Classical Particle Indistinguishability, Precisely

James Wills[*1]

[1]Philosophy, Logic and Scientific Method, London School of
Economics

Abstract

I present a new perspective on the meaning of indistinguishability
of classical particles. This leads to a solution to the problem in statisti-
cal mechanics of justifying the inclusion of a factor $N!$ in a probability
distribution over the phase space of N indistinguishable classical par-
ticles.

1   Introduction

Considerations of the identity of objects have long been part of philosophical
discussion in the natural sciences and logic. These considerations became
particularly pertinent throughout the twentieth century with the develop-
ment of quantum physics, widely recognized as having interesting and far-
reaching implications concerning the identity, individuality, indiscernibility,
indistinguishability[1] of the elementary components of our ontology. This
discussion continues in both the physics and philosophy literature[2].

---

[*]j.j.wills@lse.ac.uk

1. This menagerie of terms in the literature is apt to cause confusion, especially as there
is no clear consensus on what exactly each of these terms mean. Throughout the rest of the
paper, I will be concerned with giving a precise definition of a concept which I will label
'indistinguishability', and which, I think, matches with how most other commentators
use the word. But I believe it to be a distinct debate whether particles are 'identical',
'individual', or 'indiscernible'. The literature I address in this paper therefore does not,
for example, overlap with the debate over the status of Leibniz's PII.

2. See French (2000) for an introduction to this discussion and a comprehensive list of
references.

Such considerations in classical physics have, by comparison, been largely ignored. This trend has rightly been resisted in the philosophical and foundational literature, but discussion of the idea has remained small in comparison to the (perhaps more glamorous?) discussion in quantum theory. It is strange that this is the case, since considerations of particle indistinguishability in physics actually first made their debut in Gibbs' work on classical statistical mechanics (Darrigol 1991, p. 239). Gibbs used the notion of particle indistinguishability to argue for the introduction of a factor $N!$ in the probability distribution over the phase space for a system of N 'entirely similar' particles. It is here where the problem of giving a rigorous justification for the $N!$ factor originates.

The question of whether, and if so in what precise sense, classical particles are indistinguishable has consequences in physics and its philosophy: arguments which aim to justify the inclusion of the $N!$ factor generally assume that it is needed to derive an extensive entropy function. The basic worry is that a non-extensive entropy function yields a non-zero entropy of mixing of indistinguishable samples of gas, contradicting the prediction of thermodynamics that it is zero. This contradiction is one version of Gibbs' paradox[3], and seeking a justification for the $N!$ factor is often taken to be synonymous with seeking a solution to the paradox. The argument I will present follows the trend in the philosophical literature by taking seriously the idea of indistinguishable classical particles, but does not take as its basic motivation the derivation of an extensive entropy function. In fact, I think it is a virtue of my analysis that it shows that a rigorous derivation of the $N!$ factor can be obtained on a very general level, independent of any particular probability measure and independent of what we use the probability measure to compute. This shows that the $N!$ problem is not a problem specific to the entropy, but rather a deep puzzle in the foundations of statistical mechanics.

However, Gibbs' treatment of indistinguishable particles is beset by ambiguous and imprecisely defined terms (what exactly does it mean for classical particles to be 'indistinguishable'?). This lack of clarity has resulted in a surprisingly diverse collection of justifications for the $N!$ factor. Some prominent textbook authors like Huang (1963), ter Haar (1966), and Blundell and Blundell (2010) appeal to quantum mechanics, arguing that the insertion of the $N!$ in classical physics was necessary but not justifiable until quantum

---

3. See van Kampen (1984), Denbigh and Redhead (1989), Jaynes (1992), Dieks (2011, 2013), Saunders (2013), van Lith (2018), and Swendsen (2018) for a sample of the debate from the perspective of seeking an extensive entropy function.

mechanics came along with its associated 'fundamental indistinguishability of identical particles'. Indeed, much of the modern discussion and teaching of the physics of indistinguishable particles occurs in the context of quantum mechanics. In contrast, philosophers and those working in the foundations of physics have supplied arguments which defend the coherence of the idea of the indistinguishable classical particle. For example, Saunders (2013, 2018) supplies a justification of the $N!$ factor on the basis of a notion of classical particle indistinguishability as defined by the intrinsic properties of classical particles like mass and charge. On the other hand, Dieks (2013) argues that the 'distinguishability' of classical particles means that the $N!$ should not be included after all.

In this paper, I present an analysis of classical particle indistinguishability as 'observational indistinguishability' in a certain mathematically precise sense. I will argue that this leads to a number of interesting and welcome consequences in the foundations of statistical mechanics:

1. The identification and resolution of shortcomings in the ongoing debate concerning the solution to the $N!$ problem.

2. A reinterpretation of the quotienting procedure typically used to justify the $N!$ term and a rigorous derivation of the $N!$ factor which does not appeal to the metaphysics of particles and which rather draws only on facts about observables.

3. A reinterpretation and reconstruction of Gibbs' own argument as a special case of my analysis in which particles are observationally indistinguishable with respect to the Hamiltonian. I call this 'dynamical indistinguishability'.

In Section 2, I present the necessary formalism and concepts of classical statistical mechanics in order to state a precise formulation of the $N!$ problem. In Section 3, I critically evaluate two prominent approaches to the $N!$ factor debate. In Section 4, I demonstrate that considerations of observability are crucial to understanding the sense in which states may be said to be 'equivalent' for the purpose of gathering statistics, and I offer a new perspective on the nature of reduction via quotienting in classical mechanics. In Section 5, I offer a reconstruction of Gibbs' argument for which I develop a precise definition of classical particle indistinguishability in terms of the symmetries of the dynamics. In Section 6, I use this definition to finally derive the $N!$ factor.

2　Gibbs' Statistical Mechanics

In classical statistical mechanics[4], the main object of study is the *Hamiltonian system*. This is a triple $\langle M, \omega, h \rangle$ where M, a smooth manifold representing the possible states of the system, together with ω, the symplectic form supplying M with geometric structure, form the *phase space*. Given a *Hamiltonian*, $h : M \to \mathbb{R}$, the symplectic form ω determines a unique dynamical trajectory through each point in phase space. The dynamical trajectories are the integral curves of the vector field H generated by h via the symplectic form thus:

$$w_{ab}H^a = d_b h \tag{1}$$

where $d_b h$ is the exterior derivative of h. We interpret the different integral curves with various initial conditions as representing the possible dynamical trajectories the system may follow from the various initial states represented by those points. Dynamical trajectories may also be viewed as maps $\varphi_t : M \to M$ which take points in M and output points in M, t units along the integral curve from the initial point. The collection of these maps $\{\varphi_t\}_{t \in \mathbb{R}}$ forms a one-parameter group of diffeomorphisms and is called the *phase flow*.

When the points in the phase space for a system of n degrees of freedom are given local coordinates $(q_1, ..., q_n, p_1, ..., p_n)$, where $q_1, ..., q_n$ are the generalized configuration variables and the $p_1, ..., p_n$ are the generalized momenta, Equation (1) takes on the more familiar form of Hamilton's equations:

$$\dot{q}_i = \frac{\partial h}{\partial p_i}, \; \dot{p}_i = -\frac{\partial h}{\partial q_i} \text{ for } 1 \leq i \leq n. \tag{2}$$

Now we introduce the statistical part. In introducing his statistical mechanical framework, Gibbs invites us to

> imagine a great number of systems of the same nature, but differing in the configurations and velocities which they have at a given instant. (Gibbs 1902, p. iii)

---

4. I choose to present the Gibbs rather than the Boltzmann version of statistical mechanics. This is partly because the problem I am concerned with first originated with Gibbs. Despite this, many authors in the modern philosophical literature (like Saunders (2013) and Dieks (2013) which I discuss below) choose to state the problem in the Boltzmann framework. Happily though, this does not affect the solution of the problem since, ultimately, all the action occurs in the probability measure over the phase space which is a feature common to both Gibbs and Boltzmann versions.

This collection of systems is known as the *ensemble*: an enormous collection of copies of the system under study which contains all of the possible ways the actual system could be at an instant. It is important to recognize that the ensemble is not (for example) the collection of particles in the actual box of gas itself; it consists of possible, not actual, copies of the entire system. And it is the ensemble, not the actual system, that is the object of study in Gibbs' statistical mechanics:

> And here we may set the problem, not to follow a particular system through its succession of configurations, but to determine how the whole number of systems will be distributed among the various conceivable configurations and velocities at any required time, when the distribution has been given for some one time. The fundamental equation for this inquiry is that which gives the rate of change of the number of systems which fall within any infinitesimal limits of configuration and velocity. (Gibbs 1902, p. iii)

He is not concerned with the traditional mechanical problem of how a particular state of the system changes in time but rather with the problem of how the distribution of a number of systems over the possible states changes with time.

We are invited to imagine that the ensemble of systems is distributed in a particular way among the possible mechanical states in phase space. This distribution D on phase space tells us how many systems N in an ensemble are associated with a given region R of the phase space, through the following definition:

$$N(R) := \int_R D(p,q)\, dp_1...dq_n. \tag{3}$$

Gibbs calls $D(p,q)$ the *density-in-phase* and $dp_1...dq_n$ an *element of extension-in-phase*[5].

The total number of systems in the ensemble $N := N(M)$ is given by Equation (3) for $R = M$. We can normalize D by defining $P(p,q) := D(p,q)/N$, which Gibbs calls the probability coefficient[6]. We can then define

$$\mu(R) := \int_R P(p,q)\, dp_1...dq_n. \tag{4}$$

---

5. In modern language, this is known as the 'volume form' or 'Liouville form' associated with the symplectic manifold (M, ω), given by $\omega^n = dq_1 \wedge dp_1 \wedge \cdots dq_n \wedge dp_n$.
6. In modern language, this is a 'probability density function'.

This defines the measure of a region $R \subseteq M$. We then define

$$\mathcal{P}(R) := \frac{\mu(R)}{\mu(M)} \qquad (5)$$

to be the probability of region R. Its intended interpretation is as the probability that a system picked out of the ensemble is found in region R of the phase space. In Gibbs' words, it expresses

> [the] ratio of the number of systems falling within those limits to the whole number of systems. This is the same as the probability that an unspecified system of the ensemble [...] will lie within the given limits. (Gibbs 1902, p. 17)

This quotation suggests that the talk of ensembles of great numbers of systems is just a concrete way of interpreting the probability measure in a frequentist manner. However, in using Gibbs' formalism, we are not committed to the existence of this ensemble of systems[7]. The important message is that the central object of study in Gibbsian statistical mechanics is the probability measure over the phase space. Physics proceeds by computing expectation values for certain observables which are then compared with experiment.

Much of Gibbs' 1902 book studied the probability coefficients of arbitrary Hamiltonian systems of n degrees of freedom interacting with different environments. For example, the first probability coefficient he studied was the canonical distribution. He interpreted this to model systems which could exchange energy with their constant-temperature environment. He then introduced the micro-canonical distribution to study the properties of systems with a fixed energy. If a physicist wants to study the properties of their particular system when it is in contact with a heat bath (for example), all they have to do is take the canonical distribution, which Gibbs has already worked out and apply it to their particular system, effectively endowing the canonical coordinates with certain interpretations. For example, for a single point particle with three degrees of freedom, the generalized configuration variables represent the particle's x, y and z position coordinates and the generalized momenta represent the particle's linear momentum along the x, y and z axes. It is at this stage where the problem we are concerned with arises.

---

7. Mathematically, all we have is a probability distribution over M. How we are to interpret this probability is a matter of debate. See Myrvold (2017) for a discussion of the main issues.

While Gibbs spent most of his book studying arbitrary systems of n degrees of freedom, in the last chapter of his book he did apply his methods to a particular system: a collection of N point particles[8]. In applying his methods to this system, he expressed the following problem:

> If two phases differ only in that certain entirely similar particles have changed places with one another, are they to be regarded as identical or different phases? (Gibbs 1902, p. 187)

This quote requires some interpretation. The $6N$-dimensional phase space of a system of N point particles has phase points (also known as phases) s of the form of a $6N$ tuple:

$$s = (r_{1x}, r_{1y}, r_{1z}, p_{1x}, p_{1y}, p_{1x}, ..., r_{Nx}, r_{Ny}, r_{Nz}, p_{Nx}, p_{Ny}, p_{Nz}) \qquad (6)$$

Or, more briefly

$$s = ((\mathbf{r}, \mathbf{p})_1, ..., (\mathbf{r}, \mathbf{p})_N) \qquad (7)$$

where $\mathbf{r} = (r_x, r_y, r_z)$ and $\mathbf{p} = (p_x, p_y, p_z)$. We are to think of this tuple as consisting of N 'slots', known as *factor positions*, each of which represents one of the N particles, and into which go three position coordinates and three momentum coordinates specifying the single particle state. This phase point represents the state of the N particle system, made up of N single particle states. Now consider the family of transformations $\{\pi_{ij}\}$ for $1 \leq i, j \leq N$ whose elements swap the blocks of six canonical coordinates, which is to say they swap the single-particle states, among the factor positions i and j. For example, the transformation $\pi_{1N} : M \to M$ is defined so that

$$\pi_{1N}(s) := ((\mathbf{r}, \mathbf{p})_N, ..., (\mathbf{r}, \mathbf{p})_1). \qquad (8)$$

Let us refer to these transformations as *permutations*. The permutations π_{ij} effectively change which particle has which position and momentum. It is easily checked that these permutations are canonical transformations, a fact which will be useful to us later.

---

8. This last chapter is known for the grand canonical distribution, used in modelling systems of particles in which the particle number is allowed to vary (in chemical reactions, for example). Indeed, it is in the mathematical expression of the grand canonical distribution where the problem I discuss first appeared. However, the discussion and solution to the problem will not depend on any particular probability distribution; the problem would appear if we applied the canonical, micro-canonical or another distribution to the N particle system.

The permutation π_{ij} is sometimes described as a 'permutation of particle labels' (French and Rickles 2003) or as a 'permutation of particles among the states' (Dieks 1990; Saunders 2013). These images are tempting, but they can be misleading as to what the permutations $\{\pi_{ij}\}$ really do. Firstly, we cannot think of particles as being permuted among states since this invites the image of phase points remaining fixed, like boxes, and particles being permuted between them. In fact, the formalism indicates that we should think of the particles like boxes (they are represented by the factor positions) as remaining fixed while the states (blocks of six coordinates) are permuted among them. Nor can the maps be interpreted passively as a permutation of particle labels since, strictly speaking, there are no labels for the map to act on; all that the N tuple consists in is factor positions and single particle states. It is possible to interpret the labels as names referring to the factor positions (the label 'A' would just be shorthand for 'the first factor position in the tuple' and so on), and if we must talk of labels then this is how they should be interpreted, but this still does not capture what the maps do since, again, the factor positions remain fixed. As defined, these maps act on N tuples and return N tuples by permuting single particle states among the factor positions. The π_{ij} do not permute particle labels over particles, or particles over states: they permute states over particles[9].

Now Gibbs poses the following question: what if the particles represented by factor positions i and j are 'entirely similar' or 'indistinguishable'? The transformation π_{ij} certainly takes s to a different phase point $\pi_{ij}(s) \in M$, $s \neq \pi_{ij}(s)$, but are the states represented by $\pi_{ij}(s)$ and s, in some sense, equivalent? It is tempting to say in this case that if the particles are 'indistinguishable', it does not 'matter' which particle has which single particle state and so the states represented by the phase point s and its permuted counterpart $\pi_{ij}(s)$ are in some sense equivalent. This is a common interpretation of Gibbs' reasoning here, for he writes:

> If the particles are regarded as indistinguishable, it seems in accordance with the spirit of the statistical method to regard the phases as identical. (Gibbs 1902, p. 187)

If all N particles were 'indistinguishable' from each other, then it seems that

---

9. The distinction between permuting particles over states and permuting states over particles has made made in the context of quantum theory: French and Krause (2006, p. 136) note that we may think about permutation both ways in the context of ordinary statistics but that the distinction becomes important in the context of parastatistics. In this paper, however, we restrict ourselves strictly to classical Hamiltonian mechanics for a contained discussion.

the phase space is over-representing the physical states; each physical state is represented by $N!$ points in the phase space since N single particle states over N factor positions admit of $N!$ distinct permutations of single particle states over particles. This has a knock-on effect on the probability measure: if the physical states are over-represented in M, the standard probability measure over M developed in the earlier part of Gibbs' book for arbitrary systems would be over-counting the states. Gibbs corrects for this[10] by dividing the probability measure by $N!$. More precisely, if μ is the measure derived from the general reasoning of earlier chapters (see Equation (4)), the measure for the N particle system would be $\mu/N!$ to correct for the over-counting of the states by μ.

The terms Gibbs uses are open for interpretation: what is meant by 'entirely similar' or 'indistinguishable'? What does it mean to say that the phases are 'identical'? Is this to be read as logical identity? Or some other notion of equivalence? These ambiguities have led to a number of views regarding the $N!$ factor. In the next Section, I critically evaluate two prominent views, taking certain parts of their arguments as inspiration for my own derivation of the factor and for my answers to these questions.

3 $N!$ Justifications

In this Section, I turn to discussing and critically evaluating the arguments of Dieks and Saunders, two prominent representatives of different views regarding the $N!$ factor. This will help set the scene for the alternative view I will defend in this paper.

3.1 Dieks' Classical Particle Distinguishability

Dennis Dieks gives an argument against the inclusion of the $N!$ factor, based on the idea that classical particles cannot be indistinguishable at all because they are always distinguishable by their trajectories in configuration space. The thrust of his argument is expressed here:

> [C]lassical particles can be named and distinguished by their
> different histories. A process in which two classical particles of

---

10. In the original text, Gibbs considers the more general case in which the gas consists of numbers of particles of different kinds so his factor is $N_1!N_2!...N_h!$ where $1,...,h$ indexes the different kinds of particles in the gas such that particles of kind 1 are 'indistinguishable' from each other but 'distinguishable' from any other kind. We work with the simpler case of N particles of the same kind for simplicity and without loss of generality.

the same kind are interchanged can therefore certainly produce a different microstate. Indeed, imagine a situation in which there is one particle at position $[r_1]$ and one particle at position $[r_2]$, and in which at a later instant there is again one particle at $[r_1]$ and one at $[r_2]$; suppose that their respective momenta are the same as before. What has happened in the meantime? There are two possibilities: either the particle that was first at $[r_1]$ is later again at $[r_1]$ and the particle that was first at $[r_2]$ is later again at $[r_2]$, or the particles have exchanged their positions. The latter case would clearly be different from the former one: it corresponds to a different physical process. Although it is true that the two final situations cannot be distinguished on the basis of their instantaneous properties, their different histories show that the particle at $[r_1]$ in one final situation is not the same as the particle at $[r_1]$ in the other final situation. (Dieks 2013, p. 27)

The situation Dieks has in mind here is depicted in Figure 1. Let the initial

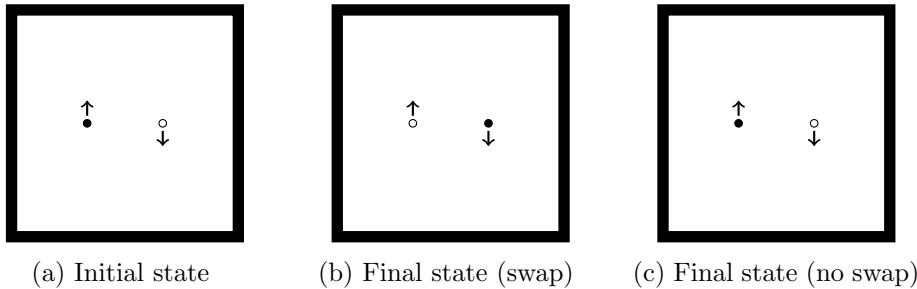(a) Initial state   (b) Final state (swap)   (c) Final state (no swap)

Figure 1: (a) is the initial state of a system consisting of two particles in a container. The arrows indicate their velocities. States (b) and (c) are the final states of particles after they follow trajectories which lead them to swapping and not swapping, respectively. The black dot depicts the particle represented by the first factor position and the white dot depicts the particle represented by the second factor position.

state (a) be represented by the phase point $s = (r_1, p_1, r_2, p_2)$. And let state (b) be the final state of the two particles if they swap, taking the system to the state $\pi_{12}(s) = (r_2, p_2, r_1, p_1)$, where $\pi_{12} : M \to M$ is the permutation, acting on the single particle states in the factor positions 1 and 2. This transformation represents the process in which the particle (represented by the first factor position, coloured black in the diagram)

which initially had (r_1, p_1) now has (r_2, p_2) and the particle (represented by the second factor position, coloured white in the diagram) which initially had (r_2, p_2) now has (r_1, p_1). The particles have swapped single particle states. State (c) shows the final state of the two particles if they do not swap, the system returning to the state s. The puzzle here is that if the particles are 'of the same kind', 'entirely similar', or 'indistinguishable' then the states $\pi_{12}(s)$ and s are in a sense equivalent, since in each there is one particle at position r_1 with momentum p_1 and another at position r_2 with momentum p_2; physically, it might not seem to matter which particle has which position and momentum. Dieks argues that although the states are not distinguishable via their 'instantaneous properties', the particles in Figure 1b followed different trajectories in configuration space to those in 1c. So, when we take the histories into account, it does seem to 'matter' which particle is which. In this sense, the particles are clearly distinguishable via their trajectories and therefore the states $\pi_{12}(s)$ and s are, in another sense, not equivalent. Which is the correct sense of equivalence to take? Dieks argues that the histories are relevant to considerations of the equivalence of the states and so he concludes that the states are distinct.

In order to evaluate the argument, let me put it in a more formal setting. Note first of all that the system in which the particles do not swap must have a different dynamics than the system in which the particles do swap. More formally, representing the dynamics as the phase flow after some fixed time $\varphi_t : M \to M$, then it must be the case that:

$$\varphi_{\neg\text{swap}} : s \mapsto s \tag{9}$$

$$\varphi_{\text{swap}} : s \mapsto \pi_{12}(s) \tag{10}$$

In other words, if the system is to evolve under the dynamics to some phase point other than s, then the phase flow must be different. Each phase flow is generated by the corresponding Hamiltonian h^{swap} and $h^{\neg\text{swap}}$. But the Hamiltonian is a function of the phase—indeed, it is the energy—and thus is exactly the sort of object that may fairly be called an 'instantaneous property' of the system. Therefore, the two final states in Figures 1b and 1c can in fact be distinguished on the basis of their instantaneous properties, since the Hamiltonian in each case is different. It is exactly because the states are different in virtue of their different Hamiltonians that the states have different trajectories in phase space.

Therefore, what Dieks has really told us is that the phase point s in the phase space of dynamical system $\langle M, \omega, h^{\neg\text{swap}} \rangle$ is a different state from

$\pi_{12}(s)$ in the phase space of dynamical system $\langle M, \omega, h^{\mathrm{swap}} \rangle$. However, this is not the question we are really trying to answer. What we want to know is whether s in the phase space of some dynamical system is in some sense equivalent to $\pi_{12}(s)$ in the phase space of *the very same dynamical system.* What are we to say about the sense in which the states represented by s and $\pi_{12}(s)$ are equivalent?

3.2 Saunders' Classical Particle Indistinguishability

Saunders (2013, 2018) supplies an answer to this question as part of an argument justifying the inclusion of the $N!$ factor in the context of Boltzmann statistical mechanics. If the particles are identical[11], understood to mean that the particles share all state-independent properties (like charge, mass and spin), then the states $\pi_{ij}(s)$ and s (for $1 \le i, j \le N$) represent one and the same physical state (the sense of equivalence used here being identity). This representational redundancy of the phase space then needs to be removed by passing to the reduced phase space, $M' = M/S_N$: the quotient of the phase space under the action of S_N, the permutation group[12]. Points in the reduced space may be understood as equivalence classes $[s]$ each containing $N!$ elements all related to the $s \in M$ by the permutations $\pi \in S_N$. Quotienting under this group effectively identifies points in the unreduced phase space related by these permutations. More informally, the unreduced phase space points are read as saying, 'the first particle has this position and momentum, the second particle has that position and momentum, and so on' while the points in the reduced phase space are read as saying 'some particle has this position and momentum and another particle has that position and momentum, and so on' While it is possible to make statements about which particle has which position and momentum in the unreduced space, it is not possible to say which is which in the reduced space.

This reduction has the desired effect on the probability measure. Let there be a measure μ on the unreduced space M such that $\mu(M) = 1$ (it is normalized) and let μ' be the measure on the reduced space M' obtained through the quotienting procedure. We may embed M' to M by mapping the points $[s] \in M'$ to some point $[s] \ni s_i \in M, 1 \le i \le N!$. This maps M' onto some region $R \subset M$. Carrying μ' over to M then allows allows us to

---

11. This term is not meant in the strict logical sense of 'one and the same', but it is widely used, so I adopt it. In the philosophical literature, particle identity has been treated with significant sophistication and precision. See, for example, Caulton and Butterfield (2011, 2012) and French (1989).

12. S_N contains $N!$ elements which are generated by the transformations π_{ij}.

compare μ' and μ on the same space. Since $\mu(R) = 1/N!$, and R can be mapped to the entire reduced space, it follows that $\mu'(M') = \mu(R) = 1/N!$ resulting in the desired $N!$ factor.

There are two key steps in this argument that I will focus on. These are:

1. Focusing on the key question of whether the points in M related by some $\pi \in S_N$ represent 'one and the same' physical state. The criterion for this was taken to be: the points s and $\pi_{ij}(s)$ represent one and the same physical state just in case the particles represented by factor positions i and j are identical.

2. If the states are equivalent in this sense, then the points s and $\pi_{ij}(s)$ are to be identified by quotienting M under the action of the permutation group.

I take issue with the first, and therefore am not convinced by the second. My critique is the subject of the next Section. In its place I will propose a new and mathematically precise perspective on the equivalence of states, which in turn gives rise to a novel interpretation of the quotienting procedure.

4   Observables and Statistics

If we assume that the job of the state space in a typical dynamical theory is to represent completely and without redundancy all the possible ways a system can be at an instant, then the question in step one of whether two states are 'one and the same' physical state is pertinent. The situation in statistical mechanics, however, is slightly different. While we are ostensibly studying a dynamical system, it is not the primary object of study. This is because statistical mechanics is a statistical theory and, as such, it is not immediately concerned with the evolution of the state of the system, but rather with studying a *probability space*. This is a triple $(\Omega, \mathcal{F}, \mu)$ where Ω is the sample space whose elements are called outcomes, $\mathcal{F}$ is a σ-algebra on Ω whose elements are called events and $\mu : \mathcal{F} \to [0,1]$ is a countably additive probability measure. This shift in focus is important: while the job of the state space M is to completely represent all the possible states of the dynamical system, the job of the subsets of Ω is to be the events to which probabilities are attached. In the previous section, it was tacitly assumed that Ω was M, but this will not, in general, be the case, as I will show shortly. It is important to realize that M and Ω are conceptually different entities with different roles vis-à-vis what they aim to represent. The importance of

noting this difference and its relevance to statistical mechanics is the topic of this section.

In a typical statistical analysis using a probability space, the job of Ω is to represent the set of possible outcomes or results of an experiment or observation while $\mathcal{F}$ represents the set of possible events which are assigned probabilities (c.f. Billingsley 1995, p. 17). I would like to draw attention to the seemingly innocuous role of observation in defining the sample space. The appearance of observation in the standard interpretation of a probability space is, I think, very telling and important. We can only collect statistical data and determine the relative frequencies of various elements of $\mathcal{F}$ when the outcomes in the sample space are observationally distinguishable from one another. Therefore, considerations of what counts as observable must be accounted for when specifying the sample space. To give a really simple example: we have two coins to be flipped at the same time, each of which can either land heads (H) or tails (T) and with one coloured red (R) and the other green (G). If the experimenter tossing the coins and recording relative frequency data is colourblind, it will only be possible for them to collect statistical information with the sample space $\{(H,H),(H,T),(T,T)\}$, even though the state space is in fact $\{(GH,RH),(GT,RH),(GH,RT),(GT,RT)\}$. That is, the sample space is constrained by the observables the experimenter has available. With respect to the observables available to the colourblind experimenter, the states $(GT,RH$ and (GH,RT) are observationally indistinguishable.

Given that the primary object of study in statistical mechanics is the probability space, we should be in the business of constructing the sample space. And given that observational considerations are inextricably bound up in the standard characterisation of a sample space, we must pay close attention to the observable in constructing the sample space for the N particle system. These considerations of what is observable will help to pin down the relevant notion of state equivalence to use in constructing our sample space: we should not be asking whether two points represent one and the same state but rather whether they are observationally equivalent. This is an important difference, since states may be observationally equivalent while not being one and the same state[13]. We therefore need a precise characterisation of observational equivalence of states in classical mechanics.

---

13. Note that, in constructing the sample space, we are not making any claims about which states are 'really' equivalent. We are in fact not modifying the phase space in any way at all, but rather building a new space, constrained by the observables, on which to subsequently define probabilities and do statistical mechanics. The argument presented here is, therefore, compatible with any empirically adequate phase space ontology.

Observable quantities in classical mechanics are represented by classical observables. These are smooth, real valued functions on the manifold $f : M \rightarrow \mathbb{R}$, assigning a real number to each state. This allows us to define precisely what it means for two states to be observationally equivalent.

Observational Equivalence. States s and s' are *observationally equivalent* with respect to a set of observables $\{f_i\}$ if and only if $f_i(s) = f_i(s')$ for all f_i.

This definition provides a schema for constructing the sample space. Simply identify your favourite set of smooth functions and build your sample space containing all and only states which are observationally distinct (inequivalent) with respect to that set.

We may consider a close cousin of Saunders' view as an instance of this schema in which we interpret his talk of particle properties in terms of observables. The next step is to interpret the claim that the particles have the same intrinsic properties in terms of classical observables. We cannot read the intrinsic properties of the particles, such as charge and mass, in the N particle system from the state since this only gives us position and momentum values. In fact, we read intrinsic properties of the particles from the classical observables. This is best illustrated with an example: if we want to say that the particles all have the same mass, we would formalize this in terms of the permutation symmetry of some classical observable, like the kinetic energy[14]. Thus, consider:

1. $K = (p_1^2 + p_2^2)/2m$;

2. $K' = p_1^2/2m + p_2^2/2M$, $m \neq M$.

The variables p_1 and p_2 pick out the momentum values in the factor positions 1 and 2 respectively. K is left invariant under the permutation π_{12} since $K(s) = K(\pi_{12}(s))$ for all s, while K' is not so invariant. We can say then that the particles represented by factor positions 1 and 2 have the same mass if and only if the kinetic energy observable is invariant under the permutation π_{12}. This is one interpretation of what it means for particles to have the same mass. Similarly, having the same value of other particle properties such as charge and spin will also be understood formally in terms of the permutation symmetry (or not) of some observable(s).

---

14. This will not be the only classical observable invariant under this transformation, but it is the obvious one to consider.

Thus, while Saunders' view says that the necessary condition for states s and $\pi_{ij}(s)$ to be identified in the quotienting procedure is that the particles share all state independent properties, the close cousin of Saunders' view would say that the necessary condition for the states to be identified in the quotienting procedure is that the states must be observationally equivalent with respect to all observables. This view is an instance of the schema with $\{f_i\}$ being the set of all smooth functions.

This instance of the schema has an undesirable consequence for Saunders with respect to step 2: it implies that no two distinct states in phase space are ever observationally equivalent:

Proposition. Let $s, s' \in M$, where M is a smooth connected manifold. If $f(s) = f(s')$ for all smooth functions $f : M \to \mathbb{R}$, then $s = s'$.

Proof. Suppose $f(s) = f(s')$ for all functions f, and let C be any curve on M containing s and s'. Now let f be any smooth function that increases strictly monotonically along C. Then since $f(s) = f(s')$ and s, s' are both on C, it follows that $s = s'$. $\qquad\square$

Therefore, if we require two states to be observationally equivalent with respect to all observables before identifying the states in the quotienting procedure, then quotienting will not actually do anything, since no two distinct states are observationally equivalent in this sense.

However, one could reasonably respond that to require observational equivalence with respect to all observables is an impossibly strong epistemic requirement. There is the inevitable experimental limitation that we are able to measure only a few observables. If we want a useful σ-algebra whose events we can actually observationally distinguish, then the sample space had better contain all and only the states observationally distinguishable with respect to some subset of the observables, such as those we can measure. This claim has two important implications.

The first implication affects step 2: the quotienting procedure. The goal of reduction is to remove representational redundancy in the phase space by identifying states (deemed physically identical) in the unreduced space as one state in the reduced space. That aim is sometimes viewed as metaphysical: it is to achieve a bijection between physically possible states of the system and points in the state space[15]. This, however, is not the goal of a statistical theory, such as statistical mechanics. All we require of our

---

15. See Castellani (2003, p. 434), Belot (2003, p. 410) and Butterfield (2007, pp. 23–4) for examples of this view. Examples of state space reduction typically occur in the relationist-substantivalist debate in spacetime physics.

sample space is that the elements of Ω be distinguishable with respect to some set of empirically accessible (measurable) observables. In fact, given our observational and experimental limitations, we know that the elements of the sample space may correspond to many genuinely distinct possibilities in the state space (recall that this was the lesson of the coin example). Since we inevitably do not have access to all the classical observables, this is our situation in statistical mechanics. So, if we wish to retain the ability to interpret phase space as representing distinct physical possibilities, then reducing the phase space is not the correct procedure to carry out, since this would identify states known to be genuinely distinct.

One way around this problem is to re-interpret the reduction procedure; instead of interpreting the reduced space as a state space in one-to-one correspondence with genuine possibilities, we could interpret it as a sample space, claiming to represent all and only the observationally distinct states without claiming to represent all and only the genuine physical possibilities. If interpreted this way, the reduction procedure can be carried out in a representationally accurate way. This interpretation is non-standard since reduction is typically used to eliminate the over-representation of possibilities. But nothing in the mathematics of reduction dictates any particular interpretation of the procedure; if we think of quotienting in terms of forming equivalence classes of states in M, we are at liberty to quotient with respect to any equivalence relation which suits our purposes. Typically, this relation is taken to be '...is identical to...' or '...is physically equivalent to...' but we may just as well take the relation to be '...is observationally equivalent to...' In the rest of the paper, however, I will adopt a position which does not require reduction under any interpretation. This position has a number of virtues and advantages over re-interpreting reduction which I explain in Section 6.

The second implication is more pressing because it also affects the possible ways out of the problem posed by the first implication. Since no distinct states are observationally equivalent with respect to all observables, which subset of observables will give us a fruitful definition of observational equivalence? A fruitful definition would be one which allows us to construct a sample space which is actually useful, that is, contains all and only the outcomes which are observationally distinct with respect to observables available to an experimenter. This set is clearly dependent on the experimental context and our epistemic situation and therefore there will be no 'one true' fruitful definition of observational equivalence.

In the next Section, I will consider one particular definition of observational equivalence resulting from choosing the subset of observables to be

the singleton set containing the Hamiltonian. I will show how this results in the $N!$ correction factor used by Gibbs and a precise definition of classical particle indistinguishability in terms of dynamical symmetries. But realize that this is only one possible way of filling in the observational equivalence schema: the division by $N!$ results from specifying a particular choice for $\{f_i\}$. Thus the '$N!$ problem' is a specific instance of the much more general 'over-counting problem'. Depending on the choice of observables, we will in general correct by a different factor.

5   Dynamical Equivalence of States and the Dynamical Indistinguishability of Particles

A hint that $\{f_i\} = \{h\}$ is the correct starting point for interpreting and reconstructing Gibbs is provided by Uffink. On introducing 'indistinguishable' particles, he makes the following comment (in a footnote) on what that might mean:

> Presumably, these terms mean (at least) that the Hamiltonian is invariant under their permutation, i.e. they have equal mass and interact in exactly the same way. (Uffink 2007, p. 997)

The hint that the indistinguishability of particles is grounded in some symmetry of the energy observable (the Hamiltonian) is very promising, but it certainly requires further argument[16]. The energy observable is but one of many smooth functions on phase space: why choose this one with respect to which to judge the (in)distinguishability of the particles and the observational equivalence of the states? In the following sections I develop in detail the line of thought inspired by this hint and I show that it results in a precise definition of classical particle indistinguishability, and a rigorous derivation of the $N!$ factor.

We shall call observational equivalence with respect to the Hamiltonian[17] h in which $\{f_i\} = \{h\}$, dynamical equivalence, to be defined more precisely below. But first, this view does require some motivation.

This position is motivated by the thought that a minimal, necessary condition for two states of a dynamical system to be observationally equivalent

---

16. This hint carries through straightforwardly to the discussion in quantum theory, in the form that the permutation operator commutes with the Hamiltonian. However, I restrict this paper to the discussion in classical Hamiltonian mechanics.

17. The Hamiltonian which is the third element of the dynamical system triple (M, ω, h).

is that they have identical dynamical evolutions. It should not be possible to count states as observationally equivalent for the purposes of constructing the sample space if they have different dynamical evolutions.

Is the condition sufficient? It is clearly not sufficient for observational equivalence, since taking $\{f_i\} = \{h\}$ allows the states to differ with respect to any number of other observables, if these differences do not make any difference to their dynamical evolutions. For example, the particles may all have different colours but if the dynamics is 'colourblind', then it should not matter if this particle is blue and that one red or this one red and that one blue. They do, nevertheless, differ with respect to the 'colour' observable. Recall, however, that our aim is to find the set of observables sufficient for our statistical aim of constructing the sample space and correct probability distribution for our Hamiltonian system.

An argument that $\{h\}$ is sufficient for this purpose comes from Chapter IV in Gibbs' book in which he arrives at the canonical distribution. He begins by singling out as worthy of study the distributions P which are stationary. That is, those which satisfy

$$\dot{P} = \{P, h\} = 0 \qquad (11)$$

where $\{P, h\}$ is the Poisson bracket of P and h. He notes that it is necessary and sufficient that P be a function only of the phase for this condition to be satisfied. We may further constrain P by requiring that they define a probability measure over M by satisfying Equations (4) and (5). In Gibbs' terminology, the conditions specified by Equations (4), (5) and (11) combine to form the condition for statistical equilibrium. This still leaves the functional form of P underdetermined, so Gibbs places a further constraint on P: based on the restriction of his study to conservative mechanical systems, he stipulates P to be some function of the energy:

$$P(s) = \alpha(h(s)) \text{ for some } \alpha : \mathbb{R} \to \mathbb{R} \text{ for all } s \in M. \qquad (12)$$

Part of the reason for this stipulation is, no doubt, to focus discussion on particular probability coefficients; taking P to be some function of the energy is an easy and natural way to satisfy the condition for statistical equilibrium. However, it is a mathematical fact that $\{P, h\} = 0$ for a conservative system does not imply that P must be some function of the energy. Gibbs's stipulation that this be so therefore appears rather arbitrary from a philosophical point of view.

However, it is a fact that Gibbs restricted his attention to stationary probability coefficients which were some function of the energy. Insofar as

we are in the business of reconstructing Gibbs' reasoning, or indeed any reasoning which restricts attention to functions of the energy, dynamical equivalence is appropriate.

I now develop a precise statement of dynamical equivalence. Consider the phase flow φ_t generated by h and the particular trajectory

$$t \mapsto \varphi_t(s) \in M \tag{13}$$

with initial condition s. Now compare the trajectory with initial condition $\pi(s)$ for some $\pi \in S_N$:

$$t \mapsto \varphi_t(\pi(s)) \tag{14}$$

We also define the permutation of the curve $t \mapsto \varphi_t(s) \in M$ as

$$t \mapsto \pi(\varphi_t(s)) \tag{15}$$

This is just like the curve $t \mapsto \varphi_t(s)$ except that every point along the curve has been acted upon by π. Intuitively, it is the very same trajectory, differing only in the factor positions of the single particle states. This motivates a definition for states s and $\pi(s)$ to be dynamically equivalent.

π-Dynamical Equivalence. States s and $\pi(s)$ are π-*dynamically equivalent* if and only if $\varphi_t(\pi(s)) = \pi(\varphi_t(s))$, for all t.

In words, this asks you to compare two trajectories and starting at some point and the permuted version of that point. If they are the same trajectory up to permutation π, then the states are π-dynamically equivalent. Abstracting away from particular trajectories, this becomes: if the transformation π leaves the hamiltonian flow invariant, then we say that the states s and $\pi(s)$ are π-dynamically equivalent.

This definition does not yet look like an instance of the observational equivalence schema in Section 4. However, we can prove that it is equivalent to one such instance using the following result: if a canonical transformation leaves the hamiltonian flow invariant, then this leaves the Hamiltonian invariant up to an additive constant. The geometric proof goes as follows[18].

Proposition. Given a symplectic manifold (M, ω), let $\varphi_t : M \to M$ be the Hamiltonian flow generated by the smooth function $h : M \to \mathbb{R}$. Let $\pi : M \to M$ be a symplectomorphism (canonical transformation), such that $\pi^*\omega = \omega$ where π^* is the pushforward of π. Then: $\varphi_t \circ \pi = \varphi_t$ iff $h \circ \pi = h + f_0$ where f_0 is a constant function.

18. An equivalent proposition and proof may be found in Corollary 8 in Arnol'd (1978, pp. 217–8)

Proof. →. A Hamiltonian flow for h is a set of integral curves threading the vector field H. So the proposition is equivalent to the statement that $\pi^* H = H$ iff $h \circ \pi = h + f_0$. H is defined to satisfy $\omega_{ab} H^a = d_b h$. Writing the inverse of ω_{ab} as ω^{ab}, we have $H^a = \omega^{ab} d_b h$. So $\pi^* H = H$ is to say $\pi^* \omega^{ab} d_b h = \omega^{ab} d_b h$. Since π^* is a symplectomorphism, $\omega^{ab} \pi^* d_b h = \omega^{ab} d_b h$. Applying ω_{ab} to both sides, we have $\pi^* d_b h = d_b h$. All diffeomorphisms commute with the exterior derivative, so this implies that $d_b \pi^* h = d_b h$. By the definition of pullback, this implies $d_b (h \circ \pi) = d_b h$. By linearity of the exterior derivative, $d_b (h \circ \pi - h) = 0$. This implies that $h \circ \pi - h = f_0$.

←. Suppose $h \circ \pi = h + f_0$. Taking the exterior derivative, we have $d_b (h \circ \pi) = d_b h$. Now, $d_b (h \circ \pi) = d_b (\pi^* h) = \pi^* d_b h$. So $\pi^* d_b h = d_b h$. Applying ω^{ab} to both sides we have $\omega^{ab} \pi^* d_b h = \omega^{ab} d_b h$. π^* is a symplectomorphism so $\pi^* \omega^{ab} d_b h = \omega^{ab} d_b h$. This is just to say that $\pi^* H = H$. □

The truth of this claim is perhaps more easily seen in the coordinate setting of Hamilton's equations. Here is a sketch of how it works. Suppose the Hamiltonian h gives rise to the flow $(q_i(t), p_i(t))$. Suppose we canonically transform the coordinates so that $h \mapsto h'$ and $(q_i(t), p_i(t)) \mapsto (Q_i(t), P_i(t))$. Since the transformation is canonical, this flow results from plugging in some Hamiltonian h' into Hamilton's equations. Now suppose that this canonical transformation actually leaves the flow invariant: $q_i(t) = Q_i(t)$ and $p_i(t) = P_i(t)$, for all i, t. Then we have that $\partial h'/\partial p_i = \partial h/\partial p_i$ and $\partial h'/\partial q_i = \partial h/\partial q_i$ respectively. It follows that $h' = h + f(q_i)$ and $h' = h + g(p_i)$ respectively. This means that $h' = h + f_0$ where f_0 is a constant.

This proposition allows us to see the definition of π-dynamical equivalence as an instance of the observational equivalence schema, with $s' = \pi(s)$, $\{f_i\} = \{h\}$: states s and $\pi(s)$ are π-dynamically equivalent if and only if $h(\pi(s)) = h(s) + f_0$, for all s.

So far, we have talked about dynamical equivalence of states in a general and abstract geometrical setting. We now need to understand how these ideas on the level of the N particle state filter down to the level of the individual particles. To this end, we now apply these general geometrical ideas in the coordinate setting to better understand the functional form of certain typical Hamiltonians and their symmetries under the permutation maps. This will in turn motivate a definition of the dynamical indistinguishability of classical particles.

Consider a system of two particles moving in one dimension along the x-axis. The four dimensional phase space of this system has points of the form $(r_{1x}, p_{1x}, r_{2x}, p_{2x})$, two of which are depicted in Figure 2. According to the proposition just proved, the states s and $\pi(s)$ are dynamically equivalent

(a) $s = (1, -1, 2, 1)$. (b) $\pi(s) = (2, 1, 1, -1)$.

Figure 2: Two states related by a permutation. In each case, there is one particle with $(r, p) = (1, -1)$ and another with $(2, 1)$ and yet the states are represented by two distinct points in the phase space. The black dot depicts the particle represented by the first factor position and the white dot depicts the particle represented by the second factor position.

when the Hamiltonian is left invariant (up to an additive constant[19]) by the transformation: $h(\pi(s)) = h(s) + f_0$.

This is satisfied by $h_1 = (p_1^2 + p_2^2)/2m$. Explicitly, the p_1 picks out the momentum value in the first factor position which is -1 for s. p_2 picks out the momentum value in the second factor position which is 1 for s. When the Hamiltonian acts on $\pi(s)$, p_1 picks out 1 and p_2 picks out -1 since these values now occupy different factor positions as a result of applying π to s. However, it is clear that it does not matter which value is inserted into p_1 or p_2 because of the symmetry of the Hamiltonian: $h_1(s) = ((-1)^2 + (1)^2)/2m = ((1)^2 + (-1)^2)/2m = h_1(\pi(s))$.

This condition is not satisfied, however, by $h_2 = (p_1^2 + p_2^2)/2m + GMm/r_1$, where GM/r_1 is the gravitational potential at r_1 (the position of the particle represented by the first factor position) of some fixed, stationary mass M. $h_2(s) = ((-1)^2 + (1)^2)/2m + GMm/(1)$ while $h_2(\pi(s)) = ((1)^2 + (-1)^2)/2m + GMm/(2)$ since r_1 picks out the value 1 for state s whereas it picks out the value 2 in state $\pi(s)$. It is clear that (in general) $h_2(s) \neq h_2(\pi(s))$, since the position and momentum variables pick out different numbers as a result of the permutation. The fact that h_2 is not symmetric in this sense means that the states s and $\pi(s)$ have different dynamical evolutions. s and $\pi(s)$ are dynamically equivalent with respect to h_1 but dynamically inequivalent with respect to h_2.

Let us consider a slightly more complicated Hamiltonian of a four particle system moving in one dimension: $h_3 = (p_1^2 + p_2^2 + p_3^2 + p_4^2)/2m + ke_1e_2/|r_{13}|$, where $ke^2/|r_{13}|$ is the electric potential energy term as a function of the relative distance between the particles represented by factor positions 1 and 3 both with charge e and k is a constant. The particles represented by factor positions 2 and 4 are neutral. In this case, the Hamiltonian satisfies

---

19. That the Hamiltonian is only invariant up to an additive constant is not a problem, since we only ever observe changes in energy anyway.

$h_3(s) = h_3(\pi_{13}(s))$ and $h_3(s) = h_3(\pi_{24}(s))$ for all s but it is not the case that, in general, $h_3(s) = h_3(\pi_{23}(s))$ nor $h_3(s) = h_3(\pi_{14}(s))$. The reason for this is that particles 1 and 3 are neutral and 2 and 4 have the same electric charge, so permuting single particle states among those factor positions makes no difference to the dynamical evolution. However permuting states among factor positions 1 and 2 for example, would lead to a different dynamical evolution, since this swaps states between a particle with an electric force acting on it with a particle with no such force.

These three examples of Hamiltonians motivate a precise definition of the 'dynamical indistinguishability' of classical particles:

Dynamical Indistinguishability. Two particles of an N particle system, represented by factor positions i and j in the N tuple, are *dynamically indistinguishable* if and only if the Hamiltonian h of the system satisfies $h(s) = h(\pi_{ij}(s)) + f_0$ for all states s.

An immediate consequence of this definition is that the dynamical indistinguishability of the particles is relative to the Hamiltonian; in fixing h, we make claims about the dynamical indistinguishability of certain particles in that system. We may link dynamical indistinguishability with the more intuitive and familiar, albeit less precise, notion of identical particles. This is best illustrated with an example.

Consider a collection of particles in a gas which are all neutral except for one charged particle. This particle will, of course, not interact with any of the neutral particles. If there is no electromagnetic field present, then there will be no electromagnetic interactions at all, and the Hamiltonian will contain no electromagnetic potential terms. Although the particle is charged, this charge is not being exploited by any field; there are no forces on this particle which depend on this charge, and so the charge is irrelevant to the dynamical evolution of the system. The charged particle would be dynamically indistinguishable from the neutral particles, even though they are not identical. We may change the Hamiltonian by turning on an electromagnetic field, inducing a force on the charged particle resulting in a different trajectory than it would have otherwise taken. In virtue of this, the particle is dynamically distinguishable from the other neutral particles.

This shows that differing intrinsic properties of the particles is a necessary but not sufficient condition for dynamical distinguishability. The particle must be charged for the field to act only on it with a force, but the mere fact that it is charged is not sufficient for it to be dynamically distinguishable, since it may be treated as neutral if there is no electromagnetic field present in the system. Thus, while it is necessary that there be some

difference for the particles to be dynamically distinguishable, it is not necessary that they share all their intrinsic properties for them to be dynamically indistinguishable. Particles may be dynamically indistinguishable even if they do in fact differ with respect to certain intrinsic properties, since the Hamiltonian may be 'blind' to those properties.

6   Solving the $N!$ Problem

In this Section, I show how we may rigorously derive the $N!$ factor on the basis of observational equivalence of states, and the observational indistinguishability of particles. I will make this argument in the particular case of dynamical equivalence, since this is the situation that Gibbs and others often assume, but it is fully generalisable. The result in this Section is a particular instance of the more general problem of deriving some over-counting correction factor on the basis of some definition of the observational equivalence of states.

Let us now consider, for simplicity, the Hamiltonian $h = \sum_i^N \mathbf{p}_i^2/2m$. In this case, all the $\pi_i \in S_N$ are dynamical symmetries. Thus, given a state $s \in M$, $\pi_i(s) \in M$ is dynamically equivalent to s for all $\pi_i \in S_N$. Therefore, relative to h, all N particles are dynamically indistinguishable from one another.

The images of s under each of the $N!$ permutation maps are in $N!$ distinct points in M representing states which are all dynamically equivalent to each other. Let us now consider instead not the points themselves, but some open region around the points. Thus, let U be the open region around $s \in U$. We may now apply the permutation maps to U to generate $N!$ distinct regions in M, $U_i := \pi_i[U]$, such that $\pi_i(s) \in U_i$. Since M is Hausdorff, for any two of these regions we have that $\pi_i(s) \in U_i$ for $i = 1, 2$ and $U_1 \cap U_2$ is empty. We may also choose the regions U_i such that they cover M, so we have that for every $s \in M$ there is a U_i such that $s \in U_i$ and $\cup_{i=1}^{N!} U_i = M$.

By this construction, the manifold has been partitioned into $N!$ dynamically equivalent regions; there is no point in M which is not contained in some U_i and no two points in any one U_i are dynamically equivalent. Therefore, any one of these regions suffices to completely represent all the observationally (dynamically, in this case) distinct states of the system. That is, any one of these regions may be considered to be the sample space: $\Omega = U_i$ for any i. Suppose that this region is equipped with any measure of the form

$$\mu(U_i) := \int_{U_i} P(p, q) \, dp_1...dq_n = 1 \text{ for some } i \tag{16}$$

where $dq...dp = \omega^n$ is the symplectic volume form (see footnote 5).

If we now follow Gibbs and let P satisfy Equation (12) (so we let P be some function of the Hamiltonian), then μ will be preserved under those canonical transformations π_i which are also dynamical symmetries: ω^n is preserved by canonical transformations and P is preserved by dynamical symmetries. For the Hamiltonian we are considering, all of the permutations are dynamical symmetries, therefore

$$\mu(U_i) = 1 \text{ for all } i \tag{17}$$

from which, recalling that $M = \cup_{i=1}^{N!} U_i$, it follows that

$$\mu(M) = \sum_{i=1}^{N!} \mu(U_i) = N!\mu(U_i) = N!. \tag{18}$$

Therefore, when we want to compute the probability of the system being in region $R \subseteq M$ using Equation (5) we substitute in Equation (18) to get $\mathcal{P}(R) = \mu(R)/N!$. This concludes the derivation of the $N!$ factor and justifies why we must divide the measure of some region R by $N!$ to correctly compute the probability of that region. But I must emphasize that this procedure is a specific case of a more general and standard statistical procedure. If we want to compute the probability that the system lies in some region $R \subset M$, we first compute the measure of that region, $\mu(R)$ and then divide by $\mu(M)$ according to Equation (5). Usually it is implicitly assumed that $\mu(M) = 1$; it is assumed that the manifold is the sample space. One of the main results of this paper is that this is not always the case and this fact can be used to solve the $N!$ problem. In the general case where the sample space is not the whole manifold, $\mu(M) \neq 1$. In this particular example, $\mu(M) = N!$.

This argument in this Section generalizes: if the Hamiltonian satisfies $h(\pi(s)) = h(s)$ for all s and for all $\pi \in S_n$ where S_n is some permutation subgroup of S_N, then $\mu(M) = n!$. It may also be the case that the Hamiltonian satisfies $h(\pi(s)) = h(s)$ for all s and for all $\pi \in S_n$ and for all $\pi \in S_m$ where S_m is an altogether different subgroup. In this case $\mu(M) = n!m!$.

The $N!$ factor is often perceived as a 'correction', indicating that some reasoning over and above the standard statistical reasoning is needed to justify its inclusion. What I have shown is that this is not the case. The $N!$ arises perfectly naturally in the course of statistical reasoning as the

normalisation factor $\mu(M)$ when all the permutations are dynamical symmetries. The reason it is perceived as a correction factor, is because division by $\mu(M) = 1$ renders itself invisible in cases when the manifold is the sample space. Thus, when the $N!$ factor appears, it looks like we have added something in that was not there before.

7 Conclusion

In this paper, I have given a solution to the $N!$ problem, based on the dynamical equivalence of the states and the dynamical indistinguishability of particles, whose precise definitions were inspired by considerations of observation, and its importance in specifying the sample space of any statistical theory.

These considerations of observation were key to my analysis. As we saw, the literature is most focused on whether two states are really 'one and the same' or whether particles are really 'identical'. This ontological view of the matter was replaced by the more epistemic notions of observational equivalence and observational indistinguishability respectively[20]. But why take this more 'epistemic' route with respect to the $N!$ problem in statistical mechanics? The first reason was that the basic object in a statistical theory is the probability space, one of whose elements being the sample space, typically characterized as containing all and only states which are observationally distinct from one another. Thus, epistemic notions enter into the very construction of the mathematical object we study in statistical mechanics. The second reason was that the notion of particles sharing intrinsic properties like mass and charge was best interpreted formally and precisely in terms of observables; the state of the system only gives us position and momentum values, so in order to access their mass or charge, we must look to certain observables like the kinetic energy or the electrical potential, for example.

The notion of observational equivalence was further narrowed to the notion of dynamical equivalence, taking the set of observables with respect to which to judge the equivalence of states to be the singleton set containing the Hamiltonian. It is granted that dynamical equivalence is not the only

---

20. I must emphasize at this point that the epistemic flavour of my analysis enters in the construction of the sample space, before any probability measures are defined on the space; the epistemic flavour does not enter into the interpretation of the probability measure. In fact, my analysis here is compatible with any interpretation of statistical mechanical probabilities: nothing in my argument requires us to make any choice of probability interpretation.

sense in which states may be said to be observationally equivalent, and hence nor is dynamical indistinguishability the only sense in which particles may be said to be indistinguishable. This is accommodated in my analysis in the fact that the different senses of particle (in)distinguishability and state (in)equivalence correspond to different choices of the set $\{f_i\}$.

Although I have restricted my discussion to classical Hamiltonian mechanics, the characterisation of observational and dynamical equivalence I develop in this paper will carry over straightforwardly to quantum theory; the characterisation of indistinguishable classical particles in terms of the permutation symmetry of the Hamiltonian takes the form of the permutation operator commuting with the Hamiltonian. The framework presented here therefore promises to make interesting and important connections between classical and quantum physics in the context of particle indistinguishability.

References

Arnol'd, V. I. 1978. *Mathematical Methods of Classical Mechanics.* Translated by K. Vogtmann and A. Weinstein. Graduate Texts in Mathematics. Springer New York.

Belot, Gordon. 2003. "Notes on Symmetries." In *Symmetries in Physics: Philosophical Reflections,* edited by Katherine A. Brading and Elena Castellani, 393–412. Cambridge University Press.

Billingsley, P. 1995. *Probability and Measure.* Wiley.

Blundell, Stephen J., and Katherine M. Blundell. 2010. *Concepts in Thermal Physics.* 2nd ed. Oxford University Press.

Butterfield, Jeremy. 2007. "On Symplectic Reduction in Classical Mechanics." In *Philosophy of Physics,* edited by Jeremy Butterfield and John Earman, 1–132. Handbook of the Philosophy of Science. Amsterdam: North-Holland.

Castellani, Elena. 2003. "Symmetry and equivalence." In *Symmetries in Physics: Philosophical Reflections,* edited by Katherine Brading and Elena Castellani, 425–436. Cambridge University Press.

Caulton, Adam, and Jeremy Butterfield. 2011. "On Kinds of Indiscernibility in Logic and Metaphysics." *The British Journal for the Philosophy of Science* 63, no. 1 (August): 27–84.

Caulton, Adam, and Jeremy Butterfield. 2012. "Symmetries and Paraparticles as a Motivation for Structuralism." *The British Journal for the Philosophy of Science* 63 (2): 233–285.

Darrigol, Olivier. 1991. "Statistics and Combinatorics in Early Quantum Theory, II: Early Symptoma of Indistinguishability and Holism." *Historical Studies in the Physical and Biological Sciences* 21 (2): 237–298.

Denbigh, K. G., and M. L. G. Redhead. 1989. "Gibbs' Paradox and Non-Uniform Convergence." *Synthese* 81 (3): 283–312.

Dieks, Dennis. 1990. "Quantum Statistics, Identical Particles and Correlations." *Synthese* 82 (1): 127–155.

———. 2011. "The Gibbs Paradox Revisited." In *Explanation, Prediction, and Confirmation,* edited by Dennis Dieks, Wenceslao J. Gonzalez, Stephan Hartmann, Thomas Uebel, and Marcel Weber, 367–377. Dordrecht: Springer Netherlands.

———. 2013. "Is There a Unique Physical Entropy? Micro versus Macro." In *New Challenges to Philosophy of Science,* edited by H. Andersen, D. Dieks, W González, T. Uebel, and G. Wheeler, 23–34. Springer Verlag.

French, S., and D Krause. 2006. *Identity in Physics: A Historical, Philosophical, and Formal Analysis.* OUP: Clarendon Press.

French, Steven. 1989. "Identity and individuality in classical and quantum physics." *Australasian Journal of Philosophy* 67 (4): 432–446.

———. 2000. "Identity and individuality in quantum theory." In *The Stanford Encyclopedia of Philosophy,* Winter 2019, edited by Edward N. Zalta. Metaphysics Research Lab, Stanford University.

French, Steven, and Dean Rickles. 2003. "Understanding Permutation Symmetry." In *Symmetries in Physics: Philosophical Reflections,* edited by Katherine Brading and Elena Castellani, 212–238. Cambridge University Press.

Gibbs, Josiah Willard. 1902. *Elementary Principles in Statistical Mechanics.* New Haven: Yale University Press.

Huang, Kerson. 1963. *Statistical Mechanics.* John Wiley & Sons.

Jaynes, E. T. 1992. "The Gibbs Paradox." In *Maximum-entropy and Bayesian methods,* edited by G. Erickson, P. Neudorfer, and C. R. Smith, 1–21. Dordrecht: Kluwer.

Myrvold, Wayne. 2017. "Probabilities in Statistical Mechanics." In *The Oxford Handbook of Probability and Philosophy,* edited by Alan Hajek and Christopher Hitchcock. Oxford University Press.

Saunders, Simon. 2013. "Indistinguishability." In *The Oxford Handbook of Philosophy of Physics,* edited by R. Batterman, 340–380. Oxford University Press.

———. 2018. "The Gibbs Paradox." *Entropy* 20 (8).

Swendsen, R. H. 2018. "Probability, Entropy, and Gibbs' Paradox(es)." *Entropy* 20 (450).

ter Haar, D. 1966. *Elements of Thermostatistics.* Holt, Reinhart & Winston.

Uffink, Jos. 2007. "Compendium of the Foundations of Classical Statistical Physics." In *Philosophy of Physics,* edited by Jeremy Butterfield and John Earman, 923–1074. Handbook of the Philosophy of Science. Amsterdam: North-Holland.

van Kampen, N. G. 1984. "The Gibbs Paradox." In *Essays in theoretical physics in honour of Dirk ter Haar,* edited by W. E. Parry, 303–312. Oxford: Pergamon Press.

van Lith, Janneke. 2018. "The Gibbs Paradox: Lessons from Thermodynamics." *Entropy* 20 (5).