Maxwell-Boltzmann Statistics and the Metaphysics of Modality

Author(s): Bruce L. Gordon

Source: *Synthese*, Dec., 2002, Vol. 133, No. 3 (Dec., 2002), pp. 393-417

Published by: Springer

Stable URL: https://www.jstor.org/stable/20117313

BRUCE L. GORDON

MAXWELL–BOLTZMANN STATISTICS AND THE METAPHYSICS OF MODALITY

ABSTRACT. Two arguments have recently been advanced that Maxwell-Boltzmann particles are indistinguishable just like Bose–Einstein and Fermi–Dirac particles. Bringing modal metaphysics to bear on these arguments shows that ontological indistinguishability for classical (MB) particles does not follow. The first argument, resting on symmetry in the occupation representation for all three cases, fails since peculiar correlations exist in the quantum (BE and FD) context as harbingers of ontic indistinguishability, while the indistinguishability of classical particles remains purely epistemic. The second argument, deriving from the classical limits of quantum statistical partition functions, embodies a conceptual confusion. After clarifying the doctrine of haecceitism, a third argument is considered that attempts to deflate metaphysical concerns altogether by showing that the phase-space and distribution-space representations of MB-statistics have contrary haecceitistic import. Careful analysis shows this argument to fail as well, leaving *de re* modality unproblematically grounding particle identity in the classical context while genuine puzzlement about the underlying ontology remains for quantum statistics.

It has been the contention of some that the mere fact that Maxwell–Boltzmann particles are capable of demonstrating non-classical (Bose–Einstein or Fermi–Dirac) statistical behavior gives sufficient evidence that they too are properly conceived to be indistinguishable.[1] There are two arguments to consider in this respect, and a third that is instructive for further probing the root of the matter. In the first instance, four Italian physicists – Costantini, Galavotti, Garibaldi, and Rosa – have argued that what is fundamental to an understanding of the classical and quantum statistics of particles is the type of correlation existing between them, and that the indistinguishability of particles need not concern us at all since it is an artifact of the probability functions, not a property of the particles themselves. On the basis of a "reconstruction" of elementary particle statistics, Costantini himself has argued that a careful analysis of the Boltzmann distribution yields the conclusion that classical particles are indistinguishable, i.e., unable to be distinguished relative to the probability distribution characteristic of MB statistics. More recently, Fujita has argued that the space-time "traceability" and "labelability" of classical particles does not entail their distinguishability, employing an argument

Synthese **133:** 393–417, 2002.

which involves taking the "classical mechanical" and "classical statistical" limits of the quantum statistical partition functions. He then suggests a reconceptualization of MB statistics under the assumption of indistinguishability. I will argue that the views of Costantini et al. on the one hand, and of Fujita on the other, rest on a conflation of epistemology with metaphysics and have no metaphysical import.

Finally, in a different vein, Nick Huggett quite recently has constructed an argument that classical particle statistics has no metaphysical implications because (so he asserts) different representational spaces have contrary metaphysical import. He first suggests that MB statistics represented in phase-space evinces an haecceitistic structure, and then argues that haecceitism *fails* in an equally viable distribution-space representation. Furthermore, contrary to popular perception, Huggett maintains that the observed statistics in fact are indifferent between the representations, rather than confirming the correctness of phase-space over distribution-space. This part of the argument is intended to remove any positive ground for thinking that MB statistics is intrinsically haecceitistic. The second part counters the opposing suggestion that MB statistics is intrinsically anti-haecceitistic by rejecting a thermodynamic argument asserting that taking classical statistical mechanics to have an haecceitistic structure leads to an empirically incorrect value for the entropy of the system. If the thermodynamic argument were correct, it would generate the conclusion that haecceitism is demonstrably (experimentally) false for MB systems. But this argument rests on there being *prima facie* evidence from thermodynamics that the distribution-space description of the system yields the correct value for equilibrium entropy, while the phase-space description does not. Huggett concludes, correctly in my opinion, that this thermodynamic argument fails to generate an anti-haecceitist conclusion, but then moves on to a conclusion which does not follow, namely that classical physics lacks metaphysical import and therefore that no new wrinkles in the metaphysics of individuality are introduced by quantum statistics.

Beginning with a metaphysical prelude, therefore, and countenancing another metaphysical discussion by way of intermission, the positions of Costantini, Fujita and Huggett will be evaluated in turn. The conclusion will be that all three are unsuccessful in their bid (implicit or explicit) to extirpate modality *de re* from the domain of classical statistical mechanics.

1. MODALITY *DE RE*: THE VERY IDEA

The task of the present section is somewhat Kantian in character. What we want to do, so to speak, is delineate the "transcendental" ground for the very possibility of material identity, i.e., to limn the circumstances under which explicating material identity and referential semantics in the categories of *de re* modality – essentialism, haecceitism, the Kripke–Putnam causal theory, Plantinga's Boethian compromise, etc. – is possible with respect to the ontology of a given physical theory. Not surprisingly, I think it is fair to say that this can happen when there is a purchase point provided by classical criteria of (material) individuation: one or more of a set of properties that are uniquely attributable to the object in question; the existence of a unique spatio-temporal location or trajectory, subject to the impenetrability assumption that no two objects can ever occupy the same spatio-temporal location; and the possibility of labeling (naming) the object in a way that uniquely fixes its reference. If classical individuative criteria of this sort are not applicable, then not only are metaphysical doctrines like essentialism and haecceitism and various theories pertaining to referential semantics otiose in this realm, but the objects themselves, if it is even possible for them to exist, do not have intrinsic (*de re*) identities. Let's elaborate on this point a bit, focusing first on particle statistics as the relevant context.

There is a well-worn tradition in the literature of the philosophy of science that makes a distinction between observational (operational) and conceptual distinguishability. Conventionally, two physical entities are said to be *observationally* (*operationally*) *distinguishable* just in case replacing one by the other leads to a measurably different state of the physical system. Particles obeying Bose–Einstein, Fermi–Dirac and various parastatistics are observationally *indistinguishable*. Two physical entities are correspondingly defined to be *conceptually distinguishable* just in case it is possible *in principle* to regard them as distinct, even if it is impossible in practice to distinguish them. Indeed, if they are not conceptually distinguishable in this manner, the question arises as to what meaningful sense can be given to the suggestion that they are "two", that is, in what their individuality or identity consists. This latter question seems to provide some motivation for Leibniz's *Principle of the Identity of Indiscernibles* (*PII*) – that it is impossible for two individuals to differ *only* numerically. We may provisionally differentiate between the notions of physical identity/individuality and physical distinguishability by regarding the former as something *intrinsic* to an object (associated with it uniquely) and the latter as an *extrinsic* relation among two or more objects. *Prima facie*, that individuality and identity can be distinguished in such manner from distinguishability minimally seems to presuppose the truth of some form of haecceitism or essentialism.

The notion of the conceptual distinguishability of physical objects may be clarified with the help of the three (potentially connected) criteria mentioned earlier. A physical object may be regarded as conceptually distinct from another just in case: (1) it can be distinguished with reference to one or more of a set of properties attributable to it uniquely; or (2) it can be distinguished by its spatio-temporal location/trajectory, under the impenetrability assumption that no two individuals can ever occupy the same spatio-temporal location; or (3) it is labelable (nameable) in a way which uniquely fixes its reference.

In the quantum mechanical case, we may distinguish between the *intrinsic* and the *state-dependent* properties of particles. Particle kinds in physics are distinguished by possession of the same intrinsic properties. Since intrinsic properties are definitive of particle *kinds*, unless some extra-dynamical individuative intrinsic properties (like particle indices/labels) are artificially introduced into the theory, *individual* particle distinctions will have to be made on the basis of contingent properties. But contingent (non-intrinsic) properties are ones that are dependent on the *state* of the particle(s) in question. Among the state-dependent properties are such particle attributes as position, momentum, and energy. Apart from the de Broglie–Bohm interpretation, the peculiarities and difficulties of which we will not discuss here, the individuative capacity of the contingent/state-dependent properties is suspect when the common occurrence of superposed states is considered, as is the idea that spatio-temporal location could be employed as an individuative concept in the quantum domain. So it seems that the most relevant "in principle" distinguishability criterion is connected with the possibility of treating individual particle labels as rigid designators or (as in Plantinga's theory) expressive of individual essences (haecceities) in the various theoretical contexts. Insofar as treating particle labels this way is possible in the quantum context, the *in principle* compatibility of quantum statistics with essentialism and haecceitism will remain secure. There are good reasons to suspect that quantum theory allows no purchase point for *de re* modality, however, but this discussion is beyond the scope of our present interests.[2]

By way of elaboration on these metaphysical themes, note that the standard Aristotelian line on the identification of material substances is that they are individuated by their matter, and at the most basic level of analysis the matter composing an entity is not a substance that is separable from it. The question of how we are to individuate matter itself then arises. Aristotle did not really provide an answer, but the medievals answered this question by postulating special properties which guarantee the uniqueness of a material individual. The first was *ubiety*, which is something like a unique absolute spatial position, and the second was *haecceity*, a primitive *this*ness of identity which precedes and serves as a substantial substratum for all of an object's qualitative properties. If we are to assert that a material object has an intrinsic (*de re*) identity, it seems pretty clear that something like these special properties must hold – something must uniquely distinguish it from other material objects, *qua* material. Spatial overlap or coincidence is therefore precluded (impenetrability must be affirmed), and in light of the standard *gedankenexperiment* involving two objects which share all of their qualitative properties in a completely spatio-temporally symmetric universe, a primitive thisness of material composition (material haecceity) must be affirmed.

When we arrive at what physics takes to be the fundamental material simples (subatomic "entities" having no parts), the issues critical to their possession of material individuality are the existence of uniquely individuating spatio-temporal trajectories (of which impenetrability is a corollary), and their possession of a material haecceity or uniquely individuating qualitative properties. These issues are not separate, since both material haecceities and material bearers of qualitative properties must have a location – there can be no irreducible or qualitatively identified *this* in a physical sense if there is no unique and definite location where *this* is found. Let me emphasize that we are speaking ontologically here, not epistemically. If a material object has an haecceity, it is conceptually possible for us to be irremediably ignorant of its unique location, even though it has one. It is not possible, however, for there to be material individuality in the ontological absence of *any* location at which *this* supposed material individual exists.

If these lessons are taken to heart, what must be said about the *properties* (qualitative or non-qualitative) of material individuals? The requisite metaphysics can be captured this way:

> All material individuals I are such that for every property P having a well-defined value or range of values, and all times t during which I may be said to exist, either I exemplifies a definite value of P at t or I does not possess any value of P at t (i.e., I does not possess P at t at all).

We might even include spatio-temporal attributes like "being at spatio-temporal location (x, y, z, t) in reference frame R" in the scope of such properties. If these various conditions are *not* met, it would be a mistake to think that we are dealing with a material individual at all, since there is no primitive substantial thisness in view, no spatio-temporal location in question, and there are no identity-conferring properties to which we have recourse. In the absence of any individuality, all labels, names or indices attaching to the purported material entities must be regarded as fictions. No material thing is named if the intended referent has no substantial thisness, no location, and no uniquely identifying properties. If a catch-phrase is desired, we could do little better than to borrow Quine's dictum that "there is no entity without identity".

2. COSTANTINI'S "RATIONAL RECONSTRUCTION" OF STATISTICAL MECHANICS

Four Italian physicists – Costantini, Galavotti, Garibaldi and Rosa – recently have attempted a Carnapian rational reconstruction of the foundations of statistical mechanics. What they have proposed is a reconceptualization of MB, BE and FD statistics on the basis of a set of "ground hypotheses" (conditions) that imply the respective probability distributions. In brief, the details of the reconstruction are as follows (cf. Costantini et al. 1983, 153–156):

The phase-space is partitioned into occupation cells $V_1, \ldots V_k$ over which particles $a_1, \ldots a_N$ are distributed by an assignment E. A statistic, S_E, on E is a probability distribution over E. A given assignment E is described by a proposition associating occupation cells as "attributes" with particles as "individuals". An atomic proposition in this context consists in the assignment of a particle a_i to a cell V_j, denoted by $a_i \in V_j$. The position in phase space of a particle a_i, $1 \leq i \leq N$, can be interpreted as a random variable, and the set of all the occupation cells V_j, $1 \leq j \leq k$, as the state space. S_E can then be constructed as a stochastic process with discrete parameters, where $\Pr\{a_i \in V_j\}$ is the initial probability of particle a_i being in cell V_j, and

$$(2.1) \quad \Pr\{a_{i+1} \in V_j \mid (a_1 \in V_{j_1}) \cap \cdots \cap (a_i \in V_{j_i})\}$$

is the "transition probability" that particle a_{i+1} is in cell V_j given that particles a_1 through a_i are in cells V_{j_1} through V_{j_i}, respectively.

With this structure in place, some definitions are put forward. The probability function Pr is called *regular* if it is always greater than zero for every $i \geq 0$; it is called *symmetric* if its value does not alter with any finite permutation of the particles. The value of a symmetric probability function thus depends solely on the number of particles in each occupation cell, not on their identity.

An MB proposition for N particles assigns a cell to each of them, that is,

$$(2.2) \quad E_{\text{MB}} = (a_1 \in V_1) \cap \cdots \cap (a_{N_1} \in V_1) \cap (a_{N_1+1} \in V_2) \cap \cdots \cap (a_N \in V_k),$$

where the number of particles in cell V_j is N_j, $1 \leq j \leq k$, and $\sum_{j=1}^{k} N_j = N$. The k-tuple characterizing E_{MB} is defined as

$$(2.3) \quad \mathbf{N} = (N_1, \ldots, N_j, \ldots, N_k), \sum_{j=1}^{k} N_j = N.$$

For every $\mathbf{N}$, define $\mathbf{N}^j = (N_1, \ldots, N_j + 1, \ldots, N_k)$, and $\mathbf{N}_0 = (0, \ldots, 0, \ldots, 0)$.

If E_{MB} and E'_{MB} are MB propositions, then E'_{MB} is said to be *isomorphic* to E_{MB} just in case E'_{MB} can be obtained from E_{MB} by a permutation of its particles. A BE proposition corresponding to a given E_{MB}, denoted E_{BE}, is constructed by taking the union of all the MB propositions which are isomorphic to that E_{MB}. The FD propositions, denoted E_{FD}, corresponding to a given E_{MB} are all of the corresponding E_{BE}'s the cells of which contain no more than one particle. Combinatorial analysis on this basis yields the conclusion that the number of E_{MB}'s is k^N, the number of E_{BE}'s is $\binom{N+k-1}{N}$, and the number of E_{FD}'s is $\binom{k}{N}$. The corresponding statistical distributions S_{MB}, S_{BE} and S_{FD} are therefore respectively $1/k^N$ on E_{MB}, $1/\binom{N+k-1}{N}$ on E_{BE}, and $1/\binom{k}{N}$ on E_{FD}.

If Pr is symmetric, it can be represented by a function having k-tuples as arguments: the distribution $p(\mathbf{N}) = (p_1(\mathbf{N}), \ldots, p_j(\mathbf{N}), \ldots, p_k(\mathbf{N}))$ is called the *representative function* of Pr just in case for all cells V_j and all $\mathbf{N}$, $p_j(\mathbf{N}) = \Pr\{a_{N+1} \in V_j \mid E_{\mathrm{MB}}\}$. If h and j are distinct state indices, and p is the representative function of Pr, then the *relevance quotient* of Pr is defined as

$$(2.4) \quad Q_j^h(\mathbf{N}) = \frac{p_j(\mathbf{N}^h)}{p_j(\mathbf{N})}.$$

The relevance quotient is thus the ratio of the transition probability to V_j at the $(i+1)$th step to the transition probability to V_j at the ith step, given that the $(i+1)$th particle is in a cell $V_h \neq V_j$. Pr is called *invariant* just in case for all $\mathbf{N}$ and $\mathbf{N}'$ such that $\mathbf{N} \neq \mathbf{N}'$, and all $j \neq h$ and $f \neq g$, $Q_j^h(\mathbf{N}) = Q_g^f(\mathbf{N}')$. To simplify the notation, stipulate that

$$\Pr\{a_i \in V_j\} = p_j(\mathbf{N}_0) = \gamma_j,$$

and

$$\frac{\Pr\{a_i \in V_j \mid a_m \in V_h\}}{\Pr\{a_i \in V_j\}} = Q_j^h(\mathbf{N}_0) = \eta_{jh}, \quad \text{for } j \neq h.$$

As a preliminary to delineating the ground hypotheses of MB, BE and FD statistics, a set of seven conditions is specified for the probability Pr:

(C1a) Pr is regular;

(C1b) $\Pr\{a_{N+1} \in V_j \mid (a_i \in V_{j_1}) \cap \cdots \cap (a_N \in V_{j_N})\} \geq 0$, if $j_i \neq j$ for every $1 \leq j \leq N$;

(C2) Pr is symmetric;

(C3) Pr is invariant;

(C4) For any j, $\gamma_j = 1/k$;

For Pr satisfying either of (C1a) or (C1b) and (C2) through (C4), there are three concomitant characteristics:

(C*) if $j \neq h$, then $0 \leq \eta_{jh} \leq 1$;

(C†) $\eta_{jh} = \eta_{gf} = \eta$; and

(C§) $p_j(N) = \dfrac{N_j + \gamma_j \lambda}{N + \lambda}$, where $\lambda = \eta/(1 - \eta)$.

(C5) $\eta = 1$;

(C6) $\eta = k/(k + 1)$; and

(C7) $\eta = k/(k - 1)$.

The key definition is now forthcoming: a set of conditions {C1, ..., Cn} is a set of *ground hypotheses* for a statistic S_E just in case {C1, ..., Cn} implies S_E on E.

From these definitions and conditions three theorems follow (for proofs, see Costantini et al. 1983, 155–156):

THEOREM 1. C1a, C2, C3, C4, and C5 constitute the ground hypotheses for S_{MB};

THEOREM 2. C1a, C2, C3, C4, and C6 constitute the ground hypotheses for S_{BE}; and

THEOREM 3. C1b, C2, C3, C4, and C7 constitute the ground hypotheses for S_{FD}.

What can be said about this "rational reconstruction"? We can see from the difference between (C1a) and (C1b), that FD particles fail to satisfy the regularity condition, which is just another way of saying that they obey the Pauli Exclusion Principle. Condition (C2) indicates that the probability distribution is dependent solely upon the number of particles in each cell, and is independent of their identity. Costantini (1979) has shown that the condition of invariance, (C3), is equivalent to the stipulation that the probability of a particle being in occupation cell V_j is a function of the number of particles in that cell, and the number of particles in other cells, independently of the way the particles in the other cells are distributed.

Conditions (C5), (C6), and (C7) give us the relevance quotients respectively characterizing MB, BE, and FD statistics. In (C5), the condition of η being equal to 1 is another way of expressing the statistical independence of the particles in the MB distribution. There is a null correlation when a particle's occupation of the jth cell is conditionalized upon another's occupation of the hth cell. In BE statistics, there is a negative correlation upon such conditionalization, indicated by the fact that η is less than one. This reflects the tendency of bosons to aggregate in cells which are already occupied rather than populating new ones, and points to the symmetry of the total wavefunction with respect to particle transpositions. On the other hand, η is greater than one for FD statistics, indicating a positive correlation reflective of the tendency of fermions to populate different cells. This tendency is linked to the antisymmetric character of the fermionic total wavefunction, expressing the essential content of the Pauli Exclusion Principle. It is worth noting that in the limit, as k increases without bound, both BE and FD statistics converge to S_{MB}, since the values of η get increasingly close to 1.

This is all well and good and perhaps even useful in some contexts, but why does Costantini (1987) think that it shows classical particles to be indistinguishable? The explanation, it turns out, is fairly simple. Costantini takes the ground hypothesis of symmetry, (C2), which characterizes all three statistics as dependent only upon the number of particles in each cell, and not upon which particles are in which cells, as the defining point of particle indistinguishability. Symmetric probability functions do not change their value when particles are permuted. Classical particles are thus indistinguishable with respect to the probability function which governs their behavior, namely the function which generates MB statistics.

Now, where does this leave the proponent of modality *de re* in respect of classical statistical mechanics? Should the partisans of essentialist or haecceitist interpretations be perturbed by these observations? In a word, no. The matter of permutation invariance in the context of classical statistics relates strictly to observational indistinguishability, not conceptual indistinguishability. For example, the probability function governing the life-expectancy of middle-aged men in actuarial calculations is also symmetric, but this hardly renders them indistinguishable in an interesting sense, let alone provides the occasion for an identity crisis.

Furthermore, the symmetry condition that Costantini regards as the well-spring of an indistinguishability held in common by all three statistical contexts, cannot really be regarded as having this consequence. His discussion fails to mention the counter-intuitive predictions quantum mechanics makes in relationship to single states as opposed to uniform mixtures, as well as the strange predictions that arise when non-commuting observables are considered. All that is embedded in the general characterization of symmetry (permutation invariance) given by Costantini et al. and mistakenly taken to be exhaustive of the notion of "indistinguishability", is de Finetti's purely classical notion of exchangeability (see Jeffrey 1988; Zabell 1988). The correlations that Costantini et al. discuss have a classical ignorance model in the quantum statistical context, because what they represent are the joint probabilities for a given set of *commuting* observables. Classical and quantum statistics can agree only under conditions of maximal ignorance, where both use a uniform statistical distribution. In conditions of less than maximal ignorance the symmetric character of the probability function generating a *classical indistinguishability* in both cases disappears, but peculiar correlations still remain in the quantum context (cf. van Fraassen 1991, 413, 417–418). So it seems that in the process of reducing classical and quantum statistics to a set of "ground hypotheses" some critical factors have been overlooked. Classical indistinguishability is purely epistemic – it arises from ignorance. Quantum indistinguishability runs deeper and gives every appearance of being ontic.

3. FUJITA ON PERMUTATIONAL SYMMETRY AND THE CLASSICAL LIMIT

Fujita (1991) begins and ends his discussion of the indistinguishability of classical particles with the physicist's standard for particle indistinguishability: the permutational symmetry of both the many-particle distribution function and the dynamic functions of a physical system. Under these conditions, no empirical property of the system discriminates among the particles. Fujita's goal is to argue that this concept of indistinguishability is every bit as necessary for classical particles as it is for quantum particles. As should be obvious by now, this criterion by itself is insufficient to pose any threat to essentialism or haecceitism, and pushing it too far does a grave injustice to some ineradicable differences between classical and quantum statistics.

From a philosophical standpoint, Fujita begins his discussion with a conceptual confusion (1991, 440):

> It is often said that classical particles are "essentially distinguishable" because we can trace their motion and label them. But "traceability" and "labelability" do not automatically lead to the distinguishability (*sic*). In order to distinguish between two particles, we must look at all of the properties of the two-particle system. By examining one-particle properties such as traceability, labeling possibility (and mass and charge) alone, we cannot distinguish between two particles. In other words, we must define the distinguishability (and indistinguishability) referring to all of the two-particle properties (*sic*).

On purely conceptual grounds, if two particles have spatio-temporal trajectories which are incapable of coincidence, and *a fortiori*, if they are individually labelable, they are *essentially distinguishable*. All the essentialist or haecceitist requires for distinguishability is a property (qualitative or non-qualitative) that uniquely individuates the object in question. If two objects are conceptually distinguishable, they are *ipso facto* essentially distinguishable. So understood, space-time traceability and particle labelability both provide a basis for distinguishability. What Fujita should have said is that, *relative to the physical system considered as a whole*, the empirical (in)distinguishability of the particles composing it is decided only after the question of permutation invariance (symmetry) is resolved for all of the particle transpositions within the system. With these constraints, it is clear that under certain conditions classical particles are empirically indistinguishable. But this is an indistinguishability arising from classical ignorance, and it gives no pause to the essentialist or haecceitist.

Another point worthy of note is that Fujita's reliance upon the classical statistical and classical mechanical limits of quantum systems in his argument for the (empirical) indistinguishability of classical particles obfuscates significant differences between classical and quantum statistics. Following Dirac's treatment of quantum particles, Fujita requires that a system of (empirically) indistinguishable classical particles satisfy the conditions

$$(3.1a)\quad P\rho(\vec{r}_1, \vec{p}_1, \ldots, \vec{r}_N, \vec{p}_N) = \rho(\vec{r}_1, \vec{p}_1, \ldots, \vec{r}_N, \vec{p}_N),$$

$$(3.1b)\quad P\xi = \xi \qquad \text{all } P,$$

where the P's are permutation operators, ρ represents the N-particle distribution function, and ξ the dynamical functions (including the Hamiltonian), the latter two being dependent upon the canonical particle variables for position and momentum. In order to achieve the requisite (empirical) indistinguishability, Fujita emphasizes that (3.1b) has to include *all* of the dynamical properties of the system, e.g., the Hamiltonian, the center of mass, the total linear and angular momenta, the mass, the momentum, the heat currents, etc.

Analogously, the requirements for the indistinguishability of quantum particles can be encapsulated in the commutator equations

(3.2a) $[P, \hat{\rho}] \equiv P\hat{\rho} - \hat{\rho}P = 0,$

(3.2b) $[P, \hat{\xi}] = 0,$

(3.2c) $[P, \hat{H}] = 0,$

with the density operator represented by $\hat{\rho}$, the operators for the dynamical observables by $\hat{\xi}$, and the permutation operators by P. (3.2c) is really subsumed under (3.2b), but Fujita notes it separately because of its significance – if the Hamiltonian is symmetric, the symmetry of the density operator is guaranteed, and particle indistinguishability is a permanent property of the system.

With the standard definitions of the Poisson and commutator brackets, the *classical mechanical limit* (CM limit) of a quantum system is defined to be

(3.3) $$\lim_{h \to 0} \frac{1}{i\hbar}[\hat{A}, \hat{B}] = \{A, B\},$$

where $A(\hat{A})$ and $B(\hat{B})$ are dynamical functions (operators). This limit is fictional, of course, because $\hbar$ is a physical constant. The limit represents the *transition* between the quantum and classical descriptions of a system. For our purposes, it is important to note that there are still residual quantum effects (dependent on Planck's constant) even *after* the classical statistical limit is taken (Fujita 1991, 443).

Statistical mechanics mathematically relates the thermodynamic properties of macroscopic objects to the motion of their microscopic constituents. Since the microscopic constituents obey quantum dynamics, the correct description must lie in principle within the domain of *quantum statistical mechanics*. Under thermodynamic conditions of high temperature (T) and low density (n), however, classical statistical mechanics serves as a useful approximation. With this in mind, we may define the *classical statistical limit* (CS limit) as the situation represented by:

(3.4) $$T \to \infty \quad \text{and} \quad n \to 0.$$

These are the same conditions as those governing the applicability of the ideal gas law (PV = RT), so (3.4) could equally well be called the ideal gas limit. Unlike the CM limit, the conditions governing the CS limit are subject to experimental control.

With the definition of these limits in hand, Fujita presses his central point (1991, 443):

> Both CM and CS are *continuous* limits. Therefore, these limits cannot alter the indistinguishability defined by the symmetry of many-body functions (operators) with respect to *discrete* permutations of particle indices. From this it follows that *classical* [*particles*] *are indistinguishable*.

This assertion is supposed to receive support from three detailed calculations: (1) In the CS limit region the Bose (or Fermi) distribution can be approximated by the Boltzmann distribution (1991, 444–446); (2) The BE and FD distribution functions both approach the MB distribution in the CM limit as well (1991: 446ff). In particular, if we consider a binary mixture of $N_1 + N_2 = N$ particles in three dimensions, the CM limit of the quantum partition functions (BE and FD) reveal that, when the quantum cells of dimension $(2\pi\hbar)^3$ are reduced to points in phase space, the indistinguishability factors $(2\pi\hbar)^3$ and $(N!)^{-1}$ arise naturally rather than being subject to *ad hoc* addition as they are in many treatments of statistical mechanics (1991, 448–451); and (3) when the Maxwell velocity distribution is derived under the assumption of indistinguishable particles, the thermodynamic perplexity referred to as "Gibbs' paradox" is obviated (1991, 451–453).

We respond briefly to these considerations. The fact that *quantum* particles, whether bosons or fermions, retain their indistinguishability in the classical limit (CM or CS), says nothing about the indistinguishability of *classical* particles in and of themselves. Therefore, to advance the indistinguishability of classical particles on the basis of (1) and (2) is a *non-sequitur* pure and simple – these considerations are irrelevant to the issue at hand. What *is* revealed by consideration of the continuity of the limits is that quantum particles retain their indistinguishability even when their behavior approximates an MB distribution. So *quantum* particles are intrinsically indistinguishable even when their behavior approximates that of classical particles. Furthermore, to make use of (1) and (2) in a case for the purported indistinguishability of classical particles suggests an assimilation of classical to quantum statistics (or vice-versa) in a manner that will not work, as we saw in the discussion of Costantini's reconstruction in the last section.

As regards the third consideration, note that the MB velocity distribution can be derived using the "Boltzmann relation" between entropy S, and occupation probability W: if we divide the phase space for N particles into k occupation cells, letting N_k represent the number of particles in each cell subject to the constraints that $\sum_k N_k = N$ and the total energy E is $\sum_k \epsilon_k N_k$, with $\epsilon_k = p_k^2/2m$, then the most probable distribution of the occupation numbers $\{N_k\}$ results from maximizing the Boltzmann relation $S = k_B \log W$ subject to $\sum_k N_k = N$. If the particles are assumed to be distinguishable, the occupation probability is given by the multinomial coefficient, that is,

$$(3.5) \quad W = \frac{N!}{\prod_k N_k!}.$$

Using (3.5), the MB distribution is derived in the large-number limit as $N \to \infty$. The difficulty with the derivation is that the entropy S, defined using Boltzmann's relation and (3.5), is not an extensive quantity (it is not additive in the sense of being proportional to the size of the system). To achieve the right experimental value, the "indistinguishability" factor $1/N!$ has to be inserted as an *ad hoc* correction to the occupation probability, yielding

$$(3.6) \quad W' = \prod_k \frac{1}{N_k!}$$

as the modified occupation probability. This *ad hoc* intervention leads to the MB velocity distribution without any further difficulties. The fact that such gerrymandering is necessary is essentially the content of Gibbs paradox.

In light of this, Fujita thinks it better to begin the derivation with an assumption of the "indistinguishability" of classical particles. If we assume this, then the factor $1/N!$ is included in the expression for the relative probability of N_k particles occupying the kth cell in phase-space right from the start. Then when all of the cells are considered, we obtain (3.6) straight away as the occupation probability for the system (1991, 452). From the vantage point of "indistinguishability", therefore, the derivation of the MB velocity distribution can be completed without any mid-course corrections.

But is this "streamlined" derivation metaphysically significant? Gibbs paradox is often thought to adumbrate quantum statistics. In my view it does no such thing, because the indistinguishability characterizing it involves the unexceptional classical notion of exchangeability. The reason that this notion is unexceptional is that it is an expression of epistemic ignorance. The same can be said about Fujita's derivation of the MB velocity distribution under the assumption of "indistinguishable" classical particles.

In cases such as these, where a classical ignorance model suffices, no metaphysical peculiarities are generated by the insensitivity of the probability distribution for a system to the individual identities of its members.

4. A TALE OF TWO HAECCEITISMS

Before we move on to discuss Huggett's views on the metaphysical commitments (or lack thereof) of MB statistics, a short metaphysical interlude is in order to get a proper feel for the doctrine of haecceitism. The term "haecceitism" was introduced into the lexicon of modal metaphysics by David Kaplan back in 1975.[3] We will take his seminal paper on the topic as our starting point. He defines the term as follows (Kaplan 1975, 722–723):

> [T]here seems to be some disagreement as to whether we can meaningfully ask whether a possible individual that exists in one world also exists in another without taking into account the attributes and behavior of the individuals that exist in the one world and making a comparison with the attributes and behavior of the individuals that exist in the other world. The doctrine that holds that it does make sense to ask – without reference to common attributes and behavior – whether *this* is the same individual in another possible world, that individuals can be extended in logical space (i.e., through possible worlds) in much the way we commonly regard them as being extended in physical space and time, and that a common "thisness" may underlie extreme dissimilarity or distinct thisnesses may underlie great resemblance, I call *Haecceitism*. ...
>
> The opposite view, *Anti-Haecceitism*, holds that for entities of distinct possible worlds there is no notion of transworld being. They may, of course, be linked by a common concept and distinguished by another concept ... but there are, in general, many concepts linking any such pair and many distinguishing them. Each, in [its] own setting, may be clothed in attributes which cause them to *resemble* one another closely. But there is no metaphysical reality of sameness or difference which underlies the clothes ... 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.
>
> Haecceitism holds that we can meaningfully speak of a thing itself – without reference either explicit, implicit, vague or precise to individuating concepts (other than being *this* thing), defining qualities, essential attributes, or any other of the paraphernalia that enable us to distinguish one thing from another. It may be that each thing has essential attributes with which it is vested at all times and in each possible world in which it exists. But that is an issue posterior to whether things have transworld being.

There are, I believe, two ways in which this doctrine can be understood. The first, and most obvious, is that there exist *haecceities* – non-qualitative properties of *this*ness that differentiate individuals. Robert Adams (1979) refers to this notion of non-qualitative identity as *primitive* thisness, where the primitiveness means that the metaphysical identity of an object precedes and underlies any qualitative properties it may have. A *this*ness, as the property of being identical with a *specific* individual, need not be primitive in this sense. If we contrast thisnesses with *suchnesses*, which are purely qualitative properties, it is a logical possibility that every thisness is analyzable into, perhaps even identical with, suchnesses. This was Leibniz's claim in relation to his principle of the identity of indiscernibles. But I take it that to assert that there are haecceities is to maintain there to be metaphysically primitive thisnesses which transcend all questions of suchness. If there are such non-qualitative thisnesses, then the principle of the identity of indiscernibles (construed as the doctrine that if A and B share all of their *qualities*, they are the same entity) is false. Furthermore, if there is a metaphysically primitive property of being identical with a specific individual, then the transworld identity of individuals is also primitive (and the supposed "problem" of transworld identity disappears). The reason for this is not far to seek: the property of being identical with a specific individual (taken to be metaphysically primitive) is the same property in every world in which it is instantiated, and (obviously enough) it is instantiated by the same individual in every world in which it is instantiated at all.

The second, and much less obvious, way of understanding haecceitism is put forth by David Lewis (1986, 220ff). Lewis takes the doctrine of haecceitism to be the denial of a supervenience thesis. Possible worlds can differ from one another in two basic ways – they can differ qualitatively by representing different patterns of instantiation of intrinsic qualities and external relations, or they can differ with respect to what they represent *de re* about various individuals. Of these two ways of differing, the second may be dependent on the first, or independent of it. One might maintain, with Leibniz, that representation *de re* supervenes on the qualitative character of the worlds, so that any two worlds differing in representation *de re* always differ qualitatively. Or one might deny this supervenience thesis, and maintain that there may exist *de re* differences among worlds without any difference in qualitative character. It is this latter option which Lewis takes as definitive of haecceitism (1986, 221):

> If two worlds differ in what they represent *de re* concerning some individual, but do not differ qualitatively in any way, I shall call that a *haecceitistic difference*. *Haecceitism*, as I propose to use the word, is the doctrine that there are at least some cases of haecceitistic difference between worlds. *Anti-haecceitism* is the doctrine that there are none.

The question now arises as to the connection between the two varieties of haecceitism. Can sense be made of an haecceitistic difference between worlds without recourse to haecceities? Conversely, can one accept haecceities and deny that there can be haecceitistic differences among possible worlds? Lewis thinks that both of these questions can be answered in the affirmative, but his reasons in the first case are far from clear. While acknowledging the logical compatibility of the two species of haecceitism, he remarks that "a haecceitist [in the second sense] need not believe in non-qualitative properties. He might even be a nominalist and reject properties altogether" (1986, 225). This is the sum and substance of his discussion about the matter. We have the illusion of a logical space being opened up here, but no sense of an intelligible position that might step in to fill it. Let us consider the options with which we are presented. Suppose that we eschew non-qualitative properties but we are not nominalists. Insofar as we speak of *this* individual as opposed to *that* one, we must presume thisness, as the property of being identical with a particular individual, to consist in the possession of, or be analyzable into, qualitative properties (suchnesses). This has to be the case since, without non-qualitative properties, there is nothing else left for thisness to be. In short, we must all be good Leibnizians. But now we cannot affirm that haecceitistic differences between worlds are possible, because every difference is a qualitative difference. Alternatively, suppose that we are nominalists and deny the existence of properties (qualitative or non-qualitative) altogether. If there are no properties to distinguish between worlds, how then are worlds distinguished? Presumably by extension (listing individuals) and by convention (definitional stipulation). Transworld identification for the nominalist is thus a matter of convention, not metaphysical connection. What sense then can be made of a nominalist using *de re* predications to distinguish worlds? None whatsoever. Any way you look at it, haecceitistic difference between worlds requires the acceptance of haecceities.

The converse is not true, however. One can embrace haecceities without admitting the possibility of haecceitistic differences between worlds. It can be asserted consistently that every individual has a metaphysically primitive (non-qualitative) thisness while also maintaining that every difference in the *de re* representation of a particular individual between worlds is qualitative in character. For example, we may follow Lewis (1986, 225) and associate with each individual the singleton containing that individual. By such an identification we obtain a non-qualitative property (identified with the singleton set) that we may take to represent the haecceity of each individual. With Plantinga, we may then insist that every individual possesses a unique essence comprised of the complete and consistent set of its world-indexed properties. This individual essence is so complete that there is no room left for an haecceitistic difference between worlds, despite each individual having a metaphysically primitive (non-qualitative) thisness. So one can be an haecceitist in the first sense without being an haecceitist in the second sense, but not vice-versa.

5. HUGGETT ON HAECCEITISTIC NEUTRALITY

We are now in a good position to evaluate Huggett's (1999) argument that the MB distribution must be seen as metaphysically neutral, and classical statistical mechanics therefore devoid of substantial metaphysical attachments, because different representations of it are subject to opposing metaphysical interpretations. In particular, he asserts that the phase-space representation is haecceitistic while the distribution-space representation is anti-haecceitistic. As I will argue, however, the metaphysical ambiguity advanced by Huggett as a basis for eschewing metaphysical commitments in MB statistics does not achieve this goal. The particle representation (phase-space) does not necessitate haecceitism any more than the occupation representation (distribution-space) necessitates anti-haecceitism. Instead, *both* representations presuppose classical criteria of material individuation, and therefore permit *every variety* of haecceitism and essentialism as a viable account of particle identity. Huggett also concludes from the purported ambiguity of classical statistics that quantum statistics will have nothing new to contribute to discussions of the metaphysics of individuality. This does not follow. Furthermore, when one moves to the quantum context, it is far from clear that the notion of an individual particle with its own properties makes sense, and insofar as it does not, the classical metaphysics of identity and individuation presupposed by MB statistics is rendered problematic.

Huggett begins his discussion by showing that it is possible to do MB statistics in the occupation number representation rather than the particle representation (Huggett 1999, 8–12). He calls these state spaces respectively the *distribution space* (Z-*space*) and the *phase space* (Γ-*space*). He contends that phase space, as the many particle state space formed from the sum of single particle state spaces (μ-*spaces*), has an intrinsically haecceitistic structure, while distribution space, which gives a description of the physical system in terms of occupation numbers for single particle states, "gives a complete account of the qualitative character of a world" without reference to specific individuals (1999, 12). This latter condition therefore renders distribution space inimical to haecceitism. Since both representations are possible and (supposedly) have contrary metaphysical import, MB statistics neither confirms nor disconfirms the classical metaphysics of individuality.

Huggett's conclusion that phase space has an haecceitistic structure while distribution space is anti-haecceitistic rests in part on an acceptance of David Lewis's definition and analysis of those terms, and in part on some argumentation which seems to fall short of its goal. We saw earlier that Lewis's treatment of haecceitism needs to be qualified. Along with Lewis, Huggett understands haecceitism as the denial of a supervenience thesis. Leibniz maintained that representation *de re* supervenes on the qualitative character of worlds, so that two worlds differing in representation *de re* always differ qualitatively. This is the Principle of the Identity of Indiscernibles (*PII*). If one denies this supervenience, i.e., denies *PII*, one can maintain that it is possible for there to be *de re* differences among worlds without any difference in qualitative character. To assert this latter position is, by Lewis' definition, to be a haecceitist. But haecceitism can also be understood as the assertion that there are *haecceities*, non-qualitative properties (properties that don't involve having any qualities) that differentiate individuals. We showed that it is *impossible* to be an haecceitist in Lewis's sense without also accepting the existence of haecceities (though the converse is possible). Huggett's endorsement of Lewis's position leads to an explicit avowal of haecceities as "murky notion(s)", and the assertion that "haecceitism is neither necessary nor sufficient for haecceities" (1999, 8). But haecceitism is *sufficient* for haecceities, and this in turn has a deleterious effect on Huggett's arguments.

Keeping this in mind, Huggett's argument for the failure of haecceitism in distribution space runs like this (1999, 10–13): There is a many-one mapping of phase space descriptions into distribution space. Distribution space, however, gives a complete account of the qualitative character of a world because any two states with the same distributions are qualitatively identical. The reason these states are qualitatively identical is that there are no "natural properties" (read: empirical consequences) in the physical system dependent upon a non-qualitative delineation of identities for individual particles. Haecceitism therefore fails in Z-spaces "if we assume that the Identity of Indiscernibles applies between states" (1999, 13). A question seems to be begged at this point because the relevant definition of haecceitism just *is* the denial of the Identity of Indiscernibles. Nonetheless, the argument proceeds from this supposition to the contention that since there are multiple phase space descriptions for each qualitatively complete description in distribution space, phase space descriptions must recognize non-qualitative differences among worlds. The doctrine of haecceitism therefore holds in Γ-spaces, but fails in Z-spaces (1999, 13).

There are some problems with this view. The first question we need to ask is whether the phase-space representation is unequivocally haecceitistic in structure. Haecceitism, while certainly *compatible* with phase-space, is *not* a necessary metaphysical accompaniment of it. In classical physics, all the particles comprising a physical system have well-defined spatio-temporal trajectories quite independently of our ability to determine them. Furthermore, I take it that spatio-temporal location is a uniquely individuating *qualitative* property for classical particles because such particles do not have superposed states and are impenetrable, hence incapable of occupying the same location at the same time. So even if one embraces *PII*, which amounts to a denial of haecceitism on the Lewis–Huggett construal of that doctrine, recourse to unique spatio-temporal properties (known or unknown) provides a basis for distinguishing individual particles in phase-space. Thus phase-space is *not* intrinsically haecceitistic, though it requires the classical individuative criteria metaphysically grounding all logically coherent doctrines of haecceitism or essentialism.

More importantly, *these remarks apply equally well to distribution space!* States with the same distribution are "qualitatively identical" only in the sense that statistical properties of the physical system are insensitive to differences among the individual particles comprising it. That the particles are indistinguishable with respect to the probability distribution describing their group behavior is hardly surprising – it is the purpose of statistical descriptions to distill systemic trends rather than get lost in the details of individual behavior. To reiterate a familiar point in slightly new vocabulary, a classical ignorance model suffices as the basis for a Z-space representation in the MB-context. *But this means that the Z-space description is anything but qualitatively complete!*

I suspect that the reason Huggett accepts the conclusion that distribution-space is inimical to haecceitism resides in a tempting analogy between the occupation representation and Lewis' characterization of (anti-)haecceitism. In the occupation representation, the system is modeled by which states are occupied, and by how many particles, but is insensitive to the question of which particles are in which state. In Lewis' definition of anti-haecceitism, a world is modelled by which qualities are instantiated, and by how many individuals, but (owing to *PII*) the world is insensitive to *de re* differences which are not manifested qualitatively (in fact the claim is that such differences do not exist). The conclusion that distribution-space is anti-haecceitistic is generated by equating physical systems with worlds, particles with individuals, and states with qualities.

As we noted earlier, assuming that *PII* holds between states in Z-space and these state-descriptions are complete begs the question against haecceitism's compatibility with distribution-space representations, and this question-begging is compounded if one does not recognize that Lewis's haecceitism actually *does* entail the existence of haecceities. One cannot speak of haecceitism in terms of worlds and the qualities instantiated within them without thereby introducing primitive thisness as a non-qualitative property of the individuals contained in those worlds. But if there are non-qualitative properties of primitive thisness, it goes without saying that worlds can differ *de re* while being qualitatively identical. Because of this, even if the occupation representation provided a complete qualitative description of a physical system in classical physics (which it does not), there could still be *de re* differences to which the representation was insensitive. In short, haecceitism is compatible with the distribution space formulation of MB-statistics. Furthermore, since the possibility of this formulation is predicated on the basis of classical ignorance, classical individuative criteria apply to the individuals subsumed by the distribution, and Z-space is therefore compatible with whatever coherent doctrines of haecceitism or essentialism one might care to espouse.

The important point to walk away with here is that both the phase-space and distribution-space representations in MB statistics presuppose classical individuative criteria for particles. It is precisely these criteria that give *de re* modality a foothold in physical theory, and it is just these criteria that quantum statistics problematizes. Huggett's conclusion (1999, 23–24) is therefore problematic:

> Our analysis has been directed at showing that there are no very heavy metaphysical implications of classical physics, and that therefore anticipated innovations in the notion of an individual in quantum mechanics will not be innovative at all.

It should be clear that this is not the case for classical statistics given that it presupposes classical individuative criteria, and even if it were, the conclusion about quantum statistics would not follow.

There remains one issue to consider in a bit more detail than we have thus far, namely whether Gibbs paradox is rightly considered a harbinger of quantum statistics. Some of Huggett's observations are a help to us here. If his conclusion that distribution-space is anti-haecceitistic were accepted, it would be reasonable to suppose that experimental evidence for the superiority of Z-space over Γ-space would constitute empirical grounds for the rejection of haecceitism. Since Huggett wants to maintain that there are no important implications for the metaphysics of individuality that attach to classical physics, it becomes necessary for him to show that Gibbs paradox does not reveal the superiority of distribution-space and anti-haecceitism. A co-belligerence with him in this cause is possible, but the motivation must be different. Our purpose is to establish that Gibbs paradox poses no threat to classical individuality, and therefore is not properly understood as foreshadowing quantum statistics.

Huggett's characterization of the thermodynamic considerations in favor of distribution-space is as follows (1999, 20):

The entropy in statistical mechanics is defined by the logarithm of the number of available ... states:

$$(*) \qquad S \equiv k \cdot \log N(n_i),$$

where k is Boltzmann's constant. Now, statistical physics has empirical import because it is intended to explain thermodynamics. In this reduction, the statistical entropy, S, is identified (at least in equilibrium) with the thermodynamical entropy σ. The main evidence in favor of S_Z (obtained when.... $N = N_Z(n_i)$) is that σ is supposed to be *extensive* or *additive*. That is, if two equilibrium systems are in mutual equilibrium, then the entropy of their union is the sum of the entropy of the two parts: $\sigma_{1\cap 2} = \sigma_1 + \sigma_2$. Via the identification we expect for the statistically defined equilibrium entropies that $S_{1\cap 2} = S_1 + S_2$. One can straightforwardly check that this equality holds for S_Z, but not for S_Γ, obtained by substituting $N_\Gamma(n_i)$ for N in (*). The nonadditivity of S_Γ thus is supposed to rule out Γ-space in favor of the extensive S_Z. Further, observation of entropy-dependent properties shows that they are in agreement with S_Z. Hence we are supposed to accept Z-space and anti-haecceitism.

The evidence in question here is none other than Gibbs paradox, of course: if two systems with maximal thermodynamic entropy are in mutual equilibrium, then it is experimentally the case that the thermodynamic entropy of their union is the sum of their individual entropies, i.e., $\sigma = \sigma_1 + \sigma_2$. When the statistical entropy is considered, however, additivity (extensivity) is satisfied in distribution-space but not in phase-space. The suggestion is that this confirms S_Z over S_Γ, and therefore gives empirical evidence that Z-space is the correct description.

Huggett demurs, however, suggesting that the inference can be blocked in two ways. The first is to resist the inference from S_Z, taken as the correct expression for statistical entropy, to Z-space as the correct representation space. The second is to maintain that the thermodynamic evidence does not warrant the conclusion of S_Z's superiority, because extensivity is not required by the Laws of Thermodynamics. We can gloss the arguments straightforwardly (see Huggett, 1999, 20–23 for a more extended discussion).

The inference from S_Z to Z-space is blocked by noting that we have some leeway in the statistical definition of entropy. If S_Z is taken as the correct function for the entropy, then we can take its domain to be either Z-space or Γ-space. In Z-space, statistical entropy will be defined as $S_Z \equiv k \log W_Z$, and in Γ-space it will be defined as $S_Z \equiv k \log W_\Gamma/n!$. The point is that both phase-space and distribution-space serve equally well as the domain for statistical entropy, so even if we assume that S_Z is the correct expression, we need not conclude that distribution-space is the correct representation. Either space will support whatever definition of entropy we require.

Huggett's second point (1999, 21–22) is that extensivity can be regarded as optional for thermodynamics, which means we need not take S_Z to be the superior entropic expression. As he notes, this argument was first given by van Kampen (1984). We start with the assumption that the Laws of Thermodynamics (conservation of energy, increasing entropy, Nernst's heat theorem) imply all of the observable consequences necessary to the theory. The question van Kampen then asks is quite simple: is observable additivity of entropy entailed by the Laws of Thermodynamics? If not, then extensivity is not a necessary concomitant of observable thermodynamics. Thermodynamical entropy receives expression in the Second Law as $d\sigma = dQ/T$, where dQ is the increase in heat and T is the temperature. For the purposes of van Kampen's argument, the important point is that the Second Law defines entropy differences not as an absolute value, but as work done. It is for this reason that the Second Law does not entail additivity. Consider first the entropies of distinct and disjoint systems. Since the Law is only descriptive of single systems at different points in their evolution, it is silent on the issue of additivity in this instance. How about the case of distinct non-disjoint systems? While the progression from separate samples into one can be regarded as the evolution of a single system, van Kampen (1984, 305) maintains that to derive extensivity in this way would rely on an assumption of the entropic equality of identically prepared distinct systems. Although this may be a useful assumption, it is not a consequence of the Second Law, which deals only with work differences in the thermodynamical trajectory of single systems. So extensivity is not an observable consequence of the Laws of Thermodynamics, and S_Z need not be regarded therefore as the superior expression for statistical entropy.

The puzzlement putatively engendered by Gibbs paradox concerning the identity of classical particles thus is dissolved. There is nothing privileged about distribution-space in regard to the statistical characterization of entropy, and no more is implicated by the "paradox" than the *classical* indistinguishability of particles in certain MB contexts. Classical indistinguishability is explained by an ignorance model that presupposes classical criteria for material identification. Gibbs paradox therefore fails to adumbrate the inadequacy of these classical criteria in the context of quantum statistics. Furthermore, at the risk of flogging a dead horse, even if the demonstrably correct representation were in Z-space (and there is no such demonstration), any accompanying indistinguishability of the particles described would be attributable to an epistemic ignorance revelatory of the incomplete character of the description. No innovative ontological consequences pertaining to particle identity follow from this at all.

6. CONCLUDING REMARKS

Despite the arguments of a number of physicists and philosophers to the contrary, we have seen no good reason to suppose that Maxwell–Boltzmann statistics partakes of an indistinguishability of particles akin to that observed in quantum statistics. Such indistinguishability as there may be in the classical statistical context is fully explained by epistemic ignorance without remainder. Any suggestion that the "indistinguishability" of Maxwell–Boltzmann particles obviates the metaphysical puzzles of quantum statistics is therefore incorrect. Huggett's attempt to deflate metaphysical questions in the context of classical statistics by way of contrasting interpretations with supposedly equal epistemic support similarly was seen to fall short of this goal. Maxwell–Boltzmann particles are metaphysically unproblematic and pose no threat to *de re* modality by way of refusing it a foothold. By contrast, it seems that quantum statistics gives rise to correlations that engender genuine puzzlement about the underlying ontology, though detailed exploration of this claim must await another occasion.

ACKNOWLEDGEMENTS

My thanks to members of the Chicagoland Philosophy of Physics Forum for helpful comments on an earlier version of this material – especially Arthur Fine, Nick Huggett, David Malament and Matthew Frank. My thanks also to the Center for Philosophy of Religion at the University of Notre Dame and to Baylor University for financial support that made this work possible.

NOTES

1 This issue can be turned on its head by regarding the indistinguishability of classical particles in some contexts as obviating the peculiarity of the indistinguishability of BE and FD particles in all contexts. But this approach assumes that the epistemic ignorance which led to BE or FD distributions for MB particles properly explains the non-classical behavior described by quantum statistics, an option at best open to de Broglie–Bohm hidden variable theorists.

2 See the essays in Elena Castellani, ed. (1998) for an exploration of problems of identity and individuation in the quantum context.

3 The word "haecceitism" derives from the Latin *haecceitas* (thisness) invented by the medieval philosopher Duns Scotus (cf. also Robert Adams 1979, 6–7).

REFERENCES

Adams, R.: 1979, 'Primitive Thisness and Primitive Identity', *Journal of Philosophy* **76**, 5–26.

Castellani, E. (ed.): 1998, *Interpreting Bodies: Classical and Quantum Objects in Modern Physics*, Princeton University Press, Princeton.

Costantini, D.: 1979, 'The Relevance Quotient', *Erkenntnis* **14**, 149–157.

Costantini, D.: 1987, 'Symmetry and the Indistinguishability of Classical Particles', *Physics Letters A* **123**, 433–436.

Costantini, D., Galavotti, M. C., and Rosa, R.: 1982, 'A Rational Reconstruction of Elementary Particle Statistics', *Scientia* **117**, 151–159.

Costantini, D., Galavotti, M. C., and Rosa, R.: 1983, 'A Set of "Ground Hypotheses" for Elementary-Particle Statistics', *Il Nuovo Cimento* **74B**, 151–158.

Costantini, D. and Garibaldi, U.: 1988, 'Elementary Particle Statistics Reconsidered', *Physics Letters A* **134**, 161–164.

Costantini, D. and Garibaldi, U.: 1989, 'Classical and Quantum Statistics as Finite Random Processes', *Foundations of Physics* **19**, 743–755.

Fujita, S.: 1991, 'On the Indistinguishability of Classical Particles', *Foundations of Physics* **21**, 439–457.

Huggett, N.: 1999, 'Atomic Metaphysics', *Journal of Philosophy* **XCVI**, 5–24.

Jeffrey, R.: 1983, *The Logic of Decision* (2nd edition), University of Chicago Press, Chicago.

Jeffrey, R.: 1988, 'Conditioning, Kinematics, and Exchangeability', in B. Skyrms and W. Harper (eds), *Causation, Chance, and Credence*, Kluwer Academic Publishers, Dordrecht, pp. 221–255.

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

Kripke, S.: 1972, *Naming and Necessity*, Harvard University Press, Cambridge, MA.

Lewis, D. K.: 1986, *On the Plurality of Worlds*, Basil Blackwell, Oxford.

Parry, W. E. (ed.): 1984, *Theoretical Physics*, Pergamon Press, Oxford.

Plantinga, A.: 1974, *The Nature of Necessity*, Clarendon Press, Oxford.

Plantinga, A.: 1978, 'The Boethian Compromise', *American Philosophical Quarterly* **15**, 129–138.

Putnam, H.: 1977a, 'Is Semantics Possible?' in S. Schwartz (ed.), *Naming, Necessity and Natural Kinds*, Cornell University Press, Ithaca, pp. 102–118.

Putnam, H.: 1977b, 'Meaning and Reference', in S. Schwartz (ed.), *Naming, Necessity and Natural Kinds*, Cornell University Press, Ithaca, pp. 119–132.

van Fraassen, B. C.:1991, *Quantum Mechanics: An Empiricist View*, Clarendon Press, Oxford.

van Kampen, N. G.: 1984, 'The Gibbs Paradox', in Parry (1984), 303–312.

Zabell, S.: 1988, Symmetry and Its Discontents', in B. Skyrms and W. Harper (eds), *Causation, Chance, and Credence*, Kluwer Academic Publishers, Dordrecht, pp. 155–190.

History and Philosophy of Physics
Baylor University
P.O. Box 97130
Baylor University
Waco, TX 76798-7130
U.S.A.
E-mail: Bruce_Gordon@baylor.edu