Permutations, redux [DRAFT]

Adam Caulton (adam.caulton@philosophy.ox.ac.uk)

July 21, 2023

Abstract

The purpose of this article is to give a general overview of permutations in physics, particularly the symmetry of theories under permutations. Particular attention is paid to classical mechanics, classical statistical mechanics and quantum mechanics. There are two recurring themes: (i) the metaphysical dispute between haecceitism and anti-haecceitism, and the extent to which this dispute may be settled empirically; and relatedly, (ii) the way in which elementary systems are individuated in a theory's formalism, either primitively or in terms of the properties and relations those systems are represented as bearing.

Section 1 introduces permutations and provides a brief outline of the symmetric and braid groups. Section 2 discusses permutations in the general setting provided by model theory, in particular providing some definitions and elementary results regarding the permutability and indiscernibility of objects. Section 3 lays some philosophical groundwork for later sections, in particular articulated the distinction between haecceitism and anti-haecceitism and the distinction between transcendental and qualitative individuation. Section 4 addresses classical mechanics and introduces the procedure of quotienting, under which permutable states are identified. Section 5 addresses classical statistical mechanics, and outlines a number of equivalent ways to implement permutation invariance. I also briefly outline how particles may be qualitatively individuated in this framework. Section 6 addresses quantum mechanics. This contains an outline of: the representation theory of the symmetric groups; the topological approach to quantum statistics, in which the braid groups become relevant; and a brief proposal for qualitatively individuating quantum particles, and its implications for entanglement. Section 7 concludes with a discussion of equilibrium ensembles in the classical and quantum theories under permutation invariance.

A (much) shorter version of this paper was published as a chapter in E. Knox & A. Wilson (eds), *the Routledge Companion to Philosophy of Physics* (Routledge, 2021), pp. 578–594.

Contents

1 Introduction

1.1 What are permutations? Elements of the symmetric groups

Given any set X, consider the set of bijections $\pi : X \to X$. This set can be endowed with a natural group structure: the identity e maps each object to itself and the binary group operation $\circ$ is just functional composition: $(\pi_2 \circ \pi_1)(a) := \pi_2(\pi_1(a))$. (In future, we will represent composition of group elements by simple concatenation; e.g. $\pi_2\pi_1$ instead of $\pi_2 \circ \pi_1$.) The resulting group is called the *symmetric group on* X, denoted S_X or $\mathrm{Sym}(X)$. The order of S_X is $|X|!$, where $|X|$ is the cardinality of X.

In the case where $X = \{1, 2, \ldots, N\}$ for some positive natural number N, the corresponding group, often called *the symmetric group on N symbols*, is denoted S_N.

Any subgroup of $\mathrm{Sym}(X)$ is called a *permutation group*. In a certain sense, the study of permutation groups encompasses *all* groups; this is due to Cayley's Theorem:

Theorem (Cayley, 1854)**.** *Any group G is isomorphic to some subgroup of Sym(G).*

The proof relies on considering the left action of G on itself: each element $g \in G$ is associated with the left action $\Phi(g, \cdot)$, which is a bijection $\Phi(g, \cdot) : G \to G$; this association preserves the group structure of G, i.e. $\Phi(h, \Phi(g, \cdot)) = \Phi(h \circ g, \cdot)$. However, in this article I will focus on the finite symmetric groups S_N, especially as they are realised or represented as groups of permutations on physical systems, or on formal labels for clusters of degrees of freedom.

Let $X = \{a_1, \ldots, a_N\}$ be any non-empty finite set; then the permutation π may be represented as a $2 \times N$ array, as follows:

$$\begin{pmatrix} a_1 & a_2 & \cdots & a_N \\ \pi(a_1) & \pi(a_2) & \cdots & \pi(a_N) \end{pmatrix}$$

In this notation, the order of the columns is redundant, since any re-ordering establishes the same bijection on X. Permutations in S_N may also be given a graphical representation. For example:

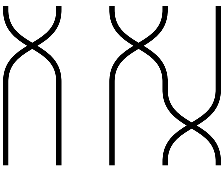

corresponds to $\begin{pmatrix} 1 & 2 & 3 & 4 & 5 \\ 2 & 1 & 5 & 3 & 4 \end{pmatrix}$

Any permutation π of the form $\pi(a_1) = a_2$; $\pi(a_2) = a_3$; $\ldots$; $\pi(a_{n-1}) = a_n$; $\pi(a_n) = a_1$ (and otherwise act as the identity) is called an *n-cycle*, and may be succinctly denoted by $(a_1 a_2 \cdots a_{n-1} a_n)$ or $(a_2 \cdots a_{n-1} a_n a_1)$, etc. Any permutation at all may be uniquely decomposed into a series of disjoint (i.e. commuting) cycles: for example, the permutation

$$\begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 \\ 5 & 4 & 3 & 7 & 1 & 6 & 2 \end{pmatrix}$$

is the permutation $(15)(247)(3)(6) = (15)(247)$. Note that each number appears here at most once, and that we may omit the 1-cycles, corresponding to *fixed points*, so long as we know the value of N. (The only exception is $e = (1)(2)\cdots(N)$, which we will continue to denote by e.)

The fact that permutations are decomposable into disjoint cycles permits a particularly succinct representation of any permutation, known as *cycle notation*. Cycle notation can be made unique by demanding that: (i) shorter cycles are written before longer cycles; (ii) if two cycles have the same length, then the cycle containing the lowest number appears first; and (iii) each cycle begins with its lowest number. So e.g. (57)(624)(31) must be written as (13)(57)(246). The unique decomposition into disjoint cycles entails that each permutation has a well defined *cycle structure* or *cycle type*, which is its number of 1-cycles, 2-cycles, 3-cycles, etc. So e.g. (13)(57)(246) in S_9 is composed of two 1-cycles (namely, (8) and (9), not explicitly written), two 2-cycles and one 3-cycle; we may represent its cycle type as $(1, 1, 2, 2, 3)$, or $1^2 2^2 3^1$.

Given any group G, two elements $g, g' \in G$ are called *conjugate* iff there is some $h \in G$ such that $g' = h \circ g \circ h^{-1}$. Conjugacy is an equivalence relation, so any group G may be partitioned into conjugacy classes $[g] := \{h \circ g \circ h^{-1} : h \in G\}$. In the case of S_N, any element's cycle type is shared by all and only its conjugates. For example, the permutation (15)(247) is conjugate to

$$[(37)(1426)](15)(247)[(37)(1426)]^{-1} = (37)(1426)(15)(247)(37)(1624) = (45)(236)$$

So cycle type is characteristic of each conjugacy class. In the example just given, the cycle type is $1^2 2^1 3^1$ (assuming we are in S_7), and so we can deduce that the corresponding conjugacy class has order $7!/(2! \cdot 2 \cdot 3) = 420$.[1] To take an example which will recur in this article: S_3 has three cycle types: 1^3, corresponding to the conjugacy class $\{e\}$; $1^1 2^1$, corresponding to conjugacy class $\{(12), (13), (23)\}$; and 3^1, corresponding to conjugacy class $\{(123), (132)\}$. All this will be important in Section 6.1, where I briefly outline the representation theory of S_N.

Finally, in anticipation of the braid groups, any element of S_N can be expressed (not necessarily uniquely) as a sequence of adjacent pairwise swaps $\sigma_i := (i, i + 1)$. For example, $(123) = (12)(23) = \sigma_1 \sigma_2$, as can be seen as follows (remember to read from right to left and top to bottom):

$$\begin{pmatrix} 1 & 2 & 3 \\ 2 & 1 & 3 \end{pmatrix} \begin{pmatrix} 1 & 2 & 3 \\ 1 & 3 & 2 \end{pmatrix} = \begin{pmatrix} 1 & 3 & 2 \\ 2 & 3 & 1 \end{pmatrix} \begin{pmatrix} 1 & 2 & 3 \\ 1 & 3 & 2 \end{pmatrix} = \begin{pmatrix} 1 & 2 & 3 \\ 2 & 3 & 1 \end{pmatrix}$$

In graphical notation, this is immediate:

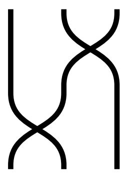

In this way, we can think of the $(N - 1)$ adjacent swaps σ_i as generating the full group S_N, where the σ_i are subject to the following conditions:

1. $\sigma_i \sigma_j = \sigma_j \sigma_i$, for all $i, j = 1, \ldots, N - 1$ such that $|i - j| \geqslant 2$;
2. $\sigma_i \sigma_{i+1} \sigma_i = \sigma_{i+1} \sigma_i \sigma_{i+1}$, for all $i = 1, \ldots, N - 2$;
3. $(\sigma_i)^2 = e$, for all $i = 1, \ldots, N - 1$.

Or, graphically:

1. 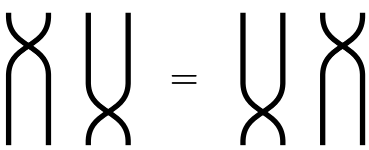 2. 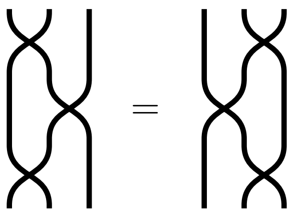 3. 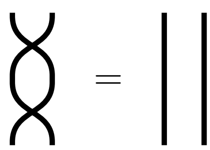

[1]In general for S_N, the cycle type $1^{c_1} 2^{c_2} \cdots N^{c_N}$, where $\sum_{n=1}^{N} n c_n = N$, has order $\frac{N!}{\prod_{n=1}^{N} c_n! n^{c_n}}$.

Even though the decomposition of a permutation into adjacent swaps is typically not unique, it is well defined whether the number of adjacent swaps is even or odd. Correspondingly, permutations are classified into even or odd. The group of even permutations in S_N, the *alternating group* A_N, is a normal subgroup of S_N. The corresponding quotient group is $S_N/A_N \cong \mathbb{Z}_2$, for any N. For more on the symmetric and permutation groups, consult Sagan (1991) and Dixon & Mortimer (1996).

1.2 The braid groups

The third condition on the generators of S_N above is equivalent to the condition that adjacent swaps are self-inverse: $(\sigma_i)^{-1} = \sigma_i$. This condition may be relaxed. We can represent this graphically by giving each swap an orientation, indicated by which strand is on top of which at crossings, as follows:

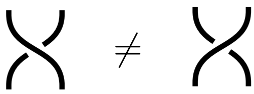

The resulting elements σ_i are called *(adjacent) braids*. The group generated by forming arbitrary sequences, known as *braid words*, of adjacent braids $\sigma_1, \ldots, \sigma_{N-1}$ (and their inverses) is called the *braid group on N symbols* and is denoted B_N. The conditions on the adjacent braids may be represented graphically as follows:

1. 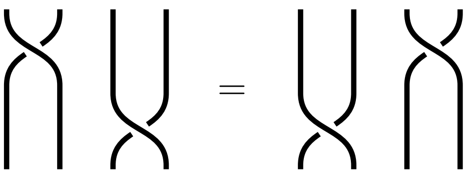

2. 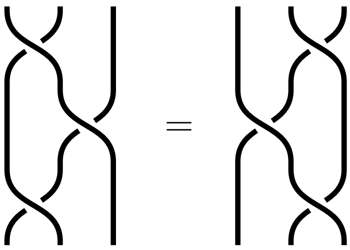

Relaxing the self-inverse condition leads to a (much) larger group. For example, S_2, generated by $\{\sigma_1\}$, has only two elements, e and $\sigma_1 = (12)$ and is isomorphic to $\mathbb{Z}_2$, the simplest non-trivial cyclic group; while B_2, generated by $\{\sigma_1\}$, has denumerably many elements (all distinct):

$$\ldots,\ (\sigma_1)^{-2},\ (\sigma_1)^{-1},\ e,\ \sigma_1,\ (\sigma_1)^2,\ \ldots$$

and is isomorphic to the additive integers $\mathbb{Z}$. The braid group will be important in Section 6.3, where I will briefly outline the quantization of reduced configuration spaces. For more on the braid groups, consult Kassel & Turaev (2008).

2 Permutations in logic and model theory

Models, in roughly the sense of first-order structures (or, more realistically, Bourbaki structures), offer a very general and systematic means of representing possibilities—that is, possible states or worlds. Equally (as we shall see in later Sections), we may take a single model's domain to comprise some naturally demarcated gamut of possibilities, as is done in formal definitions of state spaces in classical or quantum mechanics. It will therefore be helpful to briefly review the treatment of permutations in this general context. (For comprehensive treatments of model theory, see Hodges 1993 and Button & Walsh 2018.)

For the sake of simplicity, I will concentrate on first-order structures of the form $\mathfrak{A} = \langle A, \mathcal{R} \rangle$ and call them *models*. The *domain* of $\mathfrak{A}$, A, is any set. The *structure* of $\mathfrak{A}$, $\mathcal{R} = \langle R_1, \ldots, R_m \rangle$, is a sequence of relations on A; so, if R_i is an n-ary relation, then $R_i \subseteq A^n$. In the context of the semantics for a first-order language, each R_i is the extension assigned by $\mathfrak{A}$ to an n-ary predicate symbol.

Following Quine and again for the sake of simplicity, I restrict attention to models whose structures comprise only relations; distinguished elements (extensions assigned to constants) and functions (extensions assigns to functors) may be replaced without loss by relations in the usual way (see Quine 1986, pp. 25-6).

2.1 Permutations and permutability

Given any model $\mathfrak{A} = \langle A, \mathcal{R} \rangle$ with domain A and relations $\mathcal{R} = \langle R_1, \ldots, R_m \rangle$, any permutation $\pi : A \to A$ induces a lift π^* on models with the same domain A (Button and Walsh 2018, §??? call this the *Push-Through construction*):

For each n-ary $R \in \mathcal{R}$ and all $a_1, \ldots, a_n \in A$: $\langle a_1, \ldots, a_n \rangle \in \pi^* R$ iff $\langle \pi(a_1), \ldots, \pi(a_n) \rangle \in R$.

Then we may define $\pi^*\mathcal{R} := \langle \pi^*(R_1), \ldots, \pi^*(R_m) \rangle$ and $\pi^*\mathfrak{A} := \langle A, \pi^*\mathcal{R} \rangle$. π constitutes an *isomorphism* between $\mathfrak{A}$ and $\pi^*\mathfrak{A}$.

Any permutation $\pi : A \to A$ is a *symmetry*, a.k.a. an *automorphism*, of $\mathfrak{A}$ iff $\pi^*\mathfrak{A} = \mathfrak{A}$, i.e. it *fixes*, or leaves invariant, the structure of $\mathfrak{A}$. The symmetries of $\mathfrak{A}$ form a subgroup $\mathrm{Aut}(\mathfrak{A})$ of the group $\mathrm{Sym}(A)$ of all permutations on $\mathfrak{A}$'s domain. Any model $\mathfrak{A}$ is *rigid* iff its only symmetry is the identity on A (i.e. $\mathrm{Aut}(\mathfrak{A}) = \{e\}$). I will call a model *symmetric* iff it is not rigid, i.e. has a non-trivial symmetry, and *totally symmetric* iff every permutation on its domain is a symmetry (i.e. $\mathrm{Aut}(\mathfrak{A}) = \mathrm{Sym}(A)$). For example, the complex field $\mathbb{C}$ is symmetric, since $\pi(z) = z^*$ (each complex number is mapped to its complex conjugate) is a non-trivial symmetry, but not totally symmetric, since e.g. $\pi(z) = 1 - z$ is not a symmetry.[2]

Using the symmetries of $\mathfrak{A}$ we may define the relation $\sim_{\mathfrak{A}}$ of *permutability in* $\mathfrak{A}$ on $\mathfrak{A}$'s domain as follows: for all $a, b \in A$, $a \sim_{\mathfrak{A}} b$ iff there is some permutation $\pi : A \to A$ such that π is a symmetry of $\mathfrak{A}$ and $\pi(a) = b$.[3] For example, the imaginary numbers i and $-i$ are permutable in $\mathbb{C}$. The relation of permutability is the natural semantic counterpart of the syntactic notion of absolute indiscernibility—roughly, discernibility by monadic formulae that do not contain individual constants or equality. [Quine (1976), Saunders (2003a, 2003b, 2013), Ketland (2011), Caulton & Butterfield (2012), Ladyman *et al* (2012), Muller (2015).]

Permutability in $\mathfrak{A}$ is an equivalence relation, and so we can define permutability equivalence classes $[a]_{\mathfrak{A}} := \{b \in A : a \sim_{\mathfrak{A}} b\}$. Each equivalence class is the *orbit*, under the symmetries of $\mathfrak{A}$, of any one of its elements. I shall call any object $a \in A$ an *individual in* $\mathfrak{A}$ iff it is *fixed* by every symmetry of $\mathfrak{A}$, i.e. it is permutable only with itself (so $[a]_{\mathfrak{A}} = \{a\}$). For example, each rational number is an individual in the field $\mathbb{C}$.

The notion of individuality just introduced is closely connected to the notion of definability. An element $a \in A$ is *definable* in $\mathfrak{A}$ iff a uniquely satisfies some (first-order) monadic formula in the language that $\mathfrak{A}$ interprets. If a is definable in $\mathfrak{A}$, then a is an individual in $\mathfrak{A}$. In all finite models and some infinite models, definability and individuality coincide; but typically the

[2]The *field* $\mathbb{C}$ in fact has $2^{\mathfrak{c}}$ symmetries, where $\mathfrak{c}$ is the cardinality of the continuum. However, only two of these are *continuous* symmetries; these coincide with the symmetries of the *vector space* $\mathbb{C}$. See Yale (1966).

[3]It may seem that the definition is strangely asymmetric, but if $\pi : A \to A$ is a symmetry such that $\pi(a) = b$, then its inverse $\pi^{-1} : A \to A$ is a symmetry such that $\pi^{-1}(b) = a$.

expressive resources of the language will fall short of making every individual definable.

2.2 Congruence, indiscernibility and quotient models

In abstract algebra any relation is a *congruence relation* iff it is an equivalence relation and, in addition, is *compatible* with the associated algebraic structure. This notion of compatibility is normally defined for a structure with just one binary operation,[4] but it can be extended to our general models (see Ketland 2011, Ladyman, Linnebo & Pettigrew 2012, §8). Given any model $\mathfrak{A} = \langle A, \mathcal{R} \rangle$, any equivalence relation $\sim$ on the domain A is *compatible with the structure* $\mathcal{R}$ iff, for each n-ary $R \in \mathcal{R}$ and all $a_1, \ldots a_n, b_1, \ldots, b_n \in A$: if $a_i \sim b_i$ for each $i = 1, \ldots, n$, then $\langle a_1, \ldots, a_n \rangle \in R$ iff $\langle b_1, \ldots, b_n \rangle \in R$.

Given any model $\mathfrak{A} = \langle A, \mathcal{R} \rangle$, we can always find a congruence relation: I shall call it *indiscernibility in* $\mathfrak{A}$. For any $a, b \in A$, a and b are indiscernible in $\mathfrak{A}$, written $a \approx_{\mathfrak{A}} b$, iff a and b share all the same properties and bear all the same relations to the same things in $\mathfrak{A}$'s domain. More precisely, $a \approx_{\mathfrak{A}} b$ iff:

- for all 1-ary $P \in \mathcal{R}$: $a \in P$ iff $b \in P$;
- for all 2-ary $R \in \mathcal{R}$ and all $c \in A$: $\langle a, c \rangle \in R$ iff $\langle b, c \rangle \in R$, and $\langle c, a \rangle \in R$ iff $\langle c, b \rangle \in R$;
- for all 3-ary $R \in \mathcal{R}$ and all $c, c' \in A$: $\langle a, c, c' \rangle \in R$ iff $\langle b, c, c' \rangle \in R$, $\langle c, a, c' \rangle \in R$ iff $\langle c, b, c' \rangle \in R$, and $\langle c, c', a \rangle \in R$ iff $\langle c, c', b \rangle \in R$;
- and so on, for all relations in $\mathcal{R}$.

This relation is an equivalence relation, and by construction it is compatible with $\mathcal{R}$; therefore it is a congruence relation on $\mathcal{R}$. It is the natural semantic counterpart to the indiscernibility relation associated with Hilbert and Bernays (1934, §5), Quine (1970, pp. 61-4) and Saunders (2003a, p. 5). In fact, if the structure $\mathcal{R}$ contains only finitely many relations, then $\approx_{\mathfrak{A}}$ is defined in any model $\mathfrak{A}$ by the first-order sentence suggested by Quine. Indiscernibility equivalence classes are defined as expected: $[\![a]\!]_{\mathfrak{A}} := \{b \in A : a \approx_{\mathfrak{A}} b\}$.

Permutability is confused with indiscernibility at your peril! (The dialogue between Black's (1952) interlocutors A and B provides a classic cautionary tale.) In any model $\mathfrak{A}$, indiscernibility ($a \approx_{\mathfrak{A}} b$) implies permutability ($a \sim_{\mathfrak{A}} b$), and so indiscernibility classes are subsets of permutability classes ($[\![a]\!]_{\mathfrak{A}} \subseteq [a]_{\mathfrak{A}}$). However, there are models in which the converse fails: these are precisely the models in which permutability fails to be a congruence relation. This is the semantic counterpart of the celebrated fact that absolute indiscernibility (indiscernibility by monadic formulae) is necessary but typically not sufficient for utter indiscernibility (indiscernibility by arbitrary formulae) The gap between permutability and indiscernibility is opened up by the possibility of discerning permutable objects by n-ary relations, where $n \geqslant 2$. Quine appears to have been the first to notice this; Saunders (2003b) first applied the insight to physics. It has since been applied in a wide variety of treatments.

Even for models in which permutability does not imply indiscernibility, still indiscernibility has its own link to permutations. I shall call any two objects a, b *freely permutable in* $\mathfrak{A}$ iff any permutation π on $\mathfrak{A}$'s domain that swaps them (i.e. $\pi(a) = b$; $\pi(b) = a$) is a symmetry of $\mathfrak{A}$. If a and b are indiscernible in $\mathfrak{A}$, then the permutation $\pi = (ab)$, whose *only* action is to swap a and b (i.e. $\pi(a) = b$; $\pi(b) = a$; $\pi(c) = c$ for all $c \neq a, b$), is a symmetry of $\mathfrak{A}$. It

[4] The equivalence relation $\sim$ is compatible with the binary operation $\circ$, both defined on the domain A, iff: for all $a, a', b, b' \in A$, if $a \sim b$ and $a' \sim b'$, then $(a \circ a') \sim (b \circ b')$.

follows that any two elements of $[\![a]\!]_{\mathfrak{A}}$ are freely permutable in $\mathfrak{A}$. The corresponding claims for permutables are not generally true, since a and b may be permutable only by dint of symmetries that *additionally* swap objects other than a and b to which a and b are related. For example, i and $-i$ are permutable but not freely permutable in the field $\mathbb{C}$, since the simple transposition of i and $-i$ is not a symmetry. Furthermore, freely permutable objects may yet be discernible, since they may bear (symmetric) relations to each other which serve to discern them.

Given any structure $\mathfrak{A} = \langle A, \mathcal{R}\rangle$ where $\mathcal{R} = \langle R_1, \ldots, R_n\rangle$, we can define its *quotient model* $\mathfrak{A}/{\approx_{\mathfrak{A}}} = \langle A/{\approx_{\mathfrak{A}}}, \mathcal{R}^*\rangle$ where $\mathcal{R}^* = \langle R_1^*, \ldots, R_n^*\rangle$ such that:

- $A/{\approx_{\mathfrak{A}}} := \{[\![a]\!]_{\mathfrak{A}} : a \in A\}$;
- for all $R_i \in \mathcal{R}$ and all $a_1, \ldots, a_n \in A$: $\langle [\![a_1]\!]_{\mathfrak{A}}, \ldots, [\![a_n]\!]_{\mathfrak{A}}\rangle \in R_i^*$ iff $\langle a_1, \ldots, a_n\rangle \in R_i$.

The relations R_i^* are well defined precisely because indiscernibility in $\mathfrak{A}$ is a congruence relation. Despite quotienting, note that $\mathfrak{A}/{\approx_{\mathfrak{A}}}$ may still have non-trivial symmetries.

Any model $\mathfrak{A}$ is elementarily equivalent to (i.e. satisfies the same first-order sentences as) its quotient model $\mathfrak{A}/{\approx_{\mathfrak{A}}}$, so long as equality is not a logical primitive of the language. It is not generally true that a model is isomorphic to its quotient model: in fact $\mathfrak{A} \cong \mathfrak{A}/{\approx_{\mathfrak{A}}}$ iff any two distinct objects in $\mathfrak{A}$'s domain are discernible in $\mathfrak{A}$. But we may have indiscernible objects $a, b \in A$ that are nevertheless numerically distinct. This "invisibility" of indiscernible objects from first-order sentences not containing equality reflects the fact, proved by Hilbert & Bernays (1934, §5), that identity is not in general definable.

So, to summarise, we have the following chain of (generally, one-way) implications. (For each term, read 'a and b are ... in $\mathfrak{A}$'.)

$$\text{identical} \quad \Rightarrow \quad \text{indiscernible} \quad \Rightarrow \quad \text{freely permutable} \quad \Rightarrow \quad \text{permutable}$$

The last three relations coincide if (but not generally only if) $\mathfrak{A}$'s structure comprises only monadic properties (1-ary relations). All four coincide if (but not generally only if) every object in $\mathfrak{A}$'s domain uniquely satisfies some monadic property (our Quinean surrogate for every object bearing a name), for in that case $\mathfrak{A}$ is rigid. Indiscernibility is the semantic counterpart of "utter indiscernibility" and permutability is the semantic counterpart of "absolute indiscernibility", both as found in Caulton & Butterfield (2012), Ladyman *et al* (2012) and Muller (2015). Muller (2015) calls discernible permutables "relationals". Free permutability offers a *via media* between indiscernibility and permutability and may provide an explication, in semantic terms, of Ladyman & Bigaj's (2010) "witness-indiscernibility'. (For more on witness-discernibility, see Linnebo & Muller 2013 and Bigaj 2015.)

3 Related metaphysical and interpretative disputes

3.1 The identity of indiscernibles

Given any model $\mathfrak{A}$ (or better: the possibility represented by that model) we may ask whether any two distinct objects in its domain are discernible. Any such model satisfies the logically weakest non-trivial formulation of the *Principle of the Identity of Indiscernibles* (PII). A logically stronger formulation is provided by permutability: it says that any two distinct objects fail to be permutable.

Hacking (1975) provides a recipe for making this strong version of PII true no matter what: given any model $\mathfrak{A}$ which violates the Principle, we can construct a new model $\mathfrak{A}^*$, taken to be a mere notational variant of $\mathfrak{A}$ in which the Principle is upheld. The recipe, in a nutshell, is this: given $\mathfrak{A}$, let $\mathfrak{A}^*$ be the quotient model of $\mathfrak{A}$ under the equivalence relation of permutability in $\mathfrak{A}$.

The trouble with this recipe in general is that $\sim_{\mathfrak{A}}$ may fail to be a congruence relation. For example, let us consider the celebrated example of Black's (1952) two iron spheres, intrinsically identical and lying two miles apart in a relationalist space. This may be represented by the model $\mathfrak{B} = \langle B, \{R\}\rangle$, where

- $B := \{\text{Castor}, \text{Pollux}\}$;
- $R = \{\langle x, y\rangle : x \text{ is 2 miles away from } y\} = \{\langle\text{Castor}, \text{Pollux}\rangle, \langle\text{Pollux}, \text{Castor}\rangle\}$.

We have Castor $\sim_{\mathfrak{B}}$ Pollux, but Castor $\not\approx_{\mathfrak{B}}$ Pollux (Castor and Pollux are *weakly discernible*, as was pointed out by Saunders 2003a). The attempt to define a quotient structure using $\sim_{\mathfrak{B}}$ fails without further specification because $\sim_{\mathfrak{B}}$ is not a congruence relation, and so it is indeterminate whether or not we should have (for example) $\langle[\text{Castor}]_{\mathfrak{B}}, [\text{Castor}]_{\mathfrak{B}}\rangle \in R^*$. However, where permutability fails, indiscernibility succeeds: $\approx_{\mathfrak{A}}$ is guaranteed to be a congruence relation, and so given any model $\mathfrak{A}$ we may always pass to its quotient model $\mathfrak{A}/\!\approx_{\mathfrak{A}}$. In the example above, $\mathfrak{B}/\!\approx_{\mathfrak{B}}$ is isomorphic to $\mathfrak{B}$, and so has the same group of symmetries.

Alternatively, we could define the quotient model $\mathfrak{A}/\text{Aut}(\mathfrak{A}) := \langle A/\!\sim_{\mathfrak{A}}, \tilde{\mathcal{R}}\rangle$ by laying down as a general rule that, for each n-ary $R \in \mathcal{R}$ and all $a_1, \ldots, a_n \in A$: $\langle[a_1]_{\mathfrak{A}}, \ldots, [a_n]_{\mathfrak{A}}\rangle \in \tilde{R}$ iff $\langle\pi(a_1), \ldots, \pi(a_n)\rangle \in R$ for all symmetries π of $\mathfrak{A}$. This is tantamount to judicously adding to $\mathfrak{A}$'s structure until indiscernibility in $\mathfrak{A}$ coincides with permutability in $\mathfrak{A}$, and then passing to the quotient under indiscernibility $\mathfrak{A}/\!\approx_{\mathfrak{A}}$. For any model $\mathfrak{A}$, $\mathfrak{A}/\text{Aut}(\mathfrak{A})$ is rigid. In the example above, $\mathfrak{B}/\text{Aut}(\mathfrak{B})$ has only one object in its domain (since $[\text{Castor}]_{\mathfrak{B}} = [\text{Pollux}]_{\mathfrak{B}}$), which is 2 miles away from itself: i.e., $\tilde{R} = \{\langle[\text{Castor}]_{\mathfrak{B}}, [\text{Castor}]_{\mathfrak{B}}\rangle\}$.

As the example above shows, in cases where it fails to be a congruence relation, quotienting under permutability can do significant violence to the original model's structure. The procedure would be catastrophic in pure mathematics, where there is an abundance of symmetric structures. If we quotient the field $\mathbb{C}$ under its symmetries, then (among many other disasters) addition fails to be a function: $[i]_{\mathbb{C}}$ sums with itself to *both* $[2i]_{\mathbb{C}}$ and $[0]_{\mathbb{C}}$. If we quotient the symmetric group S_3 under its symmetries (S_3 happens to be its own symmetry group, and the permutability classes are the conjugacy classes), then group composition fails to be a function: $[(12)]_{S_3}$ composes with itself to produce *both* $[e]_{S_3}$ and $[(123)]_{S_3}$.

These considerations give us good reason to deny Hacking's claim that $\mathfrak{A}$ and $\mathfrak{A}/\text{Aut}(\mathfrak{A})$ can always be taken as notational variants of one another. Examples from graph theory offer good reason to deny even the corresponding claim regarding $\mathfrak{A}$ and $\mathfrak{A}/\!\approx_{\mathfrak{A}}$ (Ladyman 2007). It therefore seems sensible to conclude that PII, in either form, is not generally true of mathematical objects. However, one can still ask of any model whether or not it happens to obey PII, in either its strong ($\mathfrak{A} \cong \mathfrak{A}/\text{Aut}(\mathfrak{A})$) or weak ($\mathfrak{A} \cong \mathfrak{A}/\!\approx_{\mathfrak{A}}$) form. Furthermore—and more important for our interests here—one can ask, in any given application of some model $\mathfrak{A}$, whether $\mathfrak{A}$, $\mathfrak{A}/\!\approx_{\mathfrak{A}}$ or $\mathfrak{A}/\text{Aut}(\mathfrak{A})$ provides the most perspicuous representation of our target system. This is the topic of the next section.

3.2 Haecceitism vs. anti-haecceitism

The choice between $\mathfrak{A}$ and $\mathfrak{A}/\mathrm{Aut}(\mathfrak{A})$ as a representation of a physical target system may be posed at a higher level of abstraction: at which $\mathfrak{A}$'s domain comprises possible worlds or states. This brings us to one way of articulating the much-discussed disagreement between haecceitism and anti-haecceitism.

So let's take the model $\mathfrak{W} = \langle W, \mathcal{R}\rangle$ not as a representative of a single possibility (a possible world or state), but instead as a representative of a *space* of possibilities. The elements of $\mathfrak{W}$'s domain W are then not objects (in the everyday sense), but rather the possibilities themselves—or, at least, their mathematical representatives; call them *states*. $\mathfrak{W}$'s structure $\mathcal{R}$ then represents relations on these states: as it may be, a topology and differentiable structure, a symplectic form, vector space structure, etc., and an algebra of quantities defined on the states.

Suppose that the states themselves sufficiently resemble models, all with a common domain A and a common signature.[5] As we saw in Section 2.1, any permutation $\pi : A \to A$ induces a *lifted permutation* $\pi^* : W \to W$ on $\mathfrak{W}$'s domain. (It must be emphasised: the lifted permutations typically comprise a highly restricted subgroup of *all* permutations on W.) Even if our states do not resemble models—perhaps because they are simple points lying in a phase space, or rays lying in a Hilbert space—we may still be able to make sense of the idea that any permutation of the states' objects induces a lifted permutation on W. In classical mechanics, sense is provided by group actions, a.k.a. *group realisations* of permutation on the joint configuration space or joint phase space; in quantum mechanics, sense is provided by *group representations* of permutations on the joint Hilbert space. I explore these in some detail in Sections 4.2 and 6.1, respectively.

Take some lifted permutation $\pi^* : W \to W$. Is it a symmetry of $\mathfrak{W}$? π^* induces a lifted permutation of its own, $\pi^\star$ say, on models with the same domain W of states; so our question may be rephrased, Is $\pi^\star\mathfrak{W} = \mathfrak{W}$? Trivially, it will be if π^* acts as the identity on W, so suppose otherwise. If π^* is not a symmetry of $\mathfrak{W}$, then $\mathfrak{W}$ can tell the difference between at least one state $w \in W$ and $\pi^* w$.[6] But w and $\pi^* w$ differ only by a permutation—that is, only according to which object is which in w's structure of properties and relations. (If they happen to be representable as models, then they will be isomorphic.) So if some lifted permutation π^* is not a symmetry, then $\mathfrak{W}$ distinguishes between permuted states: it cares which object is which, deep down in the states.

For the purposes of articulating haecceitism and anti-haecceitism, I will arrange things so that $\mathfrak{W}$ *doesn't care* which object is which. Given the above, it follows that every lifted permutation π^* on states must be a symmetry of $\mathfrak{W}$, i.e. $\pi^\star\mathfrak{W} = \mathfrak{W}$. I propose that we take this as necessary and sufficient for $\mathfrak{W}$'s structure containing information only about the *qualitative character* of states.[7] Examples of qualitative character include *how many* objects (never mind *which*) bear this or that combination of properties, or take these or those values for such-and-such quantities—essentially, occupation numbers. It also includes details about the structure of the network of relations that each state attributes to its elements.

[5] Two models $\mathfrak{A} = \langle A, \langle R_1, \ldots, R_m\rangle\rangle$ and $\mathfrak{B} = \langle B, \langle R'_1, \ldots, R'_n\rangle\rangle$ have a common signature iff $m = n$ and each pair of corresponding relations R_i and R'_i have the same arity.

[6] Note that I am now using the lower case Roman letter 'w' as a variable to range over states, which may well themselves be models. This is to emphasise that we are primarily concerned here with the state space.

[7] This definition of 'qualitative character' has the counter-intuitive consequence that information about the numerical distinctness of states' objects, or about how many objects there are of various kinds, counts as part of a state's qualitative character. Yet 'qualitative' is often glossed as 'non-identity-involving' (Ladyman *et al* 2012, pp. 163, 169). The discrepancy is tolerable here, since our notion of qualitative character permits an adequate definition of haecceitism. 'Structural character' might be a preferable term—if only 'structure' weren't such a ubiquitous term!

So: for any state $w \in W$ and any lifted permutation $\pi^* : W \to W$, w and $\pi^* w$ have the same qualitative character, and will therefore be permutable in $\mathfrak{W}$. I will leave it open whether states *not* related by a lifted permutation have *different* qualitative characters. This is to allow for cases such as those which arise in quantum mechanics and classical statistical mechanics, where the structure of the state space makes permutability-in-$\mathfrak{W}$ classes closed under certain ways of combining states to produce new states. In these cases, states may have the same qualitative character, and so are permutable in $\mathfrak{W}$, even though they are not related by a *lifted* permutation.

I now have sufficient grounding to propose the following definitions:

(*Haecceitism*) Any two distinct states permutable in $\mathfrak{W}$ represent distinct possibilities.

(*Anti-haecceitism*) Any two distinct states permutable in $\mathfrak{W}$ represent the same possibility.

Anti-haecceitism is therefore a particular statement of permutation-invariance: the possibility being represented is invariant under any symmetry of $\mathfrak{W}$, which is (typically but perhaps not always) generated by a permutation of the states' elements. It follows that, for the anti-haecceitist, $\mathfrak{W}/\mathrm{Aut}(\mathfrak{W})$ provides a more perspicuous representation of the possibilities than $\mathfrak{W}$.

Haecceitism expresses a willingness to deny this permutation-invariance. $\mathfrak{W}$ must then offer a more perspicuous representation of the possibilities than $\mathfrak{W}/\mathrm{Aut}(\mathfrak{W})$. But presumably the haecceitist will hope to do better than $\mathfrak{W}$. For $\mathfrak{W}$ attributes only qualitative characters to the states (in classical and quantum mechanics, this is what can be conveyed by the sub-algebra of permutation-invariant quantities). The haecceitist, of course, takes the possible facts to surpass mere qualitative character, so we might expect them to enrich $\mathfrak{W}$'s structure (but not its domain) with the means suitable to express which-is-which facts. In classical and quantum mechanics, this is what can be conveyed by the "full" algebra of quantities, not restricted by permutation-invariance.

These definitions are very much in the spirit of Lewis (1986, p. 221), according to whom haecceitism is the denial, and anti-haecceitism the affirmation, of the claim that the what is true of any given object supervenes on the qualitative character of the global state or world. If our haecceitism is true, then distinct possibilities may have the same qualitative character. Assuming that all facts divide into which-is-which facts or facts about qualitative character, those possibilities must differ as to which object is which in the network of properties and relations. Therefore, they differ as to what is true of some object. So we may have a change in what is true of some object without a corresponding change in qualitative character, which is a failure of the former to supervene on the latter. If our anti-haecceitism is true, then any two distinct possibilities have distinct qualitative characters. So what is true of some object cannot vary without qualitative character also varying. So the former supervenes on the latter.

Are haecceitism and anti-haecceitism contraries? Not quite. First, it's clear that they agree on the totally symmetric states, since each of those states is the sole occupant of its permutability class, and so $[w]_{\mathfrak{W}} = \{w\}$ is as good as w as a representative of any possibility.[8] If *every* state in $\mathfrak{W}$'s domain is totally symmetric, then *no* two distinct states are permutable in $\mathfrak{W}$, $\mathfrak{W}$ and $\mathfrak{W}/\mathrm{Aut}(\mathfrak{W})$ are therefore isomorphic, and both haecceitism and anti-haecceitism are trivially satisfied. Moreover, the haecceitist and anti-haecceitist cannot in this case even disagree about whether $\mathfrak{W}$'s structure suffices for individuating any given state: if no two distinct states are permutable, then every one of $\mathfrak{W}$'s states is an individual in $\mathfrak{W}$, and so has a unique qualitative character.

[8]If the state w is totally symmetric, then no object in w's domain is an individual in w, and w itself is an individual in $\mathfrak{W}$. In other words: given some qualitative character, there is at most one way to be totally symmetric.

This point may seem recherché, but this apparently bizarre situation arises frequently in physics. In one way of framing classical statistical mechanics (see Section 5.3) and in the quantum mechanics of bosons and fermions (see Section 6.1), all states are totally symmetric. In these cases, it is hard to see how the dispute between haecceitism and anti-haecceitism could be adjudicated (a point frequently emphasised by French, e.g. 1989). However, a related dispute can still be formulated in these cases; that dispute is the topic of the next Section.

I'll conclude this section with a simple example to illustrate the concepts introduced. Let W be a set of simple models ("states"), all of the form $\langle A, \{P\}\rangle$, where $A = \{a_1, \ldots, a_N\}$ (So we have a fixed finite domain), and equipped with just one monadic property $P \subseteq A$. There are 2^N such states, corresponding to the 2^N choices for the extension of P. According to haecceitism, these correspond to a total of 2^N possibilities. Let $\mathfrak{W} = \langle W, \{\#\}\rangle$, where $\# : \wp(A) \to \{0, 1, \ldots N\}$ returns the cardinality of each state's property P. The value of $\#(P)$ plausibly exhausts what can be said about the qualitative character of each state $\langle A, \{P\}\rangle$. Any permutation $\pi : A \to A$ lifts to a permutation $\pi^* : W \to W$ which is a symmetry of $\mathfrak{W}$, and $\mathfrak{W}$ has no further symmetries. Any two states are isomorphic iff they agree on $\#(P)$, which is true iff they are related by a lifted permutation. So the isomorphism classes of states coincide with the permutability classes in $\mathfrak{W}$, which are the elements in the domain $W/{\sim_{\mathfrak{W}}}$ of the quotient structure $\mathfrak{W}/\mathrm{Aut}(\mathfrak{W})$, which we might here also denote by $\mathfrak{W}/\mathrm{Sym}(A)$ or even $\mathfrak{W}/S_N$. There are $N{+}1$ such classes, corresponding to the possible values in the range of # (i.e. the possible cardinalities, or "occupation numbers", of the property P). According to anti-haecceitism, these permutability classes correspond to a total of $N + 1$ possibilities. The permutability class $[\langle A, \{a_1, \ldots, a_m\}\rangle]_{\mathfrak{W}}$, an 'anti-haecceitistic state' corresponding to $\#(P) = m$, comprises $\frac{N!}{m!(N-m)!}$ distinct 'haecceitistic states'.

3.3 Transcendental vs. qualitative individuality

The dispute between haecceitism and anti-haecceitism is one about permutation-invariance. But what is being permuted, and what do the permutations represent? It is the states of $\mathfrak{W}$ that are being permuted, but the central focus is on *lifted* permutations, which are the lifts of permutations defined on the *objects* in the states' common domain. These objects are, like the states of $\mathfrak{W}$, mere representatives intended for some physical application. So what do the *objects* represent? Here I articulate what I take to be two salient proposals, which are mutually exclusive but not jointly exhaustive: *transcendental individuality* (TI) and *qualitative individuality* (QI).

(*TI*) Each element in the states' common domain denotes some object of the target system, and that element denotes the same object in all states.

(*QI*) The elements in the states' common domain denote nothing (in particular) in the target system.

The term 'transcendental individuality' was coined by Post (1963) and also appears in Redhead and Teller (1991, 1992); I mean it in roughly their sense. The guiding idea is that facts about object identity from possibility to possibility transcend the qualitative character of any possibility. This is expressed, using $\mathfrak{W}$, by taking advantage of the fact that the states have a common domain, and stipulating that the same object in that domain always stands for the same object in the target system.

TI entails haecceitism. If two distinct states are permutable in $\mathfrak{W}$, then they represent distinct possibilities by dint of representing distinct possibilities for the objects; they differ as to which object lies where in the mosaic of qualitative relational structure. Haecceitism does not entail TI, simply because it is not committed to any account of how it is that two distinct, permutable states represent distinct possibilities. However, it may well be that TI provides the only plausible motivation for haecceitism (I know of no other motivation).

The term ‘qualitative individuality’ is less commonly used in the literature, because the corresponding view is often overlooked. The guiding idea is that the objects in the states’ common domain are nothing but placeholders. They may serve as representative “hooks” on which to hang properties and relations, but any hooks will do (hence they represent nothing ‘in particular’). Specifically, we should afford no significance to the fact that the same hook appears from state to state.

The claim that these “hooks” represent *nothing* in the target system therefore stands in need of some qualification. Any hook, in a sense, represents “the” generic object, in its capacity to have properties predicated of it (in physicists’ jargon: any hook represents a ‘cluster of degrees of freedom’). But the identification of objects in the target system from possibility to possibility is to proceed, on this view, not according to the identity of hooks from state to state, but rather according to the qualitative properties and relations hung on the hooks.

We also have to be liberal in our understanding of what counts as the hooks. As we shall see in the following Sections, in classical and quantum mechanics the state spaces of joint systems are often product spaces: in classical mechanics, the Cartesian product of phase spaces; in quantum mechanics, the tensor product of Hilbert spaces. The hooks in these cases are not objects in any model’s domain, but rather the *order* of the factor spaces in the joint state space, or perhaps the label used to emphasise that ordering.

QI entails anti-haecceitism. If the objects in the states’ common domain represent nothing (beyond the general capacity to bear properties and relations), then any physically significant fact must be invariant under any permutation of those objects. In particular, any two distinct, permutable states must represent the same possibility. In physicists’ jargon, this is a familiar statement of gauge invariance: if a mathematical variable bears no physical interpretation, then all physically significant facts are represented in the formalism by gauge-invariant claims, i.e. claims which do not vary under arbitrary transformations of the corresponding mathematical variable.

Anti-haecceitism entails QI if there are distinct permutable states—at least for the objects being permuted. For, in that case states which differ only on “which-is-which facts” represent the same possibility. It follows that which-is-which facts serve no representative function, and so the associated permuted objects cannot denote anything in the target system. To put it another way: the grouping together of states into permutability classes breaks the trans-state identity relations otherwise indicated by the same object appearing from state to state.

However, if no two states are permutable—i.e., if every state in $\mathfrak{W}$’s domain is an individual in $\mathfrak{W}$ (because totally symmetric)—, then *anti-haecceitism is compatible with TI.* In this case, both haecceitism and anti-haecceitism are trivially true (remember, $\mathfrak{W}$ and $\mathfrak{W}/\mathrm{Aut}(\mathfrak{W})$ are in this case isomorphic), and one may continue to take the identity of objects from state to state as representing the identity of physical objects from possibility to possibility. This harks back to the apparently recherché observation made in the previous Section.

As we shall see in Sections 5.3 and 6.4, far from being recherché, an unholy alliance between TI and anti-haecceitism underpins counter-intuitive claims that have made in both the classical statistical mechanics and quantum mechanics—in particular, the claim that particles are indiscernible by means of monadic properties in all states. However, as will be made clear in the following Sections, these claims may coherently be denied, and an alternative is provided by qualitative individuation.

The disagreement here articulated between TI and QI mirrors an analogous dispute in the metaphysics literature, between so-called "transworld identity" and "world-bound individuals". So why the different terminology? Chiefly, I want to avoid getting into disputes about how metaphysical possibilities are best represented (by mathematical models or states, as here; or by maximal consistent sets of sentences; or by concrete worlds, as real as our own?), which tend to surround the familiar metaphysical dispute. I also want to avoid the unfortunate misapprehension that a commitment to qualitative individuation entails a commitment to the claim that all objects bear all their properties and relations essentially.[9]

4 Permutations in classical mechanics

4.1 The realisation of permutations on joint configuration spaces

We begin with $\mathcal{Q}$, the configuration space for a generic elementary system. (Note: the system being represented need not *be* elementary; it's just being treated as such.) For a particle in d-dimensional Euclidean space, $\mathcal{Q} = \mathbb{R}^d$, understood as having (at least) the structure of a smooth manifold.

Representing an assembly of such systems, say N of them, standardly proceeds as follows. First we form the joint configuration space $\mathcal{Q}^N$, often written

$$\mathcal{Q}^N := \underbrace{\mathcal{Q} \times \ldots \times \mathcal{Q}}_{N} \tag{1}$$

to indicate that the joint configuration space's points are elements in the N-fold Cartesian product of $\mathcal{Q}$'s domain and that $\mathcal{Q}^N$ has been endowed with the obvious product topology (see Willard 1970, §3.8) and differential structure.

There are broadly two routes from here. The first is to define the the joint phase space, and then consider permutations of system labels. The second is to consider permutations of system labels *first*, and then proceed to the joint phase space. I will take the first route. The joint phase space is the cotangent bundle $T^*(\mathcal{Q}^N)$, equipped with the symplectic form $\Omega = \sum_{i=1}^N \omega^{(i)}$, where each $\omega^{(i)}$ is a copy of the symplectic form ω associated with a copy of the single-system phase space $\Gamma := T^*\mathcal{Q}$. The result is equivalent (as a phase space) to the N-fold tensor product of Γ, so I will denote this joint phase space by Γ^N.

States in the joint phase space Γ^N may be denoted by $(\xi_1, \ldots, \xi_N)$, where each $\xi_i \in \Gamma$ is a single-system state. We can now define a natural *realisation*, a.k.a. *group action*, of the group S_N of permutations on N symbols, as follows. For each permutation $\pi \in S_N$, define the *lifted permutation* $\pi^* : \Gamma^N \to \Gamma^N$ such that

$$\pi^*(\xi_1, \ldots, \xi_N) := (\xi_{\pi(1)}, \ldots, \xi_{\pi(N)}) \tag{2}$$

The lift $^* : \pi \mapsto \pi^*$ is a realisation of S_N precisely because it preserves the group structure: $(\pi_1\pi_2)^* = \pi_1^*\pi_2^*$ and e^* is the identity on Γ^N. The realisation is *faithful*, i.e. distinct permutations in S_N are sent by the lift to distinct maps on Γ^N. Moreover, any lifted permutations is a symmetry of Γ^N, since any lifted permutation preserves its manifold structure and symplectic form Ω.

[9] Perhaps Kaplan (1975, p. 723) should have had the final word: 'Although the Anti-Haecceitist may seem to assert that no possible individual exists in more than one possible world, that view is properly reserved for the Haecceitist who holds to an unusually rigid brand of metaphysical determinism.' Kaplan's haecceitism and anti-haecceitism are more in line with our TI and QI respectively, being committed as they are to specifics as to *why it is* that permutable states don't or do represent the same possibility.

4.2 Classical permutation invariance

We can add more structure to Γ^N by introducing an algebra $\mathcal{A}$ of quantities. It is standard to posit the Poisson algebra of all smooth functions $f : \Gamma^N \to \mathbb{R}$, equipped, via the symplectic form Ω, with an associated Poisson bracket $\{\cdot,\cdot\} : \mathcal{A} \otimes \mathcal{A} \to \mathcal{A}$. The resulting model $\mathfrak{P} = \langle \Gamma^N, \mathcal{A} \rangle$ is the standard arena for classical N-particle mechanics.

This model is rigid, since any two distinct states differ on the values of some quantities—most obviously, the positions and momenta of the individual systems. If we wish to capture only the qualitative character of each state (in the sense of Section 3.2), then we must restrict the algebra $\mathcal{A}$ to the permutation-invariant quantities; i.e. the smooth functions $f : \Gamma^N \to \mathbb{R}$ such that

$$f(\pi^*(\xi_1, \ldots, \xi_N)) = f(\xi_1, \ldots, \xi_N) \tag{3}$$

Note that, e.g., the position of system 5 or the momentum of system 17 will not be among these quantities; but various quantities jointly able, given a specification of values, to uniquely characterise a "cloud" of N points in the single-system phase space Γ *will* be among them. Call this restricted permutation-invariant algebra $\mathcal{A}_{PI}$.

While $\langle \Gamma^N, \mathcal{A} \rangle$ is rigid, $\langle \Gamma^N, \mathcal{A}_{PI} \rangle$ has symmetries: these are precisely the lifts π^* of the permutations in S_N. We may therefore define the quotient model $\langle \Gamma^N/{\sim_{\mathfrak{P}_{PI}}}, \tilde{\mathcal{A}}_{PI} \rangle$, where $\tilde{\mathcal{A}}_{PI}$ is defined in the obvious way;[10] the phase space $\Gamma^N/{\sim_{\mathfrak{P}_{PI}}}$ is more commonly denoted by Γ^N/S_N. (See Willard 1970, §3.9 for details of defining quotient spaces.) However, Γ^N has points on which the action of S_N is not free; these are the states $(\xi_1, \ldots, \xi_N)$ such that $\xi_i = \xi_j$ for some $i \neq j$. These points form a boundary on Γ^N/S_N, and so Γ^N/S_N is not a manifold (it is an *orbifold*). This generates a host of technical issues, not least of which is the fact that tangent spaces on the boundary have the "wrong" dimension, and smooth vector fields on the bulk—such as Hamiltonian flows—cannot be defined on them.

It is standard practice to avoid this outcome by removing *collision configurations* from the joint configuration space $\mathcal{Q}^N$ *before* defining the joint phase space and its quotient under S_N. The collision configurations comprise the set $\Delta := \{(x_1, \ldots, x_N) \in \mathcal{Q}^N : x_i = x_j \text{ for some } i \neq j\}$, which has vanishing Lebesgue measure. So let $\Gamma_N := T^*(\mathcal{Q}^N \setminus \Delta)$ be our new joint phase space. The group of permutations S_N acts *freely* on this space (every state $\Xi \in \Gamma_N$ is such that $\pi_1^*\Xi \neq \pi_2^*\Xi$ for any two distinct permutations $\pi_1, \pi_2 \in S_N$), and the resulting quotient phase space Γ_N/S_N is a manifold, isomorphic as a phase space to $T^*((\mathcal{Q}^N \setminus \Delta)/S_N)$.

Note that the justification for removing the collision configurations was technical: it was to ensure a quotient phase space with nice properties. The physical justification—if there is one—is murkier: it is certainly suspect that classical systems should become discernible in all states (by their positions, or relative positions) on grounds of a technicality. Certainly, if we expect a dynamics which makes the particles impenetrable (presumably by means of some strong short-range repulsive force), then the excision is justified—but what about alternative dynamics? It must be emphasised that the quotient procedure is far from innocent: it typically leaves the resulting joint configuration space, and its quotient under S_N, topologically non-trivial. (As we shall see in Section 6.3, this is in fact *essential* in deriving the full gamut of quantum particle statistics upon quantisation.)

---

[10] We set up a bijection, in fact an algebra isomorphism, $\tilde{}\, : \mathcal{A}_{PI} \to \tilde{\mathcal{A}}_{PI}$ such that, for any $f \in \mathcal{A}_{PI}$, $\tilde{f} \in \tilde{\mathcal{A}}_{PI}$ is the unique quantity such that $f(\Xi) = \tilde{f}([\Xi])$, for all $\Xi \in \Gamma^N$. This construction mirrors the definition of $\mathfrak{A}/\mathrm{Aut}(\mathfrak{A})$, outlined in Section **??**.

4.3 Individuation in permutation-invariant classical mechanics

We appear to have two very different classical N-particle theories. The first is associated with the model $\langle \Gamma_N, \mathcal{A} \rangle$, and is suited to a proponent of transcendental individuation, and therefore haecceitism (I assume here that the particles do not possess state-independent properties, such as distinct masses, which serve to distinguish them). The second theory is associated with the model $\langle \Gamma_N/S_N, \tilde{\mathcal{A}}_{PI} \rangle$, and is suited to the proponent of qualitative individuation, and therefore anti-haecceitism. Certainly, there are distinct states in Γ^N that become permutable if we restrict to the sub-algebra of permutation-invariant quantities; so haecceitism and anti-haecceitism seem to be genuine rivals here.

However, if the Hamiltonian of the system is among the permutation-invariant quantities, then the two theories are in fact empirically equivalent, up to arbitrary stipulations. This can be seen as follows.[11] Take an arbitrary trajectory $\Xi : \mathbb{R} \to \Gamma_N$ in the haecceitistic theory. At each time $t \in \mathbb{R}$, the state $\Xi(t)$ lies in the permutability class $[\Xi(t)] \in \Gamma_N/S_N$ of the anti-haecceitistic theory. So the trajectory Ξ may be associated with a unique trajectory $[\Xi] : \mathbb{R} \to \Gamma_N/S_N$ in the anti-haecceitistic theory, where $[\Xi](t) := [\Xi(t)]$ for all $t \in \mathbb{R}$. At any time, the trajectories Ξ and $[\Xi]$ yield the same values for all qualitative (that is, permutation-invariant) quantities, where we make use of the uniquely natural association $\mathcal{A}_{PI} \leftrightarrow \tilde{\mathcal{A}}_{PI}$ between the algebras of the two theories. In particular, Ξ is solution for the Hamiltonian H iff $[\Xi]$ is a solution for the Hamiltonian $\tilde{H}$ uniquely corresponding to H, since H is permutation-invariant. Yet the permutation-invariant quantities exhaust what is experimentally determinable, precisely because the system labels constitute *transcendental* individuation criteria.

Going in the opposite direction is more involved. Take an arbitrary trajectory $\tilde{\Xi} : \mathbb{R} \to \Gamma_N/S_N$ in the anti-haecceitistic theory. At each time $t \in \mathbb{R}$, the state $\tilde{\Xi}(t)$ is a permutability class $\{\pi^*\Xi_t \in \Gamma_N : \pi \in S_N\}$ containing states of the haecceitistic theory. Now the concern arises that haecceitistic trajectories cannot be recovered: since the formation of permutability classes breaks trans-state identity relations established by system labels, there would appear to be an embarrassment of options for cross-identifying haecceitistic states at different times. (Given $\tilde{\Xi}(t)$ and $\tilde{\Xi}(t+\epsilon)$, which haecceitistic state $\pi_1^*\Xi_t \in \tilde{\Xi}(t)$ should be identified with which later state $\pi_2^*\Xi_{t+\epsilon} \in \tilde{\Xi}(t+\epsilon)$?) However, if the anti-haecceitistic trajectory $\tilde{\Xi}$ is continuous in Γ_N/S_N, then we can demand that the haecceitistic trajectories be likewise continuous in Γ_N. There will be a unique series of cross-identifications which fulfils this demand. Crucial to this uniqueness is the fact that collision points in the joint configuration space have been removed; *with* the collision points, the demand for continuity will fail to yield unique haecceitistic cross-temporal identifications if at any time the anti-haecceitistic trajectories pass through them.

Proceeding without collision points, the continuous anti-haecceitistic trajectory $\tilde{\Xi}$ defines a class of $N!$ continuous haecceitistic trajectories. And $\tilde{\Xi}$ is a solution for the Hamiltonian $\tilde{H}$ iff all of the corresponding haecceitistic trajectories are solutions for the Hamiltonian H uniquely corresponding to $\tilde{H}$ (H is unique because it is permutation-invariant). Which one of the $N!$ trajectories we choose can be fixed by arbitrary stipulation, since any two such trajectories agree, for all times, not only on the values of all permutation-invariant quantities, but also on the biography of all N particles. It must be emphasised that the same arbitrary stipulation was incumbent on the haecceitist all along: after all, the haecceitistic trajectory $\pi^*\Xi$, for any of the $N! - 1$ permutations $\pi \in S_N \setminus \{e\}$, is just as good a representative as Ξ of the same physical history.[12] The association between haecceitistic and anti-haecceitistic trajectories for

[11] What follows is a fleshing out of remarks made by Leinaas & Myrheim (1977, p. 5).

[12] This is not to say that the haecceitist is committed to *identify* all $N!$ trajectories $\pi^*\Xi$ for some $\pi \in S_N$. That would be anti-haecceitism. Rather, the point is that, while for the haecceitist the trajectories Ξ and $\pi^*\Xi$

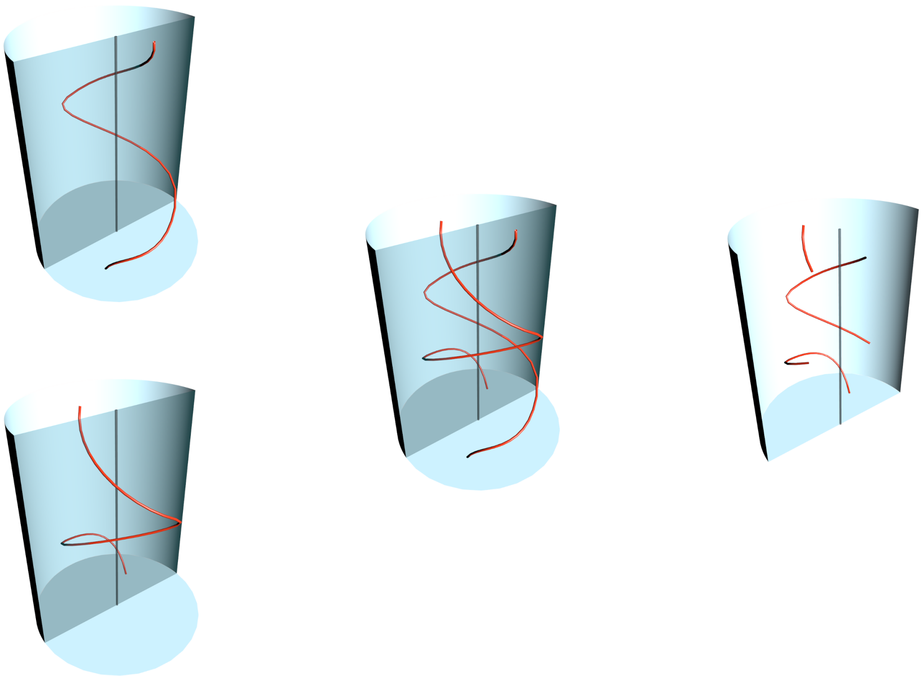

Figure 1: Each image represents the extended configuration space (with time along the vertical axis) for two particles in two-dimensional space (the centre of mass degrees of freedom have been supressed). The two images on the lefthand side represent haecceitistic trajectories (red curves) related by a permutation of the particles, realised on extended configuration space by a 180° rotation around the excised collision point (central grey pole). Each of these trajectories defines a unique anti-haecceitistic trajectory, shown on the righthand side. (This trajectory is continuous, despite appearances, by virtue of an appropriate identification of points on the boundary.) Conversely, the anti-haecceitistic trajectory defines the *pair* of anti-haecceitistic trajectories on the left, but does not choose between them. The central image represents this pair of haecceitistic trajectories in the same extended configuration space. It is clear that the trajectories do not intersect: this would involve the trajectories passing through the collision points, which have been excised.

the example of two particles in 2-dimensional space is illustrated in Figure 1.

So in the case of classical particle mechanics, we have two rival metaphysical positions, qualitative and transcendental individuation, associated with empirically equivalent theories. In fact, we may go further: the association between haecceitistic and anti-haecceitistic trajectories outlined above entails that an advocate of either metaphysical position may use either theory to adequately represent the target system. This is because the qualitative individuation criteria implied by the demand for continuous trajectories can be used, assuming the impenetrability of the particles, as a surrogate for transcendental individuation criteria, and *vice versa*.

5 Permutations in classical statistical mechanics

The equivalence outlined in the previous Section between classical particle theories conceived according to transcendental and qualitative individuation relied on the fact that the two modes of individuation can act as surrogates for one another. That fact in turn relied on the impen-

represent distinct possibilities (assuming $\pi \neq e$), it is a matter of convention *which* possibility is represented by which trajectory.

etrability of the particles, so that any two particles have distinct states—in other words: any haecceitistic joint state belongs to a permutability class on which the permutations act freely. In this Section, we will see that, when considering probability distributions over these states the two modes of individuation can produce very different verdicts.

5.1 The haecceitistic and anti-haecceitistic theories

We will begin by defining the haecceitistic state space. This is the space $\mathcal{M}(\Gamma_N) \ni \mu$ of Radon measures, which I will call *distributions*, on the haecceitistic phase space Γ_N (still assuming that collision points have been removed), whose elements I will now call haecceitistic *microstates*. The associated algebra of quantities $\langle \mathcal{A} \rangle$ is generated by those of the non-statistical particle theory by integrating over the space of microstates: for each quantity $f \in \mathcal{A}$, $\langle f \rangle \in \langle \mathcal{A} \rangle$, where

$$\langle f \rangle(\mu) := \int_{\Gamma_N} f(\Xi) \, \mathrm{d}\mu(\Xi) \ . \tag{4}$$

We can now define lifted permutations on the states $\mu \in \mathcal{M}(\Gamma_N)$. For any Borel set $\Sigma \subseteq \Gamma_N$ and any permutation $\pi \in S_N$, define the lifted permutation $\pi^\star : \mathcal{M}(\Gamma_N) \to \mathcal{M}(\Gamma_N)$ such that

$$(\pi^\star \mu)(\Sigma) := \mu(\pi^* \Sigma) \tag{5}$$

where $\pi^*\Sigma := \{\pi^*\Xi : \Xi \in \Sigma\}$ and $\pi^* : \Gamma_N \to \Gamma_N$ is the lifted permutation on microstates associated with π defined in the previous Section.

We can also define the sub-algebra $\langle \mathcal{A} \rangle_{PI}$ of permutation-invariant quantities; these are the quantities $\langle f \rangle$ such that $\langle f \rangle(\pi^\star \mu) = \langle f \rangle(\mu)$ for all distributions $\mu \in \mathcal{M}(\Gamma_N)$ and all permutations $\pi \in S_N$. It is not surprising that these are precisely the quantities generated, in the manner of equation (4), by the permutation-invariant quantities in the algebra $\mathcal{A}_{PI}$, defined in the previous Section; so we may also equally denote this algebra by $\langle \mathcal{A}_{PI} \rangle$.

The anti-haecceitistic theory may be constructed in three ways, all equivalent. The first way is to proceed as before: we consider the theory $\langle \mathcal{M}(\Gamma_N), \langle \mathcal{A}_{PI} \rangle \rangle$, identify permutability classes and quotient appropriately. In this case, as is generally true, states are permutable if they are related by a lifted permutation: i.e. $\mu \sim \mu'$ if $\mu' = \pi^\star \mu$ for some $\pi \in S_N$. However, in contrast to the microstates, states in this theory may be permutable even though they are not related by a lifted permutation. Specifically, if μ and μ' are permutable, then either one is also permutable with any convex combination $\lambda\mu + (1-\lambda)\mu'$, where $\lambda \in [0,1]$. The resulting state space $\widetilde{\mathcal{M}(\Gamma_N)/\sim}$ comprises permutability classes of distributions on Γ_N and the resulting algebra $\widetilde{\langle \mathcal{A}_{PI} \rangle}$ comprises quantities $\widetilde{\langle f \rangle}$ such that $\widetilde{\langle f \rangle}([\mu]) := \langle f \rangle(\pi^\star \mu)$, for any $\pi \in S_N$ and all $\langle f \rangle \in \langle \mathcal{A}_{PI} \rangle$.

The second way is to impose permutation-invariance on the states of the haecceitistic theory: i.e. $\pi^\star \mu = \mu$. The resulting state space $\mathcal{M}_{PI}(\Gamma_N)$ contains all and only the permutation-invariant distributions over haecceitistic microstates. The associated algebra may still be taken to be $\langle \mathcal{A} \rangle$ (where the quantities are suitably restricted to the permutation-invariant distributions), since the permutation-invariance of the states ensures permutation-invariance of the quantities: i.e. for all $\langle f \rangle \in \langle \mathcal{A} \rangle$, $\langle f \rangle(\pi^\star \mu) = \langle f \rangle(\mu)$. But note now that distinct quantities $f, g \in \mathcal{A}$ defined on microstates in Γ_N may correspond to the same quantity $\langle f \rangle = \langle g \rangle$ defined on distributions in $\mathcal{M}_{PI}(\Gamma_N)$.

The third way is to take as the states all Radon measures on the anti-haecceitistic phase space Γ_N/S_N, and to take as the algebra of quantities the algebra $\langle \tilde{\mathcal{A}}_{PI} \rangle$ generated, in a manner analogous to equation (4), from the algebra $\tilde{\mathcal{A}}_{PI}$ defined on anti-haecceitistic microstates.

The third way may appear to be the most honestly anti-haecceitistic, but all three ways lead to models $\langle \mathcal{M}(\Gamma_N)/\!\sim, \widetilde{\langle \mathcal{A}_{PI} \rangle}\rangle$, $\langle \mathcal{M}_{PI}(\Gamma_N), \langle \mathcal{A} \rangle\rangle$ and $\langle \mathcal{M}(\Gamma_N/S_N), \langle \tilde{\mathcal{A}}_{PI} \rangle\rangle$ any two of which are isomorphic.[13] That is, we may consider permutability classes of distributions on Γ_N, *or* permutation-invariant distributions on Γ_N, *or* arbitrary distributions on Γ/S_N, all equivalently. This is despite the fact that the first two theories are defined over *haecceitistic* microstates and the third is defined over *anti-haecceitistic* microstates. These equivalences are useful: of course, it means that we can use whichever formulation is the most convenient for the problem at hand. But further, these equivalences warrant scepticism towards any line of argument which relies on any one such formulation and cannot be extended to the other two. One such line of argument is discussed in Section 5.3.

5.2 Statistics and entropy in the rival theories

Do the haecceitistic and anti-haecceitistic theories make empirically inequivalent predictions? If so, we would seem to have empirical grounds for preferring one metaphysical doctrine over the other. Following Huggett (1999a), there are broadly two proposals: one, appealing to equilibrium distributions, appears to favour haecceitism; the other, appealing to Gibbs' paradox, appears to favour anti-haecceitism. Again following Huggett, and Saunders (2013), these appearances are deceptive.

The equilibrium distribution in question is the Maxwell-Boltzmann velocity distribution, derived through a combinatorial argument (see Frigg 2008, §2.2), which shows it to correspond to the macrostate with largest Lebesgue volume in the joint phase space. The details are as follows. The single-system phase space Γ is divided into c cells, each with Lebesgue volume ω and a characteristic mean energy E_i. Macrostates are then characterised by occupation numbers $(n_1, \ldots, n_c)$, where n_i is the population of particles whose states lie in the ith cell. In the haecceitistic theory, each macrostate defines a region of Γ_N consisting of the number

$$\frac{N!}{\prod_{i=1}^{c} n_i!} \tag{6}$$

of (typically disconnected) cells, each with Lebesgue volume ω^N. This volume is maximal, subject to the constraints of conservation of particle number ($\sum_i n_i = N$) and of energy ($\sum_i n_i E_i = E$), for the Maxwell Boltzmann distribution

$$n_i = N e^{-\beta(E_i - \mu)} \tag{7}$$

where β and μ are constants, determined by the two constraints.

Seemingly crucial to the derivation is the fact that the number (6) of cells in the joint phase space is maximal for the Maxwell-Boltzmann distribution (7); but in the anti-haecceitistic theory, there is only *one* cell in the joint phase space corresponding to those occupation numbers. However, an anti-haecceitistic derivation of (7) can also be given. While the haecceitistic derivation relies on a variable number of cells of constant volume, the anti-haecceitistic derivation relies on a constant number (one) of cells of variable volume—variable, due to quotienting. The cell in Γ_N/S_N corresponding to the occupation numbers $(n_1, \ldots, n_c)$ has Lebesgue volume

$$\frac{\omega^N}{\prod_{i=1}^{c} n_i!} \tag{8}$$

[13] We define the following three maps:
$\iota_1 : \mathcal{M}(\Gamma_N)/\!\sim \,\to \mathcal{M}_{PI}(\Gamma_N)$ such that $\iota_1([\mu]) := \frac{1}{N!}\sum_{\pi \in S_N} \pi^\star \mu$;
$\iota_2 : \mathcal{M}_{PI}(\Gamma_N) \to \mathcal{M}(\Gamma_N/S_N)$ such that $(\iota_2(\mu))([\Sigma]) := \mu(\Sigma)$, where $[\Sigma] := \{[\Xi] : \Xi \in \Sigma\}$;
$\iota_3 : \mathcal{M}(\Gamma_N/S_N) \to \mathcal{M}(\Gamma_N)/\!\sim$ such that $(\iota_3(\mu))(\Sigma) := [\mu](\Sigma)$.
Any combination of such maps establishes an isomorphism between the appropriate models.

The Boltzmann entropy in each case is the same, except for an additive constant $k_B \log N!$.

As for the Gibbs paradox, there is some dispute about what exactly the paradox is (see e.g. Uffink 2006, §5.2), but the broad theme is the entropy of mixing for gases. The particular topic of interest here is how one obtains a statistical mechanical entropy for an isolated gas which is extensive, i.e. scales linearly with volume and total particle number. A naïve calculation using the microcanonical ensemble goes *via* the (Lebesgue) volume, in the joint phase space Γ^N, of the energy $E = \frac{3}{2}NkT$ hypersurface, which is the product of the spatial volume V^N, where each of the N particles in the gas is confined to volume V, and the surface "area" of the $(3N-1)$-dimensional momentum hypersphere with radius $\sqrt{2mE} = \sqrt{3NmkT}$:

$$W_{\textit{naïve}} = \frac{2V^N \pi^{\frac{3}{2}N}}{\Gamma(\frac{3}{2}N)}(3NmkT)^{\frac{3N-1}{2}} \approx \frac{2V^N \pi^{\frac{3}{2}N}}{(\frac{3}{2}N)!}(3NmkT)^{\frac{3}{2}N} \; . \tag{9}$$

With the usual prescription for the entropy, one then obtains

$$S_{\textit{naïve}}(N, V, T) = k \log W_{\textit{naïve}} \approx Nk\left[\log V + \frac{3}{2}\log T + c\right] \tag{10}$$

where c is some constant. Extensivity demands that

$$S(\alpha N, \alpha V, T) = \alpha S(N, V, T) \tag{11}$$

for any $\alpha \in \mathbb{R}_+$, but this fails for $S_{\textit{naïve}}$. For example, if our gas is separated into two chambers of equal volume $\frac{1}{2}V$ and density $\frac{N}{V}$ by a partition, and we assume that the total entropy is just the sum of the entropies of the gases either side of the partition, then gently removing the partition produces a rise of entropy by $Nk \log 2$. But this is a reversible thermodynamical process: since the gas is in equilibrium after the partition is removed, i.e. it has constant temperature and density throughout the volume V, one can separate the gases again, performing negligible work, just by gently reintroducing the partition.

Gibbs' solution to this problem was to pass to the quotient phase space Γ_N/S_N. The consequence is a division by $N!$ of the naïve phase space volume $W_{\textit{naïve}}$, and one obtains the Sackur-Tetrode equation

$$S(N, V, T) \approx Nk\left[\log\left(\frac{V}{N}\right) + \frac{3}{2}\log T + c\right] \; , \tag{12}$$

which yields an extensive entropy. This seems to suggest that reconciliation with thermodynamics requires an anti-haecceitistic statistical mechanical theory. However, a haecceitistic derivation of (12) can also be given. Following a suggestion of Ehrenfest & Trkal (1921), developed by van Kampen (1984), the haecceitist demands that we take into account *all* molecules in the universe of the same kind as those in the chamber. Supposing the total number of such molecules is $M >> N$, then we must multiply the naïve phase space volume by the number of ways that N of them appear in the chamber. This number is

$$\begin{pmatrix} M \\ N \end{pmatrix} \approx \frac{M^N}{N!} \tag{13}$$

and leads to an entropy differing from (12) by an irrelevant additive constant, also a multiple of N; so the entropy is again extensive.

So it would appear that we can save the phenomena regarding equilibrium distributions and the mixing of gases under both haecceitism and anti-haecceitism. Huggett (1999a) counsels metaphysical scepticism in response. Saunders (2013) points out, in the case of Gibbs' paradox, that the haecceitistic and anti-haecceitistic solutions differ over whether the gas is properly treated as an open or closed system. Regarding this point, note that the haecceitistic solution crucially involves an appeal to the total number of molecules of the same kind in the *entire* universe, and the assumption that those in the chamber comprise a tiny fraction of this total—all this despite the fact that the molecules inside the box cannot mix with those outside.

5.3 Individuation in permutation-invariant classical statistical mechanics

Our first two anti-haecceitistic formulations, having states defined over haecceitistic microstates, retain the resources of transcendental individuation of the particles. In particular, this allows us to define a marginal distribution on the single-system phase space Γ for each of the N system labels. For example, given any joint distribution $\mu \in \mathcal{M}_{PI}(\Gamma_N)$ in the second theory, the marginal distribution μ_1 associated with system 1, whose microstate $\xi_1 \in \Gamma$ lies in the first entry of each haecceitistic microstate $(\xi_1, \ldots, \xi_N) \in \Gamma_N$, may be defined as follows. For each Borel set $\sigma \subseteq \Gamma$,

$$\mu_1(\sigma) := \int_{\sigma \times \Gamma \times \ldots \times \Gamma} \mathrm{d}\mu(\xi_1, \ldots, \xi_N) \tag{14}$$

Similarly, the marginal distribution μ_2 associated with system 2 is given by

$$\mu_2(\sigma) := \int_{\Gamma \times \sigma \times \ldots \times \Gamma} \mathrm{d}\mu(\xi_1, \ldots, \xi_N) \tag{15}$$

and so on. However, since any joint distribution in this theory is permutation invariant, we have $\mathrm{d}\mu(\xi_1, \ldots, \xi_N) = \mathrm{d}\mu(\xi_{\pi(1)}, \ldots, \xi_{\pi(N)})$ for all $\pi \in S_N$, and so the marginal distributions are all identical: $\mu_1 = \mu_2 = \ldots = \mu_N$. Therefore, *if we continue to maintain transcendental individuation of the systems*, we are forced to conclude that all systems are described by the *same* marginal distribution.[14] It follows from this that the systems, transcendentally individuated, are indiscernible by means of monadic properties, i.e. they are absolutely indiscernible, or (in the jargon of Section 2.1) permutable.

This drastic result is simply unavailable in our third, honestly anti-haecceitistic formulation. Here we cannot even be *tempted* to appeal to transcendental individuation, since the system labels have been "rubbed out" in the microstates by passing to the quotient under S_N. Proposals for qualitative individuation are thin on the ground, but one proposal might run as follows. Select N disjoint Borel sets $\sigma_i \subseteq \Gamma$ of the single-system phase space. These N sets define a Borel set

$$\Sigma(\{\sigma_1, \ldots, \sigma_N\}) := \left(\bigcup_{\pi \in S_N} \sigma_{\pi(1)} \times \ldots \times \sigma_{\pi(N)} \right) / S_N \tag{16}$$

of microstates in the anti-haecceitistic joint phase space Γ_N / S_N. If now the joint distribution μ has support in $\Sigma(\{\sigma_1, \ldots, \sigma_N\})$; i.e. $\mu(\Sigma(\{\sigma_1, \ldots, \sigma_N\})) = 1$, then we can, according to the proposal, say with certitude that:

- one of the particles has a microstate lying in σ_1;
- one of the particles has a microstate lying in σ_2;

  ⋮

[14] This is essentially the conclusion drawn by Bach (1997, pp. 7-8), discussed by Saunders (2013). Bach goes on further to draw the dizzying conclusion that 'indistinguishable particles have no trajectories.'

- one of the particles has a microstate lying in σ_N.

Since the σ_i are all disjoint, they serve to individuate all N of the particles; we may them call them *individuation criteria*. We thereby avoid the conclusion above that each particle has the same state. (How can they have the same state? Their respective microstates lie in disjoint parts of Γ.) There is no further question *which* particle's microstate lies in *which* Borel set σ_i, since we have no means by which to individuate the particles except by appeal to their states.

Developing the proposal further, we may extract single-system marginal distributions as follows. For example, associated with the set σ_1 we can define the marginal distribution μ_{σ_1} such that, for any Borel set $\tau \subseteq \Gamma$:

$$\mu_{\sigma_1}(\tau) := \int_{\Sigma(\{\tau\cap\sigma_1,\sigma_2,\ldots,\sigma_N\})} \mathrm{d}\mu([(\xi_1,\ldots,\xi_N)]) \tag{17}$$

Thus, the particle whose microstate lies in σ_1 has the marginal distribution μ_{σ_1}, and so on. And clearly, $\mu_{\sigma_1},\ldots,\mu_{\sigma_N}$ are all distinct—in fact they have mutually disjoint supports. All this relies, remember, on the joint distribution μ having support in $\Sigma(\{\sigma_1,\ldots,\sigma_N\})$. This set has a natural Cartesian product structure, which factorises into disjoint parts σ_i of the single-system phase space Γ.

The marginals calculated through qualitative individuation bear a simple relation to the marginals calculated through transcendental individuation. Let $\overline{\mu} := \mu_1 = \ldots = \mu_N$ be the single-system marginal defined in equation (14). Then, the marginal for the particle individuated e.g. by the criterion σ_1 is

$$\mu_{\sigma_1}(\tau) = \frac{\overline{\mu}(\tau\cap\sigma_1)}{\overline{\mu}(\sigma_1)} =: \overline{\mu}(\tau|\sigma_1) \tag{18}$$

where I have defined the relative measure $\overline{\mu}(\tau|\sigma_1)$ in the obvious way.

Nothing so far has placed any sane restrictions on the sets σ_i (e.g., they may be horribly miscellaneous unions of disjoint regions in Γ). I will not give an account of such restrictions here, nor investigate how individuation criteria might be extended to qualitative individuation over time, or by appeal to irreducible relations holding between the particles; these are all directions for future work. However, a reassuring result is as follows. Choose N distinct points ξ_i in Γ to define N *maximally specific* individuation criteria $\sigma_i = \{\xi_i\}$. Then $\Sigma(\{\sigma_1,\sigma_2,\ldots,\sigma_N\})$ is the singleton of the microstate $[(\xi_1,\ldots,\xi_N)] = \{(\xi_{\pi(1)},\ldots,\xi_{\pi(N)}) : \pi \in S_N\}$. If the joint distribution has support on this singleton—i.e. it is a Dirac delta distribution centred on $[(\xi_1,\ldots,\xi_N)]$, which makes it an extremal state in the state space $\mathcal{M}(\Gamma_N/S_N)$—then each associated marginal μ_i is a Dirac delta distribution centred on ξ_i. That much is to be expected. However, also in this special case, the marginals μ_i uniquely determine the joint distribution μ, since there is only one state in $\mathcal{M}(\Gamma_N/S_N)$ which can produce those marginals *via* equation (17) and its analogues. In other words: maximally specific joint states determine, and are determined by, maximally specific states possessed by the constituent particles.

6 Permutations in quantum mechanics

There are a number of routes to a permutation-invariant quantum theory for a constant number N systems. Broadly speaking, there are three main routes:

1. *Implement permutation invariance on some corresponding permutation-non-invariant quantum theory for N systems.* This route leads to the theory of group representations; specifically, the irreducible representations of S_N. I will outline this route in Section 6.1.

2. *Quantize some corresponding permutation-invariant classical theory for N systems.* This route has permutation-invariance built in at the outset, and leads to considerations of non-equivalent quantizations brought about by topological features of the underlying classical configuration space. I will outline this route in Section 6.3.

3. *Restrict to the N-system subspace in some appropriate corresponding quantum field theory.* This route also has permutation-invariance built in at the outset, and is connected to one of the great achievements of quantum field theory, the spin-statistics theorem.

I will say very little least about route 3, or about quantum field theory more generally.[15] One would be forgiven for considering this route to be the most enlightening when it comes to the origins of permutation invariance, since (regarding route 1) our world is more accurately described by quantum field theory than many-particle quantum mechanics, and (regarding route 2) we ought to be cautious about gleaning insights into a quantum theory by appeal to some classical theory of which it happens to be the quantisation. Nevertheless, I will press on with brief outlines of routes 1 and 2.

6.1 The representation of permutations on joint Hilbert spaces

We begin with the quantum theory for a single system. Standardly, this is some separable Hilbert space $\mathcal{H}$ and an associated algebra of quantities, all of which are linear operators on $\mathcal{H}$. I will not go into the details here about where $\mathcal{H}$ and $\mathfrak{a}$ "come from", but simply take them as given.

Representing an assembly of such systems, say N of them, standardly proceeds as follows. First we form the joint Hilbert space and joint algebra

$$\mathcal{H}^N := \underbrace{\mathcal{H} \otimes \ldots \otimes \mathcal{H}}_{N} ; \qquad \mathcal{A} := \underbrace{\mathfrak{a} \otimes \ldots \otimes \mathfrak{a}}_{N} . \tag{19}$$

The model $\langle \mathcal{H}^N, \mathcal{A} \rangle$ forms the arena for the haecceitistic theory of N equivalent particles.

Lifted permutations may be defined on $\mathcal{H}^N$, and here we get into the group representation theory of S_N. There is an obvious map $U : S_N \to \mathcal{U}(\mathcal{H}^N)$ from the group S_N of permutations to the unitary operators $\mathfrak{U}(\mathcal{H}^N) \subset \mathcal{A}$ such that $U(\pi)$ implements the permutation π on joint states. We define U by its action on product states and then extend by linearity to all vectors in the joint Hilbert space. Given any permutation $\pi \in S_N$ we have, for any product state $|\psi_1\rangle \otimes \ldots \otimes |\psi_N\rangle \in \mathcal{H}^N$,

$$U(\pi)|\psi_1\rangle \otimes \ldots \otimes |\psi_N\rangle := |\psi_{\pi(1)}\rangle \otimes \ldots \otimes |\psi_{\pi(N)}\rangle \tag{20}$$

(compare with equation (2) in Section 4.1). U so-defined constitutes a representation of S_N in the technical sense, which is that for any $\pi_1, \pi_2 \in S_N$, $U(\pi_2\pi_1) = U(\pi_2)U(\pi_1)$, and so U is a group homomorphism. And furthermore, since each $U(\pi)$ is a unitary operator, U is a *unitary* representation of S_N.

Since two states differing only up to a global phase factor yield the same expectation values, one may reasonably wonder why we cannot instead demand only that $U(\pi_2\pi_1) = e^{i\omega(\pi_2,\pi_1)}U(\pi_2)U(\pi_1)$, so that we obtain a group homomorphism up to a phase factor. Such

[15] For details about quantum field theory with statistics other than Bose-Einstein and Fermi-Dirac, see Greenberg *et al* (1964), Stolt & Taylor (1970), Doplicher *et al* (1974) and Ohnuki & Kamefuchi (1982); a philosophical presentation is given by Baker *et al* (2015), with further references.

representations are called *projective* unitary as opposed to just unitary (or *linear* unitary). (When considering lifted permutations of the form $U(\pi_3\pi_2\pi_1)$, we find that ω must obey the *cocycle equation* $\omega(\pi_2,\pi_1)+\omega(\pi_3,\pi_2\pi_1)=\omega(\pi_3,\pi_2)+\omega(\pi_3\pi_2,\pi_1)$.) One answer, provided by Read (2003), is that merely projective representations of S_N violate locality: i.e. the rigorous derivation of particle statistics in quantum field theory obeying a local dynamics (with a local Hamiltonian) are at odds with these merely projective representations. However, there is a related matter whether we should be thinking of permutations in terms of the group S_N at all. I will briefly return to this in Section 6.3.

For now assuming (20), let us delve briefly into the group representation theory of S_N; complete treatments may be found in e.g. Tung (1985, Chs. 3 & 5) and Sternberg (1994, Ch. 2). We find that the representation U decomposes into a direct sum of *irreducible* unitary representations D_λ, which I will call *irreps*. Irreps come in a variety of types, labelled by λ, according to how joint states transform under the lifted permutations $U(\pi)$, and typically there will be many copies of the same irrep in the decomposition of U. One type of irrep is D_+, for which $D_+(\pi)=1$ for all $\pi\in S_N$, and which corresponds by definition to *bosonic* states. Another type of irrep is D_-, for which $D_-(\pi)=(-1)^{\deg\pi}$, where $\deg\pi$ is the number of pairwise swaps involved in the permutation π, and which corresponds by definition to *fermionic* states. If the number of particles is 3 or more, we find in addition to the bosonic and fermionic irreps a variety of multi-dimensional irreps, corresponding to what are known as *paraparticle* or *parastatistical* states.

Each copy of each irrep D_λ occurring in U acts on only a small subspace of the joint Hilbert space $\mathcal{H}^N$; these subspaces are called *irreducible invariant subspaces*, or i.i.s.s. (Messiah & Greenberg 1964 call these i.i.s.s *generalised rays*.) Each i.i.s. is the smallest non-trivial subspace left invariant under action by the lifted permutations $U(\pi)$, which is what earns them and the associated irreps D_λ the designation 'irreducible'. Just as the representation U is the analogue of a group action of S_N in the classical case, the i.i.s.s are the group representation analogue of classical orbits. Bosonic and fermionic i.i.s.s are 1-dimensional, corresponding to the irreps $D_\pm$ yielding simple scalars: $D_\pm(\pi)=\pm 1$. Paraparticle i.i.s.s are multi-dimensional, and so their associated irreps D_λ are represented by matrices.

Each irrep is associated with characteristic large-particle-number behaviour. Bosons are associated with Bose-Einstein statistics; fermions with Fermi-Dirac statistics. Both are distinguished from the classical Maxwell-Boltzmann statistics, seen in Section 5. Incidentally, the reason for this difference between classical and quantum statistics is often mistakenly linked to the dispute between haecceitism and anti-haecceitism. In fact, it is due to the ways states are "counted" in the two theories: the measure over states in the classical case is continuous (phase space volume) and in the quantum case it is discrete (dimension count). This is a point emphasised by Huggett (1999b) and Saunders (2006b, 2013).

6.2 Quantum permutation invariance

We may now define what it is for any quantity in the joint algebra to be permutation-invariant. A quantity $A\in\mathcal{A}$ is permutation-invariant iff: for any permutation $\pi\in S_N$ and any joint vector state $\Psi\in\mathcal{H}^N$,

$$\langle U(\pi)\Psi, AU(\pi)\Psi\rangle = \langle\Psi, A\Psi\rangle \tag{21}$$

(compare with equation (3) in Section 4.2). This condition is widely known as the *Indistinguishability Postulate*; as far as I know, the term originates with Messiah & Greenberg (1964). Since this condition holds for all vectors in $\mathcal{H}^N$, it may be rephrased as an operator identity.

For any permutation $\pi \in S_N$:

$$U(\pi)^\dagger A U(\pi) = A \ , \qquad \text{i.e.} \qquad [A, U(\pi)] = 0 \ . \tag{22}$$

This condition provides a criterion for membership in the permutation-invariant algebra $\mathcal{A}_{PI} \subset \mathcal{A}$ of quantities.

Now appealing to a powerful theorem in group representation theory known as *Schur's Lemma*, we may deduce the following two facts. First, the eigenspaces of any permutation-invariant quantity $A \in \mathcal{A}_{PI}$ respects the decomposition of U into irreps, in the sense that A's eigenspaces are superspaces of i.i.s.s. It follows that A takes the same expectation value $\langle \Psi, A\Psi \rangle$ on any joint vector state Ψ lying in the same i.i.s. This is the quantum analogue of the fact in classical mechanics that any permutation-invariant quantity takes the same value on every joint state in the same orbit—but this is trivial in the case of the bosonic or fermionic states, whose i.i.s.s are 1-dimensional.

Second, the permutation-invariant algebra $\mathcal{A}_{PI}$ acts *reducibly* on the haecceitistic joint Hilbert space $\mathcal{H}^N$, so that $\mathcal{H}^N$ naturally decomposes into a number of sectors $\mathcal{H}^N_\lambda$, each corresponding to an irrep type λ. In other words: any permutation-invariant quantity $A \in \mathcal{A}_{PI}$ becomes block-diagonalised by irrep type, so that transition amplitudes vanish, $\langle \Psi, A\Phi \rangle = 0$, for any joint states Ψ and Φ which belong to different sectors $\mathcal{H}^N_\lambda$. So, for example, any transition amplitude between bosonic and fermionic sectors vanishes. In the physicists' jargon, *symmetry type* (boson, fermion, etc.) is *superselected* by the restriction to permutation-invariant quantities. This second fact has no classical analogue.

It may now be wondered: What is the quantum analogue of passing to the quotient of $\langle \mathcal{H}^N, \mathcal{A}_{PI} \rangle$ under permutation symmetry; i.e. what is the arena for anti-haecceitistic quantum theory? We proceed in two stages:

(i) First we restrict to the sector $\mathcal{H}^N_\lambda$ corresponding to some chosen irrep λ; i.e. we restrict attention to just one 'symmetry type', e.g. the bosonic or fermionic states. This stage has no classical analogue.

(ii) Then, if required (i.e. if the irrep is multi-dimensional, as for paraparticles), we define a new joint Hilbert space $\mathcal{H}^N_\lambda / S_N$, any ray of which corresponds to an i.i.s. in $\mathcal{H}^N_\lambda$ (so that each generalized ray in $\mathcal{H}^N_\lambda$ is mapped to a ray $\mathcal{H}^N_\lambda / S_N$). Each permutation-invariant quantity $A \in \mathcal{A}_{PI}$ has a well-defined, unique counterpart $\tilde{A}_\lambda \in \tilde{\mathcal{A}}_\lambda$ which is a linear operator on $\mathcal{H}^N_\lambda / S_N$. (Essential here is the fact that A yields the same expectation value on any state in the same i.i.s. in $\mathcal{H}^N_\lambda$.) As would be expected, any lifted permutation is represented on $\mathcal{H}^N_\lambda / S_N$ as simple multiplication by ± 1, depending on the permutation and the irrep λ.

Note that stage (ii) is unnecessary for all of the elementary particles we believe to exist, which are all either bosons or fermions, since the bosonic and fermionic i.i.s.s are already 1-dimensional. The justification for stage (i) comes from the superselection of symmetry type by the imposition of permutation invariance on the algebra of quantities: $\mathcal{A}_{PI}$ acts irreducibly when restricted to each sector $\mathcal{H}^N_\lambda$. What is surprising is that we obtain not one but several anti-haecceitistic models $\langle \mathcal{H}^N_\lambda / S_N, \tilde{\mathcal{A}}_\lambda \rangle$, each corresponding to its own symmetry type λ.

The restriction to either the bosonic sector $\mathcal{H}^N_+$ or the fermionic sector $\mathcal{H}^N_-$ is widely known as the *Symmetrisation Postulate*. All evidence so far suggests that the postulate is true, though there is an interesting history here. Before the acceptance of quark colour, it had been proposed, chiefly by O. W. Greenberg, that quarks might be paraparticles. (For more, see French 1995.) We now believe quarks to be fermions, but this raises the general question why paraparticles, perfectly allowed under permutation invarance, are not found in nature. Dürr *et al* (2007) make the case that de Broglie-Bohm theory may have the edge over its rivals here, since the requirement that corpuscle trajectories be determinate appears to rule out multi-dimensional irreps of S_N. Baker *et al* (2015) point to a result in the framework of algebraic quantum field theory, that any local quantum field exhibiting paraparticle statistics is equivalent to one with bosonic or fermionic statistics plus a new internal degree of freedom with accompanying constraints; they argue on the basis of this that the issue of particle statistics may be a matter of mere convention, rendering the Symmetrisation Postulate more a decision than a discovery.

The fact that, for bosons and fermions, the i.i.s.s under the lifted permutations are all 1-dimensional means that we have here a case where haecceitism and anti-haecceitism trivially agree: for these joint states are fixed (up to global phase) by any lifted permutation. It is commonly claimed of these states that they represent particles which are indiscernible by means of monadic properties, i.e. absolutely indiscernible. In fact this is the wide consensus in the quantum identity literature. (The consensus began with Margenau (1944) and continued with Post (1963), French & Redhead (1988), van Fraassen (1991, Ch. 11), Butterfield (1993), Saunders (2003a, 2003b, 2006a), French & Krause (2006, §4.2.1), Muller & Saunders (2008), Muller & Seevinck (2009), Ladyman & Bigaj (2010), Caulton (2012) and Huggett & Norton (2012).) However, this conclusion can be reached only on the assumption that the particles in question are denoted by the order of the factor Hilbert spaces in the tensor product $\mathcal{H}^N$, so that a permutation of the factor Hilbert spaces under each $U(\pi)$ is taken to correspond to a permutation of the particles. This is tantamount to a commitment to transcendental individuality, as defined in Section 3.3. A proponent of qualitative individuality would contend rather that the lifted permutations $U(\pi)$ correspond to no real permutation at all. Proposals along the lines of qualitative individuality may be found in Huggett & Imbo (2009), Dieks & Lubberdink (2011) and Earman (2015).

6.3 The topological approach to quantum statistics

The topological approach bypasses the "full" joint Hilbert space $\mathcal{H}^N$ by directly quantising the corresponding anti-haecceitistic classical theory, whose configuration space is $(\mathcal{Q}^N \setminus \Delta)/S_N$. Even though the haecceitistic configuration space $\mathcal{Q}^N$ may be simply connected, non-trivial topological features are picked up by the removal of collision configurations, or by passing to the quotient, or both. These topological features give rise to a number of inequivalent ways to quantise.

I will outline here only the case in which the single-system configuration space is Euclidean space of two or more dimensions, $\mathcal{Q} = \mathbb{R}^d$, where $d \geqslant 2$.[16] Due to a celebrated result by Stone and von Neumann, there is only one way to quantize the joint configuration space $\mathcal{Q}^N = \mathbb{R}^{Nd}$; i.e. there is only one representation, up to unitary equivalence, of the canonical commutation relations (CCRs):

$$[Q^i, Q^j] = [P_i, P_j] = 0; \qquad [Q^i, P_j] = i\hbar\delta^i_j \ . \tag{23}$$

This is the Schrödinger representation, whose Hilbert space is $\mathcal{L}^2(\mathbb{R}^{Nd})$, the space of equivalence classes $[f]$ of square-integrable complex functions $f : \mathbb{R}^{Nd} \to \mathbb{C}$ on the configuration space $\mathbb{R}^{Nd}$.[17]

[16]Hansson *et al* (1992) consider the $d = 1$ case. Imbo *et al* (1990) consider more exotic configurations spaces, which lead to exotic particle statistics they call *ambistatistics*.

[17]The associated equivalence relation is $f \sim g$, iff $\|f - g\| = 0$, where $\|\cdot\|$ is the norm induced by the inner product defined on the RHS of (24).

The inner product is provided by

$$\langle [f], [g] \rangle := \int_{-\infty}^{\infty} \mathrm{d}^d\mathbf{x}\ f^*(\mathbf{x})g(\mathbf{x}) \tag{24}$$

and the Q^i and P_i, each defined on a dense domain, are given by

$$(Q^i f)(\mathbf{x}) := x^i f(\mathbf{x}) \ ; \qquad (P_i f)(\mathbf{x}) := -i\hbar \frac{\partial f}{\partial x^i} \equiv -i\hbar(\partial_i f)(\mathbf{x}) \tag{25}$$

(and then uniquely extended to self-adjoint operators).

One might wonder why we cannot use $\tilde{P}_i$ instead of P_i, where $\tilde{P}_i := P_i + \alpha_i(\mathbf{Q})$ and each $\alpha_i(\mathbf{x})$ is a real-valued function over $\mathbb{R}^{Nd}$.[18] After all, these operators clearly have the desired commutation relations with Q^i. The answer is that we can, so long as the α_i obey conditions which also ensure the desired commutation relations among the $\tilde{P}_i$ themselves, namely $[\tilde{P}_i, \tilde{P}_j] = 0$ (commutation of the Q^i is already taken care of). Using the Schrödinger prescription for P_i, and the fact that the P_i commute among themselves and that the $\alpha_i(\mathbf{Q})$, being functions of the Q^i, also commute among themselves, we must demand that

$$[\tilde{P}_i, \tilde{P}_j] = [\alpha_i(\mathbf{Q}), P_j] + [P_i, \alpha_j(\mathbf{Q})] = -i\hbar\,(\partial_i\alpha_j - \partial_j\alpha_i) = 0\ . \tag{26}$$

So satisfaction of the Heisenberg CCRs is equivalent, on these assumptions, to the condition $\partial_i\alpha_j - \partial_j\alpha_i = 0$, or $\mathbf{d}\boldsymbol{\alpha} = 0$, where $\mathbf{d}$ is the exterior derivative operator on the configuration space $\mathbb{R}^{Nd}$. *Since* $\mathbb{R}^{Nd}$ *is simply-connected,* $\mathbf{d}\boldsymbol{\alpha} = 0$ implies that $\boldsymbol{\alpha} = -\hbar\mathbf{d}\xi$, i.e. $\alpha_i = -\hbar\partial_i\xi$ for some real scalar field $\xi : \mathbb{R}^{Nd} \to \mathbb{R}$ (the multiplicative constant $-\hbar$ is for convenience, and ensures that ξ is dimensionless). The field ξ may be interpreted as a "local phase" function, implementing a gauge transformation of local phase at every point in configuration space.

The local phase function ξ defines a unitary transformation $U : f(\mathbf{x}) \mapsto (Uf)(\mathbf{x}) := e^{i\xi(\mathbf{x})}f(\mathbf{x})$, which reveals our new representation to be unitarily equivalent to the Schrödinger representation, as expressed in the following commutative diagrams:[19]

$$\begin{array}{ccc} f & \xmapsto{\ Q_i\ } & Q_i f \\ \downarrow{\scriptstyle U} & & \downarrow{\scriptstyle U} \\ Uf & \xmapsto{\ Q_i\ } & Q_i U f = U Q_i f \end{array} \qquad\qquad \begin{array}{ccc} f & \xmapsto{\ P_i\ } & P_i f \\ \downarrow{\scriptstyle U} & & \downarrow{\scriptstyle U} \\ Uf & \xmapsto{\ \tilde{P}_i\ } & \tilde{P}_i U f = U P_i f \end{array}$$

So our alternative representation $(Q^i, \tilde{P}_j)$ is unitarily equivalent to the Schrödinger representation (Q^i, P_j) only so long as the configuration space is simply connected. If the configuration space is *not* simply connected, then the vanishing of $\mathbf{d}\boldsymbol{\alpha}$ does not imply that $\boldsymbol{\alpha} = -\hbar\mathbf{d}\xi$ for some local phase function ξ. So we have a recipe for generating unitarily inequivalent representations of the CCRs: find a curl-free collection of real scalar fields α_i not all of whose closed-loop integrals are zero (i.e. a closed, non-exact connection $\boldsymbol{\alpha}$).

The gamut of available options for $\boldsymbol{\alpha}$ is intimately tied to the specific way in which the configuration space $\mathfrak{Q}$ fails to be simply connected. Equivalence classes of connections $\boldsymbol{\alpha}$ (where two connections are equivalent iff they give rise to unitarily equivalent representations of the

[18] In fact $\boldsymbol{\alpha} \equiv \alpha_i \mathrm{d}x^i$ earns an interpretation as a $U(1)$ connection field. Each α_i must be real-valued to ensure that $\tilde{P}_i$ is a symmetric operator. Imaginary components may be accommodated by suitably adjusting the Lebesgue measure $\mathrm{d}^d\mathbf{x}$ in the definition of the inner product. See e.g. Jauch (1968, ??) for more details.

[19] The crucial facts are that U commutes with each Q^i and that $(\tilde{P}_i U f)(\mathbf{x}) = [-i\hbar\partial_i + \alpha_i(\mathbf{x})]e^{i\xi(\mathbf{x})}f(\mathbf{x}) = [\hbar(\partial_i\xi)(\mathbf{x}) + \alpha_i(\mathbf{x})]e^{i\xi(\mathbf{x})}f(\mathbf{x}) - i\hbar e^{i\xi(\mathbf{x})}(\partial_i f)(\mathbf{x}) = (UP_i f)(\mathbf{x})$.

CCRs) are in one-to-one correspondence with the unitarily inequivalent one-dimensional irreps of $\pi_1(\mathfrak{Q})$, the *fundamental group of* $\mathfrak{Q}$: essentially, the homotopy equivalence classes of closed loops on $\mathfrak{Q}$, endowed with a natural group structure.[20]

We are interested in the configuration space $\mathfrak{Q} = (\mathbb{R}^{Nd} \setminus \Delta)/S_N$ (for $d \geqslant 2$), which is in all cases multiply connected. However, the source of topological non-triviality is different, depending on whether $d = 2$ or $d > 2$. Let us take the $d > 2$ case first.

The configuration space $\mathbb{R}^{Nd} \setminus \Delta$, after removal of collision configurations but before quotienting, remains simply connected if $d > 2$: any closed loop encircling a collision configuration can be continuously transformed away from the missing point. Therefore the assumption of impenetrability of the particles is, in a sense, innocent. However, the quotient space is multiply connected: $\pi_1((\mathbb{R}^{Nd} \setminus \Delta)/S_N) = S_N$. Since we assuming (for now) only 1-dimensional irreps, there are two inequivalent quantisations, corresponding to bosons or fermions. The appearance of these inequivalent quantisations is due solely to quotienting, i.e. anti-haecceitism regarding the classical configurations.

In the case $d = 2$, the configuration space becomes multiply connected by the removal of collision configurations: $\pi_1(\mathbb{R}^{2N} \setminus \Delta) = B_N$, where B_N is the braid group on N strands, introduced in Section 1.2. The quotient space $(\mathbb{R}^{2N} \setminus \Delta)/S_N$ has the same fundamental group; so in this case the appearance of inequivalent quantisation is due solely to the impenetrability of the particles, whether or not we take the classical configurations haecceitistically or anti-haecceitistically. The braid group B_N has a *continuum* of one-dimensional irreps: these include the familiar bosons and fermions, but also include so-called *anyons*. Anyonic statistics are now widely taken to account for the fractional quantum Hall effect, of which Wilczek (1990) offers an account and a collection of important papers.

So far we have restricted attention to the one-dimensional irreps of the fundamental group $\pi_1(\mathfrak{Q})$. Recall that this arose because we were considering mild mutations of the Schrödinger representation, in which a curvature-free connection term is added to the definition of the momentum. What about representations which do not relate to the Schrödinger representation in this simple way? By allowing multi-dimensional irreps of $\pi_1(\mathcal{Q})$, we recover, in the case $d > 2$ and $N > 2$, all of the (by now familiar) paraparticle representations, and in the case $d = 2$ we obtain multi-dimensional (i.e. non-Abelian) irreps of the braid group.

These results reassure us that quantisation of the joint configuration space and quotienting under S_N are commuting procedures, as pointed out by Landsman (2016). However, in the case $d = 2$, we will miss the anyonic representations if we proceed in the manner of Section 6.1 and define the joint Hilbert space and algebra by forming the standard tensor products, and then impose permutation invariance.

Considerations specific to the quantum theory provide new subtleties regarding the status of the collision configurations. Bourdeau & Sorkin (1992) show that, in the case of fermions, Pauli exclusion renders the inclusion or excision of collisions from the joint configuration space moot; but that in the case of bosons or anyons, the collision points should be retained, and a unique self-adjoint Hamiltonian can be defined on the corresponding space with a boundary. However, these conclusions are confined to $d = 2$ and no internal degrees of freedom. In the even more specific context of Bohmian mechanics, Brown *et al* (1999) observe that, if $d > 2$, coincident bosonic corpuscles will, due to the symmetrisation of the joint state, remain coincident for all

[20] For all the gory details, consult Landsman (2013). The pioneers were Laidlaw & DeWitt (1971) and Leinaas & Myrheim (1977). A popular account of the connection between inequivalent quantisations and the fundamental group considers quantisations of the (simply connected) *covering space* $\tilde{\mathfrak{Q}}$ of $\mathfrak{Q}$, such that $\mathfrak{Q} \cong \tilde{\mathfrak{Q}}/\pi_1(\mathfrak{Q})$, subject to the constraint that one obtains a "sensible" quantum theory on $\mathfrak{Q}$. See Morandi (1992, Ch. 3) for more details.

time.

6.4 Individuation in permutation-invariant quantum mechanics

If we follow route 1 (implementing permutation invariance on the joint Hilbert space), but not if we follow route 2 (quantising the anti-haecceitistic configuration space), then, in the case of bosons and fermions, the sectors $\mathcal{H}^N_\pm$ retain the resources of transcendental individuation of the particles. In particular, this allows us to define reduced density operators for each of the N system labels. For example, given any joint state $\rho \in \mathfrak{D}(\mathcal{H}^N_\pm)$ the reduced density operator ρ_1 associated with system 1, whose states correspond to the first entry of the tensor product $\mathcal{H}^N$, may be defined as follows:

$$\rho_1 := \mathrm{tr}_{2,3,\ldots,N}(\rho) \ , \tag{27}$$

where we perform a partial trace over all but the first factor Hilbert space. Similarly, the reduced density operator ρ_2 associated with system 2 is given by

$$\rho_2 := \mathrm{tr}_{1,3,\ldots,N}(\rho) \tag{28}$$

where we perform a partial trace over all but the second factor Hilbert space, and so on. However, the joint state is permutation invariant, i.e. $\pi^\star\rho = \rho$ for all $\pi \in S_N$, and so the reduced density operators are all identical: $\rho_1 = \rho_2 = \ldots = \rho_N$. Therefore, *if we continue to maintain transcendental individuation of the systems*, we are forced to conclude that all systems are described by the *same* density operator. It follows from this that the systems, transcendentally individuated, are indiscernible by means of monadic properties, i.e. they are absolutely indiscernible. This is the wide consensus in the quantum identity literature.[21]

This drastic result, analogous to the drastic result met in classical statistical mechanics (Section 5.3) is unavailable for paraparticle states, or on route 2, where system labels have been "rubbed out". A proposal for qualitative individuation, by appeal to approximate classical trajectories, is given by Dieks & Lubberdink (2011). I will outline here a more general proposal, in analogy to the classical statistical mechanical case.

Select N orthogonal projectors E_i on the single-system Hilbert space $\mathcal{H}$. These N projectors define a projector

$$\mathcal{E}(\{E_1, \ldots, E_N\}) := \sum_{\pi \in S_N} E_{\pi(1)} \otimes \ldots \otimes E_{\pi(N)} \tag{29}$$

on the joint Hilbert space $\mathcal{H}^N$ and, by restriction, any of the symmetry sectors and their quotients under S_N. If now the joint state ρ has support in the range of $\mathcal{E}(\{E_1, \ldots, E_N\})$; i.e. $\mathrm{Tr}(\rho\mathcal{E}(\{E_1, \ldots, E_N\})) = 1$, then we can, according to the proposal, say with certitude that:

- one of the particles has a state lying in the range of E_1;
- one of the particles has a state lying in the range of E_2;

  ⋮

- one of the particles has a state lying in the range of E_N.

[21]A representative sample: Margenau (1944), Post (1963), French & Redhead (1988), van Fraassen (1991, Ch. 11), Butterfield (1993), Saunders (2003a, 2003b, 2006a), French & Krause (2006, §4.2.1), Muller & Saunders (2008), Muller & Seevinck (2009), Ladyman & Bigaj (2010), Caulton (2012), Huggett & Norton (2012). A notable exception is Huggett & Imbo (2009).

Since the E_i are all orthogonal, they serve to individuate all N of the particles; we may them call them *individuation criteria*. We thereby avoid the conclusion above that each particle has the same state. (How can they? Their respective states lie in orthogonal sectors of $\mathcal{H}$.) There is no further question *which* particle's state lies in the range of *which* projector E_i, since we have no other means by which to individuate the particles except by appeal to their states.

Developing the proposal further, we may extract single-system reduced density operators as follows. For example, associated with the projector E_1 we can define the reduced density operator ρ_{E_1}, *via* Gleason's Theorem, by the measures it assigns to projectors F on $\mathcal{H}$, as follows:

$$\mathrm{tr}(\rho_{E_1}F) := \mathrm{Tr}(\rho\mathcal{E}(\{E_1FE_1, \ldots, E_N\})) \tag{30}$$

Thus, the particle individuated by E_1 has the reduced density operator ρ_{E_1}, and so on. And clearly, $\rho_{E_1}, \ldots, \rho_{E_N}$ are all distinct—in fact they have mutually orthogonal supports. All this relies, remember, on the joint state ρ having support in the range of $\mathcal{E}(\{E_1, \ldots, E_N\})$. This subspace has a natural tensor product structure, which we are essentially factorising into orthogonal subspaces $\mathrm{ran}E_i$ of the single-system Hilbert space $\mathcal{H}$.

The reduced density operators calculated through qualitative individuation bear a simple relation to the reduced density operators calculated through transcendental individuation. Let $\overline{\rho} := \rho_1 = \ldots = \rho_N$ be the single-system marginal defined in equation (27). Then, the reduced density operator for the particle individuated e.g. by the criterion E_1 is

$$\rho_{E_1} = \frac{E_1\overline{\rho}E_1}{\mathrm{tr}(\overline{\rho}E_1)} \tag{31}$$

which is just the Lüders rule conditionalisation of $\overline{\rho}$ on E_1 in the single-system Hilbert space.

Nothing so far has placed any sane restrictions on the sets σ_i; e.g., they may be horribly miscellaneous unions of disjoint regions in Γ). I will not give an account of such restrictions here, nor investigate how individuation criteria might be extended to qualitative individuation over time, or by appeal to irreducible relations holding between the particles (these are areas for future work). However, a reassuring result is as follows. Choose N distinct points ξ_i in Γ to define N *maximally specific* individuation criteria $\sigma_i = \{\xi_i\}$. Then $\Sigma(\{\sigma_1, \sigma_2, \ldots, \sigma_N\})$ is the singleton of the microstate $[(\xi_1, \ldots, \xi_N)] = \{(\xi_{\pi(1)}, \ldots, \xi_{\pi(N)}) : \pi \in S_N\}$. If the joint distribution has support on this singleton—i.e. it is a Dirac delta distribution centred on $[(\xi_1, \ldots, \xi_N)]$, which makes it an extremal state in the state space $\mathcal{M}(\Gamma_N/S_N)$—then each associated marginal μ_i is a Dirac delta distribution centred on ξ_i. That much is to be expected. However, also in this special case, the marginals μ_i uniquely determine the joint distribution μ, since there is only one state in $\mathcal{M}(\Gamma_N/S_N)$ which can produce those marginals *via* equation (17). In other words: maximally specific joint states determine, and are determined by, maximally specific states possessed by the constituent particles.

7 Conclusion: what can be settled experimentally?

In quantum mechanics, the decomposition of the joint Hilbert space into symmetry sectors, corresponding to the irreps of S_N, leads to experimentally significant predictions regarding systems in equilibrium. One important topic that I should like to conclude with is the variety of canonical ensembles derived for bosons, fermions and unpermutable particles, both compared to each other and to the classical case.

The set-up is this: we have a system of N non-interacting particles. Each particle's state space is decomposed into energy levels, assumed discrete and indexed by energy E_i. In classical mechanics, discreteness is achieved by coarse-graining the single-system phase space Γ into cells of equal Lebesgue measure, say ω. In quantum mechanics, discreteness is secured by assuming that the particles occupy bound states (typically, because they are confined to some finite volume). Each energy level has a degeneracy g_i: in classical mechanics, due to coarse-graining, this means that g_i cells in Γ have the (mean) energy E_i; in quantum mechanics, it means that the energy E_i eigenspace has g_i dimensions. Let n_i be the population of particles with energy E_i, these are *occupation numbers*.

As we saw in Section 5.2, the classical measure corresponding to the occupation numbers n_i is some multiple of ω^N. That multiple is either

$$W_{MB} = N! \prod_i \frac{g_i^{n_i}}{n_i!} \qquad \text{or} \qquad W_{MB}^{(PI)} = \prod_i \frac{g_i^{n_i}}{n_i!} \ , \tag{32}$$

depending on whether we are in the haecceitistic phase space Γ_N or its anti-haecceitistic quotient Γ_N/S_N. ('MB' stands for Maxwell-Boltzmann.)

In quantum mechanics, the Lebesgue measure is replaced by a count of dimensions in Hilbert space. If permutation-invariance does not hold (so that the full Hilbert space $\mathcal{H}^N$ is available to the joint system), then the number of dimensions corresponding to the occupation numbers being given by the n_i is again W_{MB}. However, if permutation invariance does hold, and we restrict to the bosonic or fermionic symmetry sector, then the dimension count for a bosonic or fermionic system are then given by

$$W_{BE} = \prod_i \frac{(g_i + n_i - 1)!}{n_i!(g_i - 1)!} \qquad \text{or} \qquad W_{FD} = \prod_i \frac{g_i!}{n_i!(g_i - n_i)!} \ . \tag{33}$$

('BE' stands for Bose-Einstein, i.e. bosons, and 'FD' stands for Fermi-Dirac, i.e. fermions.) I do not consider the various possible paraparticle ensembles.

The canonical ensembles for all four of our W measures can be derived by maximising the Boltzmann entropy $k \log W$, subject to conservation of energy and particle number (as already outlined in Section 5.2. The results are:

$$n_i^{MB} = g_i N e^{-\beta(E_i - \mu)} \ ; \qquad n_i^{BE} = \frac{g_i N}{e^{\beta(E_i-\mu)} - 1} \ ; \qquad n_i^{FD} = \frac{g_i N}{e^{\beta(E_i-\mu)} + 1} \ ; \tag{34}$$

where β and μ are constants. Identification of the Boltzmann entropy with the thermodynamic entropy leads to the identification $\beta = \frac{1}{kT}$, and μ is determined by conservation of particle number.

The three ensembles are distinct and experimentally distinguishable, at least if the g_i are not too much greater than the n_i. Indeed, the principle of Pauli exclusion, which is responsible for the Fermi-Dirac ensemble, is crucial to our current understanding of the stability of bulk matter.[22] However, in the dilute limit ($g_i >> n_i$), all three ensembles are approximately equal. The origin of the difference between these ensembles away from the dilute limit is not haecceitism or anti-haecceitism. Rather, the difference is due to the fact that the measures are in the classical case continuous (phase space volume) and in the quantum case discrete (dimension count)—a point made by Huggett (1999b) and Saunders (2006b, 2013).

As already mentioned in Section 5.2, in the classical theory, haecceitism and anti-haecceitism are not experimentally distinguishable by particle statistics: the Maxwell-Boltzmann ensemble is recovered in both cases. The Maxwell-Boltzmann ensemble is *also* obtained in the permutation-non-invariant quantum theory. In the permutation-invariant quantum theory, we obtain either

[22] See Dyson & Lenard (1967, 1968) and Lieb (1976, 1990).

the Bose-Einstein or Fermi-Dirac ensembles (or paraparticle ensembles, not considered here). These ensembles are experimentally distinguishable, so it is in principle experimentally determinable, by appeal to particle statistics, whether or not *permutation invariance* holds. (This is also true in the classical case: after all, the invariance of the total energy under arbitrary permutations is an essential assumption in the derivation of W_{MB}.)

So permutation invariance is experimentally determinable in quantum mechanics: does it follow that anti-haecceitism is likewise experimentally determinable? The issue is subtle. The haecceitist has an alternative explanation of why it is that the accessible joint (pure) states lie in some symmetry sector $\mathcal{H}^N_\lambda$ rather than the "full" joint Hilbert space $\mathcal{H}^N$. The explanation is simply that the Hamiltonian H is permutation-invariant (i.e. S_N-equivariant, $[H, U(\pi)] = 0$). Since this entails that all matrix elements connecting inequivalent irreps are zero, any joint state lying within a symmetry sector at some time remains there for all future (and past) times. Moreover, the permutation-invariance of the Hamiltonian is easily understandable under haecceitistism: the particles are permutable by dint of their being intrinsically indistinguishable (i.e. having the same mass, spin, charge, etc.). This explanation was first suggested by French & Redhead (1988).

This alternative explanation extends even to paraparticles, for which $\mathcal{H}^N_\lambda$ and its quotient $\mathcal{H}^N_\lambda/S_N$ are not isomorphic. Certainly, the haecceitist's dimension count differs from the anti-haecceitist's by a factor d_λ, the dimension of each generalized ray, since distinct rays in the same i.i.s. represent, for the haecceitist but not the anti-haecceitist, distinct possibilities. This factor is the same for all paraparticle states of the same symmetry type, and so amounts to an irrelevant additive constant $k \log d_\lambda$ in the Boltzmann entropy.

However, the haecceitistic explanation can only be given if route 1 is taken (i.e. permutation invariance is implemented on the haecceitistic joint Hilbert space). On either route 2 (quantising the anti-haecceitistic configuration space) or route 3 (some constant particle number sector in quantum field theory), the representative resources simply *do not exist* to differentiate between states that are permutable in the haecceitistic joint Hilbert space. Therefore any advocate of haecceitism in the quantum regime must provide convincing reasons for taking this route, even though: (i) all routes lead, in the end, to equivalent theories if the number of spatial dimensions exceeds 2; (ii) route 1 excludes anyonic behaviour if the number of spatial dimensions is 2; and (iii) route 3 would appear to be our best explanation for the source of permutation invariance in quantum mechanics. It seems implausible, to say the least, that such convincing reasons could be found.

So to summarise: in both classical and quantum theories, although it is experimentally determinable whether or not permutation-invariance holds—at least for the Hamiltonian—it is not straightforwardly experimentally determinable whether haecceitism or anti-haecceitism is true *if* permutation invariance holds. Broadly speaking, this is because the haecceitist and anti-haecceitist can both give explanations for permutation invariance: under haecceitism, permutation-invariance is a result of the physics not caring which particle is which; under anti-haecceitism, it is a result of their being no question of which is which. However, in the quantum case, the haecceitistic explanation requires one to favour a particular formulation of many-particle quantum mechanics that we have no independent reason to favour and a number of good reasons to disfavour.

8 References

- Bach, A. (1997), *Indistinguishable Classical Particles.* Berlin: Springer.

- Baker, D., Halvorson, H. & Swanson, N. (2015), 'The conventionality of parastatistics', *British Journal for the Philosophy of Science* **66**, pp. 929-976.
- Bigaj, T., (2015), 'On Discernibility and Symmetries', *Erkenntnis* **80**, pp. 117-135.
- Bourdeau, M. and Sorkin, R.D. (1992), 'When Can Identical Particles Collide?', *Physical Review D* **45**, pp. 687-696.
- Brown, H., Sjöqvist, E. and Bacciagaluppi, G. (1999), 'Remarks on Identical Particles in de Broglie-Bohm Theory', *Physics Letters A* **251**, pp. 229-235.
- Butterfield, J. (1993), 'Interpretation and Identity in Quantum Theory', *Studies in History and Philosophy of Science* **24**, pp. 443-476.
- Button, T. & Walsh, S. (2018), *Philosophy and Model Theory.* Oxford: Oxford University Press.
- Dieks, D. (1990), 'Quantum Statistics, Identical Particles and Correlations', *Synthese* **82**, pp. 127-155.
- Dieks and Lubberdink (2011), 'How Classical Particles Emerge From the Quantum World', *Foundations of Physics* **41**, pp. 1051-1064.
- Dixon & Mortimer (1996), *Permutation Groups.* Berlin: Springer.
- Doplicher, S., Haag, R. & Roberts, J. E. (1974), 'Local Observables and Particle Statistics II', *Communications in Mathematical Physics* **35**, pp. 49-85.
- Dürr, D., Goldstein, S., Taylor, J., Tumulka, R. & Zanghi, N. (2007), 'Quantum mechanics in multiply-connected spaces', *Journal of Physics A* **40**, pp. 2997-3031.
- Dyson, F. J. & Lenard, A. (1967), 'Stability of matter I', *Journal of Mathematical Physics* **8**, pp. 423-434.
- Dyson, F. J. & Lenard, A. (1968), 'Stability of matter II', *Journal of Mathematical Physics*, **9**, pp. 1538-1545.
- Earman, J. (2015), 'Some Puzzles and Unresolved Issues about Quantum Entanglement', *Erkenntnis* **80**, pp. 303-337.
- Ehrenfest, P. & Trkal, V. (1921), 'Deduction of the dissociation equilibrium from the theory of quanta and a calculation of the chemical constant based on this', *Proceedings of the Amsterdam Academy* **23**, pp. 162-183.
- van Fraassen, B. C. (1991), *Quantum Mechanics: An Empiricist View.* Oxford: Clarendon Press.
- French, S. (1989), 'Identity and individuality in classical and quantum physics', *Australasian Journal of Philosophy* **67**, pp. 432-446.
- French, S. (1995), 'The esperable uberty of quantum chromodynamics', *Studies in History and Philosophy of Modern Physics* **26**, pp. 87-105.
- French, S. and Krause, D. (1995), 'Vague identity and quantum non-individuality', *Analysis* **55**: 20-26.
- French, S. and Krause, D. (2006), *Identity in Physics.* Oxford: Oxford University Press.
- French, S. and Redhead, M. (1988), 'Quantum physics and the identity of indiscernibles', *British Journal for the Philosophy of Science* **39**, pp. 233-46.
- French, S., & Rickles, D. (2003), 'Understanding Permutation Symmetry', in K. Brading & E. Castellani (eds.), *Symmetries in Physics: Philosophical Reflections* (Cambridge: Cambridge University Press), pp. 212-38.

- Frigg, R. (2008), 'A field guide to recent work on the foundations of statistical mechanics', in Rickles (ed.), *The Ashgate Companion to Contemporary Philosophy of Physics* (London: Routledge), pp. 105-202.

- Greenberg, O. W., Dell'Antonio, G. F. & Sudarshan, E. C. G. (1964), 'Parastatistics: Axiomatic Formulations, Connection with Spin and TCP Theorem for a General Field Theory', in F. Gursey (ed.), *Group Theoretical Concepts and Methods in Elementary Particle Physics* (London: Gordon and Breach).

- Greiner, W. & Müller, B. (1994), *Quantum Mechanics: Symmetries*, Second Edition. Berlin: Springer.

- Hansson, T. H., Leinaas, J. M., & Myrheim, J. (1992), 'Dimensional reduction in anyon systems', *Nuclear Physics B* **384** (3), pp. 559-580.

- Hawley, K. (2006), 'Weak Discernibility', *Analysis* **66**, pp. 300-3.

- Hawley, K. (2009), 'Identity and Indiscernibility', *Mind* **118**, pp. 102-19.

- Hilbert, D. and Bernays, P. (1934), *Grundlagen der Mathematik, Volume 1.* Berlin: Springer

- Hodges, W. (1993), *Model Theory.* Cambridge: Cambridge University Press.

- Huggett, N. (1999a), 'Atomic Metaphysics', J. Phil., 96, 5-24.

- Huggett, N. (1999b), 'On The Significance of the Permutation Symmetry', *British Journal for the Philosophy of Science* **50**, pp. 325-347.

- Huggett, N. & Imbo, T. D. (2009), 'Indistinguishability', in D. Greenberger, K. Hentschel & F. Weinert (eds.), *Compendium of Quantum Physics* (Berlin: Springer), pp. 311-317.

- Huggett, N. & Norton, J.

- Imbo, T. D., Shah Imbo, C. and Sudarshan, E. C. G. (1990), 'Identical Particles, Exotic Statistics and Braid Groups', *Physics Letters B* **234**, pp. 103-107.

- van Kampen, N. (1984), 'The Gibbs paradox', in W.E. Parry, (ed.), *Essays in Theoretical Physics in Honourof Dirk ter Haar* (Oxford: Pergamon Press).

- Kaplan, D. (1975), 'How to Russell a Frege-Church', *Journal of Philosophy* **72**, pp. 716-729.

- Kassel, C. & Turaev, V. (2008), *Braid Groups: Graduate Texts in Mathematics 247.* Berlin: Springer.

- Ketland, J. (2011), 'Identity and indiscernibility', *Review of Symbolic Logic* **4**, pp. 171-185.

- Krause, D. and French, S. (2007): 'Quantum sortal predicates', *Synthese* **143** (14) pp. 417-430.

- Ladyman, J. (2007), 'On the identity and diversity of objects in a structure', *Proceedings of the Aristotelian Society Supplementary Volume* **81**, pp. 23-43.

- Ladyman, J., & Bigaj, T. (2010), 'The principle of the identity of indiscernibles and quantum mechanics', *Philosophy of Science* **77**, pp. 117-136.

- Ladyman, J., Linnebo, Ø. & Pettigrew, R. (2012), 'Identity and discernibility in philosophy and logic', *Review of Symbolic Logic* **5**, pp. 162-186.

- Laidlaw, M. G. G. & DeWitt, C. M. (1971), 'Feynman Functional Integrals for Systems of Indistinguishable Particles', *Physical Review D* **3**, pp. 1375-1378.

- Landsman, N. P. (2016), 'Quantization and superselection sectors III: Multiply connected spaces and indistinguishable particles', *Reviews of Mathematical Physics* **28**, 1650019.

- Leinaas, J. M. & Myrheim, J. (1977), 'On the Theory of Identical Particles', *Il Nuovo Cimento B* **37**, pp. 1-23.
- Lewis, D. K. (1986), *On the Plurality of Worlds.* Oxford: Blackwell.
- Lieb, E. H. (1976), 'The stability of matter', *Reviews of Modern Physics* **48**, 553-569.
- Lieb, E. H. (1990), 'The stability of matter: from atoms to stars', *Bulletin of the American Mathematical Society* **22**, pp. 1-49.
- Linnebo, Ø. & Muller, F. A., (2013), 'On Witness-discernibility of Elementary Particles', *Erkenntnis* **78**, pp. 1133-1142.
- 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', *British Journal for the Philosophy of Science* **52**, pp. 303-30.
- Massimi, M. (2003): *Pauli's Exclusion Principle: the Origination and Validation of a Scientific Principle*, Cambridge: Cambridge University Press.
- Messiah, A. M. L., & Greenberg, O. W. (1964), 'Symmetrization Postulate and its Experimental Foundation', *Physical Review B* **136**, pp. 248-267.
- Morandi, G. (1992), *The Role of Topology in Classical and Quantum Physics.* Berlin: Springer.
- Muller, F. A. (2015), 'The Rise of Relationals', *Mind* **124**, pp. 201-237.
- 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.
- Ohnuki, Y. & Kamefuchi, S. (1982), *Quantum Field Theory and Parastatistics.* Berlin: Springer.
- Pooley, O. (2006), 'Points, particles, and structural realism', in D. Rickles, S. French & J. Saatsi (eds.), *The Structural Foundations of Quantum Gravity* (Oxford: Oxford University Press), pp. 83-120.
- Post, H. (1963), 'Individuality and physics', *The Listener* **70**, pp. 534-37.
- Quine, W. V. O. (1960): *Word and Object.* Cambridge: Harvard University Press.
- Quine, W. V. O. (1976): 'Grades of discriminability', *Journal of Philosophy* **73**, pp. 113-16.
- Quine, W. V. O. (1986), *Philosophy of Logic*, Second Edition. Cambridge, MA: Harvard University Press.
- Read, N. (2003), 'Non-Abelian braid statistics versus projective permutation statistics', *Journal of Mathematical Physics* **44**, pp. 558-563.
- Redhead, M. & Teller, P. (1991), 'Particles, Particle Labels, and Quanta: The Toll of Unacknowledged Metaphysics', *Foundations of Physics* **21**, pp. 43-62.
- Redhead, M. and Teller, P. (1992), 'Particle Labels and the Theory of Indistinguishable Particles in Quantum Mechanics', *British Journal for the Philosophy of Science* **43**, pp. 201-218.
- Sagan, B. E. (1991), *The Symmetric Group: Representations, Combinatorial Algorithms, and Symmetric Functions*, Second Edition. Berlin: Springer.

- Saunders, S. (2003a), 'Physics and Leibniz's Principles', in K. Brading and E. Castellani, (eds.), *Symmetries in Physics: Philosophical Reflections* (Cambridge: Cambridge University Press), pp. 289-307.
- 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. 151-173.
- Saunders, S. (2006a): 'Are quantum particles objects?', *Analysis* **66**, pp. 52-63.
- Saunders, S. (2006b), 'On the explanation for quantum statistics', *Studies in the History and Philosophy of Modern Physics*, **37**, pp.192-211.
- Saunders, S. (2013), 'Indistinguishability', in R. Batterman (ed.), *The Oxford Handbook in Philosophy of Physics* (Oxford: Oxford University Press), pp. 340-380.
- Souriau, J.-M. (1997), *Structure of Dynamical Systems: A Symplectic View of Physics* (trans. by C. H. Cushman-de Vries). Boston, MA: Birkhäuser.
- Sternberg, S. (1994), *Group Theory and Physics.* Cambridge: Cambridge University Press.
- Stolt, R. H. & Taylor, J. R. (1970), 'Correspondence Between the First- and Second-Quantized Theories of Paraparticles', *Nuclear Physics B* **19**, pp. l-19.
- Sudarshan, E. C. G., Imbo, T. D., & Govindarajan, T. R. (1988), 'Configuration Space Topology and Quantum Internal Symmetries', *Physics Letters B* **213** (4), pp. 471-476.
- Swendsen, R. H. (2002), 'Statistical mechanics of classical systems of distinguishable particles', *Journal of Statistical Physics* **107**, pp. 1143-1165.
- Tung, W.-K. (1985), *Group Theory in Physics.* Singapore: World Scientific.
- Uffink, J. (2006), 'Compendium of the foundations of classical statistical physics', in J. Butterfield and J. Earman (eds.), *Handbook of the Philosophy of Science: Philosophy of Physics* (Amsterdam: North Holland).
- Wilczek, F. (1990), *Fractional Statistics and Anyon Superconductivity.* Singapore: World Scientific.
- Willard, S. (1970), *General Topology.* Reading, MA: Addison-Wesley.
- Yale, P. B. (1966), 'Automorphisms of the Complex Numbers', *Mathematics Magazine* **39**, pp. 135-141.