

NOTES ON THERMODYNAMIC FORMALISM

DANNY CALEGARI

ABSTRACT. These are introductory notes on the thermodynamical formalism, with application to topics in dynamics and geometric group theory. These notes are based on a graduate course co-taught with Amie Wilkinson given at the University of Chicago in Fall 2014.

CONTENTS

1. Classical thermodynamics	1
2. Statistical mechanics	13
3. Shannon's theory of information	29
Appendix A. Classical and quantum mechanics	36
Appendix B. Method of steepest descent and the Darwin-Fowler theorem	43
Appendix C. Large deviations and Chernoff inequality for Markov chains	46
Acknowledgments	46
References	46

1. CLASSICAL THERMODYNAMICS

“Today, it is not only that our kings do not know mathematics, but our philosophers do not know mathematics and — to go a step further — our mathematicians do not know mathematics.”

– J. Robert Oppenheimer, Harper's Magazine, Volume 217, October 1958

Classical thermodynamics is concerned with properties of macroscopic bodies in equilibrium. It ignores the microscopic details of the system, and concerns itself with aggregate properties in situations where these properties make sense. It is not a dynamical theory, in the sense that time does not occur as a variable in thermodynamic equations.

Our main purpose in this section is to introduce the basic state variables and the relations between them, and to derive some of the more important identities and inequalities, which we will revisit when we come to discuss the theory of statistical mechanics in § 2. A basic reference for the material in this section is Pippard [9]. Historical sources include Gibbs [5].

1.0.1. *Equilibrium, reversible and quasistatic processes.* An isolated system is in *equilibrium* if its macroscopic properties do not change with time. This is an idealization in several important respects, since in the first place no real systems can ever be truly isolated; and in the second place, we shall see that these macroscopic properties (when described ultimately in terms of microscopic foundations) are *statistical* in nature, and are subject to detectable fluctuations over sufficiently long time intervals.

Interactions between systems are said to be *reversible* if they proceed infinitesimally slowly as a result of infinitesimal differences in temperature or pressure. Thus an infinitesimal change in the temperature or pressure of the systems would cause the interactions to occur in the opposite direction. Interactions between systems are *quasistatic* if they occur so slowly that the system can be assumed to be in equilibrium throughout the process. An example of a quasistatic interaction which is not reversible is one in which work done is dissipated (albeit infinitely slowly) through friction. A system may be assumed to be in equilibrium throughout a reversible or quasistatic process.

We are concerned with properties of the system that make sense when the system is in equilibrium. These properties are often expressed as (real) values of certain *functions* of the equilibrium state; i.e. so-called *state functions*. Examples of state functions are pressure, volume, temperature, etc. During an interaction which is not quasistatic, the system may go temporarily out of equilibrium, at which point it may not have a well-defined pressure or temperature (for example). However, some properties of the system — for example, the total internal energy U — make sense even for a system out of equilibrium.

More abstractly, we can imagine that the collection of all possible states of the system is some (very high, or infinite dimensional) manifold, and the subset of equilibrium states is a (typically low-dimensional) submanifold. Some functions (energy, volume) are well-defined everywhere; other functions (temperature, entropy) — the *state functions* — are defined only on the submanifold of equilibrium states.

1.0.2. *Temperature.* When two systems are brought into contact with each other, their individual states might change, as (for example) *heat* flows from one to the other. The systems (individually in equilibrium) are said to be in *equilibrium with each other* if no heat flows between them when they are brought into contact, in which case they are said to have the same *temperature*. The *zeroth law* of thermodynamics says that temperature is well-defined as a state function; i.e. that the property of having the same temperature is an equivalence relation; thus the temperature T of a system in equilibrium is the equivalence class that it represents. A quasistatic transformation of a system is *isothermal* if temperature remains constant throughout. Note that temperature does not make sense for systems that are not in equilibrium.

1.0.3. *Energy.* Energy is transferred between systems in two forms — “ordered” energy in the form of *work*, and “disordered” energy in the form of *heat*. A transformation of a system is *adiabatic* if no heat is transferred. The *first law* of thermodynamics says that energy is a conserved quantity. Note that heat and work are not themselves functions of the state of a system; they are measures of the amount of energy that is transferred between systems in different forms. The infinitesimal heat and work transfer to a system can be thought of as 1-forms on the space of states of the system; these 1-forms are *not* exact (or even closed)

in general. In the physics literature they are denoted δQ and δW , even though Q and W do not really exist as functions. The internal energy is denoted U ; thus $dU = \delta Q + \delta W$. Note that W here is the net energy transferred to the system in the form of work; i.e. it is the difference between the work done *to* the system, and the work done *by* the system. Internal energy makes sense for systems that are not in equilibrium, and δQ and δW make sense throughout non-quasistatic interactions.

1.0.4. *Direction of time.* Time is not a variable in thermodynamics, in that we are not concerned with the rate at which certain processes occur. However, we are crucially interested in the *direction* of time, in the sense of the distinction between the past and the future. The *second law* of thermodynamics says that — other variables being held constant — heat always flows from an object of higher temperature to an object of lower temperature, and never the other way around. Implicit in this law is the implication that the state function temperature (whose existence is promised by the zeroth law) takes values in an ordered set, so that it makes sense to talk about higher and lower temperatures.

1.0.5. *Work.* Recall that we defined *work* as the transfer of energy between systems in an ordered form. A more conventional definition of work is that it is the application of force over some distance. The intensity of the force exerted by a system along its boundary is the *pressure*. For a system in equilibrium, the pressure is constant, and denoted P .

Pressure is force exerted per unit area; i.e. $P = F/A$. Suppose the system is confined in such a way that it expands in a frictionless way, and only in one linear direction, by exerting pressure on a mass m at position x . Since $dV = A dx$ and $F = ma = m d^2x/dt^2$ we compute

$$\int P dV = \int m \frac{d^2x}{dt^2} dx = \int m \frac{d^2x}{dt^2} \frac{dx}{dt} dt = \frac{1}{2} m \left(\frac{ds}{dt} \right)^2$$

Thus if the system moves through some closed loop γ in the PV -plane (algebraically) enclosing a region Ω , the total work done is

$$-\Delta W = \int_{\gamma} P dV = \int_{\Omega} dP dV$$

i.e. the “area” of Ω .

Other forms of work — such as the work required to stretch a wire by exerting a tensile force, or to affect the electrical potential of a chemical substance by performing some reaction — can all be expressed in the form $\sum X_i dx_i$, where each X_i is some generalized force and x_i is some generalized coordinate in some configuration space. For the sake of simplicity we express all such terms in the form $-PdV$ in what follows.

1.1. **Ideal gas.** The equilibrium state of an ideal monatomic gas (i.e. a gas in which all the molecules are single atoms of one kind) in which the number of atoms is fixed, is described by only two independent quantities — pressure P and volume V . Every other state function can be expressed in terms of P and V ; thus, for example, one has Boyle’s law

$$PV = nRT$$

where n is the number of atoms (measured in units of *moles*; i.e. approximately 6.022×10^{23}), where R is a universal constant (approximately 8.314 Joules per mole Kelvin) and T

is the temperature measured in degrees Kelvin. Notice that these constants do *not* depend on the *kind* of atoms in the gas; in particular, two different ideal monatomic gases at the same pressure, temperature and volume have the same number of atoms — even if the atoms in the two gases are of different kinds!

If the number of molecules is expressed literally as an integer N then one writes Boyle's law in the form

$$PV = Nk_B T$$

where now k_B is *Boltzmann's constant*, equal to approximately $1.38 \times 10^{-23} \text{m}^2 \text{kg s}^{-2} \text{K}^{-1}$. We have not yet justified the existence of temperature as a state function taking numerical values; thus one may take Boyle's law as the *definition* of temperature for an ideal gas.

The internal energy of an ideal gas is proportional to temperature; in fact,

$$U = C_V N k_B T = C_V PV$$

where C_V is the *heat capacity* at constant volume, which is $3/2$ for a monatomic gas; thus $U = 3/2 PV$. Actually, it is conventional to write $1/C_V = \gamma - 1$ for reasons that will shortly become clear; in our case $\gamma - 1 = 2/3$ so $\gamma = 5/3$. Consequently

$$3/2(PdV + VdP) = dU = \delta Q - PdV$$

For an adiabatic process, $\delta Q = 0$ and therefore

$$0 = 5/2 PdV + 3/2 VdP$$

or equivalently,

$$PV^{5/3} = \text{const.}$$

explaining the origin of γ as the exponent of volume in this formula.

1.2. Carnot cycle. When energy in the form of heat is supplied to an engine from a reservoir, some of that heat can be used to do work, and some is dissipated into the environment. The *efficiency* of an engine is the ratio of the work done to the heat extracted from the reservoir. Suppose our engine goes through some cyclic process, so that the total change in internal energy is zero. Then by the conservation of energy, the work done $-W$ is the difference of the heat extracted from the reservoir Q_H and the heat dissipated into the environment Q_C ; i.e. $-W = Q_H - Q_C$. So the efficiency is $\eta := 1 - Q_C/Q_H$.

The *Carnot cycle* is an ideal heat engine cycle consisting of two isothermal processes and two adiabatic processes. In one cycle, the four phases are:

- (1) a quantity Q_H of heat enters at temperature T_H during an isothermal expansion;
- (2) there is an adiabatic expansion during which the gas cools to temperature T_C ;
- (3) a quantity Q_C of heat exits at temperature T_C during an isothermal compression;
- (4) there is an adiabatic compression, raising the temperature back to T_H .

where these four steps are all performed reversibly; see Figure 1.

The second law of thermodynamics has the following remarkable implications:

- (1) of all heat motors working between a heater and a refrigerator at given temperatures T_H and T_C , the Carnot cycle has the greatest efficiency; in particular
- (2) any two Carnot cycles operating between the same two temperatures have *the same efficiency as each other*, independent of the natures of the systems.

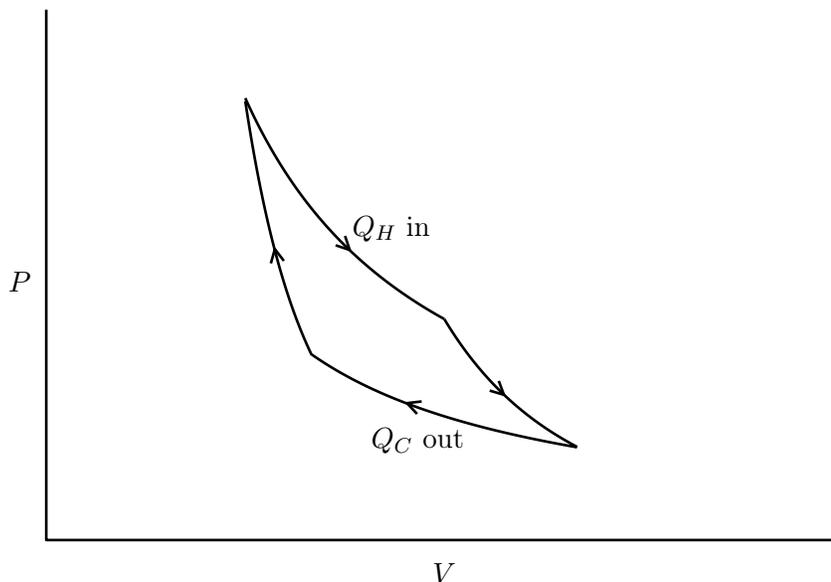


FIGURE 1. The four steps of the Carnot cycle.

To see the first implication, let A and B be two heat motors, let B be the (reversible) Carnot cycle, and let A be more efficient than B . Let A transform heat into work, and reverse B (and run it at a faster rate) to transform this work back into heat. Since A by hypothesis is more efficient than B , it takes less heat to do the same amount of work; thus the net effect of A and B combined is to transfer heat from the refrigerator to the heater. Since no net work is done, this is in violation of the second law of thermodynamics. The second implication follows immediately from the first: since any Carnot cycle (irrespective of the nature of the system) has the maximal possible efficiency, it follows that any two Carnot cycles have the same efficiency as each other.

Thus, the efficiency Q_C/Q_H of a Carnot engine operating between temperatures T_C, T_H depends only on T_C, T_H , and there is some function f for which

$$Q_C/Q_H = f(T_C, T_H)$$

In particular, f satisfies the identity

$$f(T_1, T_3) = f(T_1, T_2)f(T_2, T_3)$$

for any three temperatures T_1, T_2, T_3 . From this identity, it follows that it is possible to express f in the form $f(T_1, T_2) = g(T_1)/g(T_2)$ for some function g . Using g , we can *define* temperature (hitherto taking values in some unspecified ordered set) as the state function whose value is $g(T)$ where g is as above; by abuse of notation, we denote the value of $g(T)$ in the sequel by T . Note that g as above is well-defined up to the ambiguity of an unspecified multiplicative constant; another way of saying this is that temperature is not dimensionless, but comes in units of “degrees”; on the other hand, there is no ambiguity about the definition of an *absolute zero* temperature.

1.3. **Entropy.** In the first and third stages of the Carnot cycle the heat of the system changes by Q_H and $-Q_C$ respectively. The other two stages are adiabatic, so $\delta Q = 0$ throughout. It follows that $\int_\gamma \delta Q/T = 0$ where the integral is taken over the entire cycle γ . Since any loop in the PV -plane can be approximated by a sum of Carnot cycles, $\delta Q/T$ must be an exact differential, so that we can write it in the form

$$dS := \frac{\delta Q}{T}$$

The quantity S (*a priori* defined only up to an additive constant) is called the *entropy* of the system, and is seen by this reasoning to be a state variable. Using this definition we can express the law of conservation of energy purely in terms of state variables, in the form

$$dU = TdS - PdV$$

for reversible processes.

Entropy is *not* defined for a system which is not in equilibrium. We can imagine an irreversible transformation that takes a system from one equilibrium state x to another equilibrium state y by a path γ which takes it out of the manifold of equilibrium states in the interior of γ . Let γ_{rev} be another path in the manifold of equilibrium states with the same endpoint; thus γ_{rev} represents a *reversible* transformation between the two endstates. Now, δQ is defined as a 1-form along both paths γ , and γ_{rev} . Temperature is not well-defined along γ ; however it might be the case that all transfer of heat to the system at every point along the path γ comes from a body in equilibrium at a well-defined temperature T (e.g. a heat bath). Then it makes sense to integrate $\delta Q/T$ along γ , where T is the temperature of the body supplying the heat at each point along γ , as above (note that δQ could be negative, so that the heat is actually flowing into the bath). Since the system starts and ends in equilibrium, it has well-defined temperatures at x and y . Now, the composition of the path γ with the path $-\gamma_{\text{rev}}$ (i.e. γ_{rev} traversed in the opposite direction) is a closed loop β , whose efficiency is necessarily less than or equal to that of any reversible loop that performs the same amount of work. Thus $\int_\beta \delta Q/T \leq 0$, and we obtain the fundamental inequality

$$\int_\gamma \frac{\delta Q}{T} \leq \int_{\gamma_{\text{rev}}} \frac{\delta Q}{T}$$

In particular, for a system which is thermally isolated, so that δQ is identically zero along some irreversible path γ , we obtain the fundamental inequality

$$\Delta S \geq 0$$

i.e. *entropy is nondecreasing for irreversible processes in a thermally isolated system.* This is the form in which the second law of thermodynamics is often stated.

Remark 1.1. The second law of thermodynamics is not easy to interpret purely in the abstract; if a system in equilibrium is truly isolated, it won't change at all, so $\Delta S \geq 0$ is vacuously satisfied. In cases of genuine interest, we are actually interested in systems that are subject to *constraints* of some kind, and we care about what spontaneous changes will take place when the constraints are removed or altered. For example, consider a gas confined (by some divider) to the left half of a fixed volume of space in which the right half consists of vacuum. When the divider is removed (a change of *constraint* which in itself

does not noticeably affect the instantaneous mechanical state of the system) the gas rapidly expands to fill the vacuum, and then settles down into an equilibrium state.

1.3.1. *Reversible and irreversible processes.* We now give some examples of irreversible processes which increase the entropy of an isolated system. The first two examples concern ideal gases, the only explicit thermodynamic systems we have described up to this point.

Example 1.2 (Gas expanding into a vacuum). Consider the example (discussed in Remark 1.1) of a monatomic ideal gas confined to some subset of volume V_1 of a contained of total volume V_2 , in which the remainder consists of vacuum. When we remove a divider, the gas expands to fill the total volume. By the conservation of energy, we have $U = 3/2P_1V_1 = 3/2P_2V_2$ where P_1, P_2 are the initial and final pressures of the system. We can compute the change in entropy as a result of this expansion by considering a reversible path between the two endpoints. Since $PV = Nk_B T$ for an ideal gas, we see that the temperature does not change during the expansion (one says that the “Joule coefficient is zero”). Thus we may take a reversible path γ_{rev} to be one in which the temperature is constant throughout, i.e. a segment of the hyperbola $PV = T$ for fixed T . Then

$$\Delta S = \int_{\gamma_{\text{rev}}} \frac{\delta Q}{T}$$

Now, $\delta Q = PdV$ along γ_{rev} , because the energy of an ideal gas is proportional to its temperature, and temperature is being held constant. Thus

$$\Delta S = \int_{V_1}^{V_2} \frac{Nk_B}{V} dV = Nk_B \left(\log \frac{V_2}{V_1} \right)$$

Example 1.3 (Mixing of two or more gases). Suppose we have two quantities n_A, n_B of monatomic ideal gases A and B at the same temperature and pressure, which are brought into contact. If the divider separating them is removed, they will slowly mix with each other, eventually becoming evenly distributed throughout their combined volumes. Note that since the gases are at the same temperature and pressure, their respective volumes are proportional to their quantity: $V_A = n_A RT/P$ and $V_B = n_B RT/P$.

We would like to compute the change in total entropy of the system as a result of this mixing. As each gas expands into the common space, the temperature and thus the total energy of each individual substance stays fixed; thus we have

$$0 = dU_A = \delta Q_A - P_A dV_A$$

so that

$$dS_A = \frac{P_A dV_A}{T} = n_A R \frac{dV_A}{V_A}$$

and thus

$$\Delta S_A = n_A R \log \frac{V_A + V_B}{V_A} = n_A R \log \frac{n_A + n_B}{n_A}$$

If we define $n := n_A + n_B$ to be the total amount of gas, and $x_A := n_A/n$, $x_B := n_B/n$ to be the relative fraction of each kind, then the total increase in entropy due to mixing is

$$\Delta S = \Delta S_A + \Delta S_B = nR (-x_A \log(x_A) - x_B \log(x_B))$$

Similarly, if we mix together m different gases at the same temperature and pressure, with total quantity $n := \sum_i n_i$ and relative fractions $x_i := n_i/n$ (so that $0 < x_i < 1$ for each i and $\sum_i x_i = 1$), then the total increase in entropy due to mixing is

$$\Delta S = nR \sum_i -x_i \log(x_i)$$

Remark 1.4. Notice in this second example that we computed the total entropy of the system to be the sum of the entropies of the subsystems: $S = S_A + S_B$. In general, if we have a system which consists of several subsystems, each in isolation from each other, it makes sense to define extensive properties additively, by summing the contributions from the subsystems (thus there is a well-defined total energy, total volume, total entropy etc.) but it does *not* make sense to define intensive properties of the total system, unless the intensive properties of the subsystems are all equal, in which case we can take this common value. Thus, the union of two disjoint subsystems, each individually in equilibrium, and at temperatures T_A and T_B respectively, does not itself have a well-defined temperature unless $T_A = T_B$. Thus, properly speaking, the system as a whole cannot be said to be in thermodynamic equilibrium, even if its state is completely stationary! The key here is that a state is in thermal equilibrium (at some well-defined temperature T), not if it is unchanging when kept in isolation, but if it is unchanging when brought into thermal contact with a heat bath at that same temperature. If C consists of the union of A and B at temperatures $T_A \neq T_B$, then if C is brought in thermal contact with a heat bath at temperature T , then no matter what T is, there will be some nonzero transfer of heat between the bath and at least one of A or B .

Reversible processes are an ideal fiction; even ignoring dissipation of energy as heat due to friction, when there are nonzero differences in temperature or pressure between systems, heat or work is transmitted at a definite rate, and total entropy increases. We consider two basic examples:

Example 1.5 (Systems at different temperatures). Consider two systems A and B with energy U_A, U_B and entropy S_A, S_B . Suppose we bring the systems into contact and let heat flow from one to the other, holding volumes constant. If the combined system is isolated, the total energy is conserved, so $dU_A = -dU_B$. Since volume is held constant, no work is done, so $dU_A = \delta Q_A = -\delta Q_B$. Thus

$$dS = dS_A + dS_B = \delta Q_A \left(\frac{1}{T_A} - \frac{1}{T_B} \right)$$

The second law of thermodynamics says that if $\delta Q_A \geq 0$ (i.e. if heat flows from B to A) then $T_B \geq T_A$; i.e. $dS \geq 0$. If $T_B > T_A$ then $dS > 0$ unless the materials are perfectly insulated, so the process is irreversible.

Example 1.6 (Systems at different pressures). Consider two systems in contact at the same temperature but with different pressures. For a single system, $TdS = (dU + PdV)$. By conservation of energy, $dU_A = -dU_B$, so if the systems are at the same temperature,

$$dS = \left(\frac{P_A}{T} - \frac{P_B}{T} \right) dV_A$$

The body at higher pressure (say, A) will tend to expand, so that V_A increases. If $P_A > P_B$ then $dS > 0$, so the process is irreversible.

1.4. Enthalpy, Helmholtz free energy, Gibbs free energy. One of the crucial properties of entropy is that it can be used to predict when a system is in equilibrium with its surroundings: since the total entropy can only increase under irreversible processes, no such processes can occur when total entropy is at a (local) maximum. The drawback is that it is necessary to compute, not just the entropy of the system we are interested in, but also the entropy of the surroundings.

1.4.1. *Enthalpy.* The *enthalpy* of a system is defined to be

$$H := U + PV$$

Thus,

$$dH = \delta Q + VdP$$

so *at constant pressure*, $dP = 0$ and

$$dH|_{P=\text{const.}} = \delta Q|_{P=\text{const.}}$$

Thus, enthalpy is a state variable which can be viewed as a measure of “heat content” in a constant pressure system. This is important because many thermodynamic processes of significant interest (for example, chemical reactions, melting/freezing, etc.) occur at constant pressure.

Example 1.7 (Methane formation). The reaction which combines a mole of carbon and two moles of dihydrogen and produces a mole of methane ($C + 2H_2 \rightarrow CH_4$) is *exothermic*, producing 74 kJ of heat at 1 atmosphere of pressure. Thus $\Delta H = -74$ kJ in this case.

A comparison of the heat of formation (i.e. the change in enthalpy to produce some molecule from its constituent atoms) of the ingredients and the products of a chemical reaction can predict whether the reaction will be exothermic or endothermic; this fact (really just a restatement of the law of conservation of energy) is known as *Hess's Law*.

1.4.2. *Gibbs free energy.* Enthalpy can be used to tell whether certain processes (at constant pressure) will create or absorb heat, but it cannot be used directly to tell whether the process will occur.

The *Gibbs free energy* is a more useful quantity, measuring the amount of energy in the system available to do useful work at fixed temperature and pressure. Formally,

$$G := H - TS$$

Example 1.8 (Enthalpy, entropy and temperature). Now, consider some chemical reaction that takes place at fixed pressure and temperature. The reaction releases $-\Delta H$ heat into the surroundings, changing the entropy of the surroundings by $-\Delta H/T$. However, the reaction might change the number of internal degrees of freedom of the system (e.g. by changing the number of molecules) and therefore the internal entropy. Thus the change in *total* entropy of the system plus the surroundings

$$\Delta S_{\text{system}} + \Delta S_{\text{surroundings}} = \Delta S_{\text{system}} - \Delta H_{\text{system}}/T = -\Delta G_{\text{system}}/T$$

so that in order for the reaction to take place (without violating the second law of thermodynamics) the Gibbs free energy must *decrease*.

Informally, an exothermic reaction will take place at sufficiently low temperature, no matter what the change in entropy, whereas an endothermic reaction which increases entropy will take place at sufficiently high temperature. Finally, an endothermic reaction which decreases entropy will never take place spontaneously, whereas an exothermic reaction which increases entropy will take place at any temperature. See Figure 2.

ΔH	
no T	high T
low T	ΔS any T

FIGURE 2. Gibbs free energy determines at which temperatures a reaction will take place.

Example 1.9 (Reversible and irreversible expansions). Imagine a gas expanding in a piston at constant pressure and temperature (e.g. because the gas is adjacent to a heat reservoir). In the absence of friction, and if the expansion takes place reversibly, the amount of heat entering the system is Q_R . If the expansion is irreversible (e.g. because some heat is generated due to friction) the amount of heat entering the system Q_I satisfies $Q_R > Q_I$. Thus,

$$\Delta G = \Delta H - T\Delta S = Q_I - Q_R < 0$$

(here we compute $\Delta S = Q_R/T$ at constant temperature; because entropy is a state function, it must be computed by a *reversible* transformation of the system).

Example 1.10 (Rubber band). When a rubber band is stretched, it gives off heat, and when it snaps back from a stretched state, it absorbs heat. In particular, ΔH is positive during the process by which the rubber band snaps back. On the other hand, this process occurs spontaneously, so that $\Delta G < 0$. But $\Delta G = \Delta H - T\Delta S$, and therefore we deduce the (perhaps surprising) fact that ΔS must be positive; i.e. the entropy of the stretched rubber band is *smaller* than that of the loose one. This has to do with the molecular structure of rubber, which consists of long, tangled polymers; in the loose state, these polymers can arrange themselves in many different ways, contributing greatly to the disorder (and hence the entropy) of the system. When they are stretched tight, there is less flexibility in their arrangement, and the disorder (and hence the entropy) decreases.

1.4.3. *Helmholtz free energy*. The *Helmholtz free energy* of a system is defined to be

$$A := U - TS$$

Thus, for reversible processes,

$$dA = -SdT + \delta W$$

so *at constant temperature*, $dT = 0$ and

$$dA|_{T=\text{const.}} = \delta W|_{T=\text{const.}}$$

Thus, Helmholtz free energy is a state variable which can be viewed as a measure of the capacity of the system to do “useful work” in a constant temperature system.

Similarly to the Gibbs free energy, the Helmholtz free energy can be used to tell whether certain processes (at constant temperature and volume) will take place. If T and V are constant, then for an irreversible spontaneous process with fixed temperature and volume, $dU \leq TdS$; i.e. $dA \leq 0$. In other words, a system is in equilibrium if it (locally) minimizes Helmholtz free energy.

1.4.4. *Relation between different kinds of energy.* The four quantities U, A, H, G , all with units of energy, are related as indicated in the diagram:

$$\begin{array}{ccc} U & \xrightarrow{+PV} & H \\ \downarrow -TS & & \downarrow -TS \\ A & \xrightarrow{+PV} & G \end{array}$$

The derivatives of these quantities are given by formulae

$$dU = TdS - PdV, \quad dH = TdS + VdP, \quad dA = -SdT - PdV, \quad dG = -SdT + VdP$$

One way to express the relationships between these quantities is in terms of the *Legendre transform* (see § A.1).

For example, if we keep the volume of the system constant, then no work is done and we get the relation

$$\left. \frac{\partial S}{\partial U} \right|_{V=\text{const.}} = \frac{1}{T}$$

The *inverse temperature* is defined to be

$$\beta := \frac{1}{k_B T}$$

where k_B is Boltzmann’s constant. If we “normalize” entropy and Helmholtz free energy to the dimensionless quantities $\mathcal{S} := S/k_B$ and $\mathcal{F} := \beta A$, then β is the derivative of \mathcal{S} as a function of energy (at constant volume), and \mathcal{F} becomes the Legendre transform of \mathcal{S} ; i.e. (at constant volume) $d\mathcal{S}(U)/dU = \beta$, $d\mathcal{F}(\beta)/d\beta = U$, and

$$\mathcal{F}(\beta) + \mathcal{S}(U) = \beta U$$

1.4.5. *Chemical potential.* Up to now we have generally considered systems which are permitted to exchange energy (in the form of work or heat) with their environment, but not matter. Imagine a system whose boundary is permeable or semi-permeable, admitting the transfer (in one direction or both) of molecules of some type with the ambient environment. The chemical potential is the change in (Gibbs) free energy of the system as the result of this transfer; thus “potential” in this sense measures the potential energy of the molecules that might be released in some reaction.

The equation for the differential of the Gibbs energy is modified to the form

$$dG = -SdT + VdP + \sum \mu_i dN_i$$

where for each kind i of particle, the total number N_i is now considered an extensive thermodynamic state function, and

$$\mu_i = \left. \frac{\partial U}{\partial N_i} \right|_{S,V,N_{j \neq i} \text{ constant}} = \left. \frac{\partial G}{\partial N_i} \right|_{T,P,N_{j \neq i} \text{ constant}}$$

Note that as far as applications to chemistry are concerned, all that matters is the difference in potential energy between expressions of the form $\sum n_{i,A} \mu_{i,A}$ and $\sum n_{j,B} \mu_{j,B}$ whenever there is some reaction that takes as ingredients quantities $n_{i,A}$ of particle i and produces quantities $n_{j,B}$ of particle j .

1.4.6. *Phase transitions.* A system in equilibrium does not need to be homogeneous; it could consist of a union of different substances in equilibrium with each other. For example, we could consider a mixture of n different substances, each of which can be in one of f separate *phases* (e.g. solid, liquid, gas). Each substance in each phase obeys its own equations of state, and makes its own contribution to the (Gibbs) free energy G . Let m_{jk} be the mass of substance k in phase j , and let G_j be the total contribution to the Gibbs free energy from the substances in phase j ; thus $G = G_1 + \cdots + G_f$.

Each G_j is (neglecting effects due to electromagnetism, gravity, etc.) a function of T , P and the various m_{jk} , and is homogeneous of first order in the m_{jk} . Thus the partial derivatives of the G_j depend only on the ratios of the various m_{jk} , together with the values of T and P ; i.e. on $(n-1)f+2$ variables.

Since for each substance k the sum $\sum_j m_{jk}$ is fixed, infinitesimal variations of the system are spanned by covectors of the form $dm_{jk} - dm_{ik}$. The condition that G is at a local minimum thus implies an equation of the form $\partial G_i / \partial m_{ik} = \partial G_j / \partial m_{jk}$ for all substances k and all phases i, j . Thus we obtain $n(f-1)$ equations in $(n-1)f+2$ variables, so that the dimension v of the space of states in such a mixture of phases is

$$v := (n-1)f+2 - n(f-1) = 2 - f + n$$

known as ‘‘Gibbs’ phase rule’’.

Example 1.11 (Triple point of water). Consider the example of water, which can be in any of three phases: ice, liquid water, or steam. In this case $n = 1$ and $f = 3$ so $v = 0$, so that the three phases can coexist only at isolated values of the pair of variables temperature and pressure. In fact, under reasonable conditions, there is a unique value of T, P called the *triple point*; for water this is at approximately $T = 0.0075^\circ\text{C}$ and $P = 0.00602$ atm. The 1-dimensional manifolds where two phases can coexist (water/ice or steam/water) intersect the line $P = 1$ atm at $T = 0^\circ\text{C}$ (freezing point at atmospheric pressure) and $T = 100^\circ\text{C}$ (boiling point at atmospheric pressure) respectively. The manifolds of positive codimension representing coexistence of phases correspond to discontinuities in derivatives of the Gibbs free energy G as a function of the global state.

In fact, the liquid/gas interface does not extend to arbitrarily high temperatures, but only to a *critical point*, beyond which new phases emerge (compressible liquid, superheated vapour, supercritical fluid). The critical point for water is at $T = 647^\circ\text{K}$ and $P = 374$ atm.

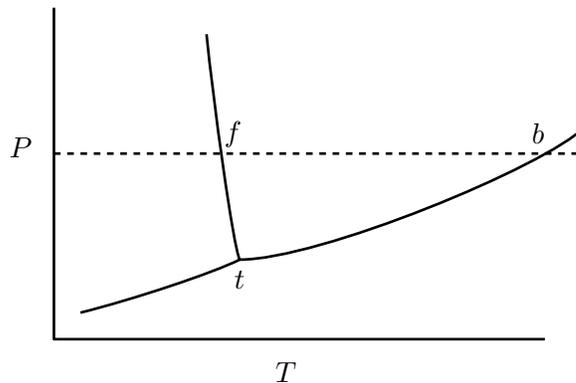


FIGURE 3. Water can coexist simultaneously in three phases in equilibrium only at the triple point t . The dashed line indicates atmospheric pressure; thus f and b are the freezing and boiling points of water at 0°C and 100°C respectively.

2. STATISTICAL MECHANICS

It is time to peer under the hood, and explain how the macroscopic thermodynamic properties of systems in equilibrium emerge from the mechanics of their microscopic parts, as seen through the lens of probability.

2.1. The kinetic theory of gases. How do macroscopic thermodynamic functions arise from the microscopic properties of a system in “equilibrium”? And what does equilibrium mean at the microscopic level anyway? We first consider several different plausible models (ranging from classical to quantum mechanics) which give roughly consistent answers.

2.1.1. Maxwell distribution of velocities. Imagine a large number N of elastic identical particles confined to a bounded region Ω in the absence of external forces. If x, y, z denote the three coordinate directions, let $\phi(x, y, z)dxdydz$ denote the probability distribution of velocities of particles, in the idealized limit that $N \rightarrow \infty$. Maxwell derives a formula for ϕ under the following assumptions:

- (1) the distribution of velocities should be invariant under the group of rotations of 3-dimensional space; and
- (2) the distribution of components of velocities in orthogonal directions should be mutually independent.

The second assumption is justified on the grounds that under the infinitesimal elastic collisions between particles, the orthogonal components of the velocities transform independently of each other.

If the x component of velocity has the distribution $f(x)dx$ then by reasons of symmetry, the y and z components should have distributions $f(y)dy$ and $f(z)dz$ respectively. Our two assumptions together give a relation of the form

$$f(x)f(y)f(z) = \phi(x, y, z) = \psi(x^2 + y^2 + z^2)$$

for suitable ψ . Taking logarithms, and using the relation $\log f(x) = 1/3 \log \psi(3x^2)$ we get

$$\log \psi(3x^2) + \log \psi(3y^2) + \log \psi(3z^2) = 3 \log \psi(x^2 + y^2 + z^2)$$

Applying $\partial/\partial x$ to both sides shows that $\log \psi$ is quadratic in x, y, z . From symmetry, and the fact that a probability measure is positive and has finite total mass, we deduce $\phi(v) = C_1 e^{-C_2 v^2}$ for some positive constants C_1, C_2 depending on the context; i.e. the distribution of velocities follows a Gaussian law on \mathbb{R}^3 . Since the kinetic energy of a particle is $\epsilon := \frac{1}{2}m(x^2 + y^2 + z^2)$, the map $\mathbb{R}^3 \rightarrow \mathbb{R}^+$ defined by $(x, y, z) \rightarrow \epsilon$ pushes forward the measure $\phi(x, y, z) dx dy dz$ to $e^{-\epsilon\beta} \epsilon^{1/2} d\epsilon$ (up to a multiplicative constant) for some β . In particular, the expected amount of kinetic energy per particle is

$$\langle k \rangle := \frac{\int_0^\infty e^{-\epsilon\beta} \epsilon^{3/2} d\epsilon}{\int_0^\infty e^{-\epsilon\beta} \epsilon^{1/2} d\epsilon} = \frac{\beta^{-5/2} \Gamma(5/2)}{\beta^{-3/2} \Gamma(3/2)} = \frac{3}{2\beta}$$

and the amount of kinetic energy “in each coordinate direction” is $\frac{1}{2\beta}$. On physical grounds one can argue that β as above should be inversely proportional to temperature, so that the average kinetic energy per degree of freedom is $\frac{1}{2}k_B T$ for some constant k_B , and the total kinetic energy is $U := \frac{3}{2}Nk_B T$.

2.1.2. *Density of gas in a constant gravitational field.* Imagine a vertical column of gas in a constant gravitational field. How does pressure (equivalently, density) of the gas vary as a function of height? In equilibrium, the temperature T should be constant throughout, and (by symmetry) the density of molecules of gas per unit volume should depend only on the height t ; call this density $n(t)$. Thus the pressure at height t should satisfy

$$P(t) = n(t)k_B T$$

Since forces balance in equilibrium, the difference in pressure $dP(t)/dt$ associated to a change dt in height should be the difference in the weight of the column of gas above some fixed region of unit area, which is $-mgn(t)dt$ where m is the mass of a molecule. Thus we obtain a differential equation

$$\frac{dn(t)}{dt} = \frac{-mgn(t)}{k_B T}$$

with solution $n(t) = n_0 \cdot e^{-mgt/k_B T}$. Since $-mgt$ is the potential energy of a molecule at height t , we see that the distribution of number of particles as a function of potential energy has the form $c \cdot e^{-\text{potential energy}(t)\beta} dt$ for $\beta = 1/k_B T$.

2.1.3. *Ensembles of identical particles.* In slightly more generality, let us imagine a large number N of identical objects, each of which has some finite number f of degrees of freedom, which can be expressed as position and momentum coordinates in phase space. Thus, the state of each object is described by a point in a phase space of dimension f , and the state of the entire system is described by a point in the product of N copies of the single-object phase space, i.e. a space of dimension fN . We suppose that the energy of each object is given by a quadratic function on phase space \mathbb{R}^f ; or, more generally, that it depends on some subset of the coordinates, and is a sum of quadratic functions of each of these.

If we know the total energy U of the system of N objects, then we know that the motion of the system in phase space is confined to a hypersurface Σ in \mathbb{R}^{fN} . If the particles do

not interact at all, the energy of each individual particle is conserved, and their motion is confined to a subsurface of codimension at least N . But if the particles can interact with each other and exchange energy, it is possible to imagine that the system could reach any point in Σ .

In the absence of more detailed knowledge, it is plausible to think of the state of the system at some unspecified future time as a random variable, associated to a probability measure on Σ . The symplectic volume form on phase space restricts to a measure of finite volume on a thin shell between two nearby level sets U and $U + \epsilon$ of the energy function; rescaling this to a probability measure, and taking the limit as $\epsilon \rightarrow 0$, we obtain a probability measure on Σ invariant under any Hamiltonian flow. This is called the *microcanonical ensemble*.

Pick one coordinate x associated to any of the N particles, and push forward the uniform measure on Σ to the x coordinate. Since energy is quadratic (by hypothesis), up to an affine change of coordinates we may suppose that the energy $\epsilon(x)$ associated to each coordinate x is just x^2 , and Σ is the sphere in \mathbb{R}^{fN} of radius \sqrt{U} , and we are looking at the pushforward of the uniform measure to the x coordinate. Evidently this pushforward is of the form $C \cdot \sqrt{U - x^2}^{(fN-3)} dx$. For N very big, this is very close to zero unless $x \ll 1$. So it is reasonable to approximate $1 - x^2$ by e^{-x^2} , and deduce that the pushforward is well-approximated by a Gaussian of the form $C \cdot e^{-x^2 \frac{fN}{2U}} dx$. If we write the average energy per degree of freedom as $U/fN = k_B T/2 = 1/2\beta$, then the pushforward measure is $C \cdot e^{-\beta\epsilon(x)} dx$.

2.1.4. Equipartition theorem. The plausible idea, derived on symmetry grounds in § 2.1.1, that the expected amount of kinetic energy is equal in each of the three coordinate directions, is part of a more general principle called the “equipartition theorem”, which argues that every independent mode (i.e. degree of freedom) of the system should get the same average energy.

So, for example, for a diatomic gas, the average amount of energy in each translational (coordinate) mode should be equal to the amount in each vibrational and rotational mode; thus the heat capacity C_V , which is $3/2$ for a monatomic gas (with contributions of $1/2$ from each of the three translational modes), should be $7/2$ for a diatomic gas (two rotational modes for inertia perpendicular to the axis of the molecule; and two vibrational modes consisting of both potential and kinetic energy of vibration).

The equipartition theorem presupposes a probability distribution on the phase space of each individual particle in which the hypersurfaces of constant energy ϵ are weighted by a factor proportional to $e^{-\beta\epsilon}$, for $\beta = 1/k_B T$. Two plausible reasons to consider such distributions are given in § 2.1.1 and § 2.1.3.

Proposition 2.1 (Equipartition theorem). *Suppose we have a classical mechanical system depending on variables p_j, q_j for which the Hamiltonian $\epsilon(p, q)$ can be expressed as a sum of quadratic functions $\epsilon = \sum_j \epsilon(p_j) + \epsilon(q_j)$ of each variable separately. Then for a probability distribution in phase space proportional to $e^{-\beta\epsilon(p, q)} dp_1 \cdots dq_n$, we have*

$$\langle \epsilon(p_j) \rangle = \langle \epsilon(q_j) \rangle = \frac{1}{2\beta} \text{ for all } j$$

Proof. By the assumption that we can write $\epsilon = \sum_j \epsilon(p_j) + \epsilon(q_j)$, the total integral can be expressed in the form

$$Z := \int e^{-\beta\epsilon} dp_1 \cdots dq_n = \prod_j \int e^{-\beta\epsilon(p_j)} dp_j \int e^{-\beta\epsilon(q_j)} dq_j = \prod_j Z(p_j)Z(q_j)$$

Thus the expected value of $\epsilon(p_j)$ (say) is

$$\langle \epsilon(p_j) \rangle = \frac{\int \epsilon(p_j) e^{-\beta\epsilon} dp_1 \cdots dq_1}{Z} = \frac{\int \epsilon(p_j) e^{-\beta\epsilon(p_j)} dp_j}{Z(p_j)} = \frac{1}{2\beta} = \frac{1}{2} k_B T$$

if $\epsilon(p_j)$ is quadratic in p_j . □

However, the equipartition theorem fails at sufficiently low temperatures because of quantum effects, depending on differences in the spacing between energy levels of the different modes; those modes with big gaps between energy levels get “frozen out” at sufficiently low temperatures, and share less of the total energy. See § 2.4. The temperatures at which the equipartition theorem fails for particular modes depends greatly on the particular atomic details of the structure in question. For many real diatomic gases, vibrational modes freeze out at temperatures on the order of thousands of degrees kelvin, while rotational modes freeze out at temperatures on the order of tens of degrees kelvin, so that in practice, the approximation $C_V = 5/2$ is reasonable for many real diatomic gases at room temperature (e.g. hydrogen, carbon monoxide, air). For translational degrees of freedom, the spacing between energy levels is controlled by the dimensions of the container, so that for macroscopic containers the translational degrees of freedom are well described by the classical model throughout a great range of temperatures.

2.2. Boltzmann distribution. The fundamental problem in statistical mechanics is to understand how a fixed amount E of energy is distributed over a large number N of identical systems. In idealizations, we consider the limit as both E and N are taken to infinity in such a way that E/N is constant.

2.2.1. Boltzmann versus Gibbs. Suppose we have a quantum mechanical system whose states have energy levels ϵ_j (possibly with multiplicity). In Boltzmann’s original formulation, the system is assumed to be composed of a very large number (N) of weakly interacting identical particles, each with few internal degrees of freedom; the canonical example is an ideal monatomic gas.

However, few realistic physical systems of interest satisfy these assumptions; thus in Gibbs’ formulation we consider an idealized ensemble consisting of N independent copies of *the system itself*, which are interacting weakly in such a way as to allow only the transfer of energy between the different copies. In this formulation, all but one of the copies of the system constitute a sort of “heat bath”, weakly coupled to the copy under consideration. In thermodynamic equilibrium, the state of a system coupled to a heat bath should depend only on the temperature of the heat bath. So the probability of finding a system in equilibrium in some particular state should be governed by the *same distribution* in Gibbs’ formulation as in Boltzmann’s, despite the fact that the physical justification is completely different! The model in which the system is allowed to exchange energy with a heat bath

at the same temperature (so that the total energy of the system is not fixed, but is rather a random variable) is called the *canonical ensemble*.

2.2.2. *Distribution of maximum probability.* In either case therefore we consider the problem of how to distribute a very large number N of objects among the possible energy levels ϵ_j . If we let n_j denote the number of objects with energy ϵ_j , we are subject to two constraints

$$\sum n_j = N, \quad \text{and} \quad \sum n_j \epsilon_j = E$$

where E is the total energy.

The number of ways of partitioning N into subsets with cardinality n_j is $\Omega := N! / \prod_j n_j!$. If all possible distributions are equally possible, the most likely partition is the one for which Ω has the highest cardinality. It will turn out for N very large ($N \rightarrow \infty$ in the idealized limit) that almost all possible partitions of N have distributions very close to those with maximum cardinality; this is a sort of “law of large numbers”, known as the Darwin-Fowler Theorem, and we give the proof in Appendix B (a different sort of justification is given in § 2.2.3). Thus we are led naturally to the problem of finding the distribution which maximizes Ω .

Stirling’s approximation for the factorial function is

$$n! \sim \sqrt{2\pi n} \left(\frac{n}{e}\right)^n$$

so that

$$\log(n!) = n \log(n) - n + O(\log(n))$$

Thus (neglecting lower order terms)

$$\log(\Omega) \sim N(\log(N) - 1) - \sum n_j(\log(n_j) - 1)$$

We can find the distribution maximizing $\log(\Omega)$ (and therefore also maximizing Ω) by the method of Lagrange multipliers. We want to maximize $\log(\Omega)$ as a function of the n_j , subject to fixed values of $\sum n_j$ and $\sum n_j \epsilon_j$. At the maximum, the gradient of $\log(\Omega)$ should therefore be a linear combination of the gradients of $\sum n_j$ and $\sum n_j \epsilon_j$; i.e. there should be constants α and β so that

$$\frac{\partial \log(\Omega)}{\partial n_j} - \alpha \frac{\partial (\sum n_j)}{\partial n_j} - \beta \frac{\partial (\sum n_j \epsilon_j)}{\partial n_j} = 0$$

for all j . In other words, $\log(n_j) = -\alpha - \epsilon_j \beta$ or equivalently, $n_j = e^{-\alpha - \epsilon_j \beta}$ for all j . If we define $Z(\beta) := \sum e^{-\epsilon_j \beta}$ then the probability $\text{Pr}(j)$ of the system being in state j is

$$\text{Pr}(j) = \frac{e^{-\epsilon_j \beta}}{Z(\beta)}$$

This is the *Boltzmann distribution*. The function $Z(\beta)$ is the *partition function* of the distribution.

2.2.3. *Heat bath.* Another derivation of the Boltzmann distribution on physical grounds comes by making certain assumptions about the nature of a large heat bath. In the model of the system coupled to a large heat bath, we make the assumption that the total energy of the system plus heat bath is much larger than that of the system itself, and that the energy levels of the heat bath in the likely range are “quasi-continuous” (i.e. the gaps between levels are of a lower order of magnitude than the size of the variations considered). If the total energy of the system plus the bath is (approximately) E_0 , then the number of ways the system can have energy approximately ϵ_j is proportional to the density of the number of states of the heat bath at energy $E_0 - \epsilon_j$. Let $\eta(t)$ be the density of states of the heat bath as a function of energy; our assumption of quasi-continuity says that η has an approximate derivative, which should be approximately constant in the domain of interest. Thus

$$\frac{\Pr(\text{energy} = \epsilon_j)}{\Pr(\text{energy} = \epsilon_k)} = \frac{\eta(E_0 - \epsilon_j)}{\eta(E_0 - \epsilon_k)} = e^{-\frac{d \log \eta}{dt}(E_0)(\epsilon_j - \epsilon_k)}$$

Writing $\beta := d \log \eta / dt(E_0)$ and normalizing so that the total probability is 1, we recover the Boltzmann distribution.

This derivation is taken from Feynman [4], pp. 1–3.

2.3. **Thermodynamic quantities from the Boltzmann distribution.** Our next task is to express the classical thermodynamic quantities of the system in terms of the partition function and the Boltzmann distribution.

2.3.1. *Independent systems in equilibrium.* Suppose we have two systems A and B with energy levels ϵ_i and ϵ'_j respectively. The system $A + B$ which is the disjoint union of A and B has states which are pairs (i, j) of states of A and B , with energy levels $\epsilon''_{(i,j)} := \epsilon_i + \epsilon'_j$. Thus the partition function for $A + B$ has the form

$$Z_{A+B}(\beta) = \sum_{i,j} e^{-\beta(\epsilon_i + \epsilon'_j)} = \left(\sum_i e^{-\beta \epsilon_i} \right) \left(\sum_j e^{-\beta \epsilon'_j} \right) = Z_A(\beta) Z_B(\beta)$$

Thus the conditional probability that system A is in state i as part of the system $A + B$ is equal to the probability that system A is in state i in isolation, providing β is equal in every case. Thus if we interpret β as a function of temperature, we recover the zeroth law of thermodynamics, that two systems individually in equilibrium at the same temperature will be in equilibrium with each other. Exactly which function of temperature β is will become clear in § 2.3.3.

2.3.2. *Energy U .* For a system coupled to a heat bath, the internal energy is subject to fluctuations, and we can only talk about its expected value:

$$U = \langle E \rangle = \sum_j \epsilon_j \Pr(j) = \frac{\sum_j \epsilon_j e^{-\epsilon_j \beta}}{\sum_j e^{-\epsilon_j \beta}} = -\frac{Z'}{Z} = -\frac{d \log Z}{d\beta}$$

Fluctuations from the average are controlled by the variance:

$$\text{var}(E) := \langle E^2 \rangle - \langle E \rangle^2 = \frac{Z''}{Z} - \left(\frac{Z'}{Z} \right)^2 = \frac{d^2 \log Z}{d\beta^2}$$

Since the variance is evidently positive for positive T whenever there are states with different energy, we see that $\log Z$ is strictly convex as a function of β .

2.3.3. *Temperature T , Entropy S .* How does U vary as a function of “volume”? Here we think of volume as a euphemism for the various constraints on the form of the system which determine the energy levels ϵ_j associated to the states. In quantum dynamics, states are eigenfunctions of the Hamiltonian operator, and energy levels are eigenvalues. For example, for a single particle confined to a rigid enclosed region, in the absence of an external potential, the states are eigenvalues of the Laplacian with Dirichlet boundary conditions (see Example A.6). Generically, the eigenspaces are simple, and if we deform the constraints of the system (and thereby deform the Hamiltonian), the eigenfunctions and eigenvalues vary continuously, and it makes sense to talk about the *same* state j in a varying family of conditions; likewise, it makes sense to talk about the derivative of the energy level ϵ_j as a function of parameters. For the sake of discussion therefore, “varying volume” means varying the ϵ_j , although we do not pursue the question of what sorts of variations are physically meaningful.

In any case, thinking of Z now as a function of β and of the ϵ_j , we can write

$$\begin{aligned} d \log Z &= \frac{\partial \log Z}{\partial \beta} d\beta + \sum_j \frac{\partial \log Z}{\partial \epsilon_j} d\epsilon_j \\ &= -U d\beta + \sum_j \frac{-\beta e^{-\beta \epsilon_j} d\epsilon_j}{\sum_j e^{-\beta \epsilon_j}} \\ &= -U d\beta - \beta \sum_j \Pr(j) d\epsilon_j \end{aligned}$$

and therefore we obtain the formula

$$d(\log Z + U\beta) = \beta \left(dU - \sum_j \Pr(j) d\epsilon_j \right)$$

We argue that it makes sense to interpret $\sum_j \Pr(j) d\epsilon_j$ as δW — the (average) work done to one of our N systems under the change of constraints. For a copy of the system in a definite fixed state j , adjusting the constraints of the system transfers energy in an ordered form (i.e. as work), in exactly the quantity $d\epsilon_j$. Thus we may interpret the difference $dU - \sum_j \Pr(j) d\epsilon_j$ as the (average) *heat* δQ transferred to a system under an infinitesimal change of constraints. We have seen that δQ (for reversible transformations) is not exact, but becomes exact when divided by temperature — in fact, this was how we *defined* the numerical value of temperature in § 1.3. Thus we are justified in interpreting β as the inverse of temperature, up to a universal constant, the *Boltzmann constant*; i.e.

$$\beta = \frac{1}{k_B T}$$

Moreover, on the same grounds, we are justified in interpreting entropy (up to a multiplicative constant) as the expression on the left-hand side; i.e.

$$S = k_B (\log Z + U\beta) = k_B \left(\log Z - \beta \frac{d \log Z}{d\beta} \right)$$

2.3.4. *Helmholtz Free energy A.* Classically, we define the Helmholtz free energy by the formula $A := U - ST$. Since $S = k_B(\log Z + U\beta)$ we derive the simple formula

$$A = -\frac{1}{\beta} \log Z$$

so that the dimensionless quantity βA we considered earlier is just $-\log Z$. Another way of expressing this formula is

$$e^{-A\beta} = \sum e^{-\epsilon_j\beta}$$

2.3.5. *Another formula for S.* It is possible to rearrange the formula for S into a form which is more familiar to mathematicians, as follows. We compute

$$\begin{aligned} -\frac{1}{\beta} \sum_j \Pr(j) \log \Pr(j) - \frac{1}{\beta} \log Z &= -\frac{1}{\beta} \left(\sum_j \frac{e^{-\epsilon_j\beta}}{Z} \log \frac{e^{-\epsilon_j\beta}}{Z} + \log Z \right) \\ &= -\frac{1}{\beta} \sum_j \frac{e^{-\epsilon_j\beta}}{Z} \log e^{-\epsilon_j\beta} \\ &= -\frac{1}{\beta} \sum \frac{-\epsilon_k\beta e^{-\epsilon_k\beta}}{Z} = U \end{aligned}$$

Since $A = -\frac{1}{\beta} \log Z$ and $U = A + TS$ we obtain the following formula for the dimensionless quantity S/k_B considered earlier:

$$\frac{S}{k_B} = - \sum_j \Pr(j) \log \Pr(j)$$

2.3.6. *Variational properties of S and A.* Recall our earlier formula $\log \Omega \sim N \log N - \sum n_j \log n_j$ where Ω is the number of assignments of N objects into states j with cardinality n_j (and $\sum n_j = N$). If $\nu(j) := n_j/N$ is the probability measure on states associated to an arbitrary assignment, then

$$-N \sum_j \nu(j) \log \nu(j) = - \sum n_j \log \frac{n_j}{N} = - \left(\sum n_j \log n_j \right) + N \log N = \log \Omega$$

For any probability measure ν , let $S(\nu)$ denote the *entropy* of ν , defined by $S(\nu) := -k_B \sum_j \nu(j) \log \nu(j)$. Then we recover the following fact:

Proposition 2.2 (Maximum entropy). *The Boltzmann distribution maximizes entropy for fixed energy amongst all probability distributions on states.*

Since $S(U)/k_B$ and $\beta A(\beta)$ are related by the Legendre transform, they carry the same information. So the statement that the Boltzmann distribution maximizes entropy for fixed energy can be recast as the statement that the Boltzmann distribution minimizes *free energy* for fixed *temperature*. Here we need a notion of free energy for an arbitrary probability distribution. Using $A = U - ST$, it makes sense to define, for an arbitrary probability measure ν , the quantity

$$\beta A_\nu := \beta(\langle E \rangle_\nu - S(\nu)T) = \sum \nu(j)(\beta\epsilon_j + \log \nu(j))$$

where A_ν is the “free energy associated to the probability distribution ν ”, at temperature $1/k_B\beta$.

Proposition 2.3 (Minimum free energy). *The Boltzmann distribution minimizes free energy for fixed temperature amongst all probability distributions on states.*

Proof. This follows immediately from the properties of the Legendre transform. Another proof follows from the concavity of the function \log :

$$\begin{aligned} -\beta A_\nu &= -\sum_j \nu(j) (\log \nu(j) + \beta \epsilon_j) \\ &= \sum_j \nu(j) \log \frac{e^{-\beta \epsilon_j}}{\nu(j)} \leq \log \sum_j \nu(j) \frac{e^{-\beta \epsilon_j}}{\nu(j)} = \log Z = -\beta A \end{aligned}$$

where the strict concavity of \log means that there is equality if and only if $\nu(j)$ is proportional to $e^{-\beta \epsilon_j}$, which is the Boltzmann distribution. \square

2.4. Quantum harmonic oscillator. Consider a single quantum harmonic oscillator, with energy levels $\epsilon_j = (j + 1/2)h\nu$ for $j = 0, 1, 2$, and so on. Changing the energy levels of every state by the same constant does not affect the probability distribution, so we ignore the “zero-point energy” $h\nu/2$, and set $\epsilon_j = jh\nu$ for simplicity. Then $\Pr(j) = e^{-j\beta h\nu}/Z$ and

$$Z = \sum_{j=0}^{\infty} e^{-j\beta h\nu} = \frac{1}{1 - e^{-\beta h\nu}}$$

and

$$U = \langle E \rangle = h\nu \frac{e^{-\beta h\nu}}{1 - e^{-\beta h\nu}}$$

where $\beta = 1/k_B T$. When T is sufficiently high (depending on ν) then $\beta h\nu$ is very small, so that U is approximately equal to $k_B T$. However, at low temperatures, the average energy goes to zero. Conversely, for fixed temperature, taking ν sufficiently high makes U go to zero.

The partition function for a collection of independent quantum harmonic oscillators is the product of the individual partition functions; thus if the oscillators have different frequencies, the energy is “frozen out” of sufficiently high frequency wavelengths, at any fixed temperature.

This is fundamentally a quantum phenomenon; in the classical limit $h \rightarrow 0$ the average energy is equal to $k_B T$ independent of frequency (this is the classical “equipartition theorem” for independent degrees of freedom for a system in equilibrium). For a black body (i.e. an opaque, non-reflective body held at constant, uniform temperature) the modes are standing waves; the number of modes per unit frequency is thus (asymptotically) proportional to the square of the frequency (this can be derived from Weyl’s asymptotic law; see § A.6.1); so equipartition leads to the absurd conclusion that more energy is radiated at higher frequencies, so that the total radiated power is infinite. This apparent paradox was known as the “ultraviolet catastrophe”, and was resolved by Einstein’s quantum theory of light, although the correct law for radiative intensity was derived earlier by Planck who postulated that energy levels for light at fixed frequency should be quantized, but without proposing a mechanism.

Figure 4 depicts the relationship between intensity of radiation and frequency for the classical Rayleigh-Jeans law

$$I_{\text{RJ}}(\nu) = \frac{8\pi\nu^2}{c^3}k_B T$$

and for Planck's law

$$I_{\text{P}}(\nu) = \frac{8\pi\nu^2}{c^3} \frac{h\nu}{e^{h\nu/k_B T} - 1}$$

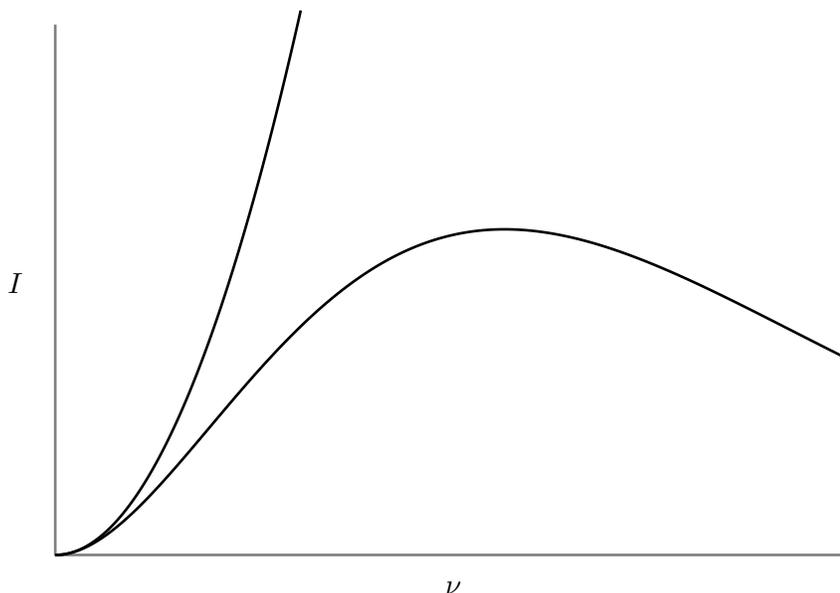


FIGURE 4. The two curves indicate the relationship between radiative intensity and frequency for a black body according to the Rayleigh-Jeans and Planck laws. The two laws agree asymptotically for low frequency.

2.5. Many-particle noninteracting quantum systems. Let us now consider quantum many-particle systems. Each system contains N identical particles, with no interactions between them. Each particle can be in any of a (discrete) set of states $\{j\}$ with energy ϵ_j . The state of the *system* therefore is a vector $\{n_j\}$ of integers, where n_j is the number of particles in state j . Because the particles are *identical*, the state is *completely specified* by the vector $\{n_j\}$.

There are restrictions on the possible values of the n_j depending on the spin of the particles:

- (1) if the particles are *bosons* (i.e. have integral spin) the n_j may be any non-negative integer $0, 1, 2, 3, \dots$ for each j ; whereas
- (2) if the particles are *fermions* (i.e. have half-integral spin) the n_j may be 0 or 1 for each j .

Thus in each case, the partition function has the form

$$Z(\beta) = \sum_{\{n_j\}: \sum n_j = N} e^{-\beta \sum n_j \epsilon_j}$$

Without the restriction $\sum n_j = N$ we could just write $\hat{Z}(\beta) := \prod_j \left(\sum_{n_j} e^{-\beta n_j \epsilon_j} \right)$; i.e.

$$\hat{Z}(\beta) = \prod_j \left(\frac{1}{1 - e^{-\beta \epsilon_j}} \right) \text{ (bosonic), or } \hat{Z}(\beta) = \prod_j (1 + e^{-\beta \epsilon_j}) \text{ (fermionic)}$$

One way to find the restriction to vectors $\{n_j\}$ with $\sum n_j = N$ is to introduce formal variables, and perform a contour integral which can be evaluated by the method of steepest descent. An analogous calculation is performed in § B; also, see Schrödinger [10].

A more physical method is to consider the system coupled to a heat bath which allows the transfer of both energy *and particles* from the reservoir. This leads to a model called the *grand canonical ensemble*, and this is the method we pursue, closely following Feynman [3].

Suppose it takes energy μ to move a particle from the box to the reservoir, where μ is a function of temperature, volume, etc. Up to sign, μ is the *chemical potential* (at least formally); compare with § 1.4.5. Since the number of particles in the system is not fixed, we concentrate instead on the expected number $\langle N \rangle$; increasing $\langle N \rangle$ will increase μ , so by adjusting μ we can adjust $\langle N \rangle$ to the desired value.

Since in the new model, particles have a “potential” to be in the reservoir, the new energy levels for a particle in the system in state j (taking this potential into account) is $\epsilon_j - \mu$. Thus we can define a new partition function, which is a function both of β and of μ , by

$$Z(\beta, \mu) := \sum_{\{n_j\}} e^{-\beta \sum n_j (\epsilon_j - \mu)}$$

where the sum is now taken over *all* allowable vectors $\{n_j\}$. If we define

$$G := -\frac{1}{\beta} \log Z(\beta, \mu)$$

then

$$e^{-\beta G} = Z(\beta, \mu) = \sum e^{-\beta \sum n_j (\epsilon_j - \mu)}$$

so that G is another kind of free energy, similar to the Helmholtz free energy; however because of its dependence on the chemical potential, it is more usual to identify it with the Gibbs free energy. Now,

$$\frac{\partial Z}{\partial \mu} = \sum \beta \left(\sum n_j \right) e^{-\beta \sum n_j (\epsilon_j - \mu)} = \sum \beta N e^{-\beta \sum n_j (\epsilon_j - \mu)}$$

Thus

$$\langle N \rangle = \frac{1}{Z} \sum N e^{-\beta \sum n_j (\epsilon_j - \mu)} = \frac{1}{\beta Z} \frac{\partial Z}{\partial \mu} = \frac{1}{\beta} \frac{\partial \log Z}{\partial \mu} = -\frac{\partial G}{\partial \mu}$$

2.5.1. *Quantum ideal Bose-Einstein gas.* For bosons (i.e. an ideal Bose-Einstein gas) we have

$$e^{-\beta G} = \sum e^{-\beta \sum n_j (\epsilon_j - \mu)} = \prod_j \frac{1}{1 - e^{-\beta (\epsilon_j - \mu)}}$$

or in other words,

$$G = \frac{1}{\beta} \sum_j \log (1 - e^{-\beta (\epsilon_j - \mu)})$$

If we consider particles restricted to a box of volume V , the density of states with momentum approximately equal to p is proportional to $V \cdot d^3p/(2\pi\hbar)^3$ (the constant of proportionality depends on the number of spin modes allowable). The energy of a particle of momentum p is $\epsilon := p^2/2m$; we approximate the sum above by an *integral*

$$G = s \cdot \frac{1}{\beta} \int_{\mathbb{R}^3} \log \left(1 - e^{-\beta p^2/2m} e^{\beta\mu} \right) \frac{d^3p}{(2\pi\hbar)^3} V$$

so the density ρ of particles is

$$\rho := \frac{\langle N \rangle}{V} = -\frac{1}{V} \frac{\partial G}{\partial \mu} = s \cdot \int_{\mathbb{R}^3} \frac{e^{-\beta p^2/2m} e^{\beta\mu}}{1 - e^{-\beta p^2/2m} e^{\beta\mu}} \frac{d^3p}{(2\pi\hbar)^3}$$

We evaluate this integral following Feynman. Define new variables

$$\alpha := e^{\beta\mu}, \quad x^2 = \frac{\beta|p|^2}{2m}$$

so that $|p|^2 = 2mx^2/\beta$ and $d|p| = (2m/\beta)^{1/2} dx$. Thus, the measure $d^3p/(2\pi\hbar)^3$ on \mathbb{R}^3 pushes forward to the measure $x^2/2\pi^2\hbar^3(2m/\beta)^{3/2} dx$ on $[0, \infty)$, and we obtain

$$\begin{aligned} \rho &= s \int_0^\infty \frac{e^{-x^2}\alpha}{1 - e^{-x^2}\alpha} \left(\frac{x^2}{2\pi^2\hbar^3} \left(\frac{2m}{\beta} \right)^{3/2} \right) dx \\ &= s \frac{1}{4\pi^2\hbar^3} \left(\frac{2m}{\beta} \right)^{3/2} \int_{-\infty}^\infty x^2 \left(\alpha e^{-x^2} + \alpha^2 e^{-2x^2} + \alpha^3 e^{-3x^2} + \dots \right) dx \\ &= s \left(\frac{mk_B T}{2\pi\hbar^2} \right)^{3/2} \left(\alpha + \frac{\alpha^2}{2^{3/2}} + \frac{\alpha^3}{3^{3/2}} + \dots \right) \end{aligned}$$

where we use the identity

$$\int_{-\infty}^\infty x^2 e^{-ax^2} dx = \frac{\sqrt{\pi}}{2} a^{-3/2}$$

which can be proved by integration by parts and a change of variables. With notation $\zeta_r(\alpha) := \sum_{n=1}^\infty \frac{\alpha^n}{r^n}$, this “simplifies” to

$$\rho = s \left(\frac{mk_B T}{2\pi\hbar^2} \right)^{3/2} \zeta_{3/2}(\alpha)$$

Thus, given ρ , we can find $\alpha = e^{\beta\mu}$ (and therefore also μ) by solving

$$\zeta_{3/2}(\alpha) = \frac{1}{2} \left(\frac{2\pi\hbar^2}{mk_B} \right)^{3/2} \frac{\rho}{T^{3/2}}$$

To compute U , we should ignore the chemical potential μ , since this does not contribute to the “internal” energy of the system, but only reflects the energy needed to move particles

between the system and the reservoir. Thus

$$\begin{aligned} U &= \frac{1}{Z} \sum \left(\sum n_j \epsilon_j \right) e^{-\beta \sum n_j (\epsilon_j - \mu)} = \frac{\partial \beta G}{\partial \beta} + \mu \langle N \rangle \\ &= s \int_{\mathbb{R}^3} \frac{(e^{-\beta p^2/2m} e^{\beta \mu}) p^2/2m}{1 - e^{-\beta p^2/2m} e^{\beta \mu}} V \frac{d^3 p}{(2\pi \hbar)^3} \\ &= s \frac{3}{2} k_B T \left(\frac{m k_B T}{2\pi \hbar^2} \right)^{3/2} V \zeta_{5/2}(\alpha) \end{aligned}$$

by a similar computation as before, using the identity

$$\int_{-\infty}^{\infty} x^4 e^{-ax^2} dx = \frac{3}{4} a^{-5/2} \sqrt{\pi}$$

In particular, for low ρ and high T (or both), the ratio $\rho/T^{3/2}$ is small, so $\zeta_{3/2}(\alpha)$ is small, and we have $\zeta_{3/2}(\alpha) \sim \alpha \sim \zeta_{5/2}(\alpha)$, so that

$$U \sim \frac{3}{2} k_B T \langle N \rangle$$

recovering the classical $C_V = 3/2$ for a monatomic ideal gas. However, for *small* T (with ρ fixed), the value of $\zeta_{3/2}(\alpha)$ increases with α , until as $\alpha \rightarrow 1$ we have

$$\zeta_{3/2}(1) = 2.612 \text{ and } \zeta_{5/2}(1) = 1.341$$

whereas at $\alpha > 1$, these functions diverge! Thus, T can only decrease to a *critical temperature* T_c where $\alpha = 1$, and

$$\frac{\langle N \rangle}{V} = \rho = s \left(\frac{m k_B T_c}{2\pi \hbar^2} \right)^{3/2} 2.612$$

For $T < T_c$ our formulae become nonsensical; but remember that we approximated the formula for G (given by the Boltzmann distribution) by an integral, which is only valid when the proportion of particles in any given state is negligible compared to $\langle N \rangle$. Below the critical temperature,

$$\langle n_0 \rangle := \frac{\partial G}{\partial \epsilon_0} = \frac{1}{e^{-\beta(\mu - \epsilon_0)} - 1}$$

is comparable in order of magnitude to $\langle N \rangle$. On the other hand, if we assume that ϵ_0 is the unique lowest energy level,

$$\langle n_1 \rangle = \frac{1}{e^{-\beta(\mu - \epsilon_1)} - 1} \ll \langle n_0 \rangle$$

since T_c is very small (and β is big). Thus it is reasonable to approximate the contribution *from terms with energy bigger than ϵ_0* by an integral. This leads to the approximation

$$U = \frac{3}{2} k_B T \cdot (0.5134) \left(\frac{T}{T_c} \right)^{3/2} \langle N \rangle$$

so that there is a *phase transition* for C_V , which becomes non-differentiable at the critical temperature.

The phenomenon whereby a definite fraction of the particles are in the lowest energy state below the critical temperature is known as *Bose-Einstein condensation*.

2.5.2. *Fermi gas.*2.5.3. *Charge screening in metals.*

2.6. Lattice models. An important class of models — in fact the most important from the point of view of applications to dynamics — are lattice models. In this case the states Ω of the system are functions $X \rightarrow \Sigma$ where Σ is finite, and X is some discrete space, usually with a group action. Here one has in mind (for instance) a crystal lattice, in which the space X (the “sites”) denote the location of the particles in the lattice, and Σ parameterizes the states of the individual particles. The group action enforces the (spatial) homogeneity of the model.

The most basic example is that X is itself a finitely generated abelian group; in the simplest case something like $(\mathbb{Z}/n\mathbb{Z})^d$ where one takes d fixed and n large (and considers the limit as $n \rightarrow \infty$). The group G acts on itself by multiplication, and hence on states Ω by $g \cdot \omega(h) = \omega(h + g)$ (we write the group operation additively since G is abelian).

The action of G on Ω induces an action on the space of probability measures on Ω , by $g_*\nu(A) = \nu(g^{-1}A)$. If G is abelian, it is amenable, and therefore (even when infinite) preserves many probability measures on Ω . If ν is invariant, then for any function f on Ω (integrable with respect to ν),

$$\langle f \circ g \rangle_\nu := \int_\Omega f \circ g \, d\nu = \int_\Omega f \, dg_*\nu = \int_\Omega f \, d\nu = \langle f \rangle_\nu$$

We consider energy functions ϵ on Ω which are obtained by *summing local contributions*. This means there is some function $\psi : \Omega \rightarrow \mathbb{R}$ so that

$$\epsilon(\omega) = \sum_{g \in G} \psi(g \cdot \omega)$$

For energy functions of this kind, and for any G -invariant probability measure ν , (assuming G finite) we have

$$\begin{aligned} -\beta A_\nu &:= \frac{S(\nu)}{k_B} - \beta \langle \epsilon \rangle_\nu \\ &= \frac{S(\nu)}{k_B} + \sum_\omega -\nu(\omega) \sum_G \beta \psi(g \cdot \omega) \\ &= \frac{S(\nu)}{k_B} + \sum_\omega -\nu(\omega) |G| \beta \psi(\omega) \end{aligned}$$

So we get a formula, valid for G -invariant ν ,

$$\frac{-\beta A_\nu}{|G|} = \frac{S(\nu)}{k_B |G|} + \langle -\beta \psi \rangle_\nu$$

and we deduce (at fixed β) that this last expression is maximized (amongst G -invariant probability measures ν) by the Boltzmann distribution (which is evidently G -invariant). The terms $S(\nu)/k_B |G|$ and $\langle -\beta \psi \rangle_\nu$ can be thought of respectively as the contribution to entropy, and the contribution to energy, *per lattice site*.

Thus the Boltzmann distribution solves a variational problem *amongst G -invariant probability measures*, which can be expressed purely in terms of *local* contributions.

Example 2.4 (Bernoulli). The simplest example of a lattice model is the Bernoulli model, in which $\psi(\omega)$ depends only on $\omega(0)$ (the value at a single site); i.e. $\psi(\omega) = f(\omega(0))$ for some $f : \Sigma \rightarrow \mathbb{R}$. There is no interaction between different sites, so the partition function factorizes into a product of local contributions. Writing

$$q_\beta(s) := \frac{e^{-\beta f(s)}}{\sum_{s \in \Sigma} e^{-\beta f(s)}}$$

we get

$$Z(\beta) = \left(\sum_{s \in \Sigma} e^{-\beta f(s)} \right)^{|G|}$$

and the Boltzmann distribution becomes

$$\Pr(\omega) = \frac{e^{-\beta \sum_g f(\omega(g))}}{Z(\beta)} = \prod_{g \in G} q_\beta(\omega(g))$$

i.e. the sites are independently and identically distributed with the distribution q_β .

2.7. Ising model. The *Ising model* is a lattice model on a d -dimensional grid $(\mathbb{Z}/n\mathbb{Z})^d$ in which the local contribution to energy has the form

$$\psi(\omega) = \sum_{d(e,0)=1} -\omega(0)\omega(e) + m\omega(0)$$

Here the notation $d(e, 0) = 1$ means the lattice points e adjacent to 0. The qualitative behaviour of this model depends very strongly on dimension, and exact solutions are known only in dimension 1 and 2.

The constant m indicates the potential in the presence of an external magnetic field. By varying this magnetic field (formally, by taking partial derivatives of the partition function with respect to m) we can determine the magnetization of the material.

2.7.1. 1-dimensional case. The exact solution of the 1-dimensional Ising model is very simple, and was accomplished by Ising himself. We enumerate the lattice points in cyclic order as $0, 1, \dots, (n-1)$, and then the partition function has the form

$$Z(\beta, m) = \sum_{\omega} e^{-\beta(\sum_{j=0}^{n-1} -\omega(j)\omega(j+1) + m\omega(j))}$$

To compute the behaviour of this function for large n , we introduce a certain matrix L called a *transfer matrix*.

We can define a directed graph Γ with two vertices, corresponding to the two possible spins -1 and 1 , with one directed edge from each vertex to each other vertex (including two self-edges at each of the vertices). States ω of the system are in bijection with (based) loops in Γ of length n : the vertices visited by the loops correspond to the spins of ω at the lattice points in cyclic order. We can decorate each directed edge with the local contribution it makes to the energy of the state ω , so that the total energy is the product of the $e^{-\beta \cdot \text{edge weight}}$ terms.

The partition is thus a sum over all loops of products associated to the edge weights. This may be elegantly expressed in terms of a *transfer matrix*. This is a 2×2 matrix L

whose entries correspond to the four directed edges:

$$L = \begin{pmatrix} e^{-\beta(-1+m)} & e^{-\beta(1+m)} \\ e^{-\beta(1-m)} & e^{-\beta(-1-m)} \end{pmatrix}$$

And then the partition function is just a trace of a power of L :

$$Z(\beta, m) = \text{Tr } L^n$$

But the trace of L^n is just the sum $\lambda_1^n + \lambda_2^n$ where the λ_i are the eigenvalues of L which are given by the quadratic formula

$$\lambda_1, \lambda_2 = e^\beta \cosh \beta m \pm e^\beta \sqrt{\sinh^2 \beta m - e^{-4\beta}}$$

2.7.2. 2-dimensional case. The Ising model in 2 dimensions is very different from the model in 1 dimension. The biggest difference is the existence of a *phase transition* between disordered states (at low β) and ordered states (at high β). At the critical $\beta = 0.440686793$ the solutions exhibit a remarkable *scale invariance*. The difference between ordered and disordered states can be expressed in terms of the correlation between spins at distinct lattice points.

Since the mean of every random variable $\omega(x)$ is zero, the correlation at two lattice points x, y , is just the expectation of $\omega(x)\omega(y)$. The correlation depends (up to order of magnitude) only on the distance from x to y . Roughly speaking, the phases are

- ordered: $\langle \omega(x)\omega(y) \rangle \sim \text{constant} > 0$;
- disordered: $\langle \omega(x)\omega(y) \rangle \sim e^{-Cd(x,y)}$; and
- critical: $\langle \omega(x)\omega(y) \rangle \sim d(x,y)^\alpha$.

The behaviour at the critical point is evidently the most interesting. The exponent α at the critical point is often independent of many features of the model, and is fixed by an emergent symmetry of the system at the critical point, known as *renormalizability*.

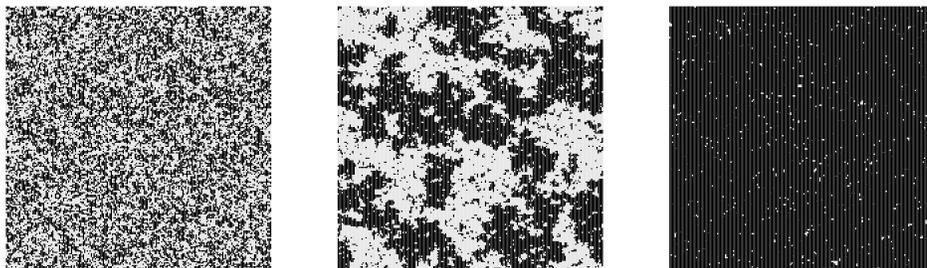


FIGURE 5. Typical states in the Ising model on a 200×200 grid at $\beta = 0.2$, at the critical $\beta = 0.440686793$, and at $\beta = 0.6$, sampled using the Metropolis algorithm run for 4×10^8 steps.

3. SHANNON'S THEORY OF INFORMATION

3.1. Entropy as information.

Example 3.1 (Bayesian statistics). One application of the variational definition of the Boltzmann distribution, as the probability distribution maximizing entropy, is to Bayesian statistics. Suppose we have an unknown probability distribution, but suppose we know some facts about the distribution; e.g. the mean, or the variance. Our *a priori* best guess for the distribution should make as few assumptions as possible. Another way of saying this is that further experiments should give as much information as possible, so that we can falsify our guess quickly if it is wrong. The expected amount of information obtained by randomly sampling a distribution is given by the entropy (in the case of a continuous probability distribution $p(t)dt$ this is given by the integral $-\int p(t) \log p(t)dt$). The distribution maximizing entropy subject to some constraint can therefore be found as the Boltzmann distribution for a suitable energy functional. Two natural examples are:

- (1) The mean of a distribution $p(t)$ is $\int t p(t)dt$. Suppose for simplicity that we consider probability measures supported on $[0, \infty)$. We can consider a system with states distributed like dt on $[0, \infty)$, with energy $\epsilon(t) = t$. Then the exponential distribution $C.e^{-\beta t} dt$ maximizes entropy for a fixed mean.
- (2) The variance of a distribution $p(t)$ with mean 0 is $\int t^2 p(t)dt$. We can consider a system with states distributed like dt on $(-\infty, \infty)$, with energy $\epsilon(t) = t^2$. Then the Gaussian distribution $C.e^{-\beta t^2} dt$ maximizes entropy for mean 0 and fixed variance.

3.2. Infinitude of primes. The perspective of information theory is useful in mathematics, and gives us new ways to look at the familiar. For example, one can use information theory to give a charming proof of the infinitude of primes.

Every positive integer has a unique representation as a product of primes; i.e. it can be written uniquely in the form

$$n = 2^{d_1(n)} \cdot 3^{d_2(n)} \cdot \dots \cdot p_k^{d_k(n)} \cdot \dots$$

where p_k is the k th prime, and $d_k(n)$ is the number of times p_k divides n . Obviously, each $d_k(n)$ satisfies $d_k(n) \leq \log_{p_k}(n) \leq \log_2(n)$. Thus, if there were only finitely many (say N) primes, we could recover any number of size at most n uniquely from a list of N non-negative integers, each of size at most $\log_2(n)$. But there are only $\log_2(n)^N$ such lists of integers, and $\log_2(n)^N \ll n$ for big n , for any fixed N , giving a contradiction.

This idea can be used to give an implicit estimate of the growth rate of $\pi(n)$, the number of primes smaller than or equal to n . If we choose a random positive integer m less than some cutoff n , then $d_k(m)$ is a random variable. The exact distribution depends on n , but taking $n \rightarrow \infty$ for fixed k , the distribution of $d_k(m)$ converges very quickly (in measure) to the geometric distribution

$$\Pr(d_k(m) = j) = (p_k - 1)p_k^{-j-1}$$

The entropy of this distribution is

$$\begin{aligned} H_k &:= \sum_j -(p_k - 1)p_k^{-j-1} \log((p_k - 1)p_k^{-j-1}) \\ &= -\log(p_k - 1) + \sum_j (j + 1)p_k^{-j-1}(p_k - 1) \log p_k \end{aligned}$$

which is of order $(\log p_k)/p_k$. The random variables $d_k(m)$ all become independent as $n \rightarrow \infty$, so the entropy of the joint distribution of the d_k is approximately the sum of the H_k . On the other hand, the entropy of the choice of $m \leq n$ is just $\log n$. So we obtain an estimate of the form

$$\sum_{k=1}^{\pi(n)} \frac{\log p_k}{p_k} \sim \log n$$

(we are imposing a cutoff on the number of terms to consider, using the naive estimate that $d_k(m) = 0$ for $k > \pi(n)$, so that we should really set $H_k = 0$ for such p_k).

As a check, note that the prime number theorem gives $p_k \sim k \log k$ and $\pi(n) \sim n/\log n$, so the left hand side is of order

$$\sum_{k=1}^{n/\log n} \frac{\log k + \log \log k}{k \log k} \sim \sum_{k=1}^{n/\log n} \frac{1}{k} \sim \log(n/\log n) \sim \log n$$

3.3. Non-negative matrices. Suppose W is a real $n \times n$ matrix with all entries positive. The most important fact about such matrices is the

Theorem 3.2 (Perron-Frobenius Theorem). *Let W be a real matrix with all entries strictly positive. Then W has a unique eigenvalue λ of biggest absolute value. Furthermore, λ is real and positive, and the generalized eigenspace is 1-dimensional, and is spanned by an eigenvector v with all entries strictly positive.*

If W merely has non-negative entries, there is still a real positive eigenvalue λ whose absolute value is *at least* as big as that of every other eigenvalue, and an eigenvector v with all entries non-negative; these facts follow from the Perron-Frobenius theorem and continuity.

Proof. Consider the linear action of W on \mathbb{R}^n . Since all the entries are strictly positive, this action takes the positive cone \mathbb{R}_+^n strictly inside itself. The projectivization of the positive cone is a simplex Δ , so the Brouwer fixed point theorem already gives a fixed point in Δ , i.e. an eigenvector v with all coordinates strictly positive.

In fact, a more careful analysis shows that W fixes a *unique* point in the simplex. To see this, it is useful to introduce the *Hilbert metric* on Δ . The Hilbert metric is defined on any bounded convex subset Δ of projective space. For any two points $p, q \in \Delta$, there is a maximal straight line segment ℓ containing p and q , meeting $\partial\Delta$ at p_∞, q_∞ so that p_∞, p, q, q_∞ appear in that order in ℓ . There is a unique projective map τ taking ℓ to $[0, \infty]$ and taking p to 1; the distance from p to q is then $\log \tau(q)$.

The essential properties of the Hilbert metric are nondegeneracy (it is a genuine metric), and a strict contraction property: if $\Delta' \subset \Delta$ is a convex subset, and d', d denote the

Hilbert metrics on Δ' and Δ respectively, then $d'(p, q) \geq d(p, q)$ for $p, q \in \Delta'$ on the line subsegment $\ell' = \ell \cap \Delta'$, with equality if and only if $\ell' = \ell$.

Now, because W takes Δ strictly inside itself, it acts as a strict contraction of the Hilbert metric, and therefore has a unique fixed point in the interior. Furthermore, the eigenvalue λ associated to the fixed point v is strictly bigger in absolute value than any other eigenvalue μ . For, if $|\mu| = \lambda$, the projective action of W would not be strictly contracting at v on the projectivization of the subspace spanned by v and the (generalized) eigenspace of μ . \square

Remark 3.3. The contraction property of the Hilbert metric is the real projective analog of the Schwarz lemma in complex analysis, and the Hilbert metric is the analog of the Kobayashi metric.

Note that if W has positive entries, the same is true of its transpose. Thus a positive matrix has *left* and *right* Perron-Frobenius eigenvectors (unique up to scale) with the same eigenvalue.

3.4. Directed graphs and regular languages. A directed graph is just a graph with a choice of orientation on each edge. We consider “generalized” graphs, possibly with self-loops, and multiple (directed) edges between two vertices.

A *topological Markov chain* is just a directed graph Γ , together with a choice of an *initial vertex* v . Associated to Γ is the set of *directed simplicial paths* in Γ starting at v . The *adjacency matrix* M is the matrix whose ij entry is the number of directed edges from i to j .

A directed graph is *transitive* if there is a nontrivial directed path from every vertex to every other vertex. It is *aperiodic* if the greatest common divisor of the lengths of closed loops is equal to 1. If it is transitive and aperiodic, we say it is *mixing*. A directed graph is transitive if and only if for every ij there is an n so that $(M^n)_{ij} > 0$. A directed graph is mixing (transitive and aperiodic) if and only if there is an n so that $(M^n)_{ij} > 0$ for all ij . In particular, M^n is a positive matrix, and has Perron-Frobenius left and right eigenvectors with the same eigenvalue. In the sequel we will generally (though not always) restrict attention to mixing graphs, for simplicity.

In general, we can define an equivalence relation \sim on vertices, where $i \sim j$ if there are directed paths from i to j and from j to i . There is a new graph associated to Γ , whose vertices are the equivalence classes of Γ , and with an edge between vertices if there was a directed edge between representatives of the equivalence classes; note that this new graph has the property that there are *no directed loops*. Each equivalence class X in Γ has its own adjacency matrix M_X with a Perron-Frobenius eigenvalue λ_X ; the *dominant* equivalence classes X are those for which $\lambda_X = \lambda$, the Perron-Frobenius eigenvalue of M .

If Γ is a directed graph, we can label the edges with symbols from a finite alphabet Σ . The labeling is *deterministic* if there is at most one outgoing edge from any given vertex with any given edge label. If the edges are labeled, each finite directed path starting at v determines a finite word in the alphabet Σ , by reading the edge labels associated to the path in order. If the labeling is deterministic, the map from paths to words is injective. We further can choose some subset of the vertices of Γ and declare them to be *accept states*; such a choice determines a subset of the directed paths — those that start at v and end at an accept state.

A deterministic labeling and choice of a set of accept states as above determines a subset $L \subset \Sigma^*$ of the set of finite strings (i.e. words) in the alphabet Σ ; it is the *language* associated to the (decorated) graph Γ . There is a bijection between words in L and directed paths in Γ starting at v and ending at an accept state. Subsets L of Σ^* of this form are called *regular languages*. Note that a given language may be produced by (infinitely) many different graphs and/or labelings. Note further that regular languages which are prefix-closed (i.e. if $w \in L$, any prefix of w is in L) are exactly those associated to directed graphs in which every vertex is an accept state.

Example 3.4. Regular languages are precisely those languages corresponding to *regular expressions* in the sense of computer science (e.g. as used widely in the UNIX operating system). The “state” of a parser which recognizes terms matching a regular expression corresponds to a vertex of a graph Γ , as it reads the letters one by one. Since Γ is finite, the parser has *bounded* memory; this limitation means that many simple languages are not regular. For example, the language $L \subset \{(,)\}$ consisting of correctly paired parentheses, is not regular; it is rather an example of a *context-free* language, at the next level of complexity in the Chomsky hierarchy.

Let L be a regular language, and let L_n denote the set of words of length n . We can form a generating function

$$S_L(z) := \sum_{n=0}^{\infty} |L_n| z^n$$

Proposition 3.5 (Regular language has rational generating function). *Suppose L is regular. Then the generating function $S_L(z)$ is a rational function $q(z)/p(z)$, whose denominator divides the characteristic polynomial of the adjacency matrix of any directed graph Γ parameterizing L .*

Proof. Let M denote the adjacency matrix of a directed graph Γ as above. Let ι denote the vector with a 1 entry correspond to the initial vertex, and 0 elsewhere, and let α denote the vector with 1 entries at the accept states, and 0 elsewhere. Evidently, $|L_n| = \iota^T M^n \alpha$. If p is the characteristic polynomial of M , then $p(M) = 0$, so we obtain the identity

$$0 = \iota^T M^{n-d} p(M) \alpha = p_0 |L_n| + p_1 |L_{n-1}| + \cdots + p_d |L_{n-d}|$$

where the p_i are the coefficients of p . It follows that the product $S_L(z)p(z)$ vanishes in degree $\geq d$, so that it is a *polynomial* $q(z)$. But then $S_L(z)$ is a ratio of polynomials, which is to say, it is rational, of the form $S_L(z) = q(z)/p(z)$, possibly not in reduced form. \square

3.5. Channels and coding. The material in this section borrows very heavily from Shannon’s seminal paper [11], especially § 1.

3.5.1. Capacity of a channel. A *discrete communications channel* is a directed graph with a deterministic edge labeling by some finite alphabet Σ , and a *transmission time* $t_{ij,s} > 0$ associated to each directed edge. For simplicity, we consider channels in which every vertex of the underlying graph is an accept state; i.e. those corresponding to prefix-closed regular languages. The channel can transmit exactly the regular language of words associated to the underlying graph, and the time taken to transmit a word is the sum of the transmission times associated to the edges in the corresponding directed path.

Definition 3.6. The *capacity* C of the channel is the limit

$$C := \lim_{t \rightarrow \infty} \frac{1}{t} \log N(t)$$

where $N(t)$ is the number of different signals that can be sent in time t .

We define a matrix $W(x)$ (a “weighted” adjacency matrix) whose ij entry is the sum

$$W(x)_{ij} = \sum_s e^{-xt_{ij;s}}$$

where x is a formal variable, whose significance we will see in a moment. The capacity of the channel can be obtained by spectral theory:

Proposition 3.7 (Capacity of a channel). *The capacity of a channel is the number x for which the largest real eigenvalue of $W(x)$ (i.e. the Perron-Frobenius eigenvalue) is 1.*

Proof. We prove this under the assumption that the underlying graph is mixing — i.e. transitive and aperiodic. For simplicity, assume that the $t_{ij;s}$ are all positive integers. We build a generating function

$$S(z) := 1 + N(1)z + N(2)z^2 + \dots$$

Then we have

$$S(z) = \iota^T \sum_{n=0}^{\infty} W(x)^n \alpha$$

where $z = e^{-x}$, and where α is the vector with every entry equal to 1 (since we are assuming for convenience that every state is an accept state). Then $S(z)$ converges for $|z|^{-1} > \limsup N(n)^{1/n}$, i.e. for $\log(|z|) < -C$ by the definition of C .

On the other hand, by Perron-Frobenius, $\sum W(x)^n$ converges provided the Perron-Frobenius eigenvalue is < 1 . Thus C is exactly the value of x for which the Perron-Frobenius eigenvalue of $W(x)$ is 1, as claimed. This proves the proposition if the $t_{ij;s}$ are all integers.

The same argument evidently works if the $t_{ij;s}$ are all rational, since we can just scale the units of time so that they are integers. Now, it is obvious that C is monotone decreasing as a function of each individual $t_{ij;s}$. In particular, both C and x (as defined in the statement of the proposition) vary continuously with the $t_{ij;s}$. Since the proposition is proved for rational $t_{ij;s}$, it follows for arbitrary $t_{ij;s}$ by continuity. \square

3.5.2. Markov chains. A *Markov chain* is an underlying topological Markov chain, together with an assignment of non-negative numbers $p_{ij;s}$ to each directed edge, so that for any vertex i , the p define a probability distribution on the set of outgoing edges that start at i ; i.e. so that $\sum_{j,s} p_{ij;s} = 1$ for all i . A Markov chain determines a random walk in the directed graph, starting at v , so that if the walk is at vertex $i(n)$ after n steps, it transitions along the edge labeled s from i to j with probability $p_{ij;s}$.

We can define a transition matrix L with $L_{ij} = \sum_s p_{ij;s}$. If every $p_{ij;s}$ is positive, and if the underlying topological Markov chain is mixing, then some power of L is a positive matrix, and there are unique left and right Perron-Frobenius eigenvectors (up to scale). Note that the Perron-Frobenius eigenvalue of L is 1, so that if P is a left PF eigenvector

(normalized to have entries summing to 1), we have $PL = P$. Note further that for *any* vector Q , the iterates QL^n converge to a constant multiple of P . We think of P as a probability distribution on vertices, and note that $P_j N$ is just the expected number of times we visit the vertex j in a random walk of length N .

If the Markov chain is in state i , the entropy of the next symbol (i.e. the next edge in the walk) is $-\sum_{j,s} p_{ij;s} \log p_{ij;s}$. Thus the average entropy per symbol in the walk is $\sum_{i,j,s} P_i p_{ij;s} \log p_{ij;s}$.

If we also assign transmission times $t_{ij;s}$ to edges, then we can talk about the *rate* at which the Markov chain generates entropy:

$$R := \frac{\text{average entropy per symbol}}{\text{average time per symbol}} = \frac{-\sum_{i,j,s} P_i p_{ij;s} \log p_{ij;s}}{\sum_{i,j,s} P_i p_{ij;s} t_{ij;s}}$$

3.5.3. *Maximum rate of entropy production.* Now let's consider a fixed channel, and ask: how can we assign transition probabilities to edges in order to maximize the rate of entropy production? There is a very satisfying answer:

Proposition 3.8 (Maximum entropy per unit time over a channel). *The maximum rate R at which a Markov chain whose topology is described by a given channel can produce information is equal to the capacity C of the channel.*

Proof. We prove the proposition under the hypothesis that the underlying topological Markov chain is mixing. Let B be a right Perron-Frobenius eigenvector for $W = W(C)$ of eigenvalue 1; i.e. such that

$$B_i = \sum_{j,s} e^{-Ct_{ij;s}} B_j$$

(note that such a B exists, since the Perron-Frobenius eigenvalue of W is 1) and define

$$p_{ij;s} := e^{-Ct_{ij;s}} B_j / B_i$$

We check that this really defines transition probabilities:

$$\sum_{j,s} p_{ij;s} = \sum_{j,s} e^{-Ct_{ij;s}} B_j / B_i = B_i / B_i = 1$$

and furthermore (under the mixing assumption) every B_j is positive, since B is a Perron-Frobenius eigenvector, so the $p_{ij;s}$ are finite and positive. We calculate the rate of entropy production:

$$\begin{aligned} R &= \frac{-\sum P_i p_{ij;s} \log (e^{-Ct_{ij;s}} B_j / B_i)}{\sum P_i p_{ij;s} t_{ij;s}} \\ &= \frac{C \sum P_i p_{ij;s} t_{ij;s} - \sum P_i p_{ij;s} \log B_j + \sum P_i p_{ij;s} \log B_i}{\sum P_i p_{ij;s} t_{ij;s}} \end{aligned}$$

Now, since $\sum_{j,s} p_{ij;s} = 1$ for each i , we have $\sum_{i,j,s} P_i p_{ij;s} \log B_i = \sum_i P_i \log B_i$. On the other hand, for each j we have $\sum_{i,s} P_i p_{ij;s} = P_j$ since P is a left eigenvalue for the matrix L with $L_{ij} = \sum_s p_{ij;s}$. Thus $-\sum P_i p_{ij;s} \log B_j + \sum P_i p_{ij;s} \log B_i = 0$ and $R = C$, as claimed. \square

This is a remarkable fact; it implies that there is a way to assign transition probabilities to edges so that the stationary Markov process (whose statistical properties are extremely easy to analyze) produces a measure on the set of signals taking an arbitrary (fixed) time to transmit which is very close (in the sense of information) to the uniform measure. Since the latter probability measure is typically the one of interest, this gives an incredibly powerful tool for its analysis.

The analog of this fact in the world of 1-dimensional (smooth) dynamics is the phenomenon that the Gibbs measure associated to a suitable potential function (the logarithm of the unstable Jacobian) is non-singular with respect to Lebesgue measure; in the theory of hyperbolic or Kleinian groups it is the phenomenon that the Patterson-Sullivan measure on the visual boundary is quasi-equivalent to Hausdorff measure.

3.5.4. *Efficient codes.* We suppose we have an information source, consisting of a mixing (i.e. transitive and aperiodic) Markov chain, with transition probabilities $p_{ij;s}$ generating words in Σ^* . If P is a left Perron-Frobenius eigenvector for the transition matrix L (i.e. $PL = P$) then the information source produces entropy at the average rate of

$$H := \sum_i P_i \left(\sum_{s,j} -p_{ij;s} \log p_{ij;s} \right)$$

per symbol. In other words, if ν_N is the probability measure on words of length N produced by the Markov chain, then $S(\nu_N)/N \rightarrow H$.

Lemma 3.9. *Given any positive ϵ, δ there is an integer N so that the sequences produced by the information source of any fixed length $n \geq N$ fall into two classes:*

- (1) *a set of bad sequences with total probability $< \epsilon$; and*
- (2) *a set of good sequences, each good sequence ξ occurring with a probability $p(\xi)$ satisfying*

$$\left| \frac{\log p(\xi)^{-1}}{n} - H \right| < \delta$$

Moreover, there are at most $e^{(H+\delta)n}$ good sequences, and at most $|\Sigma|^n$ bad sequences.

Proof. First assume that the Markov chain is transitive. Then in a random sequence of length n we expect to be in state i approximately $P_i n$ times, and to transition from i to j via an edge labeled s approximately $P_i p_{ij;s} n$ times. The law of large numbers for transitive Markov chains says that the probability of a deviation of order $\pm \delta n$ in one of these (finitely many) quantities is less than ϵ when n is sufficiently large; a sharper estimate is given by the *Chernoff inequality* — see § C.

Thus, for a subset of sequences accounting for $(1 - \epsilon)$ of the total probability (the *good* sequences), each sequence ξ occurs with probability

$$p(\xi) := \prod_{i,j,s} p_{ij;s}^{(P_i p_{ij;s} \pm \delta)n}$$

giving the desired estimate. An upper bound on the number of good sequences follows immediately from a lower bound on the probability of each sequence. An upper bound on the number of bad sequences is simply the total number of possible sequences of length n in the alphabet Σ , which is $|\Sigma|^n$.

If the Markov chain is not transitive, it is nevertheless true that all but exponentially few paths spend most of their time in the maximal transitive subgraphs of the Markov chain whose adjacency matrices have Perron-Frobenius eigenvalue equal to that of the graph as a whole. There are finitely many such subgraphs, and each one can be visited at most once. So apply the estimate to the subpaths in each individual subgraph. \square

The fundamental theorem of coding in a noiseless channel is the following.

Theorem 3.10 (Coding over a noiseless channel). *Suppose we have an information source driven by a Markov chain that produces random words in some alphabet, with an average entropy of H per symbol, and suppose we want to send this information over a fixed channel. Then for any positive ϵ , there is a coding which lets us send information supplied by the source over the channel at the average rate of $C/H - \epsilon$ symbols per unit time, but it is not possible to transmit at an average rate bigger than C/H symbols per unit time.*

Proof. The upper bound is immediate from the definition of the capacity of the channel. So it remains to show that there is a code which comes close to realizing this maximum.

We partition sequences of length n produced by the information source into good and bad sequences, as in Lemma 3.9. The channel can send $e^{(H+\epsilon)n}$ different sequences in times $(H+2\epsilon)n/C$, so we can code the good sequences arbitrarily as some subset of these, keeping one over as a “start/stop” symbol. Code the bad sequences arbitrarily as sequences which can be sent over the channel in time $(\log |\Sigma| + \epsilon)n/C$, and use the start/stop symbol as punctuation before or after the transmission of each bad sequence. \square

Remark 3.11. The closer ϵ is to zero, the bigger we must choose the size of the blocks n for the code. Likewise, the size of the blocks n depends on the constant d appearing in the Chernoff inequality, and this depends in turn on the mixing time for the Markov chain; in fact it is proportional to λ_1 , the smallest nonzero eigenvalue of the matrix $\text{Id} - (p + p^*)/2$ where p is the matrix with entries p_{ij} ; see e.g. Lezaud [6]. As a rough order of magnitude, it takes $O(1/\lambda_1)$ random steps on Γ for any initial probability distribution on states to come to within a ratio of $1/2$ of the stationary distribution (in L^1).

Remark 3.12. The essential point of Theorem 3.10 is that the only property of an information source that determines the average rate at which it can be sent over a fixed channel is the entropy; informally, at the level of information content, the entropy is the *only* essential variable. If we think of coding schemes as morphisms between information sources, and lossless (and hence invertible) coding schemes as isomorphisms, then we can think of this theorem as saying that information sources with the same entropy per symbol are isomorphic. Shannon’s coding theorem inspired work of Kolmogorov, Sinai and many others in ergodic theory, culminating in the famous *Ornstein isomorphism theorem* [8], which says that two Bernoulli shifts are measurably isomorphic if and only if they have the same entropy.

APPENDIX A. CLASSICAL AND QUANTUM MECHANICS

A.1. Legendre transform. Let V be a finite dimensional real vector space, and $F : V \rightarrow \mathbb{R}$ a strictly convex function; i.e. a function for which the Hessian $\frac{\partial^2 F}{\partial v_i \partial v_j}(p)$ is positive

definite for all $p \in V$. Equivalently, the region above the graph $\Gamma(F) := \{(F(p), p)\}$ in $\mathbb{R} \times V$ is strictly convex.

If F is strictly convex, the map $L_F : V \rightarrow V^*$ defined by $L_F(p) = dF_p \in T_p^*V \cong V^*$ is a diffeomorphism onto its image S_F , which is an open and convex subset of V^* called the *stability set*.

Definition A.1. The *Legendre transform* of F , denoted F^* , is the function $F^* : S_F \rightarrow \mathbb{R}$ whose value at $\ell \in S_F$ is equal to $-\min_{p \in V} F(p) - \ell(p)$.

Note by strict convexity and the definition of S_F , the point $p \in V$ where $F(p) - \ell(p)$ achieves its minimum is the unique point where $dF_p = \ell$.

Equivalently, $p \in V$ is the unique point at which the tangent hyperplane to $\Gamma(F)$ at $(F(p), p)$ has slope ℓ , and $-F^*(\ell)$ is the intersection of this tangent hyperplane with the line $\{v = 0\} \subset \mathbb{R} \times V$. Thus for $\ell = dF_p$ we have the identity

$$F^*(\ell) + F(p) = \langle \ell, p \rangle$$

In this symmetric form it is straightforward to see that the Legendre transform is an involution when restricted to the class of strictly convex functions F for which $S_F = V^*$.

A.2. Euler-Lagrange equations. Suppose M is a manifold, and $\mathcal{L} : TM \rightarrow \mathbb{R}$ a function (the *Lagrangian*). The *action* of a smooth map $\gamma : [0, 1] \rightarrow M$ is the integral

$$\mathcal{A}(\gamma) := \int_0^1 \mathcal{L}(\gamma'(t)) dt$$

and it is a fundamental problem to minimize (or at least find the critical points of) \mathcal{A} on the family of smooth maps with fixed boundary conditions $\gamma(0) = p$ and $\gamma(1) = q$.

Example A.2 (Riemannian metric). Suppose M is a Riemannian metric. Then we can define $\mathcal{L}(\gamma') = \frac{1}{2} \|\gamma'\|^2$, in which case $\mathcal{A}(\gamma)$ is the *energy* of γ . Critical points of energy are precisely the geodesics for the metric (with prescribed endpoints).

Suppose we choose local coordinates x_i, v_i on TM , where $v_i = \frac{\partial}{\partial x_i}$, so that we can write \mathcal{L} as $\mathcal{L}(x_i, v_i)$. In a local chart, we can write $x_i(t) := x_i(\gamma(t))$ and $\dot{x}_i(t) := v_i(\gamma'(t))$. A smooth variation $\gamma(s)$ of γ is given in local coordinates by $(x_i + s\xi_i, \dot{x}_i + s\dot{\xi}_i)$. A map γ is a critical point for the action if for every such local variation we have

$$\begin{aligned} 0 &= \left. \frac{d}{ds} \right|_{s=0} \mathcal{A}(\gamma(s)) = \int_0^1 \sum_i \xi_i \frac{\partial \mathcal{L}}{\partial x_i}(\gamma'(t)) + \sum_i \dot{\xi}_i \frac{\partial \mathcal{L}}{\partial v_i}(\gamma'(t)) dt \\ &= \int_0^1 \sum_i \xi_i \left(\frac{\partial \mathcal{L}}{\partial x_i}(\gamma'(t)) - \frac{d}{dt} \left(\frac{\partial \mathcal{L}}{\partial v_i}(\gamma'(t)) \right) \right) dt \end{aligned}$$

where the last step uses integration by parts and the boundary condition. Thus we deduce the *Euler-Lagrange equations*

$$\frac{\partial \mathcal{L}}{\partial x_i} = \frac{d}{dt} \frac{\partial \mathcal{L}}{\partial v_i}$$

along a critical trajectory γ for the action.

A.3. Hamiltonian formulation. Suppose the Lagrangian $\mathcal{L} : TM \rightarrow \mathbb{R}$ is fiberwise strictly convex. We can define a new function H called the *Hamiltonian*, defined on a subset of T^*M fiberwise equal to the stability set of \mathcal{L} on the corresponding fiber of TM , by the Legendre transform.

We have the function $L_{\mathcal{L}} : TM \rightarrow T^*M$ defined fiberwise by

$$L_{\mathcal{L}}(x, v) = \sum \frac{\partial \mathcal{L}}{\partial v_i} dv_i \in T_{(x,v)}^* T_x \cong T_x^*$$

If we choose coordinates (x_i, w_i) on T^*M dual to (x_i, v_i) then a curve $\gamma' : [0, 1] \rightarrow TM$ determines $\zeta : [0, 1] \rightarrow T^*M$ by

$$\zeta(t) := L_{\mathcal{L}}(\gamma(t), \dot{\gamma}(t))$$

and we can abbreviate $(x(\zeta(t)), w(\zeta(t)))$ by $(x(t), w(t))$. Note that $x(\zeta(t)) = x(\gamma(t))$ so the two definitions of $x(t)$ agree. Note further that $w_i(t) = \frac{\partial \mathcal{L}}{\partial v_i}(t)$.

Hamilton's equations for the function H on T^*M are a pair of equations for a curve $(x(t), w(t))$:

$$\frac{dx_i}{dt} = \frac{\partial H}{\partial w_i}(x(t), w(t)) \quad \frac{dw_i}{dt} = -\frac{\partial H}{\partial x_i}(x(t), w(t))$$

Proposition A.3. *A curve $\gamma : [0, 1] \rightarrow M$ is a critical point for action associated to a Lagrangian \mathcal{L} (equivalently, it satisfies the Euler-Lagrange equations in every coordinate chart) if and only if $\zeta := L_{\mathcal{L}}(\gamma')$ satisfies Hamilton's equations for the Hamiltonian H obtained from \mathcal{L} by fiberwise Legendre transform.*

Proof. The first of Hamilton's equations $\frac{dx_i}{dt} = \frac{\partial H}{\partial w_i}(x(t), w(t))$ is just the statement that the Legendre transform is an involution.

Now, by the definition of the Legendre transform, for $L_{\mathcal{L}}(x(t), v(t)) = (x(t), w(t))$ we have

$$H(x(t), w(t)) + \mathcal{L}(x(t), v(t)) = w(t) \cdot v(t)$$

Differentiating with respect to x , we get

$$\frac{\partial H}{\partial x} + \frac{\partial H}{\partial w} \cdot \frac{\partial w}{\partial x} = \frac{\partial w}{\partial x} \cdot v - \frac{\partial \mathcal{L}}{\partial x}$$

But by Hamilton's first equation, $\frac{\partial H}{\partial w} = \dot{x} = v$, so $\partial \mathcal{L} / \partial x = -\partial H / \partial x$ on respective curves related by $L_{\mathcal{L}}$.

Putting this together we get

$$\frac{dw}{dt} = \frac{d}{dt} \frac{\partial \mathcal{L}}{\partial v}$$

whereas

$$-\frac{\partial H}{\partial x} = \frac{\partial \mathcal{L}}{\partial x}$$

as we established above, so the second of Hamilton's equations is seen to be equivalent to the Euler-Lagrange equation. \square

This correspondence between solutions of the Euler-Lagrange and Hamilton equations is remarkable; in the first case our solutions *minimize* the value of a functional — the action. In the second case our solutions *conserve* the value of a function — the Hamiltonian.

A.4. Symplectic geometry. The cotangent bundle of a smooth manifold has a canonical symplectic structure.

A.4.1. *Canonical symplectic structure.*

Definition A.4. Let M be a smooth manifold. The *tautological 1-form* α on T^*M is defined as follows. Let $x \in M$ and $w \in T_x^*M$ be given. The projection $T^*M \rightarrow M$ induces $\pi : TT^*M \rightarrow TM$, and if $v \in T_{(x,w)}T^*M$ then we can define $\alpha(v) = w(\pi(v))$.

The *canonical symplectic form* ω on T^*M is the 2-form $\omega := -d\alpha$.

A diffeomorphism $f : M \rightarrow N$ induces $f^* : T^*N \rightarrow T^*M$ and $df^* : TT^*N \rightarrow TT^*M$. This is compatible with projection, in the sense that

$$df\pi_M df^*(v) = \pi_N(v)$$

In particular, f^* is a symplectomorphism from T^*N to T^*M . In the special case that $M = N$ we see that any diffeomorphism from M to itself induces a symplectomorphism from T^*M to itself. In particular, a smooth manifold has a *canonical* volume form on T^*M .

If x_i are local coordinates on M , and w_i are local coordinates on T_x^*M for each x so that w_i is the coefficient of dx_i in the expression of a 1-form (at a given point) in terms of the basis dx_i , then together x_i, w_i are coordinates on T^*M , and the tautological 1-form is $\sum w_i dx_i$. So we can compute $\omega = \sum dx_i \wedge dw_i$ which is nondegenerate, proving that ω is symplectic.

A.4.2. *Hamilton's equations.* Any smooth function H on a symplectic manifold determines a flow $\phi_t'(x) := X_H(x)$, where the vector field X_H is the unique vector field so that $\omega(X_H, Y) = dH(Y)$ for every other vector field Y . In local coordinates x_i, w_i we have

$$dH = \sum \frac{\partial H}{\partial x_i} dx_i + \sum \frac{\partial H}{\partial w_i} dw_i$$

so that

$$\partial H / \partial x_i = dH(\partial / \partial x_i) = \omega(X_H, \partial / \partial x_i) = -dw_i(X_H) = -\dot{w}_i$$

and

$$\partial H / \partial w_i = dH(\partial / \partial w_i) = \omega(X_H, \partial / \partial w_i) = dx_i(X_H) = \dot{x}_i$$

which are just Hamilton's equations from before.

A.4.3. *Poisson bracket.* The *Poisson bracket* $\{\cdot, \cdot\}$, defined on $C^\infty(M)$ by the formula

$$\{A, B\} := \omega(X_A, X_B)$$

makes $C^\infty(M)$ into a Lie algebra, with center the constant functions, and quotient the Lie algebra of Hamiltonian vector fields:

$$0 \rightarrow \mathbb{R} \rightarrow C^\infty(M) \rightarrow \text{Ham}(M) \rightarrow 0$$

Note that for the coordinate functions we have

$$\{x_i, x_j\} = \{w_i, w_j\} = 0, \quad \{x_i, w_j\} = \delta_{ij}$$

If A is any smooth function on T^*M , the time derivative of A under the Hamiltonian flow generated by H can be computed as

$$\frac{dA}{dt} = \sum \frac{\partial A}{\partial x_i} \dot{x}_i + \frac{\partial A}{\partial w_i} \dot{w}_i = \sum \frac{\partial A}{\partial x_i} \frac{\partial H}{\partial w_i} - \frac{\partial A}{\partial w_i} \frac{\partial H}{\partial x_i} = \{A, H\}$$

Taking A equal to x_i or w_i recovers Hamilton's equations.

A.5. Classical phase space. In classical mechanics we very often consider a Lagrangian of the form

$$\mathcal{L} = \text{kinetic energy} - \text{potential energy}$$

where the kinetic energy is a positive definite quadratic form on $T_x M$ for all x (in other words, a Riemannian metric) and the potential energy is a function only of position; i.e. $\mathcal{L}(x, v) = \langle v, v \rangle_x - \Phi(x)$. A positive definite quadratic form on a vector space V gives rise to an identification of V with V^* . The Legendre transform of the positive definite quadratic form is then equal to *the same* quadratic form (after making this identification); this is a consequence of Euler's theorem on homogeneous functions, which for quadratic forms $f(x)$ takes the form $\nabla f \cdot x = 2f$. Adding a constant to a function subtracts the same constant from its transform. So the Hamiltonian H is of the form $H(x, w) = \langle w, w \rangle_x + \Phi(x)$, or in words,

$$H = \text{kinetic energy} + \text{potential energy}$$

i.e. H is the "total energy" of the system, and Hamilton's equations confirm the law of conservation of energy.

Example A.5. Consider a configuration of a very large number N of particles, confined to some region Ω of (3-dimensional) space. The position of the N particles in time traces out a path γ in $M := \Omega^N \subset \mathbb{R}^{3N}$, and we can lift γ tautologically to a path γ' in TM .

If our particles have coordinates $(x_i(t), y_i(t), z_i(t))$ we have

$$\mathcal{L} = \frac{1}{2} \sum_i m_i (\dot{x}_i^2 + \dot{y}_i^2 + \dot{z}_i^2) - \Phi(x_1, \dots, z_N)$$

where m_i is the mass of the i th particle. The conjugate momenta are

$$p_{x_i} := \frac{\partial \mathcal{L}}{\partial \dot{x}_i} = m_i \dot{x}_i$$

and so on, so that we can write H in the form

$$H = \frac{1}{2} \sum_i \frac{p_{x_i}^2 + p_{y_i}^2 + p_{z_i}^2}{2m_i} + \Phi(x_1, \dots, z_N)$$

The symplectic form on phase space T^*M is the form $\sum_i dx_i \wedge dp_{x_i} + dy_i \wedge dp_{y_i} + dz_i \wedge dp_{z_i}$.

A.6. Quantum mechanics. In (nonrelativistic) quantum mechanics, the state of the system is described by a vector ψ of norm 1 in a Hilbert space \mathcal{H} , and each observable is represented by a Hermitian (i.e. self-adjoint) operator A acting on \mathcal{H} . The *expected value* of the observable \hat{A} , when the system is in state ψ , is the quantity

$$\langle \hat{A} \rangle = \langle \psi, \hat{A} \psi \rangle$$

For example, the Hamiltonian itself (whose value represents the energy) is an operator \hat{H} (for simplicity we consider the case that \hat{H} is time-independent). Since it is Hermitian, the eigenvalues are real, and (at least when \mathcal{H} is finite dimensional) we can decompose \mathcal{H} into a direct sum of orthonormal eigenvectors for \hat{H} ; the eigenvalue ϵ_i associated to the eigenvector ψ_i is the *characteristic energy* of the state ψ_i .

The evolution of the system in time is governed by the Schrödinger equation:

$$\hat{H}\psi = i\hbar \frac{\partial\psi}{\partial t}$$

where \hbar denotes Planck's constant h divided by 2π . Thus if we express ψ at time 0 as a linear combination of eigenfunctions $\psi = \sum a_j \psi_j$, we get the formula $\psi(t) = \sum a_j e^{-i\epsilon_j t/\hbar} \psi_j$. Since the ψ_j are orthonormal,

$$\langle \psi(t), \hat{H}\psi(t) \rangle = \sum \epsilon_j |a_j|^2$$

is independent of time, and we recover the fact that energy is conserved.

A.6.1. *Phase space.* Suppose we are interested as before in N particles confined to some region Ω in 3-dimensional space. Recall that the phase space is $(T^*\Omega)^N$. The Hilbert space \mathcal{H} is the space of L^2 complex-valued functions on Ω^N (i.e. functions of *half* of the coordinates on the full phase space). Poisson brackets of functions A on ordinary phase space are related to *commutators* of the corresponding operators \hat{A} on \mathcal{H} ; in most procedures for quantization this correspondence should satisfy

$$[\hat{A}, \hat{B}] = i\hbar \widehat{\{A, B\}} + O(\hbar^2)$$

Note that the associativity of composition of operators corresponds to the Jacobi identity for the Poisson bracket.

There are *ad hoc* rules for constructing operators corresponding to functions on phase space expressed as simple algebraic formulae in the positions x_j and momenta w_j ; these rules are sometimes ambiguous, but are easy to remember in many cases of interest. For any expression $A(x, w)$, the corresponding operator is obtained by replacing x_j by the operator “multiplication by x_j ”, and replacing w_j by the operator $-i\hbar\partial/\partial x_j$. Thus as operators,

$$[\hat{x}_j, \hat{w}_k]f = -i\hbar \left(x_j \frac{\partial f}{\partial x_k} - \frac{\partial(x_j f)}{\partial x_k} \right) = i\hbar \delta_{jk} f$$

On \mathbb{R}^3 with potential energy Φ , the Hamiltonian is

$$\hat{H} := -\hbar^2 \sum_j \frac{1}{2m_j} \left(\frac{\partial^2}{\partial x_j^2} + \frac{\partial^2}{\partial y_j^2} + \frac{\partial^2}{\partial z_j^2} \right) + \Phi(x)$$

We can effectively confine our particles to a compact region Ω by setting Φ equal to 0 on Ω and ∞ outside. If our particles are all identical with a common mass m , then \hat{H} reduces to $\hat{H} = \hbar^2/2m\Delta$ acting on $L^2(\Omega^N)$, where Δ is the “geometer's Laplacian” (whose spectrum is real and non-negative). The eigenstates of energy thus correspond to eigenvalues of the Laplacian with Dirichlet boundary conditions (i.e. they vanish on $\partial\Omega^N$).

The number of eigenvalues less than x is controlled by the famous *Weyl law*, which says that for a bounded domain Ω in \mathbb{R}^d , if $N(t)$ denotes the number of Dirichlet eigenvalues less than or equal to t (counted with multiplicity), there is a formula

$$\lim_{t \rightarrow \infty} \frac{N(t)}{t^{d/2}} = (2\pi)^{-d} \omega_d \text{vol}(\Omega)$$

where ω_d is the volume of the unit ball in \mathbb{R}^d . In particular, for Ω a bounded domain in \mathbb{R}^3 , one has $N(t) \sim t^{3/2} \text{vol}(\Omega)$.

Example A.6 (Particle confined to a cubic box). For a single particle confined to a cube Ω with side length s , the eigenfunctions are of the form

$$\psi(n_x, n_y, n_z) := C \cdot \sin\left(\frac{\pi n_x}{s} x\right) \sin\left(\frac{\pi n_y}{s} y\right) \sin\left(\frac{\pi n_z}{s} z\right)$$

for (non-negative) integers n_x, n_y, n_z and a constant C (well-defined up to phase) such that ψ has norm 1, corresponding to a state with energy $(n_x^2 + n_y^2 + n_z^2)h^2/8ms^2$. The number of states with energy bounded by t is equal to the number of integer lattice points in the intersection of the positive orthant of \mathbb{R}^3 with the ball of radius $R := \sqrt{8ms^2 t/h^2}$ which is $C \cdot s^3 t^{3/2} + O(t)$.

Example A.7 (Harmonic oscillator). Consider a particle of mass m confined to 1 dimensions and subject to a quadratic potential $\Phi(x) = \frac{1}{2}fx$; for example, classically we could consider an ideal spring with spring constant f . The Hamiltonian is

$$\hat{H} := -\frac{\hbar^2}{2m} \frac{\partial^2}{\partial x^2} + \frac{1}{2}fx^2$$

with eigenfunctions of the form $p_j(x/\sigma)e^{-\omega x^2/2\sigma^2}$ for suitable σ (depending on m and f) and polynomials p_j called the *Hermite polynomials*. The eigenvalues are of the form

$$\epsilon_j := \left(j + \frac{1}{2}\right) h\nu, \quad j = 0, 1, 2, \dots$$

where $\nu = (1/2\pi)(f/m)^{1/2}$. The quantity ϵ_0 is called the “zero point energy”.

Remark A.8. The $O(\hbar^2)$ correction term in the correspondence between Poisson bracket of functions and commutator of operators is characteristic of a procedure known as *deformation quantization*. The algebra of smooth functions on a symplectic manifold is deformed — ordinary multiplication of functions is replaced by an associative, but non-commutative operation called the *Moyal product* \star , taking values in expressions of the form

$$f \star g = fg + \sum_{n=1}^{\infty} \hbar^n C_n(f, g)$$

for suitable bidifferential operators C_n of order n , and so that

$$f \star g - g \star f = i\hbar\{f, g\} + O(\hbar^2)$$

As an abstract algebra, this algebra is isomorphic to the algebra of observables on the Hilbert space of the system; thus $\widehat{f \star g} = \widehat{f} \widehat{g}$.

APPENDIX B. METHOD OF STEEPEST DESCENT AND THE DARWIN-FOWLER THEOREM

In this section we sketch the proof of the Darwin-Fowler theorem, which says roughly that *almost all* distributions of N systems with total energy $E := UN$ are very similar (in a suitable sense) to the Boltzmann distribution. Our treatment is closely modeled on Schrödinger [10] Ch. 5–6.

For concreteness, let us first suppose we have a *finite* set Σ of states. An *assignment* (ignoring energy for the moment) is just an element of Σ^N , i.e. a function from an N element set to Σ . We let $A = \Sigma^N$ denote the set of all possible assignments. This is a finite set, and it admits a uniform probability measure in which every element has the same probability.

There is a *counting* map from A to \mathbb{R}_+^Σ , sending an assignment to the vector whose coefficients are the numbers n_j/N for the various states $j \in \Sigma$. Since each n_i is non-negative, and $\sum n_j = N$, the image is contained in the simplex $\Delta \subset \mathbb{R}_+^\Sigma$ of non-negative vectors whose coefficients sum to 1 (i.e. the simplex of probability measures on Σ).

Inside A we have the subset A_E of assignments with total energy $E = NU$, again with the uniform probability measure. The counting measure takes A_E to a polyhedron $P_E \subset \Delta$ which is the intersection of Δ with the level set of the energy function U on Δ , which is the linear function $U(t_1, t_2, \dots) = \sum t_j \epsilon_j$. Thus, each N gives us a *probability distribution* ν_N on P_E (supported on a finite lattice, which gets finer and finer as $N \rightarrow \infty$). The Darwin-Fowler theorem says (informally) that the measures ν_N converge weakly to the Dirac measure concentrated on the Boltzmann distribution; i.e. for all continuous functions f on P_E the integrals $\int_{P_E} f d\nu_N \rightarrow f(\text{Boltzmann distribution})$. Note that each ν_N is a *probability measure on the space P_E of probability measures on Σ* , and the Boltzmann distribution is a specific *point* in P_E .

One issue that we have completely ignored is the question of whether there are any assignments with total energy E *at all*. Therefore it is convenient (if physically dubious) to assume that the energy levels ϵ_j are *integers with no common factors*. We prove the theorem under these assumptions, and explain at the end what modifications must be made in the more general case of arbitrary energy levels or infinitely many states.

B.1. Generating functions. We introduce formal variables x_j associated to the states $j \in \Sigma$. The multinomial theorem gives

$$\sum_{\sum n_j = N} \frac{N!}{\prod_j n_j!} x_1^{n_1} x_2^{n_2} \cdots = (x_1 + x_2 + \cdots)^N$$

(ignoring questions of convergence if the number of states is infinite). However this does not yet take account the restriction that $\sum n_j \epsilon_j = E$. This can be achieved by introducing an additional formal variable z , defining the generating function

$$F := \sum_{\sum n_j = N} \frac{N!}{\prod_j n_j!} (x_1 z^{\epsilon_1})^{n_1} (x_2 z^{\epsilon_2})^{n_2} \cdots$$

and observing that all the data of the distribution ν_N can be obtained from the coefficient of z^E in F . If we define

$$f(z) := x_1 z^{\epsilon_1} + x_2 z^{\epsilon_2} + \dots$$

then $F = f(z)^N$.

For any numerical value of the x_i , the function $f(z)$ is a power series in z , and we are interested in the residue at 0 of the function $z^{-E-1} f(z)^N$; i.e.

$$\hat{\Omega}(x_1, x_2, \dots) := \frac{1}{2\pi i} \oint_{\gamma} z^{-E-1} f(z)^N dz$$

where the integral is taken over a sufficiently small loop γ about 0 containing no singularities of $f(z)$.

B.2. Method of steepest descent. Consider the restriction of the integrand $z^{-E-1} f(z)^N$ to the positive real axis. The first term starts out at $+\infty$ and then decreases very rapidly and monotonically. The second term $f(z)^N$ is equal to 1 at $z = 0$, and increases monotonically with z to infinity. Furthermore, comparing logarithmic derivatives,

$$\frac{d \log(z^{-E-1})}{dz} = -\frac{E+1}{z}, \quad \frac{d \log(f(z)^N)}{dz} = N \frac{f'(z)}{f(z)} = N \frac{\sum x_j \epsilon_j z^{\epsilon_j - 1}}{\sum x_j z^{\epsilon_j}}$$

one sees that the first is negative and monotonically increasing from $-\infty$ to 0, while the second is positive and monotonically increasing; this can be seen from the second derivative

$$\frac{d}{dz} \left(N \frac{f'(z)}{f(z)} \right) = N \frac{\sum x_j \epsilon_j^2 z^{\epsilon_j - 2} \sum x_j z^{\epsilon_j} - (\sum x_j \epsilon_j z^{\epsilon_j - 1})^2}{(\sum x_j z^{\epsilon_j})^2}$$

which is positive by Cauchy-Schwarz.

In particular, it follows that the integrand has a unique minimum on the positive real axis, at the point z_0 obtained by solving

$$-\frac{E+1}{z_0} + N \frac{f'(z_0)}{f(z_0)} = 0$$

Now, under our assumption that $E, N \rightarrow \infty$ with E/N fixed, the point z_0 has a well defined limit. Moreover, if we define $g(z)$ implicitly by $e^{g(z)} := z^{-E-1} f(z)^N$ then

$$g''(z_0) = \frac{E+1}{z_0^2} + N \left(\frac{f''(z_0)}{f(z_0)} - \frac{f'(z_0)^2}{f(z_0)^2} \right)$$

will grow without bound (linearly in E or N), so that the minimum is very sharply peaked at z_0 .

Near z_0 we can write the integrand as a power series

$$z^{-E-1} f(z)^N = z_0^{-E-1} f(z_0)^N e^{g''(z_0)(z-z_0)^2/2+\dots}$$

Taking $z = z_0 + iy$ for small real y , the last term in this product is well-approximated by $e^{-\frac{1}{2}y^2 g''(z_0)}$, a very sharply peaked Gaussian.

Suppose we take the contour of integration to be the circle about 0 with radius z_0 . Near z_0 the integral is well approximated by

$$\frac{1}{2\pi i} z_0^{-E-1} f(z_0)^N \int_{-\infty}^{\infty} e^{-\frac{1}{2}y^2 g''(z_0)} i dy = z_0^{-E-1} f(z_0)^N \frac{1}{\sqrt{2\pi g''(z_0)}}$$

Away from z_0 , but with $|z|$ fixed, the term $f(z) = x_1 z^{\epsilon_1} + x_2 z^{\epsilon_2} + \dots$ is the sum of a number of terms of fixed magnitude, but with different phases, and therefore we have $|f(z)| < |f(z_0)| = f(z_0)$ on the circle $|z| = z_0$ (this is where we use the assumption that the ϵ_j have no common factor). Taking N th powers makes $|f(z)^N| \ll |f(z_0)^N|$, so the integral is dominated by the contribution near z_0 , which is given above. Thus, for $N, E \gg 1$ we get the increasingly good approximation

$$\hat{\Omega}(x_1, x_2, \dots) = z_0^{-E-1} f(z_0)^N \frac{1}{\sqrt{2\pi g''(z_0)}}$$

B.3. Expected value and variance of the n_j . We now compute the expectation and variance of the n_j from the expression $\hat{\Omega}$. If we define for each distribution $\{n_j\}$ with $\sum n_j = N$ and $\sum n_j \epsilon_j = E$ the function

$$P(\{n_j\}) := \frac{N!}{\prod_j n_j!} x_1^{n_1} x_2^{n_2} \dots$$

and let $\hat{\Omega}$ denote the sum over all distributions, then the expected value of n_j is

$$\langle n_j \rangle = \frac{\sum n_j P(\{n_j\})}{\hat{\Omega}} = x_j \frac{\partial \log \hat{\Omega}}{\partial x_j}$$

and

$$\langle n_j^2 \rangle = \frac{\sum n_j^2 P(\{n_j\})}{\hat{\Omega}} = x_j \frac{\partial}{\partial x_j} \left(\frac{x_j}{\hat{\Omega}} \frac{\partial \hat{\Omega}}{\partial x_j} \right) + \left(\frac{x_j}{\hat{\Omega}} \frac{\partial \hat{\Omega}}{\partial x_j} \right)^2$$

so that the variance is

$$\langle n_j^2 \rangle - \langle n_j \rangle^2 = x_j \frac{\partial \langle n_j \rangle}{\partial x_j}$$

From our steepest descent approximation we compute

$$\langle n_j \rangle = x_j \frac{\partial \log \hat{\Omega}}{\partial x_j} = x_j \frac{\partial}{\partial x_j} \left(g(z_0) - \frac{1}{2} \log 2\pi g''(z_0) \right)$$

Note that z_0 is implicitly a function of the x_j . Thus, differentiating, this gives

$$\begin{aligned} \langle n_j \rangle &= x_j g'(z_0) \frac{\partial z_0}{\partial x_j} + x_j \frac{\partial g(z)}{\partial x_j}(z_0) - \frac{x_j}{2} \frac{\partial}{\partial x_j} \log 2\pi g''(z_0) \\ &= x_j g'(z_0) \frac{\partial z_0}{\partial x_j} + \frac{N x_j z_0^{\epsilon_j}}{f(z_0)} - \frac{x_j}{2} \frac{\partial}{\partial x_j} \log 2\pi g''(z_0) \end{aligned}$$

The first term is zero, by the definition of z_0 . The last has order $\log N$, so can be neglected. Thus, setting all $x_j = 1$ and writing $\beta = -\log z_0$ we get

$$\frac{\langle n_j \rangle}{N} = \frac{e^{-\beta \epsilon_j}}{\sum_j e^{-\beta \epsilon_j}}$$

agreeing with the Boltzmann distribution.

Using the fact that $g'(z_0)$ vanishes identically in x_j , and ignoring the order $\log N$ term, we get an estimate for the variance

$$\langle n_j^2 \rangle - \langle n_j \rangle^2 = x_j \frac{\partial \langle n_j \rangle}{\partial x_j} = x_j \frac{N z_0^{\epsilon_j}}{f(z_0)} + x_j^2 N \left(\frac{\partial}{\partial x_j} \left(-\frac{z_0^{\epsilon_j} f'(z_0)}{f(z_0)^2} + \frac{\epsilon_j z_0^{\epsilon_j - 1}}{f(z_0)} \right) - \frac{z_0^{2\epsilon_j}}{f(z_0)^2} \right)$$

Setting the $x_j = 1$, this can be seen to be of order $\langle n_j \rangle$, so that the standard deviation is of order $O(\sqrt{N})$, and for large N almost every distribution is practically indistinguishable from the Boltzmann distribution. This completes the proof of the Darwin-Fowler theorem, under the assumption of a finite set of states, and integer energy levels.

APPENDIX C. LARGE DEVIATIONS AND CHERNOFF INEQUALITY FOR MARKOV CHAINS

C.1. Chernoff inequality for independent random variables. The simplest form of the Chernoff bound, for independent random variables, is the following:

Proposition C.1. *Let X_1, X_2, \dots, X_n be independent random variables taking values in $[0, 1]$, and let $X = X_1 + X_2 + \dots + X_n$ denote their sum. Let $\mu = \langle X \rangle = \sum \langle X_i \rangle$. Then for any $\epsilon \geq 0$,*

$$\Pr(X \geq (1 + \epsilon)\mu) \leq e^{-\frac{\epsilon^2}{2+\epsilon}\mu} \text{ and } \Pr(X \leq (1 - \epsilon)\mu) \leq e^{-\frac{\epsilon^2}{2}\mu}$$

Proof. The proof is an application of Markov's inequality. We estimate

$$\Pr(X \geq (1 + \epsilon)\mu) = \Pr(e^{tX} \geq e^{t(1+\epsilon)\mu}) \leq \frac{\langle e^{tX} \rangle}{e^{t(1+\epsilon)\mu}} = e^{-t(1+\epsilon)\mu} \langle \prod_i e^{tX_i} \rangle = e^{-t(1+\epsilon)\mu} \prod_i \langle e^{tX_i} \rangle$$

□

ACKNOWLEDGMENTS

Danny Calegari was supported by NSF grant DMS 1005246.

REFERENCES

- [1] N. Davidson, *Statistical Mechanics*, McGraw-Hill, New York, 1962
- [2] E. Fermi, *Thermodynamics*, Prentice-Hall, New York, 1937
- [3] R. Feynman, *The Feynman lectures on physics*,
- [4] R. Feynman, *Statistical Mechanics; a set of lectures*, Advanced book classics, Addison-Wesley, Reading, 1998
- [5] J. Gibbs, *On the equilibrium of heterogeneous substances*, Trans. Connecticut Acad. Arts Sci. Vol. III, pp. 108–248 and 343–524
- [6] P. Lezaud, *Chernoff and Berry–Essén inequalities for Markov processes*, ESAIM Probab. Stat. **5** (2001), 183–201
- [7] J. von Neumann, *Mathematical foundations of Quantum Mechanics*, Princeton landmarks in mathematics, Princeton University Press, Princeton, 1996
- [8] D. Ornstein, *Bernoulli shifts with the same entropy are isomorphic*, Adv. Math. **4** (1970), 337–352
- [9] A. Pippard, *The elements of Classical Thermodynamics*, Cambridge University Press, Cambridge, 1957
- [10] E. Schrödinger, *Statistical thermodynamics*, Cambridge University Press, Cambridge, 1952
- [11] C. Shannon, *A mathematical theory of communication*, Bell System technical journal (July–October 1948) **27** (3): 379–423

UNIVERSITY OF CHICAGO, CHICAGO, ILL 60637 USA
E-mail address: dannyc@math.uchicago.edu