On an Unbelievably Short Proof of the 2nd Law

Daniel Parker

Philosophy Department, Virginia Tech

Blacksburg, VA 24061

parkerdn@vt.edu

Abstract

This paper investigates Jaynes' "unbelievably short proof" of the 2nd law of thermodynamics. It assesses published criticisms of the proof and concludes that these criticisms miss the mark by demanding results that either import expectations of a proof not consistent with an information-theoretic approach, or would require assumptions not employed in the proof itself, as it looks only to establish a weaker conclusion. Finally, a weakness in the proof is identified and illustrated. This weakness stems from the fact the Jaynes' assumption of unitary evolution is too strong given his perspective, rather than too weak to provide the desired results.

Acknowledgements: I would like to thank Wayne Myrvold for his comments on a draft of this paper.

1. Introduction

The cornerstone of the Jaynesian approach to statistical mechanics is the claim that the probability distribution associated with the statistical mechanical ensemble be interpreted as epistemic.[1] More to the point, the approach claims that the probability distribution should not be thought of as an objective physical part of our ontology. Rather, it is generated by a privileged way of assigning a measure to one's uncertainty or ignorance regarding the exact microstate of an SM system given a set of known constraints by means of a generalised principle of indifference known as the Maximum Entropy Principle (MEP).

Briefly, the MEP prescribes an algorithm for the generation of probabilities in statistical mechanics through the use of a standard calculation by the technique of Lagrange multipliers, where one adopts the density matrix that both satisfies a set of thermodynamic constraints (interpreted as the phase averages of the density matrix) and also maximises the von Neumann (or Shannon-Weaver) entropy

$$S_I = -Tr\left[\rho \ln \rho\right] \qquad (S_I = -\sum_i p_i \ln p_i). \tag{1}$$

over the relevant space of states.

As Jaynes repeatedly emphasised, he saw the MEP as providing the foundation for making the best inferences possible given a state of incomplete information, and statistical mechanics (SM) is but a particular example of this method. As he characterised

---

[1] This paper is primarily formulated in terms of density matrices and unitary dynamics, though the paper could equivalently written in from the perspective of probability distributions and Hamiltonian dynamics.

it, the proper objective of SM is generate inferences as to how thermodynamic systems are likely to behave, given all the knowledge one possesses regarding their current and past states.

One aspect (and indeed a very important one) of any foundational account of SM includes an interpretation and a demonstration of the 2nd law: we look to provide a statement of how the 2nd law is to be understood within a given interpretation, and to show how the 2nd law follows from this interpretation (if it follows at all). Although a clear and unproblematic statement of the content of the 2nd law is problematic (Uffink 2001), one intuitive aspect of the content of the 2nd law is the claim that the entropy increases over time. In the Jaynesian vein, this is translated as the claim that given some thermodynamic constraints that initially generate an ensemble distribution and a particular thermodynamic entropy, it is a consequence of the unitary evolution of such systems that the predicted entropy based on the evolved thermodynamic observables cannot be less that the initial values. As Jaynes characterises it, the "real reason for the second law [is that] *it is a fundamental requirement on any reproducible process that the phase volume W' compatible with the final state cannot be less that the phase volume* W_0 *which describes our ability to reproduce the initial state*" (1965, 395 emphasis original).

This paper is an attempt to investigate Jaynes' attempt to characterise and prove an 'information-theoretic' account of the 2nd law, in the limited sense just described. The next section briefly presents the proof. Section 3 discusses and evaluates some common criticisms that the proof is too weak in its results, which I argue demand more out of the proof than Jaynes either desires or feels is necessary on his conception of the 2nd law. The final section argues that in fact the proof is in some sense too strong in its assumptions, as one is not immediately entitled to assume unitary evolution from the Jaynesian perspective, and absent this assumption the proof fails.

2. The Unbelievably Short Proof

This emphasis on the manipulability of physical systems is a major theme of Jaynes (1965). In this paper, Jaynes offers an "almost unbelievably short" proof of the 2nd law, derived based on his approach to SM. His proof is indeed very short, and discussion of the proof has focused on a more detailed exposition found in Lavis and Milligan (1985), which is itself cobbled together from Robertson (1966) and Hobson and Loomis (1968).[2]

The proof considers a canonical system undergoing an adiabatic change. Define an arbitrary set of observables $\{\Omega_1(t), \Omega_2(t), \Omega_3(t), \Omega_4(t)\ldots\Omega_m(t)\}$ and suppose we have made a set of measurements made at t_0, $\{\omega_1(t_0), \omega_2(t_0), \omega_3(t_0), \omega_4(t_0)\ldots\omega_m(t_0)\}$. On the basis of these values, we find the density matrix $\rho_0(t_0)$ that maximises the information-theoretic entropy $S_I^{(0)}(t_0) = -Tr\left[\hat{\rho}_0(t_0)\ln\hat{\rho}_0(t_0)\right]$ subject to the constraints

$$\omega_k(t_0) = Tr\left[\hat{\rho}_0(t_0)\hat{\Omega}(t_0)\right], k = 1\ldots m \tag{2}$$

and the thermodynamic entropy is given as a function of this density matrix, by

$$S_e^{(0)}(t_0) = -kTr\left[\hat{\rho}_0(t_0)\ln\hat{\rho}_0(t_0)\right].^{3} \tag{3}$$

---

[2] The presentation here employs the quantum mechanical formalism, though the classical case has an exact analogue (e.g. Frigg 2008 169-171).

[3] This is the standard prescription according to Jaynes' Maximum Entropy Principle.

Although the information-theoretic entropy and the thermodynamic entropy are conceptually independent, in this case they yield identical values up to Boltzmann's constant. If we allow the system to evolve, at some later time *t*, we *predict* the values of the observables {ω_1(t), ω_2(t), ω_3(t), ω_4(t), ω_5(t)…}, using

$$\omega_k(t) = Tr\left[\hat{\rho}_0(t)\hat{\Omega}(t)\right], k = 1 \ldots m \qquad (4)$$

Given these predicted observables, we define a new density matrix according to these new constraints, and subject to maximising the information-theoretic entropy (call it ρ(t)).

$$\omega_k(t) = Tr\left[\hat{\rho}(t)\hat{\Omega}(t)\right], k = 1 \ldots m \qquad (5)$$

The new thermodynamic entropy is calculated according to the new density operator for the predicted values of the observables:

$$S_e(t) = -kTr\left[\hat{\rho}(t)\ln\hat{\rho}(t)\right] \qquad (6)$$

Since

i. $S_e^{(0)}(t)$ is invariant under unitary evolution.

ii. both ρ_0(t) and ρ(t) satisfy the constraints provided by the values of the predicted observables.

iii. ρ(t), but not necessarily ρ_0(t), maximises the information-theoretic entropy for the predicted values of the observables.

It follows that $S_e^{(0)}(t_0) \leq S_e(t)$.

3. Criticism

It is fair to say that philosophers have been unimpressed by this proof (e.g. Earman 1986, Sklar 1993, Frigg 2008), and Earman has gone so far as to claim that he agrees that this proof is "almost unbelievably short" provided that one removes the 'almost'. Generally, there appear to be three main threads of criticism concerning the proof.

Dynamical details. First, it does not provide the kind of details one would expect from a proof of the 2nd law, insofar as it cannot supply the values of transport coefficients, relaxation times, etc. Taken on its face, this objection seems to fall flat. The proof does not purport to provide these details, and it is hard to fault the proof for failing to provide them as this is not its intent. This is not to say that these details are uninteresting, just that they are well beyond its intended scope. Given that all that is assumed is that the density matrix evolves unitarily, it would be hard to imagine the details of the evolution would follow from this fact alone.

Further, there are two ways to read Jaynes' proof as demonstrating the increase of entropy. On the weak reading of the proof, all Jaynes demonstrates is that the entropy at some arbitrary later time (given unitary dynamics) will be greater than or equal to the entropy at the initial time. If this is all that Jaynes intends to demonstrate, he seems to be successful. There is no expectation that one should take the entropic values generated by this proof seriously, in the sense that this reading of the proof does not claim to (nor can it) generate *specific predictions*, but only shows that the entropy must increase (or be equal to its initial value) at some later time. But this is not the entire content of the 2nd law, and the worry is that the too much is missing from the proof to recover a useful statement of the law. This would be the intent of a stronger proof of the 2nd law.

Monotonicity. Second, the proof fails to demonstrate that the entropy increases monotonically. Rather, the proof only asserts that the entropy value at later times will be greater than the initial value of the entropy based on the values of the observables at t_0; that is, although the proof asserts that for $t_1 < t_2$, $S(t_0) < S(t_1)$ and $S(t_0) < S(t_2)$, it does not show that $S(t_1) < S(t_2)$.

Piggybacking on the previous discussion, it is hard to see how this can be a problem for the proof itself since again, the only dynamical assumption that enters into the proof is that the density matrix evolves unitarily. Further, it is even unclear as to whether a monotonically increasing entropy curve is a necessary or even desirable feature of a proof of the 2nd law. Although monotonic increase is an intuitive desideratum for the behaviour of entropy, there exists no consensus as to how to define the thermodynamic entropy in non-equilibrium contexts. The fact that the entropy described in (6) does not monotonically increase should not necessarily be seen as a problem.

Moreover, there are experimentally realisable scenarios where the entropy curve does not monotonically increase (on an intuitive understanding of the entropy), but actually decreases. In the spin-echo experiments, where the entropy is associated with the alignment of the spins, the entropy increases until a magnetic pulse is applied to the system, and the entropy begins to decrease back to its initial value, in an apparent violation of the monotonic increase of the entropy. In such a case, Jaynes' proof captures the relevant features of this phenomenon rather nicely, since at all times the entropy is greater than or equal to the initial entropy, though it does not increase monotonically.

Ridderbos (2002) and Ridderbos and Redhead (1998) suggest that the spin-echo experiments provide an argument for the correctness of the fine-grained Gibbs entropy, roughly since its constancy ensures that there is no spontaneous decrease in entropy as the spin states return to their original orientation. In the Jaynesian case, we have two density matrices to work with, ρ and ρ_o, the latter describing the initial density matrix and the former indicating the distribution that would be used if it were constructed from the predicted values of the observables at a later time. Clearly, the behaviour of the matrix ρ_o, being restricted to unitary evolutions, mirrors the relevant features of the Gibbsian fine-grained entropy to which Ridderbos and Redhead appeal. In the present context, this adds force to the undesirability of a proof that strictly requires monotonic entropic increase. Conversely, the predicted matrix identifies the entropy that would be generated had the matrix been generated only from contemporaneous predicted constraints, i.e. if one's knowledge of the state of the system were limited to the values of those predicted constraints.

Sklar (1993) exploits this distinction. He argues that the spin-echo experiments are actually problematic for the Jaynesian because the proof seems to suggest that we get an increase in entropy through the use of the matrix ρ by purposely throwing away information about the history of the system (contrary to Jaynes' repeated assertions that we should never throw relevant information away), information that is crucial to providing a correct description of the spin-echo behaviour. This seems in tension with the proof of the 2^{nd} law, since it suggests that there may be contexts where we might want to use the matrix ρ_o and those where it is preferable to use ρ to describe the system, depending on our objectives. A partial solution to this worry is to maintain a conceptual separation between the experimental, thermodynamic entropy, and the information-theoretic entropy associated with one's knowledge of the system, the latter always being described by the evolution of the matrix ρ_o, whereas the former is at least partially defined by the subjunctive reading in the paragraph above.

Time Step Dependence. Finally, a third criticism attempts to block an obvious remedy to the problem of the non-monotonic behaviour of the entropy curve (on the assumption that this is a genuine problem). Lavis and Milligan (1985) suggest that one could recover monotonic behaviour if, instead of comparing the entropy at later times with the entropy at the time of the initial measurements, we compare the entropy at t_0 to t_1, and then repeat the proof with t_1 taken as the initial time so that it ensures $S(t_1) < S(t_2)$. In Frigg's words, the problem with this approach is that "it would have the disadvantage that the entropy curve would become dependent on the sequence of instants of time chosen. This seems odd even from a radically subjectivist point of view: why should the value of S_e at a particular instant of time, depend on earlier instants of time at which we chose to make predictions, or worse, why should it depend on us having made any predictions at all?" (2008, 172) Lavis and Milligan argue that this rescue, despite these problems, appears unmotivated from the Jaynesian perspective, since it is assumed that the values of the observables at t_0 are privileged, in that this is the only time that the values are actually measured on the system itself, rather than 'merely inferred'.

Here again, it is striking that such expectations are being placed on Jaynes' proof. Again, it is not clear that the monotonic increase in entropy is a desideratum for the proof, and without any dynamical assumptions in place it is hard to see how one might expect monotonicity to be shown.[4] Further, Frigg, Lavis and Milligan are right to note that privileging certain times seems unmotivated from the Jaynesian perspective. Fortunately, as described above, monotonicity is neither a desired result nor to be expected from a proof that only relies on the unitary nature of the evolution.

Another point of concern in Frigg's evaluation of the proof is that the time dependence of this move only seems odd if we do identify the entropy predicted at later times with the *thermodynamic* entropy, which is thought to be a property of the physical system. Indeed, Frigg points to the strange possibility that making predictions, or even *not* making predictions, might have an effect on the system's entropy, which is indeed an odd claim. It would be strange for the entropy curve to depend crucially on the seemingly arbitrary times at which we make predictions about the system.

But is it the case that the entropy that figures into this proof really is the thermodynamic entropy? Although Jaynes seems to identify the *predicted* entropy with the thermodynamic experimental entropy, there is no demonstration that this is actually the case, and the entropy S_I is treated as conceptually independent in the proof. Indeed, whether the TD entropy does match the predicted entropy depends on the details of the dynamics: one should not expect detailed and correct entropy curve from the assumption of unitary evolution alone. Indeed, the immediate inference from the information-theoretic entropy to the experimental thermodynamic entropy is not licensed by Jaynes in any way. Jaynes (1963) himself argues that the information-theoretic entropy "*for some*

[4] In fact, Jaynes' original presentation of the proof only looks to establish that the final, equilibrium entropy is greater than the initial entropy, without considering any intermediate values.

distributions and *in some physical situations*, has long been recognised as representing entropy. However, we have to emphasise that the "information-theory entropy" S_I and the experimental thermodynamic entropy S_e are entirely different concepts. Our job cannot be to *postulate* any relation between them; it is rather to *deduce* whatever relations we can from known mathematical and physical facts." (187, emphasis original) Absent some set of relevant "mathematical and physical facts", such as the details of the dynamical evolution, there is no reason to expect that the two entropies should match. As such, (6) is ambiguous: the predicted entropy S_e might refer to one of two quantities. Either we take this quantity to be the actual predicted entropy given by a precise quantity, or as the weaker statement that *whatever* dynamics are in place, the expected experimental entropy will be greater than the initial entropy, without taking the proof to generate a specific numerical value for either the expectation values or the entropy. Given the analysis above, it is clear that the latter reading is to be preferred.[5] Nonetheless, in the next section we shall see that even this reading fails.

In sum, the expectations placed on the proof seem unreasonable, and they are fairly easy to diagnose. Instead of amounting to a critique of the information-theoretic approach on its own terms, these criticisms amount to a list of desiderata associated with interpreting SM as a proper physical theory, and demanding that the proof track, in all its gory detail, the specifics of the *physical* behaviour of the system. But the proof purports

[5] Spelling out the exact relation between these two entropies would surely be a difficult task that cannot be endeavoured here. However, I at least suggest that one component of this relation is given by the subjunctive, counterfactual, reading of the thermodynamic entropy mentioned in discussing Sklar.

to do no such thing, and interpreted as reflecting the *physical* necessity of the 2^{nd} law, Jaynes' proof is indeed wanting in numerous respects. Insofar as the proof only assumes unitary evolution and nothing more, it is hard to see how one could hope for a stronger result than the one proven: the premises of the proof are far too weak to give us relaxation times or a monotonic entropy curve. Without a specification of the dynamics, what more could be expected out of this proof than what has been given?

However, when interpreted as a statement regarding the manipulability of such systems based on the values of the observables at some initial time, the proof takes a somewhat different light. If these probabilities are interpreted as epistemic rather than as a physical probability distribution (whatever the relation between these might be), the object of the proof is not to describe some physical necessity, but to establish a restriction on one's ability to make inferences about and control the evolution of the system, given only the knowledge that the dynamics are unitary. Jaynes' objective in this proof is limited to this: there is no reproducible or controllable way that the values of the thermodynamic observables can change adiabatically such that the entropy associated with the predicted density matrix will decrease below its initial value.

4. Why Assume Unitarity?

This does not mean that all is well with Jaynes' proof. However, the qualm I have with the proof is not that the proof is too weak in its conclusions, but that its fundamental assumption is too strong. Jaynes' programme is intended as a framework in which one makes inferences about the future state of thermodynamic systems, based on the results of an initial set of measurements. And although we can expect the evolution of the probability distribution, construed physically, to evolve unitarily, it is unclear why the evolution of an *epistemic* probability distribution, insofar as it describes the evolution of our knowledge, should evolve in this way. Indeed, it would be odd for this distribution to evolve unitarily, mirroring the physical evolution (unless we were Laplacian demons). Although the proof assumes unitary evolution, and many of the criticisms discussed above hit on this assumption as being too weak to deliver the results desired, this assumption in fact seems unjustifiably strong.

The natural counterpoint here is to defend Jaynes' proof by observing that there is a difference between *actually being able to evolve a probability distribution through its unitary evolution*, such that for any initial distribution one can actually describe its time-evolved state[6], and the weaker claim that *we know* that the distribution, *however* it evolves, evolves unitarily. And in fact, nothing is assumed in the proof other than that the distribution evolves in a unitary fashion. So one should read the proof as not demanding a sort of Laplacian omniscience, but merely the more general observation that *we know* evolutions are unitary. Clearly, this is in the spirit of the Jaynesian approach, and perhaps this would vindicate the proof.

However, this move will not work. Even though we assume the fundamental dynamical evolution to be unitary, this is not actually something one is entitled to assume in this proof. The assumption of unitary evolution is subject to an important qualification:

---

[6] In principle if we knew how to evolve the probability distribution exactly, we would be able to demonstrate the monotonic increase in entropy (presuming it is a fact), determine relaxation times, etc. Indeed, much research and many open questions in SM would be obviated if this were feasible.

the degrees of freedom that comprise the representation of the system must be complete. If there are additional physical degrees of freedom associated with the system beyond those that are specified, it is possible that the evolution relative to the specified known degrees of freedom will not appear unitary.

From an 'objective' perspective, this hardly matters: unitary evolution is a fundamental feature of the dynamics, and if one fails to incorporate some degrees of freedom into the description of the system, then the description is wrong. However, from the Jaynesian perspective, one is only permitted to assume unitarity if one is justified in believing that all the physical degrees are represented in one's description of the system, since it is our goal to make inferences based only on the knowledge in hand.

To see how this is problematic, as a toy example we consider a discrete phase space taking on six possible microstates, and loosely based on Jaynes' (1963) Brandeis dice problem. Initially, we are given an expectation value across the states of 4.5, and for the moment can be thought as representing the average value of the pips on a die. Jaynes (1979) solves this problem through the MEP method, attaining the following probability distribution:

$$p(1)=0.054,\ p(2)=0.079,\ p(3)=0.114,\ p(4)=0.166,\ p(5)=0.24,\ p(6)=0.348 \quad (7)$$

According to Jaynes, this distribution represents the maximal entropic state for the constraints given, and the maximal entropic state for the system, absent the constraint that the expectation value is 4.5, has an expectation value of 3.5 where each outcome is assigned equal probability. Clearly, the further the expectation value from 3.5, the lower the maximum entropy assigned to the system, such that

$$S(x)>S(y) \text{ iff } |x-3.5|<|y-3.5| \quad (8)$$

where x and y are the expectation values associated with some macrostate of the system.

Now, let us specify that the system evolves by some unitary dynamics, such that for each discrete time step, a new state is achieved by permuting the probabilities of each individual microstate in some definite but unspecified way. Clearly, the informational entropy due to such dynamical evolution is unchanged, since this amounts to permuting the indices on outcomes in the Shannon-Weaver entropy. However, the expectation value may change through arbitrary permutations. Because the probabilities assigned by the MEP are monotonically increasing when $x>3.5$ and monotonically decreasing when $x<3.5$, the expectation value of the system is bounded under evolution since the permutations can only move the expectation value closer to 3.5; that is, if the initial expectation value was 4.5, the dynamics can only generate probability distributions with expectation values between 2.5 and 4.5. Thus if one constructs a new probability distribution on the basis of these new expectation values, the entropy associated with the new distribution must be greater than the entropy of the original distribution. This is the essence of Jaynes' proof.

We can see the relevant features of the proof even in this toy model. Although the initial entropy represents a lower bound on the entropy, it is entirely possible that the unspecified unitary dynamics move the system into states which move closer and further away from the mean value of 3.5, and thus the entropy cannot, by this method, be shown to increase monotonically with each time step.

Despite the above arguments that usual criticisms of this 'almost unbelievably short' proof are misguided, there is a sense in which one needs a more robust characterisation of the dynamics beyond the stipulation that they are unitary. The proof works because there is no way in which the expectation values can obtain values outside the bound described above, and this is achieved by limiting the dynamics to permutations of the initial probability distribution. However, if the state space is incompletely specified, in the sense that there are further degrees of freedom not accounted for in the description of the system, it is possible for the entropy relative to this description to decrease.

For instance, suppose that there are additional dofs not accounted for in the phase space, such that (according to some 'natural' measure or reproducible physical distribution) the probability of each outcome from 1-6 is actually the marginal sum over a more fine-grained joint distribution where each outcome is actually associated with two outcomes of unequal weight. If these probabilities are permuted by some unitary dynamics, it is entirely possible for the expectation values (from the coarse-grained perspective) to, say, exceed 4.5, since it is with respect to the coarse-grained p.d. that the initial p.d. is generated. For instance, suppose that each microstate can be bifurcated into two states with different densities, such that according to the 'natural' measure $p_1(1)=p_2(1)$, with $p_2(y)=(1/y^2)p(y)$ and $p_1(y)=p(y)-p_2(y)$ for $y\in\{2, 3, 4, 5, 6\}$. The resultant probabilities for this fine-grained state are as follows:

	$P_1(y)$	$P_2(y)$
$P_x(1)$	0.02718	0.02718
$P_x(2)$	0.05908	0.01969
$P_x(3)$	0.10148	0.01268
$P_x(4)$	0.15511	0.01034
$P_x(5)$	0.23018	0.00959
$P_x(6)$	0.33784	0.00965

If the actual dynamics are given by the rule that for each time step, the states in the left hand column move to the right hand column, and states in the right hand column move to the left hand column below it in a cyclical fashion such that $p_2(6) \rightarrow p_1(1)$, then one computes the expectation value after one time step to be 4.53, which has a lower entropy than the initial state from the perspective of the coarse-grained description where the only observable defined is Y. Such examples are easy to multiply.

This point is similar to criticisms of Jaynes' programme that suggest that he tacitly relies on systems being ergodic in order to get his programme off the ground (Sklar 1993, Frigg 2008). In the case of ergodicity, the worry is that if the space is not metrically indecomposable, the phase averages associated with the density matrix may not equal the time averages, so that the Jaynesian method will deliver the wrong predictions if it cannot demonstrate the system to be ergodic. However, this worry is different, and perhaps more fundamental, for in the case of ergodic theory the dynamics are assumed to be measure preserving. In this case, if there is no guarantee of unitary evolution because relevant degrees of freedom are missed, one cannot even assume measure preservation. Nonetheless, the best response available to the Jaynesian is likely parallel to the one offered against the ergodic criticism: provided that one is sure the dynamics are unitary, the appearance of apparent entropy decreasing phenomena is an indication that important physical features of the system have been omitted and need to be found or included. As such, it should not be understood as an indictment of the Jaynesian progamme itself. The reader is left to decide whether this response is convincing.

5. Conclusion

Some common criticisms of Jayne's unbelievably short proof have been discussed. Generally, these criticisms have pointed to the fact that the proof is too weak in its conclusions, and any attempt to strengthen them would be unjustified from the Jaynesian perspective. Conversely, I have argued that the demands placed upon this proof are too stringent, and in fact demand more from a proof of the 2nd law than Jaynes desires or needs to show. However, I have suggested that there is a sense in which the proof is actually too strong, rather than too weak. Specifically, unless the dynamical description of the system is known to be complete, one is not entitled to assume the unitary evolution of the system relative to the description given. Without the assumption of unitarity, it is possible for the entropy to decrease, even if the fundamental dynamics are unitary.

References

Earman, J. (1986). “The Problem of Irreversibility”, *Proceedings of the Philosophy of Science Association*, Vol. 2, 226-233.

Frigg, R. (2008). “A Field Guide to Recent Work on the Foundations of Statistical Mechanics”, in *Ashgate Companion to Philosophy of Physics*, ed. D. Rickles. Aldershot: Ashgate.

Hobson A. & D. N. Loomis (1968), “Exact Classical Nonequilibrium Statistical-Mechanical Analysis of the Finite Ideal Gas”, *Physical Review*, **173**, 285-295.

Jaynes, E. T. (1963), "Information theory and statistical mechanics." In *1962 Brandeis Summer Institute in Theoretical Physics,* ed. K. Ford. New York: Benjamin. Reprinted in Jaynes (1983)

Jaynes, E. T. (1965). “Gibbs vs. Boltzmann Entropies” in *Papers on Probability, Statistics and Statistical Physics*. (1983) R. D. Rosenkrantz ed. Dordrecht: Reidel, 391-398.

Jaynes, E. T. (1979). Where Do We Stand on Maximium Entropy? *Papers on Probability, Statistics and Statistical Physics*. R. D. Rosenkrantz. Dordrecht, Reidel.

Jaynes, E. T. (I983). Papers in Probability, Statistics and Statistical Physics. R. D. Rosenkrantz, ed. Dordrecht, Reidel.

Jaynes, E. T. (1992). The Gibbs Paradox. Maximum-Entropy and Bayesian Methods. P. N. G. Erickson, and C. R. Smith. Dordrecht, Kluwer**:** 1-22.

Lavis, D. & P. Milligan (1985) “Essay Review of Jaynes’ Collected Papers”, *British Journal for the Philosophy of Science*, 36:193-210.

Ridderbos, K. (2002). "The Coarse-Graining Approach to Statistical Mechanics: How Blissful is our Ignorance?" *Studies in History and Philosophy of Modern Physics* **33**: 65-77.

Ridderbos, T. M. and M. L. G. Redhead (1998). "The Spin Echo Experiments and the Second Law of Thermodynamics." *Foundations of Physics*, 28:1237-1270.

Robertson, B. (1966), "Equations of Motion in Nonequilibrium Statistical Mechanics", *Physical Review*, 144:151-161.

Sklar, L. (1993). Physics and Chance. Cambridge: Cambridge University Press.

Uffink, J. (2001). "Bluff Your Way in the Second Law of Thermodynamics", *Studies in History and Philosophy of Modern Physics*, 32:305-394.