Markov Processes Relat. Fields **5**, 357–383 (1999)

Markov Processes and Related Fields
© Polymat, Moscow 1999

(Non-) Gibbsianness and Phase Transitions in Random Lattice Spin Models*

C. Külske
Weierstraß–Institut für Angewandte Analysis und Stochastik, Mohrenstraße 39, D-10117 Berlin, Germany. E-mail: kuelske@wias-berlin.de

Received April 16, 1999, revised July 6, 1999

Abstract. We consider disordered lattice spin models with finite-volume Gibbs measures $\mu_\Lambda[\eta](d\sigma)$. Here σ denotes a lattice spin variable and η a lattice random variable with product distribution P describing the quenched disorder of the model. We ask: when will the joint measures $\lim_{\Lambda\uparrow\mathbb{Z}^d} \mathsf{P}(d\eta)\mu_\Lambda[\eta](d\sigma)$ be [non-] Gibbsian measures on the product of spin space and disorder space? We obtain general criteria for both Gibbsianness and non-Gibbsianness providing an interesting link between phase transitions at a fixed random configuration and Gibbsianness in product space: loosely speaking, a discontinuity in the quenched Gibbs expectation can lead to non-Gibbsianness, (only) if it can be observed on the spin observable conjugate to the independent disorder variables.

Our main specific example is the random field Ising model in any dimension for which we show almost sure [almost sure non-] Gibbsianness for the single- [multi-] phase region. We also discuss models with disordered couplings, including spin glasses and random bond ferromagnets, where various mechanisms are responsible for [non-] Gibbsianness.

KEYWORDS: disordered systems, Gibbs measures, non-Gibbsianness, random field model, random bond model, spin glass

AMS SUBJECT CLASSIFICATION: 82B44, 82B26, 82B20

1. Introduction

The purpose of this paper is to present a class of measures on discrete lattice spins showing a rich behavior with respect to their Gibbsianness properties. The

*Supported by the DFG Schwerpunkt 'Wechselwirkende stochastische Systeme hoher Komplexität'

examples we consider turn up in a natural context of well-studied disordered systems.

Given a random lattice system, such as the random field Ising model, we look at the *joint* distribution of spins and random variables describing the disorder. It is now very natural from a probabilistic point of view to consider the corresponding *joint measures* on the skew space resulting from the a priori distribution of the disorder variables. Taking the infinite-volume limit leads to infinite-volume measures on the skew space. We will investigate the Gibbsianness properties of such measures, for general finite range potentials. As we will see, this gives rise to a whole family of interesting examples of measures with non-trivial behavior.

Why consider these measures? Gibbs measures are the basic objects for a mathematically rigorous description of equilibrium statistical mechanics. They are characterized by the fact that their finite-volume conditional expectations can be written in terms of an absolutely summable interaction potential. The failure of the Gibbsian property is linked to the emergence of long-range correlations or hidden phase transitions.

In the theory of disordered systems on the other hand, it is important to understand the effects caused by 'exceptional realizations' of the disorder variables. It might happen e.g. that a Gibbs state (even a ground state) depends in a very non-local way on the disorder variables for untypical realizations, while it is usually locally determined. This causes difficulties for the investigation of such systems. Take as an example the analysis of the random field Ising model of [2] where the concept of a 'bad region' in space (where the realization of the random variable was untypically large) is introduced — one then needs to show that it is 'irrelevant' to show ferromagnetic order. Even if the effect of exceptional realizations is not dominant on the average, some effects result, like Griffiths' singularities. Indeed, it is the exceptionally large connected clusters of occupied sites which are responsible for their emergence in the GriSing random field [15]. (An analogy between both notions of exceptional realizations was pointed out in [3].) In any case, it is a very important theme in disordered systems to understand the potentially non-local behavior of the system that is caused by exceptional behavior of the quenched randomness.

Now, as we will see in our general investigation, the [non-] Gibbsianness of the joint measures is related in an interesting way to the [non-] locality of certain expectations of random Gibbs measures as a function of the disorder variables. Since such a non-locality can arise in a variety of different ways, there is a variety of different 'mechanisms' for non-Gibbsianness. So, the much-disputed phenomenon of non-Gibbsianness becomes related in a somewhat surprising way to continuity questions of the random Gibbs measures on the spins with respect to disorder, or, in other words, phase transitions induced by changes of the disorder variables.

The present investigation was motivated by the special recent example of the Ising ferromagnet with site dilution ('GriSing random field') that was shown to

be non-Gibbsian but almost Gibbsian in [12] where an interesting realization of the disorder variables leading to 'non-continuity' was found. One should however note that it is not just the large connected occupied clusters which are present in a.e. realization that are responsible for this non-locality. This shows that one has to be precise about the kind of nonlocality in discussion and analogies to Griffiths singularities should not be pushed to far.

Mathematically the analysis was simplified for the GriSing random field because the system considered breaks down into finite pieces. This is of course not true in most of the systems of interest. Such a 'non-decoupling' is going to be a complication of the general treatment we are going to present, as we will see. (While the decoupling argument is somewhat simpler for the random field Ising model treated in Theorem 1.1, it is more involved for the random coupling models treated in Theorem 3.1.)

Let us remark that there has been some discussion during the last years about numerous examples of non-Gibbsian measures, to what extent the failure of the Gibbsian property has to be taken seriously, and what suitable generalizations of Gibbsianness should be (see e.g. [3, 8, 9, 13, 23], references therein, and the basic paper [11]). While this discussion still does not seem to be finished, the answers seem to depend on the specific situation. Our point in this context is less a general philosophical one, but rather to provide interesting examples that show (non-) Gibbsianness in a slightly different light related to important issues in the theory of random Gibbs measures.

Basic definitions

Denote by $\Omega = \Omega_0^{\mathbb{Z}^d}$ the space of *spin configurations* $\sigma = (\sigma_x)_{x\in\mathbb{Z}^d}$, where Ω_0 is a finite set. Similarly we denote by $\mathcal{H} = \mathcal{H}_0^{\mathbb{Z}^d}$ the space of *disorder variables* $\eta = (\eta_x)_{x\in\mathbb{Z}^d}$ entering the model, where $\mathcal{H}_0$ is a finite set. Each copy of $\mathcal{H}_0$ carries a measure $\nu(d\eta_x)$ and $\mathcal{H}$ carries the product-measure over the sites, $\mathsf{P} = \nu^{\otimes_{\mathbb{Z}^d}}$. We denote the corresponding expectation by E. The space of joint configurations $\Omega \times \mathcal{H} = (\Omega_0 \times \mathcal{H}_0)^{\mathbb{Z}^d}$ is called the *skew space*. It is equipped with the product topology.

We consider disordered models whose formal *infinite-volume Hamiltonian* can be written in terms of disordered potentials $(\Phi_A)_{A\subset\mathbb{Z}^d}$,

$$H^\eta(\sigma) = \sum_{A\subset\mathbb{Z}^d} \Phi_A(\sigma,\eta) \tag{1.1}$$

where Φ_A depends only on the spins and disorder variables in A. We assume for simplicity *finite range*, i.e. that $\Phi_A = 0$ for $\operatorname{diam} A > r$. A lot of disordered models can be cast into this form.

For fixed realization of the disorder variable η we denote by $\mu_\Lambda^{\sigma^{\text{b.c.}}}[\eta]$ the corresponding *finite-volume Gibbs measures* in $\Lambda\subset\mathbf{Z}^d$ with boundary condition $\sigma^{\text{b.c.}}$. As usual, they are the probability measures on Ω that are given by the

formula

$$\mu_\Lambda^{\sigma^{\text{b.c.}}}[\eta](f) := \frac{\sum_{\sigma_\Lambda} f(\sigma_\Lambda \sigma^{\text{b.c.}}_{\mathbf{Z}^d\setminus\Lambda}) e^{-\sum_{A\cap\Lambda\neq\emptyset}\Phi_A(\sigma_\Lambda\sigma^{\text{b.c.}}_{\mathbf{Z}^d\setminus\Lambda},\eta)}}{\sum_{\sigma_\Lambda} e^{-\sum_{A\cap\Lambda\neq\emptyset}\Phi_A(\sigma_\Lambda\sigma^{\text{b.c.}}_{\mathbf{Z}^d\setminus\Lambda},\eta)}} \tag{1.2}$$

for any bounded measurable observable $f : \Omega \to \mathbf{R}$. The finite-volume summation is over $\sigma_\Lambda \in \Omega_0^\Lambda$. The symbol $\sigma_\Lambda \sigma^{\text{b.c.}}_{\mathbf{Z}^d\setminus\Lambda}$ denotes the configuration in Ω that is given by σ_x for $x \in \Lambda$ and by $\sigma_x^{\text{b.c.}}$ for $x \in \mathbf{Z}^d\setminus\Lambda$.

We look at spins and disorder variables at the same time and define *joint variables* $\xi_x = (\sigma_x, \eta_x) \in \Omega_0 \times \mathcal{H}_0$. The objects of main interest will then be the corresponding *finite-volume joint measures* $K_\Lambda^{\sigma^{\text{b.c.}}}$. They are the probability measures on the skew space $(\Omega_0 \times \mathcal{H}_0)^{\mathbf{Z}^d}$ that are given by the formula

$$K_\Lambda^{\sigma^{\text{b.c.}}}(F) := \int \mathsf{P}(d\eta) \int \mu_\Lambda^{\sigma^{\text{b.c.}}}[\eta](d\sigma) F(\sigma, \eta) \tag{1.3}$$

for any bounded measurable joint observable $F : \Omega \times \mathcal{H} \to \mathbf{R}$. We will consider the following examples in more detail.

(i) *The random field Ising model (RFIM).*
The single spin space is $\Omega_0 = \{-1, 1\}$, the Hamiltonian is

$$H^\eta(\sigma) = -J \sum_{<x,y>} \sigma_x \sigma_y - h \sum_x \eta_x \sigma_x \tag{1.4}$$

where the formal sum is over nearest neighbors $< x, y >$ and $J, h > 0$. The disorder variables are given by the random fields η_x that are i.i.d. with single-site distribution ν that is supported on a finite set $\mathcal{H}_0$.

The joint variables we will consider are given in a natural way by the Ising spin and the random field at the same site, i.e. $\xi_x = (\sigma_x, \eta_x)$; ξ_x is thus 4-valued in the case of symmetric Bernoulli distribution.

(ii) *Ising models with random couplings: random bond, EA-spin glass.*
The single spin space is $\Omega_0 = \{-1, 1\}$. The Hamiltonian is

$$H^\eta(\sigma) = -\sum_{x,\,e} J_{x,e} \sigma_x \sigma_{x+e} \tag{1.5}$$

where the formal sum is over sites $x \in \mathbf{Z}^d$ and the nearest neighbor vectors in the positive lattice directions, i.e.

$$e \in \{(1, 0, 0, \ldots, 0), (0, 1, 0, \ldots, 0), \ldots, (0, 0, \ldots, 1)\} =: \mathcal{E}.$$

The random variables $J_{x,e}$ take finitely many values, independently over the 'bonds' x, e. Specific distributions we will consider are for example as follows:

(a) random bond: $J_{x,e}$ takes values $J^1, J^2 > 0$;

(b) EA-spin glass: symmetric (non-degenerate) 3-valued, $J_{x,e}$ takes values $-J, 0, J$ with $\nu(J_{x,e} = J) = \nu(J_{x,e} = -J)$, $0 < \nu(J_{x,e} = 0) < 1$.

We define the *joint spins* by the Ising spin and the collection of adjacent couplings pointing in the positive direction, i.e. $\xi_x = (\sigma_x, \eta_x) = (\sigma_x, (J_{x,e})_{e \in \mathcal{E}})$. It is thus 16-valued in dimension 3 in case (a).

We think of the random field Ising model for a moment to motivate what we are going to do. Recall that, in two dimensions, for almost every realization of the random fields η with respect to P there exists a unique infinite-volume Gibbs measure $\mu(\eta)$ (see [1]). In three or more dimensions, for low temperatures and 'small disorder' there exist ferromagnetically ordered phases $\mu^{+,-}(\eta)$ obtained by different boundary conditions [2]. In contrast to the GriSing example of [12] we can hence consider various infinite-volume versions of the form '$\mathsf{P}(d\eta)\mu(\eta)(d\sigma)$'.

The most general thing now that we can reasonably do, is to fix *any* boundary condition $\sigma^{\text{b.c.}}$. Then, due to compactness, there are always subsequences such that the corresponding $K_\Lambda^{\sigma^{\text{b.c.}}}(d\xi)$ converges weakly to a probability measure on the skew space that we call $K(d\xi)$. Note that this measure can in general depend on the boundary condition and the particular choice of the subsequence in $d \geq 2$. It can be shown that by conditioning $K(d\xi) = K(d\sigma, d\eta)$ on the disorder variable η one obtains a (not necessarily extremal) random infinite-volume Gibbs measure, for P-almost every η.[1] The aim of this paper is to investigate the question:

When are the weak limit points of $K_\Lambda^{\sigma^{\text{b.c.}}}(d\xi)$ Gibbs measures on the skew-space? When are they almost [almost surely not] Gibbs?

To rephrase the first question in the language of [12]: is there an annealed representation of the quenched system? This investigation is about continuity properties of conditional expectations. Throughout the paper, following [23], we will use the following notion of continuity that involves only uniquely defined finite-volume events.

Definition 1.1. A point $\xi \in \Omega \times \mathcal{H}$ is called a *good configuration* for K, if

$$\sup_{\substack{\xi^+, \xi^- \\ \Lambda : \Lambda \supset V}} |K(\tilde{\xi}_x \mid \xi_{V \setminus x}, \xi^+_{\Lambda \setminus V}) - K(\tilde{\xi}_x \mid \xi_{V \setminus x}, \xi^-_{\Lambda \setminus V})| \to 0 \tag{1.6}$$

with $V \uparrow \mathbf{Z}^d$, for any site $x \in \mathbf{Z}^d$, for any $\tilde{\xi}_x \in \mathcal{H}_0$. Call ξ *bad*, if it is not good.

[1] A reader who is familiar with meta-states will recognize that this measure $K(d\sigma \mid \eta)$ is precisely the barycenter of the (corresponding) Aizenman–Wehr meta-state, see e.g. Newman [24]. For more general information about meta-states and random symmetry breaking see [25–28], [17–20].

As usual we have written $\xi_A = (\xi_x)_{x \in A}$ (and will also do so for σ_A, η_A).

In words, good configurations are the points ξ where the family of conditional expectations of K is equicontinuous with respect to the parameter Λ.

We recall: if there are no bad configurations, the measure K is Gibbsian (see [23]). If Gibbsianness does not hold, one can ask for the K-measure of the set of bad configurations. We say that K is almost Gibbsian, if this set has K-measure zero. If it has K-measure one, we say that K is almost surely non-Gibbsian. (See also the beginning of the next section.)

In the remainder of the paper we will prove criteria that ensure that a configuration (η, σ) is good or bad (see Propositions 2.1–2.4, 3.1–3.2). It might not be very intuitive at first sight to understand why such measures can ever be non-Gibbsian. Let us stress the following facts: surely, the conditional expectation of the spin variable σ_x, given the joint variable $\xi = (\sigma, \eta)$ away from x *and* η_x, is a local function, given by the local specifications. Trivially, the conditional expectation of the disorder variable η_x, given η away from x, is a local function — it is even independent. However, the conditional expectation of η_x, given η *and* σ away from x, can be highly nontrivial. Curiously, it turns out that the fact whether this conditional expectation is discontinuous in the conditioning will always depend only on η, but not on σ. Thus, it does not matter *which* but only *that* a spin configuration is given.

Rather than presenting our general results at this point, we specialize to the random field Ising model. For this model there is a complete characterization of a bad configuration in terms of the behavior of the finite-volume Gibbs measures that is particularly transparent. We obtain the following theorem.

Theorem 1.1. *Consider a random field Ising model of the form* (1.4)*, in any dimension d. A configuration $\xi = (\eta, \sigma)$ is a bad configuration for any joint measure obtained as a limit point of the finite-volume joint measures*

$$\mathsf{P}(d\eta)\mu_\Lambda^{\sigma^{b.c.}_{\partial\Lambda}}[\eta](d\tilde{\sigma})$$

if and only if

$$\lim_{\Lambda\uparrow\infty} \mu_\Lambda^+[\eta_\Lambda](\tilde{\sigma}_x = 1) > \lim_{\Lambda\uparrow\infty} \mu_\Lambda^-[\eta_\Lambda](\tilde{\sigma}_x = 1) \tag{1.7}$$

for some site x. Thus, whether ξ is good does not depend on σ. Here $\mu_\Lambda^{+,-}$ are the finite-volume Gibbs measures with $+$ (resp. $-$) boundary conditions.

Note that the theorem will also hold (say) for the joint measures corresponding to Dobrushin states that are supposed to exist in $d \geq 4$.[2] Using the known results about the random field model one immediately obtains

Corollary 1.1.

(i) $d = 1$: *K is Gibbsian, for all J, $h > 0$.*

[2] For an existence result of this model in the SOS-approximation, see [4, 16].

(ii) $d = 2$: *K is almost surely Gibbsian for all J, $h > 0$.*

On the other hand, K is not Gibbsian in the following cases:

(a) *$\nu[\eta_x = 0] > 0$, J is sufficiently large, and $h > 0$,*

(b) *J is sufficiently large and $h > 0$ sufficiently small.*

(iii) *$d \geq 3$, ν symmetric, $J > 0$ sufficiently large, $h^2\nu[\eta_x^2]$ positive and sufficiently small. Then any such K is almost surely not Gibbsian.*

Indeed, the almost sure Gibbsianness in $d = 2$ follows from the almost sure absence of ferromagnetism, proved in [1]. That we have non-Gibbsianness in $d \geq 2$ if the support of the random fields contains zero, follows from the fact that the configuration $\xi = (\eta_x \equiv 0, \sigma)$ is a bad one, if J is large enough such that there is ferromagnetic order in the homogeneous Ising ferromagnet. Ferromagnetic order also survives for a random field configuration η_x^{alt} of alternating fields of equal strength, as long as $h > 0$ is sufficiently small. (This can be seen by a Peierls argument). So, similarly, $\xi = (\eta_x^{\text{alt}}, \sigma)$ is a bad configuration for any σ, which proves (b). A.s. non-Gibbsianness under the conditions (iii) follows from the existence of ferromagnetic order, proved in [2].

The organization of the paper is as follows. In Section 2 we investigate the one-site conditional probabilities of K and prove general criteria that ensure that a configuration is good or bad. To do so, we consider the single-site variation of the Hamiltonian with respect to the disorder variable η_x. We rewrite the conditional expectations by treating the terms arising from the variation of η_x as a perturbation. This leads to expressions involving certain expectations of the 'conjugate' spin observable. In the example of the random field model this observable is just the spin σ_x; thus the corresponding criteria in Theorem 1.1 are simply formulated in terms of the magnetization.

In Section 3 we apply our results. We prove Theorem 1.1 about the RFIM. Next we comment on models with decoupling configurations, recalling the Gri-Sing random field of [12] and models with random couplings (including spin glasses) that can be zero. This provides more examples of non-Gibbsian fields. Next we specialize our criteria of Section 2 to models with random couplings, proving Theorem 3.1. Based on this we give a heuristic discussion explaining how the validity of the Gibbsian property can be linked to the absence of random Dobrushin states.

2. Criteria for joint [non-] Gibbsianness

In this section we are going to investigate whether a configuration $\xi = (\eta, \sigma)$ is good or bad for the joint states K. We will obtain criteria that are given in terms of the local specifications. To do so we introduce the single-site variation of the Hamiltonian with respect to disorder (2.2) and use the finite-volume perturbation formula (2.3) to rewrite the conditional expectations of K in the form

of Lemma 2.1. This leads to the characterization of good (resp. bad) configurations of the corollary of Proposition 2.1. As direct consequences thereof, Propositions 2.2 and 2.3 give more convenient conditions that ensure goodness (resp. badness). Under the additional assumption of almost surely convergent Gibbs measures we obtain the slightly less obvious criterion for badness of Proposition 2.4.

Before we start, let us however summarize the facts about the notion of good configuration and its relevance for Gibbsianness, for the sake of clarity.

(i) If ξ is *bad* for K, any version of the conditional expectation $\xi_{\mathbf{Z}^d} \mapsto K(\xi_x \mid \xi_{\mathbf{Z}^d \setminus x})$ must be discontinuous for some site x (use DLR-equation, see [23, Proposition 4.3]).

(ii) Conversely, assume that $\hat{\xi} \in \mathcal{G} := \{\xi; \xi \text{ is good}\}$. Then $\lim_{\Lambda \uparrow \mathbf{Z}^d} K(\xi_x \mid \hat{\xi}_{\Lambda \setminus x})$ exists for any site x and hence also $\lim_{\Lambda \uparrow \mathbf{Z}^d} K(\xi_V \mid \hat{\xi}_{\Lambda \setminus V}) =: \gamma_V(\xi_V \mid \hat{\xi}_{\mathbf{Z}^d \setminus V})$ exists for any finite volume V. If $\mathcal{G}$ has full measure with respect to K, the above limit can be (arbitrarily) extended to a measurable function of the conditioning. It is readily seen to define a version of the conditional expectation $\xi_{\mathbf{Z}^d \setminus V} \mapsto K(\xi_V \mid \xi_{\mathbf{Z}^d \setminus V})$ that is continuous within the set $\mathcal{G}$, i.e. $\xi^{(N)} \to \xi$ with $\xi^{(N)}, \xi \in \mathcal{G}$ implies $K(\xi_V \mid \xi^{(N)}_{\mathbf{Z}^d \setminus V}) \to K(\xi_V \mid \xi_{\mathbf{Z}^d \setminus V})$ (see [23, proof of Proposition 4.4]). In this situation we call K *almost Gibbs*.[3]

In particular, if *every* configuration is good, the measure K has a version of the conditional expectation that is continuous on the whole space and is thus *Gibbs*.

In the sequel it will be important to keep track of the local dependence of various quantities. It will be useful to make this explicit. We use the following

Notation. For the fixed interaction range r we introduce the r-boundary $\partial B = \{x \in \mathbf{Z}^d \setminus B; d(x, B) \leq r\}$. In the same fashion we write $\overline{B} = B \cup \partial B$ and $\partial_- B = \{x \in B; d(x, B^c) \leq r\}$, $B^o = B \setminus \partial_- B$.

In this way we will write e.g. $K_\Lambda^{\sigma^{\text{b.c.}}}(\sigma_\Lambda, \eta_{\overline{\Lambda}}) = \mathsf{P}_{\overline{\Lambda}}(\eta_{\overline{\Lambda}}) \mu_\Lambda^{\sigma^{\text{b.c.}}_{\partial\Lambda}}[\overline{\eta}_\Lambda](\sigma_\Lambda)$ to denote the corresponding probabilities.

To investigate the quantity (1.6) for the infinite-volume joint measure we will look at $K_{\Lambda_N}^{\sigma^{\text{b.c.}}_{\partial\Lambda_N}}$ with finite Λ_N. Next, to investigate the conditional distributions

[3] If $K(\mathcal{G}) = 1$ but $\mathcal{G} \neq \mathcal{H} \times \Omega$, we have: $\mathcal{G}$ is dense in $\mathcal{H} \times \Omega$, since any ball with respect to a metric for the product topology has to have positive K-mass, under the assumption of bounded interactions Φ. Thus the conditional expectation is continuous on $\mathcal{G}$ but necessarily *not* uniformly continuous (because it could be extended to the whole space otherwise).

of ξ_x, it suffices to look at the conditional distributions of η_x. Indeed, we may write (for sufficiently large Λ_N)

$$\begin{aligned} & K_{\Lambda_N}^{\sigma^{\text{b.c.}}_{\partial\Lambda_N}}[\sigma_x;\eta_x \mid \sigma_{\Lambda\setminus x};\eta_{\Lambda\setminus x}] \\ &= K_{\Lambda_N}^{\sigma^{\text{b.c.}}_{\partial\Lambda_N}}[\sigma_x \mid \sigma_{\Lambda\setminus x};\eta_x,\eta_{\Lambda\setminus x}] \times K_{\Lambda_N}^{\sigma^{\text{b.c.}}_{\partial\Lambda_N}}[\eta_x \mid \sigma_{\Lambda\setminus x};\eta_{\Lambda\setminus x}] \end{aligned}$$

where

$$\begin{aligned} & K_{\Lambda_N}^{\sigma^{\text{b.c.}}_{\partial\Lambda_N}}[\sigma_x \mid \sigma_{\Lambda\setminus x};\eta_x,\eta_{\Lambda\setminus x}] \\ &= \frac{\mathsf{E}_{\overline{\Lambda_N}\setminus\Lambda}\,\mu_{\Lambda_N}^{\sigma^{\text{b.c.}}_{\partial\Lambda_N}}[\eta_x,\eta_{\Lambda\setminus x},\tilde{\eta}_{\overline{\Lambda_N}\setminus\Lambda}](\sigma_x,\sigma_{\Lambda\setminus x})}{\sum_{\sigma'_x}\mathsf{E}_{\overline{\Lambda_N}\setminus\Lambda}\,\mu_{\Lambda_N}^{\sigma^{\text{b.c.}}_{\partial\Lambda_N}}[\eta_x,\eta_{\Lambda\setminus x},\tilde{\eta}_{\overline{\Lambda_N}\setminus\Lambda}](\sigma'_x,\sigma_{\Lambda\setminus x})} \\ &= \mu_x^{\sigma_{\partial x}}[\eta_x,\eta_{\partial x}](\sigma_x) \end{aligned} \tag{2.1}$$

where the second equality follows from the application of the compatibility relation for the μ-measures for the inner volume made of the single site x, as soon as $\Lambda\supset\overline{x}$. There is of course no non-locality in this term as a function of $\sigma_{\Lambda\setminus x},\eta_{\Lambda\setminus x}$.

On the other hand we see that, if the conditional η_x-distribution has a non-local behavior as a function of $\sigma_{\Lambda\setminus x},\eta_{\Lambda\setminus x}$, this carries over also to the σ_x-marginal $K_{\Lambda_N}^{\sigma^{\text{b.c.}}_{\partial\Lambda_N}}[\sigma_x \mid \sigma_{\Lambda\setminus x};\eta_{\Lambda\setminus x}] = \int K_{\Lambda_N}^{\sigma^{\text{b.c.}}_{\partial\Lambda_N}}[d\tilde{\eta}_x \mid \sigma_{\Lambda\setminus x};\eta_{\Lambda\setminus x}]\mu_x^{\sigma_{\partial x}}[\tilde{\eta}_x,\eta_{\partial x}](\sigma_x)$ (unless the dependence on $\tilde{\eta}_x$ of the one-site expectation under the last integral is trivial, of course).

After these simple remarks we come to the important formula that is going to be the starting point of all our analysis.

Let us define the *single-site variation of the Hamiltonian with respect to the disorder variable*[4] η_x at the site x to be

$$\Delta H_x(\sigma_{\overline{x}},\eta_x,\eta_x^0,\eta_{\partial x}) = \sum_{A;A\ni x}[\Phi_A(\sigma_{\overline{x}},\eta_x\eta_{\partial x}) - \Phi_A(\sigma_{\overline{x}},\eta_x^0\eta_{\partial x})] \tag{2.2}$$

where $(\eta_x^0)_{x\in\mathbb{Z}^d}$ is some fixed reference configuration of the disorder variable (that is independent of x). While we will later put $\eta_x^0\in\mathcal{H}^0$, one might also want to choose some other value that is not in the support of the single-site distribution in certain situations.

The trick is to use the 'finite-volume perturbation formula'

$$\begin{aligned} & \int \mu_\Lambda^{\sigma^{\text{b.c.}}_{\partial\Lambda}}[\eta_x,\eta_{\overline{\Lambda}\setminus x}](d\sigma_\Lambda)f(\sigma_\Lambda) \\ &= \frac{\int \mu_\Lambda^{\sigma^{\text{b.c.}}_{\partial\Lambda}}[\eta_x^0,\eta_{\overline{\Lambda}\setminus x}](d\sigma_\Lambda)f(\sigma_\Lambda)\exp\{-\Delta H_x(\sigma_{\overline{x}},\eta_x,\eta_x^0,\eta_{\partial x})\}}{\int \mu_\Lambda^{\sigma^{\text{b.c.}}_{\partial\Lambda}}[\eta_x^0,\eta_{\overline{\Lambda}\setminus x}](d\sigma_\Lambda)\exp\{-\Delta H_x(\sigma_{\overline{x}},\eta_x,\eta_x^0,\eta_{\partial x})\}} \end{aligned} \tag{2.3}$$

[4] A quantity of this type also plays a crucial role in [1] where the fluctuations of extensive quantities are investigated. Its Gibbs expectation could be termed 'order parameter that is conjugate to the disorder'.

which is just a rewriting of Boltzmann factors. Using this we get

Lemma 2.1. *For any reference configuration η_x^0 the conditional expectations of the one-site disorder variable η_x can be rewritten as*

$$
\begin{aligned}
&K_{\Lambda_N}^{\sigma^{\text{b.c.}}_{\partial\Lambda_N}}[\eta_x \mid \sigma_{\Lambda\setminus x};\eta_{\Lambda\setminus x}] \\
&\quad= \nu(\eta_x)\int \mu_x^{\sigma_{\partial x}}[\eta_x^0,\eta_{\partial x}](d\tilde\sigma_x)\exp\{-\Delta H_x(\sigma_{\partial x},\tilde\sigma_x,\eta_x,\eta_x^0,\eta_{\partial x})\} \qquad (2.4)\\
&\qquad\times\int K_{\Lambda_N}^{\sigma^{\text{b.c.}}_{\partial\Lambda_N}}[d\tilde\eta_{\overline{\Lambda_N}\setminus\Lambda}\mid \sigma_{\partial_\Lambda};\eta_x^0,\eta_{\Lambda\setminus x}]\\
&\qquad\times\Big[\int \mu_{\Lambda_N}^{\sigma^{\text{b.c.}}_{\partial\Lambda_N}}[\eta_x^0,\eta_{\Lambda\setminus x},\tilde\eta_{\overline{\Lambda_N}\setminus\Lambda}](d\tilde\sigma_{\overline{x}})\exp\{-\Delta H_x(\tilde\sigma_{\overline{x}},\eta_x,\eta_x^0,\eta_{\partial x})\}\Big]^{-1}\\
&\qquad\times\Big\{\sum_{\eta_x'}\nu(\eta_x')\int \mu_x^{\sigma_{\partial x}}[\eta_x^0,\eta_{\partial x}](d\tilde\sigma_x)\exp\{-\Delta H_x(\sigma_{\partial x},\tilde\sigma_x,\eta_x',\eta_x^0,\eta_{\partial x})\}\\
&\qquad\times\int K_{\Lambda_N}^{\sigma^{\text{b.c.}}_{\partial\Lambda_N}}[d\tilde\eta_{\overline{\Lambda_N}\setminus\Lambda}\mid \sigma_{\partial_\Lambda};\eta_x^0,\eta_{\Lambda\setminus x}]\\
&\qquad\times\Big[\int \mu_{\Lambda_N}^{\sigma^{\text{b.c.}}_{\partial\Lambda_N}}[\eta_x^0,\eta_{\Lambda\setminus x},\tilde\eta_{\overline{\Lambda_N}\setminus\Lambda}](d\tilde\sigma_{\overline{x}})\exp\{-\Delta H_x(\tilde\sigma_{\overline{x}},\eta_x',\eta_x^0,\eta_{\partial x})\}\Big]^{-1}\Big\}^{-1}.
\end{aligned}
$$

Proof. To compute the conditional distribution of η_x, we use the finite-volume perturbation formula to extract the variation of η_x. We use a convention to put tildes on integration variables and write

$$
\begin{aligned}
&K_{\Lambda_N}^{\sigma^{\text{b.c.}}_{\partial\Lambda_N}}[\sigma_{\Lambda\setminus x};\eta_x,\eta_{\Lambda\setminus x}] = \mathsf{P}(\eta_x)\,\mathsf{P}(\eta_{\Lambda\setminus x})\,\mathsf{E}_{\overline{\Lambda_N}\setminus\Lambda}\,\mu_{\Lambda_N}^{\sigma^{\text{b.c.}}_{\partial\Lambda_N}}[\eta_x,\eta_{\Lambda\setminus x},\tilde\eta_{\overline{\Lambda_N}\setminus\Lambda}](\sigma_{\Lambda\setminus x})\\
&\quad= \mathsf{P}(\eta_x)\,\mathsf{P}(\eta_{\Lambda\setminus x})\,\mathsf{E}_{\overline{\Lambda_N}\setminus\Lambda}\Big(\Big[\int \mu_{\Lambda_N}^{\sigma^{\text{b.c.}}_{\partial\Lambda_N}}[\eta_x^0,\eta_{\Lambda\setminus x},\tilde\eta_{\overline{\Lambda_N}\setminus\Lambda}](d\tilde\sigma_\Lambda)\\
&\qquad\qquad\times\exp\{-\Delta H_x(\tilde\sigma_{\overline{x}},\eta_x,\eta_x^0,\eta_{\partial x})\}1_{\tilde\sigma_{\Lambda\setminus x}=\sigma_{\Lambda\setminus x}}\Big]\\
&\qquad\times\Big[\int \mu_{\Lambda_N}^{\sigma^{\text{b.c.}}_{\partial\Lambda_N}}[\eta_x^0,\eta_{\Lambda\setminus x},\tilde\eta_{\overline{\Lambda_N}\setminus\Lambda}](d\tilde\sigma_\Lambda)\exp\{-\Delta H_x(\tilde\sigma_{\overline{x}},\eta_x,\eta_x^0,\eta_{\partial x})\}\Big]^{-1}\Big)\\
&\quad= \mathsf{P}(\eta_x)\,\mathsf{P}(\eta_{\Lambda\setminus x})\mu_{\Lambda^\circ}^{\sigma_{\partial_\Lambda}}[\eta_x^0,\eta_{\Lambda\setminus x}](\sigma_{\Lambda^\circ\setminus x}) \qquad (2.5)\\
&\qquad\times\int \mu_x^{\sigma_{\partial x}}[\eta_x^0,\eta_{\partial x}](d\tilde\sigma_x)\exp\{-\Delta H_x(\sigma_{\partial x},\tilde\sigma_x,\eta_x,\eta_x^0,\eta_{\partial x})\}\\
&\qquad\times\mathsf{E}_{\overline{\Lambda_N}\setminus\Lambda}\,\frac{\mu_{\Lambda_N}^{\sigma^{\text{b.c.}}_{\partial\Lambda_N}}[\eta_x^0,\eta_{\Lambda\setminus x},\tilde\eta_{\overline{\Lambda_N}\setminus\Lambda}](\sigma_{\partial_\Lambda})}{\int \mu_{\Lambda_N}^{\sigma^{\text{b.c.}}_{\partial\Lambda_N}}[\eta_x^0,\eta_{\Lambda\setminus x},\tilde\eta_{\overline{\Lambda_N}\setminus\Lambda}](d\tilde\sigma_\Lambda)\exp\{-\Delta H_x(\tilde\sigma_{\overline{x}},\eta_x,\eta_x^0,\eta_{\partial x})\}}.
\end{aligned}
$$

We have used the compatibility relations for the local specifications in the last equation and we have assumed that Λ,Λ_N are sufficiently large. To get the

conditional expectation, we need to normalize the right-hand side by its η_x-sum. To see that the claim follows, note that

$$\begin{aligned}
&\mathsf{E}_{\overline{\Lambda_N}\setminus\Lambda}\,\frac{\mu_{\Lambda_N}^{\sigma^{\text{b.c.}}_{\partial\Lambda_N}}[\eta^0_x,\eta_{\Lambda\setminus x},\tilde\eta_{\overline{\Lambda_N}\setminus\Lambda}](\sigma_{\partial_\Lambda})}{\int \mu_{\Lambda_N}^{\sigma^{\text{b.c.}}_{\partial\Lambda_N}}[\eta^0_x,\eta_{\Lambda\setminus x},\tilde\eta_{\overline{\Lambda_N}\setminus\Lambda}](d\tilde\sigma_\Lambda)\exp\{-\Delta H_x(\tilde\sigma_{\overline{x}},\eta_x,\eta^0_x,\eta_{\partial x})\}}\\
&\quad=\int K_{\Lambda_N}^{\sigma^{\text{b.c.}}_{\partial\Lambda_N}}[d\tilde\eta_{\overline{\Lambda_N}\setminus\Lambda}\mid\sigma_{\partial_\Lambda};\eta^0_x,\eta_{\Lambda\setminus x}]\\
&\qquad\times\Big[\int\mu_{\Lambda_N}^{\sigma^{\text{b.c.}}_{\partial\Lambda_N}}[\eta^0_x,\eta_{\Lambda\setminus x},\tilde\eta_{\overline{\Lambda_N}\setminus\Lambda}](d\tilde\sigma_{\overline{x}})\exp\{-\Delta H_x(\tilde\sigma_{\overline{x}},\eta_x,\eta^0_x,\eta_{\partial x})\}\Big]^{-1}\\
&\qquad\times\mathsf{E}_{\overline{\Lambda_N}\setminus\Lambda}\,\mu_{\Lambda_N}^{\sigma^{\text{b.c.}}_{\partial\Lambda_N}}[\eta^0_x,\eta_{\Lambda\setminus x},\tilde\eta_{\overline{\Lambda_N}\setminus\Lambda}](\sigma_{\partial_\Lambda})
\end{aligned}\tag{2.6}$$

where the term in the last line does not depend on η_x. □

Remark 2.1. The formula gives the modification of the conditional expectation compared with the 'free' a priori measure $\nu(\eta_x)$ that results from the non-trivial coupling of η to the spin variable σ. The second term in the second line of (2.4), a Gibbs expectation of the exponential of the single-site variation of the Hamiltonian, is of course a local function in the conditioning. Assuming the finiteness of the interaction potential it is bounded. Thus, to investigate the possible non-locality of the left-hand side one has to investigate the expression in the third and fourth line of (2.4).

Remark 2.2. The *local* Λ_N-limit of the conditional expectation $K_{\Lambda_N}^{\sigma^{\text{b.c.}}_{\partial\Lambda_N}}[d\tilde\eta_{\overline{\Lambda_N}\setminus\Lambda}\mid\sigma_{\partial_\Lambda};\eta^0_x,\eta_{\Lambda\setminus x}]$ exists. This follows from the assumption of the existence of the joint local $\lim_{\Lambda_N\uparrow\mathbb{Z}^d}K_{\Lambda_N}^{\sigma^{\text{b.c.}}_{\partial\Lambda_N}}$. Also, the Λ_N-limit of the third and fourth line of (2.4) (that involves the average of an N-dependent function of $\tilde\eta$) exists: the Λ_N-limit of the quantity in the last line of (2.5) exists by our assumption on the existence of a Λ_N-limit on the left-hand side of (2.5). The Λ_N-limit of the last line of (2.6) (the normalization needed to obtain probabilities) also exists by the hypothesis.

Sometimes it is convenient to rewrite (2.4) using that, by the finite-volume perturbation formula, we have

$$\begin{aligned}
&\Big[\int\mu_{\Lambda_N}^{\sigma^{\text{b.c.}}_{\partial\Lambda_N}}[\eta^0_x,\eta_{\Lambda\setminus x},\tilde\eta_{\overline{\Lambda_N}\setminus\Lambda}](d\tilde\sigma_{\overline{x}})\exp\{-\Delta H_x(\tilde\sigma_{\overline{x}},\eta_x,\eta^0_x,\eta_{\partial x})\}\Big]^{-1}\\
&\quad=\int\mu_{\Lambda_N}^{\sigma^{\text{b.c.}}_{\partial\Lambda_N}}[\eta_x,\eta_{\Lambda\setminus x},\tilde\eta_{\overline{\Lambda_N}\setminus\Lambda}](d\tilde\sigma_{\overline{x}})\exp\{+\Delta H_x(\tilde\sigma_{\overline{x}},\eta_x,\eta^0_x,\eta_{\partial x})\}\\
&\quad\equiv\mu_{\Lambda_N}^{\sigma^{\text{b.c.}}_{\partial\Lambda_N}}[\eta_\Lambda,\tilde\eta_{\overline{\Lambda_N}\setminus\Lambda}]\big(\exp\{\Delta H_x(\eta_x,\eta^0_x,\eta_{\partial x})\}\big).
\end{aligned}\tag{2.7}$$

The reader may also want to note that, with usual notations, (2.7) is just a ratio of two partition functions, $Z_{\Lambda_N}^{\sigma^{\text{b.c.}}_{\partial\Lambda_N}}[\eta_x^0\eta_{\Lambda\setminus x}\tilde{\eta}_{\overline{\Lambda_N}\setminus\Lambda}]/Z_{\Lambda_N}^{\sigma^{\text{b.c.}}_{\partial\Lambda_N}}[\eta_x\eta_{\Lambda\setminus x}\tilde{\eta}_{\overline{\Lambda_N}\setminus\Lambda}]$ which makes the symmetry between η_x and η_x^0 more apparent.

From Lemma 2.1. we have

Proposition 2.1.

$$\frac{K[\eta_x^1 \mid \sigma_{\Lambda\setminus x};\eta_{\Lambda\setminus x}]}{K[\eta_x^2 \mid \sigma_{\Lambda\setminus x};\eta_{\Lambda\setminus x}]} = q^{\text{local}}(\eta_x^1,\eta_x^2,\sigma_{\partial x},\eta_{\partial x})\, q^{\text{nonloc}}_{\Lambda,x}[\eta_x^1,\eta_x^2,\eta_{\Lambda\setminus x},\sigma_{\partial_\Lambda}] \tag{2.8}$$

where

$$\begin{aligned} &q^{\text{local}}(\eta_x^1,\eta_x^2,\sigma_{\partial x},\eta_{\partial x}) \\ &\quad = \frac{\nu(\eta_x^1)}{\nu(\eta_x^2)}\int \mu_x^{\sigma_{\partial x}}[\eta_x^2,\eta_{\partial x}](d\tilde{\sigma}_x)\exp\{-\Delta H_x(\sigma_{\partial x},\tilde{\sigma}_x,\eta_x^1,\eta_x^2,\eta_{\partial x})\} \end{aligned} \tag{2.9}$$

is a local function of σ,η and

$$\begin{aligned} &q^{\text{nonloc}}_{\Lambda,x}[\eta_x^1,\eta_x^2,\eta_{\Lambda\setminus x},\sigma_{\partial_\Lambda}] \\ &\quad = \lim_{\Lambda_N\uparrow\mathbf{Z}^d}\int K_{\Lambda_N}^{\sigma^{\text{b.c.}}_{\partial\Lambda_N}}[d\tilde{\eta}_{\overline{\Lambda_N}\setminus\Lambda} \mid \sigma_{\partial_\Lambda};\eta_x^2,\eta_{\Lambda\setminus x}] \\ &\qquad \times\int \mu_{\Lambda_N}^{\sigma^{\text{b.c.}}_{\partial\Lambda_N}}[\eta_x^1,\eta_{\Lambda\setminus x},\tilde{\eta}_{\overline{\Lambda_N}\setminus\Lambda}](d\tilde{\sigma}_{\overline{x}})\exp\{\Delta H_x(\tilde{\sigma}_{\overline{x}},\eta_x^1,\eta_x^2,\eta_{\partial x})\} \end{aligned} \tag{2.10}$$

is a potentially non-local function of σ,η. The last limit exists.

Corollary 2.1. *Whether a point $\xi=(\sigma,\eta)$ is a good configuration for K does not depend on σ; $\xi=(\sigma,\eta)$ is good if and only if*

$$\begin{aligned} \sup_{\eta^+,\eta^-;\sigma^+,\sigma^-\ \Lambda:\Lambda\supset V} &\big|q^{\text{nonloc}}_{\Lambda,x}[\eta_x^1,\eta_x^2,\eta_{V\setminus x},\eta^+_{\Lambda\setminus V},\sigma^+_{\partial_\Lambda}] \\ &- q^{\text{nonloc}}_{\Lambda,x}[\eta_x^1,\eta_x^2,\eta_{V\setminus x},\eta^-_{\Lambda\setminus V},\sigma^-_{\partial_\Lambda}]\big| \to 0 \end{aligned} \tag{2.11}$$

with $V\uparrow\mathbf{Z}^d$, for any site $x\in\mathbf{Z}^d$, for any pair $\eta_x^1,\eta_x^2\in\mathcal{H}_0$.

Proof. To prove the proposition, choose the reference configuration $\eta_x^0=\eta_x^2$ and use Lemma 2.1, along with (2.7). The corollary follows from the fact that q^{local} is a local function, and that it suffices to check the conditional expectations of the disorder variable by (2.1). Note to this end that both q's in Proposition 2.1 are uniformly bounded below by a quantity that is strictly bigger than zero and uniformly bounded from above. This is clear, since both quantities are defined in terms of probability averages of uniformly bounded quantities, by the assumed finiteness of ΔH_x. □

To understand the symmetry between η^1 and η^2 in this formula, we remark that q^{local} as well as the inner integral in (2.10) can be written as ratios of partitions functions, by the remark following (2.7). We will now discuss various consequences of corollary of Proposition 2.1. It is very difficult to say anything reasonable about the behavior of the conditional measure $K_{\Lambda_N}^{\sigma^{\text{b.c.}}_{\partial\Lambda_N}}[d\tilde{\eta}_{\overline{\Lambda_N}\setminus\Lambda} \mid \sigma_{\partial_\Lambda}; \eta_x^2, \eta_{\Lambda\setminus x}]$, as a function of the spin-conditioning $\sigma_{\partial_\Lambda}$. So, in our examples we will at first draw conclusions from estimates that are *uniform* with respect to the integration variable $\tilde{\eta}_{\Lambda_N\setminus\Lambda}$.

We start with a criterion for points $\xi = (\eta, \sigma)$ being good configurations, that is a pretty much straightforward consequence of Proposition 2.1. This will be employed if we want to show Gibbsianness. Below we give a slightly more complicated criterion for points $\xi = (\eta, \sigma)$ being bad configurations, needed to investigate non-Gibbsianness.

Proposition 2.2. *Suppose that η is such that for any $x \in \mathbf{Z}^d$ we have that*

$$r_{V,x}(\eta_x^1, \eta_x^2, \eta) := \sup_{\substack{\eta^+,\eta^- \\ \Lambda:\Lambda\supset V}} \Big| \int \mu_\Lambda^{\sigma^{\text{b.c.}}_{\partial\Lambda}}[\eta_x^1, \eta_{V\setminus x}, \eta^+_{\overline{\Lambda}\setminus V}]\big(\exp\{\Delta H_x(\eta_x^1, \eta_x^2, \eta_{\partial x})\}\big) - \int \mu_\Lambda^{\sigma^{\text{b.c.}}_{\partial\Lambda}}[\eta_x^1, \eta_{V\setminus x}, \eta^-_{\overline{\Lambda}\setminus V}]\big(\exp\{\Delta H_x(\eta_x^1, \eta_x^2, \eta_{\partial x})\}\big)\Big| \to 0 \tag{2.12}$$

with $V \uparrow \mathbf{Z}^d$, for any x, for any pair $\eta_x^1, \eta_x^2 \in \mathcal{H}_0$. Then the configuration η, σ is a good configuration, for any σ.

Proof. To see that the hypothesis implies (2.11), we use that

$$\Big|\mu_{\Lambda_N}^{\sigma^{\text{b.c.}}_{\partial\Lambda_N}}[\eta_x^1, \eta_{V\setminus x}, \eta^{+,-}_{\Lambda\setminus V}, \tilde{\eta}_{\overline{\Lambda_N}\setminus\Lambda}]\big(\exp\{\Delta H_x(\eta_x^1, \eta_x^2, \eta_{\partial x})\}\big) - \mu_{\Lambda_N}^{\sigma^{\text{b.c.}}_{\partial\Lambda_N}}[\eta_x^1, \eta_{\overline{\Lambda_N}\setminus x}]\big(\exp\{\Delta H_x(\tilde{\sigma}_{\overline{x}}, \eta_x^1, \eta_x^2, \eta_{\partial x})\}\big)\Big| \le r_{V,x}(\eta_x^1, \eta_x^2, \eta) \tag{2.13}$$

to compare the μ-terms under the $\tilde{\eta}$-integrals with a term that is independent of $\tilde{\eta}$ and $\eta^{+,-}$. This shows that (2.11) is bounded by $2r_{V,x}$ which converges to zero. □

Remark 2.3. To estimate $r_{V,x}(\eta_x^1, \eta_x^2, \eta)$, we can also bound the variation of the random couplings by the variation over the boundary conditions

$$r_{V,x}(\eta_x^1, \eta_x^2, \eta) \le \sup_{\sigma^1,\sigma^2} \Big|\mu_{V^o}^{\sigma^1_{\partial^-V}}[\eta_x^1\eta_{V\setminus x}]\big(\exp\{\Delta H_x(\eta_x^1, \eta_x^2, \eta_{\partial x})\}\big) - \mu_{V^o}^{\sigma^2_{\partial^-V}}[\eta_x^1\eta_{V\setminus x}]\big(\exp\{\Delta H_x(\eta_x^1, \eta_x^2, \eta_{\partial x})\}\big)\Big|. \tag{2.14}$$

Remark 2.4. We see how (2.12) parallels (1.6). The quantity that is of interest is now the Gibbs expectation of the exponential of the single-site variation as a function of the disorder variables. In words: if we have equicontinuity in the parameter Λ of these finite Λ-Gibbs expectations with respect to the disorder variable at the point η, we conclude that η, σ is a good configuration. The reader may also find it intuitive to rewrite the Gibbs expectations appearing in (2.12) in the form of ratios of partition functions, or (equivalently) as exponentials of differences of free energies taken for η_x^1 and η_x^2. In slightly different words, the criterion thus requires equicontinuity in the volume of the single-site variations of the free energies with respect to the disorder variable at the point η.

To get a criterion for bad configurations that is independent of the behavior of the outer expectation of q^{nonloc} (see (2.10)) leads to an expression that is slightly more complicated because it contains an additional supremum.

Proposition 2.3. *Put*

$$q^{\text{upper}}_{\Lambda,x}[\eta_x^1, \eta_x^2, \eta_{\Lambda\setminus x}]$$
$$:= \limsup_{\Lambda_N \uparrow \mathbb{Z}^d} \sup_{\tilde{\eta}_{\overline{\Lambda_N}\setminus\Lambda}} \mu^{\sigma^{\text{b.c.}}_{\partial\Lambda_N}}_{\Lambda_N}[\eta_x^1, \eta_{\Lambda\setminus x}, \tilde{\eta}_{\overline{\Lambda_N}\setminus\Lambda}]\big(\exp\{\Delta H_x(\eta_x^1, \eta_x^2, \eta_{\partial x})\}\big) \quad (2.15)$$

Then η, σ is a bad configuration for K, if for some site x, for some pair η_x^1, η_x^2

$$\lim_{V\uparrow\mathbb{Z}^d} \sup_{\substack{\eta^+,\eta^- \\ \Lambda:\Lambda\supset V}} \big((q^{\text{upper}}_{\Lambda,x}[\eta_x^2, \eta_x^1, \eta_{V\setminus x}, \eta^+_{\Lambda\setminus V}])^{-1}$$
$$- q^{\text{upper}}_{\Lambda,x}[\eta_x^1, \eta_x^2, \eta_{V\setminus x}, \eta^-_{\Lambda\setminus V}]\big) \;>\; 0. \quad (2.16)$$

Proof. By (2.7) and the uniform estimate of the $\tilde{\eta}$-integral we see that

$$\begin{aligned} q^{\text{upper}}_{\Lambda,x}[\eta_x^2, \eta_x^1, \eta_{\Lambda\setminus x}]^{-1} &\leq q^{\text{nonloc}}_{\Lambda,x}[\eta_x^1, \eta_x^2, \eta_{\Lambda\setminus x}, \sigma_{\partial_\Lambda}] \\ &\leq q^{\text{upper}}_{\Lambda,x}[\eta_x^1, \eta_x^2, \eta_{\Lambda\setminus x}]. \end{aligned} \quad (2.17)$$

Hence the claim (discontinuity of the left-hand side) follows from the definition of a bad configuration. □

Models with almost surely convergent Gibbs states

Suppose that we have the existence of a weak limit

$$\lim_{\Lambda\uparrow\mathbb{Z}^d} \mu^{\sigma^{\text{b.c.}}_{\partial\Lambda}}_{\Lambda}[\eta_\Lambda] = \mu_\infty[\eta_{\mathbb{Z}^d}] \quad (2.18)$$

for P-a.e. η. It follows that $\mu_\infty[\eta_{\mathbb{Z}^d}]$ is an infinite-volume Gibbs measure for P-a.e. η that depends measurably on η. Consequently the infinite-volume joint state is then just the P-integral of μ_∞. We stress that this has not been assumed

so far and is really a much stronger assumption then local convergence of the joint states. It is not expected to hold e.g. for spin glasses in the multiphase region (that is supposed although not proved to exist).

This assumption implies that the terms in the main formula of Lemma 2.1 converge individually with $\Lambda_N \uparrow \mathbf{Z}^d$. So we have that

$$q^{\mathrm{nonloc}}_{\Lambda,x}[\eta^1_x, \eta^2_x, \eta_{\Lambda\setminus x}, \sigma_{\partial_-\Lambda}] \tag{2.19}$$
$$= \int K[d\tilde{\eta}_{\mathbf{Z}^d\setminus\Lambda} \mid \sigma_{\partial_-\Lambda}; \eta^2_x, \eta_{\Lambda\setminus x}]\, \mu_\infty[\eta^1_x, \eta_{\Lambda\setminus x}, \tilde{\eta}_{\mathbf{Z}^d\setminus\Lambda}]\big(\exp\{\Delta H_x(\eta^1_x, \eta^2_x, \eta_{\partial x})\}\big).$$

Suppose we want to exhibit a bad configuration and we have estimates on the continuity of $\eta \mapsto \mu_\infty[\eta]$ for *typical* directions but *not* in all directions. For an example of a perturbation in an atypical direction think of the random field Ising model that will be discussed below. Here the Gibbs measure with plus boundary conditions can be pushed in the 'wrong phase' by choosing the random fields to be minus in a large annulus. While the RFIM can be treated by Proposition 2.3, there are examples where we would like to get away from *uniform* estimates with respect to $\tilde{\eta}$ in favor of estimates that are only true for *typical* $\tilde{\eta}$, for the a priori measure P.

To obtain the following criterion is more subtle than what we noted in Propositions 2.2 and 2.3. The trick is to show the existence of suitable 'bad' σ-conditionings using the knowledge about typical disorder variables with respect to the unbiased P-measure.

Proposition 2.4. *Assume the almost sure existence of the weak limits of finite-volume Gibbs measures* (2.18) *and denote by K the corresponding infinite-volume joint measure.*

The configuration $\xi = (\eta, \sigma)$ is a bad configuration for K, if for each cube V, centered at the origin, there exist an increasing choice of volumes $\Lambda(V)$ and configurations $\eta^V, \bar{\eta}^V$ such that for P-a.e. $\tilde{\eta}$ we have that

$$\liminf_{V\uparrow\mathbf{Z}^d} \mu_\infty[\eta^1_x, \eta_{V\setminus x}\bar{\eta}^V_{\Lambda(V)\setminus V}, \tilde{\eta}_{\mathbf{Z}^d\setminus\Lambda}]\big(\exp\{\Delta H_x(\eta^1_x, \eta^2_x, \eta_{\partial x})\}\big)$$
$$> \limsup_{V\uparrow\mathbf{Z}^d} \mu_\infty[\eta^1_x, \eta_{V\setminus x}\eta^V_{\Lambda(V)\setminus V}, \tilde{\eta}_{\mathbf{Z}^d\setminus\Lambda}]\big(\exp\{\Delta H_x(\eta^1_x, \eta^2_x, \eta_{\partial x})\}\big) \tag{2.20}$$

for some site x and some η^1_x, η^2_x.

Proof. We will show that there exist two conditionings $\bar{\sigma}$ and σ, such that

$$\liminf_{V\uparrow\mathbf{Z}^d} q^{\mathrm{nonloc}}_{\Lambda(V),x}[\eta^1_x, \eta^2_x, \eta_{V\setminus x}\bar{\eta}^V_{\Lambda(V)\setminus V}, \bar{\sigma}_{\partial_-\Lambda(V)}]$$
$$> \limsup_{V\uparrow\mathbf{Z}^d} q^{\mathrm{nonloc}}_{\Lambda(V),x}[\eta^1_x, \eta^2_x, \eta_{V\setminus x}\eta^V_{\Lambda(V)\setminus V}, \sigma_{\partial_-\Lambda(V)}]. \tag{2.21}$$

From this and the corollary of Proposition 2.1 the badness follows. To show (2.21), we proceed as follows: the left-hand side and right-hand side of (2.20) are tail

measurable, hence almost surely constant. Denote the left-hand side of (2.20) by $\bar{q}^{\infty}[\eta_x^1, \eta_x^2, \eta_{\mathbb{Z}^d \setminus x}]$ and the right-hand side by $q^{\infty}[\eta_x^1, \eta_x^2, \eta_{\mathbb{Z}^d \setminus x}]$. We will show that there exists a conditioning σ such that the right-hand side of (2.21) is bounded from above by $q^{\infty}[\eta_x^1, \eta_x^2, \eta_{\mathbb{Z}^d \setminus x}]$. (Similarly, there exists a conditioning $\bar{\sigma}$ such that the left-hand side of (2.21) is bounded from below by $\bar{q}^{\infty}[\eta_x^1, \eta_x^2, \eta_{\mathbb{Z}^d \setminus x}]$.)

We will construct this conditioning as a sequence given on the 'small' annuli $\partial_- \Lambda(V)$ (and arbitrary for other lattice sites.) To make use of the almost sure statement with respect to the product measure P, we need to produce a formula that recovers this measure. We write

$$\begin{aligned}
&\limsup_{V \uparrow \infty} \sum_{\tilde{\sigma}_{\partial_- \Lambda(V)}} \int K[\tilde{\sigma}_{\partial_- \Lambda(V)} \mid \eta_x^2, \eta_{V \setminus x} \eta^V_{\Lambda(V) \setminus V}] \\
&\qquad \times q^{\text{nonloc}}_{\Lambda(V),x}[\eta_x^1, \eta_x^2, \eta_{V \setminus x} \eta^V_{\Lambda(V) \setminus V}, \tilde{\sigma}_{\partial_- \Lambda(V)}] \\
&= \limsup_{V \uparrow \infty} \int \mathsf{P}(d\tilde{\eta}) \mu_{\infty}[\eta_x^1, \eta_{V \setminus x} \eta^V_{\Lambda(V) \setminus V}, \tilde{\eta}_{\mathbb{Z}^d \setminus \Lambda}] \big(\exp\{\Delta H_x(\eta_x^1, \eta_x^2, \eta_{\partial x})\} \big) \\
&\leq q^{\infty}[\eta_x^1, \eta_x^2, \eta_{\mathbb{Z}^d \setminus x}] \qquad (2.22)
\end{aligned}$$

where the first equality follows from (2.19) and the inequality from Fatou's lemma with respect to product integration of the $\tilde{\eta}$. From (2.22) , the existence of such a conditioning σ is easy to see. (By contradiction, if the claim were not true, for any sequence of conditionings $\sigma_{\partial_- \Lambda(V)}$, we would have that there exists a positive ε such that $\min_{\tilde{\sigma}_{\partial_- \Lambda(V)}} q^{\text{nonloc}}_{\Lambda(V),x}[\ldots, \tilde{\sigma}_{\partial_- \Lambda(V)}] \geq q^{\infty}[\ldots] + \varepsilon$ for infinitely many V's. But this would imply that also the quantity under the limsup on the left-hand side of (2.22) (which is just a $\tilde{\sigma}_{\partial_- \Lambda(V)}$-expectation) would have to be bigger than or equal to this bound, for the same infinitely many V's.) □

3. Examples

3.1. The random field Ising model

Note that the single site perturbation with respect to the random field of the Hamiltonian is very simple, i.e.

$$\begin{aligned}
\exp\{\Delta H_x(\sigma_x, \eta_x^1, \eta_x^2)\} &= \exp\{h(\eta_x^2 - \eta_x^1)\sigma_x\} \\
&= \exp\{h(\eta_x^1 - \eta_x^2)\} + 2 \sinh h(\eta_x^2 - \eta_x^1)\, 1_{\sigma_x = 1}. \qquad (3.1)
\end{aligned}$$

An application of Propositions 2.2 and 2.3 with the aid of monotonicity arguments gives Theorem 1.1, as stated in the introduction. It provides a complete characterization of good/bad configurations in terms of the behavior of the finite-volume Gibbs expectations with plus (resp. minus) boundary conditions. The interesting part, the mechanism of non-continuity, is due to the fact that we can make any random field Gibbs measure look like the plus (minus) phase

around a given site by choosing the fields in a sufficiently large annulus to be plus (minus). That this works independently of what the fields even further outside do, is crucial for the argument.

Proof of Theorem 1.1. We use the fact that the function $(\eta, \sigma^{\mathrm{bc}}) \mapsto \mu_\Lambda^{\sigma^{\mathrm{bc}}}[\eta_\Lambda](\tilde{\sigma}_x = 1)$ is monotone (with respect to the partial order of its arguments obtained by site-wise comparison.) From this follows that the limits in (1.7) exist, due to monotonicity, for any η. Denote the left-hand side of (1.7) by $m_x^+(\eta_{\mathbf{Z}^d})$ and the right-hand side of (1.7) by $m_x^-(\eta_{\mathbf{Z}^d})$. We also note that, by the finite-volume perturbation formula, one obtains that

$$e^{h(\eta_x^1-\eta_x^2)}([m_x^{+,-}(\eta_x^1, \eta_{\mathbf{Z}^d\setminus x})]^{-1} - 1) = e^{h(\eta_x^2-\eta_x^1)}([m_x^{+,-}(\eta_x^2, \eta_{\mathbf{Z}^d\setminus x})]^{-1} - 1). \quad (3.2)$$

This shows in particular that (say) $m_x^+(\eta_x^1, \eta_{\mathbf{Z}^d\setminus x})$ and $m_x^+(\eta_x^2, \eta_{\mathbf{Z}^d\setminus x})$ are strictly monotone functions of each other (when varying $\eta_{\mathbf{Z}^d\setminus x}$). In particular we see explicitly that, whether the left-hand side and the right-hand side of (1.7) coincide does of course not depend on the value of η_x.

Now, to show that a configuration is good if the two limits coincide, we apply Proposition 2.2 and the remark after it. Using (3.1), we see that

$$r_{V,x}(\eta_x^1, \eta_x^2, \eta) \to 0$$

with $V \uparrow \mathbf{Z}^d$, if

$$\sup_{\sigma^1,\sigma^2} \left| \int \mu_V^{\sigma^1_{\partial V}}[\eta_x^1 \eta_{V\setminus x}](\tilde{\sigma}_x = 1) - \int \mu_V^{\sigma^2_{\partial V}}[\eta_x^1 \eta_{V\setminus x}](\tilde{\sigma}_x = 1) \right| \to 0 \quad (3.3)$$

with $V \uparrow \mathbf{Z}^d$. Using monotonicity in the boundary condition we see that this is equivalent to the equality of the two limits in (1.7).

Now, to show that a configuration is bad if the two limits in (1.7) do not coincide, we use Proposition 2.3. We have that

$$\begin{aligned} q_{\Lambda,x}^{\mathrm{upper}}&[\eta_x^1, \eta_x^2, \eta_{\Lambda\setminus x}] \\ &= \limsup_{\Lambda_N\uparrow\mathbf{Z}^d} \sup_{\tilde{\eta}_{\Lambda_N\setminus\Lambda}} \mu_{\Lambda_N}^{\sigma^{\mathrm{b.c.}}_{\partial\Lambda_N}}[\eta_x^1, \eta_{\Lambda\setminus x}, \tilde{\eta}_{\overline{\Lambda}_N\setminus\Lambda}]\big(\exp\{\Delta H_x(\sigma_x, \eta_x^1, \eta_x^2)\}\big) \\ &= \exp\{h(\eta_x^1 - \eta_x^2)\} + \limsup_{\Lambda_N\uparrow\mathbf{Z}^d} \sup_{\tilde{\eta}_{\Lambda_N\setminus\Lambda}} 2\sinh(h(\eta_x^2 - \eta_x^1)) \\ &\quad \times \mu_{\Lambda_N}^{\sigma^{\mathrm{b.c.}}_{\partial\Lambda_N}}[\eta_x^1, \eta_{\Lambda\setminus x}, \tilde{\eta}_{\overline{\Lambda}_N\setminus\Lambda}](\tilde{\sigma}_x = 1). \end{aligned} \quad (3.4)$$

Suppose now that $\eta_x^2 \geq \eta_x^1$. Then we get from the monotonicity

$$q_{\Lambda,x}^{\mathrm{upper}}[\eta_x^1, \eta_x^2, \eta_{\Lambda\setminus x}] \leq \exp\{h(\eta_x^1 - \eta_x^2)\} + 2\sinh(h(\eta_x^2 - \eta_x^1))\mu_\Lambda^{+_{\partial\Lambda}}[\eta_x^1, \eta_{\Lambda\setminus x}](\tilde{\sigma}_x = 1). \quad (3.5)$$

Similarly we have that

$$q^{\mathrm{upper}}_{\Lambda,x}[\eta^2_x,\eta^1_x,\eta_{\Lambda\setminus x}] \le \exp\{h(\eta^2_x-\eta^1_x)\}+2\sinh(h(\eta^1_x-\eta^2_x))\mu^{-\,\partial\Lambda}_\Lambda[\eta^1_x,\eta_{\Lambda\setminus x}](\tilde\sigma_x=1). \tag{3.6}$$

Now we use the important fact that

$$\lim_{\Lambda\uparrow\mathbb{Z}^d}\mu^{-\,\partial\Lambda}_\Lambda[\eta_V,\eta_{\Lambda\setminus V}=+](\tilde\sigma_x=1)=\lim_{\Lambda\uparrow\mathbb{Z}^d}\mu^{+\,\partial\Lambda}_\Lambda[\eta_V,\eta_{\Lambda\setminus V}=+](\tilde\sigma_x=1) \tag{3.7}$$

which follows from the unicity of the Gibbs measure of a homogeneous ferromagnet in a positive magnetic field, and, consequently,

$$\begin{aligned}&\lim_{V\uparrow\mathbb{Z}^d}\lim_{\Lambda\uparrow\mathbb{Z}^d}\mu^{-\,\partial\Lambda}_\Lambda[\eta_V,\eta_{\Lambda\setminus V}=+](\tilde\sigma_x=1)\\ &\quad=\lim_{V\uparrow\mathbb{Z}^d}\lim_{\Lambda\uparrow\mathbb{Z}^d}\mu^{+\,\partial\Lambda}_\Lambda[\eta_V,\eta_{\Lambda\setminus V}=+](\tilde\sigma_x=1)\;=\;m^+_x(\eta_{\mathbb{Z}^d})\end{aligned} \tag{3.8}$$

where the right equality follows from the inequality

$$\mu^{+\,\partial\Lambda}_\Lambda[\eta_V,\eta_{\Lambda\setminus V}](\tilde\sigma_x=1)\le\mu^{+\,\partial\Lambda}_\Lambda[\eta_V,\eta_{\Lambda\setminus V}=+](\tilde\sigma_x=1)\le\mu^{+\,\partial V}_V[\eta_V](\tilde\sigma_x=1).$$

From this we have that

$$\begin{aligned}&\lim_{V\uparrow\mathbb{Z}^d}\lim_{\Lambda\uparrow\mathbb{Z}^d}q^{\mathrm{upper}}_{\Lambda,x}[\eta^1_x,\eta^2_x,\eta_V,\eta_{\Lambda\setminus V}=-]\\ &\quad\le\exp\{h(\eta^1_x-\eta^2_x)\}+2\sinh(h(\eta^2_x-\eta^1_x))m^-_x(\eta^1_x,\eta_{\mathbb{Z}^d\setminus x})\end{aligned} \tag{3.9}$$

and, similarly

$$\begin{aligned}&\lim_{V\uparrow\mathbb{Z}^d}\lim_{\Lambda\uparrow\mathbb{Z}^d}q^{\mathrm{upper}}_{\Lambda,x}[\eta^2_x,\eta^1_x,\eta_V,\eta_{\Lambda\setminus V}=+]\\ &\quad\le\exp\{h(\eta^2_x-\eta^1_x)\}+2\sinh(h(\eta^1_x-\eta^2_x))m^+_x(\eta^2_x,\eta_{\mathbb{Z}^d\setminus x})\\ &\quad=\big(\exp\{h(\eta^1_x-\eta^2_x)\}+2\sinh(h(\eta^2_x-\eta^1_x))m^+_x(\eta^1_x,\eta_{\mathbb{Z}^d\setminus x})\big)^{-1}\end{aligned} \tag{3.10}$$

where the last line follows from relation (3.2). From (3.10) it is evident that (1.7) implies (2.16). □

3.2. Models with decoupling configurations

Suppose we have a model that allows for 'non-percolating' decoupling configurations η. By this we mean that, for given η, for any site x there exists a volume $\Lambda_x(\eta)$ such that for any $\Lambda\supset\Lambda_x(\eta)$ we have that

$$\mu^{\sigma^{\mathrm{b.c.}}_{\partial\Lambda}}_\Lambda[\eta^1_x,\hat\eta_{\overline{\Lambda}\setminus x}]\big(e^{\Delta H_x(\eta^1_x,\eta^2_x,\eta_{\partial x})}\big)=\mu^{\mathrm{open}}_{\Lambda_x(\eta)}[\eta^1_x\eta_{\Lambda_x(\eta)\setminus x}]\big(e^{\Delta H_x(\eta^1_x,\eta^2_x,\eta_{\partial x})}\big) \tag{3.11}$$

independently of Λ (for any pair η^1_x,η^2_x), for any configuration $\hat\eta$ that coincides with η on $\Lambda_x(\eta)$.

Think e.g. of an Ising model with random couplings taking the value 0 with positive probability. Then an example of a non-percolating decoupling configuration is a configuration of coupling constants such that all the resulting spin clusters (with edges of non-zero coupling constants) are finite.

For a decoupling configuration η the formula for the conditional expectations simplifies considerably. A look at (2.10) gives us

$$q^{\text{nonloc}}_{\Lambda,x}[\eta^1_x,\eta^2_x,\eta_{\Lambda\setminus x},\sigma_{\partial_\Lambda}] = \mu^{\text{open}}_{\Lambda_x(\eta)}[\eta^1_x,\eta_{\overline{\Lambda_x(\eta)}\setminus x}]\big(\exp\{\Delta H_x(\eta^1_x,\eta^2_x,\eta_{\partial x})\}\big) \tag{3.12}$$

for Λ sufficiently large (depending on η). Since any perturbation of η far away from x leaves this quantity unchanged, we immediately obtain

Proposition 3.1. *A configuration $\xi = (\eta,\sigma)$ is a good configuration, if η is a decoupling configuration. Consequently, if $\mathsf{P}\{\eta \in \mathcal{H} : \eta$ is a decoupling configuration$\} = 1$, then any joint measure that is a limit of the form (1.3) is almost surely Gibbs.*

This should not be confused with the fact that a non-decoupling η can be shown to be bad with the use of (a sequence of) decoupling configurations η^+, η^-, as the following examples show.

The GriSing random field revisited (see [12]).

The spins are $\sigma_x \in \{-1,1\}$, the local disorder variable η_x takes values in $\{0,1\}$ with $\nu[\eta_x = 1] = p \in (0,1)$ and the Hamiltonian is given by $H^\eta(\sigma) = -J\sum_{<x,y>}\eta_x\sigma_x\eta_y\sigma_y$. This model was shown to be non-Gibbs for p below the percolation threshold for site percolation. Let us see, how this comes out of our framework and explain at the same time that *any*[5] weak limit $\lim_{\Lambda_N}\mathsf{P}(d\eta)\mu^{\sigma^{b.c.}}_{\Lambda_N}[\eta](d\sigma)$ will also be non-Gibbs, for any $p \in (0,1)$ (for sufficiently large J).

There is the trivial mapping that sends the pair (η_x,σ_x) to the product $\eta_x\sigma_x$; looking at new variables that are products (as it was done in [12]) is equivalent to looking at pairs, since $\eta_x = 0$ iff $\eta_x\sigma_x = 0$. The single-site variation of the Hamiltonian with respect to the occupation variable η_x now becomes $\Delta H_x(\sigma_{\overline{x}},\eta_x,\eta^0_x,\eta_{\partial x}) = -J(\eta_x - \eta^0_x)\sum_{y:d(y,x)=1}\eta_y\sigma_x\sigma_y$. In other words, the spin observable conjugate to the disorder is an occupation-number-dependent sum of spin correlations. To show non-Gibbssianness one needs to find configurations of occupation numbers for which the corresponding Gibbs expectations show a discontinuity when the occupation numbers are perturbed arbitrarily far away from x.

Recalling [12] we look at the configuration η^{disc}, that is 0 on the 'base-plane' $B = \{x \in \mathbf{Z}^d, x_d = 0\}$ and 1 otherwise. Then $(\eta^{\text{disc}},\sigma)$ is a bad configuration for any σ, for any joint infinite-volume measure that is a limit of the form (1.3). To see this, one only needs to look at conditional probabilities for special decoupling

[5] Think e.g. of the Dobrushin states that are supposed to exist for p close to 1 in $d \geq 4$!

configurations. Indeed, for a finite box $V \subset \mathbf{Z}^d$, centered at the origin, denote by $\eta^{\mathrm{disc},V}$ the configuration that coincides with η^{disc} inside V and vanishes outside V. Denote by V^+ (V^-) the occupied sites in V in the upper (lower) half-space. For $z \in V \cap B$ denote by $\eta^{\mathrm{disc},V,z}$ the configuration that has z as an additional occupied site. Denote the nearest neighbor of the origin in V^+ by x_0 and the nearest neighbor of the origin in V^- by y_0. Put $\eta_0^2 = 1$, $\eta_0^1 = 0$. Then $\exp\{\Delta H_0(\sigma_{\overline{0}}, \eta_0^1, \eta_0^2, \eta_{\partial 0}^{\mathrm{disc}})\} = \exp\{J\sigma_0(\sigma_{x_0} + \sigma_{y_0})\}$ and one obtains

$$\begin{aligned} q_{\Lambda,0}^{\mathrm{nonloc}}[\eta_0^1, \eta_0^2, \eta_{\Lambda\setminus 0}^{\mathrm{disc},V}, \sigma_{\partial_- \Lambda}] &= 2\mu^0_{V^+ \cup V^-}(\cosh J(\tilde{\sigma}_{x_0} + \tilde{\sigma}_{y_0})) \\ &= a\mu^0_{V^+ \cup V^-}(\tilde{\sigma}_{x_0}\tilde{\sigma}_{y_0}) + b, \\ q_{\Lambda,0}^{\mathrm{nonloc}}[\eta_0^1, \eta_0^2, \eta_{\Lambda\setminus 0}^{\mathrm{disc},V,z}, \sigma_{\partial_- \Lambda}] &= a\mu^0_{V^+ \cup V^- \cup z}(\tilde{\sigma}_{x_0}\tilde{\sigma}_{y_0}) + b \end{aligned} \tag{3.13}$$

for some positive constants a, b, for Λ sufficiently large. Here μ_W^0 is the ferromagnetic Ising Gibbs measure in the finite volume W with zero boundary conditions. The correlations on the right-hand side were seen in [12] to be different for large J, for arbitrarily large V, uniformly in the location of z. (Adding a site z destroys the independence and introduces a positive correlation between σ_{x_0} and σ_{y_0} once there is ferromagnetic order.) By the corollary of Proposition 2.1 this shows that $(\eta^{\mathrm{disc}}, \sigma)$ is a bad configuration for any σ.

Our point here was that while η^{disc} is not a decoupling configuration, the perturbed configurations $\eta^{\mathrm{disc},V}, \eta^{\mathrm{disc},V,z}$ are decoupling, leading to simple formulas for q^{nonloc} that are independent of the specific joint measure and independent of the value of p.

Models with random bonds that can be zero.

We note that the same mechanism as in [12] is responsible for the occurrence of bad configurations in models with random bonds. Although it is not difficult to see once the previous example is understood, this might be interesting, because it is also true e.g. for EA spin glasses of the type ((ii),(b)) from the Introduction. We have

Proposition 3.2. *Suppose that we are given a model of the form* (1.5) *in dimensions* $d \geq 2$ *where* $\nu(J_{x,e} = 0) > 0$ *and* $\nu(J_{x,e} = J^1) > 0$ *with* J^1 *sufficiently large.*

Decompose the lattice $\mathbf{Z}^d$ *into two half-spaces* $\mathbf{Z}^d_+ \cup \mathbf{Z}^d_-$ *that are separated by a hyperplane of bonds that we call* H. *Denote by* J^{disc} *the configuration of bonds that is equal to zero for bonds in* H *and equal to* J^1 *otherwise.*

Then $\xi = (J^{\mathrm{disc}}, \sigma)$ *is a bad configuration for any joint measure obtained as a limit point of* $\mathsf{P}(dJ)\mu_\Lambda^{\sigma^{\mathrm{b.c.}}_{\partial\Lambda}}[J](d\sigma)$.

Proof. Assume that the hyperplane is of the form $H = \{< x, y >: x_d = 0, y_d = 1\}$. Then we have $< 0, e_d > \in H$. In a similar fashion as above, for finite $\tilde{V} \subset (\mathbf{Z}^d)^*$ (a box on the dual lattice, centered around the origin), denote by

$J^{\mathrm{disc},\tilde{V}}$ the configuration that coincides with J^{disc} inside $\tilde{V}$ and vanishes outside $\tilde{V}$. Denote by $\tilde{V}^+$ ($\tilde{V}^-$) the occupied bonds in $\tilde{V}$ in the upper (lower) half-space. For a bond $b \in \tilde{V} \cap H$ denote by $J^{\mathrm{disc},\tilde{V},b}$ the configuration that has b as an additional non-empty coupling taking the value J^1.

To find a discontinuity, it suffices to look at pairs η_x^1 and η_x^2 that differ only by one coupling constant,

$$\eta_0^1 = (J^1_{0,e_1}, \ldots, J^1_{x,e_{d-1}}, 0)$$

and

$$\eta_0^2 = (J^1_{0,e_1}, \ldots, J^1_{x,e_{d-1}}, J^1_{x,e_d}).$$

Then the variation at the origin becomes

$$\exp\{\Delta H_0(\sigma_{\overline{0}}, \eta_0^1, \eta_0^2, \eta_{\partial 0}^{\mathrm{disc}})\} = \exp\{J^1 \sigma_0 \sigma_{e_d}\} = \exp\{-J^1\} + 2\sinh J^1 \; 1_{\sigma_0 = \sigma_{e_d}}$$

where we have written η^{disc} for the obvious configuration corresponding to J^{disc} (and will also do so for $\eta^{\mathrm{disc},\tilde{V}}$, $\eta^{\mathrm{disc},\tilde{V},b}$). So one obtains

$$\begin{aligned} q^{\mathrm{nonloc}}_{\Lambda,0}[\eta_0^1, \eta_0^2, \eta^{\mathrm{disc},\tilde{V}}_{\Lambda\setminus 0}, \sigma_{\partial_- \Lambda}] &= \exp\{-J^1\} + 2\sinh J^1 \; \hat{\mu}^0_{\tilde{V}^+ \cup \tilde{V}^-}(\tilde{\sigma}_0 = \tilde{\sigma}_{e_d}), \\ q^{\mathrm{nonloc}}_{\Lambda,0}[\eta_0^1, \eta_0^2, \eta^{\mathrm{disc},\tilde{V},b}_{\Lambda\setminus 0}, \sigma_{\partial_- \Lambda}] &= \exp\{-J^1\} + 2\sinh J^1 \; \hat{\mu}^0_{\tilde{V}^+ \cup \tilde{V}^- \cup b}(\tilde{\sigma}_0 = \tilde{\sigma}_{e_d}) \end{aligned} \tag{3.14}$$

for Λ sufficiently large. Here $\hat{\mu}^0_{\tilde{W}}$ is the ferromagnetic Ising Gibbs measure with zero boundary conditions on the vertex set of the graph whose bonds are $\tilde{W}$ with the coupling constant J^1. Now, in the very same way as in [12], the probabilities on the right-hand sides are seen to be different, for arbitrarily large $\tilde{V}$, uniformly in the location of b. By the corollary of Proposition 2.1 this shows the claim. □

3.3. Ising models with disordered nearest neighbor couplings

Denote by $(\mathbf{Z}^d)^*$ the lattice of bonds of $\mathbf{Z}^d$. We denote subsets of $(\mathbf{Z}^d)^*$ by symbols with tildes (like $\tilde{V}$). An application of Proposition 2.2 (resp. Proposition 2.4) yields the following

Theorem 3.1. *Consider an Ising model with random nearest neighbor couplings of the form* (1.5)*, in any dimension d.*

(i) *A configuration $\xi = (J, \sigma)$ is a good configuration for any joint measure obtained as a limit point of the finite-volume joint measures*

$$\mathsf{P}(dJ)\mu_\Lambda^{\sigma^{\mathrm{b.c.}}_{\partial\Lambda}}[J](d\sigma)$$

if

$$\sup_{J^+_\Lambda, J^-_\Lambda} \left| \mu_\Lambda^{\sigma^{\mathrm{b.c.}}_{\partial\Lambda}}[J_{\tilde{V}} J^+_{(\mathbf{Z}^d)^* \setminus \tilde{V}}](\tilde{\sigma}_x = \tilde{\sigma}_y) - \mu_\Lambda^{\sigma^{\mathrm{b.c.}}_{\partial\Lambda}}[J_{\tilde{V}} J^-_{(\mathbf{Z}^d)^* \setminus \tilde{V}}](\tilde{\sigma}_x = \tilde{\sigma}_y) \right| \to 0 \tag{3.15}$$

with $\tilde{V} \uparrow (\mathbf{Z}^d)^*$.

(ii) *Suppose moreover that we have the existence of a weak limit*

$$\lim_{\Lambda \uparrow \mathbf{Z}^d} \mu_\Lambda^{\sigma^{\text{b.c.}}_{\partial\Lambda}}[J] = \mu_\infty[J]$$

for a non-random boundary condition $\sigma^{\text{b.c.}}$, *for* P*-a.e.* J. *Denote by*

$$K(d\sigma, dJ) = \mathsf{P}(dJ)\mu_\infty[J](d\sigma)$$

the corresponding joint measure.

A configuration $\xi = (J, \sigma)$ *is a bad configuration for* K, *if there exists an increasing choice of volumes* $\tilde{\Lambda}(\tilde{V})$ *and configurations* $J^{\tilde{V}}, \bar{J}^{\tilde{V}}$, *such that , for* P*-a.e.* $\tilde{J}$, *we have that*

$$\begin{aligned} &\liminf_{\tilde{V}\uparrow(\mathbf{Z}^d)^*} \mu_\infty[J_{\tilde{V}} \bar{J}^{\tilde{V}}_{\tilde{\Lambda}(\tilde{V})\setminus\tilde{V}}, \tilde{J}_{(\mathbf{Z}^d)^*\setminus\tilde{\Lambda}(V)}](\tilde{\sigma}_x = \tilde{\sigma}_y) \\ &\quad > \limsup_{\tilde{V}\uparrow(\mathbf{Z}^d)^*} \mu_\infty[J_{\tilde{V}} J^{\tilde{V}}_{\tilde{\Lambda}(\tilde{V})\setminus\tilde{V}}, \tilde{J}_{(\mathbf{Z}^d)^*\setminus\tilde{\Lambda}(V)}](\tilde{\sigma}_x = \tilde{\sigma}_y) \end{aligned} \tag{3.16}$$

for some nearest neighbor pair $< x, y >$.

Proof. To check the condition of Proposition 2.2, it suffices to look at pairs η^1_x and η^2_x that differ only by one coupling constant, say

$$\eta^1_x = (J_{x,e_1}, \ldots, J_{x,e_{j-1}}, J^1, J_{x,e_{j+1}}, \ldots, J_{x,e_d})$$

and

$$\eta^2_x = (J_{x,e_1}, \ldots, J_{x,e_{j-1}}, J^2, J_{x,e_{j+1}}, \ldots, J_{x,e_d}).$$

Put $y = x + e_j$. The variations at the site x then become

$$\begin{aligned} \exp\{\Delta H_x(\sigma_{\overline{x}}, \eta^1_x, \eta^2_x)\} &= \exp\{(J^2 - J^1)\sigma_x\sigma_y\} \\ &= \exp\{(J^1 - J^2)\} + 2\sinh(J^2 - J^1)\, 1_{\sigma_x = \sigma_y} \end{aligned} \tag{3.17}$$

which is analogous to formula (3.1) for the random field model.

Writing out the condition (2.12) from Proposition 2.2 then essentially amounts to the criterion (3.15) given in the theorem, except that possibly different values J at the bond $< x, y >$ can appear. However, there is a simple formula analogous to formula (3.2) for the random field model relating the probabilities of the event $\sigma_x = \sigma_y$ for different values of $J_{<x,y>}$ that is obtained by the finite-volume perturbation formula. From this an argument like the one given for the random field model after (3.2) shows that the validity of condition (3.15) is independent of the value of $J_{<x,y>}$. This proves statement (i).

To show that (J, σ) is a bad configuration (for any σ) by means of Proposition 2.4 we have to look at

$$\limsup_{\tilde{V}\uparrow(\mathbf{Z}^d)^*} \mu_\infty[J^1_{<x,y>}, J_{\tilde{V}\setminus<x,y>} J^{\tilde{V}}_{\tilde{\Lambda}(\tilde{V})\setminus\tilde{V}}, \tilde{J}_{(\mathbf{Z}^d)^*\setminus\tilde{\Lambda}(V)}]\big(\exp\{(J^2 - J^1)\tilde{\sigma}_x\tilde{\sigma}_y\}\big)$$

and

$$\liminf_{\tilde{V}\uparrow(\mathbb{Z}^d)^*} \mu_\infty[J^1_{<x,y>}, J_{\tilde{V}\setminus<x,y>}\bar{J}^{\tilde{V}}_{\tilde{\Lambda}(\tilde{V})\setminus\tilde{V}}, \tilde{J}_{(\mathbb{Z}^d)^*\setminus\tilde{\Lambda}(V)}]\big(\exp\{(J^2-J^1)\tilde{\sigma}_x\tilde{\sigma}_y\}\big) \quad (3.18)$$

and find two sequences of conditionings $J^{\tilde{V}}$ and $\bar{J}^{\tilde{V}}$ such that the lower expression is strictly bigger than the upper one. Assuming that $J^2 > J^1$, this is true if and only if

$$\begin{aligned}\liminf_{\tilde{V}\uparrow(\mathbb{Z}^d)^*} &\mu_\infty[J^1_{<x,y>}, J_{\tilde{V}\setminus<x,y>}\bar{J}^{\tilde{V}}_{\tilde{\Lambda}(\tilde{V})\setminus\tilde{V}}, \tilde{J}_{(\mathbb{Z}^d)^*\setminus\tilde{\Lambda}(V)}](\tilde{\sigma}_x = \tilde{\sigma}_y) \\ > \limsup_{\tilde{V}\uparrow(\mathbb{Z}^d)^*} &\mu_\infty[J^1_{<x,y>}, J_{\tilde{V}\setminus<x,y>}J^{\tilde{V}}_{\tilde{\Lambda}(\tilde{V})\setminus\tilde{V}}, \tilde{J}_{(\mathbb{Z}^d)^*\setminus\tilde{\Lambda}(V)}](\tilde{\sigma}_x = \tilde{\sigma}_y). \quad (3.19)\end{aligned}$$

Using the argument presented for the RFIM we see that this is true if and only if the same strict inequality holds for any other value of $J_{<x,y>}$ replacing $J^1_{<x,y>}$. This proves statement (ii). □

Finally we would like to discuss the relevance of Theorem 3.1 on a heuristic level in application to a random bond ferromagnet.

Heuristics considerations: gibbsianness destroyed by interfaces.

Assume $d \geq 2$. Suppose that the random bonds $J_{x,e}$ take two values $0 < J^1 < J^2 < \infty$ with positive probability, independently of the bond (x, e). We assume that J^1 is smaller than the critical inverse temperature of the corresponding homogeneous Ising ferromagnet. J^2 should be large enough and $\nu[J_{x,e} = J^1]$ should be small enough such that there is ferromagnetic order in the disordered model with P-probability one.

Let us at first look at $\sigma_x^{\text{b.c.}} \equiv 1$ boundary conditions. Then we expect almost sure joint Gibbsianness. Indeed, criterion (i) of Theorem 3.1 should be satisfied for P-a.e. configuration of couplings J, for the following reason.

Let us assume that the realization J is from the full measure set of couplings for which the finite-volume Gibbs measures converge to a ferromagnetic infinite-volume Gibbs measure. Let us check the expected behavior with two 'extreme' choices of perturbations.

Consider first a typical perturbation J^+ that does have enough stronger couplings to support the ferromagnetic order. Then the state $\mu^+_\Lambda[J_{\tilde{V}}J^+_{(\mathbb{Z}^d)^*\setminus\tilde{V}}]$ should look like $\mu^+_\infty[J]$ locally, for sufficiently large inner volume $\tilde{V}$ and any (bigger) Λ.

Choosing next $J^+ \equiv J^1$ (the weaker couplings) will however destroy the ferromagnetic order in the annulus. Hence the boundary conditions should be forgotten for sufficiently large annulus and the volume V will approximately feel open boundary conditions. (This argument is of course strictly true for the case $J^1 = 0$). The corresponding state should then approximately look like $(1/2)(\mu^+_\infty[J] + \mu^-_\infty[J])$ for large $\tilde{V}$. This would of course lead to different

expectations on general observables compared to those of $\mu_\infty^+[J]$. The point is however that the expectations of the different states on the event $\{\sigma_x = \sigma_y\}$ are the same, due to spin-flip symmetry.

We expect that in general, choosing whatever annulus should result in one of the two possibilities or a linear combination of them.

This provides an example that shows that although a phase transition occurs by varying the disorder variables in a large annulus, it leads to the same expectations on the single site perturbation of the Hamiltonian. Thus the resulting state can be Gibbs.

Let us now look at Dobrushin boundary conditions, i.e. we start from finite-volume Gibbs measures in boxes centered around the origin with plus boundary conditions on the top half, and minus boundary condition on the lower half. We assume additionally that we are in dimensions $d \geq 4$, that J^2 is large enough and $\nu[J_{x,e} = J^1]$ small enough such that there are interface states (random 'Dobrushin' states [6]) in the disordered model with P-probability one. The existence of such states that are perturbations of the spin configuration that is all plus in the upper half-space and all minus in the lower half-space was proved in [4] in the SOS-approximation of the model. (For complementary information about disordered interface models, see [5, 22].)

Now we expect almost sure non-Gibbsianness for the resulting infinite-volume joint measure, different from the model with + boundary conditions. Indeed, criterion (ii) of Theorem 3.1 should be satisfied for P-a.e. configuration of couplings J, for the following reason.

We fix a nearest neighbor pair $< x, y >$ located at, and perpendicular to, the base plane (whose intersection with the boundary of Λ is the boundary between plus and minus boundary spins). Again we look first at a typical perturbation J^+. We expect that the infinite-volume Dobrushin states $\mu_\infty^\pm[J]$ have the locality property, i.e. for P-a.e. perturbation $\tilde{J}$ we have that

$$\lim_{\tilde{V}\uparrow(\mathbb{Z}^d)^*} \mu_\infty^\pm[J_{\tilde{\Lambda}}\tilde{J}_{(\mathbb{Z}^d)^*\setminus\tilde{\Lambda}}](\tilde{\sigma}_x = \tilde{\sigma}_y) = \mu_\infty^\pm[J_{(\mathbb{Z}^d)^*}](\tilde{\sigma}_x = \tilde{\sigma}_y) \tag{3.20}$$

for any nearest neighbor pair $< x, y >$. A corresponding statement could in principle be extracted from the renormalization group analysis of [4] for the corresponding SOS-model.

Choosing next the exceptional configuration $J^+ \equiv J^1$ in an annulus $\tilde{\Lambda}(\tilde{V})\setminus\tilde{V}$ that is sufficiently large will destroy the ferromagnetic order in the annulus and decouple the volume $\tilde{V}$ from the outside. This should result in

$$\begin{aligned}\lim_{\tilde{V}\uparrow(\mathbb{Z}^d)^*} & \mu_\infty^\pm[J_{\tilde{V}}J^1_{\tilde{\Lambda}(\tilde{V})\setminus\tilde{V}}, \tilde{J}_{(\mathbb{Z}^d)^*\setminus\tilde{\Lambda}(V)}](\tilde{\sigma}_x = \tilde{\sigma}_y) \\ = \; & \frac{1}{2}(\mu_\infty^+[J_{(\mathbb{Z}^d)^*}](\tilde{\sigma}_x = \tilde{\sigma}_y) + \mu_\infty^-[J_{(\mathbb{Z}^d)^*}](\tilde{\sigma}_x = \tilde{\sigma}_y)).\end{aligned} \tag{3.21}$$

Note that both terms of the right-hand side are the same, due to spin-flip symmetry. This will differ from the expectation in the interface-state (3.20), so that we believe that criterion (3.16) should be satisfied.

Let us summarize the argument informally: for a large annulus of weak couplings arbitrarily far away from the origin, the Dobrushin state should look like a symmetric mixture of the plus and minus state. Since Dobrushin state and plus (resp. minus) state will have different expectations on the nearest neighbor spin correlations (the spin observable conjugate to the disorder) we have a non-local dependence on the coupling constants that leads to non-Gibbsianness. Let us point out that, in order to reach the conclusion of non-Gibbsianness even on the heuristic level we have presented it, we really needed Theorem 3.1 (ii) that follows from Proposition 2.4, a result that involves *typical* configurations (as opposed to the criterion of Proposition 2.3 used for the random field Ising model, a result that involves *uniform* estimates).

In the random field model we produced a jump in the conditional expectations by perturbing the magnetic fields, where the form of the spin perturbation played no role. In contrast to that, the application of Proposition 2.4 gives us only the existence of a suitable spin perturbation such that, along with the perturbation of the couplings described above, we can produce a jump. Note that there is the following fundamental difference between the random field and the random bond Ising model which is responsible for this: in the random field model, one is able to select a phase by choosing the disorder variables (magnetic fields) in a large annulus, no matter what the disorder variables even further outside will look like. In contrast to that, one is not able to 'restore' a Dobrushin state in a random bond model by a suitable choice of J's in a large annulus, if the $\pm$ boundary conditions have been forgotten, because the couplings further outside were too weak.

Acknowledgments

The author thanks A. van Enter for a private explanation of reference [12] and valuable comments on the manuscript that led to an improvement of the presentation.

References

[1] M. Aizenman and J. Wehr (1990) Rounding effects of quenched randomness on first-order phase transitions. *Commun. Math. Phys.* **130**, 489–528.

[2] J. Bricmont and A. Kupiainen (1988) Phase transition in the 3d random field Ising model. *Commun. Math. Phys.* **142**, 539–572.

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

[4] A. Bovier and C. Külske (1994) A rigorous renormalization group method for interfaces in random media. *Rev. Math. Phys.* **6** (3), 413–496.

[5] A. Bovier and C. Külske (1996) There are no nice interfaces in $2+1$ dimensional SOS-models in random media. *J. Stat. Phys.* **83**, 751–759.

[6] R.L. Dobrushin (1972) Gibbs states describing a coexistence of phases for the three-dimensional Ising model. *Theory Probab. and Appl.* **17**, 582–600.

[7] R.L. Dobrushin (1995) Lecture given at the workshop 'Probability and Physics', Renkum, August 1995.

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

[9] A.C.D. van Enter (1998) The renormalization-group peculiarities of Griffiths and Pearce: what have we learned? Available as preprint 98-692 at http://www.ma.utexas.edu/mp_arc.

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

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

[12] A.C.D. van Enter, C. Maes, R.H. Schonmann and S. Shlosman (1998) The Griffiths singularity random field. To appear in *AMS Dobrushin Memorial Volume*. Available as preprint 98-764 at http://www.ma.utexas.edu/mp_arc.

[13] R. Fernández (1999) Measures for lattice systems. *Physica A* **263**, 117-130. (Invited papers from Statphys 20, Paris (1998)). Available as preprint 98-567 at http://www.ma.utexas.edu/mp_arc.

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

[15] R.B. Griffiths (1969) Non-analytic behavior above the critical point in a random Ising ferromagnet. *Phys. Rev. Lett.* **23**, 17–20.

[16] C. Külske (1993) Ph. D. Thesis, Ruhr-Universität, Bochum.

[17] C. Külske (1997) Metastates in disordered mean-field models: random field and Hopfield models. *J. Stat. Phys.* **88** (5/6), 1257–1293.

[18] C. Külske (1998) Limiting behavior of random Gibbs measures: metastates in some disordered mean field models. In: *Mathematical Aspects of Spin Glasses and Neural Networks,* A. Bovier, P. Picco (eds.), Progr. Probab. **41**, Birkhäuser, Boston, 151–160.

[19] C. Külske (1998) Metastates in disordered mean-field models II: The superstates. *J. Stat. Phys.* **91** (1/2), 155–176.

[20] C. Külske (1998) A random energy model for size dependence: recurrence vs. transience. *Probab. Theory and Relat. Fields* **111**, 57–100.

[21] C. Külske (1998) The continuous spin random field model: ferromagnetic ordering in $d \geq 3$. To appear in *Rev. Math. Phys.* Available as preprint 98-175 at http://www.ma.utexas.edu/mp_arc/.

[22] C. Külske (1998) Stability for a continuous SOS-interface model in a randomly perturbed periodic potential. Available as preprint 98-768 at http://www.ma.utexas.edu/mp_arc/.

[23] C. Maes, F. Redig and A. Van Moffaert (1999) Almost Gibbsian versus weakly Gibbsian measures. *Stoch. Process. Appl.* **79** (1), 1–15. Available as preprint 98-193 at http://www.ma.utexas.edu/mp_arc/.

[24] C.M. Newman (1997) Topics in Disordered Systems. *Lect. Math. ETH Zürich.* Birkhäuser Verlag, Basel.

[25] C.M. Newman and D.L. Stein (1996) Spatial inhomogeneity and thermodynamic chaos. *Phys. Rev. Lett.* **76** (25), 4821.

[26] C.M. Newman and D.L. Stein (1997) Metastate approach to thermodynamic chaos. *Phys. Rev. E (3)* **55** no.5, part A, 5194–5211.

[27] C.M. Newman and D.L. Stein (1998) Simplicity of state and overlap structure in finite-volume realistic spin glasses. *Phys. Rev. E (3)* **57** no.2, part A, 1356–1366.

[28] C.M. Newman and D.L. Stein (1998) Thermodynamic chaos and the structure of short-range spin glasses. In: *Mathematical Aspects of Spin Glasses and Neural Networks,* A. Bovier, P. Picco (eds.), Progr. Probab. **41**, Birkhäuser, Boston, 243–287.

[29] T. Seppäläinen (1995) Entropy, limit theorems, and variational principles for disordered lattice systems. *Commun. Math. Phys.* **171**, 233–277.

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