Weierstraß-Institut
für Angewandte Analysis und Stochastik

Leibniz-Institut im Forschungsverbund Berlin e. V.

Preprint ISSN 2198-5855

Gibbsianness and non-Gibbsianness for Bernoulli lattice fields under removal of isolated sites

Benedikt Jahnel[1], Christof Külske[2]

submitted: September 29, 2021

[1] Weierstrass Institute
Mohrenstr. 39
10117 Berlin
Germany
E-Mail: Benedikt.Jahnel@wias-berlin.de

[2] Ruhr-Universität Bochum
Fakultät für Mathematik
44801 Bochum
Germany
E-Mail: Christof.Kuelske@ruhr-uni-bochum.de

No. 2878
Berlin 2021

2020 *Mathematics Subject Classification.* 60D05, 60K35, 82B20.

Key words and phrases. Gibbsianness, Gibbs-uniqueness, Bernoulli field, local thinning, two-layer representation, Dobrushin uniqueness, Peierls' argument.

We thank Nils Engler for inspiring discussions. This work was funded by the German Research Foundation under Germany's Excellence Strategy MATH+: The Berlin Mathematics Research Center, EXC-2046/1 project ID: 390685689 and the German Leibniz Association via the Leibniz Competition 2020.

Edited by
Weierstraß-Institut für Angewandte Analysis und Stochastik (WIAS)
Leibniz-Institut im Forschungsverbund Berlin e. V.
Mohrenstraße 39
10117 Berlin
Germany

Fax: +49 30 20372-303
E-Mail: preprint@wias-Cberlin.de
World Wide Web: http://www.wias-Cberlin.de/

Gibbsianness and non-Gibbsianness for Bernoulli lattice fields under removal of isolated sites

Benedikt Jahnel, Christof Külske

Abstract

We consider the i.i.d. Bernoulli field μ_p on $\mathbb{Z}^d$ with occupation density $p \in [0,1]$. To each realization of the set of occupied sites we apply a thinning map that removes all occupied sites that are isolated in graph distance. We show that, while this map seems non-invasive for large p, as it changes only a small fraction $p(1-p)^{2d}$ of sites, there is $p(d) < 1$ such that for all $p \in (p(d),1)$ the resulting measure is a non-Gibbsian measure, i.e., it does not possess a continuous version of its finite-volume conditional probabilities. On the other hand, for small p, the Gibbs property is preserved.

1 Introduction

Random fields under local maps are defined and analyzed in different fields of probability and statistics. In these studies a random field is typically a stochastic process whose variables take values in a subset of the real numbers, and which has a geometric index set such as the infinite lattice, a graph with infinite vertex set, or the Euclidean space. It is of interest to understand the behavior of the given random field with distribution μ under application of a map T that acts on infinite-volume realizations of the process, and investigate resulting properties of the image process whose distribution we denote by μ'. Relevant deterministic maps T typically have local-dependence windows. They may also be generalized to stochastic kernels that act locally. In this general setup let us call μ the *first-layer measure*, and the measure μ', which we shall be mostly interested in, the *second-layer measure*.

This setup is of theoretical interest, but also occurs in many applications of natural sciences, engineering, and statistics. We mention thinning transformations of Poisson point processes (in which points from a realization are omitted according to local rules) [Bre79, Bro79, Ish80, RS91, Las93, Bal05, MS10, BK19], transformations in image analysis [Gem90, Guy95], and renormalization group transformations in statistical mechanics (where a physical system like a ferromagnet is considered on increasingly large scales by means of maps which forget details on small scales) [GP78, GP79, vEFS93]. Processes appearing as local maps in this way can also be viewed as generalizations of the much used *hidden Markov models* [EAM95].

Hidden Markov models are images under local kernels of an underlying *first-layer* Markov chain, and appear when noisy observations shall be modeled, and so the generalization is made to a situation with spatial index sets, and more complex dependence structures.

It was discovered first in the context of such renormalization group transformations that strictly local transformations acting on a spatially Markovian random field μ indexed by a lattice, may result in singularities in the image measure μ'. Two concrete examples for this are provided by the low-temperature Ising model under a block average transformation, or the projection to a sublattice [Sch89, vEFS93, BKL98, LN13, BC18]. The original motivation of such renormalization

DOI 10.20347/WIAS.PREPRINT.2878 Berlin 2021

group transformations was, suggested by heuristic schemes of theoretical physics, to understand the iterated *coarse-graining dynamics* of the level of Hamiltonians to investigate critical behavior [WK74].

The singularity in the second-layer measure μ' means, that μ' loses not only the spatial Markov property of μ under the map T (which is less surprising), but it even loses the more general *Gibbs property*. This is more severe, it means that the infinite second-layer system acquires internal long-range dependence, and in particular does not posses a well-behaved Hamiltonian with good summability properties of its interaction potentials anymore. The singular long-range dependence appears on the level of finite-volume conditional probabilities of the image measure, which are not quasilocal functions of their conditioning. Put equivalently, finite-volume sub-systems depend on their boundary conditions arbitrarily far away, and their behavior cannot be described by kernels that are continuous in product topology. This may cause standard theory of infinite-volume states, including the variational principle, to fail, see the examples in [KLNR04].

A variety of examples have been studied ever since, where non-Gibbsian behavior was proved to occur with different mechanisms, but always in regimes of sufficiently strong coupling, where the first-layer measure differs much from independence. (Having said this, there are known examples that show that the range of temperatures where non-Gibbsian behavior in the image system occurs, may be larger than the critical temperature for the first-layer system [HK04]) Moreover, examples have been found where, the set of discontinuity points is even of full measure w.r.t. μ' itself, which is the strongest form of singularity [KLNR04, vEEIK12, JK17b, BKK20].

Main subclasses of relevant transformations which have been studied were projections in terms of a variety of deterministic maps [vELN17, LN13, KR17, JK17a, HKK19], and stochastic time-evolutions [LNR02, vEFdHR02, KLN07, vEEIK12, RR14, FdHM14, JK17b, KRvZ21], in various underlying geometries of lattice models, mean-field models, Kac-models, and models in the continuum [JK17b].

Informal result: Even independent fields may become non-Gibbs under projections

In the present paper we provide a new and simple example that shows that a natural local transform of range 1 can produce singularities, even when it is applied to an *independent field*. In our example we chose as the first-layer field the i.i.d. Bernoulli lattice field μ_p on the integer lattice, with state space $\{0,1\}^{Z^d}$, and occupation probability $p \in [0,1)$. The Bernoulli lattice field in itself is studied in site-percolation, where one asks for existence of infinite clusters and refined connectedness properties [Gri99]. It also drives more complex processes, in statistical mechanics of disordered systems [Gri97, vEMSS00, Bov06], and elsewhere in probability [Adl91, BR06, BDC13, Kum14, JK20] and its application. We then study the second-layer measure μ'_p that appears as an image under application of the concrete range-one map T that is defined by removing from a realization of occupied sites the *occupied isolated sites*. T is a projection map as it satisfies $T^2 = T$, and we will call it *the projection to non-isolates*. Hence it keeps from a realization of occupied sites only the occupied clusters of size of at least two. This includes the infinite cluster, in case there is one, i.e., in the percolation regime of large enough p. We may also view T as a simple smoothing transformation, as isolated 'dust' of occupied sites (or 'pixels') is forgotten under the map.

What to expect for the second-layer measure μ'_p? As the μ_p-probability that a given site is isolated equals $p(1-p)^{2d}$, the map T seems non-invasive, in particular for probabilities p close to 1. In particular the removed sites do no percolate in this regime. So one might naively conjecture that the second-layer measure shall not be much affected and well-behaved, and in particular is representable as a Gibbsian distribution with quasilocal conditional probabilities.

DOI 10.20347/WIAS.PREPRINT.2878 Berlin 2021

As a main message of this paper, we prove that this is not the case: We show that μ'_p is spatially *non-Markovian* and *non-quasilocal*, when $p < 1$ is large, see Theorem 2.1. We hope that this result is of some interest for the percolation community. We complement this result by proving regularity of the projected measure when $p \geq 0$ is small enough see, Theorem 2.2. This implies the existence of a Gibbs-non Gibbs transition driven by p.

Let us finally discuss our map T that projects to non-isolates (and removes the isolates) from a dual perspective. Namely, T has a natural companion map T^* that is again a projection map. T^* does precisely the opposite, it projects to the isolated sites (and removes the non-isolates).

There is independent interest in the action of T^* to the i.i.d. Bernoulli lattice field, for the reason that it produces the *thinned Bernoulli lattice field*, in which all occupied sites are separated. This thinned Bernoulli lattice field is relevant, as it is the lattice analogue of the well-known and much studied Matérn process in the continuum [Mat60, MHW10, BB12]. The latter by definition is derived from a first-layer Poisson process in Euclidean space by removing all points in the realization that have at least one point in the Euclidean ball of radius one.

Clearly, the second-layer measures of both maps T, T^* acting jointly on the same Bernoulli lattice-field realization appear in a natural coupling. As a first guess one may conjecture from this that, either both second-layer measures are Gibbs, or both are non-Gibbs. We warn the reader that this is too naive, not only on the level of proofs, but also the statement may be false. We leave the analysis of the companion process, the thinned Bernoulli lattice field, to another study.

The paper is organized as follows. In Section 2 we present the setting and our main results and non-Gibbsianness and Gibbsianness for the Bernoulli field under the removal-of-isolates transformation. In Section 2 we present the corresponding proofs.

2 Setting and main results

To define our process we start from the $\Omega = \{0,1\}^{\mathbb{Z}^d}$-valued i.i.d. Bernoulli field μ_p with parameter $p \in [0,1]$. We consider realizations of the Bernoulli field under the application of the transformation $T : \Omega \to \Omega$ given by

$$(T\omega)_x := \omega'_x = \omega_x \Big(1 - \prod_{y \in \partial x} (1 - \omega_y)\Big),$$

where ∂x denotes the set of nearest neighbors of $x \in \mathbb{Z}^d$ in $\mathbb{Z}^d$, equipped with the usual neighborhood structure. In words, T is the projection to the non-isolates, see Figure 1 for an illustration. The image measure under the transformation

$$\mu'_p := \mu_p \circ T^{-1}$$

is supported on the subset $\Omega' := T(\Omega)$ of sites that obey the non-isolation constraint.

Intuitively the application of T should not change the measure very much at large p close to 1, where a typical configuration consists of a large percolating cluster and very few isolated sites. In this regime one may view T as a cleansing transformation that wipes away the smallest dust of isolated sites, and keeps the apparent main parts. From the definition of T as a local map it is also obvious that for variables at sites of graph distance greater equal than 3 are independent, under μ'_p, so that one could expect that μ'_p is a nicely behaved measure.

Recall that a *specification* $\gamma = (\gamma_\Lambda)_{\Lambda \Subset \mathbb{Z}^d}$ is a consistent and proper family of conditional probabilities, i.e., for all $\Lambda \subset \Delta \Subset \mathbb{Z}^d$, $\omega_\Lambda \in \Omega_\Lambda := \{0,1\}^\Lambda$ and $\hat{\omega} \in \Omega$, we have that $\int_\Omega \gamma_\Delta(\mathrm{d}\tilde{\omega}|\hat{\omega})\gamma_\Lambda(\omega_\Lambda|\tilde{\omega}) =$

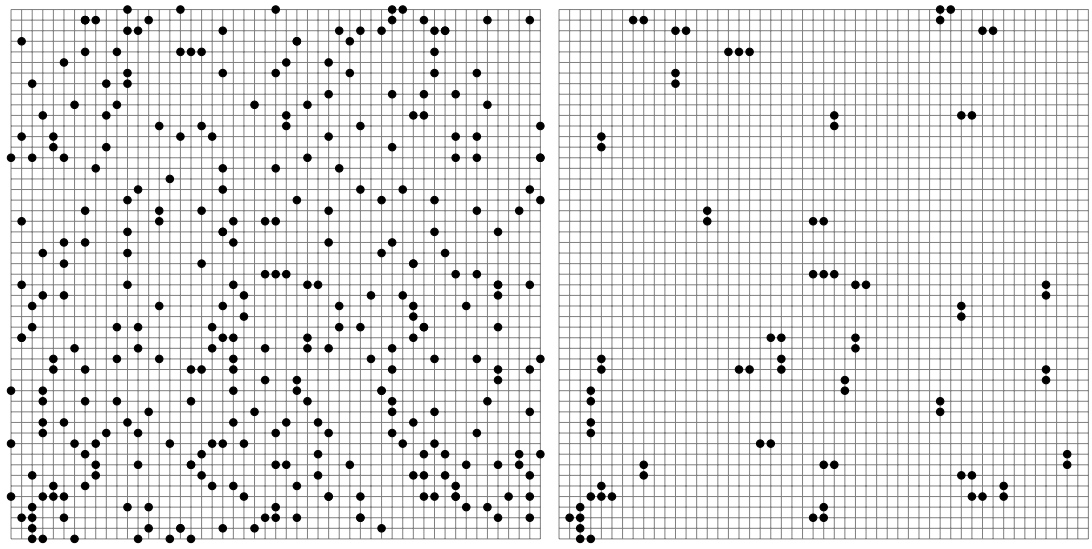

Figure 1: Realization of a Bernoulli field (left) and its image under the transformation T (right).

$\gamma_\Delta(\omega_\Lambda|\hat{\omega})$, and for all $\omega_{\Lambda^c} \in \Omega_{\Lambda^c}$ we have $\gamma_\Lambda(\omega_{\Lambda^c}|\hat{\omega}) = \mathbb{1}\{\omega_{\Lambda^c} = \hat{\omega}_{\Lambda^c}\}$. A specification is called *quasilocal*, if for all $\Lambda \Subset \mathbb{Z}^d$ and $\hat{\omega}_\Lambda \in \Omega_\Lambda$, the mapping $\omega \mapsto \gamma_\Lambda(\hat{\omega}_\Lambda|\omega)$ is continuous with respect to the product topology on Ω. We say that γ is a specification for some random field μ on Ω, if it satisfies the DLR equations, i.e., for all $\Lambda \Subset \mathbb{Z}^d$ and $\omega_\Lambda \in \Omega_\Lambda$, we have that $\int_\Omega \mu(\mathrm{d}\tilde{\omega})\gamma_\Lambda(\omega_\Lambda|\tilde{\omega}) = \mu(\omega_\Lambda)$.

Our present result shows that this is not the case in the whole parameter regime, but the following is true: For large p, μ'_p is not spatially Markovian, it is not even a Gibbs measure in the sense of existence of a version of its finite-volume conditional probabilities which is continuous with respect to the product topology on Ω'. More precisely we have the following theorem.

Theorem 2.1. (Non-Gibbsianness for large p) Consider the image measure μ'_p of the Bernoulli field on $\mathbb{Z}^d$ under the map to the non-isolates in lattice dimensions $d \geq 2$. Then, there is $p_c(d) < 1$ such that for $p \in (p_c(d), 1)$, there is no quasilocal specification γ' for μ'_p.

The result shows that Gibbsian descriptions of thinning processes of various types derived from Bernoulli or Poissonian fields are by no means obvious, and that more research on such processes in discrete and continuous setups is necessary.

In order to complement the above non-Gibbsianness result, let us also present the following theorem on the existence of a quasilocal specification for μ'_p for small values of p.

Theorem 2.2. (Gibbsianness for small p) Consider the image measure μ'_p of the Bernoulli field on $\mathbb{Z}^d$ under the map to the non-isolates in lattice dimensions $d \geq 1$. Then, for $p < 1/(2d)$, there exists a quasilocal specification γ' for μ'_p.

As we can see from the proofs, for $p < 1/(2d)$, in fact the continuity of γ' is even exponentially fast. In summary, the statements of Theorems 2.1 and 2.1 indicate a phase-diagram of Gibbsianness of thinned Bernoulli fields under the local non-isolation constraint as exhibited in Figure 2.

In the following section, we present the proofs.

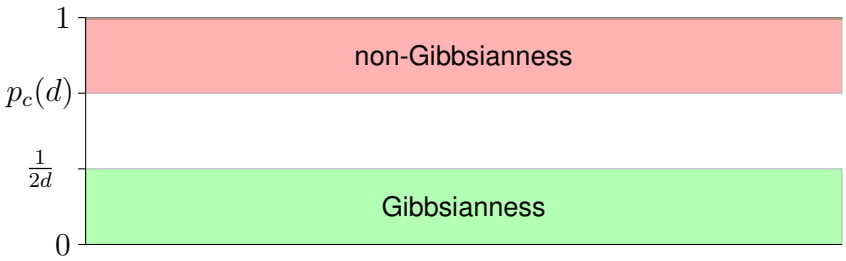

Figure 2: Illustration of Gibbs-non-Gibbs transitions in p for the thinned Bernoulli field under non-isolation constraint.

3 Proofs

The key idea of the proof is to first re-express conditional probabilities of μ'_p in finite volumes in terms of a first-layer constraint model in which occupied sites have to be isolated. Indeed, for large p, there are two distinct groundstates given by the (shifted) checkerboard configurations. We can leverage a Peierls' argument in order to show that the first-layer constraint model exhibits a phase transition of translational symmetry breaking. We note that the argument works even though there is *no* spin-flip symmetry in the system. The translational symmetry breaking gives rise to a point of discontinuity for which we subsequently show that it is present for any system of finite-volume conditional probabilities for μ'_p, i.e., it is an essential discontinuity. The converse case of small p can be handled by arguments using Dobrushin uniqueness techniques.

3.1 Proof of Theorem 2.1: Non-Gibbianness

The main ingredient for the proof is to exhibit one non-removable bad configuration for conditional probabilities. For this, we will use the so-called two-layer view, in which one needs to understand the Bernoulli field conditional on a fixed image configuration. We choose as the image configuration the all empty configuration, for which the first-layer measure becomes the Bernoulli field μ_p *conditional on isolation*.

We proceed as follows. In Section 3.1.1 we exhibit a phase transition for the latter model at large p, in which translation symmetry is broken, which can be selected via suitable shapes of loophole-volumes. The technique is based on a (slightly non-standard) Peierls argument. In Section 3.1.2 we then show how this implies non-Gibbsianness of the image measure. This is based on the proof that jumps of conditional probabilities occur for certain suitably chosen local patterns, which allow to make a transparent connection to the first-layer model in suitable connected boxes, where the Peierls argument from Section 3.1.1 was made to work.

3.1.1 Translational-symmetry breaking via a Peierls argument for the conditional first-layer model

For the purpose of showing that the empty configuration is bad for any specification, we will analyze the following particular finite-volume first-layer measures, and we will restrict to particular volumes Λ. Namely, let us consider finite volumes Λ that have a *shape of type* 0, and put fully occupied boundary conditions all 1 outside of Λ. By this we mean that Λ has a shape which allows to put the checkerboard groundstate of zeros and ones inside Λ for which the origin obtains the value 0 such that one obtains

a configuration compatible with the boundary condition, see Figure 4 for an illustration. We define

$$\nu_\Lambda(\omega_\Lambda) := \frac{\mu_{p,\Lambda}(\omega_\Lambda 1_{T(\omega_\Lambda 1_{\Lambda^c})|_\Lambda = 0_\Lambda})}{\mu_{p,\Lambda}(T(\sigma_\Lambda 1_{\Lambda^c})|_\Lambda = 0_\Lambda)},$$

where $\mu_{p,\Lambda}$ is the Bernoulli product measure in Λ. Hence, by definition ν_Λ is the Bernoulli measure conditioned on isolation of ones inside Λ, where the isolation constraint remembers also the fully occupied boundary condition.

A similar definition is made for volumes of type 1. For us, large boxes B_L centered around the origin with sidelength $2L$ with a loophole boundary, will be useful, see the illustration in Figure 3. For such type-0 boxes B_L we will show in this section that

$$\sup_L \nu_{B_L}(\omega_0 = 1) \leq \epsilon(p), \qquad \text{with} \qquad \lim_{p\uparrow 1} \epsilon(p) = 0. \tag{3.1}$$

This means that, with large probability, the origin copies the information from the boundary condition.

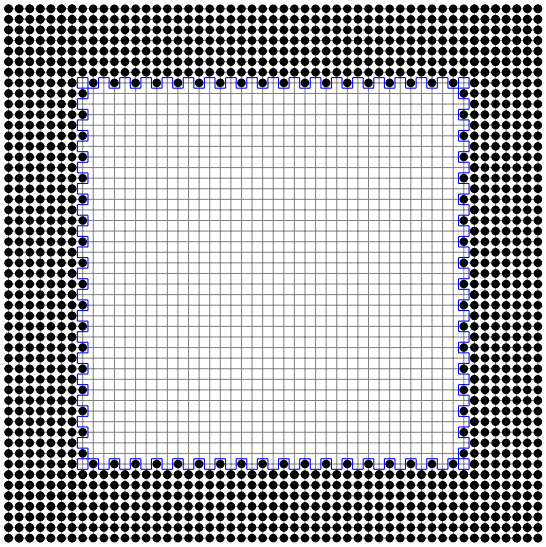

Figure 3: Illustration of a type-0 volume with loophole boundary.

Similarly, we will prove that the spin at the origin for the box shifted by a lattice unit vector e satisfies

$$\sup_L \nu_{B_L+e}(\omega_0 = 0) \leq \epsilon(p), \qquad \text{with} \qquad \lim_{p\uparrow 1} \epsilon(p) = 0. \tag{3.2}$$

This is an essential step as it proves that the shape of the volume B_L induces a phase transition for the first-layer constrained model, and there is breaking of translational symmetry. To complete the proof of essential badness of the empty configuration on the second layer, we will however need to go one step further, and connect to the measure on the second layer. This will be done in Lemma 3.2 below.

Note that configurations of the model are energetically equivalent under a lattice shift. They are not equivalent under the site-wise *spin-flip* that exchanges zeros and ones, much unlike the Ising antiferromagnet in zero external field. Therefore, the Peierls argument we are about to give has to be different from the one for the Ising ferromagnet or antiferromagnet. Namely, the Peierls argument we will present involves suitable lattice shifts of parts of configurations, while the standard more straightforward Peierls argument for the Ising model involves spin-flips.

Consider the nearest-neighbor graph with vertex set $\mathbb{Z}^d$. Consider, for a spin configuration ω, the set of sites

$$\Gamma(\omega) := \{x \in \mathbb{Z}^d : \text{ there exists } y \in \partial x \text{ such that } \omega_x = \omega_y = 0\}.$$

Note that there is a one-to-one correspondence between configurations ω that satisfy the neighborhood constraint, and sets $\Gamma(\omega)$. Note that outside of $\Gamma(\omega)$, the configuration ω looks like one of the two groundstates formed by the two possible checkerboard configurations of zeros and ones. Indeed, each site $x \notin \Gamma(\omega)$ has the property that either $\omega_x = 0$ and all the neighbors are 1 (by definition of a contour), or $\omega_x = 1$ and all the neighbors are 0 (as the model contains the isolation-constraint).

Further note that not all possible subsets of $\mathbb{Z}^d$ can occur as Γ, because of the isolation constraint of ones.

The connected (in the sense of graph-distance) components γ of these sets Γ are called *contours*. To visualize this, consider a star-shaped contour that is built from flipping the one site from one to zero starting from a checkerboard configuration, see Figure 5. This yields the minimal contour which has $2d + 1$ sites. In two dimensions, e.g., it is possible that different contours can be reached from each other via the diagonal. Note that each γ that is a contour of a configuration, must be surrounded by ones in nearest neighbor sense in the configuration. These ones must be surrounded by nearest neighbors which carry zeros, by the isolation constraint of ones. Hence the contour specifies the configuration up to sites with graph distance two.

The complement of a finite contour γ has one infinite component, and finitely many finite components (the internal ones). Each of these components are labelled by one of the two labels 1 (or 0 respectively) determined whether, given γ, the component admits a configuration obtained by substituting the infinite-volume checkerboard configurations in which the origin in $\mathbb{Z}^d$ obtains a 1 (or a 0 respectively), see Figure 4. A configuration is uniquely determined by its set of compatible contours. Contours γ

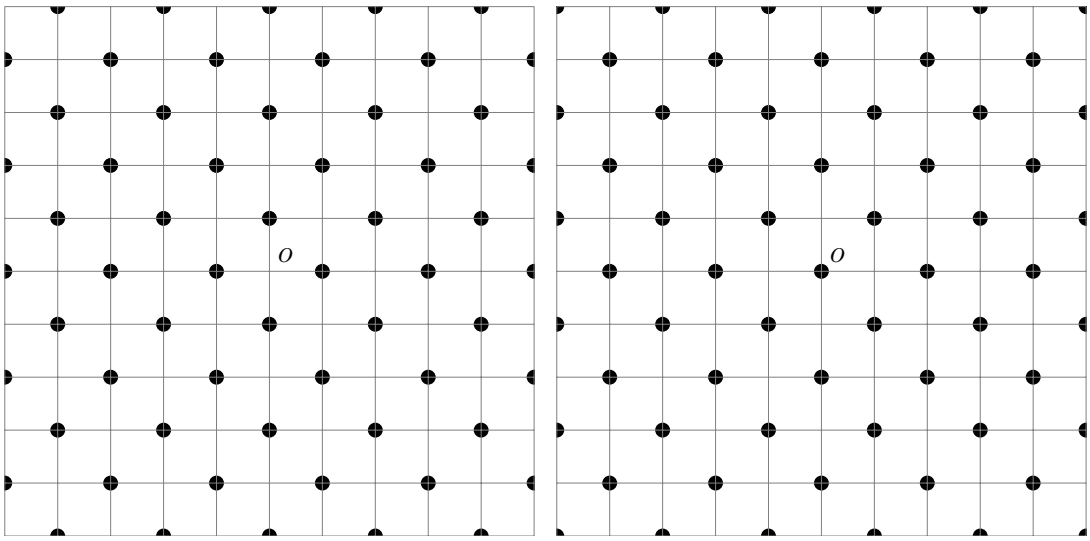

Figure 4: Illustration of the type-0 (left) and type-1 (right) groundstates, i.e., checkerboard configurations. Dots indicate occupied sites. The origin is indicated as o.

themselves are labelled by 1 (or 0) according to the two possible labels of their outer connected components. Contours are compatible when they arise from an allowed configuration, which is the case when the types of checkerboards on shared connected components of the complements match.

Now, suppose that γ is a contour in a finite volume Λ. We decompose the volume in terms of the

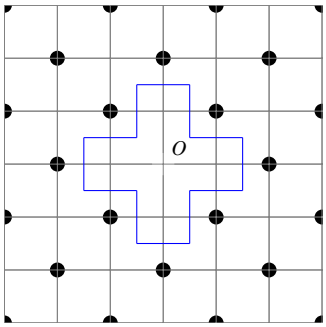

Figure 5: Illustration of a minimal contour in blue containing the origin.

contour and the connected components of its complement

$$\Lambda = \gamma \cup V_0 \cup_{i=1,\dots,k} V_i.$$

Here V_0 is the outer connected component of the complement of γ intersected with Λ. (Note that the intersection with Λ may produce several connected components, but this poses no difficulty.) The sets $(V_i)_{i\geq 1}$ are the interior connected components of the complement of γ. We also write

$$\overline{\gamma} = \gamma \cup_{i=1,\dots,k} V_i$$

for the sites that are contained in the support of γ or surrounded by γ.

We now prove the statements (3.1) and (3.2) via a Peierls estimate. It suffices to treat the first case (3.1), as the case (3.2) is similar. We start with a union bound over contours surrounding the origin

$$\nu_{B_L}(\omega_0 = 1) \leq \nu_{B_L}(\omega \colon \exists \gamma \text{ such that } \Gamma(\omega) \ni \gamma \text{ and } \overline{\gamma} \ni 0) \leq \sum_{\gamma \colon \overline{\gamma} \ni 0} \nu_{B_L}(\omega \colon \Gamma(\omega) \ni \gamma).$$

With a slight abuse of notation, we here write $\gamma \in \Gamma$ to indicate that γ is a contour in Γ.

The main point for the Peierls estimate in our non-flip invariant situation, which is formulated in the following lemma, will be the construction of compatible configurations after removal of contours. Let $|\gamma|$ denote the number of vertices in γ.

Lemma 3.1. There exists a Peierls constant $\tau = \tau(p)$ with $\lim_{p\uparrow 1} \tau(p) = \infty$, such that for all γ we have that

$$\nu_{B_L}(\omega \colon \Gamma(\omega) \ni \gamma) \leq e^{-\tau|\gamma|}.$$

Proof. Define the activity of a contour $\gamma \subset B_L$ to be the natural weight of the zeros prescribed by it in the Bernoulli measure, i.e.,

$$\rho(\gamma) := (1-p)^{|\gamma|}.$$

For any configuration ω_U in a finite volume U we write for its weight in the Bernoulli field

$$R(\omega_U) := \prod_{x \in U} p^{\omega_x}(1-p)^{1-\omega_x}.$$

In particular $\rho(\gamma) = R(0_\gamma)$, where we use the short-hand notation 0_B to indicate the all-zero configuration in the volume B. Then, we may write

$$\nu_{B_L}(\omega : \Gamma(\omega) \ni \gamma) = \frac{\rho(\gamma) Z_{V_0} \prod_{i=1}^k Z_{V_i}}{Z_{B_L}},$$

where Z_{V_i}, for $i = 0, 1, \ldots, k$ denotes the partition functions over all configurations in the volumes V_i which are compatible with γ, with the isolation constraint on the ones, with weights provided by the Bernoulli measure for all sites in V_i.

Case 1. Suppose that γ only contains interior connected components of type 0. This means that there is no typechange when going from the outside to the inside. Then we may remove γ, i.e., continue the checkerboard configuration outside of γ to where γ used to be. This means that we will assign to each ω for which $\Gamma(\omega) \ni \gamma$ the reference configuration $(\omega_{\Lambda\setminus\gamma}\omega^0_\gamma)$ which appears in the partition function Z_Λ to lower bound the latter. Here ω^0_γ denotes the type-0 checkerboard configuration on γ. It is important to note that this removal keeps all other contributions from exterior and interior components compatible. Hence, we immediately arrive at a lower bound

$$Z_{B_L} \geq R(\omega^0_\gamma) Z_{V_0} \prod_{i=1}^{k} Z_{V_i}.$$

We may write $\rho(\gamma) = R(\omega^0_\gamma)((1-p)/p)^{N^{\text{repl}}}$, where N^{repl} denotes the number of replacements of a zero by a one on the support of the contour. By definition of a contour each site in γ has a neighbor which is zero. Therefore it will replaced itself by a one, or a neighbor of it will be replaced by a one. Hence $N^{\text{repl}} \geq |\gamma|/(2d+1)$. Thus, we have the desired estimate

$$\nu_{B_L}(\omega : \Gamma(\omega) \ni \gamma) \leq \big((1-p)/p\big)^{|\gamma|/(2d+1)}.$$

Case 2. Suppose now that γ additionally also contains interior volumes of type 1 (the bad type that does not agree with the boundary outside Λ), which we will denote by W_j, $j = 1, \ldots, l$. Writing V_i for the interior components of type-0, we have

$$\nu_{B_L}(\omega : \Gamma(\omega) \ni \gamma) = \frac{\rho(\gamma) Z_{V_0} \prod_{i=l+1}^{k} Z_{V_i} \prod_{j=1}^{l} Z_{W_j}}{Z_{B_L}},$$

where all partition functions are sums over compatible configurations in the respective connected components, such that the total configuration contains the contour γ.

The difficulty of this case is that the removal of γ does not immediately create compatible configurations. However, it does so after the shift of each of the internal volumes of wrong checkerboard subtypes W_j in one of the $2d$ possible (positive or negative) lattice directions e. Let us explain the details now. Our comparison configuration will now be equal to the type-0 checkerboard on the following set γ_e which describes the appropriate modification of γ, obtained by shifts of the internal components,

$$\gamma_e := \Big(\gamma \setminus \bigcup_{j=1}^{l} (W_j + e)\Big) \cup \bigcup_{j=1}^{l} W_j \setminus (W_j + e). \tag{3.3}$$

Note that there is the volume preservation $|\gamma_e| = |\gamma|$. For each ω for which $\Gamma(\omega) \ni \gamma$, the reference configuration will then be

$$(\omega^0_{\gamma_e}, \omega_{\cup_{i=l+1}^{k} V_i}, (\theta_e \omega)_{\cup_{j=1}^{l}(W_j+e)}), \tag{3.4}$$

where θ_e represents the shift of the configuration by e. We see from the definition that ω will not be modified on the external component and the internal components of good type. It will however be

shifted by e on the internal components of bad type, and it will be a checkerboard of good type on the modified contour γ_e. Note that this configuration really occurs in Z_Λ as it satisfies the isolation constraint on the ones. Hence it can be used to lower bound the partition function which gives us

$$Z_\Lambda \geq R(\omega^0_{\gamma_e}) Z_{V_0} \prod_{i=l+1}^{k} Z_{V_i} \prod_{j=1}^{l} Z_{W_j},$$

where we have used the shift invariance for the internal partition functions Z_{W_j} (those with the bad types).

Now, the proof is finished once we show that there is a dimension-dependent constant $c_d > 0$ such that

$$\rho(\gamma) \leq R(\omega^0_{\gamma_e}) \big((1-p)/p\big)^{c_d|\gamma|}. \tag{3.5}$$

We denote by S^0 the occupied sites of ω^0 (the good checkerboard configuration).

The idea is to use the fact that any connected (in graph distance) subset of $\mathbb{Z}^d$ hits at least a positive fraction of S^0 to conclude the inequality

$$|\gamma_e \cap S^0| \geq c_d|\gamma_e| = c_d|\gamma|. \tag{3.6}$$

Then, the Inequality (3.5) would follow immediately from that.

However, there is the small problem with that argument since, while γ by definition is connected (w.r.t. graph distance), the modified set γ_e may have obtained isolated sites, see Figure 6. This may

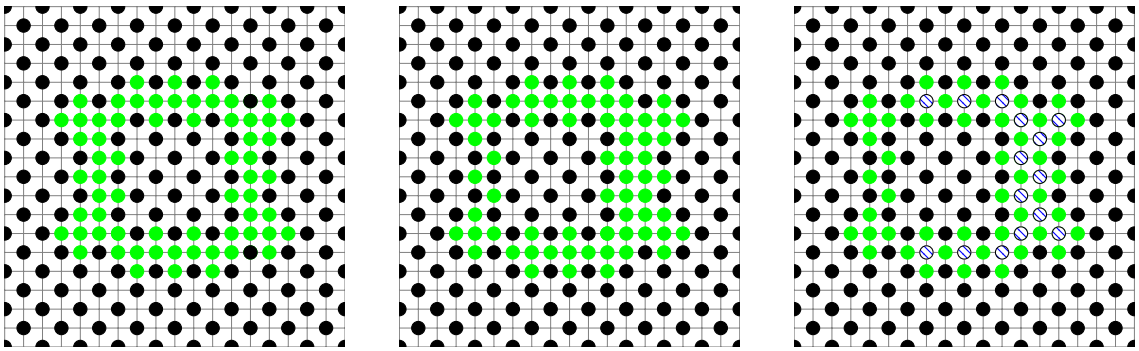

Figure 6: Illustration of a configuration with one contour γ (left in green), where the outside configuration is of type 0 and the inside configuration is of (bad) type 1. Moving the inside configuration by $-e_1$ (middle) creates a (good) configuration also inside the contour γ_e as described in (3.3) (middle in green), however the shift also creates isolated zeros. On the right, in the large connected component of γ_e, sites are indicated in dashed blue, which can be flipped from unoccupied to occupied, as in the configuration presented in (3.4), and therefore create an energetically preferable configuration.

occur if the only neighbor of such a site is eaten up by a translated interior component. We will now explain that this is not a real problem as the number of those unwanted sites is too small to spoil our desired bound (3.6).

Indeed, consider the case $\mathbb{Z}^2$ and $e = -e_1$ for visualization. Then, in each fixed row, the sets γ and γ_e contain the same number of sites. To each loss site (on the left) there corresponds a site in the set $\bigcup_{k=1}^{l} W_l \backslash (W_l + e)$ (which is added to the contour) and which is not isolated but has a nearest neighbor in γ. From this it follows that γ_e still has a nearest-neighbor-connected component $\tilde{\gamma}_e$ that is

at least of size $|\gamma|/2$. Now, already $\tilde{\gamma}_e$ being connected, it hits a fraction (which we call $\tilde{c}_d > 0$) of S^0 and we arrive at

$$|\gamma_e \cap S^0| \geq |\tilde{\gamma}_e \cap S^0| \geq \tilde{c}_d |\tilde{\gamma}_e| \geq c_d |\gamma|, \tag{3.7}$$

where $c_d = \tilde{c}_d/2$. Note that we may regard $\tilde{\gamma}_e$, as obtained in this process, as a contour, and its removal in the line of Case 1.

This proves the claim in the general form and hence proves the Peierls estimate. □

3.1.2 Non-removable discontinuities for μ'_p via the relation to Bernoulli fields conditioned on isolation

Proof of Theorem 2.1. Suppose that γ' is a any specification for μ'_p. Consider a square Q of sidelength 3 around the origin. It will play the role of an observation window. Recall that we write ω' for the second-layer variables, and we write ω for the first-layer variables. We will draw the assumption that

$$\omega'_{Q^c} \mapsto \gamma'_Q(\mathrm{d}\omega'_Q|\omega'_{Q^c})$$

is continuous at $\omega'_{Q^c} = 0'_{Q^c}$ to a contradiction, by exhibiting a finite jump size. We use the notation $1'_B$ for the all-one configuration in the volume B. A single-site observation window would not show the phenomenon, as a conditioning $0'_{0^c}$ for the model conditional on non-isolation forces the origin to be $0'$ due to the non-isolation constraint.

For the purpose of showing the persistence of jumps on Q, we consider the loophole volumes B_L and $B_L + e$ as above, and choose for each L cubes C_L that contain both B_L and $B_L + e$ with a layer of sufficiently large finite thickness. Thickness two will do.

In the first step, we consider conditional probabilities of the second-layer measure $\mu'_p = T\mu_p$ of the form

$$\mu'_p(\omega'_Q|\omega'_{C_L\setminus Q}).$$

We want to show that they are essentially discontinuous at the empty conditioning, so they cannot come from a quasilocal specification γ'. The latter argument will be discussed below in the second step, let us now discuss how to obtain the essential discontinuity. There is a slightly tricky part, as we need to go to higher volumes than single-sites, and need to take care of the constraints very carefully. In order to do consider

$$\mu'_p(\omega'_Q|0'_{B_L\setminus Q}1'_{C_L\setminus B_L})$$

where $\omega'_Q \in \{0,1\}^Q$.

It is useful to compare with the empty configuration on the second layer

$$\frac{\mu'_p(\omega'_Q|0'_{B_L\setminus Q}1'_{C_L\setminus B_L})}{\mu'_p(0'_Q|0'_{B_L\setminus Q}1'_{C_L\setminus B_L})} = \frac{\mu'_p(\omega'_Q 0'_{B_L\setminus Q}1'_{C_L\setminus B_L})}{\mu'_p(0'_Q 0'_{B_L\setminus Q}1'_{C_L\setminus B_L})}.$$

The denominators never vanish, as the configuration that appears in the denominator on the right-hand side obeys the non-isolation constraint.

Given the boundary condition $0'_{B_L\setminus Q}$, the non-isolation constraint on the second layer limits the configurations to non-isolated configurations on the cube Q in order to have a non-zero measure.

We will take as a useful local pattern on Q, the following second-layer reference configuration, which is most easily visualized in two dimensions where it looks as follows

$$\omega'^{*}_{Q} = \begin{matrix} 0 & 1 & 0 \\ 1 & 1 & 1 \\ 0 & 1 & 0 \end{matrix}.$$

It is clearly compatible with the second-layer constraint as $\omega'^{*}_{Q}0'_{Q^c}$ contains no isolated ones. In general dimensions d we choose it analogously, namely as the checkerboard groundstate of type-0 but with an additional 1 at the origin, i.e.,

$$(\omega'^{*}_{Q})_i = \begin{cases} \omega^0_i & i \in Q\backslash 0 \\ 1 & i = 0. \end{cases} \tag{3.8}$$

We have the following useful observation. Given ω'^{*}_{Q}, the underlying Bernoulli-field configuration ω_Q from which it appears as T-image, must take the same values on Q, independently of ω'_{Q^c}.

To understand the following steps it will be helpful to make a notational distinction and write $\omega^*_Q = \omega'^{*}_{Q}$ when we refer to the same configuration as a *first-layer* configuration. With this we have

$$\frac{\mu'_p(\omega'^{*}_{Q}0'_{B_L\backslash Q}1'_{C_L\backslash B_L})}{\mu'_p(0'_Q0'_{B_L\backslash Q}1'_{C_L\backslash B_L})} = \frac{\mu_{p,C_L}(T\sigma_{C_L} = \omega'^{*}_{Q}0'_{B_L\backslash Q}1'_{C_L\backslash B_L})}{\mu_{p,C_L}(T\sigma_{C_L} = 0'_Q0'_{B_L\backslash Q}1'_{C_L\backslash B_L})}.$$

Next, we are aiming for a reformulation that involves only quantities of the first-layer model with isolation constraint in the *whole volume* B_L including Q. For this we perform some manipulations. We split the numerator as follows,

$$\mu_{p,C_L}(T\sigma_{C_L} = \omega'^{*}_{Q}0'_{B_L\backslash Q}1'_{C_L\backslash B_L}) = \mu_{p,Q}(\omega^*_Q)\mu_{p,C_L\backslash Q}\Big(T(\omega^*_Q, \sigma_{C_L\backslash Q})\big|_{C_L\backslash Q} = 0'_{B_L\backslash Q}1'_{C_L\backslash B_L}\Big).$$

Now, we change the middle site on Q on the right-hand side from 1 to 0 to obtain a first-layer configuration that obeys the isolation constraint on Q. For this, we first write the simple identity

$$\mu_{p,C_L}(\omega^*_Q) = \tfrac{p}{1-p}\mu_{p,C_L}(\omega^0_Q).$$

It is important to note that we may also replace

$$T_{C_L}(\omega^*_Q, \sigma_{C_L\backslash Q})\big|_{C_L\backslash Q} = T_{C_L}(\omega^0_Q, \sigma_{C_L\backslash Q})\big|_{C_L\backslash Q},$$

which is possible as the middle site of Q has no influence on the values of the restriction of $T_{C_L}(\omega^0_Q, \sigma_{C_L\backslash Q})$ to Q^c. So we arrive at

$$\mu_{p,C_L}(T\sigma_{C_L} = \omega'^{*}_{Q}0'_{B_L\backslash Q}1'_{C_L\backslash B_L}) = \frac{p}{1-p}\mu_{p,C_L}(\sigma_Q = \omega^0_Q, T_{C_L}(\omega^0_Q, \sigma_{C_L\backslash Q})\big|_{C_L\backslash Q} = 0'_{B_L\backslash Q}1'_{C_L\backslash B_L}).$$

Now we have achieved our goal, as we may recognize that

$$\frac{\mu_{p,C_L}(\sigma_Q = \omega^0_Q, T_{C_L}(\omega^0_Q, \sigma_{C_L\backslash Q})\big|_{C_L\backslash Q} = 0'_{B_L\backslash Q}1'_{C_L\backslash B_L})}{\mu_{p,C_L}(T\sigma_{C_L} = 0'_Q0'_{B_L\backslash Q}1'_{C_L\backslash B_L})} = \nu_{B_L}(\sigma_Q = \omega^0_Q),$$

with the conditional first-layer measure ν_{B_L} that is conditioned on isolation, in the whole volume B_L, as defined in the beginning of Section 3.1.1. Indeed, the denominator on the l.h.s. is the partition function of the conditional measure, up to the terms on $C_L\backslash B_L$ on which the first-layer configuration is frozen, and which cancel against those in the numerator. In particular there is no C_L-dependence. We have thus proved the following *unfixing lemma.*

Lemma 3.2. The second-layer conditional probabilities and the first-layer model under the non-isolation constraint satisfy the following relation

$$\frac{\mu'_p(\sigma'_Q = \omega'^*_Q|0'_{B_L\setminus Q}1'_{C_L\setminus B_L})}{\mu'_p(\sigma'_Q = 0'_Q|0'_{B_L\setminus Q}1'_{C_L\setminus B_L})} = \frac{p}{1-p}\nu_{B_L}(\sigma_Q = \omega^0_Q),$$

where ω^0_Q denotes the checkerboard groundstate and ω'^*_Q is defined in (3.8). The same relation holds for the shifted volume $B_L + e$.

Note that this representation is particularly nice, as we are reduced to the discussion of the first-layer model in the full volume B_L (and not a volume reduced by Q) where we have the Peierls estimate to our disposition. In particular, by the Peierls estimate we have that

$$\nu_{B_L}(\sigma_Q = \omega^0_Q) \geq 1 - \sum_{j\in Q}\mu_{B_L}(\sigma_j \neq \omega^0_j) \geq 1 - |Q|\epsilon(p),$$

while

$$\nu_{B_L+e}(\sigma_Q = \omega^0_Q) \leq \epsilon(p).$$

In the final step, we bring the arbitrarily chosen specification γ' into play, with the aim to show that it must inherit a discontinuity at the empty configuration, too. For any pattern $\omega'_{\mathbb{Q}}$ in the observation window Q, we bound the infimum over perturbations of the empty configuration outside the volume $\Delta_L := B_L \cap B_{L+e}$ via

$$\begin{aligned}\mu'_p(\omega'_Q|0'_{B_L\setminus Q}1'_{C_L\setminus B_L}) &= \int \gamma'_Q(\omega'_Q|0'_{B_L\setminus Q}1'_{C_L\setminus B_L}\tilde{\omega}'_{C_L^c})\ \mu'_p(\mathrm{d}\tilde{\omega}'_{C_L^c}|\omega'_{C_L\setminus Q} = 0'_{B_L\setminus Q}1'_{C_L\setminus B_L})\\ &\geq \inf_{\omega'_{\Delta_L^c}} \gamma'_Q(\omega'_Q|0'_{\Delta_L\setminus Q}\omega'_{\Delta_L^c}) =: a_L(\omega'_Q).\end{aligned}$$

Similar arguments give that

$$\begin{aligned}\mu'_p(\omega'_Q|0'_{B_L+e\setminus Q}1'_{C_L\setminus B_L+e}) &\geq a_L(\omega'_Q)\\ \mu'_p(\omega'_Q|0'_{B_L\setminus Q}1'_{C_L\setminus B_L}) &\leq \sup_{\omega'_{\Delta_L^c}} \gamma_Q(\omega'_Q|0'_{\Delta_L\setminus Q}\omega'_{\Delta_L^c}) =: b_L(\omega'_Q)\\ \mu'_p(\omega'_Q|0'_{B_L+e\setminus Q}1'_{C_L\setminus B_L+e}) &\leq b_L(\omega'_Q).\end{aligned}$$

Now consider specifically the patterns $0'_Q$ and ω'^*_Q and note that

$$\frac{p}{1-p}\nu_{B_L}(\sigma_Q = \omega^0_Q) = \frac{\mu'_p(\omega'^*_Q|0'_{B_L\setminus Q}1'_{C_L\setminus B_L})}{\mu'_p(0'_Q|0'_{B_L\setminus Q}1'_{C_L\setminus B_L})} \leq \frac{b_L(\omega'^*_Q)}{a_L(0_Q)} \tag{3.9}$$

and

$$\frac{p}{1-p}\nu_{B_L+e}(\sigma_Q = \omega^0_Q) = \frac{\mu'_p(\omega'^*_Q|0'_{B_L+e\setminus Q}1'_{C_L\setminus B_L+e})}{\mu'_p(0'_Q|0'_{B_L+e\setminus Q}1'_{C_L\setminus B_L+e})} \geq \frac{a_L(\omega'^*_Q)}{b_L(0_Q)}, \tag{3.10}$$

and remark that the denominators are uniformly bounded against zero.

Now, by the Peierls estimate presented in Lemma 3.1 in previous section, we have lower bounds on the left-hand side of (3.9) and upper bounds on the left-hand side of (3.10). These contradict the continuity assumption on the specification γ', i.e., that the right-hand sides have the same limit as $L \uparrow \infty$.

This proves the discontinuity of the specification kernel γ'_Q for any arbitrary specification γ', at the fully empty configuration. □

3.2 Proof of Theorem 2.2: Gibbsianness

In this section, we construct a continuous specification γ' for μ'_p for small p. The main ingredient is an application of the Dobrushin-uniqueness bound and the backward-martingale theorem. In the first step, we construct the conditional probabilities in finite volumes.

For this we use the following notation. For $\Lambda \subset \mathbb{Z}^d$ we denote by $\Lambda^c := \mathbb{Z}^d \setminus \Lambda$ its *complement* and by $\partial_- \Lambda := \{x \in \Lambda\colon \text{there exists } y \in \Lambda^c \text{ with } y \sim_{\mathbb{Z}^d} x\}$ its *interior boundary*. The set $\Lambda^o := \Lambda \setminus \partial_- \Lambda$ denotes the *interior* and $\bar{\Lambda} := ((\Lambda^c)^o)^c$ the *extension* of Λ. Moreover, $\partial_+ \Lambda := \bar{\Lambda} \setminus \Lambda$ then denotes the *outer boundary* of Λ.

3.2.1 The specification

Let us consider, for any (large) finite volume $\Delta \Subset \mathbb{Z}^d$, conditional probabilities of μ'_p inside Δ also given a (first-layer) boundary condition ω outside Δ. More precisely, let $\Lambda \subset \Delta$, $\omega' = \omega'_\Lambda \omega'_{\Delta\setminus\Lambda} \omega'_{\Delta^c} \in \Omega'$ and let ω be such that $\omega_{\Delta^c} \in T^{-1}(\omega'_{(\Delta^o)^c})$, then

$$\begin{aligned}\gamma'_{\omega_{\Delta^c},\Lambda}(\omega'_\Lambda|\omega'_{\Delta\setminus\Lambda}) &:= \frac{\sum_{\tilde{\omega}_\Delta} \mu_p(\tilde{\omega}_\Delta) \mathbb{1}\{T_\Delta(\tilde{\omega}_\Delta \omega_{\Delta^c}) = \omega'_\Delta\}}{\sum_{\tilde{\omega}_{\Delta\setminus\Lambda^o}} \mu_p(\tilde{\omega}_{\Delta\setminus\Lambda^o}) \mathbb{1}\{T_{\Delta\setminus\Lambda}(\tilde{\omega}_{\Delta\setminus\Lambda^o}\omega_{\Delta^c}) = \omega'_{\Delta\setminus\Lambda}\}} \\ &= \frac{\sum_{\tilde{\omega}_{\Delta\setminus\Lambda^o}} \mu_p(\tilde{\omega}_{\Delta\setminus\Lambda^o}) \mathbb{1}\{T_{\Delta\setminus\Lambda}(\tilde{\omega}_{\Delta\setminus\Lambda^o}\omega_{\Delta^c}) = \omega'_{\Delta\setminus\Lambda}\} \sum_{\tilde{\omega}_{\Lambda^o}} \mu_p(\tilde{\omega}_{\Lambda^o}) \mathbb{1}\{T_\Lambda(\tilde{\omega}_{\Lambda^o}\tilde{\omega}_{\Delta\setminus\Lambda^o}) = \omega'_\Lambda\}}{\sum_{\tilde{\omega}_{\Delta\setminus\Lambda^o}} \mu_p(\tilde{\omega}_{\Delta\setminus\Lambda^o}) \mathbb{1}\{T_{\Delta\setminus\Lambda}(\tilde{\omega}_{\Delta\setminus\Lambda^o}\omega_{\Delta^c}) = \omega'_{\Delta\setminus\Lambda}\}} \\ &= \frac{\sum_{\tilde{\omega}_{\Delta\setminus\Lambda^o}} \mu_p(\tilde{\omega}_{\Delta\setminus\Lambda^o}) \mathbb{1}\{T_{\Delta\setminus\Lambda}(\tilde{\omega}_{\Delta\setminus\Lambda^o}\omega_{\Delta^c}) = \omega'_{\Delta\setminus\Lambda}\} f_{\omega'_\Lambda}(\tilde{\omega}_{\partial_-\Lambda\cup\partial_+\Lambda})}{\sum_{\tilde{\omega}_{\Delta\setminus\Lambda^o}} \mu_p(\tilde{\omega}_{\Delta\setminus\Lambda^o}) \mathbb{1}\{T_{\Delta\setminus\Lambda}(\tilde{\omega}_{\Delta\setminus\Lambda^o}\omega_{\Delta^c}) = \omega'_{\Delta\setminus\Lambda}\}},\end{aligned} \tag{3.11}$$

where we wrote $T_\Lambda(\omega)$ instead of $(T(\omega))_\Lambda$ and

$$f_{\omega'_\Lambda}(\omega_{\partial_-\Lambda\cup\partial_+\Lambda}) := \sum_{\tilde{\omega}_{\Lambda^o}} \mu_p(\tilde{\omega}_{\Lambda^o}) \mathbb{1}\{T_\Lambda(\tilde{\omega}_{\Lambda^o}\omega_{\partial_-\Lambda\cup\partial_+\Lambda}) = \omega'_\Lambda\}$$

is a local function. We have the following consistency result.

Lemma 3.3. Assume that, given $\Lambda \Subset \mathbb{Z}^d$ and $\omega' \in \Omega'$, $\lim_{\Delta\uparrow\mathbb{Z}^d} \gamma'_{\omega_{\Delta^c},\Delta}(\omega'_\Lambda|\omega'_{\Delta\setminus\Lambda}) =: \gamma'_\Lambda(\omega'_\Lambda|\omega'_{\Lambda^c})$ exists and is independent of $\omega_{\Delta^c} \in T^{-1}(\omega'_{(\Delta^o)^c})$. Then, γ' is a specification for μ'_p.

Proof. First note that for any $\omega' \in \Omega'$ we can estimate,

$$\big|\mu'_p(\omega'_\Lambda|\omega'_{\Delta\setminus\Lambda}) - \gamma'_\Lambda(\omega'_\Lambda|\omega'_{\Lambda^c})\big| \le \sup_{\omega_{\partial_+\Delta}} \big|\gamma'_{\omega_{\partial_+\Delta},\Delta}(\omega'_\Lambda|\omega'_{\Delta\setminus\Lambda}) - \gamma'_\Lambda(\omega'_\Lambda|\omega'_{\Lambda^c})\big|, \tag{3.12}$$

where the supremum is taken over suitable boundary configurations compatible with ω', since $\mu'_p(\omega'_\Lambda|\omega'_{\Delta\setminus\Lambda})$ can be written as an integral with respect to $\gamma'_{\cdot,\Delta}(\omega'_\Lambda|\omega'_{\Delta\setminus\Lambda})$. In particular, under our assumptions, the right-hand side of Equation 3.12 tends to zero as Δ tends to $\mathbb{Z}^d$.

Now, consider a cofinal sequence $\Delta_n \uparrow \Lambda^c$ and let $(\mathcal{F}'_{\Delta_n})_{n\in\mathbb{N}}$ denote the canonical filtration on Ω'. Note that the sequence of random variables $\mu'_p(\omega'_\Lambda|\mathcal{F}'_{\Delta_n})$ that are μ'_p-almost surely defined as $\mu'_p(\omega'_\Lambda|\mathcal{F}'_{\Delta_n})(\omega') = \mu'_p(\omega'_\Lambda|\omega'_{\Delta_n\setminus\Lambda})$, is a uniformly-integrable martingale adapted to $(\mathcal{F}'_{\Delta_n})_{n\in\mathbb{N}}$. Since $\sigma\left(\bigcup_n \mathcal{F}'_{\Delta_n}\right) = \mathcal{F}'_{\Lambda^c}$, by Lévy's zero-one law, $(\mu'_p(\omega'_\Lambda|\mathcal{F}'_{\Delta_n}))_{n\in\mathbb{N}}$ converges μ'_p-almost surely and in

L^1 towards $\mu'_p(\omega'_\Lambda|\mathcal{F}'_{\Lambda^c})$ as n tends to infinity. But this implies that for any $\tilde{\omega}'_\Lambda$, we can pick n sufficiently large such that

$$\begin{aligned}&\int \mu'_p(\omega')|\mu'_p(\tilde{\omega}'_\Lambda|\omega'_{\Lambda^c}) - \gamma'_\Lambda(\tilde{\omega}'_\Lambda|\omega'_{\Lambda^c})| \\ &\leq \int \mu'_p(\omega')|\mu'_p(\tilde{\omega}'_\Lambda|\omega'_{\Lambda^c}) - \mu'_p(\tilde{\omega}'_\Lambda|\omega'_{\Delta_n\setminus\Lambda})| + \int \mu'_p(\omega')|\mu'_p(\tilde{\omega}'_\Lambda|\omega'_{\Delta_n\setminus\Lambda}) - \gamma'_\Lambda(\tilde{\omega}'_\Lambda|\omega'_{\Lambda^c})| < \varepsilon,\end{aligned}$$

where we used Lévy's zero-one law in the first summand and the bound (3.12) in the second summand on the right-hand side. But this implies that $\int \mu'_p(\omega'_{\Lambda^c})\gamma'_\Lambda(\tilde{\omega}'_\Lambda|\omega'_{\Lambda^c}) = \mu'_p(\tilde{\omega}'_\Lambda)$ and hence γ' is a specification for μ'_p. □

3.2.2 Transformations into first-layer constraint models

In order to establish the conditions of Lemma 3.3 for sufficiently small p, we employ the Dobrushin uniqueness theorem for the first-layer constraint model as defined in (3.13). For this, first note that we can uniquely identify ω' with the subset of its occupied sites in $\mathbb{Z}^d$ and with some notational abuse $\bar{\omega}' \subset \mathbb{Z}^d$ of ω' is a *fixed area* in the sense that, under T, there is no choice for the Bernoulli field in how to realize ω'. Recall that ω' consist of clusters of size at least two and $\bar{\omega}'$ then consists of clusters of size at least two surrounded by unoccupied sites, see Figure 7 for an illustration.

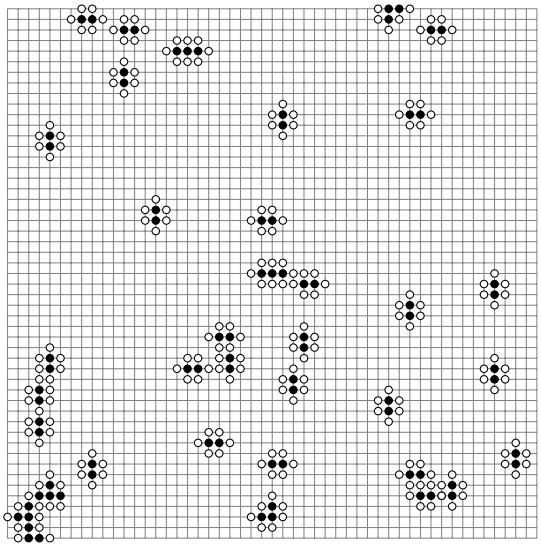

Figure 7: Illustration of the fixed area (black and white dots) based on a thinned configuration (black dots). The thinned configuration is surrounded by unoccupied sites (white dots).

In view of this, we introduce the following specification associated to the *first-layer constraint model* on Ω

$$\gamma^S_\Delta(\omega_\Delta|\omega_{\Delta^c}) := \frac{\mu_p(\omega_{\Delta\cap S})\mathbb{1}\{\omega_{\Delta\cap S}\omega_{\Delta^c\cap S} \text{ is } T\text{-feasible on } \Delta\cap S\}}{\sum_{\tilde{\omega}_{\Delta\cap S}} \mu_p(\tilde{\omega}_{\Delta\cap S})\mathbb{1}\{\tilde{\omega}_{\Delta\cap S}\omega_{\Delta^c\cap S} \text{ is } T\text{-feasible on } \Delta\cap S\}}. \tag{3.13}$$

Here, $S \subset \mathbb{Z}^d$ is an *unfixed area* that is arbitrary at this stage, $\Delta \Subset \mathbb{Z}^d$ and any configuration $\omega \in \Omega$ is called T*-feasible* on a set $\Delta \cap S$ if all occupied sites of ω in $\Delta \cap S$ have no neighboring occupied sites in $\bar{\Delta} \cap S$. In particular, with this definition,

$$\gamma'_{\omega_{\Delta^c},\Lambda}(\omega'_\Lambda|\omega'_{\Delta\setminus\Lambda}) = \gamma^S_\Delta(f_{\omega'_\Lambda}|\omega_{\Delta^c})$$

for the particular choice of the unfixed area given by $S = S(\omega'_{\Delta\setminus\Lambda}) = (\Delta \setminus \Lambda^o) \setminus \bar{\omega}'_{\Delta\setminus\Lambda}$. Here we used that, in the fixed area $\bar{\omega}'_{\Delta\setminus\Lambda}$, the Bernoulli field is completely determined by $\omega'_{\Delta\setminus\Lambda}$ and hence the corresponding factor cancels in (3.11). The following result verifies the conditions of Lemma 3.3 for sufficiently small p.

Lemma 3.4. Let $p < 1/(2d)$. Then, for any $\Lambda \Subset \mathbb{Z}^d$ and $\omega' \in \Omega'$, the limit $\lim_{\Delta\uparrow\mathbb{Z}^d} \gamma_\Delta^{S(\omega'_{\Delta\setminus\Lambda})}(f_{\omega'_\Lambda}|\omega_{\Delta^c})$ exists and is independent of $\omega_{\Delta^c} \in T^{-1}(\omega'_{(\Delta^o)^c})$.

Proof. We use the Dobrushin-uniqueness approach for the specification γ_Δ^S as defined in (3.13), where $S \subset \mathbb{Z}^d$ is any unfixed area. Consider the Dobrushin matrix

$$C_{ij}(p) = \max_{\omega_{j^c}=\tilde{\omega}_{j^c}} \|\gamma_i^S(\cdot|\omega_{i^c}) - \gamma_i^S(\cdot|\tilde{\omega}_{i^c})\|_{\mathsf{TV}}$$

for $i, j \in S$, where complements are defined in S. Note that the exterior boundary $\partial_+ S$ of S consists of unoccupied sites, see Figure 7. We have $C_{ij}(p) = 0$ unless i and j are neighbors in S. Otherwise,

$$C_{ij}(p) = \tfrac{1}{2} \max_{\omega_{j^c}=\tilde{\omega}_{j^c}} (|\gamma_i^S(0|\omega_{i^c}) - \gamma_i^S(0|\tilde{\omega}_{i^c})| + |\gamma_i^S(1|\omega_{i^c}) - \gamma_i^S(1|\tilde{\omega}_{i^c})|) = p,$$

where the maximum is realized when ω_{j^c} is unoccupied while ω_j is unoccupied and $\tilde{\omega}_j$ is occupied. In particular, for the Dobrushin criterion, we have

$$c(p) = \sup_{i\in S} \sum_{j\sim i} C_{ij}(p) \leq 2dp,$$

independent of S. By [Geo11, Theorem 8.7], for all $p < 1/(2d)$ and S, γ_Δ^S admits a unique infinite-volume Gibbs measure μ^S. Finally, using the remark made above [Geo11, Equation 8.25], $\gamma_\Delta^S(f_{\omega'_\Lambda}|\omega_{\Delta^c})$ converges uniformly in ω towards $\mu^S(f_{\omega'_\Lambda})$, which finishes the proof. □

3.2.3 Quasilocality of the specification

What remains to be done in order to finish the proof of Theorem 2.2 is to establish quasilocality for the specification. Let s denote the ℓ_∞ metric on $\mathbb{Z}^d$ and define $s(\Lambda, \Delta) = \inf\{s(i, j) \colon i \in \Lambda, j \in \Delta\}$.

Lemma 3.5. For $p < 1/(2d)$ there exist constants $C, c > 0$ such that for all $\Lambda \subset \Delta \Subset \mathbb{Z}^d$ and all configurations ω' and $\tilde{\omega}'$ with $\omega'_\Delta = \tilde{\omega}'_\Delta$ we have that

$$|\gamma'_\Lambda(\omega'_\Lambda|\omega'_{\Lambda^c}) - \gamma'_\Lambda(\omega'_\Lambda|\tilde{\omega}'_{\Lambda^c})| \leq C|\Lambda|e^{-cs(\Lambda,\Delta^c)}.$$

In particular, the specification γ' is quasilocal.

Proof. We use the representation of $\gamma'_\Lambda(\omega'_\Lambda|\omega'_{\Lambda^c})$ in terms of the unique infinite-volume Gibbs measure $\mu^S(f_{\omega'_\Lambda})$ as presented in the proof of Lemma 3.4. This representation exists since we work in the Dobrushin-uniqueness regime. Now, for the quasilocality, we use the criterion [Geo11, Remark 8.26] applied to [Geo11, Theorem 8.20]. More precisely, since $p < 1/(2d)$, by [Geo11, Theorem 8.20], for $S \cap \Delta = S' \cap \Delta$, we have that

$$|\mu^S(f_{\omega'_\Lambda}) - \mu^{S'}(f_{\omega'_\Lambda})| \leq D(\Lambda, \Delta),$$

where $D(\Lambda, \Delta) = \sum_{i\in\Lambda, j\in\Delta^c} \big(\sum_{n\geq 0} C^n\big)_{i,j}$ with $C^n = C^n(p)$ the n-th power of the Dobrushin matrix as presented in the proof of Lemma 3.4. Now choose $c > 0$ sufficiently small such that $p\mathrm{e}^c < 1/(2d)$, then, by [Geo11, Remark 8.26],

$$D(\Lambda, \Delta) \leq |\Lambda|(1 - 2dp\mathrm{e}^c)^{-1}\mathrm{e}^{-cd(\Lambda,\Delta^c)}.$$

This finishes the proof. □

References

[Adl91] J. Adler. Bootstrap percolation. *Physica A: Statistical Mechanics and its Applications*, 171(3):453–470, 1991.

[Bal05] K. Ball. Poisson thinning by monotone factors. *Electronic Communications in Probability*, 10:60–69, 2005.

[BB12] F. Baccelli and P. Bermolen. Extremal versus additive Matérn point processes. *Queueing Systems*, 71(1):179–197, 2012.

[BC18] S. Bethuelsen and D. Conache. One-sided continuity properties for the Schonmann projection. *Journal of Statistical Physics*, 172(4):1147–1163, 2018.

[BDC13] V. Beffara and H. Duminil-Copin. Planar percolation with a glimpse of Schramm–Loewner evolution. *Probability Surveys*, 10:1–50, 2013.

[BK19] B. Blaszczyszyn and H. Keeler. Determinantal thinning of point processes with network learning applications. In *2019 IEEE Wireless Communications and Networking Conference (WCNC)*, pages 1–8. IEEE, 2019.

[BKK20] S. Bergmann, S. Kissel, and C. Külske. Dynamical Gibbs–non-Gibbs transitions in Widom–Rowlinson models on trees. *arXiv preprint arXiv:2012.09718*, 2020.

[BKL98] J. Bricmont, A. Kupiainen, and R. Lefevere. Renormalization group pathologies and the definition of Gibbs states. *Communications in Mathematical Physics*, 194(2):359–388, 1998.

[Bov06] A. Bovier. *Statistical Mechanics of Disordered Systems*, volume 18 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, Cambridge, 2006.

[BR06] B. Bollobás and O. Riordan. *Percolation*. Cambridge University Press, New York, 2006.

[Bre79] P. Bremaud. Optimal thinning of a point process. *SIAM Journal on Control and Optimization*, 17(2):222–230, 1979.

[Bro79] T. Brown. Position dependent and stochastic thinning of point processes. *Stochastic Processes and their Applications*, 9(2):189–193, 1979.

[EAM95] R. Elliott, L. Aggoun, and J. Moore. *Hidden Markov Models*, volume 29 of *Applications of Mathematics (New York)*. Springer-Verlag, New York, 1995.

[FdHM14] R. Fernández, F. den Hollander, and J. Martínez. Variational description of Gibbs-non-Gibbs dynamical transitions for spin-flip systems with a Kac-type interaction. *Journal of Statistical Physics*, 156(2):203–220, 2014.

[Gem90] D. Geman. *Random Fields and Inverse Problems in Imaging*, volume 1427 of *Lecture Notes in Math.* Springer, Berlin, 1990.

[Geo11] H.-O. Georgii. *Gibbs Measures and Phase Transitions*, volume 9. Walter de Gruyter & Co., Berlin, second edition, 2011.

[GP78] R. Griffiths and P. Pearce. Position-space renormalization-group transformations: Some proofs and some problems. *Physical Review Letters*, 41(14):917, 1978.

[GP79] R. Griffiths and P. Pearce. Mathematical properties of position-space renormalization-group transformations. *Journal of Statistical Physics*, 20(5):499–545, 1979.

[Gri97] G. Grimmett. Percolation and disordered systems. In *Lectures on Probability Theory and Statistics*, pages 153–300. Springer, 1997.

[Gri99] G. Grimmett. *Percolation*, volume 321 of *Grundlehren der Mathematischen Wissenschaften*. Springer-Verlag, Berlin, second edition, 1999.

[Guy95] X. Guyon. *Random Fields on a Network*. Probability and its Applications (New York). Springer-Verlag, New York, 1995.

[HK04] O. Häggström and C. Külske. Gibbs properties of the fuzzy Potts model on trees and in mean field. *Markov Processes and Related Fields*, 10(3):477–506, 2004.

[HKK19] F. Henning, R. Kraaij, and C. Külske. Gibbs–non-Gibbs transitions in the fuzzy Potts model with a Kac-type interaction: closing the Ising gap. *Bernoulli*, 25(3):2051–2074, 2019.

[Ish80] V. Isham. Dependent thinning of point processes. *Journal of Applied Probability*, 17(4):987–995, 1980.

[JK17a] B. Jahnel and C. Külske. Sharp thresholds for Gibbs–non-Gibbs transitions in the fuzzy Potts model with a Kac-type interaction. *Bernoulli*, 23(4A):2808–2827, 2017.

[JK17b] B. Jahnel and C. Külske. The Widom–Rowlinson model under spin flip: Immediate loss and sharp recovery of quasilocality. *The Annals of Applied Probability*, 27(6):3845–3892, 2017.

[JK20] Benedikt Jahnel and Wolfgang König. *Probabilistic Methods In Telecommunications*. Compact Textbooks in Mathematics. Birkhäuser Basel, 2020.

[KLN07] C. Külske and A. Le Ny. Spin-flip dynamics of the Curie–Weiss model: Loss of Gibbsianness with possibly broken symmetry. *Communications in Mathematical Physics*, 271(2):431–454, 2007.

[KLNR04] C. Külske, A. Le Ny, and F. Redig. Relative entropy and variational properties of generalized Gibbsian measures. *The Annals of Probability*, 32(2):1691–1726, 2004.

[KR17] C. Külske and U. Rozikov. Fuzzy transformations and extremality of Gibbs measures for the Potts model on a Cayley tree. *Random Structures & Algorithms*, 50(4):636–678, 2017.

[KRvZ21] R. Kraaij, F. Redig, and W. van Zuijlen. A Hamilton–Jacobi point of view on mean-field Gibbs–non-Gibbs transitions. *Transactions of the American Mathematical Society*, 2021.

[Kum14] T. Kumagai. *Random Walks on Disordered Media and their Scaling Limits*, volume 2101 of *Lecture Notes in Mathematics*. Springer, Cham, 2014.

[Las93] G. Last. On dependent marking and thinning of point processes. *Stochastic Processes and Their Applications*, 45(1):73–94, 1993.

[LN13] A. Le Ny. Almost Gibbsianness and parsimonious description of the decimated 2d-Ising model. *Journal of Statistical Physics*, 152(2):305–335, 2013.

[LNR02] A. Le Ny and F. Redig. Short time conservation of Gibbsianness under local stochastic evolutions. *Journal of Statistical Physics*, 109(5):1073–1090, 2002.

[Mat60] B. Matérn. Spatial variation stochastic models and their application to some problems in forest surveys and other sampling investigations. *Meddelanden Fran Statens Skogsforsknings Institut*, 49, 1960.

[MHW10] J. Moller, M. Huber, and R. Wolpert. Perfect simulation and moment properties for the Matérn type III process. *Stochastic Processes and Their Applications*, 120(11):2142–2158, 2010.

[MS10] J. Moller and F. Schoenberg. Thinning spatial point processes into poisson processes. *Advances in Applied Probability*, 42(2):347–358, 2010.

[RR14] S. Rœlly and W. Ruszel. Propagation of Gibbsianness for infinite-dimensional diffusions with space-time interaction. *Markov Processes and Related Fields*, 20(4):653–674, 2014.

[RS91] T. Rolski and R. Szekli. Stochastic ordering and thinning of point processes. *Stochastic Processes and Their Applications*, 37(2):299–312, 1991.

[Sch89] R. Schonmann. Projections of Gibbs measures may be non-Gibbsian. *Communications in Mathematical Physics*, 124(1):1–7, 1989.

[vEEIK12] A. van Enter, V. Ermolaev, G. Iacobelli, and C. Külske. Gibbs-non-Gibbs properties for evolving Ising models on trees. *Annales de l'Institut Henri Poincaré Probabilités et Statistiques*, 48(3):774–791, 2012.

[vEFdHR02] A. van Enter, R. Fernández, F. den Hollander, and F. Redig. Possible loss and recovery of Gibbsianness during the stochastic evolution of Gibbs measures. *Communications in Mathematical Physics*, 226(1):101–130, 2002.

[vEFS93] A. van Enter, R. Fernández, and A. Sokal. Regularity properties and pathologies of position-space renormalization-group transformations: scope and limitations of Gibbsian theory. *Journal of Statistical Physics*, 72(5-6):879–1167, 1993.

[vELN17] A. van Enter and A. Le Ny. Decimation of the Dyson–Ising ferromagnet. *Stochastic Processes and their Applications*, 127(11):3776–3791, 2017.

[vEMSS00] A. van Enter, C. Maes, R. Schonmann, and S. Shlosman. The Griffiths singularity random field. In *On Dobrushin's way. From probability theory to statistical physics*, volume 198 of *American Mathematical Society Translations Series 2*, pages 51–58. American Mathematical Society, Providence, RI, 2000.

[WK74] K. Wilson and J. Kogut. The renormalization group and the ε expansion. *Physics Reports*, 12(2):75–199, 1974.