

1 Introduction

Statistical mechanics is the branch of physics which aims at bridging the gap between the *microscopic* and *macroscopic* descriptions of large systems of particles in interaction, by combining the information provided by the microscopic description with a probabilistic approach. Its goal is to understand the relations existing between the salient macroscopic features observed in these systems and the properties of their microscopic constituents. Equilibrium statistical mechanics is the part of this theory that deals with macroscopic systems *at equilibrium* and is, by far, the best understood.

This book is an introduction to some classical mathematical aspects of equilibrium statistical mechanics, based essentially on some important examples. It does not constitute an exhaustive introduction: many important aspects, discussed in a myriad of books (see those listed in Section 1.6.2 below), will not be discussed. Inputs from physics will be restricted to the terminology used (especially in this introduction), to the nature of the problems addressed, and to the central probability distribution used throughout, namely the **Gibbs distribution**. This distribution provides the probability of observing a particular microscopic state ω of the system under investigation, when the latter is at equilibrium *at a fixed temperature* T . It takes the form

$$\mu_\beta(\omega) = \frac{e^{-\beta \mathcal{H}(\omega)}}{\mathbf{Z}_\beta},$$

where $\beta = 1/T$, $\mathcal{H}(\omega)$ is the *energy* of the microscopic state ω and \mathbf{Z}_β is a normalization factor called the **partition function**.

Saying that the Gibbs distribution is well suited to understand the phenomenology of large systems of particles is an understatement. This book provides, to some extent, a proof of this fact by diving into an in-depth study of this distribution when applied to some of the most important models studied by mathematical physicists since the beginning of the 20th century. The many facets and the rich variety of behaviors that will be described in the following chapters should, by themselves, constitute a firm justification for the use of the Gibbs distribution for the description of large systems at equilibrium.

An *impatient reader* with some basic notions of thermodynamics and statistical mechanics, or who is willing to consider the Gibbs distribution as a postulate and is not interested in additional motivations and background, can jump directly to the

following chapters and learn about the models presented throughout the book. A quick glance at Section 1.6 might be useful since it contains a reading guide.

The rest of this introduction is written for a reader interested in obtaining more information on the origin of the Gibbs distribution and the associated terminology, as well as an informal discussion of thermodynamics and its relations to equilibrium statistical mechanics.

One of the main themes to which this book is devoted, phase transitions, is illustrated on gases and magnets. However, we emphasize that, *in this book, the main focus is on the mathematical structure of equilibrium statistical mechanics* and the many possible interpretations of the models we study will most of the time not play a very important role; they can, nevertheless, sometimes provide intuition.

Although this book is undeniably written for a mathematically inclined reader, in this introduction we will avoid delving into too much technicalities; as a consequence, it will not be as rigorous as the rest of the book. Its purpose is to provide intuition and motivation behind several key concepts, relying mainly on physical arguments. The content of this introduction is not necessary for the understanding of the rest of the book, but we hope that it will provide the reader with some useful background information.

We will start with a brief discussion of the first physical theory describing macroscopic systems at equilibrium, equilibrium thermodynamics, and present also examples of one of the most interesting features of thermodynamic systems: phase transitions. After that, starting from Section 1.2, we will turn our attention to equilibrium statistical mechanics.

1.1 Equilibrium Thermodynamics

Equilibrium thermodynamics is a phenomenological theory, developed mainly during the nineteenth century. Its main early contributors include Carnot, Clausius, Kelvin, Joule and Gibbs. It is based on a few empirical principles and does not make any assumption regarding the microscopic structure of the systems it considers (in fact, the atomic hypothesis was still hotly debated when the theory was developed).

This section will briefly present some of its basic concepts, mostly on some simple examples; it is obviously not meant as a complete description of thermodynamics, and the interested reader is advised to consult the detailed and readable account in the books of Callen [58] and Thess [329], in Wightman's introduction in [176] or in Lieb and Yngvason's paper [224].

1.1.1 On the description of macroscopic systems

Gases, liquids and solids are the most familiar examples of large physical systems of particles encountered in nature. In the first part of this introduction, for the sake of simplicity and concreteness, we will mainly consider the system made of a gas contained in a vessel.

Let us thus consider a specific, homogeneous gas in a vessel (this could be, say, one liter of helium at standard temperature and pressure). We will use Σ to denote such a specific system. The **microscopic state**, or **microstate**, of the gas is the complete microscopic description of the state of the gas. For example, from the point of view of Newtonian mechanics, the microstate of a gas of monoatomic

molecules is specified by the position and the momentum of each molecule (called hereafter *particle*). Since a gas contains a huge number of particles (of the order of 10^{22} for our liter of helium), the overwhelming quantity of information contained in its microstate makes a complete analysis not only challenging, but in general impossible. Fortunately, we are not interested in the full information contained in the microstate: the precise behavior of each single particle is irrelevant at our scale of observation.

It is indeed an empirical fact that, in order to describe the state of the gas at the **macroscopic** scale, only a much smaller set of variables, of a different nature, is needed. This is particularly true when the system is in a particular kind of state called *equilibrium*. Assume, for instance, that the gas is **isolated**, that is, it does not exchange matter or energy with the outside world, and that it has been left undisturbed for a long period of time. Experience shows that such a system reaches a state of **thermodynamic equilibrium**, in which macroscopic properties of the gas do not change anymore and there are no macroscopic flows of matter or energy (even though molecular activity never ceases).

In fact, by definition, isolated systems possess a number of *conserved* quantities, that is, quantities that do not change through time evolution. The latter are particularly convenient to describe the macroscopic state of the system. For our example of a gas, these conserved quantities are the **volume** V of the vessel¹, the **number of particles** N and the **internal energy** (or simply: **energy**) U .

We therefore assume, from now on, that the **macroscopic state** (or **macrostate**) of the gas Σ is determined by a triple

$$\mathbf{X} = (U, V, N).$$

The variables (U, V, N) can be thought of as the quantities one can control in order to alter the state of the gas. For example, U can be changed by cooling or heating the system, V by squeezing or expanding the container and N by injecting or extracting particles with a pump or a similar device. These variables determine the state of the gas in the sense that setting these variables to some specific values always yield, once equilibrium is reached, systems that are macroscopically indistinguishable. They can thus be used to put the gas in a chosen macroscopic state reproducibly. Of course, what is reproduced is the macrostate, *not the microstate*: there are usually infinitely many different microstates corresponding to the same macrostate.

Now, suppose that we split our system, Σ , into two subsystems Σ_1 and Σ_2 , by adding a wall partitioning the vessel. Each of the subsystems is of the same type as the original system and only differs from it by the values of the corresponding variables (U^m, V^m, N^m) , $m = 1, 2$. Observe that the total energy U , total volume V and total number of particles N in the system satisfy:^[1]

$$U = U^1 + U^2, \quad V = V^1 + V^2, \quad N = N^1 + N^2. \quad (1.1)$$

Variables having this property are said to be **extensive**².

¹We assume that the vessel is large enough and that its shape is simple enough (for example: a cube), so as not to influence the macroscopic behavior of the gas and boundary effects may be neglected.

²The identity is not completely true for the energy: part of the latter comes, in general, from the interaction between the two subsystems. However, this interaction energy is generally negligible compared to the overall energy (exceptions only occur in the presence of very long-range interactions, such as gravitational forces). This will be quantified once we study similar problems in statistical mechanics.

Note that the description of a system Σ composed of two subsystems Σ_1 and Σ_2 now requires 6 variables: $(U^1, V^1, N^1, U^2, V^2, N^2)$. Below, a system will always come with the set of variables used to characterize it. Note that this set of variables is not unique: one can always split a system into two pieces in our imagination, without doing anything to the system itself, only to our description of it.

One central property of equilibrium is that, when a system is at equilibrium, each of its (macroscopic) subsystems is at equilibrium too, and all subsystems are in equilibrium with each other. Namely, if we imagine that our system Σ is partitioned by an imaginary wall into two subsystems Σ_1, Σ_2 of volume V^1 and V^2 , then the thermodynamic properties of each of these subsystems do not change through time: their energy and the number of particles they contain remain constant³.

Assume now that Σ_1, Σ_2 are originally separated (far apart, with no exchanges whatsoever), each isolated and at equilibrium. The original state of the union of these systems is represented by $(\mathbf{X}^1, \mathbf{X}^2)$, where $\mathbf{X}^m = (U^m, V^m, N^m)$ is the macrostate of Σ_m , $m = 1, 2$. Suppose then that these systems are put in contact, allowing them to exchange energy and/or particles, while keeping them, as a whole, isolated from the rest of the universe (in particular, the total energy U , total volume V and total number of particles N are fixed). Once they are in contact, the whole system goes through a phase in which energy and particles are redistributed between the subsystems and a fundamental problem is to determine which new equilibrium macrostate $(\bar{\mathbf{X}}^1, \bar{\mathbf{X}}^2)$ is realized and how it relates to the initial pair $(\mathbf{X}^1, \mathbf{X}^2)$.

The core postulate of equilibrium thermodynamics is to assume the existence of a function, associated to any system Σ , which describes how the new equilibrium state is selected among the a priori infinite number of possibilities. This function is called *entropy*.

1.1.2 The thermodynamic entropy

Let us assume that Σ is the union of two subsystems Σ_1, Σ_2 and that some *constraints* are imposed on these subsystems. We model these constraints by the set \mathbb{X}_c of allowed pairs $(\mathbf{X}^1, \mathbf{X}^2)$. We expect that the system selects some particular pair in \mathbb{X}_c to realize equilibrium, in some optimal way. The main postulate of Thermostatics is that this is done by choosing the pair that maximizes the entropy:

Postulate (Thermostatics). *To each system Σ , described by a set of variables \mathbf{X} , is associated a differentiable function S^Σ of \mathbf{X} , called the **(thermodynamic) entropy**; it is specific to each system. The entropy of a system Σ composed of two subsystems Σ_1 and Σ_2 is **additive**:*

$$S^\Sigma(\mathbf{X}^1, \mathbf{X}^2) = S^{\Sigma_1}(\mathbf{X}^1) + S^{\Sigma_2}(\mathbf{X}^2). \quad (1.2)$$

Once the systems are put in contact, the pair $(\bar{\mathbf{X}}^1, \bar{\mathbf{X}}^2)$ realizing equilibrium is the one that maximizes $S^\Sigma(\mathbf{X}^1, \mathbf{X}^2)$, among all pairs $(\mathbf{X}^1, \mathbf{X}^2) \in \mathbb{X}_c$.

The principle by which a system under constraint realizes equilibrium by maximizing its entropy will be called the **extremum principle**.


³Again, this is not true from a microscopic perspective: the number of particles in each subsystem does fluctuate, since particles constantly pass from one subsystem to the other. However, these fluctuations are of an extremely small relative size and are neglected in thermodynamics (if there are of order N particles in each subsystem, then statistical mechanics will show that these fluctuations are of order \sqrt{N}).

Remark 1.1. The entropy function is characteristic of the system considered (in our example, it is the one associated to helium; if we were considering a piece of lead or a mixture of gases, such as the air, the entropy function would be different). It also obviously depends on the set of variables used for its description (mentally splitting the system into two and using 6 variables instead of 3 yields a different entropy function, even though the underlying physical system is unchanged). However, if one considers two systems Σ_1 and Σ_2 of the same type and described by the same set of variables (the latter possibly taking different values), then $S^{\Sigma_1} = S^{\Sigma_2}$. \diamond

Remark 1.2. Let us emphasize that although the thermodynamic properties of a system are entirely contained in its entropy function (or in any of the equations of state or thermodynamic potential derived later), *thermodynamics does not provide tools to determine what this function should be for a specific system* (it can, of course, be measured empirically in a laboratory). As we will see, the determination of these functions for a particular system from first principles is a task that will be devolved to equilibrium statistical mechanics. \diamond

Let us illustrate on some examples how the postulate is used.

Example 1.3. In this first example, we suppose that the system Σ is divided into two subsystems Σ_1 and Σ_2 of volume V_1 , respectively V_2 , by inserting a hard, impermeable, fixed wall. These subsystems have, initially, energy U_1 , respectively U_2 , and contain N_1 , respectively N_2 , particles. We assume that the wall allows the two subsystems to exchange energy, but not particles. So, by assumption, the following quantities are kept fixed: the volumes V_1 and V_2 of the two subsystems, the number N_1 and N_2 of particles they contain and the total energy $U = U_1 + U_2$; these form the constraints. The problem is thus to determine the values \bar{U}^1, \bar{U}^2 of the energy in each of the subsystems once the system has reached equilibrium.



Since V_1, N_1, V_2, N_2 are fixed, the postulate states that the equilibrium values \bar{U}_1 and \bar{U}_2 are found by maximizing

$$(\tilde{U}_1, \tilde{U}_2) \mapsto S(\tilde{U}_1, V_1, N_1) + S(\tilde{U}_2, V_2, N_2) = S(\tilde{U}_1, V_1, N_1) + S(U - \tilde{U}_1, V_2, N_2).$$

We thus see that equilibrium is realized when \bar{U}_1 satisfies

$$\left\{ \frac{\partial S}{\partial \tilde{U}_1}(\tilde{U}_1, V_1, N_1) + \frac{\partial S}{\partial \tilde{U}_1}(U - \tilde{U}_1, V_2, N_2) \right\} \Big|_{\tilde{U} = \bar{U}_1} = 0.$$

Therefore, equilibrium is realized when \bar{U}_1, \bar{U}_2 satisfy

$$\frac{\partial S}{\partial U}(\bar{U}_1, V_1, N_1) = \frac{\partial S}{\partial U}(\bar{U}_2, V_2, N_2).$$

The quantity ⁴

$$\beta \stackrel{\text{def}}{=} \left(\frac{\partial S}{\partial U} \right)_{V,N} \quad (1.3)$$


is called the **inverse (absolute) temperature**. The **(absolute) temperature** is then defined by $T \stackrel{\text{def}}{=} 1/\beta$. It is an empirical fact that the temperature thus defined is positive, that is, the entropy is an increasing function of the energy ⁵. We will therefore assume from now on that $S(U, V, N)$ is *increasing in U* :

$$\left(\frac{\partial S}{\partial U} \right)_{V,N} > 0. \quad (1.4)$$

We conclude that, *if two systems that are isolated from the rest of the universe are allowed to exchange energy, then, once they reach equilibrium, their temperatures (as defined above) will have equalized.*

Note that this agrees with the familiar observation that there will be a heat flow between the two subsystems, until both reach the same temperature. \diamond

Example 1.4. Let us now consider a slightly different situation, in which the wall partitioning our system Σ is allowed to slide:



In this case, the number of particles on each side of the wall is still fixed to N_1 and N_2 , but the subsystems can exchange both energy and volume. The constraint is thus that $U = U_1 + U_2$, $V = V_1 + V_2$, N_1 and N_2 are kept fixed. Proceeding as before, we have to find the values $\bar{U}_1, \bar{U}_2, \bar{V}_1, \bar{V}_2$ maximizing

$$(\bar{U}_1, \bar{U}_2, \bar{V}_1, \bar{V}_2) \mapsto S(\bar{U}_1, \bar{V}_1, N_1) + S(\bar{U}_2, \bar{V}_2, N_2).$$

We deduce this time that, once equilibrium is realized, $\bar{U}_1, \bar{U}_2, \bar{V}_1$ and \bar{V}_2 satisfy

$$\begin{cases} \frac{\partial S}{\partial U}(\bar{U}_1, \bar{V}_1, N_1) = \frac{\partial S}{\partial U}(\bar{U}_2, \bar{V}_2, N_2), \\ \frac{\partial S}{\partial V}(\bar{U}_1, \bar{V}_1, N_1) = \frac{\partial S}{\partial V}(\bar{U}_2, \bar{V}_2, N_2). \end{cases}$$

Again, the first identity implies that the temperatures of the subsystems must be equal. The quantity

$$p \stackrel{\text{def}}{=} T \cdot \left(\frac{\partial S}{\partial V} \right)_{U,N} \quad (1.5)$$

⁴In this introduction, we will follow the custom in thermodynamics and keep the same notation for quantities such as the entropy or temperature, even when seen as functions of different sets of variables. It is thus important, when writing down partial derivatives to specify which are the variables kept fixed.

⁵Actually, there are very special circumstances in which negative temperatures are possible, but we will not discuss them in this book. In any case, the adaptation of what we discuss to negative temperatures is straightforward.

is known as the **pressure** (T , in this definition, is introduced as a convention). We conclude that, *once two systems that can exchange both energy and volume reach equilibrium, their temperatures and pressures will have equalized.* \diamond

Example 1.5. For the third and final example, we suppose that the system is partitioned into two subsystems by a fixed permeable wall, that allows exchange of both particles and energy. The constraints in this case are that $U = U_1 + U_2$, $N = N_1 + N_2$, V_1 and V_2 are kept fixed. This time, we thus obtain that, at equilibrium, \bar{U}_1 , \bar{U}_2 , \bar{N}_1 and \bar{N}_2 satisfy


$$\begin{cases} \frac{\partial S}{\partial U}(\bar{U}_1, V_1, \bar{N}_1) = \frac{\partial S}{\partial U}(\bar{U}_2, V_2, \bar{N}_2), \\ \frac{\partial S}{\partial N}(\bar{U}_1, V_1, \bar{N}_1) = \frac{\partial S}{\partial N}(\bar{U}_2, V_2, \bar{N}_2). \end{cases}$$

Once more, the first identity implies that the temperatures of the subsystems must be equal. The quantity

$$\mu \stackrel{\text{def}}{=} -T \cdot \left(\frac{\partial S}{\partial N} \right)_{U,V} \quad (1.6)$$

is known as the **chemical potential** (the sign, as well as the introduction of T is a convention). We conclude that, *when they reach equilibrium, two systems that can exchange both energy and particles have the same temperature and chemical potential.* \diamond

We have stated the postulate for a very particular case (a gas in a vessel, considered as made up of two subsystems of the same type), but the postulate extends to *any* thermodynamic system. For instance, it can be used to determine how equilibrium is realized when an arbitrary large number of systems are put in contact:



We have discussed a particular case, but (1.3), (1.5) and (1.6) provide the definition of the temperature, pressure and chemical potential for any system characterized by the variables U, V, N (and possibly others) whose entropy function is known.

Several further fundamental properties of the entropy can be readily deduced from the postulate.

Exercise 1.1. Show that the entropy is **positively homogeneous of degree 1**, that is,

$$S(\lambda U, \lambda V, \lambda N) = \lambda S(U, V, N), \quad \forall \lambda > 0. \quad (1.7)$$

Hint: consider first $\lambda \in \mathbb{Q}$.

Exercise 1.2. Show that the entropy is **concave**, that is, for all $\alpha \in [0, 1]$ and any $U_1, U_2, V_1, V_2, N_1, N_2$,

$$S(\alpha U_1 + (1 - \alpha)U_2, \alpha V_1 + (1 - \alpha)V_2, \alpha N_1 + (1 - \alpha)N_2) \geq \alpha S(U_1, V_1, N_1) + (1 - \alpha)S(U_2, V_2, N_2). \quad (1.8)$$

1.1.3 Conjugate intensive quantities and equations of state

The temperature, pressure and chemical potential defined above were all defined via a partial differentiation of the entropy: $\frac{\partial S}{\partial U}, \frac{\partial S}{\partial V}, \frac{\partial S}{\partial N}$. Generally, if X_i is any extensive variable appearing in S ,

$$f_i \stackrel{\text{def}}{=} \frac{\partial S}{\partial X_i}$$

is called the **variable conjugate to X_i** . It is a straightforward consequences of the definitions that, in contrast to U, V, N , the conjugate variables are not extensive, but **intensive**: they remain unchanged under a global scaling of the system: for all $\lambda > 0$,

$$\begin{aligned} T(\lambda U, \lambda V, \lambda N) &= T(U, V, N), \\ p(\lambda U, \lambda V, \lambda N) &= p(U, V, N), \\ \mu(\lambda U, \lambda V, \lambda N) &= \mu(U, V, N). \end{aligned}$$

In other words, T, p and μ are **positively homogeneous of degree 0**.

Differentiating both sides of the identity $S(\lambda U, \lambda V, \lambda N) = \lambda S(U, V, N)$ with respect to λ , at $\lambda = 1$, we obtain

$$S(U, V, N) = \frac{1}{T}U + \frac{p}{T}V - \frac{\mu}{T}N. \quad (1.9)$$

The latter identity is known as the **Euler relation**. It allows one to reconstruct the entropy function from a knowledge of the functional dependence of T, p, μ on U, V, N :

$$T = T(U, V, N), \quad p = p(U, V, N), \quad \mu = \mu(U, V, N). \quad (1.10)$$

These relations are known as the **equations of state**.

1.1.4 Densities

Using homogeneity, we can write

$$S(U, V, N) = VS\left(\frac{U}{V}, 1, \frac{N}{V}\right) \quad \text{or} \quad S(U, V, N) = NS\left(\frac{U}{N}, \frac{V}{N}, 1\right). \quad (1.11)$$

This shows that, when using densities, the entropy can actually be considered as a density as well and seen as a function of two variables rather than three. For example, one can introduce the **energy density** $u \stackrel{\text{def}}{=} \frac{U}{V}$ and the **particle density** $\rho \stackrel{\text{def}}{=} \frac{N}{V}$, and consider the **entropy density**:

$$s(u, \rho) \stackrel{\text{def}}{=} \frac{1}{V}S(uV, V, \rho V). \quad (1.12)$$

Alternatively, using the **energy per particle** $e \stackrel{\text{def}}{=} \frac{U}{N}$ and the **specific volume** (or **volume per particle**) $v \stackrel{\text{def}}{=} \frac{V}{N}$, one can consider the **entropy per particle**:

$$s(e, v) \stackrel{\text{def}}{=} \frac{1}{N} S(eN, vN, N).$$

In particular, its differential satisfies

$$ds = \frac{\partial s}{\partial e} de + \frac{\partial s}{\partial v} dv = \frac{1}{T} de + \frac{p}{T} dv. \quad (1.13)$$

The entropy can thus be recovered (up to an irrelevant additive constant) from the knowledge of two of the three equations of state. This shows that the equations of state are not independent.

Example 1.6 (The ideal gas). Consider a gas of N particles in a container of volume V , at temperature T . An **ideal gas** is a gas which, at equilibrium, is described by the following two equations of state:

$$pV = RNT, \quad U = cRNT,$$

where the constant c , the **specific heat capacity**, depends on the gas and R is some universal constant known as the **gas constant**⁶. Although no such gas exists, it turns out that most gases approximately satisfy such relations when the temperature is not too low and the density of particles is small enough.

When expressed as a function of $v = V/N$, the first equation becomes

$$pv = RT. \quad (1.14)$$

An **isotherm** is obtained by *fixing* the temperature T and studying p as a function of v :




Figure 1.1: An isotherm of the equation of state of the ideal gas: at fixed temperature, the pressure p is proportional to $\frac{1}{v}$.

Let us explain how the two equations of state can be used to determine the entropy for the system. Notice first that the equations can be rewritten as

$$\frac{1}{T} = cR \frac{N}{U} = \frac{cR}{e}, \quad \frac{p}{T} = R \frac{N}{V} = \frac{R}{v}.$$

Therefore, (1.13) becomes

$$ds = \frac{cR}{e} de + \frac{R}{v} dv.$$

⁶Gas constant: $R = 8.3144621 \text{ J m}^{-1} \text{ K}^{-1}$.

Integrating the latter equation, we obtain

$$s(e, v) - s_0 = cR \log(e/e_0) + R \log(v/v_0),$$

where e_0, v_0 is some reference point and s_0 an undetermined constant of integration. We have thus obtained the desired fundamental relation:

$$S(U, V, N) = Ns_0 + NR \log[(U/U_0)^c (V/V_0)(N/N_0)^{-(c+1)}],$$

where we have set $U_0 \stackrel{\text{def}}{=} N_0 e_0, V_0 \stackrel{\text{def}}{=} N_0 v_0$ for some reference N_0 .

Later in this introduction, the equation of state (1.14) will be derived from the microscopic point of view, using the formalism of statistical mechanics. \diamond

It is often convenient to describe a system using certain thermodynamic variables rather than others. For example, in the case of a gas, it might be easier to control pressure and temperature rather than volume and internal energy, so these variables may be better suited for the description of the system. Not only is this possible, but there is a systematic way of determining which **thermodynamic potential** should replace the entropy in this setting and to find the corresponding extremum principle.

1.1.5 Alternative representations; thermodynamic potentials

We now describe how alternative representations corresponding to other sets of variables, both extensive and intensive, can be derived. We treat explicitly the two cases that will be relevant for our analysis based on statistical mechanics later.

The variables (β, V, N) . We will first obtain a description of systems characterized by the set of variables (β, V, N) , replacing U by its conjugate quantity β . Note that, if we want to have a fixed temperature, then the system must be allowed to exchange energy with the environment and is thus *not isolated anymore*. One can see such a system as being in permanent contact with an infinite *thermal reservoir* at fixed temperature $1/\beta$, with which it can exchange energy, but not particles or volume.

We start with some heuristic considerations.



We suppose that our system Σ is put in contact with a much larger system Σ_R , representing the thermal reservoir, with which it can only exchange energy:




Figure 1.2: A system Σ , in contact with a reservoir.

We know from Example 1.3 that, under such conditions, both systems must have the same inverse temperature, denoted by β . The total system Σ_{tot} , of energy U_{tot} , is

considered to be isolated. We denote by U the energy of Σ ; the energy of the reservoir is then $U_{\text{tot}} - U$.

We assume that the reservoir is so much larger than Σ that we can ignore the effect of the actual state of Σ on the values of the intensive parameters associated with the reservoir: $\beta_R (= \beta)$, p_R and μ_R remain constant. In particular, by the Euler relation (1.9), the entropy of the reservoir satisfies

$$S^R(U_{\text{tot}} - U) = \beta_R U_R + \beta_R p_R V_R - \beta_R \mu_R N_R = \beta(U_{\text{tot}} - U) + \beta p_R V_R - \beta \mu_R N_R.$$

Observe that, under our assumptions, the only term in the last expression that depends on the state of Σ is $-\beta U$.

To determine the equilibrium value \bar{U} of the energy of Σ , we must maximize the total entropy (we only indicate the dependence on U , since the volumes and numbers of particles are fixed): \bar{U} is the value realizing the supremum in

$$S^{\text{tot}}(U_{\text{tot}}, \bar{U}) = \sup_U \{S^\Sigma(U) + S^R(U_{\text{tot}} - U)\},$$

which, from what was said before, is equivalent to finding the value of U that realizes the infimum in

$$\hat{F}^\Sigma(\beta) \stackrel{\text{def}}{=} \inf_U \{\beta U - S^\Sigma(U)\}.$$

◇

In view of the preceding considerations, we may expect the function

$$\hat{F}(\beta, V, N) \stackrel{\text{def}}{=} \inf_U \{\beta U - S(U, V, N)\} \quad (1.15)$$

to play a role analogous to that of the entropy, when the temperature, rather than the energy, is kept fixed. The thermodynamic potential $F(T, V, N) \stackrel{\text{def}}{=} T \hat{F}(1/T, V, N)$ is called the **Helmholtz free energy** (once more, the presence of a factor T is due to conventions).

In mathematical terms, \hat{F} is, up to minor differences⁷ the **Legendre transform** of $S(U, V, N)$ with respect to U ; see Appendix B.2 for the basic definition and properties of this transform. The Legendre transform enjoys of several interesting properties. For instance, it has convenient convexity properties (see the exercise below), and it is an involution (on convex function, Theorem B.19). In other words, \hat{F} and S contain the same information.

Observe now that, since we are assuming differentiability, the infimum in (1.15) is attained when $\frac{\partial S}{\partial U} = \beta$. Since we have assumed that the temperature is positive (remember (1.4)), the latter relation can be inverted to obtain $U = U(T, V, N)$. We get

$$F(T, V, N) = U(T, V, N) - TS(U(T, V, N), V, N). \quad (1.16)$$

As a short hand, this relation is often written simply

$$F = U - TS. \quad (1.17)$$

The thermodynamic potential \hat{F} inherits analogues of the fundamental properties of S :

⁷Since S is concave, $-S$ is convex. Therefore, indicating only the dependence of S on U ,

$$\inf_U \{\beta U - S(U)\} = -\sup_U \{(-\beta)U - (-S(U))\}$$

which is minus the Legendre transform (defined in (B.11)) of $-S$, at $-\beta$.

Exercise 1.3. Show that \hat{F} is convex in V and N (the extensive variables), and concave in β (the intensive variable).


Notice that, since \hat{F} is convex in V , we have $\frac{\partial^2 \hat{F}}{\partial V^2} \geq 0$. But, by (1.16),

$$\left(\frac{\partial \hat{F}}{\partial V}\right)_{T,N} = \beta \left(\frac{\partial U}{\partial V}\right)_{T,N} - \underbrace{\left(\frac{\partial S}{\partial U}\right)_{V,N}}_{=\beta} \left(\frac{\partial U}{\partial V}\right)_{T,N} - \left(\frac{\partial S}{\partial V}\right)_{U,N} = -\frac{p}{T}.$$

Therefore, differentiating again with respect to V yields

$$\left(\frac{\partial p}{\partial V}\right)_{T,N} \leq 0, \quad (1.18)$$

a property known as **thermodynamic stability**.

 A system for which $\left(\frac{\partial p}{\partial V}\right)_{T,N} > 0$ would be unstable in the following intuitive sense: any small increase in V would imply an increase in pressure, which in turn would imply an increase in V , etc. \diamond

To state the analogue of the extremum principle in the present case, let us consider a system Σ , kept at temperature T and composed of two subsystems Σ_1 , with parameters T, V^1, N^1 , and Σ_2 , with parameters T, V^2, N^2 . Similarly as before, we assume that there are constraints on the admissible values of V^1, N^1, V^2, N^2 and are interested in determining the corresponding equilibrium values $\bar{V}^1, \bar{V}^2, \bar{N}^1, \bar{N}^2$.

Exercise 1.4. Show that \hat{F} satisfies the following extremum principle: the equilibrium values $\bar{V}^1, \bar{V}^2, \bar{N}^1, \bar{N}^2$ are those minimizing

$$\hat{F}(T, \tilde{V}^1, \tilde{N}^1) + \hat{F}(T, \tilde{V}^2, \tilde{N}^2) \quad (1.19)$$

among all $\tilde{V}^1, \tilde{N}^1, \tilde{V}^2, \tilde{N}^2$ compatible with the constraints.

The variables (β, V, μ) . We can proceed in the same way for a system characterized by the variables (β, V, μ) . Such a system must be able to exchange both energy and particles with a reservoir.

This time, the thermodynamic potential associated to the variables $\beta, V, \hat{\mu} \stackrel{\text{def}}{=} -\mu/T$ is defined by

$$\hat{\Phi}_G(\beta, V, \hat{\mu}) \stackrel{\text{def}}{=} \inf_{U,N} \{ \beta U + \hat{\mu} N - S(U, V, N) \}. \quad (1.20)$$

The function $\Phi_G(T, V, \mu) \stackrel{\text{def}}{=} T \hat{\Phi}_G(1/T, V, -\mu/T)$ is called the **grand potential**. As for the Helmholtz free energy, it can be shown that $\hat{\Phi}_G$ is concave in β and $\hat{\mu}$, and convex in V . The extremum principle extends also naturally to this case. Proceeding as before, we can write

$$\Phi_G = U - \mu N - TS \quad (1.21)$$

(with an interpretation analogous to the one done in (1.17)). Since, by the Euler relation (1.9), $TS = U + pV - \mu N$, we deduce that

$$\Phi_G = -pV, \quad (1.22)$$

so that $-\Phi_G/V$ coincides with the pressure of the system (expressed, of course, as a function of (T, V, μ)).

Generically⁸, the descriptions of a given system, in terms of various sets of thermodynamic variables, yield the same result at equilibrium. For example, if we start with a microstate (U, V, N) and compute the value of β at equilibrium, then starting with the macrostate (β, V, N) (with that particular value of β) and computing the equilibrium value of the energy yields U again.

In the following section, we leave aside the general theory and consider an example, in which an equation of state is progressively obtained from a combination of experimental observations and theoretical considerations.

1.1.6 Condensation and the Van der Waals–Maxwell Theory

Although rather accurate in various situations, the predictions made by the equation of state of the ideal gas are no longer valid at low temperature or at high density. In particular, the behavior observed for a *real* gas at low temperature is of the following type (compare Figure 1.3 with Figure 1.1): When ν (the volume per par-




Figure 1.3: An isotherm of a real gas at low enough temperature. See below (Figure 1.6) for a plot of realistic values measured in the laboratory.

ticle) is large, the density of the gas is low, it is homogeneous (the same density is observed in all subsystems) and the pressure is well approximated by the ideal gas behavior. However, decreasing ν , one reaches a value $\nu = \nu_g$, called the **condensation point**, at which the following phenomenon is observed: macroscopic droplets of liquid start to appear throughout the system. As ν is further decreased, the fraction of the volume occupied by the gas decreases while that of the liquid increases. Nevertheless, the pressures inside the gas and inside the droplets are equal and constant. This goes on until another value $\nu = \nu_l < \nu_g$ is reached, at which all the gas has been transformed into liquid. When decreasing the volume to values $\nu < \nu_l$, the pressure starts to increase again, but at a much higher rate due to the fact that

⁸The word “generically” is used to exclude first-order phase transitions, since, when the latter occur, the assumptions of smoothness and invertibility that we use to invert the relations between all these variables fail in general. Such issues will be discussed in detail in the framework of equilibrium statistical mechanics in later chapters.

the system now contains only liquid and the latter is almost incompressible⁹. The range $[\nu_l, \nu_g]$ is called the **coexistence plateau**.

The condensation phenomenon shows that, from a mathematical point of view, the equations of state of a system are not always smooth in their variables (the pressure at the points ν_l and ν_g , for instance). The appearance of such singularities is the signature of *phase transitions*, one of the main themes studied in this book.

For the time being, we will only describe the way by which the equation of state of the ideal gas can be modified to account for the behavior observed in real gases.


The Van der Waals Theory of Condensation

The first theory of condensation originated with Van der Waals' thesis in 1873. Van der Waals succeeded in establishing an equation of state that described significant deviations from the equation of the ideal gas. His analysis is based on the following two fundamental hypotheses on the *microscopic* structure of the system¹⁰:

1. The gas has microscopic constituents, the *particles*. Particles are extended in space. They interact *repulsively* at short distances: particles do not overlap.
2. At larger distances, the particles also interact in an *attractive* way. This part of the interaction is characterized by a constant $a > 0$ called *specific attraction*.

The short-distance repulsion indicates that, in the equation of state, the volume ν available to each particle should be replaced by a smaller quantity $\nu - b$ taking into account the volume of space each particle occupies. In order to deal with the attractive part of the interaction, Van der Waals assumed that *the system is homogeneous*, a drastic simplification to which we will return later. These two hypotheses led Van der Waals to his famous equation of state:

$$\left(p + \frac{a}{\nu^2}\right)(\nu - b) = RT. \quad (1.23)$$

 The term $\frac{a}{\nu^2}$ can be understood intuitively as follows. Let $\rho \stackrel{\text{def}}{=} \frac{N}{V} = \frac{1}{\nu}$ denote the density of the gas (number of particles per unit volume). We assume that each particle only interacts with particles in its neighborhood, up to some large, finite distance. By the homogeneity assumption, the total force exerted by the other particles on a given particle deep inside the vessel averages to zero. However, for a particle in a small layer along the boundary of the vessel, the force resulting from its interaction with other particles has a positive component away from the boundary, since there are more particles in this direction. This inward force reduces the pressure exerted on the boundary of the vessel. Now, the number of particles in this layer along a portion of the boundary of unit area is proportional to ρ . The total force on each particle in this layer is proportional to the number of particles it interacts with, which is also proportional to ρ . We conclude that the reduction in the pressure, compared to an ideal gas, is proportional to $\rho^2 = 1/\nu^2$. A rigorous derivation will be given in Chapter 4. ◇

⁹At even smaller specific volumes, the system usually goes through another phase transition at which the liquid transforms into a solid, but we will not discuss this issue here.

¹⁰The interaction between two particles at distance r is often modeled by a **Lennard-Jones potential**, that is, an interaction potential of the form $A r^{-12} - B r^{-6}$, with $A, B > 0$. The first term models the short-range repulsion between the particles, while the second one models the long-range attraction.

Of course, the ideal gas is recovered by setting $a = b = 0$. The analysis of (1.23) (see Exercise 1.5 below) reveals that, in contrast to those of the ideal gas, the isotherms present different behaviors depending on the temperature being above or below some **critical temperature** T_c , see Figure 1.4.




Figure 1.4: Isotherms of the Van der Waals equation of state (1.23) at low and high temperature.

For **supercritical** temperatures, $T > T_c$, the behavior is qualitatively the same as for the ideal gas; in particular, p is strictly decreasing in v . However, for **subcritical** temperatures, $T < T_c$, there is an interval over which $(\frac{\partial p}{\partial v})_T > 0$, thereby violating (1.18); this shows that this model has unphysical consequences. Moreover, in real gases (remember Figure 1.3), there is a plateau in the graph of the pressure, corresponding to values of v at which the gas and liquid phases coexist. This plateau is absent from Van der Waals' isotherms.

Exercise 1.5. Study the equation of state (1.23) for $v > b$ and show that there exists a critical temperature,

$$T_c = T_c(a, b) \stackrel{\text{def}}{=} \frac{8a}{27Rb},$$

such that the following occurs:

- When $T > T_c$, $v \mapsto p(v, T)$ is decreasing everywhere.
- When $T < T_c$, $v \mapsto p(v, T)$ is increasing on some interval.

Maxwell's Construction

Van der Waals' main simplifying hypothesis was the assumption that the system remains *homogeneous*, which by itself makes the theory inadequate to describe the inhomogeneities that must appear at condensation.

In order to include the condensation phenomenon in Van der Waals' theory, Maxwell [235] proposed a natural, albeit ad hoc, procedure to modify the low-temperature isotherms given by (1.23). Since the goal is to allow the system to split into regions that contain either gas or liquid, the latter should be at equal temperature and pressure; he thus replaced $p(v)$, on a well-chosen interval $[v_l, v_g]$, by a constant p_s , called **saturation pressure**.

From a physical point of view, Maxwell's determination of p_s , and hence of v_g and v_l , can be understood as follows. The integral $\int_{v_l}^{v_g} p(v) dv$ represents the area under the graph of the isotherm, between v_l and v_g , but it also represents the amount of *work* necessary to compress the gas from v_g down to v_l . Therefore, if one is to replace $p(v)$ by a constant value between v_l and v_g , this value should be chosen such that the work required for that compression be the *same* as the original one. That is, v_l , v_g and p_s should satisfy

$$\int_{v_g}^{v_l} p(v) dv = \int_{v_g}^{v_l} p_s dv,$$

which gives

$$\int_{v_g}^{v_l} p(v) dv = p_s \cdot (v_g - v_l). \quad (1.24)$$

This determination of p_s can also be given a geometrical meaning: it is the unique height at which a coexistence interval can be chosen in such a way that the two areas delimited by the Van der Waals isotherm and the segment are equal. For that reason, the procedure proposed by Maxwell is usually called **Maxwell's equal-area rule**, or simply **Maxwell's Construction**. We denote the resulting isotherm by $v \mapsto \text{MC } p(v, T)$; see Figure 1.5.




Figure 1.5: Maxwell's Construction. The height of the segment, p_s , is chosen in such a way that the two connected regions delimited by the graph of p and the segment (shaded in the picture) have equal areas.

Although it relies on a mixture of two conflicting hypotheses (first assume homogeneity, build an equation of state and then modify it by plugging in the condensation phenomenon, by hand, using the equal-area rule), the Van der Waals–Maxwell theory often yields satisfactory quantitative results. It is a landmark in the understanding of the thermodynamics of the coexistence of liquids and gases and remains widely taught in classrooms today.

There are many sources in the literature where the interested reader can find additional information about the Van der Waals–Maxwell theory; see, for instance, [89, 130]. We will return to the liquid-vapor equilibrium in a more systematic (and rigorous) way in Chapter 4.

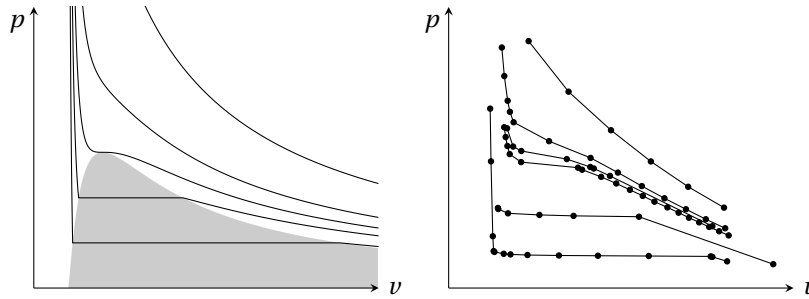


Figure 1.6: *Left*: some of the isotherms resulting from the Van der Waals–Maxwell theory, which will be discussed in detail in Chapter 4. The shaded region indicates the pairs (v, p) for which coexistence occurs. *Right*: the isotherms of carbonic acid, measured by Thomas Andrews in 1869 [12].

1.2 From Micro to Macro: Statistical Mechanics

Just as was the case for thermodynamics, the object of statistical mechanics is the description of macroscopic systems. In sharp contrast to the latter, however, statistical mechanics relies on a reductionist approach, whose goal is to derive the macroscopic properties of a system from the microscopic description provided by the fundamental laws of physics. This ambitious program was initiated in the second half of the 19th Century, with Maxwell, Boltzmann and Gibbs as its main contributors. It was essentially complete, as a general framework, when Gibbs published his famous treatise [137] in 1902. The theory has known many important developments since then, in particular regarding the fundamental problem of phase transitions.

As we already discussed, the enormous number of microscopic variables involved renders impossible the problem of deriving the macroscopic properties of a system directly from an application of the underlying fundamental theory describing its microscopic constituents. However, rather than giving up, one might try to use this at our advantage: indeed, the huge number of objects involved makes it conceivable that a *probabilistic* approach might be very efficient (after all, in how many areas does one have samples of size of order 10^{23} ?). That is, the first step in statistical mechanics is to abandon the idea of providing a complete deterministic description of the system and to search instead for a *probability distribution* over the set of all microstates, which yields predictions compatible with the observed macrostate. Such a distribution should provide the probability of observing a given microstate and should make it possible to compute the probability of events of interest or the averages of relevant physical quantities.

This approach can also be formulated as follows: suppose that the only information we have on a given macroscopic system (at equilibrium) is its microscopic description (namely, we know the set of microstates and how to compute their energy) and the values of a few macroscopic variables (the same fixed in thermodynamics). Equipped with this information, and only this information, what can be said about a typical microstate?

The main relevant questions are therefore the following:

- What probability distributions should one use to describe large systems at

equilibrium?

- With a suitable probability distribution at hand, what can be said about the other macroscopic observables? Is it possible to show that their distribution *concentrates* on their most probable value, when the system becomes large, thus yielding deterministic results for such quantities?
- How can this description be related to the one provided by equilibrium thermodynamics?
- Can phase transitions be described within this framework?

Note that the goal of equilibrium statistical mechanics is not limited to the computation of the fundamental quantities appearing in equilibrium thermodynamics. The formalism of statistical mechanics allows one to investigate numerous problems outside the scope of the latter theory. It allows for example to analyze *fluctuations* of macroscopic quantities in large finite systems and thus obtain a finer description than the one provided by thermodynamics.

In this section, we will introduce the central concepts of equilibrium statistical mechanics, valid for general systems, not necessarily gases and liquids, although such systems (as well as magnets) will be used in our illustrative examples.

As mentioned above, we assume that we are given the following basic inputs from a more fundamental theory describing the system's microscopic constituents:

1. *The set Ω of microstates.* To simplify the exposition and since this is enough to explain the general ideas, we assume that the set Ω of microstates describing the system is finite:

$$|\Omega| < \infty. \quad (1.25)$$

2. *The interactions between the microscopic constituents,* in the form of the energy $\mathcal{H}(\omega)$ associated to each microstate $\omega \in \Omega$:

$$\mathcal{H} : \Omega \rightarrow \mathbb{R},$$

called the **Hamiltonian** of the system. We will use $\mathcal{U} \stackrel{\text{def}}{=} \{U = \mathcal{H}(\omega) : \omega \in \Omega\}$.



Assumption (1.25) will require space to be discretized. The latter simplification, which will be made throughout the book, may seem rather extreme. It turns out however that many phenomena of interest can still be investigated in this setting, while the mathematical analysis becomes much more tractable. \diamond

We denote the set of probability distributions on Ω by $\mathcal{M}_1(\Omega)$. Since we assume that Ω is finite, a distribution¹¹ $\mu \in \mathcal{M}_1(\Omega)$ is entirely characterized by the collection $(\mu(\{\omega\}))_{\omega \in \Omega}$ of the probabilities associated to each microstate $\omega \in \Omega$; we will usually abbreviate $\mu(\omega) \equiv \mu(\{\omega\})$. By definition, $\mu(\omega) \geq 0$ for all $\omega \in \Omega$ and $\sum_{\omega \in \Omega} \mu(\omega) = 1$. We call **observable** the result of a measurement on the system. Mathematically, it

¹¹Although we already use it to denote the chemical potential, we also use the letter μ to denote a generic element of $\mathcal{M}_1(\Omega)$, as done very often in the literature.

corresponds to a random variable $f : \Omega \rightarrow \mathbb{R}$. We will often denote the **expected value** of an observable f under $\mu \in \mathcal{M}_1(\Omega)$ by

$$\langle f \rangle_\mu \stackrel{\text{def}}{=} \sum_{\omega \in \Omega} f(\omega) \mu(\omega),$$

although some alternative notations will occasionally be used in later chapters.

1.2.1 The microcanonical ensemble



The term “ensemble” was originally introduced by Gibbs [2]. In more modern terms, it could simply be considered as a synonym of “probability space”. In statistical mechanics, working in a specific ensemble usually means adopting either of the three descriptions described below: microcanonical, canonical, grand canonical. ♦

In Section 1.1.1, we saw that it was convenient, when describing an isolated system at equilibrium, to use extensive conserved quantities as macroscopic variables, such as U, V, N in our gas example.

We start our probabilistic description of an isolated system in the same way. The relevant conserved quantities depend on the system under consideration, but always contain the energy. Quantities such as the number of particles and the volume are assumed to be encoded into the set of microstates. Namely, we denote by $\Omega_{\Lambda;N}$ the set of all microstates describing a system of N particles located inside a domain Λ of volume $|\Lambda| = V$ (see the definition of the lattice gas in Section 1.2.4 for a specific example). For our current discussion, we assume that the only additional conserved quantity is the energy.

So, let us assume that the energy of the system is fixed to some value U . We are looking for a probability distribution on $\Omega_{\Lambda;N}$ that is concentrated on the set of all microstates compatible with this constraint, that is, on the **energy shell**

$$\Omega_{\Lambda;U,N} \stackrel{\text{def}}{=} \{\omega \in \Omega_{\Lambda;N} : \mathcal{H}(\omega) = U\}.$$

If this is all the information we have on the system, then the simplest and most natural assumption is that all configurations $\omega \in \Omega_{\Lambda;U,N}$ are *equiprobable*. Indeed, if we consider the distribution μ as a description of the *knowledge* we have of the system, then the uniform distribution represents faithfully the totality of our information. This is really just an application of Laplace’s Principle of Insufficient Reason [3]. It leads naturally to the following definition:

Definition 1.7. Let $U \in \mathcal{U}$. The **microcanonical distribution (at energy U)**, $\nu_{\Lambda;U,N}^{\text{Mic}}$, associated to a system composed of N particles located in a domain Λ , is the uniform probability distribution concentrated on $\Omega_{\Lambda;U,N}$:

$$\nu_{\Lambda;U,N}^{\text{Mic}}(\omega) \stackrel{\text{def}}{=} \begin{cases} \frac{1}{|\Omega_{\Lambda;U,N}|} & \text{if } \omega \in \Omega_{\Lambda;U,N}, \\ 0 & \text{otherwise.} \end{cases} \quad (1.26)$$

Although the microcanonical distribution has the advantage of being natural and easy to define, it can be difficult to work with, since counting the configurations on the energy shell can represent a challenging combinatorial problem, even in simple cases. Moreover, one is often more interested in the description of a system at a fixed temperature T , rather than at fixed energy U .

1.2.2 The canonical ensemble

Our goal now is to determine the relevant probability distribution to describe a macroscopic system at equilibrium *at a fixed temperature*. As discussed earlier, such a system is not isolated anymore, but in contact with a thermal reservoir with which it can exchange energy.



Once we have the microcanonical description, the problem of constructing the relevant probability distribution describing a system at equilibrium with an infinite reservoir with which it can exchange energy and/or particles becomes conceptually straightforward from a probabilistic point of view. Indeed, similarly to what we did in Section 1.1.5, we can consider a system Σ in contact with a reservoir Σ_R , the union of the two forming a system Σ_{tot} isolated from the rest of the universe, as in Figure 1.2. Since Σ_{tot} is isolated, it is described by the microcanonical distribution, and the probability distribution of the subsystem Σ can be deduced from this microcanonical distribution by integrating over all the variables pertaining to the reservoir. That is, the measure describing Σ is the marginal of the microcanonical distribution corresponding to the subsystem Σ . This is a well-posed problem, but also a difficult one. It turns out that implementing rigorously such an approach is possible, once the statement is properly formulated, but a detailed discussion is beyond the scope of this book. In Section 4.7.1, we consider a simplified version of this problem; see also the discussion (and the references) in Section 6.14.1.

Instead of following this path, we are going to use a generalization of the argument that led us to the microcanonical distribution to derive the distribution describing a system interacting with a reservoir. \diamond

When we discussed the microcanonical distribution above, we argued that the uniform measure was the proper one to describe our knowledge of the system when the only information available to us is the total energy. We would like to proceed similarly here. The problem is that the temperature is not a mechanical quantity as was the energy (that is, there is no observable $\omega \mapsto T(\omega)$), so one cannot restrict the set $\Omega_{\Lambda;N}$ to microstates with fixed temperatures. We thus need to take a slightly different point of view.

Even though the energy of our system is not fixed anymore, one might still measure its *average*, which we also denote by U . In this case, extending the approach used in the microcanonical case corresponds to looking for the probability distribution $\mu \in \mathcal{M}_1(\Omega_{\Lambda;N})$ that best encapsulates the fact that our only information about the system is that $\langle \mathcal{H} \rangle_\mu = U$.

The Maximum Entropy Principle. In terms of randomness, the outcomes of a random experiment whose probability distribution is uniform are the *least predictable*. Thus, what we did in the microcanonical case was to choose the most unpredictable distribution on configurations with fixed energy.

Let Ω be an arbitrary finite set of microstates. One convenient (and essentially unique, see below) way of quantifying the unpredictability of the outcomes of a probability distribution μ on Ω is to use the notion of *entropy* introduced in information theory by Shannon [301]:

Definition 1.8. The (Shannon) entropy ^[4] of $\mu \in \mathcal{M}_1(\Omega)$ is defined by

$$S_{\text{Sh}}(\mu) \stackrel{\text{def}}{=} - \sum_{\omega \in \Omega} \mu(\omega) \log \mu(\omega). \quad (1.27)$$

Exercise 1.6. $S_{\text{Sh}} : \mathcal{M}_1(\Omega) \rightarrow \mathbb{R}$ is concave: for all $\mu, \nu \in \mathcal{M}_1(\Omega)$ and all $\alpha \in [0, 1]$,

$$S_{\text{Sh}}(\alpha\mu + (1-\alpha)\nu) \geq \alpha S_{\text{Sh}}(\mu) + (1-\alpha)S_{\text{Sh}}(\nu).$$

The Shannon entropy provides a characterization of the uniform distribution through a variational principle.

Lemma 1.9. The uniform distribution on Ω , $\nu^{\text{Unif}}(\omega) \stackrel{\text{def}}{=} \frac{1}{|\Omega|}$, is the unique probability distribution at which S_{Sh} attains its maximum:

$$S_{\text{Sh}}(\nu^{\text{Unif}}) = \sup_{\mu \in \mathcal{M}_1(\Omega)} S_{\text{Sh}}(\mu). \quad (1.28)$$

Proof. Consider $\psi(x) \stackrel{\text{def}}{=} -x \log x$, which is concave. Using Jensen's Inequality (see Appendix B.8.1) gives

$$\begin{aligned} S_{\text{Sh}}(\mu) &= |\Omega| \sum_{\omega \in \Omega} \frac{1}{|\Omega|} \psi(\mu(\omega)) \\ &\leq |\Omega| \psi\left(\sum_{\omega \in \Omega} \frac{1}{|\Omega|} \mu(\omega)\right) = |\Omega| \psi\left(\frac{1}{|\Omega|}\right) = \log |\Omega| = S_{\text{Sh}}(\nu^{\text{Unif}}), \end{aligned}$$

with equality if and only if $\mu(\cdot)$ is constant, that is, if $\mu = \nu^{\text{Unif}}$. \square

Since it is concave, with a unique maximum at ν^{Unif} , Shannon's entropy provides a way of measuring how far a distribution is from being uniform. As shown in Appendix B.11, the Shannon entropy is the unique (up to a multiplicative constant) such function (under suitable assumptions). We can thus use it to select, among all probability distributions in some set, the one that is “the most uniform”.

Namely, assume that we have a set of probability distributions $\mathcal{M}'_1(\Omega) \subset \mathcal{M}_1(\Omega)$ representing the set of all distributions compatible with the information at our disposal. Then, the one that best describes our state of knowledge is the one maximizing the Shannon entropy; this way of selecting a distribution is called the **Maximum Entropy Principle**. Its application to statistical mechanics, as an extension of the Principle of Insufficient Reason, was pioneered by Jaynes [181].

For example, the microcanonical distribution $\nu_{\Lambda;U,N}^{\text{Mic}}$ has maximal entropy among all distributions concentrated on the energy shell $\Omega(U)$:

$$S_{\text{Sh}}(\nu_{\Lambda;U,N}^{\text{Mic}}) = \sup_{\substack{\mu \in \mathcal{M}_1(\Omega_{\Lambda;N}): \\ \mu(\Omega_{\Lambda;U,N})=1}} S_{\text{Sh}}(\mu). \quad (1.29)$$

The canonical Gibbs distribution. We apply the Maximum Entropy Principle to find the probability distribution $\mu \in \mathcal{M}_1(\Omega_{\Lambda;N})$ that maximizes S_{Sh} , under the constraint that $\langle \mathcal{H} \rangle_{\mu} = U$. From an analytic point of view, this amounts to searching

for the collection $(\mu(\omega))_{\omega \in \Omega_{\Lambda;N}}$ of nonnegative real numbers that solves the following optimization problem:

$$\text{Minimize } \sum_{\omega \in \Omega_{\Lambda;N}} \mu(\omega) \log \mu(\omega) \text{ when } \begin{cases} \sum_{\omega \in \Omega_{\Lambda;N}} \mu(\omega) = 1, \\ \sum_{\omega \in \Omega_{\Lambda;N}} \mu(\omega) \mathcal{H}(\omega) = U. \end{cases} \quad (1.30)$$

For this problem to have a solution, we require that $U \in [U_{\min}, U_{\max}]$, where $U_{\min} \stackrel{\text{def}}{=} \inf_{\omega} \mathcal{H}(\omega)$, $U_{\max} = \sup_{\omega} \mathcal{H}(\omega)$. Such problems with constraints can be solved by using the method of Lagrange multipliers. Since there are two constraints, let us introduce two Lagrange multipliers, λ and β , and define the following Lagrange function:

$$L(\mu) \stackrel{\text{def}}{=} \sum_{\omega \in \Omega_{\Lambda;N}} \mu(\omega) \log \mu(\omega) + \lambda \sum_{\omega \in \Omega_{\Lambda;N}} \mu(\omega) + \beta \sum_{\omega \in \Omega_{\Lambda;N}} \mu(\omega) \mathcal{H}(\omega).$$

The optimization problem then turns into the analytic study of a system of $|\Omega_{\Lambda;N}| + 2$ unknowns:

$$\begin{cases} \nabla L = 0, \\ \sum_{\omega \in \Omega_{\Lambda;N}} \mu(\omega) = 1, \\ \sum_{\omega \in \Omega_{\Lambda;N}} \mu(\omega) \mathcal{H}(\omega) = U, \end{cases}$$

where ∇ is the gradient involving the derivatives with respect to each $\mu(\omega)$, $\omega \in \Omega_{\Lambda;N}$. The condition $\nabla L = 0$ thus corresponds to

$$\frac{\partial L}{\partial \mu(\omega)} = \log \mu(\omega) + 1 + \lambda + \beta \mathcal{H}(\omega) = 0, \quad \forall \omega \in \Omega_{\Lambda;N}.$$

The solution is of the form $\mu(\omega) = e^{-\beta \mathcal{H}(\omega) - 1 - \lambda}$. The first constraint $\sum \mu(\omega) = 1$ implies that $e^{1+\lambda} = \sum_{\omega \in \Omega_{\Lambda;N}} e^{-\beta \mathcal{H}(\omega)}$. In conclusion, we see that the distribution we are after is

$$\mu_{\beta}(\omega) \stackrel{\text{def}}{=} \frac{e^{-\beta \mathcal{H}(\omega)}}{\sum_{\omega' \in \Omega_{\Lambda;N}} e^{-\beta \mathcal{H}(\omega')}},$$

where the Lagrange multiplier β must be chosen such that

$$\sum_{\omega \in \Omega_{\Lambda;N}} \mu_{\beta}(\omega) \mathcal{H}(\omega) = U. \quad (1.31)$$

Note that this equation always possesses exactly one solution $\beta = \beta(U)$ when $U \in (U_{\min}, U_{\max})$; this is an immediate consequence of the following

Exercise 1.7. Show that $\beta \mapsto \langle \mathcal{H} \rangle_{\mu_{\beta}}$ is continuously differentiable, decreasing and

$$\lim_{\beta \rightarrow -\infty} \langle \mathcal{H} \rangle_{\mu_{\beta}} = U_{\max}, \quad \lim_{\beta \rightarrow +\infty} \langle \mathcal{H} \rangle_{\mu_{\beta}} = U_{\min}.$$

Since β can always be chosen in such a way that the average energy takes a given value, it will be used from now on as the natural parameter for the canonical distribution. To summarize, the probability distribution describing a system at equilibrium that can exchange energy with the environment and possesses an average energy is assumed to have the following form:

Definition 1.10. The **canonical Gibbs distribution** at parameter β associated to a system of N particles located in a domain Λ is the probability distribution on $\Omega_{\Lambda;N}$ defined by

$$\mu_{\Lambda;\beta,N}(\omega) \stackrel{\text{def}}{=} \frac{e^{-\beta \mathcal{H}(\omega)}}{Z_{\Lambda;\beta,N}}.$$

The exponential $e^{-\beta \mathcal{H}}$ is called the **Boltzmann weight** and the normalizing sum

$$Z_{\Lambda;\beta,N} \stackrel{\text{def}}{=} \sum_{\omega \in \Omega_{\Lambda;N}} e^{-\beta \mathcal{H}(\omega)}$$

is called the **canonical partition function**.

We still need to provide an interpretation for the parameter β . As will be argued below, in Section 1.3 (see also Exercise 1.12), β can in fact be identified with the inverse temperature.

Exercise 1.8. Using the Maximum Entropy Principle, determine the probability distribution of maximal entropy, $\mu = (\mu(1), \dots, \mu(6))$, for the outcomes of a dice whose expected value is 4.

1.2.3 The grand canonical ensemble

Let us generalize the preceding discussion to the case of a system at equilibrium that can exchange both energy and particles with the environment. From the thermodynamical point of view, such a system is characterized by its temperature and its chemical potential.

One can then proceed exactly as in the previous section and apply the Maximum Entropy Principle to the set of all probability distributions with prescribed average energy and average number of particles. Let us denote by $\Omega_{\Lambda} \stackrel{\text{def}}{=} \bigcup_{N \geq 0} \Omega_{\Lambda;N}$ the set of all microstates with an arbitrary number of particles all located inside the region Λ . A straightforward adaptation of the computations done above (with two Lagrange multipliers β and $\hat{\mu}$) shows that the relevant distribution in this case should take the following form (writing $\hat{\mu} = -\beta\mu$).

Definition 1.11. The **grand canonical Gibbs distribution** at parameters β and μ associated to a system of particles located in a region Λ is the probability distribution on Ω_{Λ} defined by

$$\nu_{\Lambda;\beta,\mu}(\omega) \stackrel{\text{def}}{=} \frac{e^{-\beta(\mathcal{H}(\omega) - \mu N)}}{Z_{\Lambda;\beta,\mu}}, \quad \text{if } \omega \in \Omega_{\Lambda;N}.$$

The normalizing sum

$$Z_{\Lambda;\beta,\mu} \stackrel{\text{def}}{=} \sum_N e^{\beta\mu N} \sum_{\omega \in \Omega_{\Lambda;N}} e^{-\beta \mathcal{H}(\omega)}$$

is called the **grand canonical partition function**.

Similarly as before, the parameters β and μ have to be chosen in such a way that the expected value of the energy and number of particles match the desired values. In Section 1.3, we will argue that β and μ can be identified with the inverse temperature and chemical potential.

1.2.4 Examples: Two models of a gas.

We now present two examples of statistical mechanical models of a gas. The first one, although outside the main theme of this book, will be a model in the continuum, based on the description provided by Hamiltonian mechanics. The second one will be a lattice gas model, which can be seen as a simplification of the previous model and will be the main topic of Chapter 4.

The continuum gas.

We model a gas composed of N particles, contained in a vessel represented by a bounded subset $\Lambda \subset \mathbb{R}^d$. As a reader familiar with Hamiltonian mechanics might know, the state of such a system consists in the collection $(p_k, q_k)_{k=1, \dots, N}$ of the momentum $p_k \in \mathbb{R}^d$ and the position $q_k \in \Lambda$ of each particle. In particular, the set of microstates is

$$\Omega_{\Lambda;N} = (\mathbb{R}^d \times \Lambda)^N.$$

The Hamiltonian takes the usual form of a sum of a kinetic energy and a potential energy:

$$\mathcal{H}(p_1, q_1, \dots, p_N, q_N) \stackrel{\text{def}}{=} \sum_{k=1}^N \frac{\|p_k\|_2^2}{2m} + \sum_{1 \leq i < j \leq N} \phi(\|q_j - q_i\|_2),$$

where m is the mass of each particle and the potential ϕ encodes the contribution to the total energy due to the interaction between the i th and j th particles, assumed to depend only on the distance $\|q_j - q_i\|_2$ between the particles.

Let us consider the canonical distribution associated to such a system. Our discussion above does not apply verbatim, since we assumed that $\Omega_{\Lambda;N}$ was finite, while here it is a continuum. Nevertheless, the conclusion in this case is the natural generalization of what we saw earlier. Namely, the probability of an event B under the canonical distribution at inverse temperature β is also defined using the Boltzmann weight:

$$\mu_{\Lambda;\beta,N}(B) \stackrel{\text{def}}{=} \frac{1}{Z_{\Lambda;\beta,N}} \int_{\Omega_{\Lambda;N}} \mathbf{1}_B e^{-\beta \mathcal{H}(p_1, q_1, \dots, p_N, q_N)} dp_1 dq_1 \dots dp_N dq_N,$$

where $\mathbf{1}_B = \mathbf{1}_B(p_1, q_1, \dots, q_N, p_N)$ is the indicator of B and

$$Z_{\Lambda;\beta,N} \stackrel{\text{def}}{=} \int_{\Omega_{\Lambda;N}} e^{-\beta \mathcal{H}(p_1, q_1, \dots, p_N, q_N)} dp_1 dq_1 \dots dp_N dq_N.$$

Note that we cannot simply give the probability of each individual microstate, since they all have zero probability. Thanks to the form of the Hamiltonian, the integration over the momenta can be done explicitly:

$$Z_{\Lambda;\beta,N} = \left\{ \int_{\mathbb{R}} e^{-\frac{\beta}{2m} p^2} dp \right\}^{dN} \int_{\Lambda^N} e^{-\beta \mathcal{H}^{\text{conf}}(q_1, \dots, q_N)} dq_1 \dots dq_N = \left(\frac{2\pi m}{\beta} \right)^{dN/2} Z_{\Lambda;\beta,N}^{\text{conf}},$$

where we have introduced the **configuration integral**

$$Z_{\Lambda;\beta,N}^{\text{conf}} \stackrel{\text{def}}{=} \int_{\Lambda^N} e^{-\beta \mathcal{H}^{\text{conf}}(q_1, \dots, q_N)} dq_1 \dots dq_N,$$

and $\mathcal{H}^{\text{conf}}(q_1, \dots, q_N) \stackrel{\text{def}}{=} \sum_{1 \leq i < j \leq N} \phi(\|q_j - q_i\|_2)$.

In fact, when an event B only depends on the positions of the particles, not on their momenta, the factors originating from the integration over the momenta cancel in the numerator and in the denominator, giving

$$\mu_{\Lambda;\beta,N}(B) = \frac{1}{Z_{\Lambda;\beta,N}^{\text{conf}}} \int_{\Lambda^N} \mathbf{1}_B e^{-\beta \mathcal{H}^{\text{conf}}(q_1, \dots, q_N)} dq_1 \cdots dq_N.$$

We thus see that the difficulties in analyzing this gas from the point of view of the canonical ensemble come from the positions: the momenta have no effect on the probability of events depending only on the positions, while they only contribute an explicit prefactor to the partition function. The position-dependent part, however, is difficult to handle as soon as the interaction potential ϕ is not trivial.

Usually, ϕ is assumed to contain two terms, corresponding respectively to the short- and long-range part of the interaction:

$$\phi(x) = \phi_{\text{short}}(x) + \phi_{\text{long}}(x).$$

If we assume that the particles are identified with small spheres of fixed radius $r_0 > 0$, a simple choice for ϕ_{short} is the following *hard-core* interaction:

$$\phi_{\text{short}}(x) \stackrel{\text{def}}{=} \begin{cases} +\infty & \text{if } |x| \leq 2r_0, \\ 0 & \text{otherwise.} \end{cases}$$

The long-range part of the interaction can be of any type, but it should at least vanish at long distance:

$$\phi_{\text{long}}(x) \rightarrow 0 \quad \text{when } |x| \rightarrow \infty.$$

The decay at $+\infty$ should be fast enough, so as to guarantee the existence of $Z_{\Lambda;\beta,N}^{\text{conf}}$ for instance, but we will not describe this in any further detail.

Unfortunately, even under further strong simplifying assumptions on ϕ , the mathematical analysis of such systems, in particular the computation of $Z_{\Lambda;\beta,N}^{\text{conf}}$ remains as yet intractable in most cases. This is the reason for which we consider *discretized* versions of these models. The model we will now introduce, although representing a mere caricature of the corresponding continuum model, is based on an interaction embodying Van der Waals' two main assumptions: *short-range repulsion* and *long-range attraction*.

The Lattice Gas.


The lattice gas is obtained by first ignoring the momenta (for the reasons explained above) and assuming that the particles' positions are restricted to a discrete subset of \mathbb{R}^d . In general, this subset is taken to be the **d -dimensional cubic lattice**

$$\mathbb{Z}^d \stackrel{\text{def}}{=} \{i = (i_1, \dots, i_d) \in \mathbb{R}^d : i_k \in \mathbb{Z} \text{ for each } k \in \{1, \dots, d\}\}.$$

In other words, one imagines that \mathbb{R}^d is composed of small cells and that each cell can accommodate at most one particle. To describe the microstates of the model, we consider a finite region $\Lambda \subset \mathbb{Z}^d$ representing the vessel and one associates an **occupation number** η_i taking values in $\{0, 1\}$ to each cell $i \in \Lambda$: the value 0 means

that the cell is empty, while the value 1 means that it contains a particle. The set of microstates is thus simply

$$\Omega_\Lambda \stackrel{\text{def}}{=} \{0, 1\}^\Lambda.$$



Note that this model automatically includes a short-range repulsion between the particles, since no two particles are allowed to share a same cell. The attractive part of the interaction can then be included into the Hamiltonian: for any $\eta = (\eta_i)_{i \in \Lambda} \in \Omega_\Lambda$,

$$\mathcal{H}(\eta) \stackrel{\text{def}}{=} \sum_{\{i,j\} \subset \Lambda} J(j-i) \eta_i \eta_j,$$

which is completely similar to the Hamiltonian $\mathcal{H}^{\text{conf}}(q_1, \dots, q_N)$ above, the function $J: \mathbb{Z}^d \rightarrow \mathbb{R}$ playing the role of ϕ_{long} (one may assume that $J(j-i)$ depends only on the distance between the cells i and j , but this is not necessary). Note that the contribution of a pair of cells $\{i, j\}$ is zero if they do not both contain a particle.

The number of particles in Λ is given by

$$N_\Lambda(\eta) \stackrel{\text{def}}{=} \sum_{i \in \Lambda} \eta_i.$$

It will be useful to distinguish the partition functions in the various ensembles. The canonical partition function will be denoted

$$\mathbf{Q}_{\Lambda; \beta, N} = \sum_{\eta \in \Omega_{\Lambda; N}} e^{-\beta \mathcal{H}(\eta)},$$

where $\Omega_{\Lambda; N} \stackrel{\text{def}}{=} \{\eta \in \Omega_\Lambda : N_\Lambda(\eta) = N\}$, and the grand canonical one will be denoted

$$\Theta_{\Lambda; \beta, \mu} = \sum_N e^{\beta \mu N} \sum_{\eta \in \Omega_{\Lambda; N}} e^{-\beta \mathcal{H}(\eta)}.$$

Example 1.12. The simplest instance of the lattice gas is obtained by setting $J \equiv 0$, in which case the Hamiltonian is identically 0. Since its particles only interact through short-range repulsion, this model is called the **hard-core lattice gas**. \diamond

1.3 Linking Statistical Mechanics and Thermodynamics

So far, we have introduced the central probability distributions of statistical mechanics. With these definitions, the analysis of specific systems reduces to an application of probability theory and the statistical properties of any observable can

(in principle) be deduced. Nevertheless, if one wishes to establish a link with thermodynamics, then one must identify the objects in statistical mechanics that correspond to the quantities in thermodynamics that are not observables, that is, not functions of the microstate, such as the entropy or the temperature.

This will be done by making certain identifications, making one assumption (Boltzmann's Principle) and using certain analogies with thermodynamics. The real justification that these identifications are meaningful lies in the fact that the properties derived for and from these quantities in the rest of the book parallel precisely their analogues in thermodynamics. A reader unconvinced by these analogies can simply take them as motivations for the terminology used in statistical mechanics.

Since our discussion of thermodynamics mostly dealt with the example of a gas, it will be more convenient to discuss the identifications below in a statistical mechanical model of a lattice gas as well. But everything we explain can be extended to general systems.

1.3.1 Boltzmann's Principle and the thermodynamic limit

Consider the lattice gas in a region $\Lambda \subset \mathbb{Z}^d$ with $|\Lambda| = V$, composed of N particles and of total energy U . How should the entropy, function of the macrostate (U, V, N) , be defined? We are looking for an additive function, as in (1.2), associated to an extremum principle that determines equilibrium.

Let us first generalize the energy shell and consider, for each macrostate (U, V, N) , the set of all **microstates compatible with** (U, V, N) :

$$\Omega_{\Lambda; U, N} \stackrel{\text{def}}{=} \{ \eta \in \Omega_{\Lambda} : \mathcal{H}(\eta) = U, N_{\Lambda}(\eta) = N \}.$$


Definition 1.13. The **Boltzmann entropy** associated to a system of N particles in Λ with total energy U is defined by¹²

$$S_{\text{Boltz}}(\Lambda; U, N) \stackrel{\text{def}}{=} \log |\Omega_{\Lambda; U, N}|.$$

We motivate this definition with the following discussion.

Consider two lattice gases at equilibrium, in two separate vessels with equal volumes $|\Lambda_1| = |\Lambda_2| = V$, containing N_1^0 and N_2^0 particles respectively. For simplicity, assume that the particles interact only through hard-core repulsion ($\mathcal{H} \equiv 0$) and that $N \stackrel{\text{def}}{=} N_1^0 + N_2^0$ is even.

Let us now put the two vessels in contact so that they can exchange particles. To reach a new equilibrium state, the N particles are redistributed among the two vessels, using the total volume at their disposal:



¹²Physicists usually write this condition as $S_{\text{Boltz}}(\Lambda; U, N) \stackrel{\text{def}}{=} k_B \log |\Omega(\Lambda; U, N)|$, where k_B is **Boltzmann's constant**. In this book, we will always assume that the units are chosen so that $k_B = 1$.

According to the postulate of thermostatics, equilibrium is realized once the pair giving the number of particles in each vessel, (N_1, N_2) , maximizes the sum of the entropies of the vessels, under the constraint that $N_1^0 + N_2^0 = N$. Let us see why Boltzmann's definition of entropy is the most natural candidate for the function describing this extremum principle.

From the point of view of statistical mechanics, once the vessels have been put in contact, the whole system is described by $\nu_{\Lambda; N_1+N_2}^{\text{Mic}}$ (we set $\Lambda \stackrel{\text{def}}{=} \Lambda_1 \cup \Lambda_2$), under which the probability of observing N_1 particles in Λ_1 , and thus $N_2 = N - N_1$ in Λ_2 , is given by

$$\frac{|\Omega_{\Lambda_1; N_1}| \cdot |\Omega_{\Lambda_2; N_2}|}{|\Omega_{\Lambda; N}|}. \quad (1.32)$$

We are interested in the pairs (N_1, N_2) that maximize this probability, under the constraint $N_1 + N_2 = N$. In (1.32), only the numerator depends on N_1, N_2 , and $|\Omega_{\Lambda_k; N_k}| = \binom{V}{N_k}$. As can be easily verified (see Exercise 1.9 below),

$$\max_{\substack{N_1, N_2: \\ N_1 + N_2 = N}} \binom{V}{N_1} \binom{V}{N_2} = \binom{V}{\frac{N}{2}} \binom{V}{\frac{N}{2}}, \quad (1.33)$$

meaning that the most probable configuration is the one in which the two vessels have the same number of particles. In terms of the Boltzmann entropy, (1.33) takes the form

$$\max_{\substack{N_1, N_2: \\ N_1 + N_2 = N}} \{S_{\text{Boltz}}(\Lambda_1; N_1) + S_{\text{Boltz}}(\Lambda_2; N_2)\} = S_{\text{Boltz}}(\Lambda_1; \frac{N}{2}) + S_{\text{Boltz}}(\Lambda_2; \frac{N}{2}), \quad (1.34)$$

which, since this is a discrete version of the postulate of thermodynamics, makes S_{Boltz} a natural candidate for the entropy of the system.

Unfortunately, this definition still suffers from one important defect. Namely, if the system is large, although having the same number of particles in each half, $(N_1, N_2) = (\frac{N}{2}, \frac{N}{2})$, is more likely than any other repartition (N'_1, N'_2) , it is nevertheless an event with *small* probability! (Of order $\frac{1}{\sqrt{N}}$, see the exercise below.) Moreover, any other pair (N'_1, N'_2) such that $N'_1 + N'_2 = N$, $|N'_1 - \frac{1}{2}N| \ll N^{1/2}$ and $|N'_2 - \frac{1}{2}N| \ll N^{1/2}$ has essentially the same probability.

Exercise 1.9. Prove (1.33). Then, show the probability of having the same number of particles, $\frac{N}{2}$ in each vessel, is of the order $\frac{1}{\sqrt{N}}$.

What must be considered, in order to have a deterministic behavior, is not the *number* of particles in each vessel but their *densities*. Let therefore

$$\rho_{\Lambda_1} \stackrel{\text{def}}{=} \frac{N_{\Lambda_1}}{V}, \quad \rho_{\Lambda_2} \stackrel{\text{def}}{=} \frac{N_{\Lambda_2}}{V}$$

denote the random variables giving the densities of particles in each of the two halves. The constraint $N_1 + N_2 = N$ translates into

$$\frac{\rho_{\Lambda_1} + \rho_{\Lambda_2}}{2} = \frac{N}{2V} \stackrel{\text{def}}{=} \bar{\rho},$$

which is the overall density of the system.

For a large system, ρ_{Λ_1} and ρ_{Λ_2} are both close to $\bar{\rho}$, but they always undergo microscopic fluctuations around $\bar{\rho}$: the probability of observing a fluctuation of size at least $\epsilon > 0$, $\nu_{\Lambda;N}^{\text{Mic}}(|\rho_{\Lambda_1} - \bar{\rho}| \geq \epsilon)$, is always positive, even though it might be very small.

If we are after a more *macroscopic* statement, of the type “at equilibrium, the densities in the two halves are (exactly) equal to $\bar{\rho}$ ”, then some limiting procedure is necessary, similar to the one used in the Law of Large Numbers.

The natural setting is that of a large system with a fixed density. Let us thus fix $\bar{\rho}$, which is the overall density of the system. We will let the size of the system $|\Lambda| = V$ increases indefinitely, $V \rightarrow \infty$, and also let the total number of particles $N \rightarrow \infty$, in such a way that

$$\frac{N}{2V} \rightarrow \bar{\rho}.$$

This procedure is called the **thermodynamic limit**. One might then expect, in this limit, that the densities in the two subsystems concentrate on $\bar{\rho}$, in the sense that

$$\nu_{\Lambda;N}^{\text{Mic}}(|\rho_{\Lambda_1} - \bar{\rho}| \geq \epsilon) = \nu_{\Lambda;N}^{\text{Mic}}(|\rho_{\Lambda_2} - \bar{\rho}| \geq \epsilon) \rightarrow 0, \quad \text{for all } \epsilon > 0.$$

Let us see how this concentration can be obtained and how it relates to Boltzmann's definition of entropy. Keeping in mind that $N_2 = N - N_1$,

$$\begin{aligned} \nu_{\Lambda;N}^{\text{Mic}}(|\rho_{\Lambda_1} - \bar{\rho}| \geq \epsilon) &= \sum_{\substack{N_1: \\ \left| \frac{N_1}{V} - \bar{\rho} \right| \geq \epsilon}} \frac{|\Omega_{\Lambda_1;N_1}| \cdot |\Omega_{\Lambda_2;N_2}|}{|\Omega_{\Lambda;N}|} \\ &= \sum_{\substack{N_1: \\ \left| \frac{N_1}{V} - \bar{\rho} \right| \geq \epsilon}} \exp(S_{\text{Boltz}}(\Lambda_1; N_1) + S_{\text{Boltz}}(\Lambda_2; N_2) - S_{\text{Boltz}}(\Lambda; N)). \end{aligned} \quad (1.35)$$

It turns out that, in the case of the hard-core gas we are considering, the Boltzmann entropy has a well defined *density* in the thermodynamic limit. Namely, since $|\Omega_{\Lambda;N}| = \binom{2V}{N}$, by a simple use of Stirling's formula (Lemma B.3),

$$\lim \frac{1}{2V} S_{\text{Boltz}}(\Lambda; N) = s_{\text{Boltz}}^{\text{hard}}(\bar{\rho}) \stackrel{\text{def}}{=} -\bar{\rho} \log \bar{\rho} - (1 - \bar{\rho}) \log(1 - \bar{\rho}).$$

We can therefore use the **entropy density** $s_{\text{Boltz}}^{\text{hard}}(\cdot)$ in each of the terms appearing in the exponential of (1.35). Letting $\rho_k \stackrel{\text{def}}{=} \frac{N_k}{V}$ and remembering that $\frac{\rho_1 + \rho_2}{2} = \bar{\rho}$,

$$\nu_{\Lambda;N}^{\text{Mic}}(|\rho_{\Lambda_1} - \bar{\rho}| \geq \epsilon) = e^{o(1)V} \sum_{\substack{N_1: \\ |\rho_1 - \bar{\rho}| \geq \epsilon}} \exp\{s_{\text{Boltz}}^{\text{hard}}(\rho_1) + s_{\text{Boltz}}^{\text{hard}}(\rho_2) - 2s_{\text{Boltz}}^{\text{hard}}(\bar{\rho})\} V,$$


where $o(1)$ tends to 0 in the thermodynamic limit¹³. Now, $s_{\text{Boltz}}^{\text{hard}}$ is concave and so

$$\frac{s_{\text{Boltz}}^{\text{hard}}(\rho_1) + s_{\text{Boltz}}^{\text{hard}}(\rho_2)}{2} \leq s_{\text{Boltz}}^{\text{hard}}(\bar{\rho}).$$

In fact, it is *strictly* concave, which implies that there exists $c(\epsilon) > 0$ such that

$$\inf \left\{ s_{\text{Boltz}}^{\text{hard}}(\bar{\rho}) - \frac{s_{\text{Boltz}}^{\text{hard}}(\rho_1) + s_{\text{Boltz}}^{\text{hard}}(\rho_2)}{2} : \frac{\rho_1 + \rho_2}{2} = \bar{\rho}, |\rho_1 - \bar{\rho}| \geq \epsilon \right\} = c(\epsilon).$$

¹³Strictly speaking, one should treat values of ρ close to 0 and 1 separately. To keep the exposition short, we ignore this minor issue here. It will be addressed in Chapter 4.



Therefore, since the number of terms in the sum is bounded above by V , we conclude that

$$\nu_{\Lambda;U,N}^{\text{Mic}}(|\rho_1 - \bar{\rho}| \geq \epsilon) \leq V e^{-(2c(\epsilon) - o(1))V}.$$

The latter quantity tends to 0 in the limit $V \rightarrow \infty$, for any $\epsilon > 0$. We conclude that the densities in each of the two subsystems indeed concentrate on $\bar{\rho}$ as $V \rightarrow \infty$.

In other words, to make the parallel with the discrete version (1.34), we have proven that in the thermodynamic limit, the densities in the two boxes become both equal to $\bar{\rho}$; these are the *unique* densities that realize the supremum in

$$\sup_{\substack{\rho_1, \rho_2: \\ \frac{\rho_1 + \rho_2}{2} = \bar{\rho}}} \{s_{\text{Boltz}}^{\text{hard}}(\rho_1) + s_{\text{Boltz}}^{\text{hard}}(\rho_2)\} = s_{\text{Boltz}}^{\text{hard}}(\bar{\rho}) + s_{\text{Boltz}}^{\text{hard}}(\bar{\rho}).$$

The above discussion was restricted to the hard-core lattice gas, but it shows that, while the Boltzmann entropy S_{Boltz} does not fully satisfy our desiderata, considering its density in the thermodynamic limit yields a function that correctly describes the equilibrium values of the thermodynamic parameters, as the unique solution of an extremum principle.

Let us then consider a generic situation of a system with macrostate (U, V, N) . To treat models beyond the hard-core lattice gas, the definition of the thermodynamic limit must include a limit $U \rightarrow \infty$ with $U/V \rightarrow u$.

Definition 1.14. Fix u and ρ . Consider the **thermodynamic limit**, $U \rightarrow \infty$, $V \rightarrow \infty$, $N \rightarrow \infty$, in such a way that $\frac{U}{V} \rightarrow u$ and $\frac{N}{V} \rightarrow \rho$, and let Λ be increasing, such that $|\Lambda| = V$. The **Boltzmann entropy density** at energy density u and particle density ρ is defined by the following limit, when it exists:

$$s_{\text{Boltz}}(u, \rho) \stackrel{\text{def}}{=} \lim \frac{1}{V} S_{\text{Boltz}}(\Lambda, U, N).$$

Of course, two nontrivial claims are hidden in the definition of the Boltzmann entropy density: the existence of the limit and the fact that it does not depend on the chosen sequence of sets Λ . This can be proved for a very large class of models, at least for sufficiently regular sequences (cubes would be fine, but much more general shapes can be accommodated). Several statements of this type will be proved in later chapters.

In view of the above discussion, it seems natural to consider the Boltzmann entropy density as the statistical mechanical analogue of the thermodynamic entropy density $s(u, \rho) = \frac{1}{V} S(U, V, N)$.

Boltzmann's Principle. *The thermodynamic entropy density associated to the macrostate (U, V, N) , corresponding to densities $u = \frac{U}{V}$, $\rho = \frac{N}{V}$, will be identified with the Boltzmann entropy density:*

$$s(u, \rho) \leftrightarrow s_{\text{Boltz}}(u, \rho).$$

It is possible, at least for reasonably well-behaved systems, to prove that s_{Boltz} possesses all the properties we established for its thermodynamic counterpart. For this introduction, we will only give some plausibility argument to show that s_{Boltz} is concave in general (similar arguments will be made rigorous in Chapter 4).

Consider again a gas Σ contained in a cubic domain Λ , with parameters $V = |\Lambda|, U, N$. Let $u = \frac{U}{V}, \rho = \frac{N}{V}$. Fix $\alpha \in (0, 1)$ and consider u_1, u_2 and ρ_1, ρ_2 such that

$$u = \alpha u_1 + (1 - \alpha) u_2, \quad \text{and} \quad \rho = \alpha \rho_1 + (1 - \alpha) \rho_2.$$

We think of Σ as being composed of a large number M of subsystems, $\Sigma_1, \dots, \Sigma_M$, each contained in a sub-cube of volume $V' = V/M$. If one can neglect the energy due to the interaction between the subsystems (which is possible if we assume that the latter are still very large), then one way of having energy and particle densities u and ρ in Σ is to have energy and particle densities u_1 and ρ_1 in a fraction α of the subsystems, and energy and particle densities u_2 and ρ_2 in the remaining ones. We then have

$$|\Omega_{\Lambda; U, N}| \geq |\Omega_{\Lambda'; u_1 V', \rho_1 V'}|^{\alpha M} |\Omega_{\Lambda'; u_2 V', \rho_2 V'}|^{(1-\alpha)M},$$

where Λ' denotes a cube of volume V' . Therefore,

$$\frac{1}{V} \log |\Omega_{\Lambda; U, N}| \geq \alpha \frac{1}{V'} \log |\Omega_{\Lambda'; u_1 V', \rho_1 V'}| + (1 - \alpha) \frac{1}{V'} \log |\Omega_{\Lambda'; u_2 V', \rho_2 V'}|.$$

Letting first $M \rightarrow \infty$ and then taking the thermodynamic limit $V' \rightarrow \infty$ yields

$$s_{\text{Boltz}}(\alpha u_1 + (1 - \alpha) u_2, \alpha \rho_1 + (1 - \alpha) \rho_2) \geq \alpha s_{\text{Boltz}}(u_1, \rho_1) + (1 - \alpha) s_{\text{Boltz}}(u_2, \rho_2),$$

as desired.

Assuming that $s_{\text{Boltz}}(u, \rho)$ exists and satisfies the relations we have seen in thermodynamics, and using Boltzmann's principle, we will now motivate the definition of the thermodynamic potentials studied in the canonical and grand canonical ensembles of statistical mechanics, namely the *free energy* and *pressure*.

Canonical ensemble.

Observe first that the canonical partition function at parameter β of a lattice gas with N particles in a vessel of size $|\Lambda| = V$ can be rewritten as

$$\mathcal{Q}_{\Lambda; \beta, N} = \sum_{U \in \mathcal{U}} e^{-\beta U} |\Omega_{\Lambda; U, N}| = \sum_{U \in \mathcal{U}} e^{-\beta U + S_{\text{Boltz}}(\Lambda; U, N)} = e^{o(1)V} \sum_{U \in \mathcal{U}} e^{-(\beta u - s_{\text{Boltz}}(u, \rho))V},$$

where we introduced $u \stackrel{\text{def}}{=} U/V, \rho \stackrel{\text{def}}{=} N/V$ and used the definition of the Boltzmann entropy density:

$$S_{\text{Boltz}}(\Lambda; U, N) = (s_{\text{Boltz}}(u, \rho) + o(1))V.$$

One can then bound the sum from above and from below by keeping only its largest term

$$e^{-V \inf_u \{\beta u - s_{\text{Boltz}}(u, \rho)\}} \leq \sum_{U \in \mathcal{U}} e^{-(\beta u - s_{\text{Boltz}}(u, \rho))V} \leq |\mathcal{U}| e^{-V \inf_u \{\beta u - s_{\text{Boltz}}(u, \rho)\}}. \quad (1.36)$$

To be more specific, let us assume that the lattice gas has **finite-range** interactions, meaning that $J(j-i) = 0$ as soon as $\|j-i\|_2$ is larger than some fixed constant. In this case, $|\mathcal{U}|$ is bounded by a constant times V , which gives $\frac{1}{V} \log |\mathcal{U}| \rightarrow 0$. Therefore, in the thermodynamic limit ($N, V \rightarrow \infty, N/V \rightarrow \rho$), we obtain

$$\lim \frac{1}{V} \log \mathbf{Q}_{\Lambda; \beta, N} = - \inf_u \{ \beta u - s_{\text{Boltz}}(u, \rho) \} \stackrel{\text{def}}{=} -\hat{f}(\beta, \rho). \quad (1.37)$$

In order to make the desired identifications with thermodynamics, we first argue that, *under the canonical Gibbs distribution, the energy density of the system concentrates, in the thermodynamic limit, on the value $\bar{u} = \bar{u}(\beta, \rho)$ minimizing $\beta u - s_{\text{Boltz}}(u, \rho)$* . (Note the similarity between the argument below and the discussion in Section 1.1.5.) Indeed, for $\epsilon > 0$, let

$$\mathcal{U}_\epsilon \stackrel{\text{def}}{=} \{ U \in \mathcal{U} : |\beta u - s_{\text{Boltz}}(u, \rho) - \beta \bar{u} - s_{\text{Boltz}}(\bar{u}, \rho)| \leq \epsilon \}$$

denote the set of values of the energy for which $\beta u - s_{\text{Boltz}}(u, \rho)$ differs from its minimum value by at most ϵ . Then, repeating (1.36) for the sum over $U \in \mathcal{U} \setminus \mathcal{U}_\epsilon$,

$$\begin{aligned} \mu_{\Lambda; \beta, N} \left(\frac{\mathcal{H}}{V} \notin \mathcal{U}_\epsilon \right) &= \frac{\sum_{U \in \mathcal{U} \setminus \mathcal{U}_\epsilon} e^{-\beta U + S_{\text{Boltz}}(\Lambda; U, N)}}{\mathbf{Q}_{\Lambda; \beta, N}} \\ &\leq \frac{|\mathcal{U}| e^{-V \{ \beta \bar{u} - s_{\text{Boltz}}(\bar{u}, \rho) + \epsilon \}}}{e^{-V \{ \beta \bar{u} - s_{\text{Boltz}}(\bar{u}, \rho) \}}} e^{o(1)V} \leq e^{-(\epsilon - o(1))V}, \end{aligned}$$

which tends to 0 as $V \rightarrow \infty$, for any $\epsilon > 0$.

Up to now, the arguments were purely probabilistic. We are now going to use Boltzmann's Principle in order to relate relevant quantities to their thermodynamic counterparts.

First, since the energy density concentrates on the value \bar{u} , it is natural to identify the latter with the thermodynamic equilibrium energy density. Now, note that \bar{u} is also the value such that (assuming differentiability)

$$\beta = \frac{\partial s_{\text{Boltz}}}{\partial u}(\bar{u}, \rho). \quad (1.38)$$

Using Boltzmann's Principle to identify s_{Boltz} with the thermodynamic entropy density s and comparing (1.38) with the right-hand side of (1.3), we see that the parameter β of the canonical distribution should indeed be interpreted as the **inverse temperature**.

In turn, comparing (1.37) and (1.15), we see that $f(T, \rho) = T \hat{f}(1/T, \rho)$ can be identified with the Helmholtz free energy density. We conclude that, when it exists, the limit

$$-\lim \frac{1}{\beta V} \log \mathbf{Q}_{\Lambda; \beta, N} \quad (1.39)$$

is the relevant thermodynamic potential for the description of the canonical lattice gas. It will simply be called the **free energy**.

Exercise 1.10. Show that $\hat{f}(\beta, \rho)$ defined in (1.37) is concave in β , in agreement with the result of Exercise 1.3 obtained in the thermodynamical context.

Exercise 1.11. Let $\nu_{\Lambda;U,N}^{\text{Mic}}$ denote the microcanonical distribution in Λ associated to the parameters U, N . Show that its Shannon entropy coincides with the Boltzmann entropy:

$$S_{\text{Sh}}(\nu_{\Lambda;U,N}^{\text{Mic}}) = S_{\text{Boltz}}(\Lambda; U, N).$$

Therefore, Boltzmann's Principle actually identifies the Shannon entropy density associated to the microcanonical distribution with the thermodynamic entropy density.

Exercise 1.12. Let $\mu_{\Lambda;\beta(U),N}$ be the canonical Gibbs distribution associated to the parameter $\beta(U)$ for which (1.31) holds. Show that

$$\frac{\partial S_{\text{Sh}}(\mu_{\Lambda;\beta(U),N})}{\partial U} = \beta(U).$$

Therefore, identifying, in analogy with what is done in Exercise 1.11, the Shannon entropy of the canonical distribution with the thermodynamical entropy yields an alternative motivation to identify the parameter β with the inverse temperature.

Compute also $U - T(U)S_{\text{Sh}}(\mu_{\beta(U)})$ and verify that it coincides with the definition of free energy in the canonical ensemble.

Grand canonical ensemble.

We can do the same type of argument for the grand canonical partition function at parameter β, μ of a gas in a region of volume $|\Lambda| = V$:

$$\begin{aligned} \Theta_{\Lambda;\beta,\mu} &= \sum_{N,U} e^{\beta\mu N - \beta U} |\Omega_{\Lambda;U,N}| \\ &= \sum_{N,U} e^{\beta\mu N - \beta U + S_{\text{Boltz}}(\Lambda;U,N)} = e^{o(1)V} \sum_{\rho,u} e^{-\{\beta u - \beta\mu\rho - s(u,\rho)\}V}. \end{aligned}$$

Arguing as before, we conclude that, in the thermodynamic limit $V \rightarrow \infty$,

$$\lim \frac{1}{V} \log \Theta_{\Lambda;\beta,\mu} = -\inf_{u,\rho} \{\beta u - \beta\mu\rho - s(u,\rho)\} \stackrel{\text{def}}{=} -\hat{\phi}_G(\beta, -\mu/T). \quad (1.40)$$

Again, the particle and energy densities concentrate on the values $\bar{u} = \bar{u}(\beta, \mu)$ and $\bar{\rho} = \bar{\rho}(\beta, \mu)$ such that

$$\beta = \frac{\partial s}{\partial u}(\bar{u}, \bar{\rho}), \quad \beta\mu = \frac{\partial s}{\partial \rho}(\bar{u}, \bar{\rho}).$$

In view of (1.3) and (1.6), this allows us to interpret the parameters β and μ of the grand canonical distribution as the **inverse temperature** and the **chemical potential**, respectively. Moreover, comparing (1.40) with (1.20), we see that $\phi_G(T, \mu) \stackrel{\text{def}}{=} T\hat{\phi}_G(1/T, -\mu/T)$ can be identified with the density of the grand potential, which, by (1.22), corresponds to minus the pressure $p(T, \mu)$ of the model.

We thus see that, when the limit exists,

$$\lim \frac{1}{\beta V} \log \Theta_{\Lambda;\beta,\mu} \quad (1.41)$$

is the relevant thermodynamical potential for the description of the grand canonical ensemble; it will simply be called the **pressure**.

In later chapters, we will see precise (and rigorous) versions of the kind of argument used above.

1.3.2 Deriving the equation of state of the ideal gas

Computing the free energy or the pressure of a given model is not trivial, in general, and will be done for several interesting cases in later chapters. Nevertheless, if we consider the simplest possible case, the hard-core lattice gas, then some explicit computations can be done.

Fix $\beta > 0$ and $\mu \in \mathbb{R}$. Since $\mathcal{H} \equiv 0$, the grand canonical partition function is easily computed:

$$\Theta_{\Lambda;\beta,\mu} = \sum_{N=0}^V \binom{V}{N} e^{\beta\mu N} = (1 + e^{\beta\mu})^V.$$

It follows from (1.41) that, in the thermodynamic limit $V \rightarrow \infty$, the pressure is given by

$$p(T, \mu) = T \log(1 + e^{\beta\mu}). \quad (1.42)$$

The average number of particles is given by

$$\langle N_{\Lambda} \rangle_{\nu_{\Lambda;\beta,\mu}} = \frac{1}{\beta} \frac{\partial \log \Theta_{\Lambda;\beta,\mu}}{\partial \mu} = \frac{e^{\beta\mu}}{1 + e^{\beta\mu}} V.$$

In particular, in the thermodynamic limit $V \rightarrow \infty$,

$$\rho(\beta, \mu) \stackrel{\text{def}}{=} \lim_{V \rightarrow \infty} \left\langle \frac{N_{\Lambda}}{V} \right\rangle_{\nu_{\Lambda;\beta,\mu}} = \frac{e^{\beta\mu}}{1 + e^{\beta\mu}}.$$


Using this in (1.42), we obtain the equation of the isotherms:

$$p = -T \log(1 - \rho). \quad (1.43)$$

For a diluted gas, $\rho \ll 1$, a Taylor expansion gives $p = \rho T + O(\rho^2)$, which in terms of the specific volume $v = 1/\rho$ becomes

$$pv = T + O\left(\frac{1}{v}\right),$$


and we recover the equation of state for an ideal gas, see (1.14).



The reason we observe deviations from the ideal gas law at higher densities (small v) is due to the repulsive interaction between the particles, which comes from the fact that there can be at most one of them in each cell. Note, also, that we do not find a coexistence plateau in this model, see (1.43). This is due to the absence of attractive interaction between the particles. More general situations will be considered in Chapter 4.

1.3.3 The basic structure

The structure provided by the formalism presented so far, applied to the lattice gas, can be summarized as follows:



Many models will be introduced and analyzed in the rest of the book, based on this structure. These models will be used to study various aspects of equilibrium statistical mechanics: macroscopic features, first-order phase transitions, fluctuations, equivalence of the different ensembles, etc. Since it is mathematically simpler and physically often more relevant, we will mostly work at fixed temperature rather than fixed energy; the basic object for us will thus be the canonical and grand canonical Gibbs distributions.

Let us now move on to another type of phenomenon which can be studied in this formalism and which will be one of the main concerns of later chapters.

1.4 Magnetic systems


In this section, we describe another important class of macroscopic systems encountered in this book: *magnets*. We will discuss the two main types of behavior magnets can present, *paramagnetism* and *ferromagnetism*, and introduce one of the main models used for their description.

1.4.1 Phenomenology: Paramagnets vs. Ferromagnets


Consider a sample of some material whose atoms are arranged in a regular crystalline structure ^[5]. We suppose that each of these atoms carries a magnetic moment (picture a small magnet attached to each atom) called its *spin*. We assume that each spin has the tendency of aligning with its neighbors and with an external magnetic field.

If the magnetic field points in a fixed direction, the spins are globally ordered: they tend to align with the field and thus all point roughly in the same direction. If we then slowly decrease the intensity of the external field to zero, two behaviors are possible.


Paramagnetic behaviour. In the first scenario, the global order is progressively lost as the field decreases and, when the latter reaches zero, the spins' global order is lost. Such a behavior is called *paramagnetism*:




This phenomenon can be measured quantitatively by introducing the **magnetization**, which is the average of the spins, projected along the direction of the magnetic field. For a paramagnet, when the field decreases from a positive value to zero (maintaining the direction fixed), or similarly if it increases from a negative value to zero, the magnetization tends to zero:



Ferromagnetic behaviour. But another scenario is possible: as the external field decreases, the global order decreases, but the local interactions among the spins are strong enough for the material to maintain a globally magnetized state even after the external field has vanished. Such a behavior is called *ferromagnetism*:



A ferromagnet thus exhibits *spontaneous magnetization*, that is, global ordering of the spins even in the absence of an external magnetic field. The value of the spontaneous magnetization, $\pm m^*$, depends on whether the external field approached zero from positive or negative values:



Observe that as the field goes through zero, the magnetization suffers a discontinuity: it jumps from a strictly positive to a strictly negative value. This corresponds to a *first-order phase transition*.

Using the process described above, one can in principle prepare a ferromagnetic material with a spontaneous magnetization pointing in an arbitrary direction, by simply applying a magnetic field in that direction and slowly decreasing its intensity to zero.

The distinction between these two types of magnetic behavior was first made by Pierre Curie in 1895 and initiated the modern theory of magnetism. Among other important results, Curie observed that a same material can present both types of behavior, depending on the temperature: its behavior can suddenly change from ferromagnetic to paramagnetic once its temperature is raised above a well-defined, substance-specific temperature, now known as the *Curie temperature*.


1.4.2 A simple model for a magnet: the Ising model

The Ising model was introduced by Wilhelm Lenz in 1920 [221], in view of obtaining a theoretical understanding of the phase transition from ferromagnetic to paramagnetic behavior described above. The name “Ising model” (sometimes, but much less frequently, more aptly called the Lenz–Ising model, as suggested by Ising himself) was coined in a famous paper by Rudolph Peierls [266] in reference to Ernst Ising’s 1925 PhD thesis [175], which was undertaken under Lenz’s supervision and devoted to the one-dimensional version of the model.

A major concern and a much debated issue, in the theoretical physics community at the beginning of the 20th century, was to determine whether phase transitions could be described within the framework of statistical mechanics, still a young theory at that time. ^[6]

This question was settled using the Ising model. The latter is indeed *the first system of locally interacting units for which it was possible to prove the existence of a phase transition*. This proof was given in the above-mentioned paper by Peierls in 1936, using an argument that would later become a central tool in statistical mechanics. ^[7]

Its simplicity and the richness of its behavior have turned the Ising model into a preferred laboratory to test new ideas and methods in statistical mechanics. It is nowadays, undoubtedly, the most famous model in this field, and has been the subject of thousands of research papers. Moreover, through its numerous interpretations in physics as well as in many other fields, it has been used to describe qualitatively, and sometimes quantitatively, a great variety of situations. ^[8]

Figure 1.7: A spin configuration $\omega \in \Omega_{B(4)}$.

We model the regular crystalline structure corresponding to the positions of the atoms of our magnet by a finite, non-oriented graph $G = (\Lambda, \mathcal{E})$, whose set of vertices Λ is a subset of \mathbb{Z}^d . A typical example, often used in this book, is the **box of radius n** :

$$B(n) \stackrel{\text{def}}{=} \{-n, \dots, n\}^d$$

For example, $B(4)$ is represented on Figure 1.7. The edges of the graph will most often be between **nearest neighbors**, that is, pairs of vertices i, j with $\|j - i\|_1 = 1$, where the norm is defined by $\|i\|_1 \stackrel{\text{def}}{=} \sum_{k=1}^d |i_k|$. We write $i \sim j$ to indicate that i and j are nearest neighbors. So, the set of edges in the box $B(n)$ is $\{i, j\} \subset B(n) : i \sim j\}$, as depicted in Figure 1.7.

The Ising model is defined by first assuming that a spin is located at each vertex of the graph $G = (\Lambda, \mathcal{E})$. One major simplification is the assumption that, unlike the pictures of Section 1.4.1, the spins are restricted to one particular direction, pointing either “up” or “down”; the corresponding two states are traditionally denoted by $+1$ (“up”) and -1 (“down”). It follows that, to describe a microstate, a variable ω_i taking two possible values ± 1 is associated to each vertex $i \in \Lambda$; this variable will be called the **spin** at i .

A microstate of the system, usually called a **configuration**, is thus an element $\omega \in \Omega_\Lambda$, where

$$\Omega_\Lambda \stackrel{\text{def}}{=} \{-1, 1\}^\Lambda.$$

The microscopic interactions among the spins are defined in such a way that:

1. *There is only interaction between pairs of spins located at neighboring vertices.* That is, it is assumed that the spins at two distinct vertices $i, j \in \Lambda$ interact if and only if the pair $\{i, j\}$ is an edge of the graph.
2. *The interaction favors agreement of spin values.* In the most common instance of the model, to which we restrict ourselves here, this is done in the simplest possible way: a pair of spins at the endpoints i and j of an edge *decreases* the overall energy of the configuration if they agree ($\omega_i = \omega_j$) and

increases if they differ; more precisely, the spins at the endpoints of the edge $\{i, j\}$ contribute to the total energy by an amount

$$-\omega_i \omega_j.$$

Therefore, configurations in which most pairs of neighbors are aligned have smaller energy.

3. *Spins align with the external magnetic field.* Assume that a constant external magnetic field of intensity $h \in \mathbb{R}$ (oriented along the same direction as the spins) acts on the system. Its interaction with the spin at i contributes to the total energy by an amount

$$-h\omega_i.$$

That is, when the magnetic field is positive, the configurations with most of their spins equal to $+1$ have smaller energy.

The **energy** of a configuration ω is obtained by summing the interactions over all pairs and by adding the interaction of each spin with the external magnetic field. This leads to the **Hamiltonian of the Ising model**:

$$\mathcal{H}_{\Lambda;h}(\omega) \stackrel{\text{def}}{=} - \sum_{\substack{i,j \in \Lambda \\ i \sim j}} \omega_i \omega_j - h \sum_{i \in \Lambda} \omega_i, \quad \omega \in \Omega_{\Lambda}. \quad (1.44)$$

Since it favors *local alignment* of the spins, the Hamiltonian of the model is said to be **ferromagnetic**. (Note that this terminology does not necessarily imply that the models *behaves* like a ferromagnet.)

The Gibbs distribution is denoted by

$$\mu_{\Lambda;\beta,h}(\omega) = \frac{e^{-\beta \mathcal{H}_{\Lambda;h}(\omega)}}{\mathbf{Z}_{\Lambda;\beta,h}},$$

where $\mathbf{Z}_{\Lambda;\beta,h}$ is the associated partition function. The expectation of an observable $f : \Omega_{\Lambda} \rightarrow \mathbb{R}$ under $\mu_{\Lambda;\beta,h}$ is denoted $\langle f \rangle_{\Lambda;\beta,h}$.

An important observation is that, in the absence of a magnetic field (that is, when $h = 0$), even though local spin alignment is favored by the Hamiltonian, neither of the orientations ($+1$ or -1) is favored globally. Namely, if $-\omega$ denotes the spin-flipped configuration in which $(-\omega)_i \stackrel{\text{def}}{=} -\omega_i$, then $\mathcal{H}_{\Lambda;0}(-\omega) = \mathcal{H}_{\Lambda;0}(\omega)$; this implies that

$$\mu_{\Lambda;\beta,0}(-\omega) = \mu_{\Lambda;\beta,0}(\omega).$$

The model is then said to be **invariant under global spin flip**. When $h \neq 0$, this symmetry no longer holds.

1.4.3 Thermodynamic behavior

Our goal is to study the Ising model in a large region Λ and to eventually take the thermodynamic limit, for instance taking $\Lambda = B(n)$ and letting $n \rightarrow \infty$.

To simplify the discussion, we will first consider the model in the absence of a magnetic field: $h = 0$. A natural question, which will be a central theme in this book, is: under which circumstances does the ferromagnetic nature of the model, whose tendency is to align the spins locally, induce order also at the global/macroscopic

scale? To make this question more precise, we need suitable ways to quantify global order. One natural such quantity is the **total magnetization**

$$M_{\Lambda}(\omega) \stackrel{\text{def}}{=} \sum_{i \in \Lambda} \omega_i.$$

Then, the **magnetization density**

$$\frac{M_{\Lambda}(\omega)}{|\Lambda|} \in [-1, 1]$$

equals the difference between the fractions of spins that take the value +1 and −1 respectively; it therefore provides some information on the balance between the two spin values in the whole system.

As we already pointed out, the Gibbs distribution is invariant under a global spin-flip when $h = 0$. As a consequence, the *average* magnetization is zero at all temperatures:

Exercise 1.13. Show that

$$\langle M_{\Lambda} \rangle_{\Lambda; \beta, 0} = 0. \quad (1.45)$$

The interpretation of (1.45) is that, *on average*, the densities of + and − spins are equal. However, as we will see below, this does not necessarily mean that the densities of two species of spins are equal in typical configurations of the model. As a first natural step, let us study the fluctuations of M_{Λ} around this average value. Since the spins are dependent, this is a subtle question.

To approach the problem of understanding the dependence on the temperature, we will first study the fluctuations of M_{Λ} in two limiting situations, namely that of **infinite temperature** ($\beta = 1/T \downarrow 0$) and **zero temperature** ($\beta = 1/T \uparrow \infty$). Although these two cases are essentially trivial from a mathematical point of view, they will already provide some hints as to what might happen at other values of the temperature, in the infinite-volume Ising model. For the sake of concreteness, we take $\Lambda = B(n)$.

Infinite temperature. Consider the model on $B(n)$ (with a fixed n). In the limit $\beta \downarrow 0$, the Gibbs distribution converges to the uniform distribution on $\Omega_{B(n)}$: for each $\omega \in \Omega_{B(n)}$,

$$\lim_{\beta \downarrow 0} \mu_{B(n); \beta, 0}(\omega) = \mu_{B(n); 0, 0}(\omega) = \frac{1}{|\Omega_{B(n)}|}. \quad (1.46)$$

Therefore, after $\beta \downarrow 0$, $M_{B(n)}$ is a sum of independent and identically distributed random variables. Its behavior in regions of increasing sizes can thus be described using the classical limit theorems of Probability Theory. For instance, the Law of Large Numbers implies that, for all $\epsilon > 0$,

$$\mu_{B(n); 0, 0} \left(\frac{M_{B(n)}}{|B(n)|} \notin [-\epsilon, \epsilon] \right) \longrightarrow 0 \quad \text{as } n \rightarrow \infty. \quad (1.47)$$

Looking at a finer scale, the Central Limit Theorem states that, for all $a < b$,

$$\mu_{B(n); 0, 0} \left(\frac{a}{\sqrt{|B(n)|}} \leq \frac{M_{B(n)}}{|B(n)|} \leq \frac{b}{\sqrt{|B(n)|}} \right) \longrightarrow \frac{1}{\sqrt{2\pi}} \int_a^b e^{-x^2/2} dx. \quad (1.48)$$

Zero temperature. In the opposite regime, in which $\beta \uparrow \infty$ in a fixed box $B(n)$, the distribution $\mu_{B(n);\beta,0}$ concentrates on those configurations that minimize the Hamiltonian, the so-called **ground states**. It is easy to check that the Ising model in $B(n)$ has exactly two ground states: the constant configurations $\eta^+, \eta^- \in \Omega_{B(n)}$, defined by

$$\eta_i^+ \stackrel{\text{def}}{=} +1 \quad \forall i \in B(n), \quad \eta_i^- \stackrel{\text{def}}{=} -1 \quad \forall i \in B(n).$$

For any configuration ω different from η^+ and η^- , there exists at least one pair $\{i, j\}$ of nearest-neighbors in $B(n)$ such that $\omega_i \neq \omega_j$. Therefore,

$$\mathcal{H}_{B(n)}(\omega) - \mathcal{H}_{B(n)}(\eta^\pm) = \sum_{\substack{i,j \in B(n) \\ i \sim j}} (1 - \omega_i \omega_j) \geq 2. \quad (1.49)$$

Consequently,

$$\frac{\mu_{B(n);\beta,0}(\omega)}{\mu_{B(n);\beta,0}(\eta^\pm)} = \frac{e^{-\beta \mathcal{H}_{B(n)}(\omega)}}{e^{-\beta \mathcal{H}_{B(n)}(\eta^\pm)}} \leq e^{-2\beta} \rightarrow 0 \quad \text{as } \beta \uparrow \infty.$$

Since $\mu_{B(n);\beta,0}(\eta^-) = \mu_{B(n);\beta,0}(\eta^+)$, we thus get

$$\lim_{\beta \uparrow \infty} \mu_{B(n);\beta,0}(\omega) = \begin{cases} \frac{1}{2} & \text{if } \omega \in \{\eta^+, \eta^-\}, \\ 0 & \text{otherwise,} \end{cases} \quad (1.50)$$

which means that, in the limit of very low temperatures, the Gibbs distribution “freezes” the system in either of the ground states.

The two very different behaviors observed above in the limits $\beta \uparrow \infty$ and $\beta \downarrow 0$ suggest two possible scenarios for the high- and low-temperature behavior of the Ising model in a large box $B(n)$:

1. When β is small (high temperature), the global magnetization density is close to zero: with high probability,

$$\frac{M_{B(n)}}{|B(n)|} \cong 0.$$

In this scenario, in a typical configuration, the fractions of + and – spins are essentially equal.

2. When β is large (low temperature), $\mu_{B(n);\beta,0}$ concentrates on configurations that mostly coincide with the ground states η^+, η^- . In particular, with high probability,

$$\text{either } \frac{M_{B(n)}}{|B(n)|} \cong +1, \quad \text{or } \frac{M_{B(n)}}{|B(n)|} \cong -1.$$

In this scenario, **spontaneous magnetization/global order** is observed, since a majority of spins has the same sign. Observe that the Law of Large Numbers would not hold in such a regime. Namely, each spin has an average value equal to zero and, nevertheless, the observation of the system as a whole shows that $|B(n)|^{-1} \sum_{i \in B(n)} \omega_i$ is not close to 0. The symmetry under a global spin flip is **spontaneously broken**, in the sense that typical configurations favor one of the two types of spins, even though the Gibbs distribution is completely neutral with respect to both species of spins.

The main problem is to determine which of these behaviors (if any) gives the correct description of the system for intermediate values $0 < \beta < \infty$.

From the physical point of view, the question we will be most interested in is to determine whether the global alignment of the spins observed at $\beta = \infty$ survives, in arbitrarily large systems, for large but finite values of β . This is a delicate question, since the argument given above actually consisted in fixing n and observing that the ground states were dominant when $\beta \uparrow \infty$. But the true limiting procedure we are interested in is to take the thermodynamic limit at fixed temperature, that is to fix β (possibly very large) and *then* let $n \rightarrow \infty$. It turns out that, in this limit, the ground states are in fact very *unlikely* to be observed. Indeed, let us denote by $\Omega_{B(n)}^k$ the set of configurations coinciding everywhere with either η^+ or η^- , except at exactly k vertices where the spins disagree with the ground state. Such local deformations away from a ground-state are often called **excitations**. Then, for each $\omega \in \Omega_{B(n)}^k$,

$$\mathcal{H}_{B(n);0}(\omega) - \mathcal{H}_{B(n);0}(\eta^\pm) \leq 4dk.$$

(This bound is saturated when none of these k misaligned spins are located at neighboring vertices and none is located along the boundary of $B(n)$.) Observe that $|\Omega_{B(n)}^k| = \binom{|B(n)|}{k}$ and that k is always at most equal to $|B(n)|/2$. This means that, for any $k \geq 1$,


$$\begin{aligned} \frac{\mu_{B(n);\beta,0}(\Omega_{B(n)}^k)}{\mu_{B(n);\beta,0}(\eta^\pm)} &= \sum_{\omega \in \Omega_{B(n)}^k} e^{-\beta(\mathcal{H}_{B(n);0}(\omega) - \mathcal{H}_{B(n);0}(\eta^\pm))} \\ &\geq \binom{|B(n)|}{k} e^{-4d\beta k} \geq \frac{1}{k!} \left(\frac{1}{2}|B(n)|e^{-4d\beta}\right)^k \gg 1, \end{aligned}$$

for all n large enough (at fixed β). In other words, even at very low temperature, it is always much more likely to have misaligned spins in large regions. This discussion shows that there are two competing aspects when analyzing typical configurations under a Gibbs distribution at low temperature. On the one hand, configurations with low energy are favored, since the latter have a larger individual probability; this is the *energy* part. On the other hand, the number of configurations with a given number of excitations grows fast with the size of the system and rapidly outnumbers the small number of ground states; this is the *entropy* part. This *competition between energy and entropy* is at the heart of many phenomena described by equilibrium statistical mechanics (Note that it can already be witnessed in (1.37)), in particular in methods to prove the existence of a phase transition.

These questions will be investigated in detail in Chapter 3. As we will see, the *dimension* of the underlying graph \mathbb{Z}^d will play a central role in the analysis. In the next section, we discuss this dependence on d on the basis of numerical simulations.

Behavior on \mathbb{Z}^d

The one-dimensional model. The following figure shows simulations of typical configurations of the one-dimensional Ising model on $B(50)$, for increasing values of the inverse temperature β (at $h = 0$). For the sake of clarity, +, resp. −, spins are represented by black, resp. white, dots:



As the value of β increases, we see that spins tend to agree on ever larger regions; locally, a configuration looks like either of the ground states η^+ , η^- . Increasing β even more would yield, with high probability, a configuration with all spins equal. Nevertheless, for any value of β , one observes that taking the system's size sufficiently large, regions of + and – spins even out and the global magnetization is always zero on the macroscopic scale.

As seen before, global order can be conveniently quantified using $\frac{M_{B(n)}}{|B(n)|}$. Since the latter has an expectation value equal to zero (Exercise 1.13), we consider the expectation of its absolute value.




Figure 1.8: The expected value of the absolute value of the magnetization density, as a function of $p = 1 - e^{-2\beta}$, for the one-dimensional Ising model, tending to zero as $n \rightarrow \infty$ for all $p \in [0, 1)$.

This reflects the fact that, in $d = 1$, the Ising model exhibits paramagnetic behavior at all positive temperatures; see the discussion below.

The model in dimensions $d \geq 2$. In contrast to its one-dimensional version, the Ising model in higher dimensions exhibits ferromagnetic and paramagnetic behaviors, as the temperature crosses a critical value, similarly to what Curie observed in real magnets.

The phase transition is characterized by two distinct regimes (low and high temperatures), in which the large-scale behavior of the system presents important differences which become sharper as the size of the system increases. A few simulations will reveal these behaviors in $d = 2$. Consider first the model without a magnetic field ($h = 0$), in a square box $B(n)$. A few typical configurations for $n = 100$ are shown in Figure 1.9, for various values of the inverse temperature $\beta \geq 0$. For convenience, instead of varying β , we vary $p = p(\beta) \stackrel{\text{def}}{=} 1 - e^{-2\beta}$, which has the advantage of taking values in $[0, 1)$. Values of p near 0 thus correspond to high temperatures, while values of p near 1 correspond to low temperatures.




Figure 1.9: Typical configurations of the two-dimensional Ising model in the box $B(100)$, for different values of $p = 1 - e^{-2\beta}$. Black dots represent $+$ spins, white dots represent $-$ spins. When p is close to 0 (β small), the spins behave roughly as if they were independent and thus appear in equal proportions. When p is close to 1 (β large), a typical configuration is a small perturbation of either of the ground states η^+, η^- ; in particular, it has a magnetization near either $+1$ or -1 .

Figure 1.9 shows that, in contrast to what we observed in the one-dimensional case, the large-scale behavior of the system in two dimensions depends strongly on the temperature. When p is small, the symmetry under global spin flip is preserved in typical configurations and the fractions of $+$ and $-$ spins are essentially equal. When p is close to 1, this symmetry is spontaneously broken: one of the two spin types dominates the other. The simulations suggest that this change of behavior occurs when p is near 0.58, that is, when β is near 0.43. Therefore, the high- and low-temperature behaviors conjectured on the basis of the limiting cases $\beta = 0$ and $\beta \uparrow \infty$ are indeed observed, at least for a system in $B(100)$. In Figure 1.10, $\langle |\frac{M_{B(n)}}{|B(n)|}| \rangle_{B(n); \beta, 0}$ is represented as a function of p , for different values of n .




Figure 1.10: The expected value of the absolute value of the magnetization density, as a function of p , for the two-dimensional Ising model.

The simulations suggest that, as n increases, the sequence of functions $p \mapsto \langle |\frac{M_{B(n)}}{|B(n)|}| \rangle_{B(n); p}$ converges to some limiting curve. This is indeed the case and the

limiting function can be computed explicitly: ^[9]

$$p \mapsto m_p^* \stackrel{\text{def}}{=} \begin{cases} 0 & \text{if } p < p_c, \\ \left[1 - \left(\frac{2(1-p)}{p(2-p)}\right)^4\right]^{1/8} & \text{if } p \geq p_c, \end{cases} \quad (1.51)$$

where

$$p_c \stackrel{\text{def}}{=} \frac{\sqrt{2}}{1 + \sqrt{2}} \cong 0.586$$

is the critical point, to which corresponds the critical inverse temperature

$$\beta_c \stackrel{\text{def}}{=} -\frac{1}{2} \log(1 - p_c) \cong 0.441.$$

The above explicit expression implies, in particular, that the limiting magnetization density is continuous (but not differentiable) at p_c .

Exercise 1.14. Using (1.51), show that the behavior of $m_{p(\beta)}^*$ as $\beta \downarrow \beta_c$ is

$$m_{p(\beta)}^* \sim (\beta - \beta_c)^{1/8},$$

in the sense that $\lim_{\beta \downarrow \beta_c} \frac{\log m_{p(\beta)}^*}{\log(\beta - \beta_c)} = \frac{1}{8}$.

Concerning the dependence of the Ising model on the magnetic field when $h \neq 0$, two main quantities of interest will be considered: the average magnetization density $\langle \frac{M_{B(n)}}{|B(n)|} \rangle_{B(n); \beta, h}$ and the **pressure**

$$\psi_{B(n)}(\beta, h) \stackrel{\text{def}}{=} \frac{1}{\beta |B(n)|} \log Z_{B(n); \beta, h}.$$

Remark 1.15. The reader might wonder why the term *pressure* is used for the magnet. In fact, a one-to-one correspondence can be established between the microstates of the lattice gas and those of the Ising model, by

$$\omega_i \leftrightarrow 2\eta_i - 1.$$

In the Ising model, the number of spins is of course fixed and equal to the size of the region on which it is defined. But the number of + (or −) spins is not fixed and can vary. Therefore, the number of particles in the lattice gas, under the above correspondence, can also vary; it thus corresponds to a grand canonical description, in which the natural thermodynamic potential is the pressure.

The relation between the lattice gas and the Ising model will be fully described, and exploited, in Chapter 4. \diamond

The following **infinite-volume limits** will be considered:

$$m(\beta, h) \stackrel{\text{def}}{=} \lim_{n \rightarrow \infty} \left\langle \frac{M_{B(n)}}{|B(n)|} \right\rangle_{B(n); \beta, h}, \quad \psi(\beta, h) \stackrel{\text{def}}{=} \lim_{n \rightarrow \infty} \psi_{B(n)}(\beta, h).$$

Besides showing that the above limits exist, we will show that, for all $h \neq 0$,

$$\frac{\partial \psi(\beta, h)}{\partial h} = m(\beta, h).$$




Figure 1.11: Dependence of the magnetization density of the two-dimensional Ising model on the magnetic field (obtained from numerical simulations): paramagnetic behavior at high temperature (left), ferromagnetic behavior at low temperature (right).

The map $h \mapsto m(\beta, h)$ is plotted in Figure 1.11 for sub- and supercritical temperatures.

The presence of two different typical behaviors when $h = 0$ and $\beta > \beta_c$ shows how sensitive the system becomes to perturbation by an external field. On the one hand, when $\beta < \beta_c$, a small magnetic field $h > 0$ induces a positive magnetization density which is approximately proportional to h : the response of the system is linear for small h and vanishes when $h \rightarrow 0$. On the other hand, when $\beta > \beta_c$, the introduction of an infinitesimal magnetic field $h > 0$ (resp. $h < 0$) induces a magnetization density close to $+1$ (resp. -1)! This implies that, in contrast to the one-dimensional case, the pressure is not differentiable at $h = 0$: the phase transition is of first order in the magnetic field. Informally, one can say that in the absence of magnetic field, the system “hesitates” between two different behaviors and the introduction of a nonzero, arbitrarily small magnetic field is enough to tip the balance in the corresponding direction.

1.5 Some general remarks

1.5.1 The role of the thermodynamic limit

In Section 1.3.1, the thermodynamic limit has been introduced as a way of establishing a precise link between statistical mechanics and thermodynamics.

Approximating a large system by an infinite one might seem a rather radical step, since real systems are always finite (albeit quite large: a cube of iron with a sidelength of 1cm contains roughly 10^{23} iron atoms).

It turns out that taking a limit of infinite volume is important for other reasons as well.

Deterministic macroscopic behavior. As we have seen, one of the main assumptions in thermodynamics is that, once a small set of thermodynamic quantities has been fixed (say, the pressure and the temperature for an ideal gas), the values of all other macroscopic quantities are in general completely determined. In statistical mechanics, macroscopic observables associated to large finite systems are random variables which are only approximately determined: they still undergo fluctuations, although the latter decrease with the system’s size. As we will see in

Chapter 6, it is only in the thermodynamic limit that all macroscopic observables take on deterministic values (actually, we already saw concentration of some observables: the density of particles in a subsystem for the hard-core lattice gas in the microcanonical ensemble, the energy density in the canonical ensemble, etc.).

The emergence of deterministic behavior should be reminiscent of certain central results in mathematics, such as the Law of Large Numbers or the Ergodic Theorem.

Equivalence of ensembles. On the one hand, we have seen that, in thermodynamics, many choices are possible for the thermodynamic parameters used to describe a system. In the case of the gas, for example, one might use (U, V, N) , (T, V, N) , etc. If the values are suitably chosen, all these approaches lead to the same predictions for the equilibrium properties (except, possibly, at phase transitions).

On the other hand, we have seen that, in statistical mechanics, to each particular set of thermodynamic parameters corresponds an ensemble (microcanonical, canonical, grand canonical), that is, a particular probability distribution on the set of microstates. Obviously, the latter do not coincide for finite systems. It turns out that they indeed become equivalent, in general, but *only* once the thermodynamic limit has been taken: in this limit, the local behavior of the system in different ensembles generally coincide, provided that the thermodynamic parameters are chosen appropriately. In this limit, one says that there is *equivalence of ensembles*. Although equivalence of ensembles will not be described in full generality, we will come back to it in Sections 4.4, 4.7.1 and 6.14.1.

Phase transitions. One additional major reason to consider infinite-volume limits is that it is the only way the formalism of equilibrium statistical mechanics can lead to the singular behaviors thermodynamics associates to phase transitions, such as the coexistence plateau in the liquid-vapor equilibrium, or the discontinuity of the magnetization in a ferromagnet.

Notice that the dependence of a finite system on its parameters is *always* smooth. Consider, for example, the Ising model in $B(n)$. From an algebraic point of view, its partition function can be written (up to an irrelevant smooth prefactor) as a polynomial in the variables $e^{-2\beta}$ and e^{-2h} , with nonnegative (real) coefficients. It follows that the pressure $\psi_{B(n)}(\beta, h)$ is real-analytic for all values of β and h . Of course, the same is true of the magnetization in $B(n)$. An analytic singularity, such as a discontinuity of the magnetization when going from $h > 0$ to $h < 0$, can *only* occur if the thermodynamic limit is taken.

In view of the above, one might wonder how this can be compatible with our everyday experience of various types of phase transitions. The crucial point is that, although finite-volume thermodynamic quantities are always smooth, in very large systems their behavior will be closely approximated by the singular behavior of the corresponding infinite-volume quantities. This was already witnessed in Figure 1.10, in which the finite-volume magnetization of a system in a box as small as $B(250)$ already displays a near-singular behavior. For real macroscopic systems, the behavior will be experimentally indistinguishable from a genuine singularity.

Genuine long-range order vs. apparent long-range order. In our discussion of the one-dimensional Ising model, we mentioned that, for a box of arbitrary size,

with probability close to 1, configurations of this model will be perfectly ordered (all spins being equal) as soon as the temperature is low enough. Nevertheless, we will prove in Chapter 3 that the infinite one-dimensional Ising model is disordered at all positive temperatures. This shows that looking at finite systems might lead us to “wrong” conclusions. Of course, real systems are finite, so that a “real” one-dimensional Ising model would typically display order. But this ordering would be a finite-size effect only. Being able to distinguish between such effects and genuine ordering is essential to obtain a conceptual understanding of these issues (for example, the role of the dimension). An important example will be discussed in Chapter 9.

This short discussion shows that, from the point of view of statistical mechanics, thermodynamics is only an approximate theory dealing (very effectively!) with idealized systems of infinite size. In order to recover predictions from the latter in the framework of the former, it is thus necessary to take the limit of infinite systems. Of course, once a system is well understood in the thermodynamic limit, it can be of great interest to go beyond thermodynamics, by analyzing the finite-volume corrections provided by equilibrium statistical mechanics.

1.5.2 On the role of simple models

The lattice gas and the Ising model share an obvious feature: they are extremely crude models of the systems they are supposed to describe. In the case of the lattice gas, the restriction of the particles to discrete positions is a dramatic simplification and the interaction only keeps very superficial resemblance with the interactions between particles of a real gas. Similarly, in a real magnet, the mechanism responsible for the alignment of two spins is of a purely quantum mechanical nature, which the Ising model simply ignores; moreover, the restriction of the spin to one direction is also not satisfied in most real ferromagnets.

One may thus wonder about the purpose of studying such rough approximations of real systems. This was indeed a major preoccupation of physicists in the early 20th century, who believed that such models might be of interest to mathematicians, but are certainly irrelevant to physics ^[10].

Nevertheless, the point of view on the role of models and on the actual goal of theoretical physics changed substantially at that time. In the realm of statistical mechanics, the mathematical analysis of realistic models of physical systems is in general of such a degree of complexity as to be essentially hopeless. As a consequence, one must renounce to obtain, in general, a complete *quantitatively precise* description of most phenomena (for example, computing precisely the critical temperature of a real magnet). However, it is still possible and just as important to try to understand complex phenomena *at a qualitative level*: What are the mechanisms underlying some particular phenomenon? What are the relevant features of the real system that are responsible for its occurrence? For this, simple models are invaluable ^[11]. We will see all along this book that many subtle phenomena can be reproduced qualitatively in such models, without making any further uncontrolled approximation.

One additional ingredient that played a key role in this change of perspective is the realization that, in the vicinity of a critical point, the behavior of a system becomes essentially independent of its microscopic details, a phenomenon called

universality. Therefore, in such a regime, choosing a simple model as the representative of the very large class of systems (including the real ones) that share the same behavior, allows one to obtain even a *quantitative* understanding of these real systems near the critical point.

Finally, one nice side-effect of considering very simple descriptions is that they often admit many different interpretations. Already in the 1930s, the Ising model was used as a model of a ferromagnet, of a fluid, of a binary alloy and to model an adsorbed monolayer at a surface. The fact that the same model describes qualitatively a wide variety of different systems clarifies the observations made at the time that these very different physical systems exhibit very similar behavior.

1.6 About this book

We wrote this book because we believe that there does not yet exist, in the literature, a book that is self-contained, starts at an elementary level and yet provides a detailed analysis of some of the main ideas, techniques and models of the field.

The target reader we have in mind is an advanced undergraduate or graduate student in mathematics or physics, or anybody with an interest in learning more about some central concepts and results in rigorous statistical mechanics.

Let us list some of the main characteristic features of this book.

- *It is mostly self-contained.* It is only assumed that the reader has basic notions of analysis and probability (only Chapter 6 requires notions from measure theory, and the latter are summarized in Appendix B).
- *It discusses only the equilibrium statistical mechanics of classical lattice systems.* Other aspects of statistical mechanics, not treated here, can be found in the books listed in Section 1.6.2 below.
- *It favors the discussion of specific enlightening examples over generality.* In each chapter, the focus is on a small class of models that we consider to be the best representatives of the topic discussed. These are listed right below in Section 1.6.1.
- *It aims at conveying understanding and not only proofs.* In particular, the proofs given are not always the shortest, most elegant ones, but those we think best help to understand the underlying mechanisms. Moreover, the methods, ideas and concepts introduced in the course of the proof of a statement are often as important as the statement itself.

1.6.1 Contents, chapter by chapter

The first chapters are devoted mainly to the study of models whose spin variables are discrete and take values in a finite set:

- **Chapter 2: The Curie–Weiss model.** Mean field models play a useful role, both from the physical and mathematical point of view, as first approximation to more realistic ones. This chapter gives a detailed account of the Curie–Weiss model, which can be seen as the mean-field version of the Ising model. The advantage is that this model exhibits a phase transition between paramagnetic and ferromagnetic behaviors that can be described with elementary tools.

- **Chapter 3: The Ising model.** As we already said, the Ising model is the simplest “realistic” model which exhibits a non-trivial collective behavior. As such, it has played, and continues to play, a central role in statistical mechanics. This chapter uses it to introduce several very important notions for the first time, such as the notion of infinite-volume state or precise definitions of phase transitions. Then, the complete phase diagram of the model is constructed, in all dimensions, using simple mathematical tools developed from scratch.
- **Chapter 4: Liquid-vapor equilibrium.** Historically, the liquid-vapor equilibrium played a central role in the first theoretical studies of phase transitions. In this chapter, the mathematical description of the lattice gas is exposed in detail, as well as its mean-field and nearest-neighbor (Ising) versions. The mean-field (Kac) limit is also studied in a simple case, providing a rigorous justification of the van der Waals–Maxwell theory of condensation.
- **Chapter 5: Cluster Expansion.** The cluster expansion remains the most important perturbative technique in mathematical statistical mechanics. It is presented in a simple fashion and several applications to the Ising model and the lattice gas are presented. It is also used several times later in the book and plays, in particular, a central role in the implementation of the Pirogov–Sinai theory of Chapter 7.
- **Chapter 6: Infinite-volume Gibbs measures.** In this chapter, we present a probabilistic description of infinite systems of particles at equilibrium, which is known nowadays as the theory of Gibbs measures or the DLR (Dobrushin–Lanford–Ruelle) formalism. This theory is developed from scratch, using the Ising model as a guiding example. Several important aspects, such as Dobrushin’s Uniqueness Theorem, spontaneous symmetry breaking, extremal measures and the extremal decomposition, are also exposed in detail. At the end of the chapter, the variational principle is introduced; the latter is closely linked with the basic concepts of equilibrium thermodynamics.
- **Chapter 7: Pirogov–Sinai Theory.** The Pirogov–Sinai theory is one of the very few general approaches to the rigorous study of first-order phase transitions. It yields, under weak assumptions, a sharp description of such phase transitions in perturbative regimes. This theory is first introduced in a rather general setting and then implemented in detail on one specific three-phase model: the Blume–Capel model.

The last three chapters are devoted to models whose variables are of a continuous nature:

- **Chapter 8: The Gaussian Free Field.** In this chapter, the lattice version of the Gaussian Free Field is analyzed. Several features related to the non-compactness of its single-spin-space are discussed, exploiting the Gaussian nature of the model. The model has a random walk representation, whose recurrence properties are crucial in the study of the behavior of the model in the thermodynamic limit.
- **Chapter 9: Models with continuous symmetry.** An important class of models with a continuous symmetry, including the XY and Heisenberg models,

is studied in this chapter. The emphasis is on the implications of the presence of the continuous symmetry on long-range order in these models in low dimensions. In particular, a strong form of the celebrated Mermin–Wagner theorem is proved in a simple way.

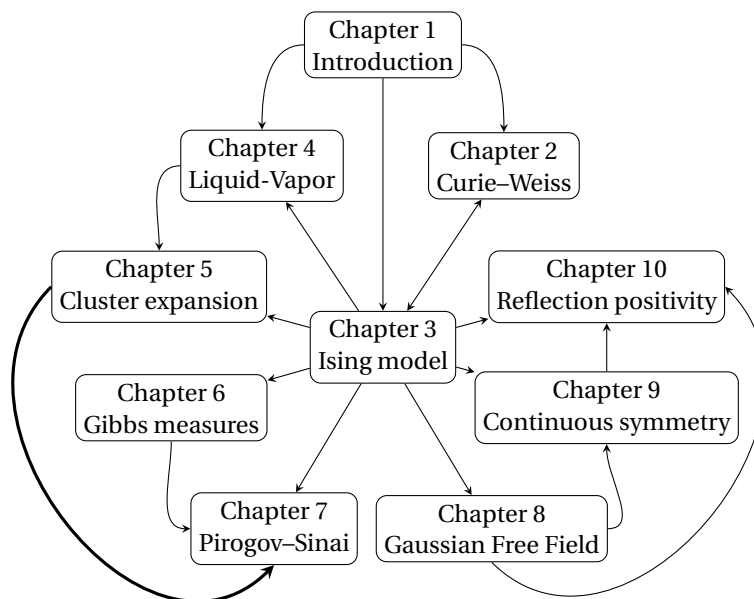
- **Chapter 10: Reflection positivity.** Reflection positivity is another tool that plays a central role in the rigorous study of phase transitions. We first expose it in detail, proving its two central estimates: the infrared bound and the chessboard estimate. We then apply the latter to obtain several results of importance. In particular, we prove the existence of a phase transition in the anisotropic XY model in dimensions $d \geq 2$, as well as in the (isotropic) $O(N)$ model in dimensions $d \geq 3$. Combined with the results of Chapter 9, this provides a detailed description of this type of systems in the thermodynamic limit.

In order to facilitate the reading of the content of each chapter, which can sometimes be pretty technical, the bibliographical references have been placed at the end of the chapter, in a section called **Bibliographical references**. Some chapters also contain a section **Complements and further reading**, in which the interested reader can find further results (usually without proofs) and suggestions for further reading. The goal of these complements is to provide information about some more advanced themes that cannot be treated in detail in the book.

The book ends with three appendices:

- **Appendix A: Notes.** This appendix regroups short **Notes** that are sometimes referred to in the text.
- **Appendix B: Mathematical appendices.** Since we want the book to be mostly self-contained, we introduce various mathematical topics used throughout the book, which might not be part of all undergraduate curricula. For example: elementary properties of convex functions, some aspects of complex analysis, measure theory, conditional expectation, random walks, etc., are briefly introduced, not always in a self-contained manner, often without proofs, but with references to the literature.
- **Appendix C: Solutions to exercises.** Exercises appear in each of the chapters, with various levels of difficulty. Hints or solutions for most of them can be found in this appendix.

We would like to emphasize that Chapter 3 plays a central role, since it introduces several important concepts that are then used constantly in the rest of the book; it should be considered as a priority for a novice reader. The only other true constraint is that Chapter 5 should be read before Chapter 7. Besides that, the chapters can mostly be read independently of each other, and any path following the arrows in the picture below represents a possible way through the book:



Warning: As we said, we have strived to make the book as self-contained as possible, and to assume as little prior knowledge from the reader as possible. Moreover, we have tried to make the chapters as independent from each other as possible, respecting when possible the conventional notations used in the field. This has had some consequence on the final form of the book.

- Together with the fact that we have avoided developing too general a theory, writing essentially independent chapters has had the inevitable consequence of introducing various repetitions: the partition function of a model, for instance, or its Gibbs distribution in finite volume, is always defined in a way suited for the particular analysis used for that model. The same holds for the pressure and other recurring quantities. We therefore warn the reader that corresponding notions might be written slightly differently from one chapter to the other.
- Like in many areas, the notational conventions in statistical mechanics are different in the mathematical and physical communities. For example, probabilists define the free energy as $\frac{1}{V} \log Z$ whereas in physics it is written as $-\frac{1}{\beta V} \log Z$, respecting the structure that appeared in the analogies with thermodynamics.

In this book, we have adopted one convention or the other, depending on the physical relevance of the theory developed in the chapter. Chapter 4, for example, was a natural place where to use the physicists' conventions, since it describes the liquid-vapor equilibrium.

The choices made are always indicated at the beginning of the chapters and we hope that this will not generate too much confusion when jumping from one chapter to another.

1.6.2 The existing literature

This book does not aim at presenting the most recent developments in statistical mechanics. Rather, it presents a set of models and methods, most of which were already known in the 1980s. However, these classical topics form the backbone of this subject and should still be learned by new researchers entering this field. The absence of an introductory text aimed at beginners was deplored by many colleagues and prompted us to write this book.

Statistical mechanics is now such a wide field that it has become impossible to cover more than a fraction of it in one book. In this section, we provide some references to other works covering the various aspects that are either not discussed at all in the present text or only very superficially. Note that we mostly restrict ourselves to books aimed at mathematicians and mathematical physicists.

Books covering similar areas. There exist several books covering some of the areas discussed in the present text. Although the distinction is a bit subjective, we split the list into two, according to what we consider to be the intended audience.

The first set of books is aimed at mathematical physicists. Ruelle wrote the first book [289] on rigorous equilibrium statistical mechanics in 1969. Discussing both classical and quantum systems, in the continuum and on the lattice, this book played a major role in the development of this field. Israel's book [176] provides an in-depth discussion of the variational principle, Gibbs measures as tangent functionals and the role of convexity in equilibrium statistical mechanics. It contains many abstract results found nowhere else in book form. Sinai's book [312] discusses the general theory of Gibbs measures on a lattice with an emphasis on phase transitions and includes perturbative expansions, the Pirogov–Sinai theory, as well as a short introduction to the renormalization group (mostly in the context of hierarchical lattices). Minlos's short book [247] covers similar grounds. The book [227] by Malyshev and Minlos deals with more or less the same topics, but with an approach based systematically on the cluster expansion. Simon's book [308] provides an extensive discussion of the pressure, Gibbs states and their basic properties, and perturbative expansions, both for classical and quantum lattice systems. Presutti's book [279] proposes an alternative approach to several of the topics covered in the present book, but with a strong emphasis on models with Kac interactions. Lavis' book [207] provides a coverage of a wide class of models and techniques.

The second set of books is aimed at probabilists. For this audience, Georgii's remarkable book [134] has become the standard reference for the theory of Gibbs measures; although less accessible than the present text, it is highly recommended to more advanced readers interested in very general results, in particular on the topics covered in our chapter 6. The shorter book [282] by Prum covers similar grounds, but in less generality. Preston's book [278] contains an interesting early account of Gibbs measures, aimed at professional probabilists and limited to rather abstract general results. Kinderman and Snell's very pedagogical monograph [192] includes a clear and intuitive exposition of the phase transition in the two-dimensional Ising model.

Disordered systems. One of the important topics in equilibrium statistical mechanics that is not even touched upon in the present book is disordered systems, in which the Gibbs measures considered depend on additional randomness (such as random interactions). In spite of the activity in this domain, there are only a limited

number of books available for mathematically-inclined readers. The book [254] by Newman discusses short-range models for spin glasses; see also [321] for an introduction. Talagrand's books [325, 326, 327] provide a comprehensive account of mean-field models for spin glasses. Bovier's book [37] starts with an introduction to equilibrium statistical mechanics (including a discussion of the DLR formalism and cluster expansion) and then moves on to discuss both mean-field and lattice models of disordered systems.

Large deviations. Large deviations theory plays an important role in equilibrium statistical mechanics, both at a technical level and at a conceptual level, providing the natural framework to relate thermodynamics and statistical mechanics. This theme will be recurrent in the present text. Nevertheless, we do not develop the general framework here. There are now many books on large deviation theory, with various levels of emphasis on the applications to statistical mechanics, such as the books by Deuschel and Stroock [77], Dembo and Zeitouni [74], den Hollander [75], Ellis [100], Rassoul-Agha and Seppäläinen [283] and Olivieri and Vares [258], as well as the lecture notes by Lanford [205], Föllmer [108] and Pfister [274]. Georgii's book also has a section on this topic [134, Section 15.5]. Let us also mention the more elementary introduction by Touchette [334].

Quantum systems. In this book, we only consider classical lattice spin systems. A discussion of quantum lattice spin systems can be found, for example, in the books by Sewell [300], Simon [308] and Bratteli and Robinson [43, 44].

Historical aspects. Except in some remarks, we do not discuss historical aspects in this book. Good references on the general history of statistical mechanics are the books by Brush [54] and Cercignani [63]; Gallavotti's treatise [130] also provides interesting information on this subject. More specific references to historical aspects of lattice spin systems are given in the articles by Brush [55], Domb [89] and Niss [255, 256, 257].

Percolation. Bernoulli percolation is a central model in probability theory, with strong links to equilibrium statistical mechanics. These links (which we only superficially address in Section 3.10.6) lead to an alternative approach to the analysis of some lattice spin systems (such as the Ising and Potts models), reinterpreting the phase transition as a percolation transition. The percolation model is discussed in detail in the books by Kesten [189], Grimmett [149, 151] and by Bollobás and Riordan [31]. The link with Ising and Potts models is explained in the books by Grimmett [150] and Werner [350], in the review paper [132] by Georgii, Häggström and Maes and in the lecture notes by Duminil-Copin [91].

Thermodynamic formalism. Some core ideas from equilibrium statistical mechanics have been successfully imported into the theory of dynamical systems, where it is usually known as the thermodynamic formalism. An excellent early reference is Bowen's book [40]. Other references are the books by Ruelle [291] and Keller [187], or the lecture notes by Sarig [293].

Stochastic dynamics. An area that is closely related to several problems studied in this book is the analysis of the stochastic dynamics of lattice spin systems. In the latter, one considers Markov chains on set Ω of microscopic configurations, under which the Gibbs distributions are invariant. The book [225] by Liggett and the lecture notes [232] by Martinelli provide good introductions to this topic.

Critical phenomena. This topic is one of the major omissions in this book. See the short discussion and the bibliographical references given in Section 3.10.11.

Exactly solvable models. A discussion of the exact (but not always necessarily rigorous) solutions of various models of statistical mechanics can be found in the books by McCoy and Wu [239], Baxter [17], Palmer [261] or Lavis [207].

Foundations of equilibrium statistical mechanics. There are several books on the foundations of statistical mechanics and its relations to thermodynamics, such as, for example, those by Gallavotti [130], Martin-Löf [230], Khinchin [190] and Sklar [314].

