On Absence of disorder chaos for spin glasses on $\mathbb{Z}^d$

Louis-Pierre Arguin*   Jack Hanson [†]

Abstract

We identify simple mechanisms that prevent a strong form of disorder chaos for edge overlap of the Ising spin glass model on $\mathbb{Z}^d$. This was first shown by Chatterjee in the case of Gaussian couplings. We present three proofs of the theorem for general couplings with continuous distribution based on the presence in the coupling realization of stabilizing features of positive density.

1 Introduction

1.1 Main result

Consider a square box Λ in $\mathbb{Z}^d$ with edge set Λ^* and exterior vertex boundary $\partial\Lambda$. The Hamiltonian of the Ising spin glass, or Edwards-Anderson model, is

$$H_{\Lambda,J}(s) = \sum_{\{x,y\} \in (\Lambda \cup \partial\Lambda)^*} -J_{xy} s_x s_y, \quad s \in \{-1,+1\}^\Lambda , \tag{1.1}$$

where the couplings $J = (J_{xy}, \{x,y\} \in (\Lambda \cup \partial\Lambda)^*)$ are IID continuous random variables under some probability $\mathbb{P}$. The distribution of the couplings is usually taken to be symmetric, but this will not be necessary for the proofs. The choice of boundary condition corresponds to setting s on $\partial\Lambda$. This choice will not play a role in the result.

The ground state $\sigma(J)$ at a realization J of the coupling is the minimizer of $H_{\Lambda,J}$:

$$\sigma(J) = \operatorname*{argmin}_{s \in \{-1,+1\}^\Lambda} H_{\Lambda,J}(s) .$$

This implies that the flip of spins in any non-empty subset $\mathcal{B} \subseteq \Lambda$ must increase the energy, yielding the equivalent characterization of the ground state:

$$\sum_{\{x,y\} \in \partial^*\mathcal{B}} J_{xy} \sigma_x(J) \sigma_y(J) > 0 \qquad \forall \mathcal{B} \subseteq \Lambda, \text{ non-empty}, \tag{1.2}$$

where $\{x,y\} \in \partial^*\mathcal{B}$ means $x \in \mathcal{B}$ and $y \notin \mathcal{B}$ (or vice-versa).

*City University of New York, Baruch College and Graduate Center
 E-mail: louis-pierre.arguin@baruch.cuny.edu
†City University of New York, City College and Graduate Center
 E-mail: jhanson@ccny.cuny.edu

In the case where the Hamiltonian admits a global spin symmetry (e.g., with periodic boundary conditions), there is a trivial degeneracy for the ground state. One can then work on spin configurations modulo the spin flip, and speak of ground state pair. Since the arguments presented below are identical in this framework, we will omit the distinction in the notation.

There are also non-trivial degeneracies at some special values of the couplings corresponding to values where one subset has a zero flip energy, i.e., the left-hand side of Equation (1.2) is 0. The set of these critical values is given by

$$\mathcal{C} = \bigcup_{s,s' \in \{-1,+1\}^{\Lambda}, s \neq s'} \left\{ J : \sum_{\{x,y\} \in (\Lambda \cup \partial\Lambda)^*} J_{xy}(s_x s_y - s'_x s'_y) = 0 \right\} . \tag{1.3}$$

The ground state is well defined on the open set $\mathbb{R}^{(\Lambda \cup \partial\Lambda)^*} \setminus \mathcal{C}$. Note that by the continuity of the distribution we have $\mathbb{P}(\mathcal{C}) = 0$.

The phenomenon of disorder chaos in spin glasses was proposed in the physics literature in [14] and in [7]. Roughly speaking, the model exhibits disorder chaos if the ground state at J and the one at a value very close to J differ substantially. The distance between two states can be quantified using the site overlap or the edge overlap (among others). This paper is concerned with the *edge overlap*

$$Q_{\Lambda}(\sigma, \sigma') = \frac{1}{|\Lambda^*|} \sum_{\{x,y\} \in \Lambda^*} \sigma_x \sigma_y \sigma'_x \sigma'_y . \tag{1.4}$$

(The edge overlap is the right measure to distinguish incongruent states, i.e., states separated by a domain wall of positive density.) For the perturbations, we also consider two IID copies $\varepsilon = (\varepsilon_{xy}, \{x,y\} \in (\Lambda \cup \partial\Lambda)^*)$, $\varepsilon' = (\varepsilon'_{xy}, \{x,y\} \in (\Lambda \cup \partial\Lambda)^*)$ of IID continuous random variables on $\mathbb{P}$, indexed by the edges, independent of J and of mean 0. A parameter $t \geq 0$ controls their magnitude. The main result is a proof of absence of disorder chaos in the sense that the average of the overlap between two ground states with slightly different couplings is bounded away from 0 uniformly in Λ.

Theorem 1.1. For any $\delta > 0$, there exists $t_0 = t_0(\delta) > 0$ such that

$$\mathbb{E}[Q_{\Lambda}(\sigma(J + t\varepsilon), \sigma(J + t\varepsilon'))] > (1-\delta)\mathbb{P}(A), \text{ for any } t < t_0(\delta),$$

where A is some explicit event with $\mathbb{P}(A) > 0$ uniformly in Λ.

The theorem was first proved by Chatterjee [8] in the case of Gaussian couplings with an explicit decay in t. Namely, if $J(t)$ and $J'(t)$ are two Ornstein-Uhlenbeck processes both starting at $J(0)$ then evolving independently, he proved for some $c > 0$ that

$$\mathbb{E}[Q_{\Lambda}(\sigma(J(t)), \sigma(J'(t)))] \geq \frac{c}{4d^2} e^{-t/(4d^2 c)} . \tag{1.5}$$

The proof of (1.5) is done in two steps. First, it is shown that the variance of the ground state energy is of the order of $|\Lambda|$. Then, the bound on the overlap follows from a relation between the variance and the overlap that is essentially a consequence of Gaussian integration by parts. The result is to be compared to the Sherrington-Kirkpatrick (SK) model on the complete graph with Gaussian coupling for which the average overlap goes to 0 as $\Lambda \to \mathbb{Z}^d$ for any fixed t [8]. The proof there is given at positive temperature. Chen, Handschy & Lerman proved the zero-temperature case, and a generalization to mixed even p-spin models [10]. Chen & Panchenko also proved a zero-temperature version for diluted spin glass models [12].

Similar results for the SK model with external field were proved with different methods by Chen in [9], and for the spherical version of the model by Chen & Sen in [13]

at zero temperature and by Chen *et al.* at positive temperature [11]. In the presence of external field, the edge overlap concentrates at some value strictly between 0 and 1. So there is a still a form of disorder chaos, though not as strong as the case without magnetic field where the overlap is 0.

The main motivation of the present paper is to pinpoint coupling features arising in finite dimension that stabilize the ground state, even at zero magnetic field. We provide three proofs of Theorem 1.1. The two proofs in Section 2.1 and 2.2 rely on the presence of strong ferromagnetic couplings. The one in Section 2.1 is simpler, but uses the assumption that 0 is in the support. The proof in Section 2.2 relies on no other assumption than the continuity of the distribution. One could hope that these features are abundant enough to prove a stronger result described in Conjecture 1.2 below.

Section 3 presents a different approach based on controlling the influence on the ground state of the coupling at a given edge. More precisely, we look at the critical droplet at an edge e, i.e., the set of vertices that flips when a coupling at e is sent to either $+\infty$ or $-\infty$. It is known, see for example [4], that the size of the critical droplet is intimately related to the number of ground states in the infinite-volume limit. This is still an important open question to be resolved related to the existence and the nature of the spin glass phase transition in finite dimension. The situation is more tractable for the model on trees, see [5], and on the half-plane, see [3, 2]. We expect that, at least for $d = 2$, the critical droplets of all edges have finite size (uniformly in Λ), in which case a much stronger version of Theorem 1.1 should hold:

Conjecture 1.2. Consider the Hamiltonian (1.1) at $d = 2$. For any $\delta > 0$ there exist $A = A(\delta)$ and $t_0 = t_0(\delta)$ (independent of Λ) such that $\mathbb{P}(A) > 1 - \delta$ uniformly in Λ, and

$$\text{on } A, \ Q_\Lambda(\sigma(J + t\varepsilon), \sigma(J + t\varepsilon')) > 1 - \delta \qquad \text{for all } t \le t_0.$$

In words, as the perturbation is turned on, all but a set of vertices of small density remain unchanged. If true, then this implies that the variance of the difference of ground state energies goes like the volume of Λ, see Theorem 1.5 in [4]. This would likely give an approach to prove uniqueness of the ground state in the infinite volume by implementing a strategy similar to the one of Aizenman & Wehr for the random field Ising model [1]. Another interesting result in $d = 2$ that may be relevant to the nature of the ground states is that the satisfied edges, where $\sigma_x \sigma_y = \text{sgn } J_{xy}$, do not percolate in an infinite-volume ground state as shown by Berger & Tessler in [6].

Our results extend to positive temperature, where σ is no longer the ground state, but rather independent random configurations σ, σ'. The configuration σ is sampled with probability proportional to the Gibbs weight $e^{-\beta H_{\Lambda, J + t\varepsilon}(\sigma)}$ for fixed realizations of J and ε, and σ' is chosen analogously (with ε' replacing ε). The distribution of such a pair is denoted by $\langle \cdot \rangle_t$. Note that, when $t = 0$, the states σ, σ' are independently sampled from the same distribution.

For simplicity, we prove the positive-temperature result only under the simplifying assumption that 0 is in the support of J (as in the zero-temperature proof of Section 2.1), though this assumption can be removed by techniques similar to those of Section 2.2.

Theorem 1.3 (Positive Temperature). Assume 0 is in the support of J, and let $\langle \cdot \rangle_t$ be defined as above. For any $\delta > 0$ there exists $t_0 = t_0(\delta)$ such that

$$\mathbb{E}[\langle Q_\Lambda(\sigma, \sigma') \rangle_t] > (1 - \delta)\delta' \mathbb{P}(A), \text{ for any } t < t_0(\delta),$$

where A is some explicit event with $\mathbb{P}(A) > 0$ uniformly in Λ and $\delta' > 0$ is an explicit constant depending on the distribution of J and β.

Notation. We write $B(x, n)$ for the set of vertices whose ℓ^∞-distance to x is less or equal to n. In other words, $B(x, n)$ is a box centered at x of sidelength $2n + 1$. We also write

$\|\cdot\|_1$ for the ℓ^1-norm on $\mathbb{Z}^d$. For the sake of conciseness, we will often use the following notation for the product $\sigma_x\sigma_y$:

$$\text{for } e = \{x, y\}, \ \sigma_e = \sigma_x\sigma_y \ .$$

1.2 Method of Proof

The three proofs of the theorem are based on the following idea. For a given edge e, we find a non-empty subset $A_e \subseteq \mathbb{R}^{(\Lambda \cup \partial\Lambda)^*}$ of realizations of couplings such that

- A_e is open and depends on a finite number of couplings uniformly in Λ and e (i.e., it is an open cylinder set);
- $\sigma_e(J)$ is constant on A_e.

To prove the theorem, we first write the overlap in terms of the event A_e

$$
\begin{aligned}
\mathbb{E}[Q_\Lambda(\sigma(J + t\varepsilon), \sigma(J + t\varepsilon'))] = {} & \frac{1}{|\Lambda^*|} \sum_{e \in \Lambda^*} \mathbb{E}[\sigma_e(J + t\varepsilon)\sigma_e(J + t\varepsilon'); J \in A_e] \\
& + \frac{1}{|\Lambda^*|} \sum_{e \in \Lambda^*} \mathbb{E}[\sigma_e(J + t\varepsilon)\sigma_e(J + t\varepsilon'); J \in A_e^c] \ .
\end{aligned}
\tag{1.6}
$$

We observe that, by conditioning on J, the independence of the perturbations $\varepsilon, \varepsilon'$ yields

$$\mathbb{E}[\sigma_e(J + t\varepsilon)\sigma_e(J + t\varepsilon'); J \in A_e^c] = \mathbb{E}[(\mathbb{E}[\sigma_e(J + t\varepsilon)|J])^2 ; J \in A_e^c] \geq 0 \ .$$

Therefore, this gives the lower bound

$$
\begin{aligned}
\mathbb{E}[Q_\Lambda(\sigma(J + t\varepsilon), \sigma(J + t\varepsilon'))] \geq {} & \frac{1}{|\Lambda^*|} \sum_{e \in \Lambda^*} \mathbb{E}[\sigma_e(J + t\varepsilon')\sigma_e(J + t\varepsilon')); J \in A_e] \\
= {} & \frac{1}{|\Lambda^*|} \sum_{e \in \Lambda^*} \Big\{ 2\mathbb{P}(\{\sigma_e(J + t\varepsilon) = \sigma_e(J + t\varepsilon')\} \cap \{J \in A_e\}) \\
& \qquad - \mathbb{P}(A_e) \Big\} \ .
\end{aligned}
$$

Since by assumption the ground state at e is constant on A_e, the summand is larger than

$$2\mathbb{P}(J \in [A_e \cap (A_e - t\varepsilon) \cap (A_e - t\varepsilon')]) - \mathbb{P}(A_e) \ ,$$

where $A_e - t\varepsilon$ stands for the translate of A_e by $t\varepsilon$. Moreover, since we assume that A_e is open and depends only on a finite number of couplings uniformly in Λ, for any $0 < \delta < 1$, there exists $t_0(\delta)$ independent of Λ such that

$$\mathbb{P}(J \in [A_e \cap (A_e - t\varepsilon) \cap (A_e - t\varepsilon')]) > (1 - \delta)\mathbb{P}(A_e) \text{ for } t < t_0(\delta) \ .$$

The conclusion of the theorem follows from this.

2 Proof of Theorem 1.1 using (anti-)ferromagnetic edges

2.1 Proof under assumptions on the support of J

We first suppose that a neighborhood of 0 is included in the support of the distribution of J.

Let $e = \{x, y\}$ be an edge. We write A_e for the set of realizations of couplings

$$A_e = \left\{ J : \ J_{xy} > \sum_{i=1}^{2d-1} |J_{xz_i}| \right\}$$

where z_i, $i = 1, \ldots, 2d - 1$, stands for the vertices neighboring x other than y. Clearly, A_e is an open set of $\mathbb{R}^{(\Lambda \cup \partial \Lambda)^*}$ that depends on only finite number of couplings. Moreover, we have $\mathbb{P}(A_e) > 0$ uniformly in Λ if a neighborhood of 0 is in the support of the distribution of J. Such edges were referred to as *super-satisfied* in [15, 3].

Proof of Theorem 1.1. In view of Section 1.2, it remains to show that the ground state is constant on A_e. In fact, it must be that $\sigma_e = +1$, otherwise the ground state property (1.2) is violated for $\mathcal{B} = \{x\}$. $\qquad\square$

2.2 Proof under no assumptions on the support of J

We present a modified construction to prove Theorem 1.1 without assumptions on the support of J. We are no longer able to force the satisfaction status of a given e in the ground state. Hence, we now construct an event A_v for a fixed vertex v on which, for suitably chosen M, many edges of $B(v, M)^*$ have stable satisfaction status. In other words, for a fixed vertex v, we show $\sigma_f(J + t\varepsilon) = \sigma_f(J)$ for "most" f near v. The argument makes clear the role of the finite-dimensionality in the absence of disorder chaos: perturbation-induced changes in boundary conditions change the energy by at most the order of the boundary size, but this cannot flip order of the volume number edges.

Fix M large and constant (to be chosen precisely). Choose some interval $I = (a, b)$ such that $\mathbb{P}(J_{xy} \in I) > 0$ and such that $0 \notin [a, b]$; we write $\operatorname{sign}(I)$ for the common sign of the elements of I. For v such that $B(v, M) \subset \Lambda$, we define A_v as follows:

$$A_v = \{J : \ J_f \in I, \text{for all } f \in B(v, M)^* \cup \partial^* B(v, M)\} \ .$$

Note that A_v is an open subset of $\mathbb{R}^{(\Lambda \cup \partial \Lambda)^*}$ that depends on a finite number of couplings. On the event A_v, the contribution to the total energy from bonds within $B(v, M)^*$ is minimized by homogeneous ferromagnetic or antiferromagnetic configurations (depending on the value of $\operatorname{sign}(I)$), which satisfy all bonds of $B(v, M)$. The contribution to the energy from bonds in $\partial^* B(v, M)$ is only of order M^{d-1}, and so it will follow that the restriction of $\sigma(J)$ to $B(v, M)$ must still satisfy a high density of bonds of $B(v, M)$. This is the content of the following proposition.

Proposition 2.1. Suppose M is large enough that $(|a| \wedge |b|) M^d > 2d(2M + 1)^{d-1}(|a| \vee |b|)$, and let $v \in \Lambda$ such that $B(v, M) \subseteq \Lambda$. Then on the event A_v, we have

$$|\{f \in B(v, M)^* : J_f \sigma_f(J) > 0\}| \geq (3/4)|B(v, M)^*| \ .$$

Proof. Suppose for a contradiction that the above cardinality bound were false; in particular, $J_f \sigma_f(J) < 0$ for at least M^d edges $f \in B(v, M)^*$ (note that $|B(v, M)^*| = d(M - 1)M^{d-1}$). Define the modification $\overline{\sigma}$ of $\sigma(J)$ obtained by satisfying all bonds in $B(v, M)^*$ and leaving the configuration outside $B(v, M)$ unchanged:

$$\overline{\sigma}_x = \begin{cases} \sigma_x(J) \ , & x \notin B(v, M); \\ (\operatorname{sign}(I))^{\|x\|_1} \ , & x \in B(v, M). \end{cases}$$

On one hand, since $\sigma(J)$ is the ground state, we have $H_{\Lambda, J}(\overline{\sigma}) - H_{\Lambda, J}(\sigma(J)) \geq 0$. On the hand, estimating this energy difference directly, we get

$$H_{\Lambda, J}(\overline{\sigma}) - H_{\Lambda, J}(\sigma(J)) = \left[\sum_{x, y \in B(v, M)} + \sum_{\substack{x \in B(v, M) \\ y \in \partial B(v, M)}} \right] J_{xy}(\overline{\sigma}_x \overline{\sigma}_y - \sigma_x(J)\sigma_y(J))$$

$$\leq -2(|a| \wedge |b|)M^d + 4d(2M + 1)^{d-1}(|a| \vee |b|) \ .$$

In estimating the first term above, we used the fact that on A_v all edges in $B(v, M)^*$ are satisfied in $\overline{\sigma}$, but (by assumption) at least M^d are unsatisfied in $\sigma(J)$. For the second term, we use the fact that $|\overline{\sigma}_f - \sigma(J)_f| \leq 2$ for each edge f. Under the assumption on M, this contradicts $H_{\Lambda,J}(\overline{\sigma}) - H_{\Lambda,J}(\sigma(J)) \geq 0$. $\qquad\square$

We now prove Theorem 1.1.

Proof of Theorem 1.1. Fix M as in the statement of Proposition 2.1. We break Λ up into blocks of sidelength $(2M + 1)$ centered at vertices $v \in \Lambda$ (with some smaller boxes at the boundary if the sidelength of Λ is not a multiple of $(2M + 1)$ — since there are $o(|\Lambda|)$ of these, we can disregard them in what follows). By Proposition 2.1, we have for all $t < t_0(\delta)$ for some $t_0(\delta)$ (independent of Λ),

$$\sum_{e \in B(v,M)^*} \mathbb{P}(\{\sigma_e(J + t\varepsilon) = \sigma_e(J + t\varepsilon')\} \cap A_v) \geq (1 - \delta)(3/4)|B(v, M)^*| \cdot \mathbb{P}(A_v) \,.$$

The claim then follows similarly to Section 1.2, since the number of boxes $B(v, M)$ is proportional to the size of Λ. $\qquad\square$

3 Proof of Theorem 1.1 using critical droplets

For a fixed edge $e = \{x, y\}$, a good measure of the sensitivity under perturbation of the ground state at e is given by the set of vertices $\mathcal{D}_e = \mathcal{D}_e(J)$ containing either x or y with the lowest flip energy in the ground state $\sigma(J)$. In other words,

$$\sum_{b \in \partial^* \mathcal{D}_e} J_b \sigma_b(J) = \min_{\mathcal{B}:e \in \partial^* \mathcal{B}} \left\{ \sum_{b \in \partial^* \mathcal{B}} J_b \sigma_b(J) \right\} \,.$$

The set $\mathcal{D}_e$ is referred to as the *critical droplet of the edge* e. More generally, we consider the spin configurations that minimize $H_{\Lambda,J}$ with a fixed configuration at e

$$\sigma^{\pm,e}(J) = \operatorname*{argmin}_{s \in \{-1,+1\}^\Lambda, s_e = \pm 1} H_{\Lambda,J}(s) \,.$$

As for the ground state, the states $\sigma^{\pm,e}$ as functions of J are well-defined on the open set $\mathbb{R}^{(\Lambda \cup \partial \Lambda)^*} \setminus \mathcal{C}$. Clearly, the ground state σ is either $\sigma^{+,e}$ or $\sigma^{-,e}$. Moreover, the critical droplet $\mathcal{D}_e$ is exactly the set of vertices where $\sigma^{+,e}$ and $\sigma^{-,e}$ differ. Similarly as in (1.2), for any set of vertices $\mathcal{B}$ such that $e \notin \partial^* \mathcal{B}$, we must have

$$\sum_{b \in \partial^* \mathcal{B}} J_b \sigma_b^{\pm,e}(J) > 0 \,. \tag{3.1}$$

The following elementary fact will be needed.

Lemma 3.1. Let $\mathcal{D}_e$ be the critical droplet of the edge $e = \{x, y\}$. Then $\mathcal{D}_e$ and $\mathcal{D}_e^c$ are a.s. connected as subgraphs of Λ.

Proof. Without loss of generality, we take $x \in \mathcal{D}_e$. Suppose $\mathcal{D}_e^c$ is not connected. Then it has at least one connected component, say $\mathcal{B}$, that does not contain y nor x. In particular, $\partial^* \mathcal{B}$ does not contain e. But by definition of the droplet, the energy of the boundary in $\sigma^{+,e}$ and $\sigma^{-,e}$ are of opposite signs. In particular, this implies

$$\sum_{b \in \partial^* \mathcal{B}} J_b \sigma^{+,e}(J) < 0 \text{ or } \sum_{b \in \partial^* \mathcal{B}} J_b \sigma^{-,e}(J) < 0,$$

thereby contradicting Equation (3.1). The same way, suppose $\mathcal{D}_e$ is not connected. Then there must exist a connected component $\mathcal{B}$ of $\mathcal{D}_e$ that does not contain x (since x is the other component) nor y (since y is in $\mathcal{D}_e^c$). In particular, e is not contained in the edge-boundary of $\mathcal{B}$, and the same reasoning as above applies. $\qquad\square$

In the next section, we explicitly construct an event with positive probability on which the critical droplet is of size one. In particular, this implies that the ground state is constant on some subset A_e of that event, thereby providing another proof of Theorem 1.1. In Section 3.2, we show how this argument can be generalized on the event that the droplet is of finite size (uniformly in Λ) with positive probability.

3.1 A critical droplet of size one

We first describe a construction which shows that $\mathcal{D}_e$ can be of order one (in fact, of cardinality exactly one) with nonvanishing probability, based on the presence of locally ferromagnetic regions. Roughly speaking, a local region of sufficiently ferromagnetic bonds encircling e causes nearby spins to strongly prefer to align, preventing the droplet from propagating outside this region.

The construction requires that the common distribution of the J_{xy}'s has a support which is not too concentrated, and so we work under the following assumption:

Assumption 3.2. There is an $r \geq 0$ such that both $\mathbb{P}(|J_e| \leq r) > 0$ and $\mathbb{P}(|J_e| > d \cdot 3^d r) > 0$.

(This obviously holds if 0 is in the support, for example.) The above assumption allows us to show a particularly strong form of bounded droplet size, as in the following proposition.

Proposition 3.3. Suppose that the distribution of J's satisfies Assumption 3.2. Let $e = \{x, y\} \in \Lambda^*$ be an edge such that both x, y are at a distance at least 2 from the boundary. There exists $c > 0$, uniform in Λ and in the choice of e, such that $\mathbb{P}(\mathcal{D}_e = \{x\}) \geq c$.
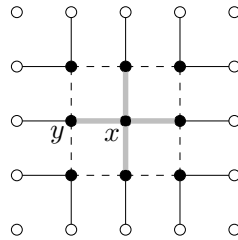

Figure 1: A depiction of the defining conditions of the event A_e, $e = \{x, y\}$. The vertices in $B(x, 1)$ are the black circles. Edges satisfying Condition (1) are the dotted lines. The ones satisfying Condition (2) are the black lines, and the ones under Condition (3) are in grey.

The set A_e of coupling realizations needed to prove Theorem 1.1, as outlined in Section 1.2, is defined by the following three conditions: for every $v \in B(x, 1) \setminus \{x\}$,

1. if $w \notin B(x, 1) \setminus \{x\}$, then $|J_{vw}| < r$;
2. if $w \in B(x, 1) \setminus \{x\}$, then $|J_{vw}| > d \cdot 3^d r$, and all such J_{vw} have the same sign;
3. $\sum_{w:\{x,w\} \in \Lambda^*} J_{xw} > 0$. (If J is negative with probability one, then take < 0.)

By construction, the set A_e is an open set that depends on a finite number of couplings. Moreover, its probability is positive and independent of e by Assumption 3.2. We prove that $A_e \subset \{\mathcal{D}_e = \{x\}\}$, thereby implying Proposition 3.3.

Proof of Proposition 3.3. We prove a stronger claim: if $u \in B(x, 1) \setminus \{x\}$, then $\sigma_u^{+,e} = \sigma_u^{-,e}$. This immediately implies Proposition 3.3: $\mathcal{D}_e$ is connected (see Lemma 3.1) and contains x but no neighbor of x, so it must be $\{x\}$.

Let $\mathcal{S} \subset \{-1, +1\}^\Lambda$ be the set of spin configurations s for which all edges between neighbors of x are satisfied, i.e., if $\{u, v\} \in \Lambda^*$ and $u, v \in B(x, 1) \setminus \{x\}$, then $J_{uv} s_u s_v > 0$. As a first step we show that $\sigma^{\pm, e} \in \mathcal{S}$ on A_e. Suppose it is not the case for $\sigma^{+, e}$, i.e., there exist u_0, v_0 as in the statement of the proposition, $J_{u_0 v_0} \sigma_{u_0}^{+, e} \sigma_{v_0}^{+, e} < 0$. We construct another spin configuration, denoted η, satisfying $\eta_e = +1$ for which $J_{uv} \eta_u \eta_v > 0$ for all u, v. We show that η has lower energy than $\sigma^{+, e}$, contradicting the definition of $\sigma^{+, e}$.

We set $\eta_x = +1$ and choose the value of η at other sites as follows. If the common sign of the J_{vw}'s in item (2) of the definition of A_e is positive, then we let $\eta_u = \sigma_u^{+, e}$ when $u \notin B(x, 1) \setminus \{x\}$; when $u \in B(x, 1) \setminus \{x\}$, we set $\eta_u = +1$. If the common sign of the J_{vw}'s in item (2) is negative, we again let $\eta_u = \sigma_u^{+, e}$ when $u \notin B(x, 1) \setminus \{x\}$, but now set $\eta_u = (-1)^{1 + \|u - x\|_1}$ when $u \in B(x, 1) \setminus \{x\}$.

We claim that $H_{\Lambda, J}(\sigma^{+, e}) > H_{\Lambda, J}(\eta)$. Indeed, because $\sigma^{+, e}$ and η may disagree only at x or on $B(x, 1) \setminus \{x\}$,

$$
\begin{aligned}
H_{\Lambda, J}(\sigma^{+, e}) - H_{\Lambda, J}(\eta) &= \sum_{u, v:\, u \in B(x, 1) \setminus \{x\}} J_{uv}(\eta_u \eta_v - \sigma_u^{+, e} \sigma_v^{+, e}) \\
&= \sum_{u, v \in B(x, 1) \setminus \{x\}} J_{uv}(\eta_u \eta_v - \sigma_u^{+, e} \sigma_v^{+, e}) \\
&\quad + \sum_{\substack{u \in B(x, 1) \setminus \{x\} \\ v \notin B(x, 1) \setminus \{x\}}} J_{uv}(\eta_u \eta_v - \sigma_u^{+, e} \sigma_v^{+, e}) \\
&\geq 2|J_{u_0, v_0}| - d \cdot (2)(3^d) \cdot \max_{\substack{u \in B(x, 1) \setminus \{x\}, \\ v \notin B(x, 1) \setminus \{x\}}} |J_{uv}| > 0 \,.
\end{aligned}
\tag{3.2}
$$

Here the first sum in (3.2) is bounded by noting that each $J_{uv} \eta_u \eta_v$ term of that sum is positive, but at least the term $J_{u_0 v_0} \sigma_{u_0}^{+, e} \sigma_{v_0}^{+, e}$ is negative. The second sum in (3.2) is controlled by lower-bounding each term by $-2|J_{uv}|$; the final inequality comes from the definition of A_e. This completes the contradiction and the proof in the case of $\sigma^{+, e}$. The proof in the case of $\sigma^{-, e}$ is similar.

Consider the bijection $\varphi : \mathcal{S} \to \mathcal{S}$ that flips the spin at x: $\varphi(s)_x = -s_x$ and $\varphi(s)_u = s_u$ for $u \neq x$. It remains to show that $\sigma^{-, e} = \varphi(\sigma^{+, e})$. Observe that the map φ maps spin configurations s, s' such that $J_e s_e = J_e s'_e = 1$ to configurations $\varphi(s), \varphi(s')$ such that $J_e \varphi(s)_e = J_e \varphi(s')_e = -1$, and

$$
H_{\Lambda, J}(s) - H_{\Lambda, J}(\varphi(s)) = H_{\Lambda, J}(s') - H_{\Lambda, J}(\varphi(s')).
\tag{3.3}
$$

On A_e, it must be that $s_x s_u$ is constant as u ranges over neighbors of x for $s \in \mathcal{S}$. This implies:

$$
\text{if } s \in \mathcal{S}, \quad H_{\Lambda, J}(\varphi(s)) - H_{\Lambda, J}(s) = 2 \sum_{u:\{u, x\} \in \Lambda^*} J_{xu} s_x s_u = 2 s_e \sum_u J_{xu} \,.
\tag{3.4}
$$

Assume without loss of generality that $\sigma(J) = \sigma^{+, e}$. Suppose it were the case that $H(\sigma^{-, e}) < H(\varphi(\sigma^{+, e}))$ — and in particular that $\sigma^{-, e} \neq \varphi(\sigma^{+, e})$. Then, it would also follow by (3.3) and (3.4) that $H(\varphi(\sigma^{-, e})) < H(\sigma^{+, e})$, contradicting the fact that $\sigma^{+, e}$ is the ground state. This completes the proof. $\qquad \square$

To prove Theorem 1.1, it remains to show that $\sigma_e(J)$ is constant on A_e.

Proof of Theorem 1.1. By Proposition 3.3, the critical droplet $\mathcal{D}_e$ is simply $\{x\}$ on A_e. In particular, the ground state is solely determined between $\sigma^{+, e}$ and $\sigma^{-, e}$ by the condition

$$
\sum_{u:\{u, x\} \in \Lambda^*} J_{xu} \sigma_x \sigma_u > 0 \,.
$$

We know from the proof of Proposition 3.3 that, on A_e, $\sigma_x^{\pm,e}\sigma_u^{\pm,e}$ is constant as u ranges over the neighbors of x. Therefore the above condition is reduced to

$$\sigma_e \cdot \left(\sum_{u:\{u,x\}\in\Lambda^*} J_{xu} \right) > 0 \ .$$

Condition 3 and the above ensure that σ_e is constant on A_e as claimed. $\qquad\square$

3.2 A general argument for finite critical droplet

In this section, we prove:

Proposition 3.4. Let A_e be an open set of $\mathbb{R}^{(\Lambda\cup\partial\Lambda)^*}$ depending on a finite number of edges such that: $A_e \subset \{|\mathcal{D}_e| < K\}$ for some $K > 0$ and $\mathbb{P}(A_e) > 0$ uniformly in Λ. Then there exists an open set $\widetilde{A}_e \subset A_e$ depending on a finite number of edges such that $\mathbb{P}(\widetilde{A}_e) > 0$ uniformly in Λ, and the ground state is constant on $\widetilde{A}_e$. In particular, Theorem 1.1 holds for the event $\widetilde{A}_e$.

The fact that the theorem holds is again by the reasoning of Section 1.2. It remains to prove that the ground state can be made constant on a subset of A_e. For that purpose, we define the *flexibility at the edge* e

$$\mathcal{F}_e(J) = \left| H_{\Lambda,J}(\sigma^{+,e}) - H_{\Lambda,J}(\sigma^{-,e}) \right| = 2 \sum_{\substack{b=\{u,v\}: \\ u\in\mathcal{D}_e(J),v\in\partial\mathcal{D}_e(J)}} J_b\sigma_b(J) > 0 \ .$$

In other words, the flexibility is the minimal energy of all surfaces going through e. The definition given above makes sense whenever J is not in the critical set $\mathcal{C}$. However, it can be extended to a continuous function on all of $\mathbb{R}^{\Lambda^*}$.

Lemma 3.5. The flexibility $\mathcal{F}_e(J)$ extends to continuous function on $\mathbb{R}^{(\Lambda\cup\partial\Lambda)^*}$. More precisely, it is a piecewise affine function of J with

$$\frac{\partial \mathcal{F}_e(J)}{\partial J_b} = \begin{cases} 2\sigma_b(J) & \text{if } b\in\partial^*\mathcal{D}_e(J), \\ 0 & \text{else,} \end{cases}$$

whenever $J \in \mathbb{R}^{(\Lambda\cup\partial\Lambda)^*} \setminus \mathcal{C}$.

Proof. Note that by definition we have

$$H_{\Lambda,J}(\sigma^{+,e}) = \min_{s\in\{-1,+1\}^\Lambda:s_e=+1} H_{\Lambda,J}(s), \qquad H_{\Lambda,J}(\sigma^{-,e}) = \min_{s\in\{-1,+1\}^\Lambda:s_e=-1} H_{\Lambda,J}(s) \ .$$

Clearly, the function $J \mapsto H_{\Lambda,J}(s)$ is a continuous function of J for a fixed s. Therefore the minimum of such functions over a finitely many values of s is itself a continuous function. This shows that $J \mapsto \mathcal{F}_e(J)$ extends to a continuous function on $\mathbb{R}^{(\Lambda\cup\partial\Lambda)^*}$.

For the derivatives, observe that $\sigma^{\pm,e}(J)$ are locally constant on the complement of $\mathcal{C}$. Therefore if the ground state is given by $\sigma^{+,e}(J)$, say, we have

$$\mathcal{F}_e(J) = H_{\Lambda,J}(\sigma^{-,e}(J)) - H_{\Lambda,J}(\sigma^{+,e}(J)) = 2 \sum_{\substack{b=\{u,v\}: \\ u\in\mathcal{D}_e(J),v\in\partial\mathcal{D}_e(J)}} J_b\sigma_b(J) \ .$$

The claim on the derivatives is then obvious. $\qquad\square$

By putting the coupling J_e apart, the flexibility can also be written as

$$\mathcal{F}_e(J) = 2|J_e - \mathcal{C}_e(J)| \ , \tag{3.5}$$

where $\mathcal{C}_e$ does not depend on J_e as a function of J. In particular, $\mathcal{C}_e(J)$ seen as a random variable is independent of J_e.

Proof of Proposition 3.4. Consider, for some $\delta > 0$ (to be fixed later), the set

$$\widetilde{A}_e = A_e \cap \{\mathcal{F}_e(J) > \delta\} \ .$$

The set $\widetilde{A}_e$ is open, since $\mathcal{F}_e$ is continuous, and depends only a finite number of edges since $\mathcal{F}_e$ is the flip energy of the droplet, the size of which is bounded by K on A_e. Moreover, by Equation (3.5), the parameter δ can be taken small enough so that $\mathbb{P}(\mathcal{F}_e > \delta)$ is arbitrarily close to 1. This implies that for δ small enough $\mathbb{P}(\widetilde{A}_e) > 0$ uniformly in Λ. By considering a subset of $\widetilde{A}_e$ if necessary, we can assume without loss of generality that $\widetilde{A}_e$ is a cylinder set whose cross-section is a finite-dimensional ball in K^* coordinates, having finite radius.

We now prove that the flexibility is strictly positive on $\widetilde{A}_e$. For $J, J' \in A_e$, Lemma 3.5 implies the following representation

$$\mathcal{F}_e(J') = \mathcal{F}_e(J) + \int_{J \to J'} \nabla \mathcal{F}_s(\mathbf{r}) \cdot \mathrm{d}\mathbf{r} \ . \tag{3.6}$$

By the last paragraph, the above only depends on K^* coordinates. In particular, the difference can be bounded by

$$|\mathcal{F}_e(J') - \mathcal{F}_e(J)| \leq \|\nabla \mathcal{F}_e\| \cdot \|J - J'\|_{\mathbb{R}^{K^*}} \ .$$

This can be made smaller than $\delta/2$ (uniformly in Λ). This is because $\|\nabla \mathcal{F}_e\| \leq 2K^*$ on $\widetilde{A}_e$ by the assumption on the size of the droplet on A_e and Lemma 3.5, and $\|J - J'\|_{\mathbb{R}^{K^*}}$ can be made as small as we wish by reducing the radius of $\widetilde{A}_e$. Together with (3.6), this implies that $\mathcal{F}_e(J) > 0$ on $\widetilde{A}_e$.

We now conclude that the ground state is constant on $\widetilde{A}_e$. Suppose it is not. Then there must exist J and J' in $\widetilde{A}_e$ such that $\sigma^{+,e}$ is the ground state at J and $\sigma^{-,e}$ is the ground state at J'. In particular, this implies

$$\mathcal{F}_e(J) = H_{\Lambda,J}(\sigma^{-,e}(J)) - H_{\Lambda,J}(\sigma^{+,e}(J), \qquad \mathcal{F}_e(J') = H_{\Lambda,J'}(\sigma^{+,e}(J')) - H_{\Lambda,J}(\sigma^{-,e}(J')) \ .$$

By continuity of $\mathcal{F}_e(J)$, this implies that any path from J to J' in $\widetilde{A}_e$ will contain at least one point J_0 with $\mathcal{F}_e(J_0) = 0$. This contradicts the previous result. $\qquad\square$

4 Proof of the positive-temperature Theorem 1.3.

We define more explicitly the positive-temperature Gibbs specification $\langle \cdot \rangle_t$: given functions f, g on $\{-1, +1\}^\Lambda$ and fixed joint disorder realization $J, \varepsilon, \varepsilon'$, we set

$$\begin{aligned}
\langle f(\sigma)g(\sigma') \rangle_t &= \sum_{\sigma, \sigma' \in \{-1,+1\}^\Lambda} \frac{f(\sigma)g(\sigma')}{Z_t Z_t'} \exp\left(-\beta \left[H_{\Lambda, J+t\varepsilon}(\sigma) + H_{\Lambda, J+t\varepsilon'}(\sigma')\right]\right) \\
&= \langle f(\sigma) \rangle_t \langle g(\sigma') \rangle_t \ ,
\end{aligned}$$

where the partition function $Z_t = \sum_{\sigma \in \{-1,+1\}^\Lambda} \exp(-\beta H_{\Lambda, J+t\varepsilon})$ (Z_t' is defined analogously, replacing ε with ε'). We prove the theorem, similarly to the proof in Section 2.1, using the assumption on the support of J to "super-satisfy" individual edges. To start, we again decompose to isolate the contribution to Q_Λ from each edge:

$$\mathbb{E}[\langle Q_\Lambda(\sigma, \sigma') \rangle_t] = \frac{1}{|\Lambda^*|} \sum_{e \in \Lambda^*} \mathbb{E}[\langle \sigma_e \sigma_e' \rangle_t] \ ; \tag{4.1}$$

Once we show that each term of the above is lower-bounded by $(1-\delta)\mathbb{P}(A_e)$ for $t < t_0$ (for appropriate choices of $\mathbb{P}(A_e)$ and t_0), the theorem will be proved.

For simplicity, we assume that $\mathbb{P}(J_e > 0) > 0$; the adaptations needed to treat the case $\mathbb{P}(J_e < 0) = 1$ are straightforward. We make nearly the same definition of A_e as in Section 2.1, namely: $A_e = \left\{ J_{xy} > \sum_{i=1}^{2d-1} |J_{xz_i}| + a \right\}$, where the z_i's are the neighbors of x other than y and where $a > 0$ is chosen such that $\mathbb{P}(A_e) > 0$. We compute

$$\mathbb{E}[\langle \sigma_e \sigma_e' \rangle_t] = \mathbb{E}[\langle \sigma_e \sigma_e' \rangle_t; A_e] + \mathbb{E}[\langle \sigma_e \sigma_e' \rangle_t; A_e^c] . \tag{4.2}$$

Conditioning on J and using independence, we see that the second term of (4.2) is positive:

$$\mathbb{E}[\langle \sigma_e \sigma_e' \rangle_t; A_e^c] = \mathbb{E}[\langle \sigma_e \rangle_t \langle \sigma_e' \rangle_t; A_e^c] = \mathbb{E}[\mathbb{E}[\langle \sigma_e \rangle_t \mid J]^2; A_e^c] \geq 0 .$$

It thus suffices to lower-bound the first term of (4.2) for a particular $e = \{x, y\}$.

We recall from Section 3.1 that $\varphi : \{-1, +1\}^\Lambda \to \{-1, +1\}^\Lambda$ denotes the bijection that flips the spin at x: $\varphi(s)_x = -s_x$ and $\varphi(s)_u = s_u$ for $u \neq x$. Since φ maps configurations s such that $J_e s_e > 0$ to configurations such that $J_e \varphi(s)_e < 0$, we can use φ to compare the contribution of such pairs $s, \varphi(s)$ to the expectation in (4.1). On the event $A_e \cap (A_e - t\varepsilon) \cap (A_e - t\varepsilon')$, we note as in Section 2.1 that there is a nonnegative energy cost for failing to satisfy edge e. More explicitly: for each s such that $s_e = 1$, both $H_{\Lambda, J+t\varepsilon}(s) \leq H_{\Lambda, J+t\varepsilon}(\varphi(s)) - a$ and $H_{\Lambda, J+t\varepsilon'}(s) \leq H_{\Lambda, J+t\varepsilon'}(\varphi(s)) - a$ whenever $0 \leq t < t_0(\delta)$ for an appropriate choice of $t_0(\delta)$.

On this event, we estimate the thermal average $\langle \sigma_e \rangle_t$:

$$
\begin{aligned}
\langle \sigma_e \rangle_t &= Z_t^{-1} \sum_{\sigma \in \{-1,+1\}^\Lambda} \sigma_e \exp(-\beta H_{\Lambda, J+t\varepsilon}(\sigma)) \\
&= Z_t^{-1} \sum_{\substack{\sigma \in \{-1,+1\}^\Lambda \\ \sigma_e = 1}} \exp(-\beta H_{\Lambda, J+t\varepsilon}(\sigma)) - \exp(-\beta H_{\Lambda, J+t\varepsilon}(\varphi(\sigma))) \\
&\geq (1 - e^{-\beta a}) Z_t^{-1} \sum_{\substack{\sigma \in \{-1,+1\}^\Lambda \\ \sigma_e = 1}} \exp(-\beta H_{\Lambda, J+t\varepsilon}(\sigma)) \\
&\geq (1 - e^{-\beta a})(2 Z_t)^{-1} \sum_{\substack{\sigma \in \{-1,+1\}^\Lambda \\ \sigma_e = 1}} \exp(-\beta H_{\Lambda, J+t\varepsilon}(\sigma)) - \exp(-\beta H_{\Lambda, J+t\varepsilon}(\varphi(\sigma))) \\
&= (1 - e^{-\beta a})/2 ,
\end{aligned}
$$

where the inequality in the third line is valid on the event $A_e \cap (A_e - t\varepsilon) \cap (A_e - t\varepsilon')$, for $0 \leq t < t_0$.

An identical argument shows the same lower bound for $\langle \sigma_e' \rangle_t$ under the same conditions, so

$$\langle \sigma_e \sigma_e' \rangle_t = \langle \sigma_e \rangle_t \langle \sigma_e' \rangle_t \geq \delta' := ((1 - e^{-\beta a})/2)^2 \text{ on } A_e \cap (A_e - t\varepsilon) \cap (A_e - t\varepsilon'), \text{ for } 0 \leq t < t_0;$$

applying this in the first term of (4.2) and inserting our estimates into (4.1) completes the proof of the theorem. $\square$

References

[1] M. Aizenman and J. Wehr. Rounding effects of quenched randomness on first-order phase transitions. *Comm. Math. Phys.*, 130(3):489–528, 1990. MR-1060388

[2] L.-P. Arguin and M. Damron. On the number of ground states of the Edwards-Anderson spin glass model. *Ann. Inst. H. Poincaré Probab. Statist.*, 50(1):28–62, 02 2014. MR-3161521

[3] L.-P. Arguin, M. Damron, C. M. Newman, and D. L. Stein. Uniqueness of ground states for short-range spin glasses in the half-plane. *Comm. Math. Phys.*, 300(3):641–657, 2010. MR-2736957

[4] L.-P. Arguin, C. M. Newman, and D. L. Stein. A relation between disorder chaos and incongruent states in spin glasses on $\mathbb{Z}^d$. *Communications in Mathematical Physics*, 367(3):1019–1043, May 2019. MR-3943487

[5] J. Bäumler. Uniqueness and non-uniqueness for spin-glass ground states on trees. *Electron. J. Probab.*, 24:17 pp., 2019. MR-3991114

[6] N. Berger and R. J. Tessler. No percolation in low temperature spin glass 1.2. *Electron. J. Probab.*, 22:Paper No. 88, 19, 2017. MR-3718716

[7] A. J. Bray and M. A. Moore. Chaotic nature of the spin-glass phase. *Phys. Rev. Lett.*, 58:57–60, Jan 1987.

[8] S. Chatterjee. *Superconcentration and related topics*. Springer Monographs in Mathematics. Springer, Cham, 2014. MR-3157205

[9] W.-K. Chen. Disorder chaos in the Sherrington-Kirkpatrick model with external field. *Ann. Probab.*, 41(5):3345–3391, 2013. MR-3127885

[10] W.-K. Chen, M. Handschy, and G. Lerman. On the energy landscape of the mixed even p-spin model. *Probability Theory and Related Fields*, 171(1):53–95, Jun 2018. MR-3800830

[11] W.-K. Chen, H.-W. Hsieh, C.-R. Hwang, and Y.-C. Sheu. Disorder chaos in the spherical mean-field model. *Journal of Statistical Physics*, 160:417–429, 2015. MR-3360467

[12] W.-K. Chen and D. Panchenko. Disorder chaos in some diluted spin glass models. *Ann. Appl. Probab.*, 28(3):1356–1378, 06 2018. MR-3809466

[13] W.-K. Chen and A. Sen. Parisi formula, disorder chaos and fluctuation for the ground state energy in the spherical mixed p-spin models. *Comm. Math. Phys.*, 350(1):129–173, 2017. MR-3606472

[14] D. S. Fisher and D. A. Huse. Ordered phase of short-range ising spin-glasses. *Phys. Rev. Lett.*, 56:1601–1604, Apr 1986.

[15] C. M. Newman and D. L. Stein. Are there incongruent ground states in 2D Edwards-Anderson spin glasses? *Comm. Math. Phys.*, 224(1):205–218, 2001. Dedicated to Joel L. Lebowitz. MR-1868997