-\mathbf{q})\delta_{r,s}, \tag{8.23}$$

where the subscript '$\pm$' indicates an anti-commutator or commutator, according to whether the quanta are fermions or bosons, respectively. And every annihilation operator annihilated the vacuum, i.e. for all $\mathbf{k}$, s:

$$a_{\mathbf{k},s}|0\rangle = 0. \tag{8.24}$$

If we now define a family of operators

$$N_{\mathbf{k},s} := a^\dagger_{\mathbf{k},s}a_{\mathbf{k},s}, \tag{8.25}$$

then from the (anti-) commutation relations (8.23),

$$\left[N_{\mathbf{p},r}, a^\dagger_{\mathbf{q},s}\right]_\pm = (2\pi)^3\delta^{(3)}(\mathbf{p}-\mathbf{q})\delta_{r,s}a^\dagger_{\mathbf{q},s}. \tag{8.26}$$

This, combined with the vacuum condition (8.24), entails that the integral

$$N(\Omega, S) = \int_\Omega \frac{\mathrm{d}^3\mathbf{k}}{(2\pi)^3}\sum_{s\in S} N_{\mathbf{k},s} \tag{8.27}$$

has a natural number spectrum on Fock space. This fact, combined with the identifications (8.21) and (8.22), entails that the operator $N(\Omega, S)$ is the quantity that *counts* the number of quanta whose momenta lie in Ω and whose internal states lie in S.

But $N(\Omega, S)$ also counts the number of varietist particles that are maximally specific *à la* states each of whose momentum lies in Ω, and internal state lies in S. To see this we only need to write $N(\Omega, S)$ out explicitly in a way that is more familiar from the point of view of elementary quantum mechanics, and our previous discussions in Sections 7.1.2 and 7.2.1. We have

$$N(\Omega, S) = \bigoplus_{n=1}^{\infty} \left[\sum_{k=1}^{n-1} \left(\bigotimes^{k-1} \mathbb{1} \otimes E(\Omega, S) \otimes \bigotimes^{N-k} \mathbb{1} \right) \right] \tag{8.28}$$

where $E(\Omega, S)$ is the single-particle projector defined by

$$E(\Omega, S)\psi_s(\mathbf{k}) := \begin{cases} \psi_s(\mathbf{k}), & \mathbf{k} \in \Omega, s \in S \\ 0 & \text{otherwise,} \end{cases} \tag{8.29}$$

and $\mathbb{1}$ is the identity on the single-particle Hilbert space. Thus each summand (in square brackets) of the right-hand side of (8.28) is an operator that acts on the appropriately (anti-) symmetrized n-particle Hilbert space. It may be checked that the operator on the right-hand side of Equation (8.28) is the unique operator that satisfies the requirements placed on it by the conditions from (8.21) to (8.27).

Note that each n-particle summand of $N(\Omega, S)$ is of precisely the same form as a single-system quantity associated with a qualitatively individuated system, where $E(\Omega, S)$ is the criterion of individuation; cf. Equation (7.55) in Section 7.2.3 (where $\alpha \equiv (\Omega, S)$). In this case the single-system quantity in question is the identity, which yields a quantity which we interpret as counting the average number of systems maximally specific *à la* states which lie in (Ω, S). But these maximally specific systems are, for the varietist, the branch-bound particles.

Thus $N(\Omega, S)$ counts both the number of Fock space quanta, and, in each n-particle subspace, the number of varietist branch-bound particles, which are in states circumscribed by (Ω, S). Moreover, this number is that same for quanta as for branch-bound particles. The correct response is clear: the Fock space quanta *are* the varietist's branch-bound particles.

To end this Subsection, we note two significant consequences of the identification of Fock space quanta with the varietist's branch-bound particles.

1. First, the lessons of anti-factorism *already* recommend something like a QFT ontology, even without having to consider superpositions of variable quantum number, difficulties in localization or the Unruh effect.

2. Second, any criticism facing varietism equally faces any interpretation of QFT that takes particles as the basic entities. So, in particular: the basis arbitrariness problem is a problem for a particle ontology in QFT as much as for varietism.

This concludes our survey of the desiderata for varietist's proposed target concept of particle. I have argued that varietist particles satisfy these desiderata well enough, despite possible problems facing trans-temporal individuation. But varietism's viability relies on there being a satisfactory solution to the basis arbitrariness problem, which I am yet to address. I will come to the problem soon, but first allow me a brief digression about discernibility.

8.2.6 Are varietist particles discernible?

The varietist's particles are composed of branch-bound particles (or are identical to them), and no two discernible objects may be composed of utterly discernible parts; so if any two particles are discernible, then they will be discernible by their parts. So our question comes to: Are branch-bound particles discernible?

Branch-bound particles are "individuated" by single-particle states. Fermionic assemblies are characterized, due to the total anti-symmetrization of their available states, by Pauli exclusion: that is, any single-particle state is occupied at most once. Therefore, *any* two distinct fermionic branch-bound particles occupy different states, and are therefore discernible. Moreover, they are always *absolutely* discernible (cf. Section 2.3), since the occupation of a particular single-particle state corresponds to a monadic property.[15] This vindicates Weyl's (1928, p. 241) claim, that 'the Leibnizian principle of *coincidentia indiscernibilium* holds in quantum mechanics [for fermions],' and his naming the exclusion principle as the 'Leibniz-Pauli exclusion principle' (Weyl (1949, p. 247)). In contrast, *factorist* fermions (and bosons and paraparticles) are always absolutely indiscernible (cf. Section 6.3), though they may be *weakly* discerned.

If bosonic or paraparticle assemblies are in states where no single-particle state

[15] If we take a single-particle state as an *intrinsic* property of a branch-bound particle, then fermionic branch-bound particles are not just absolutely, but intrinsically discernible.

is occupied more than once, then any two branch-bound particles for those assemblies are absolutely discernible too. However, bosons and paraparticles are not subject to Pauli exclusion, so single-particle states may be multiply occupied. Thus there are states for bosonic or paraparticle assemblies in which two bosons or two paraparticles may be *in*discernible by their single-particle states.

Does this make bosonic and paraparticle branch-bound particles utterly indiscernible or just absolutely indiscernible? Recall (Section 6.3) that factorist particles may be discerned, even though they occupy the same single-particle states, since they may be *weakly* discerned by some physical symmetric and irreflexive relation. Can these results for factorist particles be carried over for the varietist's branch-bound particles?

They cannot. Recall that the weak discernment of factorist particles relies on taking advantage of anti-correlations in the assembly's state (cf. Section 6.3.5). Even a state with no anti-correlations in some single-particle basis may be expressed in some new basis in which anti-correlations are guaranteed to arise. So the discernment of factorist particles relies on these particles *surviving the basis change*.

The varietist's branch bound particles do *not* survive single-particle basis changes. This is because, unlike factorist particles, they are individuated by their single-particle states. So if one changes the single-particle states in which the assembly's state is expressed, then one changes the branch-bound particles about which one is talking. (Or at least, this is so for any reasonable attempt to solve the basis arbitrariness problem; cf. Section 8.3.) Factorist particles are individuated by their factor Hilbert space labels, and these are of course preserved between basis changes. Varietist branch-bound particles just don't have anything similar to preserve them.

Thus branch-bound bosons and fermions are *utterly indiscernible* in states of the assembly in which the same single-particle state is multiply occupied. Recall too that the varietist is anti-haecceitistic about her particles (cf. Section 8.1.1). Therefore the varietist endorses the metaphysical thesis which in Part I we labelled 'QII'. If the varietist wishes to countenance only fermions, then any of the anti-haecceitistic theses, from SPII to QII, are consistent.

To conclude this Section, I will make a brief observation about how an adherence to varietism affects a well-established claim in the discernibility literature (this literature is more thoroughly discussed in Section 6.3).

The claim is that a certain argument for anti-haecceitism about particles from quantum statistics, perhaps first criticised by Redhead (1987, p. 12) and French and Redhead (1988, pp. 235-8), does not work. The argument for anti-haecceitism runs as follows. Consider two two-state quantum systems. Choose the single-system orthobasis $\{|H\rangle, |T\rangle\}$ ("head" and "tails"). Then the tensor product Hilbert space for the two-system assembly is spanned by the four states

$$|H\rangle \otimes |H\rangle; \quad |T\rangle \otimes |T\rangle; \quad |H\rangle \otimes |T\rangle; \quad |T\rangle \otimes |H\rangle. \tag{8.30}$$

But if the systems are bosons, then the observed statistics seem to show that only three distinct possibilities exist—namely: two heads, two tails, or one head and one tail—since an equal probability of $\frac{1}{3}$ is given to each. If the systems are fermions, then only *one* possibility exists: one head and one tail. These statistics can be explained (so the argument goes) if the last two states in (8.30) are not genuinely physically distinct. Yet the two states differ only haecceitistically. Therefore, there are no haecceitistic differences.

The error in this argument, as pointed out by Redhead and French, is that the *wrong* basis states are appealed to. Under IP, the four *pure* states are

$$|H\rangle \otimes |H\rangle; \ |T\rangle \otimes |T\rangle; \ \frac{1}{\sqrt{2}} \left(|H\rangle \otimes |T\rangle + |T\rangle \otimes |H\rangle\right); \ \frac{1}{\sqrt{2}} \left(|H\rangle \otimes |T\rangle - |T\rangle \otimes |H\rangle\right). \tag{8.31}$$

Now the last two of these states are certainly distinct: one is bosonic; the other is fermionic. What explains the statistics is that, if the assembly is bosonic, then only the first three states are dynamically *accessible*; and if it is fermionic, then only the last state is accessible.

The switch to varietist particles does not contradict this wisdom. For the varietist, as for anyone, the last two states in (8.31) are undeniably distinguishable—indeed they are *qualitatively* distinguishable, since an assembly's symmetry type is an eigenvalue of a symmetric quantity. But nevertheless, varietism revives the *spirit* of the argument for anti-haecceitism from statistics. The rehabilitated argument runs as follows.

Recall from Section 8.1.1 that the quantum formalism compels anti-haecceitism about varietist particles. That is, the varietist describes, for example, the state $\frac{1}{\sqrt{2}}(|H\rangle \otimes |T\rangle + |T\rangle \otimes |H\rangle)$ as one in which there is one boson maximally specific *à la* $|H\rangle$ and one boson maximally specific *à la* $|T\rangle$, and there is *no further question* which boson is which.

But the quantum formalism may be incomplete. So suppose it is, and suppose *there is* the further question: Which boson is which? Then there would have to be *two* physical states corresponding to the mathematical state $\frac{1}{\sqrt{2}}(|H\rangle \otimes |T\rangle + |T\rangle \otimes |H\rangle)$, related by a permutation of haecceities among the states *à la* which they are maximally specific. This permutation is *not* represented by the operation $|\phi\rangle \otimes |\psi\rangle \mapsto |\psi\rangle \otimes |\phi\rangle$, since varietist particles are not represented by factor Hilbert space labels. The permutation cannot be represented in the quantum formalism *at all*, since we supposed that the quantum formalism is incomplete in precisely this way.

On the other hand, on the assumption of haecceitistic differences, there would still be only *one* physical state corresponding to each of the product states $|H\rangle \otimes |H\rangle$ and $|T\rangle \otimes |T\rangle$. Thus, if haecceitism were true of varietist particles, one would expect the physical states (plural!) represented by $\frac{1}{\sqrt{2}}(|H\rangle \otimes |T\rangle + |T\rangle \otimes |H\rangle)$ to have *twice* the statistical weight assigned to $|H\rangle \otimes |H\rangle$ or $|T\rangle \otimes |T\rangle$ alone. Yet, in fact, they all have the same statistical weight, namely $\frac{1}{3}$. The best explanation of this is that our original assumption, namely that it made sense to ask which boson is which, beyond the distribution of single-system states, was incorrect. But this just is the assumption of haecceitistic differences. Thus, the quantum formalism is not incomplete, and anti-haecceitism is true of varietist particles.

8.3 A preferred basis problem for varietism

I have argued that varietism satisfies, to a sufficient degree, the desiderata for any target concept of particle. But varietism's viability depends on solving the problem first pointed out in Stage C of Section 7.1.2—the basis arbitrariness problem. This

Section argues for pessimism about varietism's prospects for solving this problem. Because of this, I will be led, in Chapter 9, to advocate a rival target concept of particle—emergentism.

The structure of this section is as follows. First, in Section 8.3.1, I will outline again the problem to be solved, and introduce five proposals, in increasing order of feasibility, to overcome it. Each of these five proposals will then be assessed, in Sections 8.3.2 to 8.3.6.

8.3.1 The problem of basis arbitrariness

In Section 8.1.3, the problem of basis arbitrariness was discussed in particular for non-GM-entangled states for assemblies of fermions and paraparticles. The problem is that there are continuum-many single-particle bases in which such states are manifestly non-GM-entangled, so for each non-GM-entangled state it is under-determined which branch-bound particles the varietist is to say compose the assembly in that state. From Stage C of Section 7.1.2 we know that, for an assembly of N fermions, this arbitrariness is parameterized by the $(N-1)!2^{N-1}$-real-dimensional manifold

$$\left(\mathbb{CP}^{N-1} \times \mathbb{CP}^{N-2} \times \cdots \times \mathbb{CP}^{1}\right)/S_N; \tag{8.32}$$

i.e., that points in this manifold correspond one-to-one to a choice of a single-particle basis in which a given non-GM-entangled state of the fermionic assembly is manifestly non-GM-entangled.

However, the problem—that the branch-bound particles out of which (the varietist will say) the assembly is composed are under-determined—is not limited to non-GM-entangled states of fermions and paraparticles. We also saw in Section 8.1.3 that it is under-determined, at least for some states, which non-GM-entangled "branches" the varietist should say superpose to yield a given GM-entangled state, even for bosonic assemblies. Again, this boils down to an under-determination of the single-particle basis which dictates which branch-bound particles the varietist says compose the assembly in that state.

Although in Section 8.1.3, I made this point by giving an example of a GM-entangled state which takes the same form for many choices of a single-particle basis, it may even be argued that the varietist's problem plagues *any* state whatsoever: for we can always re-express the same state using a different single-particle basis. I did not take this strict line, since it is a familiar fact—in classical mechanics too—that a state space may be co-ordinatized in a multitude of different ways. This in no way compromises the claim that certain co-ordinatizations are nevertheless privileged in virtue of aligning with natural ontological divisions, so long as such natural divisions exist.[16] The problem for the varietist is that there are plenty of states for which, given that natural divisions do exist, several "co-ordinatizations" (i.e., several single-particle bases) appear to have equal claim to align with them.

It must also be emphasized that, given the clean meshing between the varietist's branch-bound particles in elementary quantum mechanics and the quanta of QFT (cf. Section 8.2.5), these basis arbitrariness problems face those who seek a particle ontology of QFT. And, of course, this is all notwithstanding the additional problems facing a particle interpretation found there, such as the Unruh effect.

In the following Sections, we investigate five potential responses to the basis ambiguity problem. The responses are not always exclusive: one may be implemented for some states, and others for other states, strategically. Nor do we claim that these responses are exhaustive. Perhaps a decent proposal exists that we have overlooked, in which case varietism would be saved. But we know of no such proposal.

By way of introduction, we list the five responses here. They may be categorised into two groups: the responses which attempt to overcome basis arbitrariness by finding, for each state, a uniquely privileged basis; and the responses which

[16] This general way of looking at the matter is illustrated by what is now the consensus regarding the empirical content of diffeomorphism covariance in general relativity, in the light of Kretschmann's objection to Einstein. Despite the undeniable fact that the spacetimes in Newtonian mechanics and special relativity may too be expressed in arbitrary co-ordinates, general relativity is distinguished by the fact that no *natural* (in the sense of aligning with the inertial structure) co-ordinate system may be specified independently of a specification of the distribution of mass-energy. For a better discussion of this response to the Kretschmann objection, see e.g. Misner, Thorne and Wheeler (1973, §17.7) and Brown (2006, pp. 154-6, 178-81).

attempt to overcome the arbitrariness by somehow accommodating all natural bases at once, without privileging any one over the other. The first two responses fall under the former category; the second three responses fall under the latter. The responses are written as claims for vividness.

1. *One size fits all.* There is a uniquely natural single-particle basis for each state of the assembly. It is the same basis for *every* state. That is, there is a categorically privileged single-particle basis.

2. *Complicate.* In realistic cases, there is more than one degree of freedom under consideration. These extra degrees of freedom provide the extra structure needed to determine a uniquely natural single-particle basis, for each state.

3. *Coalesce.* All of the rival single-particle bases may be reconciled by reifying *all* of the corresponding branch-bound particles. But in each non-GM-entangled branch, each particle associated with one single-particle basis is identical to some particle associated with any other single-particle basis.

4. *Multiply.* All of the rival single-particle bases may be reconciled by reifying *all* of the corresponding branch-bound particles. The particles in all bases are all distinct one from another.

5. *Overlap.* All of the rival single-particle bases may be reconciled by reifying *all* of the corresponding branch-bound particles. But particles associated with different bases are not wholly distinct. In fact particles in different bases overlap in such a way that for each non-GM-entangled branch, the sum of all branch-bound particles in one single-particle basis are *jointly identical* to the sum of all branch-bound particles in any other single-particle basis.

8.3.2 The 'One size fits all' response

The *One size fits all* response is undeniably simple; but its simplicity issues from its rigidity, which is also the source of its drawbacks. It is clear that *One size fits all* easily and satisfactorily solves the basis ambiguity problem for non-GM-entangled fermions and paraparticles, whenever one of the rival decompositions belongs to the categorically privileged basis. But if *none* of the rival decompositions belong to this basis, *One size fits all* comes into conflict with Sections 8.1.1 and 8.1.3's prescription for deciding of which branch-bound particles the assembly in any given state is composed.

For example, consider the two-fermion state

$$\frac{1}{\sqrt{2}}\left(|a\rangle \otimes |b\rangle - |b\rangle \otimes |a\rangle\right), \tag{8.33}$$

where $|a\rangle$ and $|b\rangle$ are both states with compact support centred at locations a and b, respectively, and $\langle a|b\rangle = 0$. According to (V2) in Section 8.1.3, the possible pairs of branch-bound particles that could be said compose the assembly in this state are those for which (8.33) is manifestly non-GM-entangled. These pairs are maximally specific *à la*

$$\frac{1}{\sqrt{1+|z|^2}}\left(|a\rangle + z|b\rangle\right) \quad \text{and} \quad \frac{1}{\sqrt{1+|z|^2}}\left(|b\rangle - z^*|a\rangle\right), \tag{8.34}$$

where $|z| \leqslant 1$, and $0 \leqslant \arg(z) < \pi$ if $|z| = 1$.[17] If for one value of z (and it will be for *at most* one value) the single-particle states in (8.34) belong to the categorically privileged basis under *One size fits all*, then the proponent of *One size fits all* may say that the assembly is composed of the pair corresponding to that value of z. So far, so good. But what if *none* of the rival pairs have states in the privileged basis? (What if, in our example, the privileged basis is position, or momentum?) What ought the proponent of *One size fits all* say then?

A categorically privileged basis is categorically privileged, whether the state is manifestly non-GM-entangled in that basis or not. Therefore the proponent of *One size fits all* rejects the recommendations to the varietist I made in Sections 8.1.1 and 8.1.2 for all non-GM-entangled states, except those that happen to be mainfestly non-GM-entangled in the categorically privileged basis. According to this proposal, the objects of the quantum ontology are all branch-bound particles associated with the same privileged basis (or else they are composed from these

[17] This restriction is a convenient way to identify antipodal points on the Riemann sphere; cf. Stage C of Section 7.1.2.

branch-bound particles). Thus a state that is non-GM-entangled, but not manifestly so in the privileged basis, is treated like any GM-entangled state: namely, as a superposition of "branches" of suitable branch-bound particles.

The *One size fits all* response is clearly *ad hoc*: there is no reason to countenance a categorically privileged single-particle basis, except that it solves the basis ambiguity problem. The proposed basis is not empirically accessible, and it conflicts with the principle that non-GM-entangled states are composed of unsuperposed branch-bound particles. And, of course, there is no uniquely natural suggestion for what the privileged basis would be. (The two proposals that are perhaps the most intuitive, namely position and momentum, create particular trouble, since eigenstates for position or momentum do not exist in the standard Hilbert spaces.) Therefore we turn to the next proposal.

8.3.3 The 'Complicate' response

The next proposal seeks to single out a preferred single-particle basis in a way that better reflects the physics. It takes advantage of the fact that, in realistic scenarios, particles have more than just the locational degree of freedom. Single-particle Hilbert spaces for all actual particles include an internal spin space, and possibly other degrees of freedom, such as flavour and colour. To simplify the discussion, let us ignore these other degrees of freedom and consider only location and spin.

The utility of the spin degree of freedom in solving the basis arbitrariness problem lies in the extra structure it provides in breaking the under-determination of single-particle bases. For, just as two "distinguishable" particles may be entangled, and two "indistinguishable" particles may be GM-entangled, so the individual degrees of freedom associated with a *single* particle may be entangled. Entanglement between the degrees of freedom of a single system is exactly like entanglement between distinct distinguishable systems: the state is entangled iff it is non-separable.

The proponent of *Complicate* stipulates that branch-bound particles only possess states in which distinct degrees of freedom are not entangled; i.e. they may be attributed both a *pure* spatial state and a *pure* spin state. The advantage of this proposal over *One size fits all*, if it succeeds, is that it would be using the physical phenomenon of entanglement to determine a uniquely preferred single-particle basis, so it could not be accused of being *ad hoc*.

Apart from that, *Complicate* shares two important features with *One size fits all*. First, it breaks the under-determination of bases when one of the rivals is non-entangled in its degrees of freedom. For example, in the state

$$\frac{1}{\sqrt{2}}(|L,\uparrow\rangle \otimes |R,\downarrow\rangle - |R,\downarrow\rangle \otimes |L,\uparrow\rangle), \tag{8.35}$$

the single-particle basis $\{|L,\uparrow\rangle, |R,\downarrow\rangle\}$ is uniquely preferred over its rivals, such as $\{\frac{1}{\sqrt{2}}(|L,\uparrow\rangle + |R,\downarrow\rangle), \frac{1}{\sqrt{2}}(|L,\uparrow\rangle - |R,\downarrow\rangle)\}$.[18]

Second, however, the requirement that the single-particle basis be non-entangled may conflict with the principle that non-GM-entangled states are composed of unsuperposed particles. Define an *uncomplicated particle* as a branch-bound particle whose state is non-entangled in its separate degrees of freedom; i.e. a branch-bound particle that may be ascribed a pure state in every degree of freedom. Then there are non-GM-entangled states whose maximally specific particles are *not* uncomplicated. (Call not uncomplicated particles *complicated*.)

For example, the single-particle states

$$|\phi\rangle := \alpha|L,\uparrow\rangle + \beta|R,\downarrow\rangle; \qquad |\chi\rangle := \beta^*|L,\uparrow\rangle - \alpha^*|R,\downarrow\rangle \tag{8.36}$$

(where $|\alpha|^2 + |\beta|^2 = 1$) exhibit entanglement between the spatial and spin degrees of freedom so long as $0 < |\alpha|, |\beta| < 1$; so any particle that is maximally specific *à la* $|\phi\rangle$ or $|\chi\rangle$ is complicated. Now consider the non-GM-entangled (bosonic) state

$$\begin{aligned} |\psi\rangle \;&:=\; \frac{1}{\sqrt{2}}\left(|\phi\rangle \otimes |\chi\rangle + |\chi\rangle \otimes |\phi\rangle\right) && (8.37)\\ &\equiv\; \sqrt{2}\alpha\beta^*|L,\uparrow\rangle \otimes |L,\uparrow\rangle \;-\; \sqrt{2}\alpha^*\beta|R,\downarrow\rangle \otimes |R,\downarrow\rangle \\ &\quad +\; \left(|\beta|^2 - |\alpha|^2\right)\frac{1}{\sqrt{2}}\left(|L,\uparrow\rangle\rangle \otimes |R,\downarrow\rangle + |R,\downarrow\rangle \otimes |L,\uparrow\rangle\right). && (8.38)\end{aligned}$$

[18] Ghirardi *et al* (2002, p. 86) say of these rivals that they are 'of no practical interest'. But I am trying to solve an ontological, not a practical, problem.

This state is a *superposition* of non-GM-entangled branches whose constituent particles are uncomplicated.

Thus *Complicate*, like *One size fits all*, must suspend the recommendations to the varietist that I made in Sections 8.1.1 and 8.1.2 when the non-GM-entangled state is of the wrong type—in this case when the constituent particles could not be uncomplicated.

Here we pick up where our discussion in Section 5.2.1, which applied to the "single-particle" Hilbert space, left off. The idea, under this proposal, is that non-entanglement of separate degrees of freedom trumps non-GM-entanglement of systems, so even a "single-particle" Hilbert space may contain states for which (it may be better to say) there is more than a single particle.

There are two main objections to the *Complicate* response. The first, which is less serious, is that the demand that branch-bound particles be uncomplicated appears to conflict with the fact that complicated "single-particle" states may be used to qualitatively individuate particles (cf. Section 7.2.1). This objection is easily overcome by pointing out that it is already accepted that an individuation criterion may well succeed in picking out more than one branch-bound particle. True, the difference for the proponent of *Complicate* is that the criterion in question is associated with a *one*-dimensional projector, but to claim that branch-bound particles ought to be associated with one-dimensional projectors is simply to assert what was to be proved.

The second, more serious, objection facing *Complicate* is that it will not work for every state, since states exist which continue to suffer a basis arbitrariness even when entanglement between degrees of freedom is taken into account. One such state is the ground state for the two electrons in a Helium atom:

$$|\phi_{1s}\rangle_1 \otimes |\phi_{1s}\rangle_2 \otimes (|\uparrow\rangle_1 \otimes |\downarrow\rangle_2 - |\downarrow\rangle_1 \otimes |\uparrow\rangle_2)\,, \tag{8.39}$$

in which the states for each degree of freedom for *both* particles factorize. (Factor Hilbert space labels serve only to associate the right spin state with the right spatial wavefunction.) Here the demand that the constituent branch-bound particles be uncomplicated helps not one bit in narrowing down the options.[19]

So the varietist must either appeal to one of the other four responses to mop up these cases, or else reject *Complicate* altogether. In fact, *Complicate* will play an important role later, in the discussion of emergentism in Section 9.

8.3.4 The 'Coalesce' response

An alternative strategy in solving the basis arbitrariness problem is to somehow accommodate all of the single-particle bases in which the assembly's state may be expressed. The first response that adopts this strategy is *Coalesce*, which stipulates that branch-bound particles that are maximally specific *à la* states from one basis are identical to branch-bound particles that are maximally specific *à la* states in other bases. Thus branch-bound particles are not only maximally specific *à la* one state, but many.

For example, for the singlet state

$$\frac{1}{\sqrt{2}}\left(|{\uparrow}\rangle \otimes |{\downarrow}\rangle - |{\downarrow}\rangle \otimes |{\uparrow}\rangle\right) = \frac{1}{\sqrt{2}}\left(|{\rightarrow}\rangle \otimes |{\leftarrow}\rangle - |{\leftarrow}\rangle \otimes |{\rightarrow}\rangle\right) \tag{8.40}$$

the proponent of *Coalesce* claims that the branch-bound spin-up particle is identical to *either* the branch-bound spin-left particle *or* the branch-bound spin-right particle, and that the branch-bound spin-down particle is identical to whichever of these is not identical to the spin-up particle.

This may be generalised for any particle number N and any symmetry type. Thus, according to the proponent of *Coalesce*, a non-GM-entangled assembly is composed of N branch-bound particles which are in fact maximally specific *à la* continuum-many single-particle states. GM-entangled assemblies, as per the discussion in Section 8.1.3, are then superpositions of collections of such particles.

Now the objection. Let us ask: In the state (8.40), is the spin-up particle identical to the spin-right or the spin-left particle? That is, is the state composed

[19]This attempt to avoid the basis arbitrariness problem—and the example of the Helium ground state as a counterexample—has also recently been discussed by Bigaj in his lecture to the CLMPS 2011.

of a spin-up-spin-left particle and a spin-down-spin-right particle, or a spin-up-spin-right particle and a spin-down-spin-left particle? The proponent of *Coalesce* must allow both possibilities, so far as probabilities are concerned, since quantum mechanics tells us that $p(\leftarrow|\uparrow) = p(\rightarrow|\uparrow) = \frac{1}{2}$, etc.

Thus the proponent of *Coalesce* is forced to claim that the state (8.40)—and, with it, the entire quantum formalism—is incomplete, since it does not give us complete information about the constituents of the assembly. Rather, the state (8.40) must be interpreted as representing a statistical ensemble of collections of particles of both kinds. Indeed, since there are continuum-many bases in which the state is manifestly non-GM-entangled, the state must represent, for the proponent of *Coalesce*, a statistical ensemble of continuum-many kinds of particle collections.

It would be enough to reject *Coalesce* that, at the outset (Chapter 1), we disavowed any interpretation that entails that the quantum formalism is incomplete. But even despite this *Coalesce* could not possibly work. The problems facing any attempt to interpret quantum probabilities as epistemic are well known; and it can be seen that *Coalesce* runs afoul the Kochen and Specker (1967) no-go theorem for non-contextual hidden-variable theories.

The types of theories addressed by the Kochen-Specker theorem seek to attribute a definite true/false value to every ray in Hilbert space such that the quantum probabilities may be interpreted as arising from statistical ensembles of states corresponding to such attributions. This problem is equivalent to colouring the entire unit sphere in the Hilbert space in black and white so that, for any family of perpendicular points, all but one are painted black. As is now well known, this cannot be done for dimensions of three or more.

Let us now consider a non-GM-entangled N-particle state in which no single-particle state is occupied more than once. For this state, the mathematical problem facing the proponent of *Coalesce* is to attribute one of N particle labels to every ray in the N-dimensional subspace spanned by the component single-particle states. Again, this must be done in a way that is consistent with the quantum probabilities arising from averages over statistical ensembles of assembly states corresponding to such attributions. This problem is equivalent to colouring the entire unit sphere in this subspace with N colours so that, for any family of perpendicular points, all get painted one colour each. But if we label just one of the colours 'white' and the remaining $N-1$ 'black' (consider them as shades of black, as it were), then it is clear that this problem can be solved only if Kochen and Specker's problem can be solved—which it cannot, for $N \geqslant 3$.

The only case for which *Coalesce* escapes the no-go theorem is $N = 2$. But this should come as no consolation, since we were seeking a general solution to the basis arbitrariness problem. We must therefore look elsewhere.[20]

8.3.5 The 'Multiply' response

We could not solve the basis arbitrariness problem by stipulating that branch-bound particles from different bases are identical, so the natural next suggestion is to try the opposite: i.e. to say that any two branch-bound particles from different bases are distinct. Aside from the ontological extravangance of this response, one immediately wonders why it is that the branch-bound particles in different bases are *always* correlated in the same way, to accord with the quantum probabilities. (Why, for example, is there always one spin-left and one spin-right fermion whenever there is one spin-up and one spin-down?) In short, it seems we have a multitude of necessary connections between distinct existences.

Now Hume's dictum, that there are no such connections,[21] has recently been subjected to some serious criticism.[22] It may be argued that necessary connections between distinct existences are perfectly in order, so long as the objects in question are related in the right way, despite being distinct. Being related in the 'right way' might include: one of the objects composing the other (if one believed that composition was *not* identity); or being related by a difference in logical form, like object to universal.

[20]It is perhaps surprising to note that *Coalesce* does not run afoul of Bell's Theorem (1964)—despite (8.40) being the very state that Bell used in his proof. The reason is that Bell's Theorem assumes that the full algebra of quantities on $\mathbb{C}^4$ is available; whereas the varietist imposes the Indistinguishability Postulate, which greatly restricts the available algebra. In particular, joint probabilities for outcomes for different directions of spin cannot meaningfully be calculated, since the individuation criteria for the two particles would not be orthogonal.

[21]Hume (1740, Book I, Part III, §VI) wrote, 'There is no object, which implies the existence of any other if we consider these objects in themselves.'

[22]MacBride (1999), Cameron (2008), Wilson (2008).

Now all these cases provide interesting challenges to Hume's dictum. They may even enforce a limitation on its application. But the current case is clearly not of this type: the branch-bound particles are all taken to be objects, and they are all taken to be wholly distinct. Why should it be that they always appear together as they do?

Another response is available to the proponent of *Multiply*, and it was suggested by Lewis (1992): the connection between the particles may not be absolutely necessary, but only physically necessary. The claim of *Multiply* would then be that, somewhere in the full expanses of logical space, the branch-bound particles appear *without* particles from other bases, and it is only in the QM-worlds (the worlds that obey quantum mechanics) that they, due to physical necessity, appear together.

One objection to this attempt to escape the Humean problem is that the branch-bound particles simply don't exist in non-QM-worlds. This claim will be congenial to those who consider theoretical terms to be implicitly defined by the theory in which they are embedded (e.g. Carnap (1966), Lewis (1970b), Sneed (1971)). Entities that are so defined simply don't exist in worlds in which the theory is not true. However, this clever response fails because, as we have seen in Section 8.2.5, ontological identifications may be made between classical and quantum mechanics. So the proponent of *Multiply* must accept that branch-bound particles can be found in non-QM-worlds.

However, there is a more serious objection arising from considerations of the classical limit, which suggest that a wrong move has been made. As we have seen, in the classical limit, branch-bound particles at precise locations are *identified* with branch-bound particles with precise momenta: these "particles" are precisely the temporal stages of the familiar classical particles. Yet how can two groups of wholly distinct objects suddenly become identical in the classical limit? *Multiply* entails a sudden change where we need a continuous transition. This, together with worries about ontological extravagance and Hume's dictum, points to a natural solution.

8.3.6 The 'Overlap' response

The solution, and the best hope at solving the basis arbitrariness problem, is to retain the absolutely necessary connections between branch-bound particles from different bases, but deny that the particles are wholly distinct. This overcomes the Humean problem, since the particles now fall outside the domain of application of Hume's dictum. It quells concerns about ontological extravagance, since particles from different bases may compose the same assembly, and are therefore jointly identical. Finally, it gets the classical limit right: we may say that branch-bound particles that are maximally specific *à la* location and those that are maximally specific *à la* momentum go from partial overlap to total coincidence in the classical limit.

To illustrate the *Overlap* response, consider again the singlet state

$$\frac{1}{\sqrt{2}}\left(|{\uparrow}\rangle \otimes |{\downarrow}\rangle - |{\downarrow}\rangle \otimes |{\uparrow}\rangle\right) = \frac{1}{\sqrt{2}}\left(|{\rightarrow}\rangle \otimes |{\leftarrow}\rangle - |{\leftarrow}\rangle \otimes |{\rightarrow}\rangle\right). \tag{8.41}$$

The advocate of *Overlap* agrees with *Coalesce*, and against *Multiply*, that the assembly is in the same state no matter with what single-particle basis it is described. But *Overlap* disagrees, both with *Coalesce* that the constituent branch-bound particles are identical, and with *Multiply* that they are wholly distinct. Rather, they partly overlap, so that e.g. the spin-up particle is itself composed of parts of the spin-left and spin-right particles, and such that the sum of the spin-up and spin-down particles is exactly identical to the sum of the spin-left and spin-right particles (and every pair corresponding to every other decomposition). Similarly, the sphere has continuum-many decompositions into hemispheres; yet despite this, there remains just one sphere because hemispheres from different decompositions party overlap.

The move made here is somewhat like the move made in Section 8.1.3, where the varietist accepted the overlap between less-than-maximally specific particles. Here, as there, we may endorse multiple but seemingly conflicting descriptions of the same state by declaring them equivalent. Each description has simply decomposed the same object in diverse ways. Thus the proponent of *Overlap* seeks to totally endorse (V2) in Section 8.1.3 with the reassuring addition that *it* *does not matter* in which single-particle basis the assembly's state is described.

Following on from the calculations in Equation (8.10) and Table 8.1 in Section 8.1.3, we might even propose the squared inner product of two states as a *measure* of overlap for two branch-bound particles in those states. Since, e.g. $|\langle\uparrow|\leftarrow\rangle|^2 = |\langle\uparrow|\rightarrow\rangle|^2 = \frac{1}{2}$, we might say that the spin-up particle is composed of *exactly half* of the spin-left and spin-right particles. This measure has the desirable property that, given any single-particle state, its degree of overlap with all of the single-particle states in any chosen basis sums to unity.

Mereological overlap has been marshalled in a variety of other areas to solve similar problems. Armstrong (1978, pp. 120-4) declares resembling universals to be part-identical as a way of accounting for their resemblance without having to posit higher-order universals, which would lead to an infinite regress. Lewis (1993) attempts to solve Unger (1980) and Geach's (1980) 'Problem of the Many'—namely, that any (putatively!) single macroscopic object may be identified with a surfeit of distinct microphysical sharpenings with clearly demarcated boundaries—by pointing out that the various sharpenings are all 'almost identical', and claiming that 'almost identity' is good enough for macroscopic objects, for most purposes. So: has mereological overlap come to the rescue again?

There is an important difference between, on the one hand, the uses of overlap in Section 8.1.3, by Lewis (1993), and perhaps even by Armstrong (1978); and on the other hand, the use of overlap in this case. The difference is that, in all the former cases *it is clear what the parts are* whose shared composition entails the overlap.

Overlap is part-identity, and part-identity is identity of parts. So any claim of overlap must be supported by a specification of the parts that are shared between the putatively overlapping entities. In the case of Section 8.1.3, the parts of the overlapping particles are the maximally specific, branch-bound particles in some basis.[23] In the case of Lewis (1993), the parts of the overlapping sharpenings are

[23] Pleasingly, the claim that the less-than-maximally specific particles overlap at certain branch-bound particles is reflected perfectly in the geometry of Hilbert space, in which multi-dimensional subspaces associated with the the less-than-maximally specific particles overlap precisely at the corresponding one-dimensional subspaces that are associated with branch-bound particles.

the mereological sums of microphysical constituents, naïvely conceived. In the case of Armstrong (1978), the parts of the overlapping universals are the simpler universals of which the original universals are logical constructions. What are the parts of the putatively overlapping branch-bound particles from different bases?

Branch-bound particles cannot have other particles as parts, since they are already maximally specific. What else is there? One remaining suggestion is that the parts are the parts of wavefunctions that define the branch-bound particles. But this suggestion has the wrong results, since there are many even functions $f(x)$ that overlap some odd function $g(x)$; yet their inner product, as defined by $\int_{-\infty}^{\infty} f^*(x)g(x)\,\mathrm{d}x$, is always zero, if it exists.

From here there appear to be only two routes ahead. Either we reject *Overlap*, which, since it is the final hope to save varietism, entails a rejection of varietism; or else we attempt to make sense of the overlap between branch-bound particles from different bases as something other than mereological. The problem with pursuing the latter route is that it is only *mereology* that, by being a part of logic,[24] can overcome the worries of ontological extravagance, the violation of Hume's dictum, and the discontinuity in the classical limit, all of which faced *Multiply.* Therefore, I reluctantly take the former route, and reject varietism. However, its merits were undeniable, so we might hope to find a nearby alternative that can overcome varietism's fatal defect. With this, I turn to the final proposal for the target concept of particle, *emergentism.*

[24]Of course, it is controversial that mereology *is* logic. But even if mereology is not logic, it is hard to see that any other notion of composition (if any exists) can be claimed to be logical. These matters, of course, deserve a far more thorough study than there is space for here.

Chapter 9

Emergentism: winner in a poor field?

In this short, final Chapter, I present an anti-factorist alternative to varietism. Its main feature is that is gives up on the desideratum that particles always compose the assembly. That is, it denies the strong version of compositionality (cf. Section 5.1.4); i.e., the claim that the assembly is a "derived" or "secondary" object, in the sense of being identical to the mereological sum of its constituent particles. And it also denies the weak version of compositionality; i.e., the claim that the properties of the assembly supervene on the properties and relations of the particles. For it claims that there are states in which the assembly exists but the particles do not.

Instead, the assembly is taken as "fundamental" or "primary", in the specific sense that each state in the assembly Hilbert space is construed as representing an ascription of properties to the assembly taken as a whole. Therefore, particles, if they exist, are construed as "derived" or "secondary" entities, in the sense of being emergent properties of *some* of the assembly's states. For this reason, the proposal is called *emergentism*. This will be a somewhat disappointing *dénouement*: I reluctantly plump for emergentism for want of a better alternative.

In Section 9.1, emergentism is more precisely defined. I will argue that its most attractive version—what I call *assembly realism*—may usefully be conceived as a version of a field ontology found in the philosophy of QFT, albeit a version restricted to the limit of conserved total particle number. Section 9.2 concludes the Chapter and the dissertation. After briefly recalling the main results of previous Chapters, I address the question whether particles under emergentism really are particles worthy of the name, especially given that some form of compositionality (cf. Section 5.1.4) is denied of them. I conclude that this question has no clear cut answer, but that some version of emergentism is the best proposal available for what particles are in quantum mechanics.

9.1 Emergentism defined

The problem of basis arbitrariness facing varietism, discussed in Section 8.3, may be summarised in the following way. *Prima facie*, there is a single object—namely, the assembly—whose properties are represented by the states in Hilbert space. The varietist (and the factorist) wishes to consider this object to be complex, i.e. composed of simpler objects—namely, particles, as functionally defined by the desiderata in Section 5.1. However, *several* (indeed, continuum-many) putatively natural decompositions of the assembly exist, where the naturalness of decompositions is governed by considerations of some version of entanglement between the constituent systems, whose own states must be meaningfully ascribable under the imposition of the Indistinguishability Postulate.

The existence of several decompositions is problematic, because it seems impossible to interpret the particles issuing from each decomposition as part-identical, which would have permitted the interpretation that they are alternative decompositions of the same assembly (cf. Section 8.3.6). Yet *unless* we can make sense of rival decompositions amounting to the same thing, we are not licensed to claim what is *prima facie* obvious: that there is a *single* entity, the assembly, being decomposed.

Emergentism arises out of two possible retreats in the face of this problem. I define emergentism as the disjunction of the two, because they concur about one significant feature of particles: namely, that particles exist as (higher order) properties of some other object or objects.

The first retreat, which I discuss in Section 9.1.1, is to deny the original assumption above: namely, that the assembly is a single object, of which properties may be predicated. On this view, the assembly—and, with it, particles—are the (higher-order) properties of some other objects, perhaps spatial regions. I call this position *mode realism*. However, as I will argue, this position fares no better than the varietist's *One size fits all* response (cf. Section 8.3.2), since it unjustifiably privileges one single-system basis over another.

The second retreat, which I discuss in Section 9.1.2, is to accept the seemingly undeniable assumption that the assembly is a single object, but to deny that it may be categorically decomposed into simpler objects. On this view, particles are (higher-order) properties of the assembly, but they exist only for certain states of the assembly—in short, those for which basis arbitrariness can be overcome. I call this position *assembly realism*. I will argue that, though *prima facie* strange, assembly realism meshes best with the ontology of quantum field theory.

9.1.1 Mode realism

The idea that, in elementary quantum mechanics, one might treat the modes as the basic objects is not unfamiliar (cf. e.g. Saunders 2006). But it is more familiar in the philosophy of quantum field theory. This is easily understandable, once one considers the formal properties of the Fock space.

To illuminate these properties, consider a Fock space formed from a *finite*-dimensional Hilbert space. (So, strictly speaking, I am not considering a quantum *field*, which has infinitely many degrees of freedom. Nor am I considering a Fock space generated from a *particle*'s Hilbert space; since a particle must have location (cf. Section 5.1.2), and so has an infinite-dimensional Hilbert space.) For example, the Fock space for fermions whose single-system Hilbert space is $\mathbb{C}^d$ is

$$\begin{aligned}
\mathfrak{F}(\mathbb{C}^d) &= \mathbb{C} \oplus \mathbb{C}^d \oplus \mathcal{A}(\mathbb{C}^d \otimes \mathbb{C}^d) \oplus \mathcal{A}(\mathbb{C}^d \otimes \mathbb{C}^d \otimes \mathbb{C}^d) \oplus \ldots \\
&= \mathbb{C} \oplus \mathbb{C}^d \oplus \mathbb{C}^{\frac{1}{2}d(d-1)} \oplus \mathbb{C}^{\frac{1}{6}d(d-1)(d-2)} \oplus \ldots \\
&= \bigoplus_{N=0}^{\infty} \mathbb{C}^{\binom{N}{d}C} \qquad (9.1)
\end{aligned}$$

where ${}^{N}_{d}C := \frac{N!}{d!(N-d)!}$. The Nth summand in (9.1) corresponds to the Hilbert space for an assembly of N fermions. But this tensor sum may be re-expressed as follows:

$$\mathfrak{F}(\mathbb{C}^d) = \bigoplus_{N=0}^{\infty} \mathbb{C}^{\binom{N}{d}C} = \mathbb{C}^{\sum_{N=0}^{\infty}\binom{N}{d}C} = \mathbb{C}^{2^d} = \bigotimes^{d} \mathbb{C}^2 \ . \tag{9.2}$$

So a Fock space for d-level fermions is equivalent to the Hilbert space for d *distinguishable* 2-level quantum systems. These systems correspond one-to-one to the rays of the d-dimensional single-fermion Hilbert space, in some basis. In that basis, each ray represents the possession by a fermion of a particular eigenvalue (associated with that ray) for some single-fermion quantity whose eigenbasis is the basis in question. Thus, each of the d Hilbert spaces $\mathbb{C}^2$ is taken to represent, as it were, the space of possibilities associated with an eigenvalue. These possibilities are given by the possible number of fermions which may possess that eigenvalue. And since fermions are subject to Pauli exclusion, there are only two possibilities for each eigenvalue: 0 or 1; thus the Hilbert space for each is $\mathbb{C}^2$.

Since the Hilbert space on the right-hand side of Equation (9.2) is a tensor product Hilbert space, it is tempting to think of it as having a natural decomposition into its factor Hilbert spaces. After all, recall (Section 6.1) that I endorse factorism for distinguishable systems, i.e. systems whose joint Hilbert space has a tensor product structure. Thus the temptation is to reify the eigenvalues mentioned in the previous paragraph. The states in Fock space $\mathfrak{F}(\mathbb{C}^d)$ are then construed, not as representing property attributions to a variable (indeed, possibly *unsharp*) number of systems; rather, they are construed as representing property attributions—more specifically, occupation numbers (which may be unsharp)—to a *fixed* number of these reified eigenvalues. Modes are simply these reified eigenvalues.

If we now return to elementary quantum mechanics, the mode ontology becomes rather messy. This is because, in elementary quantum mechanics, the assembly's Hilbert space is a *single summand* of the full Fock space in (9.1), and the breaking up of the Fock space into these summands is not particularly natural from the point of view of its decomposition into factor Hilbert spaces for each mode. Indeed, the demand that the "total system number" (the operator $N(\Omega, S)$, from

Section 8.2.5, where now (Ω, S) encompasses all single-system states) be conserved means that the states of the modes must be strictly correlated so as to ensure that the sum of all occupation numbers is a constant integer.

But this is no argument against mode realism: after all, elementary quantum mechanics *is* the conserved "total system number" limit of quantum field theory. The correlation between the modes that ensures a constant "total system number" need not be interpreted as anything more than a *dynamical* phenomenon, whose contingent occurrence underpins the practical convenience of using elementary quantum mechanics over a more complicated theory that uses Fock space.

Nor can it be objected against mode realism that its ontology is too *weird.* For, although the phrase 'reified eigenvalue' might cause more than a little hesitation, reified eigenvalues are actually much more familiar than it may seem. A compelling, but troublesome,[1] example is location. The modes associated with the position quantity $\mathbf{Q}$ are nothing but *locations*; i.e., roughly speaking, spatial points.

Thus one particularly vivid instance of mode realism would be what Wallace and Timpson (2010) have recently presented as *spacetime state realism.* According to this view, any state in the assembly's Hilbert space is interpreted as representing the attribution of properties and relations (encoded in density operators) to spatial (or spacetime) regions. A typical state containing a single particle is then a (typically entangled) state of the various spatial regions in which the sum of all occupation numbers is 1 (Wallace and Timpson (2010, p. 711)). More generally, particles are construed as certain patterns of "excitation" (related to an increase of 1 in an occupation number) in the joint state of the various location. That is: particles are properties of the properties and relations between locations.

I will not go into the details of spacetime state realism any further. Instead, I turn to an objection against mode realism (and hence spacetime state realism) that I cannot see a good response to. The objection is that mode realism, like varietism, suffers from its own basis arbitrariness problem.

[1]Location is troublesome because Hilbert spaces do not contain eigenstates for location. Nevertheless, location *does* have a spectral decomposition into a family projectors, all of which *do* act on the particle Hilbert space.

For consider again the right-hand side of Equation (9.2). Now, while it is true that the Fock space may be decomposed into d 2-level Hilbert spaces; there are, in fact, continuum-many ways of doing this. Specifically: a choice of a complete orthobasis in the single-system Hilbert space $\mathbb{C}^d$ determines a unique decomposition of the Fock space into d "single-mode" Hilbert spaces, such that each vector in the chosen single-system basis corresponds to one of the d single-mode Hilbert spaces. But there are continuum-many complete orthobases in $\mathbb{C}^d$.

For example, let us consider a Fock space for 2-level fermions. For the sake of vividness, suppose the degree of freedom in question is spin for a spin-$\frac{1}{2}$ system. Then the Fock space is $\mathbb{C}^4$. A choice of modes now takes the form of a choice of direction of spin. A change of spin direction then induces a transformation *between* the single-mode Hilbert spaces.

Writing $|n\rangle_\xi$ for the state in which the mode ξ is occupied n times,

$$|1\rangle_\uparrow \otimes |0\rangle_\downarrow = |\!\uparrow\rangle = \frac{1}{\sqrt{2}}\left(|\!\rightarrow\rangle + |\!\leftarrow\rangle\right) = \frac{1}{\sqrt{2}}\left(|1\rangle_\leftarrow \otimes |0\rangle_\rightarrow + |0\rangle_\leftarrow \otimes |1\rangle_\rightarrow\right) \ . \tag{9.3}$$

So a state that is non-entangled for one choice of modes is entangled for another choice of modes (cf. Zanardi (2001, p. 1)). This adds a new twist to the old basis arbitrariness problem of Section 8.3; since GM-entanglement, unlike the entanglement between modes shown here, is a basis-*independent* matter.

One might now attempt to redeploy one or more of the five responses to the basis arbitrariness problem discussed in Section 8.3. However, *Complicate*—which relies on many degrees of freedom—could not possibly work, since modes have occupation number as their *only* degree of freedom. Apart from that, the results are the same as for varietism: i.e., the responses all fail to solve the problem.

Incidentally, I note that Wallace and Timpson (2010) effectively opt for the *One size fits all* response (Cf. Section 8.3.2), since they favour the position representation. Their argument is based on the intelligibility that the 'spatial arena' affords (Wallace and Timpson (2010, p. 724)). But it is not clear (to me, at least) why the momentum representation is any less intelligible; and the existence of more than one natural decomposition into modes—whether there are continuum-many or just two—is enough to be problematic for the mode realist. However, I should emphasise, in defence of Wallace and Timpson, that they present spacetime state realism, not as the *single best* ontology for quantum mechanics, but rather as a rival to a particular interpretation they wish to criticise, namely wave-function realism (Wallace and Timpson (2010, pp. 702-6)).

This concludes my discussion of mode realism. Since it runs afoul of its own basis arbitrariness problem, I turn to the remaining option, assembly realism.

9.1.2 The assembly is the object

I claim that, like mode realism, assembly realism has a more familiar counterpart in quantum field theory. Specifically, my claim is that assembly realism is the conserved "total particle number" limit of *field realism*, the view that the quantum field is the basic object.

Is it correct to liken an assembly to a quantum field? If a 'field' is defined as having infinite degrees of freedom (Huggett and Weingard, (1994b, p. 295)), then we may think that an assembly is unequivocally *not* a field. But why can we not identify the assembly with the quantum field *under the condition* of conserved "total particle number"? Just as modes appear to have strange correlations in the quantum-mechanical limit—which may be interpreted as a dynamically-induced phenomenon—couldn't the assembly's limited degrees of freedom be seen too as nothing but a dynamical restriction?

An objection to this that can be dismissed out of hand is that the assembly is the mereological sum of its particles, and so could not have more degrees of freedom than are allowed by those of the particles. We can dismiss this objection because it assumes what the assembly realist has already denied: namely, that the assembly is the mereological sum of more basic objects. Assembly realism means treating the assembly as basic, so it is still up for grabs whether or not its degrees of freedom are constrained dynamically or by some stronger form of necessity.

A second objection is that the identification of the assembly with a dynamically restricted quantum field could not be supported by a consideration of quantum mechanics alone, since the very idea of a quantum field is external to quantum mechanics. But this objection does too much; for we *want* to be able to appeal to a theory's "neighbours" when interpreting it (cf. our desideratum of inter-theoretic applicability in Section 5.1.5). And it must be noted that a positive case *for* the identification of the assembly with a dynamically restricted quantum field is, of course, ontological continuity in the QFT limit.

Furthermore, it is not surprising that the metaphysics issuing from an interpretation of a formalism should point to possibilities, intelligible within that metaphysics, that the formalism neglects to represent. We saw an example of this very phenomenon in Section 8.1.3, where I claimed that superpositions of variable particle number are easily accommodated by the varietist.

Treating the quantum *field* as the basic, or "fundamental" object in quantum field theory, is a doctrine with many adherents.[2] For example Wald (1994, p. 46): 'quantum field theory is the quantum theory of a field, not a theory of "particles".' And Clifton and Halvorson (2001, pp. 459): 'quantum field theory is "fundamentally" a theory of a field, not particles ... this view makes room for the reality of quanta, but only as a kind of epiphenomenon of the field associated with certain functions of the field operators.' Also Malament (1996, p. 1):

> ... in the attempt to reconcile quantum mechanics with relativity theory ... one is driven to a field theory; all talk about particles has to be understood, at least in principle, as talk about the properties of, and interactions among, quantized fields.

(And cf. also Huggett and Weingard (1994a, 1994b, 1996).)

I emphasise that all the above quotes express support for a field ontology as a response to distinctly *quantum-field-theoretic* phenomena: namely, the Unruh effect (in the case of Wald and Clifton and Halvorson) and failures of localizability (in the case of Malament). Yet we have been pushed to a similar position—we may call it the quantum mechanical limit of field realism—not directly because of any quantum-field-theoretic phenomena, but for want of a better alternative.

[2] Though, of course, opinion is not unanimous. Notable critics of the field ontology are Teller (1995, Ch. 3) and Baker (2009).

Of course, the idea that the assembly has no decomposition makes the very use of the term 'assembly' Pickwickian, but that is no objection. A far more serious objection is that the claim that quantum assemblies have no proper parts is simply incredible. We have an understanding of single- and multi-particle states, or at least we thought we did. What has become of particles?

9.1.3 Regaining particles under assembly realism

One of the major advantages of assembly realism is that one need not be committed to the idea that particles exist in every state. I combine this fact with what was a promising response to varietism's basis arbitrariness problem—namely *Complicate* (cf. Section 8.3.3)—to give an account of the emergence of particles under assembly realism.

Recall that, for the proponent of *Complicate*, particles are 'uncomplicated' maximally specific systems. That is, they are maximally specific systems for which their separate degrees of freedom are not entangled. The idea is to use the requirement of non-entanglement of separate degrees of freedom to overcome the under-determination of single-particle bases. This response works for many states, but fails to overcome the under-determination for *all* states of the assembly.

But assembly realists, unlike varietists, are not required to recover particles for all states of the assembly. Thus the assembly realist is at liberty to stipulate that, *whenever particles exist*, they are the *unique*, uncomplicated, maximally specific systems; and that in all other states of the assembly, there simply aren't any particles.

An enormous advantage of this proposal is that it manages to incorporate the varietist's 'branch-bound' particles, whose merits I discussed in Section 8.2, without succumbing to the basis arbitrariness problem. Thus we may carry over, *mutatis mutandis*, *almost* all of the considerations from Section 8.2, which applied to varietist particles, to the assembly realist's emergent particles.

One important such result is the recovery of classical particles in the classical limit. For, in the classical limit, the varietist's particles are not only non-GM-entangled (cf. Section 8.2.5); they are also non-entangled in their separate degrees of freedom; that is, they are uncomplicated, as *Complicate* requires. This follows from the fact that, in the usual study of the $\hbar \to 0$ limit, we consider a series of ever-narrowing Gaussians centred at a single point in the (reduced) configuration space (Landsman (2007, §5)).

Of course, it must be emphasised that, even in the classical limit of assembly realism, particles do not become *objects*, i.e. the subjects of predication. Rather, they become localized "spikes" in the assembly's density function (like those discussed by Redhead (1987, pp. 10-11)). Besides, particles exist in *all classical* limit states—even for assembly realism—since in the classical limit the basis arbitrariness problem vanishes (cf. the end of Section 8.3.5), and so there is a unique collection of uncomplicated maximally specific systems. Furthermore, in classical limit states of the assembly, the assembly state (begin non-GM-entangled) is determined by—i.e. it supervenes on—the properties of these "spikes". So, in the classical limit, assembly realism even manages to recover weak compositionality.

However, two issues remain for assembly realism. The first issue is that, by denying that the assembly is composed of constituent particles, assembly realism contradicts one of our desiderata for the concept of particle, namely compositionality (cf. Section 5.1.4). I deal with this issue in the next, and final, Section.

The second issue is that the claim that particles exist in some states of the assembly and not others seems to make them ontologically redundant. For the assembly exists in all states, and its properties suffice to make true or false any statement made in the quantum formalism. What work are the particles doing?

We met this issue in Section 8.1.3, in the discussion of 'cagey' varietism. I make the same reply here: the particles are not idle objects, but ontological free-riders, once the assembly and its properties are taken into consideration. For the particles *are* certain features of the assembly's state; they are not additions to the field ontology.

9.2 Conclusion: Is emergentism good enough?

In this Section, I will consider whether emergentism—in the specific form of assembly realism—really does provide a good enough concept of particle. But first allow me to recapitulate the main results of this dissertation.

1. Following Quine (e.g. 1976), instances of discernment may be separated into four kinds: internal, external, relative and weak. The disjunction of internal and external discernibility—absolute discernibility—is related to the notion of individuality: specifically, an individual is an object that is absolutely discernible from every other object in the domain. Discernibility is linked to the existence of symmetries on the domain of quantification. For any structure: if a permutation π leaves invariant the indiscernibility equivalence classes, then π is a symmetry (but not necessarily *vice versa*); and if π is a symmetry, then it leaves invariant the *absolute* indiscernibility equivalence classes (but not necessarily *vice versa*; cf. Theorem 1, Section 2.4). Furthermore, for any finite structure: any two objects are absolutely indiscernible iff there is some symmetry that maps one to the other (Theorem 2, Section 2.4).

2. According to a weak version of the identity of indiscernibles (WPII) and a new metaphysical thesis called QII, individuality is conceptually distinct from identity; for (according to those two theses) there may be non-individual objects, whereas *every* object is self-identical. And according to QII and haecceitism, identity is distinct from indiscernibility; for (according to *these* theses) there may be utterly indiscernible, yet distinct objects. (Cf. Chapter 3.)

3. Following Haslanger (2006), a term may be associated with three concepts: its avowed concept (the concept we take it to connote); its operative concept (the concept which our use of the term reveals); and its target concept (what would be a good thing to mean by the term). The project of explication, or conceptual reform, for a given term may be characterised as bringing that term's avowed and operative concepts in line with its target concept. (cf. Section 4.1.)

This framework may be adapted for the purposes of conceptual reform in physical theories in the following way. Interpreting a physical theory is a matter of laying down a representation relation between the theory's formalism and physical items (including, perhaps, non-actual physical items). Thus the formalism is afforded physical content by referring to mathematical items that, in turn, represent physical items. A precise concept may be identified with an intension; i.e., roughly speaking, a function from possible physical worlds to extensions, which are physical items, or sets of them. Thus laying down a representation relation associates concepts with elements of the theory's formalism. Conceptual reform is thereby linked to finding the best representation relation between the mathematics and the physics. (Cf. Section 4.2.)

A theory's formalism provides its own standard of naturalness with which to identify the target concept for a given term. But this claim does not commit one to a Lewisian (1983) ontology of perfectly natural properties and relations. The identification of a target concept for a term is constrained, not only by naturalness, but also by that term's operative concept. For otherwise the target concept could not count as an explication of that term. (Cf. Section 4.1.2.)

4. The operative concept of particle, applied generally, has five main strands. (1) A particle is a physical item. (2) A particle has location; so, inspired by Wigner (1939), its state space ought to provide a basis for a representation of the spacetime symmetry group. (3) A particle persists over time, but momentary particle-stages may exist even if there are no uniquely natural trans-temporal particles. (4) Particles compose assemblies, in the weak sense that an assembly's properties supervene on the properties and relations possessed by its particles; or perhaps in the stronger sense that an assembly is a mereological sum of its particles. (5) Particles are trans-theoretic entities; so particles from "neighbouring" theories ought to coincide in the limit that the mathematics of one theory tends to another. (cf. Section 5.1.)

   These strands of meaning are consistent with identifying *several* particles in certain states of what is usually called a 'single-particle' Hilbert space. This occurs whenever the Hilbert space accommodates internal, as well as spatial, degrees of freedom. (Cf. Section 5.2.)

5. The "local" operative concept of particle, as applied in quantum mechanics, associates particles with factor Hilbert spaces; this conflicts with the general operative concept. The view that the "local" operative concept is also the target concept of particle is called *factorism*. Factorism is distinct from haecceitism, since factorism is purely a thesis about what particles *are*, while haecceitism is a thesis about whether permutations of particles generate a physical difference. (Cf. Section 6.2.)

   If the Indistinguishability Postulate (IP) is imposed, then factorist particles of the same species are all absolutely indiscernible one from another, and are therefore all non-individuals. However, they may still be weakly discerned, by appealing to multi-particle quantities which register anti-correlations between singe-particle states. (Cf. Section 6.3.)

   Factorism is false for so-called "indistinguishable" particles—i.e. particles for which IP is imposed—since it makes an error analogous to reifying the average taxpayer. This is shown by their non-individuality, which entails that they cannot be picked out in language or in thought, and they cannot be associated with definite spatio-temporal trajectories in the classical limit. (Cf. Section 6.4.)

6. Anti-factorism prompts a new look at the notions of entanglement, assembly decomposition, and system individuation. The definition of entanglement proposed by Ghirardi *et al* (2002)—what I call *GM-entanglement*—is congenial to the anti-factorist. According to their account, an assembly's state is non-GM-entangled iff it is obtained by (anti-) symmetrizing a product state. If an assembly is non-GM-entangled, its state supervenes on the single-particle states of which it is constructed, together with that assembly's symmetry type (boson or fermion). (Cf. Section 7.1.)

   So-called "indistinguishable" systems may in fact be individuated (that is, picked out in language) using single-system projectors—what I call *qualitative individuation*. The Hilbert space for an assembly of "indistinguishable" systems cannot be naturally decomposed *in toto*; but *subspaces* of it *can* be naturally decomposed. The resulting constituents are precisely the qualitatively individuated systems. The failure of each qualitatively individuated system to exist in all states in the assembly's Hilbert space, and the flexibility in choosing individuation criteria, point to a quantum version of Lewis's (1968) counterpart theory for such systems. (Cf. Section 7.2.)

7. A promising anti-factorist proposal for the target concept of particle is *varietism*, the "average" of whose particles may be identified with the factorist's particles. The basic objects of the varietist ontology are branch-bound particles: particles which possess pure states and which compose non-GM-entangled branches of any state of the assembly. The qualitatively individuated systems of Chapter 7 may be identified with certain natural mereological sums of branch-bound particles. (Cf. Section 8.1.)

   Varietist particles satisfy the five strands of meaning of the operative concept of particle, laid out in Chapter 5. In the classical limit, branch-bound particles may be identified with classical particle-stages, and branch-bound particles may be identified with the quanta in QFT. Branch-bound fermions are always absolutely discernible, but branch-bound bosons and paraparticles may be utterly indiscernible. That is: the weak discernibility results established for factorist particles in Chapter 6 *cannot* be carried over to varietist branch-bound particles. (Cf. Section 8.2.)

   Varietism fails because it cannot escape a basis ambiguity problem. That is, for fermionic and paraparticle states, it is under-determined *which* branch-bound particles compose the assembly in a given non-GM-entangled state. The most promising attempt to escape the problem—to declare that branch-bound particles from different bases are part-identical—fails, since no account can be given of what these parts are supposed to be. (Cf. Section 8.3.)

8. Thus we are led to our second and last anti-factorist proposal for the target concept of particle, *emergentism*. On this view, particles do not compose the assembly, even in the weak sense. Indeed, particles are not even *objects*,

according to emergentism. This proposal bifurcates: either the modes are treated as the basic objects (mode realism), or else the entire assembly is treated as the basic object (assembly realism). Mode realism runs afoul of its own basis arbitrariness problem, since it is under-determined *which* modes ought to be reified. Assembly realism suffers no such problem, and has ontological continuity with a popular way of interpreting quantum field theory. According to assembly realism, particles are like the varietist's branch-bound particles, with the added restriction that they are non-entangled in their separate degrees of freedom; except that they are construed as higher-order *properties* of the assembly. Thus the assembly realist's particles obey most of the strands of meaning of the operative concept of particle; with the notable exception that they can no longer be said to compose the assembly.

So assembly realism provides the best hope for an explication of the concept of particle in quantum mechanics. But is it good enough? That is: is the target concept proposed by the assembly realist sufficiently similar to the operative concept of particle, discussed in Chapter 5, to warrant counting it as an explication of 'particle'?

A mark against is that, even though assembly realism manages to recover weak compositionality in the classical limit, the stationary states of electrons in molecular orbitals do not belong to the classical limit. And in these states there is no unique collection of uncomplicated, maximally specific systems. (Recall the example of the ground state of Helium, at the end of Section 8.3.3.) Thus assembly realism is committed to the claim that, in realistic cases, Austin's 'moderate-sized specimens of dry goods' are not composed of particles.

A mark in favour is that, like the varietist's particles, the assembly realist's particles satisfies many of the operative concept's strands of meaning. Specifically: they are undeniably physical; they may be attributed a location (at least as a degree of freedom); and they enjoy ontological continuity with both quantum field theory and classical mechanics. They also satisfy some (at least Weyl's (1928)) intuitions regarding discernibility. That is: fermions, whenever they exist, are always absolutely discernible; while bosons and paraparticles may be utterly indiscernible.

Another consideration in favour, of course, is that no viable alternative seems to exist. The factorist cannot offer particles that are anything other than non-individuals; and the varietist cannot tell us which branch-bound particles to decompose non-GM-entangled states into. Assembly realism, even if it is not good enough, is the best we can get.

I submit that we have here a case of semantic indecision (what I called in (v) of Section 4.1.2 a *hard* case). Therefore all I can do is present the unhappy facts: and it has been the aim of this dissertation to establish those facts. Whether we now decide to apply the term 'particle' to the assembly realists' offerings is, as Lewis (e.g. 1995) might have said, a purely political matter. To pretend to decide the matter either way would be to succumb to linguistic legislation; and there is no philosophy in that.

References

Adams, R. M. (1979), 'Primitive Thisness and Primitive Identity', *The Journal of Philosophy* **76**, pp. 5-26.

Adams, R. M. (1981), 'Actualism and Thisness', *Synthese* **49**, pp. 3-41.

Armstrong, D. M. (1978), *Universals and Scientific Realism.* Cambridge: Cambridge University Press.

Austin, J. L. (1961), 'The Meaning of a Word', in his *Philosophical Papers*, ed. by J. Urmson and G. Warnock, Oxford: Clarendon Press, pp. 55-75.

Austin, J. L. (1962), *Sense and Sensibilia.* Oxford: Oxford University Press.

Baker, D. (2009), 'Against Field Interpretations of Quantum Field Theory, *British Journal for the Philosophy of Science* **60**, pp. 585-609.

Bassi, A. and Ghirardi, G. (2003), 'Dynamical Reduction Models' *Physics Reports* **379**, p. 257-426.

Bell, J. S. (1964), 'On the Einstein-Podolsky-Rosen paradox', *Physics* **1**, pp. 195-200.

Belnap, N. (1993), 'On Rigorous Definitions', *Philosophical Studies*, **72**, pp. 115-146.

Belot, G. (2001), 'The Principle of Sufficient Reason', *The Journal of Philosophy* **98**, pp. 55-74.

Belot, G. (2003), 'Notes on Symmetries', in K. Brading and E. Castellani (eds), *Symmetries in Physics: Philosophical Reflections*, Cambridge: Cambridge University Press, pp. 393-412.

Bengtsson, I. and Życzkowski, K. (2006), *Geometry of Quantum States: An Introduction to Quantum Entanglement.* Cambridge: Cambridge University Press.

Bigelow, J. (1988), *The Reality of Numbers*, Oxford: Oxford University Press.

Black, M. (1952), 'The Identity of Indiscernibles', *Mind* **61**, pp. 153-64.

Black, R. (2000), 'Against Quidditism', *Australasian Journal of Philosophy* **78**, pp. 87-104.

Boolos, G. S. and Jeffrey, R. C. (1974), *Computability and Logic.* 4th edition. London: Cambridge University Press.

Braddon-Mitchell, D. and Nola, R. (2009), 'Introducing the Canberra Plan', in Braddon-Mitchell and Nola (eds.), *Conceptual Analysis and Philosophical Naturalism*, Cambridge: MIT Press, pp. 1-20.

Brown, H. R. (2006), *Physical Relativity.* Oxford: Oxford University Press.

Butterfield, J. N. (1993), 'Interpretation and identity in quantum theory', *Studies in the History and Philosophy of Science* **24**, pp. 443-76.

Butterfield, J. N. (2001), 'Some Worlds of Quantum Theory', in R. Russell *et al* (eds.), *Quantum Physics and Divine Action*, Vatican Observatory Publications.

Butterfield, J. N. (2006a), 'On Symplectic Reduction in Classical Mechanics', in J. Earman and J. N. Butterfield (eds), *The Handbook of Philosophy of Physics*, Amsterdam: North-Holland, pp. 1-131.

Butterfield, J. N. (2006b), 'The Rotating Discs Argument Defeated', *The British Journal for the Philosophy of Science*, **57**, pp. 1-45.

Butterfield, J. N. (2011), 'Against Pointillisme: A Call to Arms', in D. Dieks, W. J. Gonzalez, S. Hartmann, T. Uebel and M. Weber (eds.),*Explanation, Prediction, and Confirmation: The Philosophy of Science in a European Perspective* Vol. 2, Dordrecht: Springer, pp. 347-365.

Button, T. (2006), 'Realistic Structuralism's Identity Crisis: A Hybrid Solution', *Analysis* **66**, pp. 216-22.

Cameron, R. P. (2008), 'Recombination and Intrinsicality', *Ratio* **21**, pp. 1-12.

Carnap, R. (1934), *The Logical Syntax of Language.* London: Routledge & Kegan

Paul.

Carnap, R. (1950), *The Logical Foundations of Probability.* Chicago: Chicago University Press.

Carnap, R. (1956), *Meaning and Necessity: A Study in Semantics and Modal Logic.* Chicago: University of Chicago Press.

Carnap, R. (1966), *Philosophical Foundations of Physics: An Introduction to the Philosophy of Science.* New York: Basic Books Inc.

Castellani, E. (ed.) (1998), *Interpreting Bodies: Classical and Quantum Objects in Modern Physics.* Princeton: Princeton University Press.

Caulton, A. and Butterfield, J. N. (2011), 'Symmetries and Paraparticles as a Motivation for Structuralism', forthcoming in the *British Journal for the Philosophy of Science.*

Chang, H. (1996), 'The Quantum Counter-Revolution: Internal Conflicts in Scientific Change', *Studies in the History and Philosophy of Modern Physics* **26**, pp. 121-136.

Clifton, R. and Halvorson, H. (2001), 'Are Rindler Quanta Real? Inequivalent Particle Concepts in Quantum Field Theory', *The British Journal for the Philosophy of Science* **52**, pp. 417-470.

da Costa, N. and Krause, D. (2007), 'Logical and Philosophical Remarks on Quasi-set Theory', *Logic Journal of the IGPL* **15**, pp. 1-20.

da Costa, N. and Rodrigues, A. (2007), 'Definability and Invariance', *Studia Logica* **86**, pp. 1-30.

Dalla Chiara, M. L., Giuntini, R. and Krause, D. (1998), 'Quasiset Theories for Microobjects: A Comparison', in E. Castellani (ed.), *Intepreting Bodies: Classical and Quantum Objects in Modern Physics,* Princeton: Princeton University Press, pp. 142-52.

Davidson, D. (1970), Mental Events, in L. Foster and J. W. Swanson (eds.), *Experience and Theory*, London: Duckworth, pp. 79-101.

Demopoulos, W. (2007) 'Carnap on the Rational Reconstruction of Scientific Theories', in Friedman and Creath (eds.), *The Cambridge Companion to Carnap*, Cambridge: Cambridge University Press, pp. 248-272.

Dennett, D. (1981), *Brainstorms: Philosophical Essays on Mind and Psychology.* Cambridge: MIT Press.

Dieks, D. and Lubberdink, A. (2011), 'How Classical Particles Emerge from the Quantum World', *Foundations of Physics* **41**, pp. 1051-1064.

Dummett, M. (1991), *The Logical Basis of Metaphysics.* Cambridge: Harvard University Press.

Earman (ms.), 'Understanding permutation invariance in quantum mechanics'. Unpublished manuscript.

Einstein, A., Podolsky B. and Rosen, N (1935), 'Can Quantum-Mechanical Description of Physical Reality be Considered Complete?', *Physical Review* **47**, pp. 777780.

Esfeld, M. (2004), 'Quantum Entanglement and a Metaphysics of Relations', *Studies in History and Philosophy of Modern Physics* **35**, pp. 601-17.

Esfeld, M. and Lam, V. (2006), 'Modest Structural Realism', *Synthese* **160**, pp. 27-46.

Field, H. (1980), *Science Without Numbers: A Defence of Nominalism.* Oxford: Blackwell.

Fleming, G. and Butterfield, J. (1999), 'Strange positions', in J. Butterfield and C. Pagonis (eds.), *From Physics to Philosophy*, New York: Cambridge University Press, pp. 108-165.

Føllesdal, D. (1968), 'Quine on Modality', *Synthese* **19**, pp. 147-57.

Fraser, D. (2008), 'The fate of 'particle' in quantum field theories with interactions', *Studies in the History and Philosophy of Modern Physics*, **39**, pp. 841-859.

French, S. (2006), 'Identity and Individuality in Quantum Theory', in E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2007 Edition), <plato.stanford.edu/entries/qt-idind>.

French, S. and Krause, D. (2006), *Identity in Physics: A Historical, Philosophical and Formal Analysis.* Oxford: Oxford University Press.

French, S. and Ladyman, J. (2003), 'Remodelling Structural Realism: Quantum

Physics and the Metaphysics of Structure', *Synthese*, **136** (1), pp. 31-56.

French, S. and Redhead, M. (1988), 'Quantum physics and the identity of indiscernibles', *British Journal for the Philosophy of Science* **39**, pp. 233-46.

Geach, P. T. (1980), *Reference and Generality.* 3rd edition. Ithaca: Cornell University Press.

Ghirardi, G. and Marinatto, L. (2003), 'Entanglement and Properties', *Fortschritte der Physik* **51**, pp. 379387.

Ghirardi, G. and Marinatto, L. (2004), 'General Criterion for the Entanglement of Two Indistinguishable Particles', *Physical Review A* **70**, 012109-1-10.

Ghirardi, G. and Marinatto, L. (2005), 'Identical Particles and Entanglement', *Optics and Spectroscopy* **99**, pp. 386-390.

Ghirardi, G., Marinatto, L. and Weber, T. (2002), 'Entanglement and Properties of Composite Quantum Systems: A Conceptual and Mathematical Analysis', *Journal of Statistical Physics* **108**, pp. 49-122.

Gisin, N. (1991), 'Bell's Inequality Holds for all Non-Product States', *Physics Letters A* **154**, pp. 201-202.

Hacking, I. (1975), 'The Identity of Indiscernibles', *The Journal of Philosophy* **72**, pp. 249-56.

Hacking, I. (1999), *The Social Construction of What?* Cambridge: Harvard University Press.

Halvorson, H. (2001), 'Reeh-Schlieder defeats Newton-Wigner: On alternative localization schemes in relativistic quantum field theory', *Philosophy of Science* **68**, pp. 111-133.

Halvorson, H. and Clifton, R. (2002), 'No Place for Particles in Relativistic Quantum Theories?', emphPhilosophy of Science **69**, pp. 1-28.

Haslanger, S. (1989), 'Endurance and Temporary Intrinsics', *Analysis*, **49**, pp. 119-25.

Haslanger, S. (2006), 'How Good Are Our Intuitions? Philosophical Analysis and Social Kinds', *Aristotelian Society Supplementary Volume* **80**, pp. 89-118.

Hawley, K. (2001), *How Things Persist.* Oxford: Oxford University Press.

Hawley, K. (2009), 'Identity and Indiscernibility', *Mind* **118**, pp. 101-119.

Healey, R. (1978), 'Physicalist Imperialism', *Proceedings of the Aristotelian Society*, **79**, pp. 191-211.

Hilbert, D. and Bernays, P. (1934), *Grundlagen der Mathematik*, Volume 1. Berlin: Springer.

Hintikka, M. J. and Hintikka, J. (1983), 'How can Language be Sexist?', in S. Harding and M. Hintikka (eds), *Discovering Reality: Feminist Perspectives on Epistemology, Metaphysics, Methodology, and Philosophy of Science*, Dordrecht: Kluwer, pp. 139-48.

Horgan, T. (1993), 'From Supervenience to Superdupervenience: Meeting the Demands of a Material World', *Mind* **102**, pp. 556-586.

Horsten, L. (2007), 'Philosophy of Mathematics', *The Stanford Encyclopedia of Philosophy (Fall 2008 Edition)*, E. N. Zalta(ed.),
URL = <http://plato.stanford.edu/archives/fall2008/entries/philosophy-mathematics/>.

Huggett, N. (1999), 'On the significance of permutation symmetry', *British Journal for the Philosophy of Science* **50**, pp. 325-47.

Huggett, N. (2003), 'Quarticles and the Identity of Indiscernibles', in K. Brading and E. Castellani (eds.), *Symmetries in Physics: New Reflections*, Cambridge: Cambridge University Press, pp. 239-249.

Huggett, N. and Imbo, T. (2009), 'Indistinguishability', in D. Greenberger, K. Hentschel and F. Weinert (eds.), *Compendium of Quantum Physics*, Berlin: Springer, pp. 311-317.

Huggett, N. and Weingard, R. (1994a), 'Interpretations of Quantum Field Theory', *Philosophy of Science* **61**, pp. 370-388.

Huggett, N. and Weingard, R. (1994b), 'On the field aspect of Quantum fields ', *Erkenntnis* **40**, pp. 293-301.

Hume, D. (1740/1975), *A Treatise of Human Nature,* edited by L. A. Selby-Bigge, Oxford: Clarendon Press.

Jardine, N. (1986), *The Fortunes of Inquiry.* Oxford: Clarendon Press.

Jordan, T. F. (1969), *Linear Operators for Quantum Mechanics.* Duluth: Thomas F. Jordan.

Kaplan, D. (1966), 'Transworld Heir Lines', in M. J. Loux (ed.), *The Possible and the Actual: Readings in the Metaphysics of Modality*, New York: Cornell University Press, pp. 88-109.

Kaplan, D. (1975), 'How to Russell a Frege-Church', *The Journal of Philosophy* **72**, pp. 716-729.

Ketland, J. (2006), 'Structuralism and the Identity of Indiscernibles', *Analysis* **66**, pp. 303-15.

Ketland, J. (2011), 'Identity and Indiscernibility', *Review of Symbolic Logic* **4**, pp. 171-185.

Knobe, J. and Nichols, S. (eds.) (2008). *Experimental Philosophy*, Oxford: Oxford University Press.

Kochen, S. and Specker, E. P. (1967), 'The Problem of Hidden Variables in Quantum Mechanics', *Journal of Mathematics and Mechanics* **17**, pp. 59-87.

Krause, D. (1992), 'On a Quasi-set Theory', *Notre Dame Journal of Formal Logic* **33**, pp. 402-11.

Krause, D. and French, S. (2007), 'Quantum Sortal Predicates', *Synthese* **143**, pp. 417-430.

Kusch M. (1995), *Psychologism: A Case Study in the Sociology of Philosophical Knowledge.* London: Routledge.

Ladyman (1998), 'What is structural realism?' *Studies in History and Philosophy of Science* **29**, pp. 409-424.

Ladyman, J. (2002), 'Science, metaphysics and structural realism.' *Philosophica* **67**, pp. 57-76.

Ladyman, J. (2005), 'Mathematical Structuralism and the Identity of Indiscernibles', *Analysis* **65**, pp. 218-21.

Ladyman, J. (2007a), 'On the Identity and Diversity of Objects in a Structure', *Proceedings of the Aristotelian Society* **81**, pp. 23-43.

Ladyman, J. (2007b), 'Structural Realism', *The Stanford Encyclopedia of Philosophy (Summer 2009 Edition)*, Edward N. Zalta(ed.), URL = <http://plato.stanford.edu/archives/sum2009/entries/structural-realism/>.

Ladyman, J. and Ross, D. (2007), *Every Thing Must Go: Metaphysics Naturalized.* Oxford: Oxford University Press.

Landsman, N. P. (2007), 'Between Classical and Quantum', in J. Earman and J. Butterfield (eds.), *Handbook in the Philosophy of Science*, Vol 2: *Philosophy of Physics Part A*, Amsterdam: North Holland, pp. 417-554.

Leitgeb, H. and Ladyman, J. (2008), Criteria of Identity and Structuralist Ontology, *Philosophia Mathematica* **16**, pp. 388-96.

Leonard, H. S. and Goodman, N. (1940), 'The Calculus of Individuals and Its Uses', *The Journal of Symbolic Logic* **5**, pp. 45-55.

Lewis, D. (1966), 'An Argument for the Identity Theory', *Journal of Philosophy* **63**, pp. 17-25.

Lewis, D. (1968), 'Counterpart Theory and Quantified Modal Logic', *Journal of Philosophy* **65**, pp. 113-126.

Lewis, D. (1970a), 'General Semantics', *Synthese* **22**, pp. 18-67.

Lewis, D. (1970b), 'How to Define Theoretical Terms', *The Journal of Philosophy* **67**, pp. 427-446.

Lewis, D. (1972), 'Psychophysical and Theoretical Identifications', *The Australasian Journal of Philosophy* **50**, pp. 249-258.

Lewis, D. (1973), 'Causation', *Journal of Philosophy* **70**, pp. 556-567.

Lewis, D. (1983), 'New Work for a Theory of Universals', *The Australasian Journal of Philosophy* **61** pp. 343-77.

Lewis, D. (1984), 'Putnam's Paradox', *The Australasian Journal of Philosophy* **62**, pp. 221-236.

Lewis, D. (1986), *On the Plurality of Worlds.* Oxford: Blackwell.

Lewis, D. (1988), 'Statements Partly About Observation', *Philosophical Papers* **17**, pp. 1-31.

Lewis, D. (1992), 'Armstrong on Combinatorial Possibility', *The Australasian Journal of Philosophy* **70**, pp. 211-224.

Lewis, D. (1993), 'Many but almost one', in K. Campbell, J. Bacon and L. Reinhardt (eds.), *Ontology, Causality, and Mind: Essays on the Philosophy of D. M. Armstrong*, Cambridge: Cambridge University Press.

Lewis, D. (1995), 'Should a materialist believe in qualia?', *The Australasian Journal of Philosophy* **73**, pp. 140-144.

Lewis, D. (2002), 'Tensing the Copula', *Mind* **111**, pp. 1-13.

MacBride, F. (1999), 'Could Armstrong Have Been a Universal?', *Mind* **108**, pp. 471-501.

MacBride, F. (2005), 'Structuralism Reconsidered', in S. Shapiro (ed.), *The Oxford Handbook of Philosophy of Logic & Mathematics*, Oxford: Oxford University Press, pp. 563-89.

Maggiore, M. (2005), *A Modern Introduction to Quantum Field Theory.* Oxford: Oxford University Press.

Malament, D. (1996), 'In defense of dogma: Why there cannot be a relativistic quantum mechanics of (localizable) particles', in R. Clifton (ed.), *Perspectives on Quantum Reality*, Dordrecht: Kluwer, pp. 1-10.

Manton, N. and Sutcliffe, P. (2004), *Topological Solitons.* Cambridge: Cambridge University Press.

Margenau, H. (1944), 'The Exclusion Principle and its Philosophical Importance', *Philosophy of Science* **11**, pp. 187-208.

Massimi, M. (2001), 'Exclusion Principle and the Identity of Indiscernibles: a Response to Margenau's Argument', *The British Journal for the Philosophy of Science* **52**, pp. 303-330.

Mattuck, R. D. (1992), *A Guide to Feynman Diagrams in the Many-Body Problem.* New York: Dover.

Mellor, D. H. (1977), 'Conscious Belief', *Proceedings of the Aristotelian Society* **78**, pp. 87-101.

Merricks, T. (1999), 'Persistence, Parts and Presentism', *Noûs* **33**, p. 421-438.

Messiah, A. and Greenberg, O. W. (1964), 'The symmetrization postulate and its experimental foundation', *Physical Review* **136**, pp. 248-67.

Misner, C. W., Thorne, K. S. and Wheeler, J. A. (1973), *Gravitation.* San Francisco: Freeman & Co.

Morison, B. (2002), *On Location: Aristotle's Concept of Place.* Oxford: Oxford University Press.

Muller, F. A. (forthcoming), 'How to Discern Space-time Points Structurally', forthcoming in *Philosophy of Science.*

Muller, F. A. and Saunders, S. (2008), 'Discerning Fermions', *British Journal for the Philosophy of Science*, **59**, pp. 499-548.

Muller, F. A. and Seevinck, M. (2009), 'Discerning Elementary Particles', *Philosophy of Science*, **76**, pp. 179-200.

Murdoch, I. (1970), *The Sovereignty of Good.* New York: Schocken Books.

Oliver, A. (1996), 'The Metaphysics of Properties', *Mind*, **105**, pp. 1-80.

Parfit, D. (1971), 'Personal Identity', *Philosophical Review*, **80**, pp. 3-27.

Poincaré, H. (1905/1952), *Science and Hypothesis.* New York: Dover.

Pooley, O. (2006), 'Points, Particles, and Structural Realism', in S. French, D. Rickles and J. Saatsi (eds.), *The Structural Foundations of Quantum Gravity*, Oxford: Oxford University Press, pp. 83-120.

Pooley, O. (ms.), *Substantivalism and Haecceitism.*

Psillos, S. (2006), 'The Structure, the Whole Structure and Nothing But the Structure?' *Philosophy of Science* **73**, pp. 560-570.

Putnam, H. (1977), 'Realism and Reason', *Proceedings of the American Philosophical Association* **50**, pp. 483-498.

Putnam, H. (1980), 'Models and Reality', *Journal of Symbolic Logic* **45**, pp. 464-482.

Quine, W. V. O. (1960), *Word and Object.* Cambridge: Harvard University Press.

Quine, W. V. O. (1969), *Ontological Relativity and Other Essays.* New York: Columbia University Press.

Quine, W. V. O. (1970), *Philosophy of Logic.* Cambridge: Harvard University Press.

Quine, W. V. O. (1976a), 'Grades of Discriminability', *The Journal of Philosophy* **73**, pp. 113-6.

Quine, W. V. O. (1976b), 'Worlds Away', *Journal of Philosophy* **73**, pp. 859-863.

Quine, W. V. O. (1977), 'Facts of the Matter', in R. W. Shahan and K. R. Merrill (eds.), *American Philosophy: From Edwards to Quine*, Norman: University of Oklahoma Press.

Redhead, M. (1987), 'A Philosopher Looks at Quantum Field Theory', in H. Brown and R. Harré (eds.), *Philosophical Foundations of Quantum Field Theory*, Oxford: Oxford University Press, pp. 9-24.

Reichenbach, H. (1958), *The Philosophy of Space and Time.* New York: Dover.

Robinson, D. (2000), 'Identities, Distinctness, Truthmakers, and Indiscernibility Principles', *Logique & Analyse* **43**, pp. 145-83.

Russell, B. (1905), 'On Denoting', *Mind* **14**, pp. 479-493.

Saunders, S. (2003a), 'Physics and Leibnizs Principles', in K. Brading and E. Castellani (eds), *Symmetries in Physics: Philosophical Reflections,* Cambridge: Cambridge University Press.

Saunders, S. (2003b), 'Indiscernibles, covariance and other symmetries: the case for non-reductive relationism', in A. Ashtkar, D. Howard, J. Renn, S. Sarkar and A. Shimony (eds.), *Revisiting the Foundations of Relativistic Physics: Festschrift in Honour of John Stachel*, Amsterdam: Kluwer, pp. 289-307.

Saunders, S. (2006), 'Are Quantum Particles Objects?', *Analysis* **66**, pp. 52-63.

Sider, T. (2001), *Four-Dimensionalism: An Ontology of Persistence and Time.* Oxford: Clarendon Press.

Scerri, E. R. and Worrall, J. (2001), 'Prediction and the Periodic Table', *Studies in History and Philosophy of Science* **32**, pp. 407-452.

Schrödinger, E. (1926), 'Der stetige Übergang von der Mikro-zur Makromekanik. *Die*

Naturwissenschaften **14**, pp. 664-668.

Sneed, J. D. (1971), *The Logical Structure of Mathematical Physics.* Dordrecht: D. Reidel.

Strawson, P. F. (1950), 'On Referring', *Mind* **59**, pp. 320-344.

Taylor, B. (1993), 'On Natural Properties in Metaphysics', *Mind* **102**, pp. 81-100.

Teller, P. (1986), 'Relational Holism and Quantum Mechanics', *The British Journal for the Philosophy of Science* **37**, pp. 71-81.

Teller, P. (1995), *An Interpretive Introduction to Quantum Field Theory.* Princeton: Princeton University Press.

Tung, W.-K. (1985), *Group Theory in Physics.* River Edge: World Scientific.

Unger, P. (1980) 'The Problem of the Many', *Midwest Studies in Philosophy* **5**, pp. 411?467.

van Dalen, D. (1994), *Logic and Structure*, 3rd edition. Berlin: Springer.

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

Wallace, D. (2003), 'Everett and structure', *Studies in the History and Philosophy of Modern Physics* **34**, pp. 87 105. Available online at arXiv.org/abs/quant-ph/0107144 or from philsci-archive.pitt.edu.

Wallace, D. (2010), 'Decoherence and Ontology (or: How I Learned to Stop Worrying and Love FAPP)', in in Saunders, Barrett, Kent, and Wallace (ed.), *Many Worlds? Everett, Quantum Theory, and Reality* Oxford: Oxford University Press, pp. 53-72.

Wallace, D. and Timpson, C. (2010), 'Quantum Mechanics on Spacetime I: Spacetime State Realism', *British Journal for the Philosophy of Science* **61**, pp. 697-727.

Weyl, H. (1928/1931), *The Theory of Groups and Quantum Mechanics.* New York: Dover.

Weyl, H. (1949), *Philosophy of Mathematics and Natural Science.* Princeton: Princeton University Press.

Whitehead, A. N. (1925/1997), *Science and the Modern World.* Free Press (Simon & Schuster).

Wiggins, D. (1980), *Sameness and Substance.* Cambridge: Harvard University Press.

Wiggins, D. (2001), *Sameness and Substance Renewed.* Cambridge: Cambridge University Press.

Wigner, E. P. (1939), 'On Unitary Representations of the Inhomogeneous Lorentz Group', *Annals of Mathematics* **40**, pp. 149-204.

Wilson, J. (2008), 'What is Hume's Dictum, and Why Believe It?', *Philosophy and Phenomenological Research* **80**, pp. 595-637.

Wittgenstein, L. (1922), *Tractatus Logico-Philosophicus.* London: Kegan Paul.

Wittgenstein, L. (1953). *Philosophical Investigations.* Oxford: Blackwell Publishing.

Worrall, J. (1989), 'Structural Realism: The Best of Both Worlds?', *Dialectica* **43** 1-2, pp. 99-124.

Worrall, J. and Zahar, E. (2001), 'Ramsification and structural realism.' Appendix in E. Zahar, *Poincar's Philosophy: From Conventionalism to Phenomenology*, Court Publishing Co., pp. 236-251.

Yi, B.-U. (1999), 'Is Mereology Ontologically Innocent?', *Philosophical Studies* **93**, pp. 141-160.

Zanardi, P. (2001), 'Virtual Quantum Subsystems', *Physical Review Letters* **87**, 077901-1-4.

Zanardi, P., Lidar, D. A. and Lloyd, S. (2004), 'Quantum Tensor Product Structures are Observable-Induced', *Physical Review Letters* **92**, 060402-1-4.

Zimmerman, D. (1998), 'Temporal Parts and Supervenient Causation: the Incompatibility of Two Humean Doctrines', *Australasian Journal of Philosophy* **76**, pp. 265-288.

Żukowski, M., Brukner, Č., Laskowski, W. and Wieśniak, M. (2002), 'Gisin's Theorem Cannot Be Generalized to All Multi-Particle States', *Physics Review Letters* **88**, 210402-1-5.