Comment on "Mind the Gap: Boltzmannian versus Gibbsian Equilibrium"

Dustin Lazarovici*
Université de Lausanne, Faculté des Lettres,
Section de Philosophie, 1015 Lausanne, Switzerland

February 10, 2018

In a recent paper, Werndl and Frigg discuss the relationship between the Boltzmannian and Gibbsian framework of statistical mechanics, addressing in particular the question when equilibrium values calculated in both frameworks coincide. In this comment, I point out serious flaws in their work and try to put their results into proper context. I also clarify the concept of Boltzmann equilibrium, the status of the "Khinchin condition" and their connection to the law of large numbers.

1 Boltzmann vs. Gibbs

In a recent paper[1], Charlotte Werndl and Roman Frigg discuss the relationship between the Boltzmannian and Gibbsian framework of statistical mechanics, addressing in particular the question, if and when the two formulations yield equivalent predictions for equilibrium values of macroscopic observables.

In the Boltzmannian framework, the relevant set of macro-variables takes (approximately) constant values on the equilibrium region of phase space which are thus revealed by an appropriate measurement on a system in equilibrium (a system, that is, whose actual micro-state is in the equilibrium region). The role of the stationary phase space

*Dustin.Lazarovici@unil.ch

[1] C. Werndl and R. Frigg, "Mind the Gap: Boltzmannian versus Gibbsian Equilibrium". Forthcoming in *Philosophy of Science*, doi: 10.1086/694088.

measure is to establish that it is very likely to find the system in an equilibrium state. In the Gibbsian framework, equilibrium is a property of an *ensemble*, represented by a stationary distribution ρ on phase space, and it is often (though maybe somewhat carelessly) said that the prediction for a measurement of a macro-variable f on an individual ensemble system is given by the *phase average*

$$\langle f \rangle = \int f(x)\rho(x)\,\mathrm{d}x\,, \tag{1}$$

where x is the phase space variable. This quantity is also called the *ensemble average* or *expectation value* of f.

There are many situation in which the Gibbsian phase average agrees – within relevant error bounds – with the Boltzmannian equilibrium value. Werndl and Frigg mention a criterion which they call the "Khinchin condition" and which they characterize, very briefly, as the phase function having "small dispersion". In less technical terms, this is to say that the macro-variable f is approximately constant, except on a possible set of small measure. Yet another and most pertinent way to formulate the relevant condition – now from a Boltzmannian perspective – is to say that there exists a unique Boltzmann equilibrium whose corresponding macro-region exhausts almost the entire phase space volume. For then, the macro-variable f takes an (approximately) constant value – the Boltzmannian equilibrium value – on a set of measure close to 1 – the Boltzmannian equilibrium region. Hence, the phase integral (1) will be dominated by the Boltzmannian equilibrium value.

It is important to emphasize that, unless one considers a thermodynamic limit, "the Boltzmannian equilibrium value" refers in general to a small range of values of f (hence "approximately constant"). This kind of coarse-graining, often left implicit in physical discussions and routinely ignored by Werndl and Frigg, is crucial and of course legitimate, since the relevant measurements have limited accuracy. We will return to this point in section 3.

The existence of a dominant Boltzmann equilibrium is in fact the *generic* case in statistical mechanics, giving rise to the claim that Boltzmann and Gibbs make (in general) equivalent predictions for systems in the respective equilibria. If one agrees that the Boltzmannian formulation is the more fundamental one, this also explains *why* Gibbsian phase averaging yields (in general) accurate predictions for individual measurements.

On the other hand, there are special cases in which this "Khinchin condition" doesn't hold. For instance, in the two-dimensional Ising model without external field, it makes sense to speak of *two* Boltzmann equilibria below the critical temperature, corresponding to a positive or negative magnetization, respectively. The partition function and thus the Gibbsian distribution are however symmetric under a flip of all spins, hence yielding an average magnetization of zero. There is nothing inconsistent or mysterious about this fact, as long as one keeps in mind that the Gibbsian value refers, in the first place, to an ensemble average. In particular, it is well understood in the physical literature that the interesting conclusions about the Ising model cannot be drawn from such phase averages. Instead, one can for instance study phase transitions at the critical temperature by fixing either $+1$ or -1 boundary conditions (referring to the polarization of spins at the edge of the lattice) thus implicitly picking one of the two magnetization states.

In their paper, Werndl and Frigg mention the magnetization in the Ising model only briefly but present instead other (alleged) examples for a disagreement between Boltzmann and Gibbs, which are either based on inappropriate and physically irrelevant choices of macro-variables or a misidentification of the relevant macro-states. We will discuss one of their examples in the final section. Still, the basic point that the Khinchin condition, respectively the coincidence of Boltzmann and Gibbs equilibrium values, cannot be assumed *a priori* is a valid one.

On these grounds, Werndl and Frigg declare that is "[a]n important task of the foundations of SM [statistical mechanics] ... to classify under which conditions the two frameworks lead to the same results and under which conditions they do not" an go on to present a "new theorem specifying a set of conditions" under which Boltzmannian and Gibbsian are indeed equivalent. I quote their result in full:

> **Equilibrium Equivalence Theorem (EET).** Suppose that the system (X, T_t, μ_X) is composed of $N \geq 1$ constituents. That is, the state $x \in X$ is given by the N coordinates $x = (x_1, ..., x_N)$; $X = X_1 \times X_2 \ldots \times X_N$, where $X_i = X_{oc}$ for all $1 \leq i \leq N$ (X_{oc} is the one-constituent space). Let μ_X be the product measure $\mu_{X_1} \times \mu_{X_2} \ldots \times \mu_{X_N}$, where $\mu_{X_i} = \mu_{X_{oc}}$ is the measure on X_{oc}. Suppose that an observable κ is defined on the one-particle space X_{oc} and takes the values $\kappa_1, \ldots, \kappa_k$ with equal probability $1/k, k \leq N$. Suppose that the macro-variable K is the sum of the one-component observable, i.e. $K(x) = \sum_{i=1}^{N} \kappa(x_i)$. Then the value corresponding to the largest macro-region as well as the value obtained by phase space averaging is $\frac{N}{k}(\kappa_1 + \kappa_2 + \ldots \kappa_k)$.

2 The law of large numbers

What the authors don't seem to realize, and certainly don't acknowledge, is that the relevant result, under their assumptions, is a straight-forward application of the *law of* *large numbers* (LLN) – arguably the most basic theorem in all of probability theory. Since they assume a family of independent and identically distributed random variables, the (weak) law of large numbers yields

$$\mu\left(\left\{x : \left|\frac{1}{N}\sum_{i=1}^{N}\kappa(x_i) - \frac{1}{k}(\kappa_1 + \kappa_2 + \ldots + \kappa_k)\right| < \epsilon\right\}\right) \geq 1 - \frac{\sigma^2}{\epsilon^2 N}\,, \tag{2}$$

for any $\epsilon > 0$, where σ^2 is the variance of κ. For a proof, I refer to any textbook on probability theory.

As Werndl and Frigg consider an *extensive* macro-variable that grows with N, we can set $\epsilon = N^{-\delta}$ for $\delta \in [0, \frac{1}{2}]$ so that, in terms of K (and writing $\overline{K} := \frac{N}{k}(\kappa_1 + \kappa_2 + \ldots + \kappa_k)$ as a "phase average"), the LLN estimate becomes

$$\mu\left(\left\{x : \left|K(x) - \int K(x')\,\mathrm{d}\mu(x')\right| < N^{1-\delta}\right\}\right) \geq 1 - \frac{\sigma^2}{N^{1-2\delta}}\,. \tag{3}$$

Here, it is important to note that it's the *relative* deviation $\left|\frac{K-\overline{K}}{K}\right| \lesssim \frac{N^{1-\delta}}{N} = N^{-\delta}$ that becomes vanishingly small for large particle numbers. In particular, for a macroscopic system, we have $N \sim 10^{24}$, and setting $\delta = \frac{1}{3}$, we can conclude that K deviates from $\overline{K}$ by less than one millionth of a percent on a set of measure (approximately) $0,999999$.[2]

To emphasize: if the macro-variable K is extensive, i.e. of order N, fluctuations of order $N^{1-\delta}$ (with, let's say, $\delta = \frac{1}{3}$ as in the above estimate,) are negligibly small for most practical purposes. Another way to think about this is that the typical values of K determine the relevant *unites* in which the system is described macroscopically. For instance, if the macro-variable has the dimension of energy, the characteristic scale of the one-constituent variables κ_i may be *electronvolts* (eV), while the total energy K is measured in *joules* (J). Thus, fluctuations of a few billion or even trillion electronvolts barely even amount to a rounding error on the macroscopic scale ($1\text{eV} \approx 1.602 \times 10^{-19}\text{J}$).

Summing up in less technical terms, the weak law of large numbers states precisely that for large N (which is the relevant case in statistical mechanics), phase space is dominated by an "equilibrium region", on which the value of the macro-variable K is very close to the expectation value (= phase average). But this is nothing else than the Khinchin condition in the form explicated above, implying the equivalence of Boltzmannian and Gibbsian predictions under the present assumptions.

[2] A tacit assumption, generally made, is that the variance σ^2 of the *one-constituent variables* κ_i is of order 1. If σ is already extremely large, or somehow chosen to increase with N, the LLN may fail to provide relevant estimates. But while some of Werndl's and Frigg's arguments seem to hint in that direction, such examples are either completely artificial and unphysical, or actually indicate a violation of statistical independence, so that neither the LLN nor the EET are appropriate arguments.

This leads me to the following assessment of Werndl's and Frigg's "Ensemble Equivalance Theorem":

1. Contrary to what the authors claim in the paper, their sufficient conditions for the equivalence of Boltzmannian and Gibbsian equilibrium values are not different from the Khinchin condition. They have merely stated obvious sufficient conditions under which the Khinchin condition holds.

2. The relevant result can be straightforwardly generalized (maybe to a "Generalized Equilibrium Equivalence Theorem (GEET)"). In the form of eq. (3) it holds, in fact, for any sum of uncorrelated and identically distributed random variables with finite variance. This is to say that neither the uniform distribution, nor the discreteness of the random variables is required, and they don't even have to be defined on a one-constituent space. For a proof, I refer again to any textbook on probability theory.

3. Even if the conditions stated by Werndl and Frigg are thus relaxed, they are very strong and hardly apply to anything by highly idealized models. The authors don't seem to realize where the strength of their assumptions lies. They write:

   > The crucial assumptions of the theorem are (i) that the macro-variable is a sum of the observable on the constituent space and (ii) that the macro-variable on the constituent space corresponds to a partition with cells of equal probability. (p. 16)

   But assumption (i) is often justified, while the "equal probability" in (ii) plays no important role at all. The crucial and very strong assumption is the *statistical independence* of the constituent variables (their joint distribution factorizes!), which doesn't hold, in most relevant cases, due to interactions or conserved quantities. (It is precisely this loss of statistical independence that becomes significant in the Ising below the critical temperature.) Controlling correlations – or justifying in some other way a "law of large numbers" – is in fact the key issue in many of the hard problems of statistical mechanics.

4. The authors' insistence that "we need other conditions next to the Khinchin condition" is ill-motivated to begin with. They present examples (however ill-conceived) in which the Khinchin condition fails to obtain, but those are also examples in which they find the Boltzmannian and Gibbsian values to disagree. In fact, a good case can be made that the Khinchin condition – in the sense of "uniqueness and

dominance of the Boltzmann equilibrium" – is not only sufficient, but also necessary. Suppose we agree on a definition of "Boltzmann equilibrium" according to which 1) the phase space volume associated to an equilibrium region is many orders of magnitude larger than the phase space volume associated to any non-equilibrium region and 2) the phase space regions associated to all equilibrium states together exhausts almost the entire phase space volume. Then, a violation of the above condition means that there is either *no* Boltzmann equilibrium (hence no macro-value that is particularly likely to be observed on an individual trial) or that phase space is dominated by two or more equilibrium regions corresponding to significantly different macro-values – in which case the phase average will correspond to an average of these values rather than any one in particular (unless this average happens to equal one of the Boltzmannian equilibrium values).

However, instead of arguing this point in greater detail, a more concise and productive remark is maybe the following: If the Khinchin condition doesn't hold, it means that there's a high probability of finding macro-values that differ significantly from the Gibbsian phase average, so that the phase average, as a prediction for individual measurements, is highly dubious in the first place.[3]

In conclusion, even if one concurs with Werndl and Frigg that specifying new conditions for the equivalence of Boltzmannian and Gibbsian equilibrium values is a serious foundational problem in statistical mechanics (and I don't), their "Ensemble Equivalence Theorem" is no progress whatsoever. The sufficiency of their conditions is a mathematical commonplace, a simple application of the law of large numbers, and thus certainly known to the cognescenti. And while it may be legitimate for a philosopher of physics to present even the most basic results to a less knowledgable audience, if the aim is to clarify their meaning and put the results into proper context, I worry that the authors do quite the opposite.

3 Macro-states confused

I want to return to one of my initial remarks and discuss what I believe to be the fundamental misconception underlying the discussion of Werndl and Frigg.

It is important to note that while the LNN follows immediately from their assumptions, the "theorem" stated and proven by Werndl and Frigg is strictly speaking not a LLN statement, in that their Boltzmann equilibrium value and its equality to the phase

[3] It is common misconception that one measures "ergodic time-averages". This is wrong since ergodic time-scales are *much* too long (cf. Goldstein (2001)).

average is supposed to be *exact*. In particular, when they speak of "the largest macro-region", as in the formulation of the "Ensemble Equivalence Theorem", they refer to the subset of X on which K takes *precisely* the value $\frac{N}{k}(\kappa_1 + \kappa_2 + \ldots + \kappa_k)$.

In my opinion, this doesn't make their result more original from a mathematical point of view, but misses the point entirely. For one, the measure of that particular set goes to *zero* for large N. Even more basic, it is neither reasonable nor useful (except in the idealization of the thermodynamic limit) to define macro-states in terms of arbitrarily precise values of physical quantities. For these reasons, Werndl's and Frigg's "largest macro-region" has little to do with a Boltzmann equilibrium and little physical relevance in general.

The law of large numbers is in fact the paradigm that one *should* have in mind when one thinks about the Boltzmann equilibrium: there is a certain *range* of typical values for the relevant quantities. And the larger the particle number, the more weight (phase space measure or probability, if you will,) is concentrated on an ever smaller range of values around the mean.

Alternatively, one can think of the *central limit theorem* and imagine that, for large N, the distribution of $\frac{1}{N}K$ is approximately a normal distribution with variance $\sim \frac{1}{N}$. Werndl and Frigg are essentially pointing at the peak of the bell curve, saying: "this is the most likely value!" But this is missing the point. The relevant observation is that the weight of the distribution is concentrated within a few standard deviations of order $\frac{1}{\sqrt{N}}$. Since N is very large, such variations are tiny and empirically negligible on macroscopic scales. It is this range of typical and for all practical purposes indistinguishable values – not just the peak of the distribution – that one identifies with the Boltzmann equilibrium.[4]

To be clear: a "macro-variable", qua mathematical object, is usually some nice function of the microscopic variables – think for instance of the energy $H(q, p)$ as a function of the particles' positions and momenta in a canonical ensemble. But such a variable is in general too fine-grained to consider all its different values as "macroscopically distinct". A possible source of confusion is that the term "macro-variable" is used both synonymously with "random variable" (in a technical sense) and with "macroscopic observable" (in a conceptual one). If used in the latter sense, however, we must take it with a grain of salt, keeping in mind that the relevant macroscopic observations have a limited accuracy – both in practice and by definition.

I believe the reason why this point is not more regularly emphasized in the scien-

[4] The authors may believe that their macro-variable is already properly coarse-grained from partitioning the one-constituent space into cells. But this is a coarse-graining on the microscopic scale (order $\frac{K}{N}$), and thus *much* too fine.

tific literature, is that physicists are trained to think the "error bars" by default, while mathematicians (and some mathematically minded physicists) usually work in the thermodynamic limit of infinitely many degrees of freedom ($N \to \infty$) where the typical values become, indeed, sharp. (In view of the central limit theorem, think of a sequence of more and more peaked normal distributions converging to a delta-distribution.)

Another potential source of confusion is that the Gibbsian equilibrium value, if identified with the phase average (1), is a definite real number by definition. But here too, the common physicist – maybe only after being pressured on it – will find it obvious that this can hardly be read as an "infinitely precise" prediction of the theory. Instead, one can for instance compute the ensemble variance

$$(\Delta f)^2 := \int (f(x) - \langle f \rangle)^2 \rho(x) \, \mathrm{d}x \tag{4}$$

and identify the Gibbsian prediction, to be compared with the Boltzmannian one, with $\langle f \rangle \pm \Delta f$.[5]

Considering their recent work, Frigg's and Werndl's literal-mindedness about macro-values seems to be at the core of their misconceptions about Boltzmannian statistical mechanics in general and its relationship to the Gibbsian framework in particular. One of many such instances is their discussion of the "baker's gas" (a model for the ideal gas based on the baker's transformation), that the authors have used in multiple publications to demonstrate how the Boltzmann equilibrium fails to be dominant, i.e. to cover the great majority of phase space volume. To this end, they partition the one-particle phase space (corresponding to the unit square in the baker's model) into k cells and claim that the macro-state corresponding to the Boltzmann equilibrium is the "uniform distribution" for which each cells contains exactly $\frac{N}{k}$ particles. It is then easy to see that while the phase space volume associated with the uniform distribution is greater than the phase space volume associated with any other particular arrangement of particles, it fails to exhaust a majority of phase space volume for large N. In the present paper, the authors exploit this fact – amounting to an apparent failure of the Khinchin condition – by introducing a (very artificial) macro-variable, weighing the previously defined "macro-regions" in such a way that the Gibbsian phase average differs significantly from the Boltzmannian equilibrium value.

In such discussions, however, the authors simply fail to consider an appropriate coarse-graining into macro-states, repeating the mistakes of Lavis (2005), already pointed out by

[5]Note that our first formulation of the Khinchin condition was precisely that the dispersion Δf is small. If it isn't, the "Gibbsian equilibrium value" is not a useful or reliable predictions for individual measurements to begin with, cf. point 4 above.

Lazarovici and Reichert (2015). *Exact* uniform distributions, where each cell contains precisely $\frac{N}{k}$ particles, are very special configurations, their measure actually goes to zero for large N. But configurations for which the density of particles in each cells differs only slightly from $\frac{1}{k}$ are macroscopically indistinguishable and coarse-grain to the same macro-state. Otherwise, we would have to say, for instance, that a gas is "out of equilibrium" if the left-hand-side of the volume contains even a single particle more than the right-hand-side.[6]

More precisely, it is an elementary result in probability theory (part of the standard proof of the LLN, actually,) that the variance (= dispersion squared) in a sum of N independent random variables – such as the occupation number of each cell in an ideal gas model – is additive. This is to say, in particular, that typical fluctuations are of the order $\sqrt{N}$ and we will not have a dominant equilibrium region if the coarse-graining into macro-states is finer than that. This is a trivial mathematical fact, not a deep foundational issue.

In the present case, a dominant macro-state with respect to the considered uniform measure would be the one containing all configurations for which the relative number of particles in each cell is within $\frac{1}{k} \pm \frac{1}{\sqrt{N}}$ (I'm neglecting a k-dependent constant that would yield somewhat better bounds). Since $\frac{1}{\sqrt{N}}$ is a tiny number for macroscopic N, corresponding to density fluctuations of less than one tenth of a billionth of a percent, these configurations look *macroscopically* uniform and constitute the relevant Boltzmann equilibrium.[7]

Werndl and Frigg may maintain that their discussion is simply assuming a different notion of "macro-state". But if this is so, it is assuming a notion of "macro-state" that is clearly inappropriate, a notion that no one, who has understood the basics of probability theory and the scope of statistical mechanics, could have ever had in mind. Nevertheless, it seems to be a general methodology of Frigg and Werndl to define macro-states as it pleases them and expound on the problems they encounter in return. A charitable reading is that this approach is just taken from the standard playbook of analytical philosophy: arguing by means of counterexamples – however far-fetched – to uncover foundational issues and push for conceptual clarity. But even then I believe that the subject matter (and it's protagonists, some of which the authors like to criticize on a

[6]If the relevant macro-variable is not actually "occupation numbers", but whatever corresponds to the function that Werndl and Frigg introduce, the physical reasoning is slightly different but the need for appropriate coarse-graining remains the same.

[7]And what if we had *extremely* precise measurement devices? Well, then we would actually observe density fluctuations in an "ideal" gas; but this doesn't put into question any of the results or concepts of statistical mechanics. One possible way of speaking, if one is worried about terminology, would be to say that such observations cease to be "macroscopic" in the relevant sense.

regular basis,) deserve a little more respect, a little more consideration given to basic results, insights and intentions. No Boltzmannian – least of all Boltzmann himself – ever claimed that you could partition phase space in any arbitrary and physically irrelevant manner and end up with a dominant equilibrium state.

In general, I am rather skeptical of this *ad absurdum* approach to the foundations of statistical mechanics. In my opinion, it misses the point of the discipline, which is not an "axiomatic theory" but an effective framework for the description of complex systems that requires some degree of pragmatism and good physical sense. A crucial difference, to my mind, is that in an axiomatic theory, any counterexample can point to foundational issues, while in statistical mechanics, some counterexamples point merely to incompetent use.

References

Goldstein, S. (2001). Boltzmann's Approach to Statistical Mechanics. In J. Bricmont, D. Dürr, M. C. Galavotti, G. Ghirardi, F. Petruccione, and N. Zanghì (Eds.), *Chance in Physics: Foundations and Perspectives*, pp. 39–54. Berlin: Springer.

Lavis, D. A. (2005, June). Boltzmann and Gibbs: An attempted reconciliation. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics 36*(2), 245–273.

Lazarovici, D. and P. Reichert (2015). Typicality, irreversibility and the status of macroscopic laws. *Erkenntnis 80*(4), 689–716.

Werndl, C. and R. Frigg (2017). Mind the Gap: Boltzmannian versus Gibbsian Equilibrium. *Philosophy of Science 84*(5), 1289–1302.