Gibbsianness and non-Gibbsianness in lattice random fields

Roberto Fernández
Laboratoire de Mathématiques Raphaël Salem
UMR 6085 CNRS-Université de Rouen
Avenue de l'Université, BP 12
F-76801 Saint Étienne du Rouvray, France
roberto.fernandez@univ-rouen.fr

March 15, 2006

Contents

1 Historical remarks and purpose of the course

The notion of Gibbs measure, or Gibbs random field, is the founding stone of mathematical statistical mechanics. Its formalization in the late sixties, due to Dobrushin, Lanford and Ruelle [6, 34], marked the beginning of two decades of intense activity that produced a rather complete theory [55, 18] which has been exploited in many areas of mathematical physics, probability and stochastic processes, as well as for example in dynamical systems.

Despite its diverse applicability, the Gibbsian description was developed specifically to describe *equilibrium* statistical mechanics. Limitations were bound to show up when this framework was transgressed. In fact, the wake-up call came from work within the original equilibrium setting. Indeed, around 1980 Griffiths and Pearce [21, 22, 20] pointed out that some measures obtained as a result of renormalization transformations —a technique developed to study critical points [16, 19]— showed "pathologies" that contradicted Gibbsian intuition. Israel [26] quickly pointed out the cause. These measures lacked the *quasilocality* property which, as we discuss below, is one of the (two) key properties of Gibbsianness. The measures were, thus, *non-Gibbsian*.

After these early examples, almost a decade had to pass before the topic really took off. At this time there appeared a second wave of examples showing that non-Gibbsianness was rather ubiquitous; it was present in spin "contractions" [35, 9], lattice projections [56] and stationary measures of stochastic evolutions [36]. These examples motivated us to write a mini-treatise [64] where we tried to explain the relevant notions and systematize existing examples based on different non-Gibbsianness symptoms —discontinuities or zeroes of the conditional distributions, large deviations that are too large or too small. Next to our article, in the same issue of Journal of Statistical Physics, a much shorter paper by Martinelli and Olivieri [47] initiated, in fact, the next stage of the non-Gibbsianness saga, namely the efforts in the direction of *Gibbsian restoration*.

These efforts have progressed in two complementary directions. On the one hand, criteria have been put forward to determine how severely non-Gibbsian a measure is. Measures have been classified according to (i) the effect of further decimation [47, 48, 39, 60], (ii) the size of the set of points of discontinuity of the conditional probabilities [40, 14, 68] and (iii) the size of the set of configurations for which a Boltzmann description is possible [5, 44, 7, 1, 8, 2, 29]. The reader is directed to [67] for a concise comparison of these classification schemes, which, however, does not include a more recently introduced fourth category of non-Gibbsianness [69]. On the other hand, some features of Gibbsian measures have been proven to hold also for different classes of non-Gibbsian fields. They include parts of the thermodynamic formalism [36, 54] and the variational approach [37, 11, 10, 33, 69].

At present, after almost 25 years of non-Gibbsian studies, the state of affairs is the following. On the positive side, we have a rather extensive catalogue of instances of the phenomenon. The more recent, and surprising, manifestations include the non-Gibbsianness of joint measures for disordered systems [28, 66, 31] —contradicting a well-known assumption in physics literature— and the appearance and disappearance of non-Gibbsianness during dynamical evolutions of the type used in Monte Carlo simulations [62]. We also have a pretty good knowledge of mechanisms leading to non-Gibbsianness. By this I mean both the physical mechanisms (hidden variables, phase transitions of restricted systems) and the mathematical tools to provide rigorous proofs. Finally, the work on the "thermodynamic" properties of non-Gibbsian measures has brought further insight in the limitations of such an approach and the different components of the usual Gibbsian variational approach.

On the negative side, we still owe concrete answers to practitioners. It is still unclear to what extent, if any, the lack of Gibbsianness of renormalized measures compromises widely accepted calculations of critical exponents. Likewise, nothing is known on possible observable consequences of non-

Gibbsianness of simulation or sampling schemes. This situation was to be expected. Non-Gibbsianness is an elusive phenomenon, involving extremely unlikely events and very special (perverse!) boundary conditions. There is no question that we deal with a phenomenon that is widespread. It shows up, for example, in intermittent dynamical systems [42] and in problems of technological relevance [69]. Nevertheless, we seem to be still at the stage of mostly mathematical finesse.

But this finesse has been very beneficial. In at least two instances it has helped clarify an important paradoxical situation. First, through the "second fundamental theorem" in [64] which dispelled, to a certain extent, the threat of discontinuities in the renormalization-group paradigm. Second, it was instrumental in reconciling contradictory hypotheses with succesful predictions in Morita's approach to the study of disordered systems [32]. Furthermore, non-Gibbsianness has forced some healthy reconsideration of known results, especially those related to the thermodynamical and variational characterization of measures. The discontinuities, often associated only to a measure-zero set of bad configurations, rended the traditional treatment invalid. Putting it dramatically, proofs were destroyed by a few very unlikely events. It is natural to enquire whether this is due to a limitation of the techniques of proof, or whether continuity is really essential. The meticulous work on Gibbsian reconstruction is teaching us how to isolate and bring into light the different ingredients of each Gibbsian result, and to appreciate the subtle balance between topology and probability theory which supports mathematical statistical mechanics.

This course can be roughly divided in two parts. The first part is an introduction of the main concepts and notions. To make it reasonably self-contained, I will start with a rather detailed exposition of the definiton and benchmark properties of Gibbsianness. In particular, I will include a hopefully pedagogical proof of Kozlov's theorem, which has been our main tool to detect non-Gibbsianness. This will lead me, quite naturally, to an early presentation of the different non-Gibbsianness classification schemes. The second part reviews examples of non-Gibbsianness. These examples show up through violations of either the non-nullness or the quasi-locality of some conditional probability. I will try to convey at least an intuitive understanding of some of the mechanisms behind these two types of violations for the case of renormalization transformations, as well as for the case of spin-flip evolutions [62] and the case of disordered systems [28].

Most of the exposition is in the style of an overview. I will try, on the one hand, to clarify the main conceptual issues and, on the other hand, to transmit the ideas and intuitions that helped develop my own understanding of the subject. In particular, the non-quasilocality instances will be organized around three "surprising" manifestations: renormalization-group pathologies, non-Gibbsianness in Glauber evolutions and non-Gibbsianness of the joint measures of disordered systems. I hope these situations are surprising enough to convince the audience that the phenomenon is important indeed. Due to time constraints, I am leaving aside the variational and thermodynamical treatments of non-Gibbsian measures. A pedagogical self-contained exposition of these issues requires a course of its own. I refer the reader respectively to [33] and [54] for state-of-the-art presentations of these topics.

2 Setup, notation and basic notions

General definitions and notation. We consider a countable set $\mathbb{L}$, called the *lattice*, formed by *sites* x and whose subsets will be called *regions*. At each site of $\mathbb{L}$ sits a copy of a *single-spin space* S. For our pedagogical purposes it is enough to consider *finite* spins, that is $2 \leq \operatorname{card}(S) < \infty$. Most of the examples below correspond to the cases $S = \{0, 1\}$ (lattice-gas models), $S = \{-1, 1\}$ (Ising spins) or $S = \{1, 2, \ldots, q\}$ (Potts spins with q colors). Gibbs measures are defined on the *configuration* *space* $\Omega = S^{\mathbb{L}}$ which represents a large array of microscopic systems, each described by S. Thus, each configuration $\omega \in \Omega$ is a collection of values $(\omega_x)_{x\in\mathbb{L}}$ where, for concreteness, each $\omega_x \in S$ will be called the value of the *spin at x*. To fix ideas, we can take the canonical case $\mathbb{L} = \mathbb{Z}^d$, but the following presentation is written so as to make a certain generality apparent. In fact, for the purely statistical mechanical theory most of the time we will consider sets $\mathbb{L}$ endowed with a distance dist such that the parallelepipeds

$$\Lambda_n(x) \;:=\; \{y \in \mathbb{L} : dist(x, y) \leq n\} \tag{2.1}$$

have a (external) boundary whose cardinality grows slower than the cardinality of Λ_n. In addition, for the thermodynamical treatment and to study ergodicity and large-deviation properties (topics that are omitted here), an *action* of $\mathbb{Z}^d$ on $\mathbb{L}$ by homeomorphisms is needed. This means that there must exist a family of bijections indexed by $\mathbb{Z}^d$ —the *translations*— $T_x : \mathbb{L} \to \mathbb{L}$, $x \in \mathbb{L}$, such that (i) $T_x^{-1} = T_{-x}$, (ii) they are continuous and measurable (with respect to the topology and σ-algebra introduced below) and (iii) they leave invariant the distance in the sense that $\mathrm{dist}(T_x y, T_x z) = \mathrm{dist}(x, y)$. The largest d for which such an action exists is the *dimension* of the lattice $\mathbb{L}$.

Another, almost gratuitous, generalisation of the setting is to consider site-dependent single-spin spaces Ω_x. For practical purposes, the only consequence of such a more general framework is to make the notation heavier, so I will not adopt it here.

Let us fix some notational conventions. The symbol "$|\ |$" will be used in several senses: cardinality of an ensemble, absolute value of a (complex) number and, if x is a site in $\mathbb{Z}^d$, $|x| = \max_{1\leq i\leq d} |x_i|$. In particular, we shall use the distance $\mathrm{dist}(x, y) = |x - y|$. We shall denote configurations by lower case Greek letters, $\omega, \sigma, \eta \in \Omega$ and finite subsets of the lattice by uppercase letters, which will be Greek when associated to lattice regions and Latin when they are to be thought of as bonds (see below). The finiteness property will be emphasized by the symbol "$\Subset$": $\Lambda, \Gamma \Subset \mathbb{L}$. Finite-region configurations will show the region as a subscript: $\omega_\Lambda \in \Omega_\Lambda := S^\Lambda$. Configurations defined by regions will be denoted in a factorized form; an omitted subscript indicating completion to the rest of the lattice: $\omega_\Lambda \eta_{\Lambda^c} = \omega_\Lambda \eta$. A configuration σ will be said *asymptotically equal* to another configuration η if it is of the form $\sigma = \omega_\Lambda \eta$ for some $\Lambda \Subset \mathbb{L}$ and $\omega \in \Omega$. Alternatively, σ will be called a *finite-region modification* of η. The r-external boundary of a region $\Lambda \subset \mathbb{L}$ $(0 < r < \infty)$ will be denoted

$$\partial_r \Lambda(x) \;:=\; \bigl\{ y \in \mathbb{L} \setminus \Lambda : d(x, y) \leq r \bigr\} \,. \tag{2.2}$$

If $\mathcal{A}$ is a set endowed with a σ-algebra Σ, we shall denote by $\mathcal{P}(\mathcal{A}, \Sigma)$ the corresponding space of probability measures. All measurable functions (random variables) will be real-valued. To avoid overcharged notation, we shall resort to well-established abuses. For instance, the Σ-measurability of a function f will be stated as $f \in \Sigma$. Likewise, if $\mu \in \mathcal{P}(\mathcal{A}, \Sigma)$ and $f \in \Sigma$, we shall denote by $\mu(f)$ its corresponding expectation.

Measure-theoretical setup. The states in classical statistical mechanics are probability measures on a product measurable space. For finite S, the starting block is the discrete σ-algebra $\mathcal{F}_0$ formed by all the subsets of S. That is, all sets of single-spin configurations are measurable. The σ-algebra for the whole lattice, $\mathcal{F} := \mathcal{F}_0^{\mathbb{L}}$, however, leaves out some (non-measurable) sets of configurations (which are difficult to construct and do not show up in the usual civilized handling of the theory). By elementary probability theorems, the family $\mathcal{F}$ can, alternatively, be defined by one of the following equivalent characterizations:

(S1) It is the σ-algebra generated by (the smallest σ-algebra containing) the set of cylinders $\{C_{\sigma_\Lambda} : \Lambda \Subset \mathbb{L}, \sigma_\Lambda \in \Omega_\Lambda\}$, where

$$C_{\sigma_\Lambda} \;=\; \{\omega \in \Omega : \omega_\Lambda = \sigma_\Lambda\} \;. \tag{2.3}$$

(S2) It is the σ-algebra generated by (the smallest σ-algebra that makes measurable) the projections

$$\begin{array}{cccc} X_x : & \Omega & \longrightarrow & S \\ & \omega & \longmapsto & \omega_x \end{array} \tag{2.4}$$

for $x \in \mathbb{L}$.

A *random field* is a probability measure on a space $(\Omega, \mathcal{F})$.
The σ-algebra $\mathcal{F}$ contains some important sub-σ-algebras:

(i) The family of events in $\Gamma \Subset \mathbb{L}$, denoted $\mathcal{F}_\Gamma$, which is defined either by (2.3) or (2.4) but considering $\Lambda \subset \Gamma$ and $x \in \Gamma$. Note that this is *not* equal to $\mathcal{F}_0^\Gamma$. The definition adopted allows us to work within a single overall measure space $(\Omega, \mathcal{F})$.

(ii) The family of *tail* or *asymptotic* events

$$\mathcal{F}_\infty \;=\; \bigcap_{\Gamma \Subset \mathbb{L}} \mathcal{F}_{\Gamma^c} \;. \tag{2.5}$$

also called σ-algebra at infinity.

The best way to understand these σ-algebras is through the functions measurable with respect to each of them. By yet another well-known theorem in probability theory a function f is measurable with respect to $\mathcal{F}_\Gamma$ if, and only if, it "depends only on the spins in Γ" in the following equivalent senses

(L1) There exists a function $\widetilde{f}_\Gamma$ on Ω_Γ such that $f = \widetilde{f}_\Gamma\big(\{X_x\}_{x\in\Gamma}\big)$.

(L2) There exists a function f_Γ on Ω_Γ such that $f(\omega) = f_\Gamma(\omega_\Gamma)$.

(L3) $f(\omega) = f(\sigma)$ whenever $\omega_\Gamma = \sigma_\Gamma$.

(Often, no notational distinction is made between $\widetilde{f}_\Gamma$, f_Γ and f).

These facts explain why a function is called *local* or *microscopic* if it is $\mathcal{F}_\Gamma$-measurable for some finite region Γ. The same theorem implies that a function g is $\mathcal{F}_\infty$-measurable if it "does *not* depend on the spins of any finite region". That is, if $g(\omega) = g(\sigma)$ whenever there exists a finite $\Gamma \Subset \mathbb{L}$ such that $\omega_{\Gamma^c} = \sigma_{\Gamma^c}$. Typically, this means that g is defined by some kind of limit. For instance,

$$g(\omega) \;=\; \begin{cases} \displaystyle\lim_{n\to\infty} \frac{1}{|\Lambda|_n} \sum_{x\in\Lambda_n} \omega_x & \text{if the limit exists} \\ 0 \text{ (or anything)} & \text{otherwise} \end{cases} \tag{2.6}$$

is $\mathcal{F}_\infty$-measurable. These *observables at infinity* or *asymptotic observables* correspond, then, to *macroscopic* averages of a system.

Topological setup. The notion of Gibbs state requires, in addition, a topology $\mathcal{T}$ compatible with the σ-algebra $\mathcal{F}$. Compatibility means any one of the following equivalent properties: (i) open sets are measurable, (ii) continuous functions are measurable, (iii) $\mathcal{T}$ generates $\mathcal{F}$ (that is, $\mathcal{F}$ is the Borel σ-algebra corresponding to the topology $\mathcal{T}$). To achieve this compatibility, $\mathcal{T}$ is endowed with the same generators as those of $\mathcal{F}$. That is, $\mathcal{T}$ is chosen as the smallest topology on Ω containing the cylinders C_{σ_Λ} or making the projections X_x continuous. This amounts to taking as single-site topology $\mathcal{T}_0$ the set of subsets of S (all single-space sets are open), and adopting the product topology $\mathcal{T} = \mathcal{T}_0^{\mathbb{L}}$ for the whole configuration space. This topology has a number of properties which follow rather directly from its definition:

(T1) (Ω, τ) is a *compact* space, because of Tychonov's theorem. This is helpful when proving existence and boundedness results.

(T2) The topology is metrizable, for instance by $d(\omega, \sigma) = \sum_i 2^{-i} \rho(\omega_i, \sigma_i)/[1+\rho(\omega_i, \sigma_i)]$ with $\rho(a, b) = 1$ if $a \neq b$ and 0 otherwise. This means that we can use sequences to determine closed sets and to test continuity.

(T3) Open sets are very big: they are products of sets which are all equal to S except at a finite number of sites, that is, they are finite unions of cylinders. In particular, the family of cylinders $\{C_{\omega_{\Lambda_m}}\}_n$ gives a base of open neighborhoods of a given configuration ω.

Property (T3) means that two configurations are "close", topologically speaking, if they coincide over large finite regions (the larger the region the closer they are). In combination with (T2) this implies that $\mathcal{T}$ can, alternatively, be characterized as the one giving "convergence through freezing". A sequence of configurations $(\omega^{(n)})$ converges to a configuration ω iff for every cube Λ there exists a natural number N such that

$$\omega_\Lambda^{(n)} = \omega_\Lambda \quad \text{for all } n \geq N \; . \tag{2.7}$$

Furthermore, this condition is also necessary and sufficient for a sequence $(\omega^{(n)})$ to be Cauchy [in the metric given in (T2), for instance]. Hence a sequence of configurations is Cauchy if, and only if, it is convergent. The space $(\Omega, \mathcal{F})$ is, therefore, *complete*. Another consequence of (T3) is that asymptotic events are dense, because they are insensitive to changes in finite regions. In particular, the set of configurations that are asymptotically constantly equal to some fixed $s \in S$ is dense. This is a countable set, thus the configuration space is *separable*. To summarize,

(T4) $(\Omega, \mathcal{F})$ is a *Polish space*, that is, metrizable, separable and complete.

From (2.7) and (T2) we conclude that a function $f : \Omega \to \mathbb{R}$ is continuous at $\omega \in \Omega$ iff for each $\epsilon > 0$ there exists $n \in \mathbb{N}$ such that

$$\sup_{\sigma \in \Omega} \Big| f(\omega_{\Lambda_n} \sigma) - f(\omega) \Big| < \epsilon \; . \tag{2.8}$$

The compactness of Ω implies that functions continuous everywhere are uniformly continuous. Hence, the continuity of a function f on the whole of Ω is equivalent to any of the following properties:

(C1) For each $\epsilon > 0$ there exists $n \in \mathbb{N}$ such that

$$\sup_{\omega \in \Omega} \sup_{\sigma \in \Omega} \Big| f(\omega_{\Lambda_n} \sigma) - f(\omega) \Big| < \epsilon \; . \tag{2.9}$$

(C2) f can be uniformly approximated by local functions: For each $\epsilon > 0$ there exists a local function f_ϵ such that

$$\|f_\epsilon - f\|_\infty < \epsilon \,. \tag{2.10}$$

We immediately see that all local functions are continuous, while all asymptotic observables are discontinuous.

In more general settings, where Ω is not compact, a function satisfying (C1) or (C2) is termed *quasilocal.* In our case, then, continuity [that is, the validity of (2.8) for all ω] is equivalent to quasilocality. We shall use both terms almost interchangeably, with a slight preference for the latter. This is in part for historical reasons, but also to emphasize the fact that in more general settings, quasilocality, rather than continuity, is the key property. In particular, I may refer to property (2.8) as *quasilocality at* ω.

A weaker notion of quasilocality will be relevant below.

Definition 2.11 *Let $\omega, \theta \in \Omega$. A function f on Ω is* **quasilocal at ω in the direction θ** *if*

$$\Big| f(\omega_{\Lambda_n}\theta) - f(\omega) \Big| \xrightarrow[n\to\infty]{} 0 \,. \tag{2.12}$$

f is **quasilocal in the direction θ** *if it satisfies (2.12) for all $\omega \in \Omega$.*

Due to the "sup" in (2.8), a function can be quasilocal at ω in *every* direction without being continuous at ω (see the example after Remark 2.5 in [10]).

Interplay between topology and measure theory. The notion of weak convergence (=weak*-convergence in functional analysis), is perhaps the most elementary concept needing a combined topological and measure-theoretical framework. A sequence of measures μ_n on $(\Omega, \mathcal{F})$ *converges weakly* to a measure μ if its continuous expectations converge, that is, if

$$\mu_n(f) \xrightarrow[n\to\infty]{} \mu(f) \quad \text{for every } \textit{continuous} \text{ function } f \,. \tag{2.13}$$

(In more general situations, the convergence is required for functions that are continuous *and bounded.* This last condition is automatic if Ω is compact, as is the case here.) Due to the density result (C2) above, weak convergence is therefore equivalent to either of the following equivalent conditions:

(W1) $\mu_n(f) \xrightarrow[n\to\infty]{} \mu(f)$ for every *local* function f.

(W2) $\mu_n\big(C_{\sigma_\Lambda}\big) \xrightarrow[n\to\infty]{} \mu\big(C_{\sigma_\Lambda}\big)$ for every cylinder C_{σ_Λ}.

In words, weak convergence means convergence of expectations of microscopic observables. It gives no information whatsoever as to the convergence of the means of the discontinuous macroscopic (asymptotic) observables. In our examples below such a convergence will fail: Infinite-region measures μ_n are typically singular with respect to each other, as well as with respect to their weak limits μ, precisely because the respective supports are disjoint asymptotic events.

Weak convergence is, indeed, an extremely weak notion of convergence (strictly weaker than other popular modes of convergence, like convergence in probability, almost-surely, in L^p sense, in total variation ...). There is both a physical and a mathematical justification for its use. From the physical point of view, it corresponds to the idea of *infinite-volume* limit, that is, on the construction of (infinite-volume) "states" by working on finite, but progressively larger volumes. Mathematically, it is the type of convergence involved in basic limit theorems (like central-limit theorems) and, moreover, it leads to a *compact* space of probability measures (Banach-Alaoglu theorem). This is an invaluable property that, for instance, reduces to a marginal comment the potentially difficult problem of existence of Gibbs measures for a given interaction.

Another notion needed later on is the following.

Definition 2.14 *A probability measure μ on a Borel measurable space is* **non-null** *if $\mu(O) > 0$ for every open set O.*

Two more instances of topology-measure theory interplay will be found later on. First, the reference to regular conditional probabilities in Polish spaces. Second, the very notion of Gibbs measure!

3 Probability kernels, conditional probabilities and statistical mechanics

3.1 Probability kernels

We turn now to more specific notions that are not always learnt in elementary probability courses. I start with the definition of a *probability kernel* which, informally, is an object with two "slots", being a probability measure with respect to one of them and a measurable function with respect to the other one. It represents a family of probability measures which depend, in a measurable fashion, on a random parameter. Two applications are of interest here: (i) conditional probabilities —measurable functions of the conditioning configuration— and (ii) stochastic transformations —measurable with respect to the initial configuration. Kernels for the second application are usually denoted in an "operator" fashion, while the "conditioning" notation is reserved for the first case. I'll adopt this last "bar" notation for both, because I always tend to think these kernels as conveying conditioning information.

Definition 3.1 *A* **probability kernel** Ψ *from a probability space $(\mathcal{A}, \Sigma)$ to another probability space $(\mathcal{A}', \Sigma')$ is a function*

$$\Psi(\,\cdot\,|\,\cdot\,) : \Sigma' \times \mathcal{A} \longrightarrow [0, 1] \tag{3.2}$$

such that

(i) $\Psi(\,\cdot\,|\omega)$ *is a probability measure on $(\mathcal{A}', \Sigma')$ for each $\omega \in \mathcal{A}$;*

(ii) $\Psi(A'|\,\cdot\,)$ *is Σ-measurable for each $A' \in \Sigma'$.*

A, perhaps familiar, illustration of this concept is given by the transition probabilities defining a (discrete-time, homogeneous) stochastic process. In this case, $\mathcal{A} = S^{-\mathbb{N}}$, $\mathcal{A}' = S^{-\mathbb{N}\cup\{0\}}$, Σ and Σ' the respective product σ-algebras, and $\Psi(A'|\omega)$ is the probability that the event A' happens at the next instant given a history ω. The specifications discussed below constitute a multi-dimensional generalization of this example.

Probability kernels can be combined by a "convolution" in the following natural way. Suppose Ψ is a kernel from $(\mathcal{A}, \Sigma)$ to $(\mathcal{A}', \Sigma')$ and Ψ' is a kernel from $(\mathcal{A}', \Sigma')$ to $(\mathcal{A}'', \Sigma'')$. Then $\Psi\Psi'$ is the kernel from $(\mathcal{A}, \Sigma)$ to $(\mathcal{A}'', \Sigma'')$ defined by

$$\big(\Psi\Psi'\big)(A''|\omega) \;=\; \Psi\Big(\Psi'(A''|\,\cdot\,)\;\Big|\;\omega\Big)\;, \tag{3.3}$$

for $A'' \in \Sigma''$ and $\omega \in \mathcal{A}$. In more detail,

$$\big(\Psi\Psi'\big)(A''|\omega) \;=\; \int_{\mathcal{A}'} \Psi(d\omega'|\omega)\,\Psi'(A''|\omega')\;. \tag{3.4}$$

This convolution leads to a map between probability measures, which are particularly simple examples of probability kernels. Indeed, a kernel Ψ from $(\mathcal{A}, \Sigma)$ to $(\mathcal{A}', \Sigma')$ defines the map

$$\begin{array}{ccc} \mathcal{P}(\mathcal{A}, \Sigma) & \longrightarrow & \mathcal{P}(\mathcal{A}', \Sigma') \\ \mu & \longmapsto & \mu' = \mu\Psi \end{array} \tag{3.5}$$

that is,

$$\mu'(A') \;=\; \int_{\mathcal{A}} \mu(d\omega)\,\Psi(A'|\omega)\;, \tag{3.6}$$

for all $A' \in \Sigma'$. This is how renormalization transformations, discussed below, are defined. In particular, *deterministic* transformations correspond to kernels concentrated on single points:

$$\Psi(A' \mid \omega) \;=\; \delta_{\psi(\omega)}(A') \tag{3.7}$$

for a certain function $\psi : \mathcal{A} \to \mathcal{A}'$. Then, $\mu'(A') = (\mu\Psi)(A') = \mu\big[\psi^{-1}(A')\big]$ and

$$\mu'(f') \;=\; \int_{\mathcal{A}} f'\big(\Psi(\omega)\big)\,\mu(d\omega)\;, \tag{3.8}$$

for $f' \in \Sigma'$.

3.2 Conditional probabilities

Marc Kac said that probability theory is measure theory with a soul. This soul —which makes probability into a full field of its own and not just a mere chapter of finite-measure theory— is the notion of conditional expectation . It is not a simple concept, though, due to the need of conditioning with respect to events of *zero* probability. I intend to review here definitions and properties of this object, so as to explain the full mathematical meaning of the crucial notion of *specification* to be introduced shortly. Readers who are impatient or reluctant to abstract considerations may prefer to jump to the next subsection and accept Definition 3.20 through the more "physical" arguments given below. Most of these subtleties would be avoidable if we were dealing only with Gibbsian measures, but they are unavoidable for a proper understanding of non-Gibbsianness.

A very popular exercise in elementary probability courses consists in showing that two events are independent if, and only if, all the events of the σ-algebras generated by them are. This observation generalizes to the fact that the information related to conditional expectations is best encoded through *functions* that correspond to conditioning with respect to whole σ-algebras. Kolmogorov taught us the right axiomatic way to define this concept.

Definition 3.9 *Let $(\mathcal{A}, \Sigma, \mu)$ be a probability space, τ a σ-algebra with $\tau \subset \Sigma$ and f a μ-integrable Σ-measurable function. A* **conditional expectation function** *of f given τ is a function*

$$E_\mu(f \mid \tau)(\,\cdot\,) : \mathcal{A} \longrightarrow \mathbb{R} \tag{3.10}$$

such that

(i) $E_\mu(f \mid \tau)$ *is* τ*-measurable.*

(ii) For any τ*-measurable bounded function* g*,*

$$\int d\mu\, g\, E_\mu(f \mid \tau) \;=\; \int d\mu\, g\, f\;. \tag{3.11}$$

Such a function $E_\mu(f \mid \tau)$ is interpreted as the expected value of f if we have access only to the information contained in τ, that is, if we can only perform an experiment determining occurrence of events in τ rather than the more detailed events in Σ. It is the "best predictor" of f, in square-integrable sense, among the τ-measurable functions. (The reader can, for instance, have a look to Chapter 9 of the book by Williams [70] for a short but clear motivation of the previous definition and its interpretation.) Identity (3.11) is the ultimate version of the quintessential probabilistic technique of decomposing an expectation into a sum of conditioned averages weighted by the probabilities of the conditioning events ("divide-and-conquer" technique).

Several remarks are in order. First, the existence of such conditional expectations is assured by the Radon-Nikodým theorem. Second, as condition (ii) involves a μ-integral, $E_\mu(f \mid \tau)$ can be modified on a set of μ-measure zero while still satisfying the definition. Thus, Definition 3.9 does not define a unique function. Measure-zero modifications, however, are the *only* ones possible. That is, $E_\mu(f \mid \tau)$ *is defined* μ*-almost surely.* Often, more appropriately, the symbol $E_\mu(f \mid \tau)$ is reserved for the whole class of functions determined by the previous definition. Here it is being used,by abuse of notation, for any particular choice —*realization*— within this class. In this way we gain concreteness but we have to remember to include a "μ-almost surely" clause in each expression relating conditional expectations. Third, conditional expectations enjoy a number of important properties, most of which are very easy to prove (nicely summarized in Section 9.7 and the inner back cover of [70]). We highlight two of them for immediate use. First, for each bounded $g \in \tau$,

$$E_\mu(g\, f \mid \tau) \;=\; g\, E_\mu(f \mid \tau) \quad \mu\text{-almost surely}\;. \tag{3.12}$$

Second, if $\widetilde{\tau}$ is an even smaller σ-algebra, that is $\widetilde{\tau} \subset \tau \subset \Sigma$, then

$$E_\mu\Big(E_\mu(f \mid \tau)\;\Big|\;\widetilde{\tau}\Big) \;=\; E_\mu(f \mid \widetilde{\tau}) \quad \mu\text{-almost surely}\;. \tag{3.13}$$

This is the well known "tower property" of conditioning.

A highly non-trivial, somehow hidden, aspect of the previous presentation, is the fact that the functions $E_\mu(f \mid \tau)$ are constructed on an "f-to-f basis". The conditional expectation for each f is constructed without any regard for the conditional expectations of other functions f. The full-measure sets granting properties like (3.12) and (3.13) are f-dependent. The question arises whether a coordinate choice of conditional expectations is possible such that there is a full-measure set where all properties work simultaneously for *all* measurable and integrable functions. This amounts to constructing a μ-full set of $\omega \in \mathcal{A}$ for which the f-dependence $f \mapsto E_\mu(f \mid \tau)(\omega)$ corresponds to a *measure* that "explains" these conditional expectations. Of course, this is not always possible (in fact, Kolmogorov's seminal contribution consisted in showing that it is largely irrelevant; most of probability theory can be developed only with conditional expectation functions —which always exist— whether or not they come from conditional probability measures). The next definition covers the case when it happens to be possible.

Definition 3.14 *Let $(\mathcal{A}, \Sigma, \mu)$ be a probability space and τ a σ-algebra with $\tau \subset \Sigma$. A* **regular condition probability** *of μ given τ is a probability kernel $\mu_{|\tau}(\,\cdot \mid \cdot\,)$ from $(\mathcal{A}, \Sigma)$ to $(\mathcal{A}, \tau)$ such that for each μ-integrable $f \in \Sigma$*

$$\mu_{|\tau}(f \mid \cdot\,) \;=\; E_\mu(f \mid \tau)(\,\cdot\,) \quad \mu\text{-almost surely}\,. \tag{3.15}$$

We can, of course, state a more direct definition by transcribing property (3.11) at the level of kernels. However, it is more convenient for our purposes to decompose such a property with the aid of identity (3.12). In this way the following proposition is obtained.

Proposition 3.16 *Let $(\mathcal{A}, \Sigma, \mu)$ be a probability space and τ a σ-algebra with $\tau \subset \Sigma$. A regular conditional probability of μ given τ is a probability kernel $\mu_{|\tau}(\,\cdot \mid \cdot\,)$ from $(\mathcal{A}, \Sigma)$ to itself such that*

(i) $\mu_{|\tau}(f \mid \cdot\,) \in \tau$ *for each μ-integrable $f \in \Sigma$.*

(ii) *μ-almost surely, $\mu_{|\tau}(g\, f \mid \cdot\,) = g\, \mu_{|\tau}(f \mid \cdot\,)$ for each bounded $g \in \tau$ and each μ-integrable $f \in \Sigma$.*

(iii) $\mu\, \mu_{|\tau} = \mu$.

[The last identity uses the compact notation introduced in (3.5)/(3.6) for the composition of a kernel with a measure.]

In our case, (and in most of the cases encountered in day-to-day probability studies) we are saved by a remarkable theorem stating that *every measure on a Polish space has regular conditional probabilities.* As this regularity holds for every choice of conditioning σ-algebra τ, the tower property (3.13) can be transcribed in terms of kernels. To make the connection with the notion of specification, let me formalize the kernel version of the tower property for families of σ-algebras.

Definition 3.17 *Let $(\mathcal{A}, \Sigma, \mu)$ be a probability space and $\{\tau_i : i \in I\}$ a family of σ-algebras with $\tau_i \subset \Sigma$, $i \in I$. A* **system of regular conditional probabilities** *of μ given the family $\{\tau_i\}$ is a family of probability kernels $\mu_{|\tau_i}(\,\cdot \mid \cdot\,)$, $i \in I$, from $(\mathcal{A}, \Sigma)$ to itself such that*

(i) *For each $i \in I$, $\mu_{|\tau_i}$ is a regular conditional probability of μ given τ_i.*

(ii) *If $i, j \in I$ are such that $\tau_i \subset \tau_j$,*

$$\mu_{|\tau_i}\, \mu_{|\tau_j} \;=\; \mu_{|\tau_i} \quad \mu\text{-almost surely}\,. \tag{3.18}$$

[The last identity looks so admirably brief thanks to the convolution notation (3.3)/(3.4).]

This definition embodies a rather central problem in probability theory: given a measure and a family of σ-algebras, find the corresponding system of regular conditional probabilities. Such a system gives complete knowledge of the measure in relation to the experiments in question. As we discuss next, the central problem in statistical mechanics goes precisely in the *opposite* direction.

3.3 Specifications. Consistency

In physical terms, statistical mechanics deals with the following problem: Given the finite-volume (microscopic) behavior of a system in equilibrium, determine the possible infinite-volume equilibrium states to which such behavior leads. The mathematical formalization of this question (in the classical = non-quantum case) passes by the following tenets:

(SM1) Equilibrium state = probability measure

(SM2) Finite regions = finite parts of an infinite system

The description of a system in a finite region $\Lambda \Subset \mathbb{L}$ is given, thus, by a probability kernel $\pi_\Lambda(\cdot \mid \cdot)$, where $\pi_\Lambda(f \mid \omega)$ represents the equilibrium value of f when the configuration outside Λ is ω. To emphasize this last fact, and for further mathematical convenience, $\pi_\Lambda(\cdot \mid \omega)$ should be considered a probability measure on the whole of Ω acting as $\delta_{\omega_{\Lambda^c}}$ outside Λ. These kernels must obey certain constraints if they are to describe *equilibrium.* To the very least they have to be consistent with the following principle:

(SM3) A system is in equilibrium in Λ if it is in equilibrium in every box $\Lambda' \subset \Lambda$.

This means that the equilibrium value of any f in Λ can also be found through expectations in Λ' with configurations in $\Lambda \setminus \Lambda'$ distributed according to the Λ-equilibrium. That is,

$$\pi_\Lambda(f \mid \omega) \;=\; \pi_\Lambda\Big(\pi_{\Lambda'}(f \mid \cdot\,) \,\Big|\, \omega\Big) \qquad (\Lambda' \subset \Lambda \Subset \mathbb{L}) \tag{3.19}$$

Putting all this together, we arrive at the notion of specification, first introduced by Preston.

Definition 3.20 *A* **specification** *on* $(\Omega, \mathcal{F})$ *is a family* $\Pi = \{\pi_\Lambda : \Lambda \Subset \mathbb{L}\}$ *of probability kernels from* $(\Omega, \mathcal{F})$ *to itself such that*

(i) $\pi_\Lambda(f \mid \omega) \in \mathcal{F}_{\Lambda^c}$ *for each* $\Lambda \Subset \mathbb{L}$ *and bounded measurable* f*.*

(ii) Each π_Λ *is* proper*:*

$$\pi_\Lambda(g\, f \mid \omega) \;=\; g(\omega)\, \pi_\Lambda(f \mid \omega) \tag{3.21}$$

for all $\omega \in \Omega$*,* $g \in \mathcal{F}_{\Lambda^c}$ *and bounded measurable* f*.*

(iii) The family Π *is* consistent*:*

$$\pi_\Lambda\, \pi_{\Lambda'} \;=\; \pi_\Lambda \tag{3.22}$$

if $\Lambda' \subset \Lambda$*.*

[Recall the convolution notation (3.3)/(3.4).]

A specification is a physical model, a complete description of how a system in equilibrium behaves at the microscopic level, the information that will be given to you, for instance, by your experimental physicist friend. Your task, as statistical mechanics specialist, is to come up with the resulting infinite-volume (i.e. macroscopic) states. According to the previous tenets, these are measures satisfying the consistency property (3.21) when Λ becomes $\mathbb{L}$.

Definition 3.23 *A measure* $\mu \in \mathcal{P}(\Omega, \mathcal{F})$ *is* **consistent** *with a specification* $\Pi = \{\pi_\Lambda : \Lambda \Subset \mathbb{L}\}$ *if*

$$\mu\, \pi_\Lambda \;=\; \mu \tag{3.24}$$

for each $\Lambda \Subset \mathbb{L}$*. Let* $\mathcal{G}(\Pi)$ *denote the set of probability measures consistent with* Π*.*

[Recall the convolution notation (3.5)/(3.6).]

The concept of specification is very general. Systems at non-zero temperature are described by the Gibbsian specifications discussed in the next section, but models with exclusions and systems at zero temperature require more singular specifications. Conditions (3.24) are often called *DLR equations* in reference to Dobrushin, Lanford and Ruelle who first set them up for Gibbsian models. The set $\mathcal{G}(\Pi)$ can be empty [18, Example (4.16)]; otherwise it is a simplex. Its extremal points have physically appealing properties (trivial tail field, short-range correlations) associated to macroscopic behavior. The existence of several consistent measures corresponds to the existence of "multiple phases", and its indeed signals the presence of a first-order phase transition.

A comparison with the preceding subsection shows that the definition of specification collects all the properties of a system of regular kernels that do not refer to the initial measure μ. Thus, a specification can be interpreted as *a system of regular conditional probabilities* defined *without reference to an underlying measure*. In fact, the goal is precisely to find measures having each π_Λ as its $\mathcal{F}_{\Lambda^c}$-conditional probability. This observation is made precise by the following proposition whose proof should be immediate.

Proposition 3.25 *Let $\Pi = \{\pi_\Lambda : \Lambda \Subset \mathbb{L}\}$ be a specification and μ a probability measure on $(\Omega, \mathcal{F})$. The following properties are equivalent:*

(i) *μ is consistent with* Π.

(ii) *$\{\pi_\Lambda : \Lambda \Subset \mathbb{L}\}$ is a system of regular conditional probabilities of μ given the family of σ-algebras $\{\mathcal{F}_{\Lambda^c} : \Lambda \Subset \mathbb{L}\}$; i.e. $\mu_{\mathcal{F}_{\Lambda^c}}(\,\cdot \mid \omega) = \pi_\Lambda(\,\cdot \mid \omega)$ for μ-almost all $\omega \in \Omega$.*

(iii) *$\pi_\Lambda(f \mid \cdot\,) = E_\mu(f \mid \mathcal{F}_{\Lambda^c})(\,\cdot\,)$ μ-almost surely for each $\Lambda \Subset \mathbb{L}$ and each μ-integrable function f.*

Thus, while in probability one usually starts with a measure and searches for its conditional probabilities, in statistical mechanics one starts with the conditional probabilities and searches for the measure. The existence of first-order phase transitions shows that finite-volume conditional expectations, unlike finite marginal distributions, *do not uniquely determine a measure.* This explains, in part, the richness of the resulting theory.

There is, nevertheless, an important difference between specifications and systems of regular conditional probabilities brought by the absence of "μ-almost surely" clauses in the former. Indeed, in the case of specifications there is no initial privileged measure and, moreover, consistency will in general lead to infinitely many relevant measures. In such a situation there is no clear way to give meaning to almost sure statements. Hence, while (ii) of Proposition 3.16 and the tower property (3.18) hold μ-almost surely, the analogous conditions of being proper and consistent —(ii) and (iii)[=(3.6)] of Definition 3.20— hold for all $\omega \in \Omega$. Thus, not every system of regular conditional probabilities forms a specification and it is natural to wonder whether each measure admits a specification or, almost equivalently, whether a regular system can always be modified so as to obtain a specification. The answer, somehow surprisingly, is a rather general "yes" [55, 57]. A more subtle question is whether such a modification can be done so as to acquire some additional properties, like continuity with respect to the external condition. This turns out to be a deep issue that is at the heart of the non-Gibbsianness phenomenon to be studied later.

In our finite-spin setting, each proper kernel π_Λ is absolutely continuous with respect to the product of the counting measure in Ω_Λ and a delta measure on Ω_{Λ^c}.

Definition 3.26 *The* **specification densities** *of a specification $\Pi = \{\pi_\Lambda : \Lambda \Subset \mathbb{L}\}$ are the functions $\gamma_\Lambda(\,\cdot \mid \cdot\,) : \Omega_\Lambda \times \Omega_{\Lambda^c} \to [0,1]$ defined by*

$$\gamma_\Lambda(\sigma_\Lambda \mid \omega_{\Lambda^c}) \;:=\; \pi_\Lambda(C_{\sigma_\Lambda} \mid \omega)\;, \tag{3.27}$$

that is, the functions such that,

$$\pi_\Lambda(f \mid \omega) \;=\; \sum_{\sigma_\Lambda \in \Omega_\Lambda} f(\sigma_\Lambda \omega)\, \gamma_\Lambda(\sigma_\Lambda \mid \omega_{\Lambda^c}) \tag{3.28}$$

for every bounded measurable f.

These densities will be the main characters of the presentation below. They enjoy a number of useful properties. The consistency relation (3.22) applied to $f = 1\!\!1_{C_{\sigma_{\Lambda'}}}$ yields

$$\gamma_\Lambda(\sigma_\Lambda \mid \omega_{\Lambda^c}) \;=\; \sum_{\eta_\Lambda \in \Omega_\Lambda} \gamma_{\Lambda'}(\sigma_{\Lambda'} \mid \sigma_{\Lambda \setminus \Lambda'}\, \omega_{\Lambda^c})\, \gamma_\Lambda(\eta_{\Lambda'}\, \sigma_{\Lambda\setminus\Lambda'} \mid \omega_{\Lambda^c})\;. \tag{3.29}$$

From this we readily obtain a key *bar-displacement property* that will be intensively exploited in the proof of Kozlov's theorem:

Proposition 3.30 *Let $\{\gamma_\Lambda : \Lambda \Subset \mathbb{L}\}$ be a family of densities of a specification on $(\Omega, \mathcal{F})$. Consider regions $\Lambda' \subset \Lambda \Subset \mathbb{L}$ and configurations α, σ and ω such that $\gamma_{\Lambda'}(\alpha_{\Lambda'} \mid \sigma_{\Lambda\setminus\Lambda'}\, \omega_{\Lambda^c}) > 0$. Then,*

$$\frac{\gamma_\Lambda(\beta_{\Lambda'}\, \sigma_{\Lambda\setminus\Lambda'} \mid \omega_{\Lambda^c})}{\gamma_\Lambda(\alpha_{\Lambda'}\, \sigma_{\Lambda\setminus\Lambda'} \mid \omega_{\Lambda^c})} \;=\; \frac{\gamma_{\Lambda'}(\beta_{\Lambda'} \mid \sigma_{\Lambda\setminus\Lambda'}\, \omega_{\Lambda^c})}{\gamma_{\Lambda'}(\alpha_{\Lambda'} \mid \sigma_{\Lambda\setminus\Lambda'}\, \omega_{\Lambda^c})} \tag{3.31}$$

for every configuration β.

In words: the conditioning bar can be freely moved, as long as the external configurations of numerator and denominator remain identical. In fact, this condition amounts to an alternative way to define specifications in our finite-spin setting (this way is particularly popular within the Russian school.)

Exercise 3.32 *Show that a family of strictly positive density functions γ_Λ defines a specification if, and only if,*

(i) they are normalized: $\sum_{\sigma_\Lambda \in \Omega_\Lambda} \gamma_\Lambda(\sigma_\Lambda \mid \omega_{\Lambda^c}) = 1$ for every configuration ω, and

(ii) they satisfy relation (3.31) for all configurations α, β, σ and ω.

A double application of the key relation (3.31) yields the telescoping formula

$$\frac{\gamma_\Lambda(\beta_\Lambda \mid \omega_{\Lambda^c})}{\gamma_\Lambda(\alpha_\Lambda \mid \omega_{\Lambda^c})} \;=\; \frac{\gamma_{\{x\}}(\beta_x \mid \beta_{\Lambda\setminus\{x\}}\, \omega_{\Lambda^c})}{\gamma_{\{x\}}(\alpha_x \mid \beta_{\Lambda\setminus\{x\}}\, \omega_{\Lambda^c})}\, \frac{\gamma_{\Lambda\setminus\{x\}}(\beta_{\Lambda\setminus\{x\}} \mid \alpha_x\, \omega_{\Lambda^c})}{\gamma_{\Lambda\setminus\{x\}}(\alpha_{\Lambda\setminus\{x\}} \mid \alpha_x\, \omega_{\Lambda^c})}\;, \tag{3.33}$$

which implies that the single-site densities *characterize* the specification. That is, we are led to the following proposition, which is a particular case of [18, Theorem (1.33)].

Proposition 3.34 *A specification with strictly positive densities can be reconstructed, in a unique way, from its single-site densities through (3.33). As a consequence, two specifications with strictly positive densities are equal if, and only if, their single-site densities coincide.*

To benefit from this result we need a family of singletons that are known to come from a specification. A more involved question is the *construction* or *extension* issue, namely under which conditions a family of single-site densities can be extended to a full specification. To see that some conditions are needed let us apply (3.33) for $\Lambda = \{x, y\}$:

$$\frac{\gamma_{\{x,y\}}(\beta_{\{x,y\}} \mid \omega_{\{x,y\}^{\rm c}})}{\gamma_{\{x,y\}}(\alpha_{\{x,y\}} \mid \omega_{\{x,y\}^{\rm c}})} \;=\; \frac{\gamma_{\{x\}}(\beta_x \mid \beta_y\, \omega_{\{x,y\}^{\rm c}})}{\gamma_{\{x\}}(\alpha_x \mid \beta_y\, \omega_{\{x,y\}^{\rm c}})}\, \frac{\gamma_{\{y\}}(\beta_y \mid \alpha_x\, \omega_{\{x,y\}^{\rm c}})}{\gamma_{\{y\}}(\alpha_y \mid \alpha_x\, \omega_{\{x,y\}^{\rm c}})}\,. \tag{3.35}$$

The normalization $\sum_{\beta_{\{x,y\}}} \gamma_{\{x,y\}}(\beta_{\{x,y\}} \mid \omega_{\{x,y\}^{\rm c}}) = 1$ yields

$$\gamma_{\{x,y\}}(\alpha_{\{x,y\}} \mid \omega_{\{x,y\}^{\rm c}}) \;=\; \frac{\gamma_{\{y\}}(\alpha_y \mid \alpha_x\, \omega_{\{x,y\}^{\rm c}})}{\displaystyle\sum_{\beta_{\{x,y\}}} \frac{\gamma_{\{x\}}(\beta_x \mid \beta_y\, \omega_{\{x,y\}^{\rm c}})}{\gamma_{\{x\}}(\alpha_x \mid \beta_y\, \omega_{\{x,y\}^{\rm c}})}\, \gamma_{\{y\}}(\beta_y \mid \alpha_x\, \omega_{\{x,y\}^{\rm c}})} \tag{3.36}$$

This expression is, indeed, an algorithm to construct a two-site density starting from single-site functions. A similar formula holds, of course, interchanging x with y. For the algorithm to make sense, both (3.36) and its $x \leftrightarrow y$ permutation must be equal. It is not hard to check that this equality is, in fact, a necessary and sufficient condition for single-site kernels to yield unique consistent two-site kernels. In fact, as we point out in [12, Appendix], this is just the condition needed to construct consistent kernels for *all* finite regions.

Proposition 3.37 *Let $\{\gamma_{\{x\}} : x \in \mathbb{L}\}$ be a family of strictly positive functions $\gamma_{\{x\}}(\,\cdot \mid \cdot\,) : \Omega_x \times \Omega \longrightarrow (0, 1]$ satisfying*

(i) the normalization condition

$$\sum_{\sigma_x} \gamma_{\{x\}}(\sigma_x \mid \omega_{\{x,\}^{\rm c}}) \;=\; 1 \tag{3.38}$$

for all $\omega \in \Omega$, and

(ii) the order-consistency condition

$$\frac{\gamma_{\{y\}}(\alpha_y \mid \alpha_x\, \omega_{\{x,y\}^{\rm c}})}{\displaystyle\sum_{\beta_{\{x,y\}}} \frac{\gamma_{\{x\}}(\beta_x \mid \beta_y\, \omega_{\{x,y\}^{\rm c}})}{\gamma_{\{x\}}(\alpha_x \mid \beta_y\, \omega_{\{x,y\}^{\rm c}})}\, \gamma_{\{y\}}(\beta_y \mid \alpha_x\, \omega_{\{x,y\}^{\rm c}})} \;=\; \frac{\gamma_{\{x\}}(\alpha_x \mid \alpha_y\, \omega_{\{x,y\}^{\rm c}})}{\displaystyle\sum_{\beta_{\{x,y\}}} \frac{\gamma_{\{y\}}(\beta_y \mid \beta_x\, \omega_{\{x,y\}^{\rm c}})}{\gamma_{\{y\}}(\alpha_y \mid \beta_x\, \omega_{\{x,y\}^{\rm c}})}\, \gamma_{\{x\}}(\beta_x \mid \alpha_y\, \omega_{\{x,y\}^{\rm c}})} \tag{3.39}$$

for all $\alpha_x, \alpha_y \in S$ and $\omega \in \Omega$.

Then, there exists a unique specification with strictly positive densities having the $\gamma_{\{x\}}$ as its single-site densities. Furthermore, a probability measure μ is consistent with this specification if and only if it is consistent with the single-site kernels defined by the densities $\gamma_{\{x\}}$.

The proof of this proposition does not use the product structure of Ω, hence it also works for non-strictly positive specifications whose zeros come from local exclusion rules (one must just declare Ω to be the set of allowed configurations). Extension conditions, and construction algorithms when the kernels have zeros determined by asymptotic events, are discussed in [3, 4, 12, 13].

4 What it takes to be Gibbsian

4.1 Boltzmann prescription. Gibbs measures

Spin systems at equilibrium at non-zero temperature are described through Gibbsian specifications. They are defined through the Boltzmann prescription $\gamma_\Lambda \sim \mathrm{e}^{-\beta H_\Lambda}$, where H_Λ —the Hamiltonian— is a function in units of energy and β a constant in units of inverse energy. It is the inverse of the product of the temperature times the Boltzmann constant, but it is briefly called the *inverse temperature* which is the correct name if the temperature is measured in electron-volts. Of course, every non-null γ_Λ can be written as the exponential of something, but not everything has the right to be called a Hamiltonian in statistical mechanics. To model microphysics, the Hamiltonian must be a sum of local terms representing interaction energies among finite (microscopic) groups of spins. The set of these interaction energies is, thus, the basic object of the prescription.

Definition 4.1 *An* **interaction** *or* **interaction potential** *or* **potential** *is a family $\Phi = \{\phi_A : A \Subset \mathbb{L}\}$ of functions $\phi_A : \Omega \to \mathbb{R}$ such that $\phi_A \in \mathcal{F}_A$ (that is, ϕ_A depends only on the spins in the finite set A), for each $A \Subset \mathbb{L}$. Furthermore:*

- *The* **bonds** *of Φ are those finite sets A for which $\phi_A \neq 0$. Let us denote by $\mathcal{B}_\Phi$ the set of bonds.*
- *Φ is of* **finite range** *if the diameter of the bonds of Φ does not exceed a certain $r < \infty$ (the* **range***).*

Alternatively, interactions are specified writing the formal sum $H = \sum_{A\in\mathcal{B}} \phi_A$. Such an expression must be interpreted just as a bookkeeping expression.

The pair (Ω, Φ) constitute a Gibbsian *model*. The *Ising model* is, perhaps, the most popular one. It is defined by $\mathbb{L} = \mathbb{Z}^d$, $S = \{-1, 1\}$ and

$$\phi_A(\omega) \;=\; \begin{cases} -J_{\{x,y\}}\, \omega_x\omega_y & \text{if } A = \{x, y\} \text{ with } |x-y| = 1 \\ -h_x\, \omega_x & \text{if } A = \{x\} \\ 0 & \text{otherwise} \end{cases} \tag{4.2}$$

or, alternatively, $H = -\sum_{\langle x,y\rangle} J_{\{x,y\}}\, \omega_x\omega_y - \sum_x h_x\, \omega_x$. The constants $J_{\{x,y\}}$ are the nearest-neighbor *couplings*, and h_x is the *magnetic field* at x (these parameters are constant in the translation-invariant case). The notation "$\langle x,y\rangle$" is a standard way to indicate pairs of nearest-neighbor sites x, y. The minus signs are a concession to physics that demands that energy-lowering operations be alignment with the field and alignment, resp. anti-alignment, of neighboring spins in the ferromagnetic ($J_{\{x,y\}} \geq 0$), resp. anti-ferromagnetic ($J_{\{x,y\}} \leq 0$) case. The change of variables $\xi_x = (\omega_x + 1)/2$ produces the *lattice gas* model. Another well-studied model is the *Potts model* with q colors: $\mathbb{L} = \mathbb{Z}^d$, $S = \{1, 2, \ldots, q\}$ and $H = -\sum_{\langle x,y\rangle} J_{\{x,y\}}\, \mathbb{1}_{\{\omega_x=\omega_y\}}$.

The definitions of Hamiltonian and Boltzmann weights require the specification of conditions assuring the existence of the relevant series. The following definition refers to the weakest of such conditions.

Definition 4.3 *Let Φ be an interaction.*

- *The* **Hamiltonian** *for a region $\Lambda \Subset \mathbb{L}$ with frozen external condition ω is the real-valued function defined by*

$$H^\Phi_\Lambda(\sigma_\Lambda \mid \omega_{\Lambda^c}) \;=\; \sum_{A\Subset\mathbb{L}:A\cap\Lambda\neq\emptyset} \phi_A(\sigma_\Lambda\omega) \tag{4.4}$$

  for $\sigma, \omega \in \Omega$ such that the sum exists.

- Φ *is* **summable at** $\omega \in \Omega$ *if* $H_\Lambda^\Phi(\sigma_\Lambda \mid \omega_{\Lambda^c})$ *exists for all* $\Lambda \Subset \mathbb{L}$ *and all* $\sigma_\Lambda \in \Omega_\Lambda$. *Let us denote* $\Omega_{\rm sum}^\Phi$ *the set of configurations at which the interaction is summable.*

[Let me recall that $\sum_{A \ni x} \phi_A(\omega)$ exists iff the sequence $S_n(\omega) = \sum_{A: x\in A \subset V_n} \phi_A(\omega)$ is Cauchy.]

Definition 4.5 *The* **Boltzmann weights** *for an interaction* Φ *are the functions defined for all* $\Lambda \Subset \mathbb{L}$ *and all* $\omega \in \Omega_{\rm sum}^\Phi$ *by*

$$\gamma_\Lambda^\Phi(\sigma_\Lambda \mid \omega_{\Lambda^c}) \;=\; \frac{{\rm e}^{-H_\Lambda^\Phi(\sigma_\Lambda|\omega_{\Lambda^c})}}{Z_\Lambda^\Phi(\omega)}\;, \tag{4.6}$$

where $Z_\Lambda^\Phi(\omega)$ *is the* **partition function**

$$Z_\Lambda^\Phi(\omega) \;=\; \sum_{\omega_\Lambda \Subset \Omega_\Lambda} {\rm e}^{-H_\Lambda^\Phi(\sigma_\Lambda|\omega_{\Lambda^c})}\;. \tag{4.7}$$

Notice that the β factor has been absorbed into H_Λ, which amounts to a redefinition of the interaction. This stresses the fact that this factor plays no role in the discussion of general properties of Gibbs measures. It is, however, essential for the study of phase transitions. Keeping to tradition, I reserve the right to include it explicitly or absorb it according to needs.

Gibbsianness demands summability in a very strong sense.

Definition 4.8 *An interaction* Φ *on* $(\Omega, \mathcal{F})$ *is* **uniformly absolutely summable** *if*

$$\sum_{A \ni x} \|\Phi_A\|_\infty \;<\; \infty \quad \textit{for each } x \in \mathbb{L}\;. \tag{4.9}$$

The set of such uniformly absolutely summable interactions will be denoted $\mathcal{B}_1$.

This is much more than just demanding $\Omega_{\rm sum}^\Phi = \Omega$.

Definition 4.10 *On* $(\Omega, \mathcal{F})$:

- *The* **Gibbsian specification defined by an interaction** $\Phi \in \mathcal{B}_1$ *is the specification* Π^Φ *having the* Φ*-Boltzmann weights as densities, that is, defined by (3.28) for the weights* γ_Λ^Φ.
- *The* **Gibbs measures** *for an interaction* $\Phi \in \mathcal{B}_1$ *are the measures consistent with* Π^Φ.
- Π *is a* **Gibbsian specification** *if there exists an interaction* $\Phi \in \mathcal{B}_1$ *such that* $\Pi = \Pi^\Phi$.
- μ *is a* **Gibbsian measure** *(or Gibbsian random field) if there exists a* $\Phi \in \mathcal{B}_1$ *such that* $\mu \in \mathcal{G}(\Pi^\Phi)$.

Exercise 4.11 *Prove that* Π^Φ *is a specification if* $\Omega_{\rm sum}^\Phi = \Omega$.

Exercise 4.12 *Summability conditions weaker than (4.9) are also in the market. An interaction is*

- **absolutely summable** *if* $\sum_{A \ni x} \|\Phi_A(\omega)\|_\infty$ *converges for each* $\omega \in \Omega$ *and each* $x \in \mathbb{L}$;
- **uniformly summable** *if* $\sum_{A \ni x} \Phi_A(\omega)$ *converges uniformly on* $\omega \in \Omega$ *for each* $x \in \mathbb{L}$.

Find :

(i) An interaction that is uniformly but not absolutely summable. [Hint: *Consider* $\Phi_A = (-1)^n c_n$ *if* $A = \Lambda_n$ *and zero otherwise, for suitable functions* c_n.]

(ii) An interaction that is absolutely but not uniformly summable.

4.2 Properties of Gibbsian (and some other) specifications

The Gibbsian formalism —random fields consistent with specifications defined by Boltzmann weights— leads to an extremely successful description of physical reality. It provided a unified explanation of many experimental facts and phenomenological recipes, and it has been an infallible tool to study new phenomena. It explains thermodynamics, that is, the emergence of state functions like entropy and free energy, related by Legendre transforms, which contain the information needed to determine the thermal properties of matter systems. Furthermore, it provides a detailed description of phase transitions, and leads to the prediction of universal critical exponents generalizing the law of corresponding states.

Here we are interested in the mathematical properties of Gibbsian objects. Let me start by the observation that the map $\Phi \to \Pi^{\Phi}$ is far from one-to-one. Interactions can be redefined, by combining local terms, in infinitely many ways without changing the corresponding Boltzmann weights. All such interactions should be identified.

Definition 4.13 *Two interactions Φ and $\widetilde{\Phi}$, on the same space $(\Omega, \mathcal{F})$ are* **physically equivalent** *if $\pi^{\Phi}_{\Lambda} = \pi^{\widetilde{\Phi}}_{\Lambda}$ for each $\Lambda \Subset \mathbb{L}$. In our finite-spin setting, this is equivalent to $\gamma^{\Phi_\Lambda} = \gamma^{\widetilde{\Phi}}_{\Lambda}$ for each $\Lambda \Subset \mathbb{L}$*

While interactions are the right way to encode the physical information —and an economic way to parametrize families of measures—, specifications are the determining mathematical objects. Traditionally interactions have taken the center of the stage, but a specification-based approach has the advantage of avoiding the multi-valuedness problem associated to physical equivalence, which can lead to rather confusing situations [61]. Such an approach is, in fact, essential for a comparative study of Gibbsian and non-Gibbsian fields. The very beginning of this "interaction-free" program is the detection of the key features of Gibbsian specifications that single them out from the rest. This is the object of the rest of the section.

We start by determining important properties of specifications that follow from basic attributes of an underlying interaction. Given our focus on the finite-spin situation, we write them in terms of the density function. Foreseeing our non-Gibbsian needs, we shall distinguish among configurational, directional and uniform versions of each property. First, we notice that Boltzmann densities are never zero, furthermore, if $\Phi \in \mathcal{B}_1$, this non-nullness is uniform.

Definition 4.14 *A specification Π on $(\Omega, \mathcal{F})$ with densities $\{\gamma_\Lambda : \Lambda \Subset \mathbb{L}\}$ is:*

- **Non-null at** $\omega \in \Omega$ *if*
$$\gamma_\Lambda(\sigma_\Lambda \mid \omega_{\Lambda^c}) \;>\; 0 \tag{4.15}$$
for each $\Lambda \Subset \mathbb{L}$ and $\sigma_\Lambda \in \Omega_\Lambda$. Due to (3.29), this property is equivalent to **non-nullness in direction** *ω, that is, non-nullness at all configurations asymptotically equal to ω.*

- **Non-null** *if it is non-null at all $\omega \in \Omega$.*

- **Uniformly non-null** *if for each $\Lambda \Subset \mathbb{L}$*
$$\inf_{\sigma_\Lambda \in \Omega_\Lambda, \omega_{\Lambda^c} \in \Omega_{\Lambda^c}} \gamma_\Lambda(\sigma_\Lambda \mid \omega_{\Lambda^c}) \;=:\; c_\Lambda > 0 \,. \tag{4.16}$$

The most immediate consequence of uniform non-nullness is the following.

Proposition 4.17 *A measure consistent with an uniform non-null specification is non-null.*

We next observe that a range-r interaction produces weights —or kernels — that are insensitive to spins beyond the r-boundary of the region. This motivates the following definition.

Definition 4.18 *A specification* Π *on* $(\Omega, \mathcal{F})$ *with densities* $\{\gamma_\Lambda : \Lambda \Subset \mathbb{L}\}$ *is:*

- r-**Markovian in direction** $\theta \in \Omega$ *if*

$$\gamma_\Lambda(\sigma_\Lambda \mid \omega_{\partial_r \Lambda} \eta) - \gamma_\Lambda(\sigma_\Lambda \mid \omega_{\partial_r \Lambda} \widetilde{\eta}) \;=\; 0 \tag{4.19}$$

  for all $\Lambda \Subset \mathbb{L}$ *and all* $\sigma, \omega, \eta, \widetilde{\eta} \in \Omega$ *such that* η *and* $\widetilde{\eta}$ *are asymptotically equal to* θ.

- r-**Markovian** *if (4.19) holds for all* ω, η *and* $\widetilde{\eta}$ *in* Ω*, or, equivalently, if* $\pi_\Lambda(A \mid \cdot\,) \in \mathcal{F}_{\partial_r \Lambda}$ *for all* $\Lambda \Subset \mathbb{L}$ *and all* $A \in \mathcal{F}_\Lambda$.

- **Markovian** *(resp.* **Markovian in direction** $\theta \in \Omega$*) if it is* r*-Markovian (resp.* r*-Markovian in direction* $\theta \in \Omega$*) for some* $r \geq 0$.

For general, possibly infinite-range, interactions in $\mathcal{B}_1$ a simple calculation shows that strict Markovianness becomes "almost Markovianness" in the sense that the difference (4.19) becomes zero only in the limit $r \to \infty$. In our setting, this corresponds to continuity with respect to the external condition [recall the discussion around and following display (2.8)]. The corresponding definitions are as follows.

Definition 4.20 *A specification* Π *on* $(\Omega, \mathcal{F})$ *with densities* $\{\gamma_\Lambda : \Lambda \Subset \mathbb{L}\}$ *is:*

- **Quasilocal at** ω **in the direction** θ *iff*

$$\Big|\gamma_\Lambda(\sigma_\Lambda \mid \omega_{\Lambda_n} \theta) - \gamma_\Lambda(\sigma_\Lambda \mid \omega)\Big| \;\xrightarrow[n\to\infty]{}\; 0 \tag{4.21}$$

  for each $\Lambda \Subset \Lambda$ *and each* $\sigma_\Lambda \in \Omega_\Lambda$.

- **Quasilocal at** ω *iff it is quasilocal at* ω *in all directions, that is, iff*

$$\sup_{\eta, \widetilde{\eta} \in \Omega} \Big|\gamma_\Lambda(\sigma_\Lambda \mid \omega_{\Lambda_n} \eta) - \gamma_\Lambda(\sigma_\Lambda \mid \omega_{\Lambda_n} \widetilde{\eta})\Big| \;\xrightarrow[n\to\infty]{}\; 0 \tag{4.22}$$

  for each $\Lambda \Subset \Lambda$ *and each* $\sigma_\Lambda \in \Omega_\Lambda$.

- **Quasilocal** *iff*

$$\sup_{\omega, \eta, \widetilde{\eta} \in \Omega} \Big|\gamma_\Lambda(\sigma_\Lambda \mid \omega_{\Lambda_n} \eta) - \gamma_\Lambda(\sigma_\Lambda \mid \omega_{\Lambda_n} \widetilde{\eta})\Big| \;\xrightarrow[n\to\infty]{}\; 0 \tag{4.23}$$

  for each $\Lambda \Subset \Lambda$ *and each* $\sigma_\Lambda \in \Omega_\Lambda$.

In more general settings continuity is not equivalent to uniform continuity. In such situations we therefore obtain weaker definitions by replacing "quasilocal" by "continuous" and removing the "sup" in (4.22) and (4.23). A continuous specification is also called *Feller*. For our finite-spin models, Feller and quasilocality are synonymous. Let me also observe that, given the compactness of our configuration space, for a quasilocal specification non-nullness is equivalent to uniform non-nullness (the minimum is achieved).

With these definitions we can now state the easy part of Kozlov theorem.

Proposition 4.24 (Necessary conditions for Gibbsianness) *If a specification is Gibbsian, then it is uniformly non-null and quasilocal.*

Corollary 4.25 *Every Gibbsian measure is non-null.*

Exercise 4.26 *Prove Proposition 4.24. Start by proving that for* $\Phi \in \mathcal{B}_1$ *the functions* $\omega \to H_\Lambda^\Phi(\sigma_\Lambda \mid \omega_{\Lambda^c})$ *are continuous.*

Thanks to Proposition 3.34, all the preceding properties are inherited from single-site kernels.

Proposition 4.27 *Let* Π *be a specification in* Ω *with densities* $\{\gamma_\Lambda : \Lambda \Subset \mathbb{L}\}$ *and* $\omega, \theta \in \Omega$.

(a) Π *is non-null at* ω*, respectively non-null, uniformly non-null, iff the corresponding property in Definition 4.14 is satisfied for all single-site densities* $\gamma_{\{x\}}$.

(b) If Π *is non-null at* θ*, then it is* r*-Markovian in direction* θ *iff the corresponding property in Definition 4.18 is satisfied for all single-site densities* $\gamma_{\{x\}}$.

(c) If Π *is non-null at* θ*, then it is quasilocal at* ω *in the direction* θ*, respectively quasilocal at* ω *iff (4.21), respect. (4.22), is satisfied for all single-site densities* $\gamma_{\{x\}}$ *and all finite-region modifications of* ω.

(d) If Π *is uniformly non-null, then it is quasilocal iff (4.23) is satisfied for all single-site densities* $\gamma_{\{x\}}$.

Exercise 4.28 *Given a specification on* $\{-1,1\}^{\mathbb{L}}$*, consider the* spin-flip relative energies h_x *defined by the identity*

$$\frac{\gamma_{\{x\}}(\sigma_x \mid \omega)}{\gamma_{\{x\}}(-\sigma_x \mid \omega)} \;=\; \exp\Bigl\{-h_x(\sigma_x \mid \omega)\Bigr\} . \tag{4.29}$$

(i) Rewrite the previous proposition in terms of properties of h_x.

(ii) Write an analogous result for arbitrary spins, replacing the spin-flip by a permutation of S.

(The use of h_x *is favored by the Flemish school.)*

We finish this subsection with an illustration of how topology and measure theory combine to match physics.

Theorem 4.30 *A non-null probability measure on* $(\Omega, \mathcal{F})$ *is consistent with at most one quasi-local specification.*

In particular this means that a Gibbs measure can be Gibbsian for only *one* quasilocal specification, only one interaction modulo physical equivalence, only one temperature, ... A very rewarding result.
Proof. Let μ be a measure consistent with two quasilocal specifications Π, $\widetilde{\Pi}$ of respective kernels and densities π_Λ, $\widetilde{\pi}_\Lambda$, γ_Λ and $\widetilde{\gamma}_\Lambda$, $\Lambda \Subset \mathbb{L}$. For each such Λ and each $\sigma_\Lambda \in \Omega_\Lambda$, let

$$A_n \;=\; \Bigl\{w \in \Omega : \gamma_\Lambda(\sigma_\Lambda \mid \omega_{\Lambda^c}) - \widetilde{\gamma}_\Lambda(\sigma_\Lambda \mid \omega_{\Lambda^c}) > \frac{1}{n}\Bigr\} . \tag{4.31}$$

We have

$$\begin{aligned} 0 \;&=\; \mu\Bigl(\pi_\Lambda(\mathbb{1}_{A_n} C_{\sigma_\Lambda} \mid \cdot\,) - \widetilde{\pi}_\Lambda(\mathbb{1}_{A_n} C_{\sigma_\Lambda} \mid \cdot\,)\Bigr) \\ &=\; \mu\Bigl(\mathbb{1}_{A_n}\bigl[\gamma_\Lambda(\sigma_\Lambda \mid \cdot\,) - \widetilde{\gamma}_\Lambda(\sigma_\Lambda \mid \cdot\,)\bigr]\Bigr) \\ &>\; \frac{1}{n}\,\mu(A_n) . \end{aligned} \tag{4.32}$$

Hence $\mu(A_n) = 0$ and, by the σ-additivity of μ,

$$\gamma_\Lambda(\sigma_\Lambda \mid \cdot\,) \;\geq\; \widetilde{\gamma}_\Lambda(\sigma_\Lambda \mid \cdot\,) \qquad \mu\text{-almost surely.} \tag{4.33}$$

But, as μ is non-null, the set of points where (4.33) holds must be dense, and the continuity of both $\gamma_\Lambda(\sigma_\Lambda \mid \cdot\,)$ and $\widetilde{\gamma}_\Lambda(\sigma_\Lambda \mid \cdot\,)$ implies that, in fact,

$$\gamma_\Lambda(\sigma_\Lambda \mid \omega_{\Lambda^c}) \;\geq\; \widetilde{\gamma}_\Lambda(\sigma_\Lambda \mid \omega_{\Lambda^c}) \qquad \text{for all } \sigma_\Lambda \in \Omega_\Lambda \text{ and } \omega_{\Lambda^c} \in \Omega_{\Lambda^c}\,. \tag{4.34}$$

This argument also proves the opposite inequality through the interchange $\gamma_\Lambda \leftrightarrow \widetilde{\gamma}_\Lambda$. □

4.3 The Gibbsianness question

We turn now to the inverse of Proposition 4.24, namely the determination of sufficient conditions for Gibbsianness. This is a key step towards the development of a specification-based theory not relying on explicit choices of potentials. The issue is: *Which conditions grant that for a specification* Π *there exists some* $\Phi \in \mathcal{B}_1$ *such that* $\Pi = \Pi^\Phi$.

Historically, this question was first addressed —and solved— for Markovian fields. The simplest and most informative solution was proposed by Grimmett [23] who gave an explicit form of the potential using Möbius transform. Kozlov [27] proved the general version by generalizing this argument. An alternative proof was given simultaneously by Sullivan [58], but using a slightly different space of interactions. In the sequel I try to present a pedagogical exposition of Kozlov's proof and its consequences for the non-Gibbsianness typology.

Kozlov answered the Gibbsianness question by actually reconstructing a potential out of the given specification. From all the physically equivalent interactions he chose those with the *vacuum* property.

Definition 4.35 *An interaction* Φ *in* Ω *has* **vacuum** $\theta \in \Omega$ *if*

$$\phi_A(\omega) \;=\; 0 \qquad \text{if } \omega_i = \theta_i \text{ for some } i \in A \tag{4.36}$$

for all $A \Subset \mathbb{L}$.

The detailed proof of Kozlov's theorem involves a number of stages.

4.3.1 Construction of the vacuum potential

As a first step, let us obtain the formulas proposed by Kozlov (and Grimmett before him). This is actually not hard. We are presented with an initial specification with kernels π_Λ and densities γ_Λ, we choose a vacuum configuration θ and we search a potential Φ satisfying the vacuum condition (4.36) and such that the Boltzmann prescription (4.6)–(4.7) leads to the initial densities. We follow the natural strategy: We pretend that such a potential exists and see what we get analyzing first one-site regions, then two-site regions, and so on. In this way we obtain its only possible expression. This expression involves ratios of densities, thus some degree of non-nullness is required.

The first observation is that

$$H_\Lambda(\theta_\Lambda \mid \theta_{\Lambda^c}) \;=\; 0 \tag{4.37}$$

due to the vacuum condition (4.7), hence

$$\gamma_\Lambda(\theta_\Lambda \mid \theta_{\Lambda^c}) \;=\; \frac{1}{Z_\Lambda(\theta)} \tag{4.38}$$

for all $\Lambda \Subset \mathbb{L}$. For one-site regions the vacuum condition implies that

$$H_{\{x\}}(\sigma_x \mid \theta_{\{x\}^c}) \;=\; \phi_{\{x\}}(\sigma_x)\,. \tag{4.39}$$

Thus, the Boltzmann prescription and (4.38) imply

$$\mathrm{e}^{-\phi_{\{x\}}(\sigma_x)} \;=\; \frac{\gamma_{\{x\}}(\sigma_x \mid \theta_{\{x\}^c})}{\gamma_{\{x\}}(\theta_x \mid \theta_{\{x\}^c})} \tag{4.40}$$

for all $x \in \mathbb{L}$. Two-site regions come next. By the vacuum condition,

$$H_{\{x,y\}}(\sigma_{\{x,y\}} \mid \theta_{\{x,y\}^c}) \;=\; \phi_{\{x,y\}}(\sigma_{\{x,y\}}) + \phi_{\{x\}}(\sigma_x) + \phi_{\{y\}}(\sigma_y)\,. \tag{4.41}$$

Therefore, the Boltzmann prescription plus the preceding one-site calculations lead us to

$$\begin{aligned}
\mathrm{e}^{-\phi_{\{x,y\}}(\sigma_{\{x,y\}})} \;&=\; \frac{\gamma_{\{x,y\}}(\sigma_{\{x,y\}} \mid \theta_{\{x,y\}^c})}{\gamma_{\{x,y\}}(\theta_{\{x,y\}} \mid \theta_{\{x,y\}^c})} \times \mathrm{e}^{\phi_{\{x\}}(\sigma_x)} \times \mathrm{e}^{\phi_{\{y\}}(\sigma_y)}\\
&=\; \left[\frac{\gamma_{\{x,y\}}(\sigma_{\{x,y\}} \mid \theta_{\{x,y\}^c})}{\gamma_{\{x,y\}}(\theta_{\{x,y\}} \mid \theta_{\{x,y\}^c})}\right] \left[\frac{\gamma_{\{x\}}(\sigma_x \mid \theta_{\{x\}^c})}{\gamma_{\{x\}}(\theta_x \mid \theta_{\{x\}^c})}\right]^{-1} \left[\frac{\gamma_{\{y\}}(\sigma_y \mid \theta_{\{y\}^c})}{\gamma_{\{y\}}(\theta_y \mid \theta_{\{y\}^c})}\right]^{-1}
\end{aligned} \tag{4.42}$$

We begin to see alternating $+1$ and -1 exponents. To confirm this feature, let's work out the term corresponding to a three-site region $A = \{x_1, x_2, x_3\}$. As the Hamiltonian with θ external conditions is the sum of the three-site interaction plus all the two-site and one-site terms, we obtain

$$\mathrm{e}^{-\phi_A(\sigma_A)} \;=\; \frac{\gamma_A(\sigma_A \mid \theta_{A^c})}{\gamma_A(\theta_A \mid \theta_{A^c})} \times \prod_{\substack{B\subset A\\ |B|=2}} \mathrm{e}^{\phi_B(\sigma_B)} \times \prod_{x\in A} \mathrm{e}^{\phi_{\{x\}}(\sigma_x)} \tag{4.43}$$

which, by (4.40) and (4.42), implies

$$\mathrm{e}^{-\phi_A(\sigma_A)} \;=\; \left[\frac{\gamma_A(\sigma_A \mid \theta_{A^c})}{\gamma_A(\theta_A \mid \theta_{A^c})}\right] \left[\prod_{\substack{B\subset A\\ |B|=2}} \frac{\gamma_B(\sigma_B \mid \theta_{B^c})}{\gamma_B(\theta_B \mid \theta_{B^c})}\right]^{-1} \left[\frac{\gamma_{\{x\}}(\sigma_x \mid \theta_{\{x\}^c})}{\gamma_{\{x\}}(\theta_x \mid \theta_{\{x\}^c})}\right]. \tag{4.44}$$

We are ready to propose an inductive formula: If $A \Subset \mathbb{L}$,

$$\mathrm{e}^{-\phi_A(\sigma_A)} \;=\; \left[\prod_{\substack{B\subset A\\ B\neq\emptyset}} \frac{\gamma_B(\sigma_B \mid \theta_{B^c})}{\gamma_B(\theta_B \mid \theta_{B^c})}\right]^{(-1)^{|A\setminus B|}}. \tag{4.45}$$

Its log leads us to the following definition.

Definition 4.46 *Let $\theta \in \Omega$ and let Π be a specification with densities $\{\gamma_\Lambda : \Lambda \Subset \mathbb{L}\}$ that is non-null in the direction θ. The θ-***vacuum potential** *for Π is the interaction defined by*

$$\phi_A^{\gamma,\theta}(\sigma_A) \;=\; -\sum_{\substack{B\subset A\\ B\neq\emptyset}} (-1)^{|A\setminus B|} \log\left[\frac{\gamma_B(\sigma_B \mid \theta_{B^c})}{\gamma_B(\theta_B \mid \theta_{B^c})}\right] \tag{4.47}$$

for each $A \Subset \mathbb{L}$ and each $\sigma \in \Omega$.

The proof that, indeed, such a potential gives us back the original densities turns out to be a simple application of Möbius transforms.

Theorem 4.48 *If a specification Π is non-null at $\theta \in \Omega$, then the vacuum potential (4.45) verifies*

$$\sum_{\substack{B\subset\Lambda\\ B\neq\emptyset}} \phi_B^{\gamma,\theta}(\sigma_B) \;=\; -\log\left[\frac{\gamma_\Lambda(\sigma_\Lambda \mid \theta_{\Lambda^c})}{\gamma_\Lambda(\theta_\Lambda \mid \theta_{\Lambda^c})}\right] \tag{4.49}$$

and, thus, its densities with external condition asymptotically equal to θ can be written as Boltzmann weights for $\Phi^{\gamma,\theta}$:

$$\gamma_\Lambda(\sigma_\Lambda \mid \omega_{\Gamma\setminus\Lambda}\,\theta_{\Gamma^c}) \;=\; \gamma_\Lambda^{\Phi^{\gamma,\theta}}(\sigma_\Lambda \mid \omega_{\Gamma\setminus\Lambda}\,\theta_{\Gamma^c}) \tag{4.50}$$

for all $\Lambda \subset \Gamma \Subset \mathbb{L}$ and all $\sigma, \omega \in \Omega$.

Proof. Due to the bar-displacement property (3.31), it is enough to prove (4.50) for $\omega = \theta$ (recall that non-nullness at θ implies non-nullness at configurations asymptotically equal to θ). In this case it is clear that (4.49) implies (4.50), because the normalization of the densities then yields the normalization (4.41). But the equivalence between (4.47) and (4.49), supplemented with the conventions $\gamma_\emptyset = 1$ and $\phi_\emptyset^{\gamma,\theta} = 0$, is a particular case of the following well-known result. $\square$

Theorem 4.51 (Möbius transform) *Let $\mathcal{E}$ be a finite set, $\mathcal{F}$ a commutative group and F and G functions from the subsets of $\mathcal{E}$ to $\mathcal{F}$. We write $F = (F_A)_{A\subset\mathcal{E}}$, $G = (G_A)_{A\subset\mathcal{E}}$. Then,*

$$\left[\forall\, A \subset \mathcal{E}\;,\; F_A = \sum_{B\subset A} (-1)^{|A\setminus B|} G_B\right] \quad\Longleftrightarrow\quad \left[\forall\, A \subset \mathcal{E}\;,\; G_A = \sum_{B\subset A} F_B\right]. \tag{4.52}$$

Let us discuss its (elementary) proof. The argument will be useful to extract other properties of the vacuum potential. It all follows from the following, equally elementary, lemma:

Lemma 4.53 *Let E be any non-empty finite set. Then*

$$\sum_{D\subset E} (-1)^{|E|} \;=\; 0\;. \tag{4.54}$$

Proof. Let us choose some $x \in E$ and decompose

$$\sum_{D\subset E} (-1)^{|E|} \;=\; \sum_{\substack{D\subset E\\ x\in D}} (-1)^{|D|} + \sum_{\substack{C\subset E\\ x\notin C}} (-1)^{|C|}\;. \tag{4.55}$$

The substitution $D = \{x\} \cup C$ shows that both term cancel out. $\square$

Proof of Theorem 4.51. Necessity:

$$\sum_{B\subset A} F_B \;= \sum_{B\subset A}\sum_{C\subset A} (-1)^{|A\setminus C|} G_C \;=\; \sum_{C\subset A} G_C \sum_{D\subset A\setminus C} (-1)^{|D|} \;= G_A\;. \tag{4.56}$$

Sufficiency:

$$\sum_{B\subset A} (-1)^{|A\setminus B|} G_B \;=\; \sum_{B\subset A} (-1)^{|A\setminus B|} \sum_{C\subset B} F_C \;=\; \sum_{C\subset A} F_C \sum_{D\subset A\setminus C} (-1)^{|D|} \;= F_A\;. \tag{4.57}$$

In both lines, the second equality follows from the substitution $D = B \setminus C$ and the last one from the previous lemma. $\square$

4.3.2 Summability of the vacuum potential

In order to pass to the limit $\Gamma \to \mathbb{L}$ in (4.49) we need to verify that the vacuum potential is summable in some sense. Of course, that requires suitable properties of the specification. As a warm-up, let us verify that Markovianness implies finite range.

Theorem 4.58 *Let Π be a specification that is non-null and r-Markovian in direction $\theta \in \Omega$. Then the range of the θ-vacuum potential does not exceed r.*

Proof. To simplify the writing we adopt the conventions $\gamma_\emptyset = 1$, $\phi_\emptyset^{\gamma,\theta} = 0$. Let $A \Subset \mathbb{L}$ be a set of sites with diameter exceeding r, and let $x, y \in A$ such that $|x - y| > r$. We decompose the sum defining the vacuum potential in (4.47) according to the location of x and y:

$$\phi_A^{\gamma,\theta}(\sigma_A) \;=\; -\Big[\sum_{\substack{B\subset A\\ B\ni x,y}} + \sum_{\substack{B\subset A\\ B\ni x\,,B\not\ni y}} + \sum_{\substack{B\subset A\\ B\ni y\,,\,B\not\ni x}} + \sum_{\substack{B\subset A\\ B\not\ni x,y}}\Big](-1)^{|A\setminus B|}\log\Big[\frac{\gamma_B(\sigma_B \mid \theta_{B^c})}{\gamma_B(\theta_B \mid \theta_{B^c})}\Big]\,. \tag{4.59}$$

In the first three sums, let us respectively substitute $C = B \setminus \{x, y\}$, $C = B \setminus \{x\}$ and $C = B \setminus \{y\}$. Alternating signs appear which leads to

$$\begin{aligned}\phi_A^{\gamma,\theta}(\sigma_A) \;&=\; -\sum_{C\subset A\setminus\{x,y\}} (-1)^{|A\setminus C|}\;\log\Big[\frac{\gamma_{C\cup\{x,y\}}(\sigma_C\,\sigma_x\,\sigma_y \mid \theta)}{\gamma_{C\cup\{x,y\}}(\theta_C\,\theta_x\theta_y \mid \theta)}\\ &\qquad\times\frac{\gamma_{C\cup\{x\}}(\theta_C\,\theta_x \mid \theta)}{\gamma_{C\cup\{x\}}(\sigma_C\,\sigma_x \mid \theta)}\,\frac{\gamma_{C\cup\{y\}}(\theta_C\,\theta_y \mid \theta)}{\gamma_{C\cup\{y\}}(\sigma_C\,\sigma_y \mid \theta)}\,\frac{\gamma_C(\sigma_C \mid \theta)}{\gamma_C(\theta_C \mid \theta)}\Big]\end{aligned} \tag{4.60}$$

We displace the bar in the last three ratios, thanks to (3.31), so as to incorporate the whole set $C \cup \{x, y\}$ inside the conditioning. All the terms $\gamma_{C\cup\{x,y\}}(\theta_C\,\theta_x\theta_y \mid \theta)$ cancel out and we obtain

$$\begin{aligned}\phi_A^{\gamma,\theta}(\sigma_A) \;&=\; -\sum_{C\subset A\setminus\{x,y\}} (-1)^{|A\setminus C|}\;\log\Big[\frac{\gamma_{C\cup\{x,y\}}(\sigma_C\,\sigma_x\,\sigma_y \mid \theta)}{\gamma_{C\cup\{x,y\}}(\sigma_C\,\sigma_x\,\theta_y \mid \theta)}\,\frac{\gamma_{C\cup\{x,y\}}(\sigma_C\,\theta_x\,\theta_y \mid \theta)}{\gamma_{C\cup\{x,y\}}(\sigma_C\,\theta_x\,\sigma_y \mid \theta)}\Big]\\ &=\; -\sum_{C\subset A\setminus\{x,y\}} (-1)^{|A\setminus C|}\;\log\Big[\frac{\gamma_{\{y\}}(\sigma_y \mid \sigma_C\,\sigma_x\,\theta)}{\gamma_{\{y\}}(\theta_y \mid \sigma_C\,\sigma_x\,\theta)}\,\frac{\gamma_{\{y\}}(\theta_y \mid \sigma_C\,\theta_x\,\theta)}{\gamma_{\{y\}}(\sigma_y \mid \sigma_C\,\theta_x\,\theta)}\Big]\,,\end{aligned} \tag{4.61}$$

where we have used (3.31) again in each ratio. But the r-Markovianness hypothesis implies that $\gamma_{\{y\}}(\cdot \mid \sigma_C\,\sigma_x\,\theta)$ equals $\gamma_{\{y\}}(\cdot \mid \sigma_C\,\theta_x\,\theta)$, thus the argument of the logarithm is equal to one. This implies $\phi_A^{\gamma,\theta} = 0$. $\square$

We see that in the proof, Markovianness is used only at the level of single-site densities. This is, of course, not a surprise in view of Proposition (4.27). As mentioned above, this theorem (in its "directionless" version) is associated to a number of known probabilists —Averintsev, Spitzer, Hammersley and Clifford, Preston, and Grimmett. Historical notes can be found in the introduction to the last author's contribution [23], which is also the genesis for the preceding proof. The strategy of this proof can be used to prove the first of the following overdue observations.

Exercise 4.62

(i) Prove that $\Phi^{\gamma,\theta}$ is indeed a vacuum potential, that is, prove that it satisfies property (4.36).

(ii) Formalize the obvious fact that a θ-vacuum potential is unique.

It is even easier to prove a similar theorem but with *Markovian* replaced by *quasilocal*, We only need the following identity. If $\Lambda \subset \widetilde{\Lambda} \Subset \mathbb{L}$ and $\sigma, \omega \in \Omega$,

$$\begin{aligned} \log\left[\frac{\gamma_\Lambda(\omega_\Lambda \mid \omega_{\widetilde{\Lambda}\setminus\Lambda}\,\theta)}{\gamma_\Lambda(\theta_\Lambda \mid \omega_{\widetilde{\Lambda}\setminus\Lambda}\,\theta)}\right] &= \log\left[\frac{\gamma_{\widetilde{\Lambda}}(\omega_{\widetilde{\Lambda}} \mid \theta)}{\gamma_{\widetilde{\Lambda}}(\theta_\Lambda\,\omega_{\widetilde{\Lambda}\setminus\Lambda} \mid \theta)}\right] \\ &= \log\left[\frac{\gamma_{\widetilde{\Lambda}}(\omega_{\widetilde{\Lambda}} \mid \theta)}{\gamma_{\widetilde{\Lambda}}(\theta_{\widetilde{\Lambda}} \mid \theta)}\right] - \log\left[\frac{\gamma_{\widetilde{\Lambda}\setminus\Lambda}(\omega_{\widetilde{\Lambda}\setminus\Lambda} \mid \theta)}{\gamma_{\widetilde{\Lambda}\setminus\Lambda}(\theta_{\widetilde{\Lambda}\setminus\Lambda} \mid \theta)}\right] \qquad (4.63) \\ &= \sum_{\substack{B\cap\Lambda\neq\emptyset \\ B\subset\widetilde{\Lambda}}} \phi_B^{\gamma,\theta}(\omega)\,, \end{aligned}$$

$$(4.64)$$

where the first two equalities follow from the bar-displacement property (3.31) and the last one from (4.49). This immediately implies the following theorem.

Theorem 4.65 *Let Π be a specification that is non-null at ω and $\theta \in \Omega$ and quasilocal at ω in direction θ. Then its θ-vacuum potential is summable at ω. In fact,*

$$H_\Lambda^{\Phi^{\gamma,\theta}}(\sigma_\Lambda \mid \omega_{\Lambda^c}) \;=\; -\lim_{n\to\infty} \log\left[\frac{\gamma_\Lambda(\omega_\Lambda \mid \omega_{\Lambda_n\setminus\Lambda}\,\theta)}{\gamma_\Lambda(\theta_\Lambda \mid \omega_{\Lambda_n\setminus\Lambda}\,\theta)}\right] \qquad (4.66)$$

for every $\Lambda \Subset \mathbb{L}$ and $\sigma_\Lambda \in \Omega_\Lambda$, and, thus, the densities of Π with external condition ω can be written as Boltzmann weights:

$$\gamma_\Lambda(\,\cdot \mid \omega_{\Lambda^c}) \;=\; \gamma_\Lambda^{\Phi^{\gamma,\theta}}(\,\cdot \mid \omega_{\Lambda^c}) \qquad (4.67)$$

for all $\Lambda \subset \Gamma \Subset \mathbb{L}$.

4.3.3 Kozlov theorem

Gibbsianness requires uniform and absolute summability of the interaction. Absolute summability seems, in principle, not to be much of a problem. Indeed, due to our freedom to pass to physically equivalent interactions, we can use partial sums to define an equivalent, absolutely convergent interaction. There is, however, a rather subtle obstacle to this strategy (I owe this observation to Frank Redig): If we do not insist on uniformity, the resummation procedure becomes ω-dependent, and it is not clear whether the resulting potential would remain *measurable*. Therefore, we shall combine, from the outset, absoluteness with uniformity, that is, we shall place absolute value and "sup" signs all over the place. Our hypotheses will be accordingly strengthened: We shall now assume quasilocality (that is, uniform continuity) and (uniform) non-nullness.

Non-nullness implies (is equivalent to) the strict positivity of the numbers

$$m_x \;=\; \inf_\omega \gamma_{\{x\}}(\omega_x \mid \omega)\,. \qquad (4.68)$$

for all $x \in \mathbb{L}$. Quasilocality says that for each $x \in \mathbb{L}$ the function

$$g_x(r) \;=\; \sup_\omega \Big|\gamma_{\{x\}}(\omega_x \mid \omega) - \gamma_{\{x\}}(\omega_x \mid \omega_{\Lambda_r}\,\theta)\Big| \qquad (4.69)$$

converges to zero, as $r \to \infty$ (in the presence of non-nullness such a condition is equivalent to quasilocality).

To understand the basic algorithm to pass from a vacuum potential to an absolute and uniformly summable one, let us first discuss how to gain summability for bonds containing the origin. We resort to the inequality

$$|\ln a - \ln b| \;\leq\; \frac{|a-b|}{\min(a,b)} \tag{4.70}$$

valid for $a, b > 0$ (the proof is immediate from the integral definition of the logarithm), to obtain, from (4.66), the bound

$$\sup_\omega \Bigl| \sum_{\substack{B \ni 0 \\ B \subset \Lambda_r}} \phi_B^{\gamma,\theta}(\omega) \Bigr| \;\leq\; \frac{g_0(r)}{m_0} \;. \tag{4.71}$$

As $g_0(r) \to 0$, we can choose a sequence of integers r_i, $i = 1, 2, \ldots$ diverging with i and such that

$$\sum_i g_0(r_i) \;<\; \infty \;. \tag{4.72}$$

The idea is now to group the bonds within the regions

$$L_i^0 = \Lambda_{r_i} \;,\; i = 1, 2, \ldots \tag{4.73}$$

that is, within the families

$$S_1^0 = \bigl\{ B \subset L_1^0 \bigr\} \;,\; \ldots S_i^0 = \bigl\{ B \subset L_i^0 : 0 \in B \bigr\} \setminus S_{i-1}^0 \;,\; \ldots \tag{4.74}$$

The interaction

$$\varphi_A \;=\; \begin{cases} 0 & \text{unless } A = L_i^0 \text{ for some } i \geq 1 \\[2mm] \displaystyle\sum_{B \in S_i^0} \phi_B^{\gamma,\theta} & \text{if } A = L_i^0 \;, \end{cases} \tag{4.75}$$

is physically equivalent to the θ-vacuum potential $\Phi^{\gamma,\theta}$ and by (4.71)

$$\begin{aligned} \sup_\omega \Bigl| \varphi_{L_i^0}(\omega) \Bigr| \;&=\; \sup_\omega \Bigl| \sum_{\substack{B \in \Lambda_{r_i} \\ B \ni 0}} \phi_B^{\gamma,\theta}(\omega) - \sum_{\substack{B \in \Lambda_{r_{i-1}} \\ B \ni 0}} \phi_B^{\gamma,\theta}(\omega) \Bigr| \\ &\leq\; \frac{g_0(r_i) + g_0(r_{i-1})}{m_0} \end{aligned} \tag{4.76}$$

$[g_0(r_0) \equiv 0]$. Therefore, by (4.72),

$$\sum_{A \ni 0} \|\varphi_A\|_\infty \;\leq\; \frac{2}{m_0} \sum_{i \geq 1} g_0(r_i) \;<\; \infty \;. \tag{4.77}$$

To obtain an analogous summability around every site x, the preceding strategy has to be pursued so as to visit, in some fixed order, the different sites of the lattice while grouping the relevant bonds as above, taking care of counting each bond only once. In this consists the proof of the following crucial theorem.

Theorem 4.78 (Kozlov [27]) *A specification is Gibbsian if, and only if, it is uniformly non-null and quasilocal.*

Proof. We only need to prove sufficiency, as necessity has been proven in Proposition (4.24). Let us choose a vacuum $\theta \in \Omega$ (any choice will do), and consider the corresponding θ-vacuum potential. We fix an order for the sites of the lattice, $\mathbb{L} = \{x_1, x_2, \ldots\}$ and choose sequences r_i^ℓ, $i, \ell = 1, 2, \ldots$ such that

$$\sum_i g_{x_\ell}(r_i^\ell) \;<\; \infty \tag{4.79}$$

for each $x_\ell \in \mathbb{L}$ [the functions g_x have been defined in (4.69)]. We then choose "rectangles" around each of the sites x_ℓ: For $i, \ell = 1, 2, \ldots$,

$$L_i^1 = \{x_j : 1 \le j \le \widetilde{r}_i^1\}\;,\; L_i^2 = \{x_j : 2 \le j \le \widetilde{r}_i^2\}\;, \ldots\;, L_i^\ell = \{x_j : \ell \le j \le \widetilde{r}_i^\ell\}\;,\; \ldots \tag{4.80}$$

where the $\widetilde{r}_i^\ell$ are chosen so that

$$r_i^\ell \;=\; \operatorname{diam} L_i^\ell\;, \tag{4.81}$$

and we assign each bond, in a unique way, to one of such rectangles by defining

$$S_i^\ell = \big\{B \subset L_i^\ell : x_\ell \in B\big\} \setminus S_{i-1}^\ell \tag{4.82}$$

for $i, \ell = 1, 2, \ldots$ ($S_0^\ell \equiv 0$). We observe that:

(F1) the families S_i^ℓ are disjoint, and

(F2) If $B \ni x_\ell$, then $B \in \cup_{j=1}^{\ell} \cup_{i\ge 1} S_i^j$.

Finally we define

$$\varphi_A \;=\; \begin{cases} 0 & \text{unless } A = L_i^\ell \text{ for some } i, \ell \ge 1 \\[2mm] \displaystyle\sum_{B \in S_i^\ell} \phi_B^{\gamma,\theta} & \text{if } A = L_i^\ell\,. \end{cases} \tag{4.83}$$

As for (4.76),

$$\sup_\omega \Big|\varphi_{L_i^\ell}(\omega)\Big| \;\le\; \frac{g_{x_\ell}(r_i^\ell) + g_{x_\ell}(r_{i-1}^\ell)}{m_{x_\ell}}\;, \tag{4.84}$$

hence, by observation (F2) above,

$$\begin{aligned} \sum_{A \ni x_\ell} \|\varphi_A\|_\infty \;&\le\; \sum_{j=1}^{\ell} \sum_{i \ge 1} \Big\|\varphi_{L_i^j}\Big\|_\infty \\ &\le\; \sum_{j=1}^{\ell} \frac{2}{m_{x_j}} \sum_{i \ge 1} g_{x_j}(r_i^j) \end{aligned} \tag{4.85}$$

which is finite by (4.79). $\square$

The interaction φ constructed in the preceding proof is no longer a vacuum potential, and furthermore, its summability bound worsens with the order ℓ of the site. So, there is no hope of proving site-uniformly summability, that is with a supremum over x in (4.9). Another particularly annoying feature of the proof is that a translation-invariant specification does not lead to a translation-invariant interaction. The algorithm can be modified, by fixing the radii r_i^ℓ in a ℓ-independent fashion, so as to produce a translation-invariant potential. But summability is recovered only if the continuity-rate function g_0 decreases at sufficient speed. It is not known whether this extra condition is only technical (the suspicion is that it is not). There is an alternative Gibbsianness theorem by Sullivan [58] which has the advantage of yielding translation invariance without additional hypotheses. But this theorem refers to a space of interactions different from $\mathcal{B}_1$, and is thus slightly less adapted to current Gibbsian theory.

4.4 Less Gibbsian measures

Kozlov theorem leaves us with a rather simple symptomatology of non-Gibbsianness, based on only two properties. While non-nullness is not a property to be ignored, it is not usually the main problem. Furthermore, already models with exclusions and grammars have given us some familiarity with the effects of its absence. The absence of quasilocality, on the other hand, leads to more subtle, or at least less familiar, phenomena. In physical terms, the conditions of Definition (4.20) correspond to situations in which the intermediate configuration ω effectively shields the interior of the region Λ from the influence of far away regions. The failure of such type of properties would place us in an extremely unphysical situation, as it would correspond to the uncontrollability of local experiments. Mathematically, non-quasilocality causes the breakdown of proofs of a number of important properties that are behind our understanding of phase diagrams and properties of the extremal phases.

For these reasons there has been a systematic effort to determine a *taxonomy* of non-quasilocal measures, with the hope that of restoring, within each category, a different set of Gibbsian properties. While this hope has been only partially realized, the classification scheme is well established by now. To present it we need some notation.

Definition 4.86 *For a specification* Π *on* $(\Omega, \mathcal{F})$ *and* $\theta \in \Omega$*, let us denote*

$$\Omega_{\mathrm{q}}^{\theta}(\Pi) \;=\; \Big\{ w \in \Omega : \Pi \textit{ is quasilocal at } \omega \textit{ in the direction } \theta \Big\} \tag{4.87}$$

$$\Omega_{\mathrm{q}}(\Pi) \;=\; \Big\{ w \in \Omega : \Pi \textit{ is quasilocal at } \omega \Big\} \tag{4.88}$$

and, for an interaction Φ*, let us recall the set* $\Omega_{\mathrm{sum}}^{\Phi}$ *given in Definition 4.3. Then, a probability measure* μ *on* $(\Omega, \mathcal{F})$ *is*

- **quasilocal** *if it is consistent with a quasilocal specification,*
- **almost quasilocal** *or* **almost Gibbs** *if it is consistent with a specification* Π *such that*

$$\mu\big[\Omega_{\mathrm{q}}(\Pi)\big] \;=\; 1\;, \tag{4.89}$$

- **intuitively weakly Gibbs** *if it is consistent with a specification* Π *for which there exist a set* $\Omega_{\mathrm{reg}}(\Pi)$ *such that*

$$\mu\big[\Omega_{\mathrm{reg}}(\Pi)\big] \;=\; 1 \quad \textit{and} \quad \omega \in \Omega_{\mathrm{q}}^{\theta}(\Pi)\;,\; \forall\, \omega, \theta \in \Omega_{\mathrm{reg}}(\Pi)\;, \tag{4.90}$$

- **weakly Gibbs** *if it is consistent with a specification* Π*, with density functions* $\{\gamma_\Lambda : \Lambda \Subset \mathbb{L}\}$*, for which there exists an interaction* Φ *such that*

$$\mu\big[\Omega_{\mathrm{sum}}^{\Phi}\big] = 1 \quad \textit{and} \quad \gamma_\Lambda(\,\cdot \mid \omega) \;=\; \gamma_\Lambda^{\Phi}(\,\cdot \mid \omega)\;,\; \forall\, \omega \in \Omega_{\mathrm{sum}}^{\Phi}\;. \tag{4.91}$$

[In our setting, (almost) quasilocality = (almost) Feller.] Weakly Gibbs measures arose from an effort to extend an interaction-based description of non-Gibbsian measures. In contrast, almost quasilocality ignores the Boltzmann prescription and focuses on specification properties. Nevertheless, due to Theorem 4.65 both almost quasilocal and intuitively weakly Gibbs measures are weakly Gibbs as well. The configurations in $\Omega_{\rm reg}$ are the *regular points* of the corresponding interaction. We refer the reader to [67, 69] for a comparison among the different notions. I content myself with the following remarks summarizing known facts.

Proposition 4.92 *Let $\mu \in \mathcal{P}(\Omega, \mathcal{F})$.*

(i) If μ is consistent with a specification Π and there exists a $\theta \in \Omega$ such that

$$\mu\big[\Omega_{\rm q}^{\theta}(\Pi)\big] \;=\; 1 \tag{4.93}$$

then μ is weakly Gibbs.

(ii) If μ is intuitively weakly Gibbs, then it is consistent with a specification Π such that

$$\mu\Big\{\theta \in \Omega : \mu\big[\Omega_{\rm q}^{\theta}(\Pi)\big] = 1\Big\} \;=\; 1 \;, \tag{4.94}$$

The first item follows from Theorem 4.65, the second one is immediate from the definition of IWG measures. The opposite implications in both items are probably false.

In [69] the following inclusions have been pointed out:

$$\mathrm{G} \underset{\neq}{\subset} \mathrm{AQL} \underset{\neq}{\subset} \mathrm{IWG} \subset \mathrm{WG} \underset{\neq}{\subset} \mathcal{P}(\Omega, \mathcal{F}) \;, \tag{4.95}$$

where the acronyms represent the obvious families of measures. Examples of measures that are almost quasilocal but not Gibbsian include the random-cluster model when there is an almost surely unique infinite cluster [53, 24], the modified "avalanche" model of [41], the sign-fields of the SOS model [68] and the Grising random field [67] below the critical value of site-percolation. Measures that are intuitively weakly Gibbs but not almost quasilocal are constructed in [69]. They include measures absolutely continuous with respect to a product of Bernoulli measures on the positive integers and the invariant measure for the Manneville-Pomeau map whose non-Gibbsianness was determined in [42]. In this last example discontinuities appear together with lack of non-nullness. The only known example of a probability measure that is not even weakly Gibbs is the avalanche model worked out in [41]. On the other hand, the inclusions AQL $\subset$ WG and AQL $\subset \mathcal{P}(\Omega, \mathcal{F})$ are rather strict. Indeed, convex combinations of Gibbs measures for different potentials are quasilocal at no configuration [64], and measures associated to dependent (Fortuin-Kasteleyn) percolation on trees have discontinuities at a set of full-measure [25]. It is not known whether these measures are weakly Gibbsian. The combinations of Bernoulli measures with different densities, studied in [41] are such that there exists no specification Π and no configuration θ for which (4.93) is true. But this falls short of showing lack of weak Gibbsianness. There are, on the other hand, examples of measures associated to disordered systems (see Section 5.6) that are weakly Gibbsian but almost surely *not* quasilocal [28].

The proof that a measure is weakly Gibbs involves sophisticated techniques, usually coarse-graining arguments combined with cluster expansions. Nevertheless, practically all known examples of non-Gibbsian measures have been proven to be weakly Gibbsian [5, 44, 7, 8, 1, 2, 29]. In fact, if you allow me to play with words, this proven weak Gibbsianness turns out often to be rather strong in that it is associated to *absolutely* summable interactions, that, moreover, decay at a (configuration-dependent) exponential rate. Nevertheless, the existence of these strong weak potentials seems to be too weak a condition to restore useful Gibbsian properties. In particular, only very limited results hold [37, 33] regarding the extension of the variational approach to these measures.

In contrast, much more of the variational approach can be restored for almost quasilocal measures. This has been done, through relatively simple proofs [11, 10] —no coarse graining, no expansion—, in cases where FKG monotonicity can be invoked. The argument shows at the same time that some of the weak-Gibbsian measures cited above are in fact almost quasilocal. The discussion in [69] strongly indicates that these good variational-approach results may extend to the larger class of intuitively weakly measures.

The best description of the differences between the classes introduced above is contained in a remark in [69]:

- For a quasilocal measure, *every* configuration shields a finite region from *every* far away influence.
- For an almost quasilocal measure, *almost every* configuration shields a finite region from *every* far away influence.
- For an intuitively Gibbs measure, *almost every* configuration shields a finite region from *almost every* far away influence.

In practical terms, the difference between *every* and *almost every* seems impossible to detect as it refers to events that will never be measured or appear in a simulation. Nevertheless, these differences show up through distinctive mathematical properties. This contrast explains the challenge posed by the study of non-Gibbsian measures.

5 What it takes to be non-Gibbsian

5.1 Linear transformations of measures

Most of the instances of non-Gibbsianness discussed in the literature refer to measures obtained as transformations of Gibbs measures through probability kernels as defined in (3.5)/(3.6). The only exceptions are the joint measures for disordered systems. briefly presented in Section 5.6. The setting is, then, defined by a probability kernel τ from one configuration space $(\Omega = S^{\mathbb{L}}, \mathcal{F})$ to another, possibly different, space $(\Omega' = S'^{\mathbb{L}'}, \mathcal{F}')$. Non-Gibbsian studies focus on three types of measures obtained from a measure $\mu \in \mathcal{P}(\Omega, \mathcal{F})$:

(NG1) The measure obtained after a single transformation, $\mu' = \mu\tau$.

(NG2) Measures obtained after a number (sufficiently small or sufficiently large) of iterations of the transformation, $\mu^{(n)} = \mu\tau^n = (\mu\tau^{n-1})\tau$.

(NG3) Measures obtained through an infinite iteration of the transformation or invariant measures: $\mu^\infty = \lim_n \mu\tau^n$ or μ such that $\mu = \mu\tau$

[For the last two types, $(\Omega', \mathcal{F}') = (\Omega, \mathcal{F})$.]

The kernels of interest here have all a product structure

$$\tau(d\omega' \mid \omega) \;=\; \prod_{x'\in\mathbb{L}'} \tau_{x'}(d\omega'_{x'} \mid \omega) \;, \tag{5.1}$$

where each $\tau_{x'}(\,\cdot \mid \omega)$ is measure on S', and hence defined by a density

$$T_{x'}(\omega_{x'} \mid \omega) \;=\; \tau_{x'}(\{\omega'_{x'}\} \mid \omega) \;. \tag{5.2}$$

Hence, the transformed measure of $\mu \in \mathcal{P}(\Omega, \mathcal{F})$ satisfies for the weight of a cylinder

$$\mu'(C_{\omega'_{\Lambda'}}) \;=\; \int_\Omega \prod_{x'\in\Lambda'} T_{x'}(\omega'_{x'} \mid \omega)\, \mu(d\omega) \;, \tag{5.3}$$

for any $\Lambda' \Subset \mathbb{L}'$ and $\omega'_{\Lambda'} \in \Omega'_{\Lambda'}$. In particular, a deterministic transformation is defined by *functions* $T_{x'} : \Omega \to S'$ such that

$$T_{x'}(\omega_{x'} \mid \omega) \;=\; \begin{cases} 1 & \text{if } \omega'_{x'} = T_{x'}(\omega) \\ 0 & \text{otherwise} \end{cases} \tag{5.4}$$

The transformations used in physics and probability can be classified into various categories:

- A **block transformation** is such that for each $x' \in \mathbb{L}'$ there exists a *block* $B_{x'} \Subset \mathbb{L}$ such that $T_{x'}(\omega'_{x'} \mid \cdot\,) \in \mathcal{F}_{B_{x'}}$ for all $\omega'_{x'} \in S'$. Hence $T_{x'}(\omega_{x'} \mid \cdot\,)$ equals a function on $\mathcal{F}_{B_{x'}}$ which I will denote with the same symbol, namely $T_{x'}(\omega'_{x'} \mid \omega_{B_{x'}})$.
- In general terms, **renormalization transformations** are characterized by at least one of the following properties:
  - The blocks $B_{x'}, x' \in \mathbb{L}'$, form a partition of $\mathbb{L}$ (that is, they are disjoint and their union is the whole of $\mathbb{L}$).
  - The functions $T(\omega'_{x'} \mid \cdot\,)$ are continuous for all $\omega'_{x'} \in S'$.
- Transformations with overlapping blocks are typical of **stochastic evolutions**. These include **cellular automata** (discrete-time) and **spin-flip** (continuous-time) dynamics.

Here are a few examples of renormalization transformations that have played a benchmark role in non-Gibbsian studies and Gibbs-restoration projects. If necessary, the reader can suppose that $\mathbb{L} = \mathbb{Z}^d$ and $\mathbb{L}' = \mathbb{Z}^{d'}$ but, of course, lattices with a notion of $\mathbb{Z}^d$-translations (=action of $\mathbb{Z}^d$ by isomorphisms) do equally well.

Deterministic block renormalization transformations:

b^d-Decimation: $\mathbb{L}' = \mathbb{L}$, $S' = S$, $B_{x'} = \Lambda_{b-1} + bx'$, $T_{x'}(\omega_{B_{x'}}) = \omega_{bx'}$.

Spin contractions: $\mathbb{L}' = \mathbb{L}$, $S' \underset{\neq}{\subset} S$, $B_{x'} = \{x'\}$; two species:

- *Sign fields:* $S \subset \mathbb{R}$ symmetric, $T_{x'}(\omega_{x'}) = \mathrm{sign}(\omega_{x'})$.
- *"Fuzzy" spins:* $S = \cup_{i\in I} S_i$ (partition), $S' = I$, $T_{x'}(\omega_{x'}) = i$ if $\omega_{x'} \in S_i$.

Block average: $\mathbb{L}' = \mathbb{L}$, $S' \underset{\neq}{\supset} S$, $T_{x'}(\omega_{B_{x'}}) = |B_{x'}|^{-1} \sum_{y\in B_{x'}} \omega_y$.

Majority rule (odd blocks): $\mathbb{L}' = \mathbb{L}$, $S' = S = \{-1, 1\}$, ($|B_{x'}|$ odd), $T_{x'}(\omega_{B_{x'}}) = \mathrm{sign}\big[\sum_{y\in B_{x'}} \omega_y\big]$.

Stochastic block renormalization transformations:

Majority rule (even blocks): $\mathbb{L}' = \mathbb{L}$, $S' = S = \{-1, 1\}$, ($|B_{x'}|$ even), $\omega_{B_{x'}} = \mathrm{sign}\big[\sum_{y \in B_{x'}} \omega_y\big]$ if this last sum is non-zero, and $+1$ or -1 with probability $1/2$ otherwise. [*Exercise:* write the kernel densities $T_{x'}(\omega'_{x'} \mid \omega_{B_{x'}})$.]

p-Kadanoff transformation: $\mathbb{L}' = \mathbb{L}$, $S' = S$,

$$T_{x'}(\omega'_{x'} \mid \omega_{B_{x'}}) \;=\; \frac{\exp\Big[p\, \omega'_{x'} \sum_{y \in B_{x'}} \omega_y\Big]}{\text{Norm.}}\;. \tag{5.5}$$

Non-block renormalization transformations (deterministic):

Projections: $\mathbb{L}' \underset{\neq}{\subset} \mathbb{L}$, $S' = S$, $T(\omega) = \omega_{\mathbb{L}'}$. This is a generalization of decimation. The most important case is *Schonmann's example*: $S = \{-1, 1\}$, $\mathbb{L} = \mathbb{Z}^d$, $\mathbb{L}' = \mathbb{Z}^{d-1} \times \{0\}$.

Momentum transformations: $\mathbb{L}' = \mathbb{L}$, $S' \underset{\neq}{\supset} S$, $T_{x'}(\omega) = \sum_y F(x' - y)\, \omega_y$ for F summable, defined through a Fourier transform with an appropriate smooth cut-off.

The following exercise applies to all the preceding examples.

Exercise 5.6 *Let τ be either a renormalization transformation with strictly positive densities $T_{x'}(\,\cdot \mid \cdot\,)$ or a deterministic renormalization transformation such that $T_{x'}^{-1}(\omega'_{x'}) \neq \emptyset$ for all $\omega'_{x'} \in S'$.*

(i) Prove that if μ is non-null, then $\mu\tau$ gives positive measure to any cylinder $C_{\omega'_{\Lambda'}}$.

(ii) Conclude that τ maps non-null measures into non-null measures.

The situation is dramatically different if "non-null" is replaced by "Markovian": A Markovian measure, subjected to a "Markovian" (= block-renormalization) transformation may, in fact, not even be quasilocal.

5.2 Absence of uniform non-nullness

Kozlov's theorem implies two main causes of non-Gibbsianness: lack of non-nullness and lack of quasilocality. The manifestation of the former comes from the negation of the following *alignment-suppression property.*

Proposition 5.7 *If a measure μ on $(\Omega, \mathcal{F})$ is consistent with an uniformly non-null translation-invariant specification, then there exists $\delta > 0$ such that*

$$\sup_{\omega_\Lambda \in \Omega_\Lambda} \mu\big(C_{\omega_\Lambda}\big) \;\le\; \mathrm{e}^{-\delta|\Lambda|} \tag{5.8}$$

for all $\Lambda \Subset \mathbb{L}$.

Proof. Let γ_x be the single-site specification densities and $\epsilon = \inf_\sigma \gamma_{\{x\}}(\sigma_x \mid \sigma) > 0$. Then, by consistency,

$$\begin{aligned} \mu\big(C_{\omega_\Lambda}\big) \;&=\; \int \gamma_{\{x\}}(\omega_x \mid \sigma)\, \mathbf{1}_{C_{\omega_{\Lambda\setminus\{x\}}}}(\sigma)\, \mu(d\sigma) \\ &\le\; (1-\epsilon)\, \mu\big(C_{\omega_{\Lambda\setminus\{x\}}}\big)\;. \end{aligned} \tag{5.9}$$

Induction finishes the proof. $\square$

The failure of this property means that the (exponential) cost of inserting a "defect" ω_Λ is sub-volumetric. This is what happens, for instance, for some sign-field measures [35, 9, 68, 38], where there are defects that can be placed by paying only a surface-area cost. The non-Gibbsianness of the invariant measures of some cellular automata also shows up in this way. Alignment needs to appear only in some lower-dimensional manifold —a surface for the voter model [36] or the non-local dynamics of [49], a "spider" for the non-reversible automata of [15]— and the dynamics propagates it to a whole volume. In fact, as proposed in [43], the detection of this alignment propagation can be a numerical test for non-Gibbsianness. Such a test has indeed been applied [45, 46] to the invariant measure of the Toom model, with inconclusive results.

As seen in Exercise 5.6, a measure that is non-Gibbsian due to lack of non-nullness can not be the image of any non-null measure —Gibbsian or not— through strictly positive renormalization transformations.

For those that know a bit about the variational principle, I comment that (5.8) means that

$$\liminf_n \frac{-1}{|\Lambda_n|} \log \mu\big(C_{\omega_{\Lambda_n}}\big) \;\geq\; \delta \;>\; 0\,. \tag{5.10}$$

When both μ and ω are translation-invariant violation of (5.10) implies that the entropy density of the δ_ω relative to μ is zero. This is how the violation of (5.8) is usually presented, linking non-Gibbsianness to a failure of the variational principle. Furthermore, as the relative entropy is a large-deviation rate function, this failure indicates the presence of large deviations that are "too large" (its probability is penalized less than the volume exponential typical of Gibbsianness). Nevertheless, the argument based on (5.8) is more general, as it requires neither translation invariance nor the existence of relative entropy densities.

5.3 Absence of quasilocality

Let me explain, in some detail, the subtleties involved in proving that a measure μ is not quasilocal. To make the notation slightly lighter (and to acquaint the reader with yet another usual notation), let us denote by μ_Λ the kernel $\mu_{|\tau_{\Lambda^c}}$ of Definition 3.14, that is

$$\mu_\Lambda(f \mid \omega) \;=\; E_\mu(f \mid \mathcal{F}_{\Lambda^c})(\omega) \tag{5.11}$$

is a realization of the corresponding conditional expectation for bounded $f \in \mathcal{F}$, $\Lambda \Subset \mathbb{L}$ and $\omega \in \Omega$. [From our long discussion of Sections 3.2 and 3.3, the reader should retain at least this fact: conditional expectations admit an infinite number of versions (=realizations) all differing on measure-zero sets.] Let me reserve the right to denote sometimes this object as $\mu_\Lambda(f \mid \omega_{\Lambda^c})$ to emphasize its $\mathcal{F}_{\Lambda^c}$-measurability.

The measure μ is not quasilocal if it is consistent with *no* quasilocal specification. To prove this (recall that every measure is consistent with *some* specification), it is enough to find a *single*, nonremovable, point of discontinuity for a *single* μ_Λ for a *single* quasilocal f [By Proposition 4.24 this will already happen for $\Lambda = \{x\}$, $f = \mathbf{1}_{\sigma_x}$ for some $x \in \mathbb{L}$, $\sigma_x \in S$.] Let us precise this fact.

Definition 5.12 *A measure $\mu \in \mathcal{P}(\Omega, \mathcal{F})$ is* **not quasilocal at** $\omega \in \Omega$ *if there exists $\Lambda \Subset \mathbb{L}$ and f (quasi)local such that no realization of $\mu_\Lambda(f \mid \cdot)$ is quasilocal at ω.*

In other words, any realization of $\mu_\Lambda(f \mid \cdot\,)$ must exhibit an *essential discontinuity* at ω; one that survives zero-measure modifications. Let us understand what this means, for a general measurable function g. As we shall assume μ non-null (otherwise it would already be non-Gibbsian), "essential" can be associated to "supported on open sets". Thus, we are led to consider the following twin notions.

Definition 5.13 *Let g be a measurable function and μ a probability measure on $(\Omega, \mathcal{F})$. Let $\omega \in \Omega$.*

*(a) g is μ-***essentially discontinuous** *at ω if every function continuous at ω differs from g in a set of non-zero μ-measure.*

(b) g is **strongly discontinuous** *at ω if every function continuous at ω differs from g in a set having non-empty interior.*

[That is: if f is continuous at ω, then the set $\{\omega : g(\omega) \neq f(\omega)\}$ has μ-measure non-zero in (a) and contains an open set in (b).]

Remark 5.14 *If μ is non-null, every strong discontinuity is essential.*

Conditional expectations are bounded, hence they can only have jump discontinuities, caused by the presence of different limits coming from different directions. In order for such a discontinuity to be essential or strong, the set of directions from which each of the different limits is achieved should be sufficiently thick. This yields the following basic criteria.

Proposition 5.15 *Let $\mu \in \mathcal{P}(\Omega, \mathcal{F})$, g a bounded measurable function and $\omega \in \Omega$. Then g is μ-essentially discontinuous [resp. strongly discontinuous] at ω iff there exists a $\delta > 0$ such that for every neighborhood $\mathcal{N}$ of ω there exist two sets $\mathcal{N}^+$ and $\mathcal{N}^-$, with $\omega \in \mathcal{N}^\pm \subset \mathcal{N}$, such that $\mu(\mathcal{N}^\pm) > 0$ [resp. $\mathcal{N}^\pm$ open] and*

$$\bigl|g(\sigma^+) - g(\sigma^-)\bigr| > \delta \tag{5.16}$$

for every $\sigma^\pm \in \mathcal{N}^\pm$.

As the cylinders are a basis of the topology of Ω (every open set is a union of such), open neighborhoods of ω are (unions of) cylinders of the form C_{ω_Γ} for $\Gamma \Subset \mathbb{L}$. Thus, condition (5.16) is equivalent to

$$\bigl|g(\omega_{\Lambda_N}\,\sigma^+) - g(\omega_{\Lambda_N}\,\sigma^-)\bigr| > \delta \tag{5.17}$$

for N large enough, for $\sigma^\pm \in \mathcal{N}^\pm_{\Lambda_N^{\rm c}} \in \mathcal{F}_{\Lambda_N^{\rm c}}$ of non-zero measure or open, according to the case. After a little thought we see that we can rewrite Proposition 5.15 in the following equivalent form.

Proposition 5.18 *Let $\mu \in \mathcal{P}(\Omega, \mathcal{F})$, g a bounded measurable function and $\omega \in \Omega$.*

(a) g is μ-essentially discontinuous at ω iff there exist a diverging sequence $(N_i)_{i\geq 1}$ of natural numbers and real numbers δ_+ and δ_- with $\delta^+ - \delta^- > 0$ such that for each $i \geq 1$ there exist sets $\mathcal{N}_i^+, \mathcal{N}_i^- \in \mathcal{F}_{\Lambda_{N_i}^{\rm c}}$ with

$$\limsup_{i\to\infty} \mu\Bigl(g\,\mathbf{1}_{C_{\omega_{\Lambda_{N_i}}}}\,\mathbf{1}_{\mathcal{N}_i^+}\Bigr) \;>\; \delta^+ \quad \textit{and} \quad \liminf_{i\to\infty} \mu\Bigl(g\,\mathbf{1}_{C_{\omega_{\Lambda_{N_i}}}}\,\mathbf{1}_{\mathcal{N}_i^-}\Bigr) \;<\; \delta^- \tag{5.19}$$

(b) *g is strongly discontinuous at ω iff there exists a diverging sequence $(N_i)_{i\geq 1}$ of natural numbers such that for each $i \geq 1$ there exist a natural number $R_i > N_i$ and two configurations η^+, η^- such that*

$$\limsup_{i\to\infty} \Bigl| g(\omega_{\Lambda_{N_i}}\, \eta^+_{\Lambda_{R_i}\setminus\Lambda_{N_i}}\, \sigma^+) - g(\omega_{\Lambda_{N_i}}\, \eta^-_{\Lambda_{R_i}\setminus\Lambda_{N_i}}\, \sigma^-) \Bigr| \;\geq\; \delta \tag{5.20}$$

for every $\sigma^\pm \in \Omega$.

To settle our non-quasilocality issue we now apply these considerations to functions of the form $g(\,\cdot\,) = \mu_\Lambda(f \mid \cdot\,)$. From Definition 5.12 and the previous proposition we obtain:

Proposition 5.21 *Let $\mu \in \mathcal{P}(\Omega, \mathcal{F})$. Then:*

(a) *μ is not quasilocal at ω iff there exist a diverging sequence $(N_i)_{i\geq 1}$ of natural numbers and real numbers δ_+ and δ_- with $\delta^+ - \delta^- > 0$ such that for each $i \geq 1$ there exist sets $\mathcal{N}_i^+, \mathcal{N}_i^- \in \mathcal{F}_{\Lambda^{\rm c}_{N_i}}$ with*

$$\limsup_{i\to\infty} \mu\Bigl(f\, \mathbf{1}_{C_{\omega_{\Lambda_{N_i}}}}\, \mathbf{1}_{\mathcal{N}_i^+}\Bigr) \;>\; \delta^+ \quad \text{and} \quad \liminf_{i\to\infty} \mu\Bigl(f\, \mathbf{1}_{C_{\omega_{\Lambda_{N_i}}}}\, \mathbf{1}_{\mathcal{N}_i^-}\Bigr) \;<\; \delta^- \tag{5.22}$$

(b) *If μ is non-null, then it is not quasilocal at ω if there exists a diverging sequence $(N_i)_{i\geq 1}$ of natural numbers such that for each $i \geq 1$ there exist a natural number $R_i > N_i$ and two configurations η^+, η^- such that*

$$\limsup_{i\to\infty} \Bigl| \mu(f \mid \omega_{\Lambda_{N_i}}\, \eta^+_{\Lambda_{R_i}\setminus\Lambda_{N_i}}\, \sigma^+) - \mu(f \mid \omega_{\Lambda_{N_i}}\, \eta^-_{\Lambda_{R_i}\setminus\Lambda_{N_i}}\, \sigma^-) \Bigr| \;\geq\; \delta \tag{5.23}$$

for every $\sigma^\pm \in \Omega$.

As we have seen, condition (5.23) is a stronger form of non-quasilocality [(b) of Definition 5.13]. In this case it is appropriate to say that μ is *strongly non-quasilocal*, or *strongly non-Feller* [64, Definition 4.14]. To obtain (5.22) we have used consistency.

In practice, the lack of quasilocality has been detected by proving (5.23) for functions of the form $f(\sigma) = \sigma_\Lambda$. Furthermore, only single-site regions need to be checked due to Proposition 4.24. In the presence of translation invariance, then, non-quasilocality proofs typically refer to (5.23) for $\Lambda = \{0\}$ and $f(\sigma) = \sigma_0$. (This is not always the case, see for instance Section 4.3.5 in [64].)

After all these mathematical considerations, it is natural to wonder about the *physical* reasons for non-quasilocality. In quasilocal measures instead of (5.23) we get a limit zero as $\Gamma \to \mathbb{L}$. This means that the influence of $\sigma^\pm$ is shielded off if the intermediate spins are frozen in some configuration ω. In heuristic terms, in quasilocal measures the influence of far away regions is carried by the fluctuations of the spins in between; if these fluctuations are stopped so is the connection between the regions. Non-quasilocality means, thus, that there is some mechanism connecting distant regions that remains active even in the absence of fluctuations.

For measures obtained as images of a transformation the mechanism is clear; it goes under the keywords "hidden variables". While the measure acts on the space of "unhidden" *image variables* Ω', it is also determined by the "hidden" *object variables* in Ω acting through the transformation. In such a situation, the freezing of an image spin configuration acts as a conditioning on the object spin variables, under which the latter may still keep a certain amount of freedom to fluctuate. For some choice ω' of image variables, the conditioned object system may exhibit a *phase transition* which causes a long-range order that correlates local behavior to what happens at infinity. This produces non-quasilocality —that is nonzero differences (5.23)— for this particular ω'.

This "hidden variables" scenario explains non-quasilocality for renormalized measures and for measures obtained through cellular automata or spin-flip dynamics. While the non-quasilocality of joint measures of disordered models is of a different nature, still phase transitions are behind it [28, 66], as we shall discuss in Section 5.6 below.

The actual proofs of the failure of quasilocality are typically very technical. They combine a number of analytical tools (correlation inequalities, Pirogov-Sinai theory, strict convexity of thermodynamical potentials,...) with particular properties of each model in question. A systematic exposition of them is well beyond the scope of this course and may not be pedagogically useful. I prefer to discuss, instead, the overal strategy of the proof of non-quasilocality for block-renormalized measures, and illustrate other mathematical features through examples. These examples are relatively simple to analyze, and, in part due to its simplicity, have played a benchmark role in the understanding of the different manifestations of non-Gibbsianness.

5.4 Surprise number one: renormalization maps

5.4.1 The scenarios

Physicists define and work with renormalization transformations at the level of interactions (they speak of Hamiltonians, but they are really referring to interactions). Formally, they consider maps $\mathcal{R}$ related to our measure transformations τ according to the following diagram:

$$\begin{array}{ccc} \mu & \xrightarrow{\tau} & \mu' \\ \uparrow & & \downarrow \\ \Phi & \xrightarrow{\mathcal{R}} & \Phi' \end{array} \tag{5.24}$$

The diagram gives hints as to the possible mathematical complications of computing $\mathcal{R}$. While the upwards arrow on the left roughly corresponds to an exponential (Boltzmann prescription), the downwards arrow on the right corresponds to a log. This step is at the origin of the complicated diagrammatics associated to renormalization transformations. In contrast, the transformation τ is a linear object, much cleaner and straightforward at the mathematical level. In fact, from a computational standpoint, τ and $\mathcal{R}$ have complementary disadvantages: $\mathcal{R}$ involves logarithms, but τ acts on spaces of much larger dimensions. Conceptually, however, τ has the advantage of being always well defined while the status of $\mathcal{R}$ is less clear.

Renormalization transformations were initially devised to study critical points, approaching them from the high-temperature side where there is only one measure to contend with. But quickly physicists started to apply the successful renormalization ideas to first-order phase transitions, where there are several measures consistent with the same interaction. In these cases it is natural to wonder whether the different renormalized measures are associated to the same or different potentials:

$$\begin{array}{ccc} \{\mu_1,\cdots\} & \substack{\longrightarrow\\ \longrightarrow\\ \longrightarrow} & \{\mu'_1,\cdots\} \\ \uparrow\uparrow\uparrow & & \searrow\!\downarrow\!\swarrow \\ \Phi & \longrightarrow & \Phi' \end{array} \qquad \text{or} \qquad \begin{array}{ccc} \{\mu_1,\cdots\} & \substack{\longrightarrow\\ \longrightarrow\\ \longrightarrow} & \{\mu'_1,\cdots\} \\ \uparrow\uparrow\uparrow & & \downarrow\downarrow\downarrow \\ \Phi & \substack{\longrightarrow\\ \longrightarrow\\ \longrightarrow} & \{\Phi'_1,\cdots\} \end{array} \quad ? \tag{5.25}$$

While the leftmost scenario would be consistent with the renormalization paradigm, the rightmost one would indicate a *multivalued* map $\mathcal{R}$ quite contradictory to usual ideas. In fact, some numerical evidence did suggest the actual occurrence of this last scenario. To add to the confusion, the celebrated work of Griffiths and Pearce [22] pointed to the possible presence of "peculiarities" that would prevent any reasonable definition of $\mathcal{R}$. (The reader is referred to [64, Section 1.1] for historical references.)

Non-Gibbsian theory provided the necessary clarifications. It led to the following conclusions:

(a) The "multivaluedness scenario" [rightmost possibility in (5.25)] is impossible within reasonable spaces of interactions [64, Theorem 3.6].

(b) In many instances, however, as initially shown by Israel [26], renormalized measures may fail to be quasilocal. That is, the downwards arrows in (5.25) may fail to exist.

(c) If the interaction Φ and measures μ_i are translation invariant, either the renormalized measures μ_i' are all Gibbsian for the *same* interaction, or they are *all* non-Gibbsian [64, Theorem 3.4].

In conclusion, instead of those in (5.25), the two competing scenarios are

$$\begin{array}{ccc} \{\mu_1,\cdots\} & \substack{\longrightarrow\\ \longrightarrow\\ \longrightarrow} & \{\mu_1',\cdots\} \\ \uparrow\uparrow\uparrow & & \searrow\downarrow\swarrow \\ \Phi & \longrightarrow & \Phi' \end{array} \qquad \text{or} \qquad \begin{array}{ccc} \{\mu_1,\cdots\} & \substack{\longrightarrow\\ \longrightarrow\\ \longrightarrow} & \{\mu_1',\cdots\} \\ \uparrow\uparrow\uparrow & & \not\uparrow \\ \Phi & \not\longrightarrow & ?? \end{array} \tag{5.26}$$

Both of these scenarios occur —the left one probably more often— but I will concentrate on the general strategy to prove the validity of the rightmost scenario. I will only sketch the different steps, relying on two examples as an illustration: 2×2-decimation and Kadanoff transformations of the translation-invariant states of the two-dimensional Ising model in zero magnetic field at low enough temperature. The decimation example is the first and simplest example of non-quasilocal renormalized measure, which carefully analyzed by Israel [26], and is the genesis of the non-Gibbsianness work in [64]. Kadanoff transformations, on the other hand, illustrate transformations with strictly positive kernels and they were already considered by Griffiths and Pearce as sources of "pathologies". I will skip all fine calculational details —which are fully given in [64, Section 4.1.2]— and concentrate on the main brush strokes (which are already complicated enough). The strategy, which is naturally divided into four steps, in fact shows that the non-quasilocality of the renormalized measures $\mu' = \mu\tau$ satisfies the stronger property (5.23).

5.4.2 Step zero: Understanding the conditioned measures

To understand the meaning of $\mu'_{\Lambda'}(\,\cdot \mid \omega')$, for $\Lambda' \subset \mathbb{L}'$, $\omega' \in \Omega'$, we introduce the measure on $(\Omega \times \Omega', \mathcal{F} \times \mathcal{F}')$ with marginals (=projections on Ω and Ω') μ and μ'. Explicitly,

$$\widetilde{\mu}(\widetilde{F}) \;=\; \int \widetilde{F}(\omega,\omega')\,\mu(d\omega)\,\tau(d\omega' \mid \omega) \tag{5.27}$$

for every function $\widetilde{F}$ that is $\mathcal{F} \times \mathcal{F}'$-measurable and bounded. It is useful to visualize $\Omega \times \Omega'$ as configurations on two parallel "slices", $\mathbb{L}$ and $\mathbb{L}'$. The spins on the former are the *original*, *object* ot *internal* spins and those on the latter the *renormalized* or *image* spins. A simple verification of the properties determining Definition 3.9 shows that

$$\mu'_{\Lambda'}(\,\cdot \mid \omega') \;=\; \widetilde{\mu}_{\Lambda' \times \mathbb{L}}(\,\cdot \mid \omega'_{\Lambda'^{\mathrm{c}}})\;. \tag{5.28}$$

We see, then, that $\mu'_{\Lambda'}(\,\cdot \mid \omega')$ is a measure on an infinite spin space formed by the spins in $\mathbb{L}$ plus those in the finite region Λ'. The proper definition of measures for unbounded regions needs some care. In our case we count on the help of specifications.

Indeed, we are interested in a measure μ that is Gibbsian for some interaction Φ and in product transformations (5.1)/(5.2). The measure $\widetilde{\mu}$ must then be consistent with a specification defined by the interaction Φ on the slice $\mathbb{L}$ and "conic bonds" connecting $\mathbb{L}$ and $\mathbb{L}'$ defined by the functions $T_{x'}$. Rather that writing the full details for $\widetilde{\mu}$ let us focus on our target measure $\widetilde{\mu}_{\Lambda'\times\mathbb{L}}(\,\cdot\mid\omega'_{\Lambda'^{c}})$. To simplify matters still further, let me advance the fact that the addition of the finitely many spins in Λ', being only a local modification, does not produce any major change in the properties we are after (we shall come back to this in step 3 below). Hence, we look at $\widetilde{\mu}_{\mathbb{L}}(\,\cdot\mid\omega')$ which we interpret as the measure on Ω obtained by conditioning the original spins to be "compatible" with the image configuration ω'. Our previous comment on a Φ-T interaction is formalized, even more generally, as follows.

Proposition 5.29 *Let μ be consistent with a specification Π whose density functions are $\{\gamma_\Lambda : \Lambda \Subset \mathbb{L}\}$ and let τ be a block transformation defined by densities $\{T_{x'}(\omega'_{x'}\mid\omega_{B_{x'}})\}$. For each $\omega'\in\Omega'$ let*

$$\Omega^{\omega'} \;=\; \Bigl\{\omega\in\Omega : T_{x'}(\omega'_{x'}\mid\omega_{B_{x'}})>0\;,\; x'\in\mathbb{L}'\Bigr\}\;. \tag{5.30}$$

The, the measure $\widetilde{\mu}_{\mathbb{L}}(\,\cdot\mid\omega')$ is consistent with the specification $\Pi^{\tau,\omega'}$ on $\Omega^{\omega'}$ defined by the density functions

$$\gamma_\Lambda^{\omega'}(\sigma_\Lambda\mid\omega_{\Lambda^c}) \;=\; \frac{1}{\mathit{Norm.}}\,\gamma_\Lambda(\sigma_\Lambda\mid\omega_{\Lambda^c})\prod_{x'\in B'_\Lambda} T_{x'}\Bigl(\omega'_{x'}\Bigm|(\sigma_\Lambda\,\omega)_{B_{x'}}\Bigr) \tag{5.31}$$

where $B'_\Lambda=\{x'\in\mathbb{L}' : B_{x'}\cap\Lambda\neq\emptyset\}$ and "Norm" stands for the sum over σ_Λ of the numerator. The pair $(\Omega^{\omega'},\Pi^{\tau,\omega'})$ is the ω' **constrained internal-spin system***.*

Exercise 5.32 *Prove this proposition. (*Hint: *The shortest route to prove that (5.31) indeed defines a specification is through property (3.31).)*

If Π is defined by an interaction Φ and the functions $T_{x'}$ are strictly positive, then $\Omega^{\omega'}=\Omega$ and $\Pi^{\tau,\omega'}$ is Gibbsian for the interaction

$$\phi_B^{\tau,\omega'}(\omega) \;=\; \phi_B(\omega)+\begin{cases} -\log T_{x'}(\omega'_{x'}\mid\omega_{B_{x'}}) & \text{if } B=B_{x'} \text{ for some } x'\in\mathbb{L}'\\ 0 & \text{otherwise}\,.\end{cases} \tag{5.33}$$

Observe that if the temperature is considered, then the factor β multiplies only the terms ϕ_B, but *not* the last logarithm. For example, for the p-Kadanoff transformation of the Ising model with magnetic field h at inverse temperature β, the measure $\widetilde{\mu}_{\mathbb{L}}(\,\cdot\mid\omega')$ is Gibbsian for the interaction with formal Hamiltonian

$$-\beta\Bigl\{\sum_{\langle x,y\rangle}\omega_x\,\omega_y-h\sum_x\omega_x-p\,\beta^{-1}\sum_x\omega'_x\sum_{y\in B_x}\omega_y+\beta^{-1}\sum_x\log\Bigl[2\cosh\Bigl(p\sum_{y\in B_x}\omega_y\Bigr)\Bigr]\Bigr\} \tag{5.34}$$

which corresponds to an Ising model with an additional magnetic field that is positive, block-dependent, and also temperature-dependent, plus a multispin non-linear antiferromagnetic term with temperature-dependent couplings.

If the renormalization transformation is not strictly positive, for instance if it is deterministic, we fall into the framework of models with exclusions. Its analysis depends on the type of exclusion. The example of decimation transformations is particularly simple, as the constraints determined by each ω' amount to fixing the spins at the sites in $b\mathbb{Z}^d$. In such a situation it is better simply to ignore these decimated sites and consider the measure $\widetilde{\mu}_{\mathbb{Z}^d\setminus b\mathbb{Z}^d}(\,\cdot\mid\omega')$ on the remaining internal spins. That is, we take as internal spin system $\Omega_{\mathbb{Z}^d \setminus b\mathbb{Z}^d}$ and the interaction obtained from the original one by fixing the decimated spins. For the decimation of the Ising model, this internal-spin interaction corresponds to the same original Ising interaction plus a field on neighbors to decimated sites induced by the links to them. In Israel's example, three sublattices arise naturally. The decimated spins are on the *even sublattice* $\mathbb{L}_{\rm even}$ formed by sites with both coordinates even. The neighbors to decimated sites occupy the *odd/even sublattice* $\mathbb{L}_{\rm odd/even}$ where the two coordinates have different parity. The remaining sites, with both coordinates odd, form the *odd sublattice* $\mathbb{L}_{\rm odd}$. The interaction defining $\widetilde{\mu}_{\mathbb{Z}^2\setminus 2\mathbb{Z}^2}(\,\cdot \mid \omega')$ corresponds to the formal Hamiltonian

$$-\beta\Big\{\sum_{x\in\mathbb{L}_{\rm odd}}\sum_{\substack{y\in\mathbb{L}_{\rm odd/even}\\ |x-y|=1}}\omega_x\,\omega_y \quad + \sum_{x\in\mathbb{L}_{\rm odd/even}}\Big(\sum_{\substack{y\in\mathbb{L}_{\rm even}\\ |x-y|=1}}\omega'_y\Big)\,\omega_x\Big\}\;. \tag{5.35}$$

In conclusion, step zero teaches us that each conditioned measure in question —$\widetilde{\mu}_{\Lambda'\times\mathbb{L}}(\,\cdot \mid \omega'_{\Lambda'^c}) = \mu'_{\Lambda'}(\,\cdot \mid \omega')$ or $\widetilde{\mu}_{\mathbb{L}}(\,\cdot \mid \omega')$— is determined through consistency with some interaction. If the interaction presents a first-order phase transition, there are infinitely many measures to choose from. The proof of non-Gibbsianness, in fact *needs* the presence of such phase transitions. Let us now move to the remaining steps.

5.4.3 The three steps of a non-quasilocality proof

Step 1: Choice of an image configuration producing a phase transition on the internal spins We need to choose some special configuration $\widehat{\omega}'$ for which the constrained internal spins undergo a first order phase transition. That is, $\widehat{\omega}'$ must be such that there exist two different measures $\mu_+^{\widehat{\omega}'}, \mu_-^{\widehat{\omega}'} \in \mathcal{G}(\Pi^{\tau,\widehat{\omega}'})$. Those being different means that there exists a local observable f such that

$$\big|\mu_+^{\widehat{\omega}'}(f) - \mu_-^{\widehat{\omega}'}(f)\big| \;=:\; \delta \;>\; 0\;. \tag{5.36}$$

In such a situation, one may wonder which measure has the right to be denoted $\widetilde{\mu}_{\mathbb{L}}(\,\cdot \mid \widehat{\omega}')$. While we do not answer this, the rest of the argument shows that whichever the choice, it leads to a discontinuity at $\widehat{\omega}'$.

The choice of $\widehat{\omega}'$, of course, depends on the problem. If the original model already exhibits multiple phases, then the rule of thumb is to choose $\widehat{\omega}'$ so as not to favor any of these phases. For the Kadanoff and decimation examples this means that $\widehat{\omega}'$ must be "magnetically neutral". The simplest choice, the alternating configuration $\widehat{\omega}'_x = (-1)^{|x|}$, is already suitable.

For Israel's example, this choice causes the cancellation of the effective field due to neighboring decimated spins, which corresponds to replacing the even spins by holes. Formally, the second sum in the internal-spin interaction (5.35) disappears. This corresponds to an Ising model on the *decorated* lattice $\mathbb{L}\setminus\mathbb{L}_{\rm even}$, formed by sites with four neighbors —those in $\mathbb{L}_{\rm odd}$— and the "decorations" —sites in $\mathbb{L}_{\rm odd/even}$— linked only to two other sites. If we are only interested in observables on the odd lattice we can sum first over the spins at the decorations. A little bit of algebra shows that

$$\sum_{\sigma_d=\pm1}\exp\big(\beta\,\sigma_1\sigma_d + \beta\,\sigma_d\sigma_2\big) \;=\; C\exp\big(\beta'\sigma_1\sigma_2\big) \tag{5.37}$$

where C is an uninteresting constant and

$$\beta' \;=\; \tfrac{1}{2}\log\cosh 2\beta\;. \tag{5.38}$$

This means that the internal-spin system, constrained by the alternating configuration, becomes equivalent to an Ising model at a lower temperature. If the initial model is at a low enough temperature, β' exceeds the critical Onsager value and the internal-spin model acquires two different pure phases, respectively supported on configurations formed by a percolating sea of "+1" and a sea of "−1", with fluctuations on finite and isolated "islands". These are our measures $\mu_+^{\widehat{\omega}'}$ and $\mu_-^{\widehat{\omega}'}$; they are characterized by the fact that

$$0 \;<\; m(\beta') \;:=\; \mu_+^{\widehat{\omega}'}(\sigma_0) \;=\; -\mu_-^{\widehat{\omega}'}(\sigma_0)\,. \tag{5.39}$$

The analogous proof for the Kadanoff transformed measure is much more involved. It demands a technical, but widely used, perturbative approach starting from the zero-temperature phase diagram. Let me describe it briefly, while referring the reader to [64, Appendix B] for a detailed presentation and all relevant definitions and references. In a nutshell, the approach has two stages:

Stage 1: Determination of the translation-invariant *ground states* of the model. These are the translation-invariant measures consistent with the specification obtained as the zero-temperature ($\beta \to \infty$) limit of the specification under study. Two type of conditions must be met for the approach to be applicable. First, the extremal points of such set of measures (*pure phases*) must be δ-like, that is, supported by single configurations. Second, the resulting phase diagram (that is, the catalogue of ground states for different values of parameters like h) must be *regular*, in some precise sense, or have appropriate symmetry properties. In particular the number of extremal translation-invariant ground states must be finite throughout a whole region of parameter values.

Stage 2: Stability of the zero-temperature phase diagram. This is proven through a very powerful and sophisticated theory due to Pirogov and Sinai. Its hypotheses include the regularity features mentioned above plus the so-called *Peierls condition* which roughly means that local fluctuations of ground states are suppressed exponentially in its volume. This allows to show stability of phases by suitably generalizing the Peierls contour argument for the Ising model.

It is relatively simple to verify that the translation-invariant ground states interaction (5.34) with an alternating block-field ω'_x are (δ-measures on) the all-"+1" or the all-"−1" configurations, depending on h, with coexistence for $h = 0$ (by symmetry reasons). The validity of the Peierls condition follows, by continuity considerations, from that of the Ising model. Some subtleness arises from the fact that (5.34) has β-dependent parameters (the last two). This requires a stronger (uniform) version of Pirogov-Sinai theory. The conclusion of all this analysis is that the interaction (5.34) admits two different consistent measures $\mu_+^{\widehat{\omega}'}$ and $\mu_-^{\widehat{\omega}'}$, with properties similar to those of the decimation case. In particular they satisfy (5.39).

Step 2: Choice of discontinuity neighborhoods To prove (5.23) the measures consistent with $\Pi^{\tau,\widehat{\omega}'}$ need to be approximated by measures obtained similarly for image spins fixed in configurations of the form $\widehat{\omega}'_{\Lambda'}\eta'_{\Gamma'\setminus\Lambda'}\sigma'$. The idea is to find configurations ${\eta'}^{\pm} \in \Omega'$ and a sequence of natural numbers N_R, with $N_R > R$, such that *all* measures $\mu^{R,{\sigma'}^{+}}$ and $\mu^{R,{\sigma'}^{-}}$, respectively consistent with the specifications $\Pi^{\tau,\widehat{\omega}'_{\Lambda'_R}{\eta'}^{+}_{\Lambda'_{N_R}\setminus\Lambda'_R}\sigma'^{+}}$ and $\Pi^{\tau,\widehat{\omega}'_{\Lambda'_R}{\eta'}^{-}_{\Lambda'_{N_R}\setminus\Lambda'_R}\sigma'^{-}}$, satisfy that, for any choice of $\sigma'^{+}, \sigma'^{-} \in \Omega'$

$$\mu^{R,{\sigma'}^{\pm}}(f) \;\xrightarrow[R\to\infty]{}\; \mu_{\pm}^{\widehat{\omega}'}(f) \tag{5.40}$$

where f is the observable satisfying (5.36). Combining the latter with (5.40) we thus obtain that

$$\lim_{R\to\infty} \left| \widetilde{\mu}_{\mathbb{L}}\Big(f \;\Big|\; \widehat{\omega}'_{\Lambda'_R}\, \eta'^{+}_{\Lambda'_{N_R}\setminus\Lambda'_R}\, \sigma'^{+}\Big) - \widetilde{\mu}_{\mathbb{L}}\Big(f \;\Big|\; \widehat{\omega}'_{\Lambda'_R}\, \eta'^{-}_{\Lambda'_{N_R}\setminus\Lambda'_R}\, \sigma'^{-}\Big) \right| \;\geq\; \delta \tag{5.41}$$

for any $\sigma'^{+}, \sigma'^{-} \in \Omega'$ for R large enough. In view of (5.28), this almost proves (5.23) for the renormalized measure μ'. The existence of configurations $\eta'^{\pm}$ with the above properties is, as a matter of fact, a further condition for the choice of $\widehat{\omega}'$.

For the 2×2-decimation of the Ising model it is relatively simple to prove (5.40). Indeed, a short calculation shows that if the decimated spins are fixed in the alternating configuration inside a region Λ_R and equal to $+1$ in the annulus immediately outside, the internal spins in the region Λ_{R+1}, are subjected to an Ising interaction with a magnetic field at the boundary. This field is at least equal to β' *whatever the configuration of the spins further out* (internal or otherwise). Hence, regardless of the image configuration on Λ^{c}_{R+1}, the expected magnetization at the origin is (by Griffiths inequalities) no smaller than that of an Ising model on a square with "+" boundary conditions, which converges to that of the "+" Ising measure when the size of the square diverges. An analogous argument can be done for a "-1" boundary condition. We conclude that (5.40) is verified for $N_R = R+1$, $\sigma'^{\pm} = \pm$ and $f(\sigma) = \sigma_0$, and thus

$$\lim_{R\to\infty} \left| \widetilde{\mu}_{\mathbb{L}}\Big(\sigma_0 \;\Big|\; \widehat{\omega}'_{\Lambda'_R} (+1)_{\Lambda'_{R+1}\setminus\Lambda'_R}\, \sigma'^{+}\Big) - \widetilde{\mu}_{\mathbb{L}}\Big(\sigma_0 \;\Big|\; \widehat{\omega}'_{\Lambda'_R} (-1)'^{-}_{\Lambda'_{R+1}\setminus\Lambda'_R}\, \sigma'^{-}\Big) \right| \;=\; 2\, m(\beta') \tag{5.42}$$

which is nonzero if the temperature of the original Ising model is low enough.

The argument for all other cases (including decimation in higher dimensions) is less simple. The standard strategy involves finding configurations $\eta'^{+}, \eta'^{-} \in \Omega'$ such that:

(i) The specifications $\Pi^{\tau,\eta'^{+}}$ and $\Pi^{\tau,\eta'^{-}}$ admit *unique* consistent measures respectively denoted by $\mu^{\eta'^{+}}$ and $\mu^{\eta'^{-}}$.

(ii) For any $R > 0$, all measures $\mu^{R,\eta'^{+}}$ consistent with $\Pi^{\tau,\widehat{\omega}'_{\Lambda'_R}\eta'^{+}}$ and all measures $\mu^{R,\eta'^{-}}$ consistent with $\Pi^{\tau,\widehat{\omega}'_{\Lambda'_R}\eta'^{-}}$ satisfy

$$\left|\mu^{R,\eta'^{+}}(f) - \mu^{R,\eta'^{-}}(f)\right| \;\geq\; \left|\mu^{\widehat{\omega}'}_{+}(f) - \mu^{\widehat{\omega}'}_{-}(f)\right| \tag{5.43}$$

(this is often done with the help of correlation inequalities).

Property (i) implies that each of the specifications $\Pi^{\tau,\widehat{\omega}'_{\Lambda_R}\eta'^{\pm}}$, $R > 0$, has also a single consistent measure because it is obtained from $\Pi^{\tau,\eta'^{\pm}}$ by a local change. We also make use of the following fact: Let (μ_n) and (Φ_n) be respectively sequences of measures and interactions on $(\Omega, \mathcal{F})$ such that $\mu_n \in \mathcal{G}(\Phi_n)$. Then, if Φ_n converges (in $\mathcal{B}_1$) to an interaction Φ, every convergent subsequence of (μ_n) is consistent with Π^{Φ}. We apply this to the sequence of interactions $\Phi^{\widehat{\omega}'_{\Lambda'_R}\eta'^{\pm}_{\Lambda'_N}}$ which converges, as $N\to\infty$, to $\Phi^{\widehat{\omega}'_{\Lambda'_R}\eta'^{+}}$, to conclude, from (5.43) and (5.36), that for each $R > 0$ one can choose N_R sufficiently large so that (5.41) is valid for a δ slightly smaller than that in (5.36).

Step 3: "Unfreezing" of Λ' The last step consists in showing that, as a consequence of the previous steps, we can actually find an observable $f' \in \mathcal{F}'_{\Lambda'}$, somehow related or inspired by f, so that the analogue of (5.41) holds for $\widetilde{\mu}_{\Lambda'\times\mathbb{L}}(f' \mid \cdot)$. In fact, for each $\omega' \in \Omega'$ the specification $\Pi^{\tau,\Lambda',\omega'}$

defining the measures $\widetilde{\mu}_{\Lambda'\times\mathbb{L}}(\,\cdot\mid\omega'_{\Lambda'^c})$ is obtained from the specification $\Pi^{\tau,\omega'}$ defining $\widetilde{\mu}_{\mathbb{L}}(\,\cdot\mid\omega')$ by "unfreezing" the factors $T_{x'}(\,\cdot\mid\omega_{B_{x'}})$ for $x'\in\Lambda'$. This corresponds to a multiplication of the kernels of $\Pi^{\tau,\omega'}$ by a local density, or to the addition of a finite number of bonds to the interaction defining the latter. Therefore there is a canonical bijection between $\mathcal{G}(\Pi^{\tau,\Lambda',\omega'})$ and $\mathcal{G}(\Pi^{\tau,\omega'})$ for each fixed ω'. In particular, the existence of unique or multiple phases in one of them implies the same feature in the other one. We conclude that the configurations $\widehat{\omega}'$ and $\eta^{\pm}$ chosen above allow also the successful completion of steps 1 and 2 for $\widetilde{\mu}_{\Lambda'\times\mathbb{L}}$ for every $\Lambda'\Subset\mathbb{L}'$. We only need to show the existence of f' such that

$$\left|\widetilde{\mu}_{+}^{\widehat{\omega}',\Lambda'}(f')-\widetilde{\mu}_{-}^{\widehat{\omega}',\Lambda'}(f')\right| \;>\; 0\,. \tag{5.44}$$

where

$$\widetilde{\mu}_{\pm}^{\widehat{\omega}',\Lambda'}(f') \;=\; \sum_{\sigma'_{\Lambda'}}\int f'(\sigma'_{\Lambda'})\prod_{x'\in\Lambda'}T_{x'}\big(\sigma'_{x'}\mid\omega_{B_{x'}}\big)\,\mu_{\pm}^{\widehat{\omega}'}(d\omega)\,. \tag{5.45}$$

The properties of $mu_{\pm}^{\widehat{\omega}'}$ must now be exploited. For the decimation and Kadanoff examples, we have to consider

$$2\,m' \;:=\; \sum_{\sigma'_0}\int \sigma'_0\,T_0\big(\sigma'_0\mid\omega_{B_0}\big)\,\Big[\mu_{+}^{\widehat{\omega}'}(d\omega)-\mu_{-}^{\widehat{\omega}'}(d\omega)\Big]\,. \tag{5.46}$$

At low enough temperatures the measure $\mu_{+}^{\widehat{\omega}'}$ favors "+1" spins, while $\mu_{-}^{\widehat{\omega}'}$ favors "−1" (this can be seen by correlation inequalities or contour arguments: the probability that a finite region be inside or intersecting a contour goes to zero as temperature decreases). The transformation density T_0, on the other hand, favors alignment of σ'_0 with the majority of the spins in ω_{B_0}. Both effects combined lead to $m'>0$.

5.4.4 Non-quasilocality throughout the phase diagram

Following the preceding argument, non-quasilocality has been exhibited for renormalizations of the Ising model at low temperature and zero field, for all of the block transformations described in Section 5.1. The renormalized measures have subsequently been shown to be weakly Gibbs [1], while decimated measures are, in fact, almost quasilocal [11, 10].

We see, however, that the above argument relies on the existence of phase transitions *for the constrained internal spin system* rather than for the original system. Therefore, non-quasilocality should be expected also outside the coexistence region of the original model, in situations where the constraints produced by the renormalized spins act like fields that bring the internal system into a phase coexistence region. So we must add the scenario

$$\begin{array}{ccc} \mu & \xrightarrow{\;\mathcal{T}\;} & \mu' \\ \uparrow & & \not\uparrow \\ \Phi & \not\longrightarrow & ?? \end{array} \tag{5.47}$$

in competition with scenario (5.24). Israel [26] already exhibited such a phenomenon in his 2×2-decimation example: A small but non-zero magnetic field of the original Ising model can be compensated by the (non translation-invariant) field created by a suitable $\widehat{\omega}'$ so that, at low (original) temperatures, the non-decimated spins undergo a phase transition and the decimated measure becomes non-quasilocal. This measure is, however, almost-quasilocal [14], and its quasilocality can be restored by further decimations [47]. More dramatic examples include block-averaging [64, Section 4.3.5] and majority [63] transformations of the Ising model at high magnetic field, and decimations of high-q-Potts models above the critical temperature [63]. One can even design, for each temperature, a perverse transformation such that the renormalization of the Ising measure at this temperature is non-quasilocal [59].

There is a clear message coming from these examples: The choice of a renormalization transformations is a touchy business. Top-of-the-shelf choices may lead to non-Gibbsian renormalized measures for which calculations of Hamiltonian parameters —renormalized temperature, renormalized couplings— have a doubtful meaning. The transformation must be well-adapted to the problem, and the questions, at hand. In particular, block-spin transformations may not be a good idea at low temperatures, where long-range order pervades. Rather, renormalization ideas should be applied at the level of collective variables, like contours [17].

5.5 Surprise number two: spin-flip evolutions

Metropolis and heath-bath algorithms have been instrumental for the simulation of statistical mechanical systems. They are processes in which each spin of a finite lattice is visited according to a certain routine (sequentially, randomly, by random shuffling) and updated stochastically by comparing energies before and after the proposed flip. Their continuous-time counterpart are the *Glauber spin-flip dynamics* in which the updates are attempted according to independent Poissonian clocks attached to each site. The dynamics are tailored so as to converge to a target spin measure which is the object of the simulation. Each simulation realization is started as some initial configuration, and a sample configuration is collected after a number of steps. If this number is sufficiently large, the samples are distributed almost like in the target measure. Often, the initial configuration is a ground state, or zero-temperature measure, and the simulation acts as a numerical furnace that heats it up ("unquenches" it) so to bring it to a typical configuration at the intended temperature.

These simulation schemes define a sequence of transformations of measures as considered in Section 5.1. Actual simulations apply these transformations to Boltzmann measures in finite regions (usually with periodic boundary conditions), but ideally they should be applied to measures on the whole lattice. An ideal "unquenching" transformation is, then, a high-temperature Metropolis or Glauber dynamics (that is, a dynamics converging to a high temperature Gibbs state) applied to a low-temperature Gibbs state. We were surprised by the fact that, if the temperature difference between the initial and final states is big enough, non-Gibbsianness enters into the picture [62].

To see how, let us consider a very simple updating process for Ising spins, in which at successive time units each spin is flipped independently with probability $\epsilon \in (0,1)$. The invariant measure for this process gives equal probability to each spin configuration, thus the process can be interpreted either as an infinite-temperature parallel Metropolis algorithm, or an infinite-temperature discrete time Glauber dynamics. Mathematically, this process is a block transformation with $\Omega' = \Omega = \{-1,1\}^{\mathbb{Z}^d}$, single-site blocks and kernel densities

$$\begin{aligned} T_{\{x\}}(\omega_x \mid \omega_x) &= 1-\epsilon \\ T_{\{x\}}(-\omega_x \mid \omega_x) &= \epsilon \end{aligned} \tag{5.48}$$

Such densities are better expressed as a matrix $(T_x)_{\sigma\,\eta} := T_{\{x\}}(\sigma \mid \eta)$ which takes the form

$$T_x = \begin{pmatrix} 1-\epsilon & \epsilon \\ \epsilon & 1-\epsilon \end{pmatrix} = \mathbb{I} - \epsilon \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix} =: \mathbb{I} - \epsilon\,\mathbb{J}\,. \tag{5.49}$$

The n-th iteration of such a transformation corresponds, thus, to single-site kernels $T^n_{\{x\}}(\sigma_x \mid \eta_x) = (T^n_x)_{\sigma_x\,\eta_x}$ where T^n_x is the n-th power of the matrix T_x. Given that $\mathbb{J}^\ell = 2^{\ell-1}\mathbb{J}$ if $\ell \geq 1$ (and equal to $\mathbb{I}$

if $\ell = 0$), we obtain

$$T_x^n = \sum_{\ell=0}^{n} \binom{n}{\ell} (-\epsilon)^\ell \, \mathbb{J}^\ell = \mathbb{I} + \frac{1}{2} \sum_{\ell=1}^{n} \binom{n}{\ell} (-2\epsilon)^\ell \mathbb{J} = \mathbb{I} + \frac{1}{2} \big[(1-2\epsilon)^n - 1\big] \mathbb{J} \tag{5.50}$$

$$= \frac{1}{2} \begin{pmatrix} 1+a_n & 1-a_n \\ 1+a_n & 1-a_n \end{pmatrix} \tag{5.51}$$

with

$$a_n = (1-2\epsilon)^n. \tag{5.52}$$

Therefore

$$T_{\{x\}}^n(\omega_x' \mid \omega_x) = \frac{1}{2} + \frac{a_n}{2} \omega_x' \omega_x = A_n \, \mathrm{e}^{h_n \omega_x' \omega_x} \tag{5.53}$$

where the factor $A_n = \big[2\cosh h_n\big]^{-1}$ will be eaten up by normalizations and

$$h_n = \log\Big(\frac{1+a_n}{1-a_n}\Big) . \tag{5.54}$$

[In fact, T^n is a Kadanoff transformation with single-site blocks and $p = h_n$.] Let me observe that

$$h_n \underset{n\to\infty}{\searrow} 0 \qquad \text{and} \qquad h_n \underset{\epsilon\to 0}{\nearrow} \infty . \tag{5.55}$$

We can now make use of the analysis of the previous section. For fixed n we look to the pair of slices $\Omega \times \Omega'$ respectively formed by the initial configurations and those at "time" n, that is at the n-th iteration of the process. The non-quasilocality of the transformed measure μ' is related to the existence of some $\widehat{\omega}'$ for which the resulting constrained initial-spin system exhibits multiple phases. Such a system corresponds to an interaction which includes, as additional terms, the bonds (5.53). Therefore, if we start with an Ising measure, the condition of observing a configuration ω' at time n is seen by the initial spins as a correction to the magnetic field leading to a formal Hamiltonian

$$-\beta \Big\{ \sum_{\langle x,y\rangle} \omega_x \, \omega_y - \sum_x \Big(h + \frac{h_n}{\beta} \omega_x'\Big) \, \omega_x \Big\} . \tag{5.56}$$

We can distinguish three regimes:

(i) *Short times:* For n small, the effective magnetic field $\big|h + \frac{h_n}{\beta}\omega_x'\big|$ is large if ϵ is sufficiently small [rightmost observation in (5.55)]. Hence no phase transition is present and the time-n measure is expected to be quasilocal. This can be proven, for n small enough, through an argument that relies on the existence of "global" specifications from which the specifications $\Pi^{\tau,\omega'}$ are derived. The argument exploits FKG monotonicity and Dobrushin uniqueness criterion. If the initial model is itself at high temperature, then the measure remains Gibbsian throughout the evolution.

(ii) *Long times:* For n large, $h + \frac{h_n}{\beta}\omega_x' \sim h$ [rightmost observation in (5.55)] hence no phase transition, and thus the quasilocality of μ', is expected (and proven) if $h > 0$, while for large β and $h = 0$ a phase transition makes the transformed measure discontinuous at $\widehat{\omega}'$ =alternating configuration.

(iii) *Intermediate times:* If $h > 0$ and ϵ small, then for large enough β there is a range of n for which a configuration $\widehat{\omega}'$ exists such that $\frac{h_n}{\beta}\widehat{\omega}_x'$ effectively compensates h. The resulting phase transition leads to the non-quasilocality of the evolved measure.

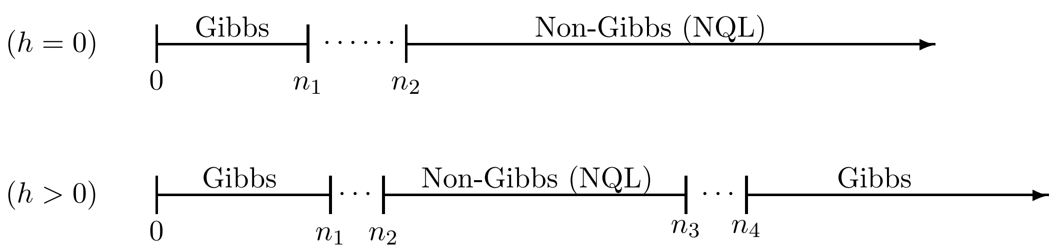

Figure 5.1: Proven regimes of Gibbsianness and non-Gibbsianness for a low-temperature Ising measure subjected to fast heating

The situation is summarized in Figure 5.5 when the stirring probability ϵ is small. Larger values of ϵ lead to larger changes in each time unit and some of the initial regions may disappear (some n_i may turn out to be smaller than one). Through a more complicated but similar analysis the same statements are proven for general high-temperature stochastic dynamics both in discrete and continuous time [62]. In these cases the effective Hamiltonian for the evolved measure acquires some long-range terms that decay exponentially with the diameter of the bond. They must be controlled by perturbative arguments (cluster expansions, Pirogov-Sinai theory).

The heuristic explanation of these results is as follows. For short enough times the evolution causes only a few changes. Therefore the evolved measure differs little of the initial measure and, in particular, preserves its Gibbsian character. This is true for more general reversible dynamics, for instance for dynamics of Kawasaki type —which conserve the total number of spins in each value— or mixtures of Glauber and Kawasaki dynamics [51]. The onset of non-Gibbsianness at later times —and of the subsequent Gibbsianness if $h > 0$— corresponds to a transition in the *most probable history of an improbable configuration* (the expression is of Aernout van Enter). There are two competing mechanisms to explain the presence of a droplet ω_Λ at some instant of the evolution: (i) It has been created by the dynamics, and (ii) the droplet was already present initially and it survived the evolutionary period. The probabilistic cost of the first event increases, roughly, exponentially with the volume $|\Lambda|$ of the droplet. The second mechanism is even more costly if the droplet is atypical for the initial measure, because its initial presence costs already a volume exponential. But if ω is typical of any of the phases of the initial system, this factor becomes exponential only in the surface area $|\partial\Lambda|$. As droplets at worst shrink at constant velocity, the second mechanism is more probable for such a droplet for intermediate times.

Suppose now that at some not-too-short time we observe a configuration $\widehat{\omega}'_\Lambda \sigma'_{\Gamma\setminus\Lambda}$ with Λ large and Γ enormous, $\widehat{\omega}'$ atypical of any of the phases of the initial system and σ' typical of one of them. The most likely explanation is, thus, that $\widehat{\omega}'_\Lambda$ was formed during the evolution, while σ'_Γ remains of the initial configuration. The initial gigantic σ' droplet causes a bias on the evolved configuration around the origin. In this way, through the original ("hidden") spins, the far-away annulus $\sigma'_{\Gamma\setminus\Lambda}$ determines the evolved measure close to the origin; quasilocality is lost. For non-zero magnetic field, the initial system has only one phase. If the elapsed time is large enough, only droplets typical of this phase are able to survive, any other σ' must have been created by the evolution. This creation is a local phenomenon, so quasilocality is recovered.

Whereas in a renormalization context, lack of quasilocality implies that a renormalization group map does not exist, here the physical interpretation is that the evolved (fastly heated) measure cannot be described by a temperature, after some time. This phenomenon has been the object of a numerical study [52].

5.6 Surprise number three: disordered models

A statistical mechanical system is *disordered* if there are parameters in the interaction that are themselves random variables. Its mathematical framework is as follows. Besides the space of spin configurations $(\Omega = S^{\mathbb{L}}, \mathcal{F})$ there is another space of *disorder variables* $\big(\Omega^* = (S^*)^{\mathbb{L}^*}, \mathcal{F}^*\big)$, where S^* is some space that need not be finite or discrete, $\mathbb{L}^*$ is a countable set and $\mathcal{F}^*$ is the product σ-algebra of some natural Borel measure structure of S^*. The disorder variables come equipped with some *disorder measure* $\mathbb{P}$ that is often extremely simple, typically a product measure. A *disordered interaction* is a family of functions $\big\{\phi_A(\,\cdot \mid \cdot\,) \in \mathcal{F} \times \mathcal{F}^* : A \Subset \mathbb{L}\big\}$ such that $\phi_A(\,\cdot \mid \eta^*) \in \mathcal{F}_A$ for each $A \Subset \mathbb{L}$ and $\eta^* \in \Omega^*$. Often, the disorder dependence is also local in the sense that for each $A \Subset \mathbb{L}$ there exists $A^* \Subset \mathbb{L}^*$ such that $\phi_A(\sigma \mid \cdot\,) \in \mathcal{F}^*_{A^*}$ for each $\sigma \in \Omega$. A disordered interaction defines for each value η^* an interaction $\Phi(\,\cdot \mid \eta^*) = \big\{\phi_A(\,\cdot \mid \eta^*) : A \Subset \mathbb{L}\big\}$ on $(\Omega, \mathcal{F})$ which, under the $\mathcal{B}_1$-summability condition $\sum_{A \ni x} \|\phi_A(\,\cdot \mid \eta^*)\| < \infty$, $x \in \mathbb{L}$, leads to Gibbsian specifications $\Pi^{\Phi(\,\cdot|\eta^*)}$ on $(\Omega, \mathcal{F})$. The study of *quenched disorder* amounts to the determination of features of the phase diagram and properties of extremal measures of the models defined by these specifications for *fixed* typical choices of the disorder. More precisely, the interest focuses on features and properties valid $\mathbb{P}$-almost surely, that is, for almost all disorder configurations η^*. In contrast, in the analysis of *annealed disorder* there is a previous $\mathbb{P}$-average over the Gibbs weights of the magnitude in question. This averaging makes annealed disorder much easier to study than its quenched counterpart.

Let me mention three well-known examples.

Random-field Ising model: It represents an Ising model with a random independent magnetic field at each site. That is, $\mathbb{L}^* = \mathbb{L}$, $S^* \subset \mathbb{R}$, $\mathbb{P}$ is the product of reasonable single-site distributions (ex. Gaussian or of bounded support) and the disordered interaction yields the formal Hamiltonians

$$-\sum_{\langle x,y\rangle} \sigma_x \sigma_y - h \sum_x \eta^*_x \, \sigma_x \; . \tag{5.57}$$

Edwards-Anderson spin glass: It corresponds to a zero-field Ising model with random independent couplings. Therefore the disorder acts on the bond-lattice, $\mathbb{L}^* = \big\{\{x,y\} : x, y \in \mathbb{L}, |x-y| = 1\big\}$, $S^* \subset \mathbb{R}$, $\mathbb{P}$ is a product of reasonable single-bond distributions and the formal disordered Hamiltonians are

$$-\sum_{\langle x,y\rangle} \eta^*_{\{x,y\}} \, \sigma_x \sigma_y \; . \tag{5.58}$$

Griffiths singularity (GriSing) random field: It describes an Ising model on the random lattice determined by independent site-percolation. Thus $\mathbb{L}^* = \mathbb{L}$, $S^* = \{0,1\}$ and $\mathbb{P}$ is the product of Bernoulli variables taking value 1 with probability p and 0 with probability $1-p$. The formal Hamiltonians are

$$-\sum_{\langle x,y\rangle} \eta^*_x \eta^*_y \, \sigma_x \, \sigma_y - h \sum_x \eta^*_x \sigma_x \; . \tag{5.59}$$

This model was introduced by Griffiths to illustrate the appearance of singularities, now known as Griffiths singularities, that prevent the infinitely derivable disordered free energy from being analytic.

A natural approach to the study of quenched disorder is to place spin and disorder variables on the same footing and consider a "grand-canonical ensemble" on the product space $\Omega \times \Omega^*$ from which quenched measures are obtained as projections on Ω. In this way quenching is incorporated within the grand-canonical average and hence constitutes an "annealed approach to quenching disorder". Such an approach was first advocated by Morita in the sixties [50]. Formally, this corresponds to considering the *skew space* $(\Omega \times \Omega^*, \mathcal{F} \times \mathcal{F}^*)$ and joint-variable measures K obtained as weak limits

$$K(d\omega, d\eta^*) \;=\; \lim_{n\to\infty} \lim_{m\to\infty} \mathbb{P}_{\Lambda^*_{r_m}}(d\eta^* \mid \alpha^*)\, \pi_{\Lambda_{s_n}}^{\Phi(\,\cdot\,|\eta^*_{\Lambda^*_{r_m}}\alpha^*)}(d\omega \mid \sigma) \tag{5.60}$$

where (r_n) and (s_m) are diverging sequence of box sizes and α^* and σ are disorder and spin boundary conditions. Such limits always exist, by compactness, if S^* is compact.

Morita's theory supposed the existence of an effective Hamiltonian for the joint variables, that is, the Gibbsianness of these measures K. It is now known that this assumption is false in general [28, 66, 65]. A rough explanation of this fact comes from the fact that a joint effective Hamiltonian should deal with terms of the form

$$\log\left(\frac{\mathbb{P}(\eta^*_{\Lambda^*})}{Z_\Lambda^{\Phi(\,\cdot\,|\eta^*)}}\right) \tag{5.61}$$

which become ill-defined, in the limit $\Lambda \to \mathbb{L}$, precisely when there are Griffiths singularities (or other phase transitions).

As an illustration, let us consider the conditional probability at the origin of a measure of type (5.60) for the GriSing model. After a brief verification we see that

$$\begin{aligned}
&K_{\{0\}}\big(\eta^*_0 = +1 \,\big|\, \sigma\,\eta^*\big)\\
&\quad= \lim_{n\to\infty} \frac{\mathbb{P}(\eta^*_0 = 1)\, \gamma_{\Lambda_n}^{\Phi(\,\cdot\,|1_{\{0\}}\eta^*_{\Lambda_n\setminus\{0\}})}(\sigma_{\Lambda_n} \mid \sigma_{\Lambda_n^{\mathrm{c}}})}{\mathbb{P}(\eta^*_0 = 1)\, \gamma_{\Lambda_n}^{\Phi(\,\cdot\,|1_{\{0\}}\eta^*_{\Lambda_n\setminus\{0\}})}(\sigma_{\Lambda_n\setminus\{0\}} \mid \sigma_{\Lambda_n^{\mathrm{c}}}) \;+\; \mathbb{P}(\eta^*_0 = 0)\, \gamma_{\Lambda_n}^{\Phi(\,\cdot\,|0_{\{0\}}\eta^*_{\Lambda_n\setminus\{0\}})}(\sigma_{\Lambda_n\setminus\{0\}} \mid \sigma_{\Lambda_n^{\mathrm{c}}})}\\
&\quad= \frac{p}{p\,\Delta_{\mathrm{QL}} + (1-p)\,\Delta_{\mathrm{NQL}}}
\end{aligned} \tag{5.62}$$

(if necessary, Λ_n should be replaced by Λ_{r_n}). The term

$$\Delta_{\mathrm{QL}} \;=\; \lim_{n\to\infty} \frac{\gamma_{\Lambda_n}^{\Phi(\,\cdot\,|1_{\{0\}}\eta^*_{\Lambda_n\setminus\{0\}})}(\sigma_{\Lambda_n\setminus\{0\}} \mid \sigma_{\Lambda_n^{\mathrm{c}}})}{\gamma_{\Lambda_n}^{\Phi(\,\cdot\,|1_{\{0\}}\eta^*_{\Lambda_n\setminus\{0\}})}(\sigma_{\Lambda_n} \mid \sigma_{\Lambda_n^{\mathrm{c}}})} \;=\; 1 + \frac{\gamma_{\{0\}}^{\Phi(\,\cdot\,|1_{\{0\}}\eta^*_{\Lambda_n\setminus\{0\}})}(-\sigma_0 \mid \sigma_{\{0\}^{\mathrm{c}}})}{\gamma_{\{0\}}^{\Phi(\,\cdot\,|1_{\{0\}}\eta^*_{\Lambda_n\setminus\{0\}})}(\sigma_0 \mid \sigma_{\{0\}^{\mathrm{c}}})} \tag{5.63}$$

is perfectly continuous with respect to both η^* and σ. The discontinuity appears in

$$\Delta_{\mathrm{NQL}} \;=\; \lim_{n\to\infty} \frac{\gamma_{\Lambda_n}^{\Phi(\,\cdot\,|0_{\{0\}}\eta^*_{\Lambda_n\setminus\{0\}})}(\sigma_{\Lambda_n\setminus\{0\}} \mid \sigma_{\Lambda_n^{\mathrm{c}}})}{\gamma_{\Lambda_n}^{\Phi(\,\cdot\,|1_{\{0\}}\eta^*_{\Lambda_n\setminus\{0\}})}(\sigma_{\Lambda_n\setminus\{0\}} \mid \sigma_{\Lambda_n^{\mathrm{c}}})} \tag{5.64}$$

because of the presence of the ratio

$$\frac{Z_{\Lambda_n}^{\Phi(\,\cdot\,|1_{\{0\}}\eta^*_{\Lambda_n\setminus\{0\}})}(\sigma_{\Lambda_n^{\mathrm{c}}})}{Z_{\Lambda_n}^{\Phi(\,\cdot\,|0_{\{0\}}\eta^*_{\Lambda_n\setminus\{0\}})}(\sigma_{\Lambda_n^{\mathrm{c}}})} \;=\; \pi_{\Lambda_n\setminus\{0\}}^{\Phi(\,\cdot\,|0_{\{0\}}\eta^*_{\Lambda_n\setminus\{0\}})}\Big(2\cosh\Big[\beta\sum_{|y|=1}\eta^*_y\,\sigma_y\Big] \,\Big|\, \sigma_{\Lambda_n^{\mathrm{c}}}\Big)\,. \tag{5.65}$$

The discontinuity takes place at disorder configurations η^* with more than one percolation clusters, all of them excluding the origin. Then, it is not hard to see [66] that a local modification causing a connection between clusters produces a finite change en the expectation (5.65). The absolute value of this change is is bounded below by a positive constant that does not depend on the distance at which the connection is established. Quasilocality is thereby lost. The point of discontinuity depends only on the disorder variable η^*; the conditioning spin configuration σ is irrelevant.

This type of non-quasilocality is, in my opinion, more subtle and surprising than those analyzed in the previous sections. It appears for values of the disorder that are close to those for which the quenched system has a phase transition. These are precisely the disorder configurations triggering the presence of arbitrary long, but finite, order that leads to Griffiths singularities. In the present model, these configurations are unlikely if p is smaller than the critical percolation probability. Thus, the model is almost quasilocal for those p [66]. The random-field Ising model in three or more dimensions has a more dramatical feature [28]. At low temperature there is a full measure set of random fields for which the quenched model has a phase transition. Hence the joint measure is *almost non-quasilocal* that is, the set of discontinuities has full measure. On the other hand, the joint measures of finite-range disordered models can be proven to be weakly Gibbsian [30], hence we have here the largest possible divorce between the notions of almost quasilocality and weak Gibbsianness.

There is another sense in which the non-Gibbsianness of joint disorder measures is complementary to that caused by renormalization transformations or spin dynamics. In the previous cases there was a two-slice system, defined on $\Omega \times \Omega'$ that was Gibbsian, and non-Gibbsianness appeared upon projection to the Ω' variables. In the present case, the two-slice model on $\Omega \times \Omega^*$ is non-Gibbsian, while projections to each of the slices can restore Gibbsianness. [The Ω projection is the quenched average of Gibbsian measures which can be Gibbsian, while the Ω^*-projection is the disorder measure $\mathbb{P}$ which is usually a product measure, and thus trivially Gibbsian.]

In fact, the non-Gibbsianness of the joint measures turned out to be beneficial to Morita's approach. Indeed, besides the hypothetical joint Hamiltonian, Morita's theory included other assumptions equally contradictory with Gibbsianness. And yet, the approach was undeniably successful. Non-Gibbsianness does solve such a paradox [32]. First, it removes inconsistencies related with the untenable Gibbsian hypothesis, and second, it allows for a rigorous justification of the equations solving the model. This has been a remarkable achievement of non-Gibbsianness theory.

Acknowledgments

Our field is blessed by the fact that its founding fathers set up a friendly, open an unassuming style of work, where ideas are discussed generously and freely. I thank wholeheartedly Aernout, Anton, François and Frank for the immense task of setting up a school fully inscribed in such a tradition. I also thank Aernout van Enter for an invaluable critical reading of the manuscript.

References

[1] J. Bricmont, A. Kupiainen, and R. Lefevere. Renormalization group pathologies and the definition of Gibbs states. *Comm. Math. Phys.*, 194(2):359–388, 1998.

[2] J. Bricmont, A. Kupiainen, and R. Lefevere. Renormalizing the renormalization group pathologies. *Phys. Rep.*, 348(1-2):5–31, 2001. Renormalization group theory in the new millennium, II.

[3] S. Dachian and B. S. Nahapetian. Description of random fields by means of one-point conditional distributions and some applications. *Markov Proc. Rel. Fields*, 7:193–214, 2001.

[4] S. Dachian and B. S. Nahapetian. Description of specifications by means of probability distributions in small volumes under condition of very weak positivity. *J. Stat. Phys.*, 117:281–300, 2004.

[5] R. L. Dobrushin. A Gibbsian representation for non-Gibbs fields. Lecture given at the workshop "Probability and Physics", Renkum, The Netherlands, 1995.

[6] R. L. Dobrushin. The description of a random field by means of conditional probabilities and conditions of its regularity. *Theor. Prob. Appl.*, 13:197–224, 1968.

[7] R. L. Dobrushin and S. B. Shlosman. Gibbsian description of "non-Gibbsian" fields. *Russian Math. Surveys*, 52:285–97, 1997.

[8] R. L. Dobrushin and S. B. Shlosman. "Non-Gibbsian" states and their Gibbs description. *Comm. Math. Phys.*, 200(1):125–179, 1999.

[9] T. C. Dorlas and A. C. D. van Enter. Non-Gibbsian limit for large-block majority-spin transformations. *J. Stat. Phys.*, 55:171–181, 1989.

[10] R. Fernández, A. Le Ny, and F. Redig. Restoration of Gibbsianness for projected and FKG renormalized measures. *Bull. Braz. Math. Soc. (N.S.)*, 34(3):437–455, 2003.

[11] R. Fernández, A. Le Ny, and F. Redig. Variational principle and almost quasilocality for renormalized measures. *J. Statist. Phys.*, 111(1-2):465–478, 2003.

[12] R. Fernández and G. Maillard. Chains with complete connections and one-dimensional Gibbs measures. *Electron. J. Probab.*, 9:145–76, 2004.

[13] R. Fernández and G. Maillard. Construction of a specification from its singleton part, 2005. Paper 05-288 at `http://www.ma.utexas.edu/mp_arc`.

[14] R. Fernández and C.-Ed. Pfister. Global specifications and non-quasilocality of projections of Gibbs measures. *Ann. Probab.*, 25:1284–1315, 1997.

[15] R. Fernández and A. Toom. Non-gibbsianness of the invariant measures of of non-reversible cellular automata with totally asymmetric noise. *Asthérisque*, 287:71–87, 2003.

[16] M. E. Fisher. Scaling, universality and renormalization group theory. In F. J. W. Hahne, editor, *Critical Phenomena (Stellenbosch 1982)*, pages 1–139. Springer-Verlag (Lecture Notes in Physics #186), Berlin–Heidelberg–New York, 1983.

[17] K. Gawędzki, R. Kotecký, and A. Kupiainen. Coarse-graining approach to first-order phase transitions. In *Proceedings of the symposium on statistical mechanics of phase transitions—mathematical and physical aspects (Trebon, 1986)*, volume 47, pages 701–724, 1987.

[18] H.-O. Georgii. *Gibbs Measures and Phase Transitions.* Walter de Gruyter (de Gruyter Studies in Mathematics, Vol. 9), Berlin–New York, 1988.

[19] N. Goldenfeld. *Lectures on Phase Transitions and the Renormalization Group.* Addison-Wesley (Frontiers in Physics 85), 1992.

[20] R. B. Griffiths. Mathematical properties of renormalization-group transformations. *Physica*, 106A:59–69, 1981.

[21] R. B. Griffiths and P. A. Pearce. Position-space renormalization-group transformations: Some proofs and some problems. *Phys. Rev. Lett.*, 41:917–920, 1978.

[22] R. B. Griffiths and P. A. Pearce. Mathematical properties of position-space renormalization-group transformations. *J. Stat. Phys.*, 20:499–545, 1979.

[23] G. Grimmett. A theorem about random fields. *Bull. London Math. Soc.*, 5:81–4, 1973.

[24] G. Grimmett. The stochastic random-cluster process and the uniqueness of random-cluster measures. *Ann. Prob.*, 23:1461–510, 1995.

[25] O. Häggström. Almost sure quasilocality fails for the random-cluster model on a tree. *J. Stat. Phys.*, 84:1351–61, 1996.

[26] R. B. Israel. Banach algebras and Kadanoff transformations. In J. Fritz, J. L. Lebowitz, and D. Szász, editors, *Random Fields – Rigorous Results in Statistical Mechanics and Quantum Field Theory*, volume II, pages 593–608. North-Holland, Amsterdam, 1981. Colloquia Mathematica Societatis Janos Bolyai **27**, Esztergom (Hungary), 1979.

[27] O. K. Kozlov. Gibbs description of a system of random variables. *Probl. Inform. Transmission*, 10:258–65, 1974.

[28] C. Külske. (Non-) Gibbsianness and phase transitions in random lattice spin models. *Markov Process. Related Fields*, 5(4):357–383, 1999.

[29] C. Külske. On the Gibbsian nature of the random field Kac model under block-averaging. *J. Statist. Phys.*, 104(5-6):991–1012, 2001.

[30] C. Külske. Weakly Gibbsian representations for joint measures of quenched lattice spin models. *Probab. Theory Related Fields*, 119(1):1–30, 2001.

[31] C. Külske. Analogues of non-Gibbsianness in joint measures of disordered mean field models. *J. Statist. Phys.*, 112(5-6):1079–1108, 2003.

[32] C. Külske. How non-Gibbsianness helps a metastable Morita minimizer to provide a stable free energy. *Markov Process. Related Fields*, 10(3):547–564, 2004.

[33] C. Külske, A. Le Ny, and F. Redig. Relative entropy and variational properties of generalized Gibbsian measures. *Ann. Probab.*, 32(2):1691–1726, 2004.

[34] O. E. Lanford III and D. Ruelle. Observables at infinity and states with short range correlations in statistical mechanics. *Commun. Math. Phys.*, 13:194–215, 1969.

[35] J. L. Lebowitz and C. Maes. The effect of an external field on an interface, entropic repulsion. *J. Stat. Phys.*, 46:39–49, 1987.

[36] J. L. Lebowitz and R. H. Schonmann. Pseudo-free energies and large deviations for non-Gibbsian FKG measures. *Prob. Th. Rel. Fields*, 77:49–64, 1988.

[37] R. Lefevere. Variational principle for some renormalized measures. *J. Statist. Phys.*, 96(1-2):109–133, 1999.

[38] J. Lörinczi. Non-Gibbsianness of the reduced SOS-measure. *Stoch. Proc. Appl.*, 74:83–8, 1998.

[39] J. Lörinczi and K. Vande Velde. A note on the projection of Gibbs measures. *J. Stat. Phys.*, 77:881–7, 1994.

[40] J. Lörinczi and M. Winnink. Some remarks on almost Gibbs states. In N. Boccara, E. Goles, S. Martinez, and P. Picco, editors, *Cellular Automata and Cooperative Systems*, pages 423–432, Dordrecht, 1993. Kluwer.

[41] C. Maes, A. Van Moffaert, and F. Redig. Almost Gibbsian versus weakly Gibbsian measures. *Stoch. Proc. Appl.*, 79:1–15, 1998.

[42] C. Maes, F. Redig, F. Takens, A. van Moffaert, and E. Verbitskiy. Intermittency and weak Gibbs states. *Nonlinearity*, 13(5):1681–1698, 2000.

[43] C. Maes and K. Vande Velde. The (non-)Gibbsian nature of states invariant under stochastic transformations. *Physica A*, 206:587–603, 1994.

[44] C. Maes and K. Vande Velde. Relative energies for non-Gibbsian states. *Commun. Math. Phys.*, 189:277–86, 1997.

[45] D. Makowiec. Gibbsian versus non-Gibbsian nature of stationary states for Toom probabilistic cellular automata via simulations. *Phys. Rev. E*, 55:6582–8, 1997.

[46] D. Makowiec. Stationary states of Toom cellular automata in simulations. *Phys. Rev. E*, 60:3787–95, 1999.

[47] F. Martinelli and E. Olivieri. Some remarks on pathologies of renormalization-group transformations. *J. Stat. Phys.*, 72:1169–1177, 1993.

[48] F. Martinelli and E. Olivieri. Instability of renormalization-group pathologies under decimation. *J. Stat. Phys.*, 79:25–42, 1995.

[49] F. Martinelli and E. Scoppola. A simple stochastic cluster dynamics: rigorous results. *J. Phys. A*, 24:3135–57, 1991.

[50] T. Morita. Statistical mechanics of quenched sollid solutions with applications to magnetically diluted alloys. *J. Math. Phys.*, 5:1402–5, 1964.

[51] A. Le Ny and F. Redig. Short time conservation of Gibbsianness under local stochastic evolutions. *J. Statist. Phys.*, 109(5-6):1073–1090, 2002.

[52] M. Oliveira and A. Petri. Boltzmann temperature in out-of-equilibrium lattice gas. `ArXiv cond-mat/0511263`, 2005.

[53] C.-E. Pfister and K. Vande Velde. Almost sure quasilocality in the random cluster model. *J. Stat. Phys.*, 79:765–74, 1995.

[54] C.-Ed. Pfister. Thermodynamical aspects of classical lattice systems. In *In and out of equilibrium (Mambucaba, 2000)*, volume 51 of *Progr. Probab.*, pages 393–472. Birkhäuser Boston, Boston, MA, 2002.

[55] C. Preston. *Random Fields*. Springer-Verlag (Lecture Notes in Mathematics #534), Berlin–Heidelberg–New York, 1976.

[56] R. H. Schonmann. Projections of Gibbs measures may be non-Gibbsian. *Commun. Math. Phys.*, 124:1–7, 1989.

[57] A. D. Sokal. Existence of compatible families of proper regular conditional probabilities. *Z. Wahrscheinlichkeitstheorie verw. Geb.*, 56:537–548, 1981.

[58] W. G. Sullivan. Potentials for almost Markovian random fields. *Commun. Math. Phys.*, 33:61–74, 1973.

[59] A. C. D. van Enter. Ill-defined block-spin transformations at arbitrarily high temperatures. *J. Stat. Phys.*, 83:761–5, 1996.

[60] A. C. D. van Enter. A remark on the notion of robust phase transitions. *J. Statist. Phys.*, 98(5-6):1409–1416, 2000.

[61] A. C. D. van Enter and R. Fernández. A remark on different norms and analyticity for many-particle interactions. *J. Stat. Phys.*, 56:965–972, 1989.

[62] A. C. D. van Enter, R. Fernández, F. den Hollander, and F. Redig. Possible loss and recovery of Gibbsianness during the stochastic evolution of Gibbs measures. *Comm. Math. Phys.*, 226(1):101–130, 2002.

[63] A. C. D. van Enter, R. Fernández, and R. Kotecký. Pathological behavior of renormalization-group maps at high fields and above the transition temperature. *J. Stat. Phys*, 79:969–92, 1995.

[64] A. C. D. van Enter, R. Fernández, and A. D. Sokal. Regularity properties and pathologies of position-space renormalization-group transformations: Scope and limitations of Gibbsian theory. *J. Stat. Phys.*, 72:879–1167, 1993.

[65] A. C. D. van Enter, C. Külske, and C. Maes. Comment on: Critical behavior of the randomly spin diluted 2d ising model: A grand ensemble approach, by r. kühn. *Phys. Rev. Lett.*, 84:6134, 2000.

[66] A. C. D. van Enter, C. Maes, R. H. Schonmann, and S. B. Shlosman. The Griffiths singularity random field. In *On Dobrushin's way. From probability theory to statistical physics*, volume 198 of *Amer. Math. Soc. Transl. Ser. 2*, pages 51–58. Amer. Math. Soc., Providence, RI, 2000.

[67] A. C. D. van Enter, C. Maes, and S. B. Shlosman. Dobrushin's program on Gibbsianity restoration: weakly Gibbs and almost Gibbs random fields. In *On Dobrushin's way. From probability theory to statistical physics*, volume 198 of *Amer. Math. Soc. Transl. Ser. 2*, pages 59–70. Amer. Math. Soc., Providence, RI, 2000.

[68] A. C. D. van Enter and S. B. Shlosman. (Almost) Gibbsian description of the sign fields of SOS fields. *J. Statist. Phys.*, 92(3-4):353–368, 1998.

[69] A. C. D. van Enter and E. A. Verbitskiy. On the variational principle for generalized Gibbs measures. *Markov Process. Related Fields*, 10(3):411–434, 2004.

[70] D. Williams. *Probability with Martingales*. Cambridge University Press, Cambridge, 1991.