Dancing with frustration

Short-range Ising spin glasses:

what is proved, what is believed, and where proofs break

Yan Ru Pei

Independent researcher

yanrpei@gmail.com ORCID 0000-0002-6999-3691

November 2026

doi:10.5281/zenodo.24061128

Abstract

The Edwards–Anderson model, Ising spins on $\mathbb{Z}^d$ with independent symmetric random nearest-neighbour couplings, was introduced in 1975 as the simplest short-range model of a spin glass. Physicists expect a spin-glass transition at positive temperature in three dimensions and above, and none in two. Mathematically, no proof of spin-glass order is known for symmetric independent couplings in any finite dimension, and for undiluted laws such as $\pm J$ or Gaussian there is no proof that the two-dimensional model is disordered at all positive temperatures.

This review records, as of November 2026, what is proved about short-range Ising spin glasses, what is believed, and where proofs break. It covers the competing physical pictures (replica symmetry breaking, droplets, trivial–nontrivial and chaotic pairs) and the numerics that calibrate them; metastates and the two-dimensional fluctuation bounds; ground states, disorder chaos and Chatterjee's zero-temperature theorems; the cluster representations proposed for frustrated systems, and two-replica blue percolation; high-temperature methods, the comparison with the random-field Ising model, and a comparison with spanning forests and electrical networks; and hierarchical lattices, diluted lattices and the Nishimori line. Results from several programmes are restated in one notation, and the arguments that teach something are proved or sketched. Route warnings record attractive lines of argument together with the lemma each one would need and the cases in which that lemma fails. A final part collects the barriers that any proof of order or non-order must respect, small test cases for candidate inequalities, and open problems. Results of the author's own programme are marked with their status, and each computer-assisted proof states what its certificate checks.

Reader's guide

The question

Take Ising spins on $\mathbb{Z}^d$, put an independent symmetric random coupling (for example $\pm J$ or Gaussian) on every nearest-neighbour edge, and lower the temperature. Does the system freeze into a state with long-range order in the overlap of two replicas, a *spin-glass phase*, and if so in which dimensions? This is the Edwards–Anderson (EA) question. Physicists have a fairly confident answer from numerics and renormalization-group reasoning: no transition at positive temperature in $d = 2$, a transition in $d = 3$ and above. Mathematically almost nothing of this is proved. There is no proof of spin-glass order for symmetric independent couplings in any finite dimension at any temperature, and, for undiluted laws such as $\pm J$ or Gaussian, no proof that the two-dimensional model is disordered at all positive temperatures (for diluted laws just above the percolation threshold this is proved, Chapter 13).

The model was introduced by Edwards and Anderson in 1975 [81] as the simplest short-range model of the dilute magnetic alloys then being called spin glasses; its mean-field version followed the same year [224]. In the half century since, the companion models of the subject have been transformed by an exchange of methods with other fields (Table 1). Percolation and the random-cluster model drew on Boolean analysis, the theory of algorithms and complex analysis; the arboreal gas on supersymmetry; the Sherrington–Kirkpatrick model on interpolation, ultrametric combinatorics, random-matrix geometry and the theory of computation. The short-range model has had no comparable exchange. The metastate framework, the fluctuation bounds and the zero-temperature chaos results described in this review are real progress, and some of them import tools from the same fields, but none of them decides whether the model orders: for $\pm J$ and Gaussian couplings the structural question posed in 1975 is open in every dimension $d \geq 2$.

We expect the short-range spin glass to be one of the holy grails of mathematical physics in the coming years. It is also a natural test bed for frontier AI systems that do mathematics: the questions are precise and long open, partial results come with checkable hypotheses, and the barriers collected in Chapter 15 provide concrete test cases that any proposed argument must pass.

The subtitle, *what is proved, what is believed, and where proofs break*, names the three threads that run through the review. *What is proved*: the theorems about states, metastates, ground states, chaos, graphical representations and special lattices, each stated exactly and, where it teaches something, proved or sketched. *What is believed*: the competing physical pictures and the numerics that calibrate them, kept visibly separate from the theorems. *Where proofs break*: the routes that look attractive and the precise lemma each one would need, together with the cases in which that lemma is known to fail. The review is restricted to short-range Ising spin glasses; the Sherrington–Kirkpatrick model and other mean-field models appear only where one of their tools or identities is used.

Model	Fields that fed in	Landmark results
Bernoulli and random-cluster (FK) percolation	Boolean analysis (noise sensitivity, sharp thresholds); theory of algorithms (decision trees, the OSSS inequality); discrete complex analysis and SLE	$p_c = \frac{1}{2}$ on $\mathbb{Z}^2$ [137]; noise sensitivity of crossings [25, 105]; sharp transitions for every $q \geq 1$ [78]; the self-dual point is critical [24]; conformal invariance [147, 225]
Ising and ϕ^4	random-current (graph) representations	marginal triviality of critical scaling limits in $d = 4$ [4]
Arboreal gas (random spanning forests)	supersymmetry and hyperbolic sigma models; electrical networks	fermionic representation [41]; hyperbolic symmetry and its consequences [23]; uniqueness of the infinite tree [113]
Sherrington–Kirkpatrick and mean-field spin glasses	interpolation and convexity; ultrametric combinatorics; random matrices and the geometry of landscapes; theory of computation (message passing, optimization, hardness); Gaussian analysis	the Parisi formula [111, 210, 231]; ultrametricity [107, 207]; landscape complexity [18]; algorithms that follow the ground state [163, 228] and the overlap-gap barrier [98]; superconcentration and chaos [48]
Short-range EA model on $\mathbb{Z}^d$	ergodic theory (metastates); Gaussian analysis; graphical representations; gauge identities	the structure of states and metastates (Chapters 3 and 4); zero-temperature chaos (Chapter 6); high-temperature and special-lattice theorems (Chapters 9 and 12 to 14); no theorem deciding order or non-order at low temperature for $\pm J$ or Gaussian couplings in any $d \geq 2$

Table 1: Companion models and the fields whose methods transformed them, with a few landmark results. The list is illustrative, not exhaustive.

Why another review? The standard accounts of spin-glass theory are mostly about mean field. The physics texts of Mézard, Parisi and Virasoro [161] and of Castellani and Cavagna [43], and the mathematical monographs of Talagrand [232] and Panchenko [208], centre on the Sherrington–Kirkpatrick model and its relatives; the books of Bovier [34] and of Contucci and Giardinà [62] give the short-range model a chapter or two. The substantive mathematical treatments of the short-range model are Newman's lectures [171] and the surveys and book of Newman and Stein [183, 185, 186], with a recent chapter on metastates by Newman, Read and Stein [190] in [45]. Much has happened since those accounts were written: zero-temperature chaos and droplet theorems, two-dimensional fluctuation bounds and their limits, new graphical representations, classifications on hierarchical lattices, results on the Nishimori line, and computer-assisted theorems. This review is meant partly as a progress tracker: a record, as of November 2026, of where the mathematical understanding of short-range Ising spin glasses stands, what is proved, what is believed, and what any proof must get past.

How statements are labelled

- **Theorem, Proposition, Lemma, Corollary**: proved results, with the source in brackets. A proof in the text is complete unless it is labelled *Proof sketch*, whose first sentence says what it omits. Proofs that are long without being illuminating are cited or moved to Appendix A.

- **Conjecture**: a precise mathematical statement that is open.
- **Physics prediction**: a statement supported by physical reasoning or numerics, not by proof.
- **Observation**: a fact that is not a theorem of a published paper, often a small computation. The text says how each was checked (by a proof, by an exact computation, or numerically), so not every observation is a theorem.
- **Route warning**: an attractive line of argument together with the reason it cannot work as stated, marked as proved false, failing even under optimistic assumptions, or as a failure only suggested by numerics or physics.
- **Open problem**: a question we think is worth a year of effort, usually quoted from the literature.

Our own results are cited by their arXiv or Zenodo identifiers where they are public; those not yet published carry no citation and are proved or sketched in the text or in Appendix A, except for a few results of Chapter 8 that belong to a revised preprint not yet public and are marked as such where they appear. A computer-assisted proof is marked as such, and Appendix B lists what each certificate checks. Every theorem is stated with its quantifiers in one order: the coupling law and the dimension, the temperature range, "for almost every J" (or the disorder-averaged form), the state, boundary condition or metastate the conclusion is about, and then the conclusion. Results are restated in one notation (Appendix C translates it into the notations of the main programmes).

A map of the review

Part I: the question
Chapter 1 sets up the model, gauge symmetry and what "order" means, proves the easy facts (existence and self-averaging of the quenched free energy, high-temperature uniqueness) and shows why the usual correlation inequalities fail. Chapter 2 explains the physical pictures (replica symmetry breaking, droplets, trivial–nontrivial (TNT), chaotic pairs) and the numerics.

Part II: states, metastates and two dimensions
Chapter 3 develops metastates, the tool that makes "many states" a precise question. Chapter 4 is the fluctuation programme in two dimensions and the table of which pictures survive which theorems.

Part III: zero temperature
Chapter 5 treats ground states, domain walls and zero-temperature dynamics; Chapter 6 treats disorder chaos and superconcentration, including Chatterjee's zero-temperature theorems.

Part IV: graphical representations
Chapter 7 puts the proposed cluster representations of frustrated Ising systems side by side. Chapter 8 follows one of them, two-replica blue percolation (percolation of a random set of edges satisfied in both of two replicas), as far as it goes toward order, and reads the blue count as a statement about periodic metastates.

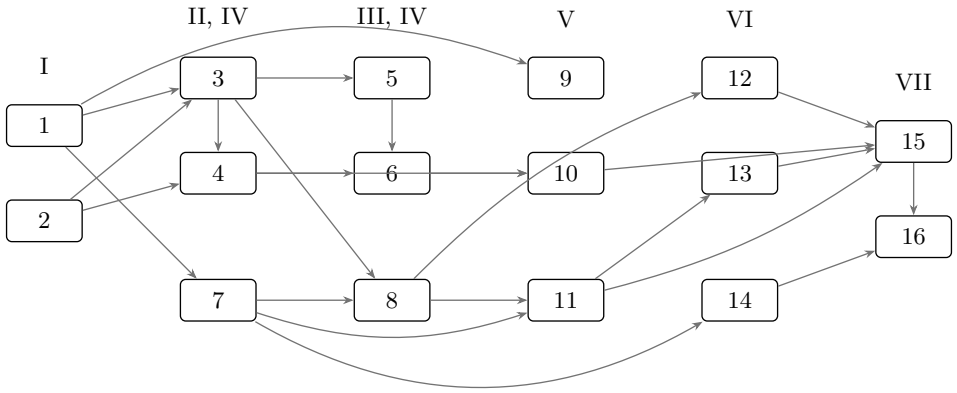


Figure 1: Main dependencies between chapters (an arrow means "uses definitions or results of"). Chapter 1 is needed everywhere; the parts after it can be read in several orders. Chapters 15 and 16 draw on all of them.

Part V: toward non-order in low dimension
Chapter 9 is what high-temperature methods give; Chapter 10 compares with the random-field Ising model, whose analogous questions are solved, and isolates the step that does not transfer; Chapter 11 develops a comparison with spanning forests and electrical networks.

Part VI: testbeds and the Nishimori line
Chapter 12 covers hierarchical and long-range proxies, Chapter 13 diluted and decorated lattices, and Chapter 14 gauge symmetry, the Nishimori line and biased disorder.

Part VII: where proofs break, and what to aim at
Chapter 15 collects the barriers that a proof of order or non-order must respect, with a list of small test cases on which a candidate inequality should be checked first; Chapter 16 lists open problems and the programmes attacking them.

Reading paths

- *For the open problems only*: the first sections of Chapter 1, Chapter 2, then Chapter 16, with Chapter 15 for what a solution must avoid.
- *For techniques*: metastates and fluctuation bounds (Chapters 3 and 4), chaos and superconcentration (Chapter 6), graphical representations (Chapter 7) and the random-field comparison (Chapter 10).
- *For the cluster and comparison routes*: Chapters 1, 7, 8 and 11, then the warnings of Chapter 15.
- *For readers coming from physics*: Chapter 2, the picture-by-constraint table at the end of Chapter 4, then Chapter 14.

What is known, at a glance

d	Rigorous results (nearest-neighbour, symmetric iid couplings unless stated)	Physics expectation
1	No spin-glass order at any positive temperature: a gauge transformation maps the chain to a ferromagnetic chain (Chapter 1).	no transition
2	Unique Gibbs state for T above the critical temperature of the $\|J\|$-ferromagnet (the ferromagnet with couplings $\|J_{xy}\|$; $T > 2.269$ for $\pm J$; Chapter 1); certified L^2 non-order for $\pm J$ for $T \geq 1.90926$ (Chapter 9); for continuous laws (not $\pm J$), fluctuation constraints on metastates and on ground states (Chapters 4 and 5); all-temperature non-order conditional on a comparison conjecture (Chapter 11).	no transition at $T > 0$
3	Uniqueness for $T > T_c$ of the $\|J\|$-ferromagnet (≈ 4.51 for $\pm J$; Chapter 1); certified L^2 non-order for $\pm J$ for $T \geq 2.73214$ (Chapter 9).	transition at $T_c \approx 1.10$ ($\pm J$)
large; 7–12	Two-replica blue percolation in both overlap sectors: for all large d at $\beta = c/(2d\,\mathbb{E}\|J\|)$, each fixed $c > 1$, by a cavity argument, and computer-assisted at seventeen explicit temperatures for $\pm J$ in $d = 7, \dots, 12$ (from a revised preprint not yet public); it does not by itself imply order (Chapter 8).	transition
every d	For Gaussian couplings, zero-temperature disorder chaos and rough droplet boundaries (Chapter 6); site chaos extends to every symmetric law and every temperature, with the perturbation acting on underlying Gaussian variables (for free boundary conditions this is a theorem of Chen, Kim and Sen); no proof of spin-glass order at any temperature for symmetric iid couplings.	transition for $d \geq 3$
diluted	For the diluted $\pm J$ law with the density of nonzero couplings slightly above the percolation threshold, in every $d \geq 2$: a unique Gibbs state at every temperature, with an explicit, computer-assisted window in $d = 2$ (Chapter 13).	
biased	For $\pm J$ couplings on $\mathbb{Z}^2$ that are positive with probability p: ferromagnetic order at the Nishimori temperature for $p > 0.9634$; overlap order and Gibbs non-uniqueness at all temperatures up to twice the Nishimori temperature for $p > 0.9634$ and, computer-assisted, for $p \geq 0.9537$ at $\beta \geq 10^6$ (Chapter 14).	

Table 2: A coarse summary; each entry is stated exactly, with its hypotheses, in the chapter cited. "L^2 non-order" means that $\mathbb{E}\langle\sigma_x\sigma_y\rangle_J^2$ decays as $|x-y| \to \infty$, uniformly over finite subgraphs with free boundary conditions and over tori of side at least 11, which rules out EA order (Section 1.4) for these boundary conditions; "certified" means computer-assisted, with certificates checked in exact rational or interval arithmetic (Appendix B). Physics values are numerical estimates (Chapter 2).

The gap between the two columns is the subject of the review. The most informative partial results come from three directions that rarely meet: the metastate, fluctuation and ground-state programme of Newman, Stein and collaborators (Part II and Chapter 5), next to which sit the zero-temperature chaos theorems of Chatterjee and of Arguin and Hanson (Chapter 6); exact structures, gauge identities and special lattices, including the Nishimori–Ohzeki line of work (Part VI); and graphical and comparison methods (Parts IV and V), including the blue-percolation and forest-comparison routes that we develop (Chapters 8 and 11). Chapter 15 explains why each of them stops where it does, and Chapter 16 what it would take to go further.

Contents

Reader's guide ii

I The question 1

1 The model and what order would mean 2
1.1 The Edwards–Anderson model . . . 2
1.2 Gauge symmetry and frustration . . . 3
1.3 Infinite-volume states . . . 4
1.4 What order would mean . . . 5
1.5 The free energy: existence and self-averaging . . . 7
1.6 High temperature: comparison with the ferromagnet . . . 8
1.7 Why the ferromagnetic toolbox fails . . . 11
1.7.1 No second Griffiths inequality . . . 11
1.7.2 No quenched positivity with a ferromagnetic reference . . . 12
1.7.3 No FKG for the averaged overlap . . . 13
1.8 Where the threads go . . . 13

2 What physicists believe 15
2.1 The mean-field template . . . 15
2.2 The stiffness argument . . . 16
2.3 Four pictures of the low-temperature phase . . . 17
2.3.1 Droplets and scaling . . . 17
2.3.2 Replica symmetry breaking in short range . . . 18
2.3.3 Trivial–nontrivial . . . 19
2.3.4 Chaotic pairs . . . 19
2.3.5 What distinguishes the pictures . . . 20
2.4 The numbers . . . 20
2.5 From beliefs to theorems . . . 22

II States, metastates and two dimensions 23

3 How many states? Metastates 24
3.1 Finite volumes and their limits . . . 24
3.2 Metastates . . . 26
3.3 What translation covariance forbids . . . 28

3.4 Interfaces are space-filling . . . 30
3.5 Inside a Gibbs state: weights and structure . . . 31
3.6 The four classes . . . 32
3.7 Overlap identities . . . 34
3.8 Summary and outlook . . . 36

4 Fluctuations and the two-dimensional programme **37**
4.1 Imry–Ma and Aizenman–Wehr . . . 38
4.2 A volume-order variance bound . . . 39
4.3 Two dimensions . . . 41
4.3.1 Zero temperature . . . 43
4.4 Upper bounds and disorder chaos . . . 44
4.5 Bond observables in two dimensions . . . 45
4.6 Route warnings . . . 46
4.7 The pictures against the constraints . . . 46
4.8 Summary and outlook . . . 47

III Zero temperature **49**

5 Ground states and their stability **50**
5.1 Ground states on infinite graphs . . . 50
5.2 Two dimensions: at most one domain wall . . . 53
5.3 The half-plane . . . 55
5.4 Frustration, matching and unsatisfied edges in the plane . . . 56
5.5 Zero-temperature dynamics and fixation . . . 58
5.6 Where the two-dimensional methods stop . . . 59

6 Chaos and superconcentration **60**
6.1 Perturbing the couplings . . . 60
6.2 The Hermite basis and the Ornstein–Uhlenbeck semigroup . . . 61
6.3 Superconcentration and the edge overlap . . . 62
6.4 Site chaos: the gauge identity . . . 64
6.5 Consequences: valleys, droplets and critical droplets . . . 66
6.6 Boundary sensitivity . . . 68
6.7 The spectral sample . . . 69
6.8 What the chaos theorems do and do not show . . . 70

IV Graphical representations **71**

7 Cluster representations of frustrated Ising systems **72**
7.1 One plaquette, many pictures . . . 72
7.2 The Edwards–Sokal coupling with frustration . . . 73
7.3 Averaging over the couplings restores insertion tolerance . . . 74
7.4 Frustrated FK clusters percolate in the plane . . . 76
7.5 The Nishimori line: averaged FK bonds are independent . . . 77
7.6 Two replicas: double FK and the CMR colouring . . . 78

7.7 Signed FK percolation . . . 80
7.8 Two copies and disagreement percolation . . . 82
7.9 Frustration-adjusted random currents . . . 82
7.10 Cluster moves for simulation . . . 83
7.11 How the representations fit together . . . 83
7.11.1 One recipe: allowance matrices . . . 83
7.11.2 One probability space . . . 86
7.12 A tangent: cluster representations beyond edges . . . 87
7.13 Summary . . . 89

8 Blue percolation and imbalance **90**
8.1 The setting: periodic joint limits and the two sectors . . . 90
8.2 The crossing lemma: correlation forces blue connection . . . 92
8.3 At most two infinite blue clusters . . . 93
8.4 Both sectors percolate in high dimension . . . 94
8.4.1 The leaf identity and the cost of restoring an edge . . . 94
8.4.2 One attempt per target, and thinning . . . 95
8.4.3 Proof of the high-dimensional theorem . . . 96
8.4.4 Explicit dimensions . . . 97
8.5 Percolation without order . . . 98
8.6 Imbalance would give order . . . 99
8.7 What the blue count says about the periodic metastate . . . 100
8.7.1 Two periodic replicas sample one state twice . . . 100
8.7.2 Distinct states force blue percolation where they disagree . . . 101
8.7.3 The blue count reads the pair of pure states . . . 102
8.7.4 Consequences for the four classes . . . 103
8.8 Why the technique stops short of three dimensions . . . 105
8.9 Summary . . . 106

V Toward non-order in low dimension **107**

9 High temperature and its edge **108**
9.1 Two notions of a high-temperature phase . . . 108
9.2 The high-temperature expansion and Griffiths singularities . . . 109
9.3 Comparison with the ferromagnet, and a strict improvement . . . 110
9.4 The cavity L^2 inequality . . . 111
9.5 Certified windows in two and three dimensions . . . 115
9.6 Where the cavity method stops . . . 116
9.7 The Kac limit . . . 117
9.8 What is established, and what is open . . . 117

10 The random-field Ising model as a mirror **119**
10.1 The model and the Imry–Ma picture . . . 119
10.2 Two dimensions: no order, and how fast it is lost . . . 120
10.3 Three and more dimensions: order by a Peierls map . . . 121
10.4 The mirror image: the joint flip is a gauge transformation . . . 123

10.5 The transfer table . . . 125
10.6 Route warnings . . . 126
10.7 What the mirror shows . . . 127

11 Comparison with forests and electrical networks 128
11.1 Fixed magnitudes, fair signs, and electrical networks . . . 128
11.1.1 Electrical tools . . . 129
11.2 One edge, one bundle, one cut . . . 130
11.3 The arboreal gas and its electrical bound . . . 132
11.4 The comparison conjectures and what they would give . . . 135
11.5 Series–parallel networks . . . 136
11.6 What a comparison model must contain . . . 139
11.7 Where the comparison stands . . . 140

VI Testbeds and the Nishimori line 142

12 Hierarchical and long-range proxies 143
12.1 Migdal–Kadanoff lattices and their exact recursion . . . 143
12.2 No order by exact dilution . . . 145
12.3 Order by coupling growth . . . 147
12.4 What does not transfer to the cubic lattice . . . 150
12.5 Dyson's hierarchical spin glass . . . 151
12.6 One-dimensional long-range chains . . . 152
12.6.1 Exotic states . . . 153
12.6.2 A transition on the Nishimori line . . . 154
12.7 Summary . . . 155

13 Diluted and decorated lattices 157
13.1 Exact cancellation and finite effective clusters . . . 157
13.2 Decorated lattices . . . 159
13.2.1 Forest information cannot decide order . . . 161
13.3 Diluted $\pm J$ couplings on $\mathbb{Z}^d$ . . . 162
13.4 One mechanism in four settings . . . 164
13.5 Route warnings . . . 165
13.6 What the testbeds show . . . 165

14 Gauge symmetry and the Nishimori line 167
14.1 Biased couplings and the gauge average . . . 167
14.2 Magnetization and overlap . . . 169
14.3 Griffiths inequalities on the line . . . 170
14.4 Order near full bias: random contours . . . 171
14.5 Below the multicritical point: vertical or reentrant? . . . 173
14.6 The duality conjecture . . . 174
14.7 The inference reading . . . 175
14.8 Planar models and free fermions . . . 176
14.9 Summary . . . 176

VII Where proofs break, and what to aim at **178**

15 Barriers **179**
15.1 What a barrier is . . . 179
15.2 The fair average must be used exactly . . . 180
15.3 Forest information cannot decide order . . . 180
15.4 Percolation is not order . . . 182
15.5 The qualitative hypotheses do not imply order . . . 182
15.6 No positivity survives the quenched average . . . 183
15.7 Scale-uniform control is the wall . . . 184
15.8 What a comparison reference must contain . . . 186
15.9 Disorder-conjugacy tools are blind to order . . . 187
15.10 Biased order does not transfer to fair disorder . . . 188
15.11 A kill-test suite . . . 189
15.12 Summary . . . 189

16 Open problems and programmes **193**
16.1 Ranked targets . . . 193
16.2 A map of conditional reductions . . . 196
16.3 Programmes, frontiers and attack points . . . 198
16.3.1 Newman, Stein and collaborators: metastates, fluctuations, $d = 2$ . . . 198
16.3.2 Nishimori, Ohzeki and Okuyama: gauge identities, hierarchical lattices, long range . . . 198
16.3.3 Chatterjee and zero-temperature chaos . . . 199
16.3.4 Random-field-style multiscale methods . . . 199
16.3.5 Blue percolation, imbalance and forest comparison . . . 200
16.4 Newman's research problems . . . 200
16.5 Summary . . . 200

A Proofs moved out of the main text **202**
A.1 Proof of the quenched free energy theorem, parts (b) and (c) . . . 202
A.2 The bridge from the finite-volume target to the pressure target . . . 202
A.3 Proof of Newman's uniqueness theorem . . . 203
A.4 The rung computation . . . 204
A.5 The Aizenman–Wehr construction . . . 204
A.6 Ergodicity of iid couplings . . . 205
A.7 Stochastic stability . . . 205
A.8 Concentration of the energy . . . 205
A.9 The Gaussian lower bound in two dimensions . . . 206
A.10 Energy bookkeeping for the dynamics of Chapter 5 . . . 207
A.11 Hermite facts used in Chapter 6 . . . 207
A.12 Proof of Theorem 6.7 for a symmetric law with 0 in its support . . . 208
A.13 Proof of Theorem 6.21 . . . 208
A.14 The paths in the proof of Theorem 6.21 . . . 209
A.15 Annealed insertion tolerance in periodic limits (Chapter 7) . . . 209
A.16 The one-attempt exploration (Chapter 8) . . . 210
A.17 Almost-sure coexistence of the two blue sectors (Chapter 8) . . . 210

A.18 Two routine proofs for the cavity bounds (Chapter 9) . . . 211
A.19 Exponential L^2 decay in fixed dimension (Chapter 9) . . . 211
A.20 The reduction behind the certified windows (Chapter 9) . . . 212
A.21 Sharpness of the cavity loss (Chapter 9) . . . 215
A.22 Tilted sign laws on three parallel edges (Chapter 11) . . . 216
A.23 Restoring one edge (Chapter 11) . . . 216
A.24 The closure theorem (Chapter 11) . . . 216
A.25 The edge of the random-cluster family (Chapter 11) . . . 218
A.26 The random-cluster electrical factor (Chapter 11) . . . 218
A.27 The uniform-flux XY comparison on a ring of squares (Chapter 11) . . . 218
A.28 Exotic states of long-range chains (Chapter 12) . . . 219
A.29 The certified bound for biased couplings (Chapter 14) . . . 219

B Computer-assisted certificates **221**
B.1 Certified theorems . . . 221
B.2 Exact finite computations inside proofs . . . 222

C Notation across programmes **225**
C.1 Translation tables . . . 225
C.2 The author's papers . . . 226
C.3 Framings . . . 226

D Cluster-representation formulas **229**
D.1 One replica: the frustrated Edwards–Sokal coupling . . . 229
D.2 The Nishimori line . . . 230
D.3 Signed FK . . . 230
D.4 Two replicas: double FK and the CMR colouring . . . 230
D.5 Two copies and random currents . . . 231
D.6 Constants of the blue-percolation proofs . . . 231

E Sources of the numerics **233**

Part I

The question

Chapter 1

The model and what order would mean

Take the nearest-neighbour Ising model on $\mathbb{Z}^d$ and give every edge an independent coupling of random sign. Does the resulting system freeze into an ordered state at low temperature, as a ferromagnet does, and if so, in which dimensions? This question, behind the whole review, is open in every dimension $d \geq 2$: for symmetric couplings on $\mathbb{Z}^d$ no proof is known that any form of spin-glass order occurs at any temperature, and for undiluted laws such as $\pm J$ or Gaussian no proof is known that it is absent at all positive temperatures in $d = 2$ (for diluted laws just above the percolation threshold absence is proved, Chapter 13).

This chapter defines the Edwards–Anderson model, its Gibbs states and its gauge symmetry, computes three small examples (a chain, a tree, one frustrated square), and makes precise several inequivalent meanings of "order"; later chapters prove different ones. Two easy theorems follow: the quenched pressure has a self-averaging limit (a convexity argument of Contucci and Lebowitz), and at high temperature the Gibbs state is unique (Newman's comparison with the associated ferromagnet). Explicit finite counterexamples then show why second Griffiths inequalities, the FKG lattice condition and monotonicity under adding bonds are unavailable. By the end the reader can state the central open problem precisely, prove the two theorems, and say which inequality fails, on which small graph, when one tries to copy a ferromagnetic proof.

1.1 The Edwards–Anderson model

The Edwards–Anderson (EA) model [81] keeps the geometry of the Ising ferromagnet and replaces its couplings by independent random variables of both signs. Some edges then prefer their endpoints aligned and some prefer them opposite, and on a cycle with an odd number of negative edges no spin configuration can satisfy them all.

Definition 1.1 (The EA model)**.** Let $G = (V, E)$ be a graph; our main case is $V = \mathbb{Z}^d$ with the nearest-neighbour edge set E_d. Write $\Lambda_L = \{-L, \dots, L\}^d$, and for finite $\Lambda \subset \mathbb{Z}^d$ let $E(\Lambda)$ be the set of edges with both endpoints in Λ, $\partial\Lambda$ the set of sites outside Λ adjacent to Λ, and $\partial_e\Lambda$ the set of edges joining Λ to $\partial\Lambda$.

(i) The *couplings* $J = (J_e)_{e \in E}$ are independent with common law ν on $\mathbb{R}$. The law is *symmetric* if J_e and $-J_e$ have the same law; this is the EA model proper. The main

examples are $\pm J$, $\nu = \frac{1}{2}(\delta_{+1} + \delta_{-1})$, and *Gaussian*, $\nu = N(0,1)$. We also use the *biased* law $\nu_p = p\,\delta_{+1} + (1-p)\,\delta_{-1}$ and the *diluted* law $\mu_p = (1-p)\,\delta_0 + \frac{p}{2}(\delta_{+1} + \delta_{-1})$.

(ii) The *Hamiltonian* of $\sigma \in \{-1,+1\}^V$ is $H_J(\sigma) = -\sum_{xy\in E} J_{xy}\sigma_x\sigma_y$, and $\beta \geq 0$ is the inverse temperature. An edge xy is *satisfied* by σ if $J_{xy}\sigma_x\sigma_y > 0$.

(iii) We write $\mathbb{P}$ and $\mathbb{E}$ for the law of J and its expectation (the *disorder average*), and $\langle\cdot\rangle_J$ for a Gibbs average at fixed J, with the volume and boundary condition added when needed.

The finite-volume Gibbs measures are the usual ones. For finite $\Lambda \subset \mathbb{Z}^d$ and a boundary condition $\eta \in \{-1,+1\}^{\mathbb{Z}^d}$ let

$$\mu^{\beta,\eta}_{\Lambda,J}(\sigma) = \frac{1}{Z^{\eta}_{\Lambda,J}(\beta)} \exp\Big(\beta \sum_{xy\in E(\Lambda)} J_{xy}\sigma_x\sigma_y + \beta \sum_{\substack{xy\in\partial_e\Lambda \\ x\in\Lambda}} J_{xy}\sigma_x\eta_y\Big), \qquad \sigma \in \{-1,+1\}^{\Lambda}. \tag{1.1}$$

Dropping the boundary sum gives the *free* boundary condition $\mu^{\beta,\mathrm{f}}_{\Lambda,J}$; on a box, identifying opposite faces gives the *periodic* one. Free and periodic boundary conditions are invariant under the global flip $\sigma \mapsto -\sigma$, so all their odd correlations vanish. A boundary condition is *coupling-independent* if it does not depend on J; free, periodic and any fixed η are of this kind. Coupling-dependent boundary conditions will matter later, because they are the only way to select an ordered state when there are many (Chapter 3).

1.2 Gauge symmetry and frustration

A ferromagnet has one obvious ordered state, all plus. The EA model has no such distinguished configuration, because of a symmetry of the joint law of couplings and spins.

Definition 1.2 (Gauge transformation)**.** For $\tau \in \{-1,+1\}^V$ let J^τ and $\sigma\tau$ be given by $J^\tau_{xy} = J_{xy}\tau_x\tau_y$ and $(\sigma\tau)_x = \sigma_x\tau_x$.

Lemma 1.3 (Gauge invariance)**.** *For every τ, $H_{J^\tau}(\sigma\tau) = H_J(\sigma)$. Consequently, for finite Λ, with $\tau_x = 1$ off Λ, the map $\sigma \mapsto \sigma\tau$ carries $\mu^{\beta,\eta}_{\Lambda,J}$ to $\mu^{\beta,\eta}_{\Lambda,J^\tau}$, and the same holds for free and periodic boundary conditions with arbitrary τ. If ν is symmetric, then J^τ has the same law as J for each fixed τ, and for every coupling-independent boundary condition b and every nonempty $A \subset \Lambda$,*

$$\mathbb{E}\,\langle\sigma_A\rangle^{b}_{\Lambda,J} = 0, \qquad \sigma_A = \textstyle\prod_{x\in A}\sigma_x. \tag{1.2}$$

Proof. $J^\tau_{xy}(\sigma\tau)_x(\sigma\tau)_y = J_{xy}\tau_x\tau_y\sigma_x\tau_x\sigma_y\tau_y = J_{xy}\sigma_x\sigma_y$ term by term, and the boundary terms are unchanged because $\tau = 1$ off Λ. For symmetric ν the map $J \mapsto J^\tau$ flips the sign of an independent symmetric coordinate on each edge with $\tau_x\tau_y = -1$, so it preserves the product law. For (1.2) pick $x \in A$ and $\tau = -1$ at x only; then $\langle\sigma_A\rangle_{\Lambda,J} = -\langle\sigma_A\rangle_{\Lambda,J^\tau}$, and averaging over J gives $\mathbb{E}\langle\sigma_A\rangle = -\mathbb{E}\langle\sigma_A\rangle$. □

So averaged correlations carry no information for symmetric couplings; every interesting quantity must be gauge invariant, such as $\langle\sigma_x\sigma_y\rangle^2_J$ or the overlap of two replicas introduced in Section 1.4. The gauge-invariant content of the signs is their *frustration*.

Definition 1.4 (Frustration)**.** A cycle C of G is *frustrated* for J if $\prod_{e\in C} \operatorname{sgn} J_e = -1$.

Frustration is gauge invariant, since each vertex of a cycle contributes $\tau_x^2 = 1$ to the product. Conversely, if G is connected and no cycle with nonzero couplings is frustrated, then J is gauge equivalent to $|J|$: choose τ on a spanning forest of the edges with $J_e \neq 0$ so that the forest edges become positive; every other edge with $J_e \neq 0$ closes a cycle of nonzero couplings with forest edges and is then positive too. In particular *on a tree, and in $d = 1$, the EA model is a ferromagnet in disguise*, and so is the Mattis model [157], $J_{xy} = \xi_x \xi_y$ with arbitrary signs ξ. Random signs alone do not make a spin glass; frustrated cycles do.

Example 1.5 (The chain)**.** On $\mathbb{Z}$ the gauge $\tau_x = \prod_{0<y\le x} s_{y-1,y}$, with $s_e = \operatorname{sgn} J_e$ if $J_e \neq 0$ and $s_e = 1$ otherwise (and symmetrically for $x < 0$), turns any J into $|J|$ (Fig. 1.1(c)). With free boundary conditions the high-temperature expansion of a tree gives $\langle \sigma_0 \sigma_n \rangle_J = \prod_e \tanh(\beta J_e)$ over the edges between 0 and n, so for iid couplings

$$\mathbb{E}\langle \sigma_0 \sigma_n \rangle_J^2 = m^{|n|}, \qquad m = \mathbb{E} \tanh^2(\beta J_e) < 1, \qquad \sum_{n \in \mathbb{Z}} \mathbb{E}\langle \sigma_0 \sigma_n \rangle_J^2 = \frac{1+m}{1-m} < \infty \tag{1.3}$$

at every $\beta < \infty$, and the Gibbs state is unique at every $\beta < \infty$ (Corollary 1.17 below with $p_c(1) = 1$).

Example 1.6 (Trees)**.** On an infinite tree T with $\pm J$ couplings the gauge again removes all signs, so the free-boundary state is the gauge image of the free-boundary state of the Ising ferromagnet on T. Its two-point functions $\langle \sigma_\rho \sigma_v \rangle_J = \pm(\tanh \beta)^{|v|}$ decay exponentially at every temperature. Nevertheless Evans, Kenyon, Peres and Schulman [84, Theorem 1.1] (for regular trees, Bleher, Ruiz and Zagrebnov [30]) show that the free-boundary state is extremal when $\tanh \beta < \operatorname{br}(T)^{-1/2}$ and is not extremal when $\tanh \beta > \operatorname{br}(T)^{-1/2}$, where $\operatorname{br}(T)$ is the branching number ($\operatorname{br} = b$ for the $(b{+}1)$-regular tree). In the second regime their reconstruction bound [84, Theorem 1.1(i)] (flip probability $\frac{1}{2}(1 - \tanh \beta)$) gives $\inf_n \mu^{\mathrm{f}} |\mu^{\mathrm{f}}(\sigma_\rho \mid \sigma_{T_n})| > 0$, with T_n the nth level of T and μ^{f} the free state (a gauge-invariant quantity). By the Markov property and reverse martingale convergence, $\mu^{\mathrm{f}}(\sigma_\rho \mid \sigma_{T_n})$ tends to the conditional expectation of σ_ρ given the tail σ-field, which is therefore not a.s. zero: the free state decomposes into extremal states whose root magnetizations have a positive mean square, although $\langle \sigma_\rho \rangle = 0$ in the free state. The transition is invisible in the two-point function of the chosen state and visible only in its decomposition, a theme of this review.

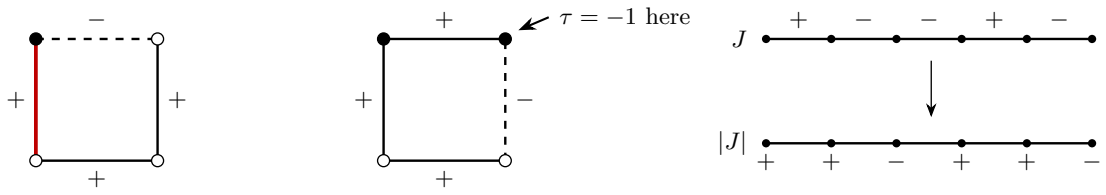
Example 1.7 (One frustrated plaquette)**.** Take the square with vertices $1, 2, 3, 4$ in cyclic order, unit couplings, and $F = \prod_e J_e$. By gauge invariance everything depends only on F. With $t = \tanh \beta$ the high-temperature expansion has two terms, the empty set and the whole cycle:

$$Z_J = 2^4 \cosh^4 \beta \, (1 + F\, t^4). \tag{1.4}$$

If $F = -1$ every configuration violates at least one edge; the ground states violate exactly one, which can be any of the four, and the global flip doubles the count to eight ground states (Fig. 1.1(a)). The frustrated square has the *smaller* partition function, $1 - t^4 < 1 + t^4$. This innocent fact drives the failure of FKG for the averaged two-replica law in Proposition 1.23.

1.3 Infinite-volume states

Order concerns what survives of (1.1) as $\Lambda \uparrow \mathbb{Z}^d$, that is, the infinite-volume Gibbs states of Dobrushin, Lanford and Ruelle.



(a) $F = -1$: one edge unsatisfied

(b) gauge at one vertex: two signs and one spin flip, F kept

(c) chain: τ (below the sites) maps J to $|J|$

Figure 1.1: Frustration and gauge. Solid edges are positive couplings, dashed negative; white and black discs are ± 1 spins. (a) A frustrated square ($F = -1$): in a ground state exactly one edge (red) is unsatisfied, and it can be any of the four. (b) Flipping the gauge at one corner flips the two incident couplings and the spin there, so energies and the frustration F are unchanged. (c) On a chain, or any tree, a gauge transformation removes every sign.

Definition 1.8 (Gibbs states)**.** Fix J and β. A probability measure μ on $\{-1,+1\}^{\mathbb{Z}^d}$ is an *infinite-volume Gibbs state* if for every finite Λ the conditional law of σ_Λ given the spins off Λ is $\mu^{\beta,\eta}_{\Lambda,J}$ evaluated at $\eta = \sigma$, μ-almost surely. The set of Gibbs states is denoted $\mathcal{G}_J(\beta)$. Its extreme points are the *pure* states; non-extremal states are *mixed*.

We use the following standard facts [95, Lemma 6.27 and Theorems 6.26, 6.56, 6.58, 6.63, 6.72], [106]. $\mathcal{G}_J(\beta)$ is a nonempty convex set, compact for the topology of local convergence. A Gibbs state is pure if and only if it is trivial on the tail σ-field, equivalently if its correlations between a fixed local event and events far away vanish. Every Gibbs state is a unique mixture of pure states, and every pure state μ is the local limit of $\mu^{\beta,\eta}_{\Lambda_L,J}$ for μ-almost every η. The global flip maps $\mathcal{G}_J(\beta)$ to itself, so pure states come in flip-related pairs (possibly coinciding). Finally [171, Proposition 3.1], there is a unique Gibbs state if and only if, for every two sequences of fixed boundary conditions $\eta^{(L)}$ and $\eta'^{(L)}$, the measures $\mu^{\beta,\eta^{(L)}}_{\Lambda_L,J}$ and $\mu^{\beta,\eta'^{(L)}}_{\Lambda_L,J}$ have the same local limit.

For random couplings, Newman introduces two critical inverse temperatures [171, (3.58)–(3.59) and (4.1)–(4.2)]:

$$\begin{aligned} \beta_c(d,\nu) &= \sup\{\beta' : \text{ a.s. } |\mathcal{G}_J(\beta)| = 1 \text{ for all } \beta \in [0,\beta']\}, \\ \bar{\beta}_c(d,\nu) &= \sup\{\beta : \text{ a.s. } |\mathcal{G}_J(\beta)| = 1\}. \end{aligned} \tag{1.5}$$

Both are deterministic and $\beta_c \le \bar{\beta}_c$; for ferromagnets they coincide, but for spin glasses this is not known [171, p. 47]. The number of pure states is a.s. constant (Chapter 3).

Spin-glass finite-volume states have no monotonicity, and Newman and Stein [172] (see also [171, pp. 54–55]) argued that when many pure states compete, those with a fixed coupling-independent boundary condition need not converge at all (*chaotic size dependence*); metastates, which tame this, are the subject of Chapter 3. Here infinite-volume statements are uniqueness statements or statements about finite-volume averages.

1.4 What order would mean

For symmetric couplings with coupling-independent boundary conditions every averaged correlation vanishes (1.2), so order must be detected by gauge-invariant quantities. Edwards and Anderson [81] proposed to measure whether spins are frozen in *some* pattern, through the squared correlation $\langle \sigma_0 \sigma_x \rangle_J^2$ or the agreement between two independent copies.

Definition 1.9 (Overlaps, EA order, spin-glass susceptibility)**.** Let σ, τ be two *replicas*: independent samples from the same Gibbs measure at fixed J (we say each time which measure). For finite Λ the *site overlap* and *edge overlap* are

$$q_\Lambda = \frac{1}{|\Lambda|}\sum_{x\in\Lambda}\sigma_x\tau_x, \qquad q^e_\Lambda = \frac{1}{|E(\Lambda)|}\sum_{xy\in E(\Lambda)}\sigma_x\sigma_y\tau_x\tau_y.$$

For a given state or boundary prescription, *EA order* means that $\mathbb{E}\langle\sigma_0\sigma_x\rangle_J^2$ does not tend to 0 as $|x|\to\infty$. For a measurable choice $J\mapsto\alpha_J$ of pure states the *EA order parameter* is $q_{\mathrm{EA}} = \mathbb{E}\,\alpha_J(\sigma_0)^2$, as in Newman's discussion of the two-state picture [171, (4.10)]. The *spin-glass susceptibility* is $\chi_{\mathrm{SG}} = \sum_x \mathbb{E}\langle\sigma_0\sigma_x\rangle_J^2$, and its finite-volume version is $\chi_{\mathrm{SG}}(\Lambda) = |\Lambda|^{-1}\sum_{x,y\in\Lambda}\mathbb{E}\langle\sigma_x\sigma_y\rangle_J^2$.

The two descriptions are linked by an identity that we use throughout. Since τ is an independent copy, $\langle\sigma_x\tau_x\sigma_y\tau_y\rangle = \langle\sigma_x\sigma_y\rangle\langle\tau_x\tau_y\rangle$, and so

$$\langle q_\Lambda^2\rangle = \frac{1}{|\Lambda|^2}\sum_{x,y\in\Lambda}\langle\sigma_x\sigma_y\rangle_J^2, \qquad \mathbb{E}\langle q_\Lambda^2\rangle = \frac{\chi_{\mathrm{SG}}(\Lambda)}{|\Lambda|}. \tag{1.6}$$

In particular a uniformly bounded finite-volume susceptibility forces $\mathbb{E}\langle q_\Lambda^2\rangle\to 0$: finite χ_{SG} is the cleanest certificate of the high-temperature phase, and Chapter 9 proves bounds of this kind.

"Order" can be made precise in several ways not known to be equivalent, and different results of the review prove different ones. We fix three targets, to which later chapters refer by name.

Definition 1.10 (Three order-parameter targets)**.** Fix d, a symmetric law ν with $\mathbb{E}|J_e|<\infty$, and β.

(T1) *Pressure target.* Let $\Psi^{(2)}_\Lambda(\beta,\lambda) = \mathbb{E}\log\sum_{\sigma,\tau}\exp\bigl(-\beta H^{\mathrm{f}}_{\Lambda,J}(\sigma)-\beta H^{\mathrm{f}}_{\Lambda,J}(\tau)+\lambda\sum_{x\in\Lambda}\sigma_x\tau_x\bigr)$ be the two-replica quenched log-partition function with free boundary conditions and replica coupling λ, and $p^{(2)}(\beta,\lambda) = \lim_L |\Lambda_L|^{-1}\Psi^{(2)}_{\Lambda_L}(\beta,\lambda)$ (the limit exists by Corollary 1.13). The target is a cusp at $\lambda=0$: $\partial^+_\lambda p^{(2)}(\beta,0) > 0$.

(T2) *Stationary target.* Let $\mathbf{P}$ be a law of (J,σ,τ) whose J-marginal is $\nu^{\otimes E_d}$, which is invariant under the joint translations of $\mathbb{Z}^d$, and under which, given J, (σ,τ) has law $\int\Gamma\otimes\Gamma\,\kappa_J(\mathrm{d}\Gamma)$ for some probability measure κ_J on $\mathcal{G}_J(\beta)$ (a single state, or a metastate as in Chapter 3). By the ergodic theorem $Q = \lim_L q_{\Lambda_L}$ exists $\mathbf{P}$-a.s. The target is $\mathbf{E}[Q^2] > 0$ for the stated $\mathbf{P}$.

(T3) *Finite-volume target.* For a stated coupling-independent boundary prescription b (free, periodic, ...), with $\langle\cdot\rangle$ the product of two copies of $\mu^{\beta,b}_{\Lambda_L,J}$: $\liminf_{L\to\infty}\mathbb{E}\langle q^2_{\Lambda_L}\rangle > 0$.

The review does not identify the three targets silently: Corollary 1.13 below proves that (T3) with free boundary conditions implies (T1), and other bridges appear where they are used (Chapter 8). None of the targets distinguishes spin-glass from ferromagnetic order (the low-temperature Ising ferromagnet satisfies all three); for symmetric ν this is harmless by (1.2), but for biased laws (Chapter 14) the two must be told apart. And (T2), (T3) refer to a specified state or boundary prescription: by Example 1.6 a state can have decaying correlations while its pure components are ordered.

1.5 The free energy: existence and self-averaging

For ferromagnets the free energy per site has a limit by the first Griffiths inequality: adding a ferromagnetic bond can only increase the pressure, which is therefore superadditive. Contucci and Lebowitz observed that a quenched version of this monotonicity holds for centred random couplings, for a reason that has nothing to do with ferromagnetism: convexity. We write Ψ for quenched log-partition functions (the "quenched pressure" of [63]; the free energy per site is $-\Psi/(\beta|\Lambda|)$).

Lemma 1.11 (Contucci–Lebowitz [63, Theorem 1 and the remark after it])**.** *Let Λ be finite, $\mathcal{X}$ a finite family of subsets of Λ, $(J_X)_{X\in\mathcal{X}}$ independent random variables with $\mathbb{E}J_X = 0$ and $\mathbb{E}|J_X| < \infty$, and $\bar U$ a non-random function of $\sigma \in \{-1,+1\}^\Lambda$. For $\lambda \in \mathbb{R}^{\mathcal{X}}$ set*

$$\Psi(\lambda) = \mathbb{E}\log \sum_{\sigma} \exp\Big(\sum_{X\in\mathcal{X}} \lambda_X J_X \sigma_X - \bar U(\sigma)\Big).$$

Then for each $A \in \mathcal{X}$, with the other λ_X fixed, $\lambda_A \mapsto \Psi$ is convex on $\mathbb{R}$ and nondecreasing on $[0,\infty)$; equivalently $\partial\Psi/\partial\lambda_A = \mathbb{E}[J_A\langle\sigma_A\rangle] \geq 0$ for $\lambda_A \geq 0$.

Proof. For fixed J, $\lambda_A \mapsto \log Z$ is convex, with second derivative $J_A^2(1-\langle\sigma_A\rangle^2) \geq 0$, and its first derivative $J_A\langle\sigma_A\rangle$ is bounded by $|J_A|$; so Ψ is convex and differentiable with $\partial_{\lambda_A}\Psi = \mathbb{E}[J_A\langle\sigma_A\rangle]$. At $\lambda_A = 0$ the Gibbs average $\langle\sigma_A\rangle$ does not involve J_A, so by independence the derivative equals $\mathbb{E}[J_A]\,\mathbb{E}[\langle\sigma_A\rangle|_{\lambda_A=0}] = 0$. A convex function whose derivative vanishes at 0 is nondecreasing on $[0,\infty)$. □

The proof uses only that σ_A is bounded and enters the exponent linearly in $\lambda_A J_A$, so it applies verbatim to any bounded $f_A(\sigma)$ in place of σ_A. The positive quantity is $\mathbb{E}[J_A\langle\sigma_A\rangle]$, not $\mathbb{E}\langle\sigma_A\rangle$, which has no sign (it vanishes for symmetric couplings).

Theorem 1.12 (Quenched free energy)**.** *Let ν be symmetric with $\mathbb{E}|J_e| < \infty$, $d \geq 1$, $\beta \geq 0$, and $\Psi_\Lambda(\beta) = \mathbb{E}\log Z^{\mathrm{f}}_{\Lambda,J}(\beta)$.*

(a) (Superadditivity [63, Theorem 2].) *If Λ is the disjoint union of $\Lambda_1,\dots,\Lambda_K$, then $\Psi_\Lambda \geq \sum_h \Psi_{\Lambda_h}$.*

(b) (Limit [63, Corollary 2.1].) *With $B_\ell = \{1,\dots,\ell\}^d$, the limit $p(\beta) = \lim_{L\to\infty}|\Lambda_L|^{-1}\Psi_{\Lambda_L}(\beta)$ exists and equals $\sup_{\ell\geq 1}\ell^{-d}\Psi_{B_\ell}(\beta)$. Periodic boundary conditions give the same limit.*

(c) (Self-averaging.) *If in addition $\mathbb{E}J_e^2 < \infty$, then $|\Lambda_L|^{-1}\log Z^{\mathrm{f}}_{\Lambda_L,J}(\beta) \to p(\beta)$ for ν-almost every J and in L^2.*

Almost-sure existence of the quenched free energy is classical (Vuillermot, van Enter and van Hemmen [238, 240]). The convexity proof of (a)–(b) needs no estimate at all, and (c) is a short second-moment argument; we deduce (a) directly from Lemma 1.11, where Contucci and Lebowitz phrase the same step as an interpolation.

Proof. (a) Apply Lemma 1.11 with $\mathcal{X}$ the edges of $E(\Lambda)$, $\bar U = 0$ and the coupling strength β written as λ_e. Lowering λ_e from β to 0 on the edges that join different Λ_h decreases Ψ, and at $\lambda_e = 0$ on those edges the partition function factorizes over the blocks.

(b) and (c) are routine and are proved in Section A.1: (b) combines (a) for disjoint translates of a box with the a priori bounds $|R|\log 2 \leq \log Z^{\mathrm{f}}_{R,J} \leq |R|\log 2 + \beta\sum_{e\in E(R)}|J_e|$, and (c) follows from the Efron–Stein inequality, since resampling one coupling changes $\log Z$ by at most $\beta|J_e - J'_e|$. □

Corollary 1.13 (From finite-volume overlap to a pressure cusp; after van Enter–Griffiths [237]). *Let ν be symmetric with $\mathbb{E}|J_e| < \infty$ and $\beta \geq 0$. For every $\lambda \in \mathbb{R}$ the limit $p^{(2)}(\beta,\lambda) = \lim_L |\Lambda_L|^{-1}\Psi^{(2)}_{\Lambda_L}(\beta,\lambda)$ of* (T1) *exists; it is even and convex in λ, and*

$$\partial^+_\lambda p^{(2)}(\beta,0) \;\geq\; \limsup_{\ell\to\infty} \mathbb{E}\langle |q_{B_\ell}|\rangle^{\mathrm{f}}_{B_\ell} \;\geq\; \limsup_{\ell\to\infty} \mathbb{E}\langle q^2_{B_\ell}\rangle^{\mathrm{f}}_{B_\ell}, \tag{1.7}$$

where $\langle\cdot\rangle^{\mathrm{f}}_{B_\ell}$ is the product of two copies of $\mu^{\beta,\mathrm{f}}_{B_\ell,J}$. In particular the finite-volume target (T3) *with free boundary conditions implies the pressure target* (T1).

Proof sketch. This sketch omits the bookkeeping, which is in Section A.2. The two-replica weight has the form of Lemma 1.11 (bounded functions $\sigma_x\sigma_y+\tau_x\tau_y$ times the centred couplings, site-local reference term $-\lambda\sum_x \sigma_x\tau_x$), so the proof of Theorem 1.12 gives the limit as $\sup_\ell \ell^{-d}\Psi^{(2)}_{B_\ell}(\beta,\lambda)$. On one box, flip symmetry and Jensen's inequality give $\langle e^{\lambda\ell^d q}\rangle \geq \frac{1}{2}e^{\lambda\ell^d\langle|q|\rangle}$, so one box with a large overlap already lifts the supremum at $\lambda > 0$ linearly in λ. □

The converse of Corollary 1.13 is not known: convex functions with smooth minima can converge to a function with a cusp, so a pressure cusp alone carries no finite-volume information at $\lambda = 0$. None of this monotonicity extends to correlations (Section 1.7).

1.6 High temperature: comparison with the ferromagnet

At high temperature one expects a unique Gibbs state whatever the signs. The cleanest proof compares the spin glass with the ferromagnet of the same coupling magnitudes through the random-cluster representation of Fortuin and Kasteleyn [90], set up for couplings of arbitrary sign by Kasai and Okiji, Swendsen and Wang, and Edwards and Sokal [80, 132, 230], as presented in [170, 171]; Chapter 7 develops it further.

Fix J, β, a finite Λ and a boundary condition η. Let $E^*_\Lambda = E(\Lambda)\cup\partial_e\Lambda$ and $p_e = 1-e^{-2\beta|J_e|}$; edges with $J_e = 0$ have $p_e = 0$, are always closed, and can be discarded. On $\{-1,+1\}^\Lambda\times\{0,1\}^{E^*_\Lambda}$, with $\sigma = \eta$ off Λ, let

$$P^\eta_{\Lambda,J}(\sigma,\omega) \;\propto\; \prod_{e\in E^*_\Lambda}\Big[(1-p_e)\,\mathbf{1}\{\omega_e=0\} + p_e\,\mathbf{1}\{\omega_e=1,\ e \text{ satisfied by } \sigma\}\Big]. \tag{1.8}$$

Open edges are allowed only where the spins satisfy the coupling. Let $k_\Lambda(\omega)$ be the number of ω-clusters (on the vertex set $\Lambda\cup\partial\Lambda$) that do not meet $\partial\Lambda$, and let U^η_J be the set of ω for which some σ with $\sigma = \eta$ off Λ satisfies every open edge; a configuration outside U^η_J contains an open frustrated cycle, or an open path between boundary sites whose sign product disagrees with η.

Lemma 1.14 (Frustrated Edwards–Sokal coupling [171, Propositions 3.2–3.3]).

(i) The σ-marginal of $P^\eta_{\Lambda,J}$ is the Gibbs measure $\mu^{\beta,\eta}_{\Lambda,J}$.

(ii) The ω-marginal is $\varphi^\eta_{\Lambda,J}(\omega) \propto \prod_e p_e^{\omega_e}(1-p_e)^{1-\omega_e}\,2^{k_\Lambda(\omega)}\,\mathbf{1}\{\omega\in U^\eta_J\}$.

(iii) Given ω, the spins of clusters meeting $\partial\Lambda$ are determined by η and the signs of J, and every other cluster independently takes one of its two J-compatible colourings with probability $\frac{1}{2}$. Given σ, the edges are independent: each satisfied edge is open with probability p_e, each unsatisfied edge is closed.

For the ferromagnet $|J|$ with $\eta \equiv +1$ the indicator in (ii) is identically 1; this is the wired *random-cluster measure, denoted $\varphi^{\mathrm{w}}_{\Lambda,|J|}$.*

Proof. Summing (1.8) over ω, a satisfied edge contributes $(1-p_e)+p_e = 1$ and an unsatisfied edge contributes $1-p_e = e^{-2\beta|J_e|}$; since $J_e\sigma_x\sigma_y = |J_e|$ or $-|J_e|$ according as e is satisfied or not, the product is a constant times the Gibbs weight. Summing over σ for fixed $\omega \in U^\eta_J$, the relative signs inside each cluster are fixed by J and the clusters meeting $\partial\Lambda$ are fixed by η, which leaves $2^{k_\Lambda(\omega)}$ choices; for $\omega \notin U^\eta_J$ there is none. Part (iii) is read off from the product form. □

For free boundary conditions (iii) gives the representation $\langle\sigma_x\sigma_y\rangle^{\mathrm{f}}_{\Lambda,J} = \mathbb{E}_\varphi[\mathrm{sgn}_J(x,y;\omega)]$, where the *relative parity* $\mathrm{sgn}_J(x,y;\omega)$ is the sign product of J along any open path from x to y (well defined because $\omega \in U$) and 0 if $x \not\leftrightarrow y$ [171, (3.10)]. For a ferromagnet this is the connection probability; for a spin glass the signs can cancel, so percolation of ω does not by itself imply order (Chapter 7). The marginal (ii) is the wired ferromagnetic random-cluster measure *conditioned* on the decreasing event U^η_J (closing edges removes constraints; Fig. 1.2(a)). This is the whole mechanism of the comparison.

Lemma 1.15 (Domination [171, Proposition 3.6 and (3.25)])**.** *With the notation above:*

(i) $\varphi^{\mathrm{w}}_{\Lambda,|J|}$ satisfies the FKG lattice condition, and $\varphi^{\mathrm{w}}_{\Lambda,|J|} \preccurlyeq \bigotimes_e \mathrm{Bernoulli}(p_e)$;

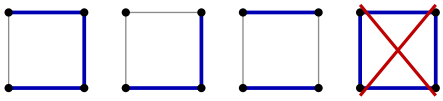
(ii) for every set $E' \subseteq E^*_\Lambda$ of edges and configurations $\xi \le \xi'$ on E' with $\varphi^\eta_{\Lambda,J}(\omega_{E'} = \xi) > 0$, $\varphi^\eta_{\Lambda,J}(\,\cdot \mid \omega_{E'} = \xi) \preccurlyeq \varphi^{\mathrm{w}}_{\Lambda,|J|}(\,\cdot \mid \omega_{E'} = \xi')$ on the remaining edges. In particular $\varphi^\eta_{\Lambda,J} \preccurlyeq \varphi^{\mathrm{w}}_{\Lambda,|J|}$.

Proof. (i) is Newman's Proposition 3.6 and (3.25): the density of φ^{w} with respect to the Bernoulli measure is proportional to 2^{k_Λ}, which is decreasing, and the lattice condition is checked edge pair by edge pair. (ii) A strictly positive measure with the lattice condition is *monotonic*: its conditional laws given the values on a set of edges are stochastically increasing in those values [109, Theorem 2.27]. So $\varphi^{\mathrm{w}}(\cdot \mid \xi) \preccurlyeq \varphi^{\mathrm{w}}(\cdot \mid \xi')$, and $\varphi^{\mathrm{w}}(\cdot \mid \xi)$ is positively associated. By (ii) of Lemma 1.14, $\varphi^\eta(\cdot \mid \xi)$ is $\varphi^{\mathrm{w}}(\cdot \mid \xi)$ conditioned on the decreasing event $\{\omega : (\xi,\omega) \in U^\eta_J\}$, and conditioning a positively associated measure on a decreasing event lowers it stochastically [171, Proposition 3.4]. □

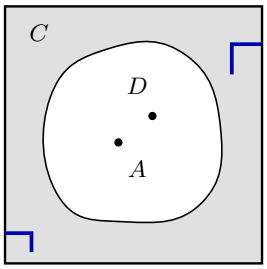
Theorem 1.16 (Newman [170, Theorem 2], [171, Theorem 3.11])**.** *Let $J = (J_e)_{e\in E_d}$ be any real couplings, not necessarily random, and $\beta \ge 0$. If the ferromagnet with couplings $|J|$ has a unique infinite-volume Gibbs state at β, then so does J. More precisely, for finite $A \subset \Lambda$ and all boundary conditions η, η',*

$$\left|\langle\sigma_A\rangle^\eta_{\Lambda,J} - \langle\sigma_A\rangle^{\eta'}_{\Lambda,J}\right| \;\le\; 2\sum_{x\in A} \varphi^{\mathrm{w}}_{\Lambda,|J|}(x \leftrightarrow \partial\Lambda). \tag{1.9}$$

Proof sketch. This sketch omits the construction of the two bond measures on one probability space (Section A.3). Two boundary conditions can influence σ_A only through open paths of the frustrated cluster measure, and by Lemma 1.15 these lie inside the clusters of the wired ferromagnet, which do not reach A from far away in its uniqueness phase. Explore the wired cluster C of $\partial\Lambda$ edge by edge, revealing each edge in both bond measures with a common uniform variable; by Lemma 1.15(ii) the frustrated configuration stays below the wired one, so every revealed edge from C to $D = \Lambda \setminus C$ is closed in both. Given what was revealed, σ_D has the free law $\mu^{\beta,\mathrm{f}}_{D,J}$ whatever η, and the law of D does not depend on η (Fig. 1.2(b)); so η affects $\langle\sigma_A\rangle$



(a) frustrated square: any proper subset of edges may be open, the full cycle may not



(b) C: wired cluster of $\partial\Lambda$; inside D the state is free

Figure 1.2: The comparison behind Theorem 1.16. (a) For a frustrated square the Edwards–Sokal bond measure is the ferromagnetic random-cluster measure with the configurations containing an open frustrated cycle removed, a decreasing event; so the frustrated measure is stochastically smaller. (b) Exploring the wired ferromagnetic cluster C of the boundary (grey) also explores the frustrated one inside it; on the unexplored region D the conditional state is free whatever the boundary condition, which can therefore affect σ_A only if C reaches A.

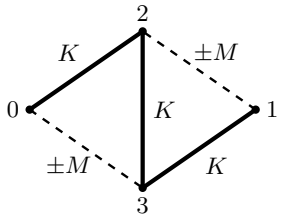
only on $\{A \cap C \neq \emptyset\}$, which gives (1.9). Integrating it against two Gibbs states, and using that $\varphi^{\mathrm{w}}_{\Lambda,|J|}(x \leftrightarrow \partial\Lambda)$ decreases to the plus-state magnetization of the ferromagnet, which vanishes in its uniqueness phase, gives uniqueness. □

For random couplings one combines Theorem 1.16 with the comparison of the wired measure with Bernoulli percolation, averaged over the couplings [9]. Let $p_c(d)$ be the critical probability of Bernoulli bond percolation on $\mathbb{Z}^d$, with $p_c(1) = 1$.

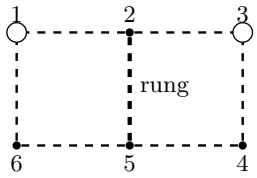
Corollary 1.17 (High-temperature uniqueness [171, Theorem 3.11, Proposition 4.1])**.** *Let ν be any law on $\mathbb{R}$ (not necessarily symmetric), $d \geq 1$, and $\beta \geq 0$ with $\mathbb{E}(1 - e^{-2\beta|J_e|}) < p_c(d)$. Then for ν-almost every J the Gibbs state is unique. Hence $\beta_c(d,\nu) > 0$ for every law; $\beta_c(1,\nu) = \infty$; $\beta_c(d,\nu) = \infty$ if $\nu(\mathbb{R} \setminus \{0\}) \leq p_c(d)$; and for $\pm J$ couplings $\beta_c(d)$ is at least the critical inverse temperature of the Ising ferromagnet on $\mathbb{Z}^d$.*

Proof. By Lemma 1.15(i), $\varphi^{\mathrm{w}}_{\Lambda,|J|}(x \leftrightarrow \partial\Lambda) \leq P_{(p_e)}(x \leftrightarrow \partial\Lambda)$, where $P_{(p_e)}$ is independent percolation with parameters $p_e = p_e(J)$. Averaging over J gives independent percolation with parameter $\mathbb{E}p_e < p_c(d)$, under which $x \leftrightarrow \partial\Lambda_L$ has probability tending to 0; since these events decrease in L, for almost every J the limit of $\varphi^{\mathrm{w}}_{\Lambda_L,|J|}(x \leftrightarrow \partial\Lambda_L)$ vanishes for every x, and (1.9) gives uniqueness as in the proof of Theorem 1.16. The percolation bound is nondecreasing in β, so one null set of couplings serves all smaller β at once, as the definition of β_c requires. The remaining statements follow: $\mathbb{E}p_e \to 0$ as $\beta \to 0$; $\mathbb{E}p_e < 1$ for every $\beta < \infty$; $\mathbb{E}p_e < \nu(\mathbb{R} \setminus \{0\})$ unless $\nu = \delta_0$, when $\mathbb{E}p_e = 0$; and for $\pm J$ the associated ferromagnet is the Ising model, to which the first part of Theorem 1.16 applies directly. □

For $\pm J$ this gives uniqueness for $T > T_c^{\mathrm{Ising}}(d)$ (≈ 2.269 in $d = 2$, ≈ 4.51 in $d = 3$), far above the believed spin-glass critical temperatures (Chapter 2). De Santis showed that the $\pm J$ uniqueness region is strictly larger [70], and other methods reach further (Chapter 9).



(a) ferromagnetic path 0–2–3–1, fair chords: $\mathbb{E}\langle\sigma_0\sigma_1\rangle < 0$



(b) adding the fair rung 25 lowers $\mathbb{E}\langle\sigma_1\sigma_3\rangle^2$

Figure 1.3: Two small counterexamples to ferromagnetic intuition. Thick solid edges are deterministic ferromagnetic couplings, dashed edges carry independent fair signs. (a) Hard random chords on a ferromagnetic path reverse the quenched correlation of the endpoints (Proposition 1.21). (b) On the 2×3 box, where every edge carries a fair sign, inserting the middle rung (thick) decreases the quenched squared correlation of the circled corners (Proposition 1.19); on this graph the mixed pressure derivative of Contucci and Lebowitz also changes sign (Proposition 1.18).

1.7 Why the ferromagnetic toolbox fails

Beyond high temperature, the theory of the Ising ferromagnet rests on positivity: the second Griffiths inequality (giving monotonicity in the volume and comparisons between lattices), FKG for spins and random-cluster edges, and reflection positivity (infrared bounds). For the spin glass, Lemma 1.11 is a true analogue of the *first* Griffiths inequality, but only for the pressure; the natural quenched analogues of the others are false, already on one or two plaquettes.

1.7.1 No second Griffiths inequality

Contucci and Lebowitz asked whether the mixed derivative $\partial^2\Psi/\partial\lambda_A\partial\lambda_B = \mathbb{E}\big[J_AJ_B(\langle\sigma_A\sigma_B\rangle - \langle\sigma_A\rangle\langle\sigma_B\rangle)\big]$, in the setting of Lemma 1.11, has a definite sign for $A \neq B$; their computations suggested it is negative [63, p. 1466]. For independent $N(0,1)$ couplings J_X and arbitrary $\bar{U}$ this is a statement about squared correlations: Gaussian integration by parts, $\mathbb{E}[J_Bg(J_B)] = \mathbb{E}[g'(J_B)]$, applied to $g = \langle\sigma_B\rangle$ with $\partial\langle\sigma_B\rangle/\partial J_B = \lambda_B(1 - \langle\sigma_B\rangle^2)$, gives $\partial\Psi/\partial\lambda_B = \lambda_B\,\mathbb{E}[1 - \langle\sigma_B\rangle^2]$, and differentiating in λ_A,

$$\frac{\partial^2\Psi}{\partial\lambda_A\,\partial\lambda_B} = -\lambda_B\,\frac{\partial}{\partial\lambda_A}\mathbb{E}\langle\sigma_B\rangle^2 \qquad (A \neq B).$$

So, for Gaussian couplings, a negative sign for $A \neq B$ would say that adding any interaction increases every quenched squared correlation, and in particular, as Contucci and Lebowitz note, that squared correlations increase with the volume. For $\pm J$ couplings both the sign of the mixed derivative and the monotonicity of squared correlations under adding bonds fail, on the smallest box of $\mathbb{Z}^2$ that contains two plaquettes.

Proposition 1.18 (Contucci–Unguendoli–Vernia [64, pp. 3–5])**.** *Let G_6 be the six-cycle $1, 2, \dots, 6$ with the chord 25 (the 2×3 box of $\mathbb{Z}^2$), with independent fair ±1 couplings. Then $\mathbb{E}\big[J_{12}J_{23}(\langle\sigma_1\sigma_3\rangle - \langle\sigma_1\sigma_2\rangle\langle\sigma_2\sigma_3\rangle)\big]$ is negative for small β and positive for large β; it changes sign near $\beta \approx 0.695$.*

Contucci, Unguendoli and Vernia compute the expectation in closed form by gauge-fixing all couplings except J_{12}, J_{23} (we confirmed the sign change numerically by enumerating all

configurations); their simulations of $\mathbb{E}[J_e\langle\sigma_x\sigma_y\rangle]$ on boxes $\{1,\dots,L\}^d$, $d=2,3$, for $\pm J$ and Gaussian couplings, show oscillation in L [64, Figs. 1–2] (numerical evidence). For squared correlations the failure can be exhibited directly.

Proposition 1.19 (Adding a fair bond can lower a squared correlation)**.** *On the* 2×3 *box of Fig. 1.3(b), with vertices* $1,2,3$ *on top and* $6,5,4$ *below, give the six perimeter edges independent fair* ±1 *couplings and the rung* 25 *an independent fair sign times* $\lambda\ge 0$*. Let* $t=\tanh\beta$*,* $u=\tanh\beta\lambda$*,* $Q(t,u)=\mathbb{E}\langle\sigma_1\sigma_3\rangle^2$*, and* $A=t^2+t^4$*,* $B=2t^3$*,* $C=1+t^6$*. Then*

$$Q(t,u)=\frac{1}{4}\Big(\frac{A+Bu}{C+Bu}\Big)^2+\frac{1}{4}\Big(\frac{A-Bu}{C-Bu}\Big)^2+\frac{1}{2}\Big(\frac{t^4-t^2}{1-t^6}\Big)^2. \tag{1.10}$$

In particular, at $t^2=\frac{1}{2}$ *a rung of the same strength as the perimeter edges* ($u=t$) *gives* $Q(t,t)-Q(t,0)=\frac{47356}{207025}-\frac{116}{441}=-\frac{1304}{38025}<0$*; and whenever* $\tanh^2\beta>2-\sqrt{3}$*, a sufficiently weak rung already decreases* Q*.*

The proof gauges a spanning tree to positive signs, which leaves two fair signs on the edges 12 and 23; the high-temperature expansions of Z and of $Z\langle\sigma_1\sigma_3\rangle$ then have four terms each, and the rest is algebra, given in Section A.4.

Route warning 1.20 (Monotone restoration of bonds)*.* A tempting strategy is to prove order on a convenient subgraph and then restore the deleted couplings, arguing that adding independent symmetric bonds can only increase $\mathbb{E}\langle\sigma_x\sigma_y\rangle^2$. Proposition 1.19 shows that this is false, already on a planar box with free boundary conditions; so no representation of $\mathbb{E}\langle\sigma_x\sigma_y\rangle^2$ as a connection probability can be monotone under adding symmetric bonds. The example is for ±1 couplings (and nearby symmetric laws, by continuity); we know no Gaussian counterexample.

1.7.2 No quenched positivity with a ferromagnetic reference

Contucci and Lebowitz also asked [63, p. 1466] whether, for a ferromagnetic reference interaction $\bar U$, adding symmetric random interactions preserves the first Griffiths inequality in quenched form, $\mathbb{E}\langle\sigma_A\rangle\ge 0$; they wrote that this "seems very reasonable to expect". It fails on four spins.

Proposition 1.21 (A negative answer on the diamond)**.** *On the vertices* $0,1,2,3$ *let the path* 0–2–3–1 *carry the deterministic ferromagnetic coupling* $K>0$ *and the chords* 03 *and* 12 *carry* MJ_{03} *and* MJ_{12} *with independent fair signs* J_{03},J_{12} *(Fig. 1.3(a)). Let* $t=\tanh\beta K$*. Then*

$$\lim_{M\to\infty}\mathbb{E}\langle\sigma_0\sigma_1\rangle=-\frac{2t^3}{1+3t^2}<0,$$

so $\mathbb{E}\langle\sigma_0\sigma_1\rangle<0$ *for all sufficiently large* M*, at every* $\beta>0$*. For instance, at* $e^{2\beta K}=3$ *and* $e^{2\beta M}=9$*,* $\mathbb{E}\langle\sigma_0\sigma_1\rangle=-461/13832$*. The same holds, by continuity, if each chord is given a small positive mean.*

Proof. As $M\to\infty$ each chord forces its endpoints: $\sigma_3=J_{03}\sigma_0$ and $\sigma_2=J_{12}\sigma_1$. If both chords are positive, the three path edges each contribute $\sigma_0\sigma_1$ to the energy and $\langle\sigma_0\sigma_1\rangle\to\tanh(3\beta K)$. In each of the other three sign patterns the path contributes $\sigma_0\sigma_1$ once and $-\sigma_0\sigma_1$ twice, so $\langle\sigma_0\sigma_1\rangle\to -t$. Averaging, $\frac{1}{4}(\tanh 3\beta K-3t)$, and $\tanh 3x=(3t+t^3)/(1+3t^2)$ gives the limit. The finite instance is an exact evaluation of the 16×4 terms. Finally $\mathbb{E}\langle\sigma_0\sigma_1\rangle$ is a continuous function of the atoms of the chord law, so a small shift of both atoms keeps it negative. □

A single edge cannot reverse the sign: for a law symmetric about a nonnegative mean, $\mathbb{E}\tanh(\beta J) \geq 0$, because $\tanh a + \tanh b$ has the sign of $a + b$. On the Nishimori line positivity holds for a different reason, $\mathbb{E}\langle\sigma_A\rangle = \mathbb{E}\langle\sigma_A\rangle^2$ (Chapter 14).

Open problem 1.22 (Quenched positivity for iid biased couplings)**.** For iid biased $\pm J$ couplings (ν_p with $p > \frac{1}{2}$), or Gaussian couplings with positive mean, on boxes $\Lambda_L \subset \mathbb{Z}^d$ with free boundary conditions, is $\mathbb{E}\langle\sigma_x\sigma_y\rangle \geq 0$ for all x, y?

1.7.3 No FKG for the averaged overlap

A natural replacement for spin positivity is the coupling-averaged law of the overlap field $q_x = \sigma_x\tau_x$, gauge invariant and often treated heuristically as an attractive Ising-like field. Its lattice condition fails on a single square, at every temperature.

Proposition 1.23 (The averaged overlap law is not lattice-FKG)**.** *On the free square with vertices $1, 2, 3, 4$ in cyclic order and independent fair ± 1 couplings, let $P(q)$ be the law of $q = (\sigma_x\tau_x)_x$ under $\mathbb{E}[\mu_J^{\beta,\mathrm{f}} \otimes \mu_J^{\beta,\mathrm{f}}]$ and $t = \tanh\beta$. Writing configurations in the order (q_1, q_2, q_3, q_4),*

$$P(++++)\,P(-+-+) - P(+++-)\,P(-+++) = -\frac{t^8(1+t^8)}{8(1+t^2)^4(1+t^4)^4} < 0 \qquad (\beta > 0),$$

so P violates the FKG lattice condition for the pair $(++-+)$, $(-+++)$.

Proof. Let $F = \prod_e J_e$ and $q_e = q_x q_y$ for $e = xy$. Up to a common factor, $e^{\beta J\sigma_x\sigma_y}e^{\beta J\tau_x\tau_y} \propto (1 + t^2 q_e) + tJ_e\sigma_x\sigma_y(1 + q_e)$. Summing over σ, only the empty set and the full cycle survive, and $\prod_e(1 + q_e) = 16 \cdot \mathbf{1}\{q \text{ constant}\}$; with (1.4) this gives

$$P(q) = \mathbb{E}_F\,\frac{\prod_e(1 + t^2 q_e) + 16F\,t^4\,\mathbf{1}\{q \text{ constant}\}}{16(1 + Ft^4)^2} = S\prod_e(1 + t^2 q_e) + R\,\mathbf{1}\{q \text{ constant}\},$$

with $S = \frac{1+t^8}{16(1-t^8)^2} > 0$ and $R = -\frac{2t^8}{(1-t^8)^2} < 0$; R is negative because the frustrated square, with its smaller partition function, receives the larger weight Z^{-2}. The four configurations in the minor have $\prod_e(1 + t^2 q_e)$ equal to $(1 + t^2)^4$, $(1 - t^2)^4$, $(1 - t^4)^2$, $(1 - t^4)^2$, and only $(++++)$ is constant. The product terms cancel exactly, leaving $R(1 - t^2)^4 S$, which is the stated value since $1 - t^8 = (1 - t^2)(1 + t^2)(1 + t^4)$. □

Without the frustration correction R the law would have product form, and the minor would vanish. The same square law is nevertheless positively associated at every $\beta > 0$ (computer-assisted, Appendix B: the covariances of all 13 861 pairs of nontrivial increasing events are checked in exact arithmetic), so the failure concerns conditional positivity, which is what attractive comparison arguments use. The same mechanism breaks bond reflection positivity of the averaged two-replica law (six-cycle, every $\beta > 0$; 2×4 box, $\beta = \log 2$; Chapter 15).

1.8 Where the threads go

The quenched free energy exists and self-averages; the Gibbs state is unique at high temperature; and the second Griffiths inequality, quenched positivity with a ferromagnetic reference, monotonicity of squared correlations under adding bonds and the FKG lattice condition for the averaged overlap law all fail on graphs with at most two plaquettes. The central question remains open.

Open problem 1.24 (The central question; Newman's Research Problems 4.2 and 4.3 [171, p. 48])**.** For symmetric couplings on $\mathbb{Z}^d$, determine for all ν, and in particular for $\pm J$ and Gaussian couplings, whether $\beta_c(2,\nu) = \infty$, or at least whether $\bar{\beta}_c(2,\nu) = \infty$; and determine whether for some d and some ν, $\bar{\beta}_c(d,\nu) < \infty$, or at least whether $\beta_c(d,\nu) < \infty$.

Newman also asks the first question "at least for some ν with $\nu(\mathbb{R} \setminus \{0\}) > p_c(2) = 1/2$"; that form is answered affirmatively by the diluted laws μ_p with p slightly above $\frac{1}{2}$ (Chapter 13). Chapter 2 describes what physicists believe. Lemma 1.14 and the law behind Proposition 1.23 start Chapters 7 and 8; Chapter 11 uses the exact fair average that the counterexamples above show cannot be avoided.

Chapter 2

What physicists believe

Large simulations locate a spin-glass transition in the three-dimensional Edwards–Anderson model with an error below one percent, and most physicists take its existence for granted, although, as Newman and Stein note, there is as yet no proof "or even a strong physical argument" for it [185, p. 3]. What physicists have debated for forty years is the *nature* of that phase: how many pure states it has, how they are organized, and how the system responds to a small change of temperature, of couplings or of an applied field. Four pictures dominate the debate: the replica-symmetry-breaking picture inherited from the mean-field solution, the droplet or scaling picture, the trivial–nontrivial picture suggested by ground-state numerics, and the chaotic-pairs picture constructed by Newman and Stein.

This chapter describes each picture by the physics argument behind it and by its predictions, stated in the language of Chapter 1. The arguments are of three kinds: scaling hypotheses about the energies and shapes of low-lying excitations; expansions of the replica field theory around the mean-field solution, controlled only in high dimension; and simulations or exact ground-state computations at finite sizes, extrapolated to infinite volume. None of them is a proof, and we label every prediction as such ("Physics prediction"); a reader who wants only the theorems can read this chapter as a list of targets. We also record, with their sources and error bars, the numerical values on which the rigorous programme can lean as calibration: critical temperatures, stiffness exponents, and the lower and upper critical dimensions. The chapter follows the framing of Newman's lectures [171, Chapter 4] and of Newman and Stein's survey of results and speculations [185], updated by the metastate interpretation of replica symmetry breaking due to Read [221] and the review of Newman, Read and Stein [190]. Which parts of these pictures rigorous work has ruled out, which it has constrained and which it leaves untouched is the subject of Chapters 3 and 4.

2.1 The mean-field template

The Sherrington–Kirkpatrick (SK) model [224] puts N spins on a complete graph with Hamiltonian $H_N(\sigma) = -N^{-1/2}\sum_{i<j} J_{ij}\sigma_i\sigma_j$, J_{ij} iid standard Gaussian. Parisi's solution [210, 211] of this model, through the replica method, predicted a phase in which the Gibbs measure splits into many "pure states" not related by symmetry. Two replicas drawn from the Gibbs measure then have an overlap whose law $P_J(q)$ is nontrivial. Its four essential features, in Newman's formulation [171, pp. 52–54], are: (1) P_J depends on the disorder even as $N \to \infty$ (*non-self-averaging*); (2) P_J is a countable sum of point masses; (3) its disorder average P is continuous between two atoms at

$\pm q_{\mathrm{EA}}$; (4) the overlaps of pure states are *ultrametric*: among three pure states the two largest of the three mutual distances $q_{\mathrm{EA}} - q_{\alpha\gamma}$ are equal.

Much of this is now a theorem for the SK model and its mixed p-spin relatives: Guerra's broken-replica bound and Talagrand's proof of the Parisi formula for the free energy [111, 231], its extension by Panchenko [209], and Panchenko's proof that the Ghirlanda–Guerra identities force ultrametricity [207]. In a magnetic field the mean-field theory predicts a phase transition along the de Almeida–Thouless (AT) line [69], below which replica symmetry is broken.

The question for short-range models is how much of this template survives. A complete graph has no geometry: no boundary, no interfaces, no notion of a pure state as an infinite-volume object. On $\mathbb{Z}^d$ all three exist, and Section 1.3 defines pure states unambiguously [171, p. 53]. The pictures below differ in how they transplant the mean-field picture into this setting, or whether they do.

2.2 The stiffness argument

The oldest argument about the existence of the phase is a domain-wall (or stiffness) argument, due to McMillan and to Bray and Moore [35, 160]. Impose on a box of side L first periodic, then antiperiodic boundary conditions in one direction; the second choice flips the signs of the couplings on one seam S of L^{d-1} edges crossing a hyperplane, which forces a relative domain wall across the box. Let ΔF_L be the free-energy difference (Fig. 2.1(a)).

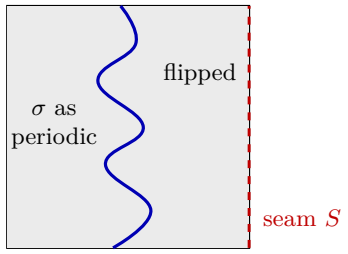
Physics prediction 2.1 (Stiffness exponent and lower critical dimension; McMillan, Bray—Moore)**.** The typical size of ΔF_L grows as ΥL^{θ} for a stiffness exponent $\theta = \theta(d)$ and a stiffness modulus $\Upsilon > 0$. If $\theta > 0$, large-scale excitations become ever more costly and the spin-glass order of the ground state survives at small positive temperature; if $\theta < 0$, they become cheap, order is destroyed at every $T > 0$, and $T_c = 0$ with a correlation length diverging as a power of $1/T$. The lower critical dimension d_l is the dimension where θ changes sign.

For a ferromagnet the same argument gives $\theta = d - 1$: the wall costs its area. In a spin glass the wall can find a path through the frustration and the signs of the couplings it crosses partly cancel, so $\theta < d - 1$. Fisher and Huse argued from numerics and scaling that $\theta \le (d-1)/2$ (see [227, p. 9] and [17, §5.1(v)]). That bound, read as a bound on the variance, is a theorem with a half-page proof. It is stated by Stein and by Arguin, Newman and Stein, who attribute it to Aizenman and Fisher and to Newman and Stein (unpublished) [227, Theorem 1], [17, Theorem 5.1]; we give the short argument with an explicit constant.

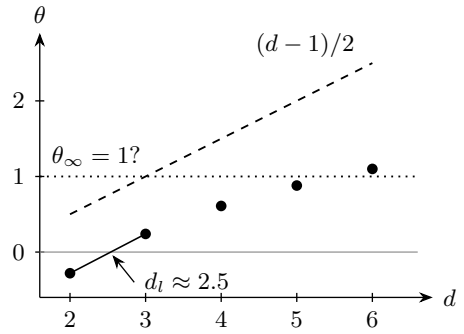
Proposition 2.2 (Surface bound on interface fluctuations)**.** *Let ν be symmetric with $\mathbb{E}J_e^2 < \infty$, let Λ be the discrete torus of side L in $\mathbb{Z}^d$, and let S be the set of L^{d-1} edges crossing a fixed hyperplane. Let J^S agree with J off S and equal $-J$ on S, and $X = \log Z_{\Lambda,J}(\beta) - \log Z_{\Lambda,J^S}(\beta)$ (periodic minus antiperiodic). Then*

$$\mathbb{E}X = 0, \qquad \operatorname{Var} X \;\le\; 4\beta^2\, \mathbb{E}J_e^2\, L^{d-1}.$$

Proof. Write $X = f(J) - f(J^S)$ with $f = \log Z_{\Lambda,\cdot}$, so that $X(J^S) = -X(J)$. Condition on the couplings off S. Given them, the couplings on S are independent and symmetric, so J and J^S have the same conditional law, and $\mathbb{E}[X \mid J_{S^c}] = -\mathbb{E}[X \mid J_{S^c}] = 0$. Hence $\mathbb{E}X = 0$ and $\operatorname{Var} X = \mathbb{E}\operatorname{Var}(X \mid J_{S^c})$: the bulk couplings contribute nothing. For the seam, $|\partial X/\partial J_e| \le 2\beta$ for $e \in S$, since each of the two terms has derivative $\pm\beta\langle\sigma_x\sigma_y\rangle$. The Efron–Stein inequality over the $|S|$ seam couplings, as in the proof of Theorem 1.12(c), gives $\operatorname{Var}(X \mid J_{S^c}) \le \frac{1}{2}\sum_{e\in S} 4\beta^2 \cdot 2\mathbb{E}J_e^2$. □



(a) antiperiodic: a relative domain wall, $\Delta F_L \sim \Upsilon L^\theta$



(b) numerical stiffness exponents and the rigorous ceiling $(d-1)/2$

Figure 2.1: The stiffness argument. (a) Antiperiodic boundary conditions (couplings flipped on the seam S) force a relative domain wall; its free-energy cost scales as L^θ. Only the seam couplings contribute to the variance of the cost, which gives the rigorous ceiling of Proposition 2.2. (b) Numerical estimates of θ for $d = 2, \dots, 6$ (Table 2.2) against the ceiling $(d-1)/2$; θ crosses zero near $d_l \approx 2.5$, and the dotted line marks the mean-field value θ_∞ suggested in [31].

So in free-energy units the standard deviation of ΔF_L is at most $2(\mathbb{E}J_e^2)^{1/2}L^{(d-1)/2}$: if ΔF_L scales as L^θ in the sense of its standard deviation, then $\theta \le (d-1)/2$. The bound says nothing about the sign of θ, which is the whole question. Lower bounds of volume order are known only for a different quantity, the free-energy difference between two putative incongruent states restricted to a box (Chapter 4); there the role of the upper bound is played by a deterministic bound of surface order, and fluctuations of size $L^{d/2}$ contradict a bound of size L^{d-1} only in $d = 2$. This is the mechanism of the two-dimensional results of Chapter 4, and the reason it does not extend to $d \ge 3$.

Numerically θ is negative in $d = 2$ and positive from $d = 3$ on (Table 2.2, Fig. 2.1(b)); interpolation puts d_l near $5/2$, a value also obtained from a replica computation of the interface free energy [93] and from analyses of stiffness and barrier data [32, 154]. At the other end, the zero-field transition is believed to have mean-field critical exponents above the upper critical dimension $d_u = 6$ [115], see also [31]; and the replica field theory of the low-temperature phase is believed to be valid for $d > 6$ [190, p. 22].

2.3 Four pictures of the low-temperature phase

The pictures below assume that a spin-glass phase exists, with broken spin-flip symmetry, and describe its structure. They are phrased in terms of pure states and, where needed, the metastate, a probability measure κ_J on $\mathcal{G}_J(\beta)$ that records which Gibbs state Γ a large box with coupling-independent boundary conditions looks like in its interior (Chapter 3); the metastate is *trivial* if it is a point mass, and a Gibbs state is *trivial* if it is a single pure state or a single flip-related pair. Figure 2.2 is the cartoon to keep in mind.

2.3.1 Droplets and scaling

The droplet picture of McMillan, Bray and Moore, and Fisher and Huse [36, 86, 88, 160] assumes that the low-temperature phase is controlled by a zero-temperature fixed point, and that its

excitations are *droplets*: compact clusters of coherently flipped spins. We use the explicit list of assumptions given by Arguin, Newman and Stein [17, §5.1]. At scale L around a point, (i) the energy of the cheapest droplet of linear size between L and $2L$ has a distribution of width ΥL^θ with positive density at zero energy; (ii) droplet boundaries have fractal dimension d_s with $d-1 < d_s < d$; (iii)–(iv) the stiffness exponent of Physics prediction 2.1 governs both interfaces and droplets, $\theta > 0$; (v) $\theta \le (d-1)/2$. The physics argument is then short. At temperature T a droplet of size L is thermally active only if its energy is of order T, which by (i) has probability of order $(T/\Upsilon)L^{-\theta}$; large active droplets are rare, so the phase is essentially frozen in one configuration and its global flip.

Physics prediction 2.3 (Droplet–scaling picture)**.** For $\beta > \beta_c$: there is exactly one flip-related pair of pure states, a pure state α and its global flip $\bar\alpha$; the metastate is trivial and coupling-independent boundary conditions converge to $\frac12(\alpha+\bar\alpha)$; the overlap distribution is $P(q) = \frac12\delta_{q_{\mathrm{EA}}} + \frac12\delta_{-q_{\mathrm{EA}}}$, with finite-volume weight at $|q| < q_{\mathrm{EA}}$ vanishing like $L^{-\theta}$; interfaces have dimension $d_s < d$; the phase is *chaotic*: a change δJ of the couplings (or δT of the temperature) rearranges the state beyond an overlap length $\ell \sim (\delta J)^{-1/\zeta_c}$ with $\zeta_c = d_s/2 - \theta$ [36, 88], [17, §5.1(vi)]; and there is no AT line in a field [156, §4].

The two-state conclusion is the content of Fisher and Huse's argument against many states [87, 121]. Arguin, Newman and Stein show that, granting assumptions (i)–(vi), their rigorous lower bound on energy fluctuations between incongruent ground states forces the zero-temperature metastate built from coupling-independent boundary conditions onto a single ground-state pair [17, Claim 5.5]; the bound itself is discussed in Chapters 4 and 6.

2.3.2 Replica symmetry breaking in short range

The most direct transplant of the SK picture, which Newman and Stein call the *standard SK picture* [171, pp. 53–56], posits that the periodic finite-volume states converge in a natural way to an infinite-volume state $\rho_J = \sum_\alpha W_\alpha\,\alpha$, a mixture of pure states α with J-dependent weights W_α, whose overlap distribution has the four features of Section 2.1, including non-self-averaging. That version cannot hold: for a translation-covariant, measurable state, the overlap distribution is a translation-invariant function of iid couplings and is therefore almost surely constant [173], [171, Proposition 4.6]. The proof, and exactly what it does and does not exclude, is in Chapter 3.

The version that survives is the *nonstandard* or metastate interpretation, proposed by Newman and Stein [171, pp. 62–64], [185, pp. 7–8], and derived from the replica field theory by Read [221]: replicas are taken in large finite volumes, a typical volume shows a state Γ drawn from a dispersed metastate, and non-self-averaging lives in the variability of Γ, not in a J-dependence of the averaged overlap law. The physics argument is the replica field theory around the Parisi saddle point, which is under control at tree level for $d > 6$, with the assumption that loop corrections do not destroy the replica-symmetry breaking pattern down to the lower critical dimension.

Physics prediction 2.4 (Replica symmetry breaking, metastate interpretation; Parisi, Newman–Stein, Read)**.** For $\beta > \beta_c$: there are infinitely many pure states; the metastate κ_J is nontrivial (dispersed over uncountably many Gibbs states), and a typical Γ drawn from it is itself a nontrivial mixture of many pure states with ultrametric overlaps; the overlap law $P_{J,\Gamma}$ of a sampled Γ varies with Γ, while the metastate-averaged law P_J is self-averaging; interfaces between pure states are space-filling, and system-size excitations cost free energy of order one [227, p. 9]; the metastate-averaged state has power-law correlations with exponent $\zeta = 4$ in high dimension,

and the logarithm of the number of pure states distinguishable in a window of size W is of order $W^{d-\zeta}$ [221]; and there is an AT line in a field [156, §4].

Read's analysis also gives a simple criterion in the language of the mean-field theory: if Parisi's order-parameter function $q(x)$ on $[0,1]$ (the inverse of the distribution function of the overlap, with overlaps taken nonnegative) satisfies $q(0) \neq \int_0^1 q(x)\,\mathrm{d}x$, the metastate is nontrivial [221]. Replica symmetry breaking, read through the metastate, is therefore a statement about the variability of the state seen in large boxes.

Newman and Stein also argued, from the invariance of the metastate under gauge-related boundary conditions, that the nonstandard picture is implausible [183, 185]. That step is heuristic, and Read has argued that it is inconclusive [221]; Newman, Read and Stein list the corresponding class of scenarios as open [190, p. 13]. The review treats the nonstandard picture as not excluded.

2.3.3 Trivial–nontrivial

Ground-state computations in three and four dimensions by Krzakala and Martin and by Palassini and Young [145, 205] found, at accessible sizes, excitations that flip a finite fraction of the box at an energy cost of order one, yet whose surface has fractal dimension $d_s < d$. This is not what the droplet picture predicts (system-size excitations should cost L^θ), nor what the mean-field picture predicts (they should be space filling).

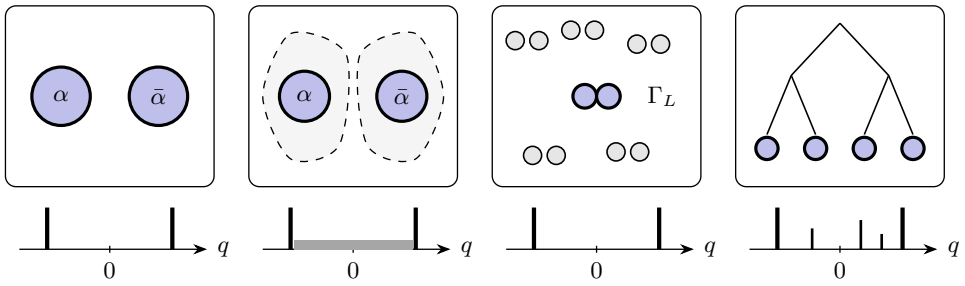
Physics prediction 2.5 (Trivial–nontrivial (TNT) picture; Krzakala–Martin, Palassini–Young)**.** System-size excitations of order-one energy have surface dimension $d_s < d$. Hence the spin-overlap distribution of large finite boxes is nontrivial, with weight between $\pm q_{\mathrm{EA}}$, while the edge-overlap distribution is trivial, since such excitations cross only a vanishing fraction of the edges. In the thermodynamic limit there is a single flip-related pair of pure states, as in the droplet picture.

The last sentence reflects a general fact: excitations whose interfaces have zero density do not create new infinite-volume states in covariant constructions (Chapter 3); Newman and Stein observe that such pictures describe excitations that leave the pure-state structure unchanged [185, p. 3], and Newman, Read and Stein group TNT with the droplet picture as single-pair scenarios [190, p. 3].

2.3.4 Chaotic pairs

The chaotic-pairs picture of Newman and Stein [183, 185] is the many-state scenario that, by the heuristic argument just mentioned, they regard as the only one compatible with the metastate theory [185, p. 11]. Each large box looks like a single flip-related pair of pure states, but the pair changes chaotically with the box size, so the metastate is dispersed over infinitely many trivial Gibbs states.

Physics prediction 2.6 (Chaotic pairs; Newman–Stein)**.** For $\beta > \beta_c$: there are infinitely many flip-related pairs of pure states; the periodic metastate is dispersed over Gibbs states of the form $\Gamma = \frac{1}{2}(\alpha + \bar{\alpha})$; in each large box the overlap distribution is approximately $\frac{1}{2}\delta_{q_{\mathrm{EA}}} + \frac{1}{2}\delta_{-q_{\mathrm{EA}}}$, as in the droplet picture, while overlaps of replicas taken after the infinite-volume limit concentrate at 0; interfaces between different pairs are space filling, with free energies of order L^θ, $0 < \theta < (d-1)/2$ [227, p. 9]. The picture makes no specific prediction about the dimension dependence [185, p. 5].



(a) droplet (b) TNT (finite box) (c) chaotic pairs (d) RSB (one Γ)

Figure 2.2: Four pictures of the low-temperature phase as cartoons of state space; shaded discs are the pure states present in the interior of a large box, and the small histograms show the overlap of two replicas there. (a) Droplet: one flip-related pair. (b) TNT: one pair, but large finite boxes contain cheap, sponge-like excitations (dashed), so the finite-volume spin-overlap law has weight between $\pm q_{\mathrm{EA}}$. (c) Chaotic pairs: infinitely many pairs, but each box shows only one, which changes chaotically with the box size. (d) Replica symmetry breaking: each box shows many pure states, organized ultrametrically (the tree), with a sample-dependent overlap law; the state seen varies from box to box.

The prediction for replicas taken after the limit is the one of [185, Fig. 1(d)], where it is labelled conjectural. A model with this structure at zero temperature is known heuristically: the strongly disordered model of Newman and Stein, whose ground states are minimum spanning trees, is believed to have many ground states for $d > 6$ and of order one below [124]. We know of no direct numerical test of chaotic pairs in the Edwards–Anderson model.

2.3.5 What distinguishes the pictures

Table 2.1 collects the predictions. The pictures agree that a phase exists for $d \geq 3$ and that spin-flip symmetry is broken in it; they differ in the number of states, in whether boxes show one pair or many, in the dimension of interfaces, and in the response to a field. Note that the finite-volume overlap distribution does not separate them cleanly: TNT and RSB both predict a nontrivial $P_L(q)$ at accessible sizes, while chaotic pairs and droplets both predict a trivial one. The metastate (Chapter 3) is the device that separates "many states seen in one box" from "many states seen across boxes".

Chaos does not discriminate either. It is a central prediction of the droplet picture, but it is not specific to it: for the SK model, disorder chaos at every temperature in zero field is a theorem of Chatterjee [46], who also proves that the *edge* overlap of the Edwards–Anderson model is not chaotic [46, §1.4]. Chapter 6 shows that the zero-temperature chaos theorems for the Edwards–Anderson model hold already in $d = 1$, where the ground state is unique, so they are not evidence of many states.

2.4 The numbers

The rigorous programme cannot use numerics as input, but it can use them for calibration: a proof strategy that would give order in $d = 2$, or no order in $d = 3$, contradicts every simulation. Table 2.2 lists the values we use, with their sources; all are heuristic, from finite-size scaling of

	Droplet	TNT	Chaotic pairs	RSB (metastate form)
Pure-state pairs	one	one	infinitely many	infinitely many
Metastate / sampled state	trivial / one pair	trivial / one pair	dispersed / one pair	dispersed / many pure states
Finite-box spin overlap $P_L(q)$	two atoms; weight off $\pm q_{\mathrm{EA}}$ decays like $L^{-\theta}$	nontrivial at accessible L	two atoms	nontrivial, sample dependent
Edge overlap	trivial	trivial	trivial in a box	nontrivial, a function of q
Large-scale excitations or interfaces	$d_s < d$, cost L^{θ}	$d_s < d$, cost $O(1)$	space filling between pairs, cost L^{θ}	space filling, cost $O(1)$
AT line in a field	no	—	—	yes

Table 2.1: Predictions of the four pictures (heuristic). A dash means that the picture, as formulated in the cited sources, makes no specific claim. "Dispersed" means that the metastate is not a point mass; for chaotic pairs and RSB it is believed to be supported on uncountably many states.

Quantity	Estimate	Source
$d = 2$, Gaussian	$T_c = 0$; $\theta = -0.2793(3)$, wall dimension 1.27319(9)	[139]
	earlier: $\theta = -0.266(2)$ and $-0.282(2)$ (two boundary changes)	[116]
$d = 2$, $\pm J$	$T_c = 0$; wall energy saturates ($\theta = 0$); entropic stiffness exponent 0.50(1), order at $T = 0$ only	[116, 233]
$d = 3$, $\pm J$	$T_c = 1.1019(29)$, $\nu = 2.562(42)$, $\eta = -0.3900(36)$	[127]
	$T_c = 1.109(10)$, $\nu = 2.45(15)$, $\eta = -0.375(10)$	[118]
$d = 3$, Gaussian	$T_c = 0.951(9)$, $\nu = 2.44(9)$, $\eta = -0.37(5)$	[135]
$d = 4$, $\pm J$	$T_c = 2.03(3)$, $\nu = 1.00(10)$	[155]
Stiffness θ, $d = 3, 4, 5, 6$	0.24(1), 0.61(2), 0.88(5), 1.1(1)	[31]
Lower critical dimension	$d_l \approx 5/2$	[32, 93, 154]
Upper critical dimension ($h = 0$)	$d_u = 6$	[115]
$d = 3$, replicon exponent	0.38(2) (droplet: 0 asymptotically)	[125, eq. (19), recalling earlier Janus data]
Interface dimension	$< d$ for $d \leq 5$, close to d at $d = 6$ (approximate algorithms)	[243]

Table 2.2: Calibrated numerical values (heuristic). Error bars are those quoted by the authors; "wall dimension" is the fractal dimension of the domain wall induced by a change of boundary conditions.

simulations or exact ground states at finite L. Temperatures are in units of the coupling scale ($|J_e| = 1$ for $\pm J$, unit variance for Gaussian).

The evidence on the nature of the phase is split, and the split sits in the extrapolation to large L rather than in the data. In favour of replica symmetry breaking in $d = 3$: equilibrium simulations on the special-purpose computer Janus at $L = 32$ down to $0.64\,T_c$ find spin- and link-overlap statistics consistent with it at the length scales studied [125]; the $d = 4$ $\pm J$ data of [155] were read the same way; a numerical construction of the Aizenman–Wehr metastate in $d = 3$ gives evidence for a dispersed metastate [29]; and recent low-energy-excitation studies prefer the RSB fit for the excitation dimension [57]. In favour of a single pair: overlap-peak statistics of the three-dimensional model differ from those of SK [246]; stiffness measured with thermal

boundary conditions grows with L [242]; interface dimensions computed with approximate algorithms are below d for $d \leq 5$ [243]; and Moore argues that the droplet picture predicts an apparent replicon exponent $2\theta \approx 0.48$ below a crossover length that may be hundreds of lattice spacings in $d = 3$, against the measured 0.38(2) [164]. The ground-state data that motivated TNT come with the authors' own caveat of a possible crossover at larger sizes [205]. A review of the numerics concludes that in zero field "the most economic conclusion" is an RSB phase in three dimensions, while in a field the lower critical dimension for the AT transition seems to be larger than three [156].

The field question is the sharpest split. For an AT transition: evidence in $d = 4$ from Janus [126] and, recently, in $d = 3$ from a modified susceptibility analysis [162]. Against: one-dimensional long-range proxies for $d = 3$ and $d = 4$ are consistent with no AT line or give conflicting results [146], and a Heisenberg proxy in a random vector field suggests that the AT line disappears below six dimensions [239]. Temperature chaos is observed to be much harder to see than disorder chaos [134], and to be driven by rare events [85].

2.5 From beliefs to theorems

Every picture of this chapter assumes that the transition exists, and none of them is a theorem for the Edwards–Anderson model on $\mathbb{Z}^d$ in any dimension; the rigorous status is that of Open problem 1.24. What rigorous work has done instead is to constrain the pictures. The first such constraint is already proved in this chapter: Proposition 2.2 caps interface fluctuations at surface order. The metastate theory of Chapter 3 excludes the standard SK picture in every dimension, restricts the possible structures of states drawn from a metastate, and shows that interfaces between states drawn from covariant constructions are either empty or of positive density; Chapter 4 turns fluctuation bounds into strong restrictions in $d = 2$ and ends with a table of each picture against each rigorous constraint, next to the numerics of Table 2.2. Zero temperature, where the droplet and chaos predictions can be tested on ground states, is Chapters 5 and 6.

Open problem 2.7 (Newman's Research Problems 4.5 and 4.14 [171, pp. 50, 64])**.** Determine whether the number of pure states is infinite, or at least larger than two, for some d, ν and β. Determine whether for some d, ν and β the metastate is, for ν-almost every J, supported on many Gibbs states, or on Gibbs states that decompose into more than two pure states.

Part II

States, metastates and two dimensions

Chapter 3

How many states? Metastates

Fix a dimension d, a symmetric coupling law ν and an inverse temperature β. Whether the Edwards–Anderson model orders at low temperature is a question about infinite-volume Gibbs states: order in the sense of Edwards and Anderson (Chapter 1) requires a state in which $\langle\sigma_0\sigma_x\rangle_J^2$ does not decay, and in a symmetric construction this already forces at least two pure states, a pure state α and its global flip $\bar{\alpha}$. The pictures of Chapter 2 disagree about how many pure states there are, how they are weighted, and which of them a large finite box actually sees. This chapter makes these questions precise and proves the constraints that hold in every dimension and at every temperature.

Two difficulties shape everything: finite-volume Gibbs measures need not converge as the box grows (chaotic size dependence), and no measurable covariant rule picks one ordered state out of a flip pair. Both are resolved by passing to probability measures on states, the *metastates* of Aizenman–Wehr and Newman–Stein, whose tools are those of ergodic theory: covariance and ergodicity turn many statements into dichotomies ("a.s. constant", "empty or of positive density"). By the end the reader can construct a metastate, prove that covariant overlap distributions self-average and that covariant interfaces are absent or space-filling, state the structure theorems and the Newman–Read–Stein classification, and derive the Ghirlanda–Guerra identities for the link overlap, with the reasons why they cannot by themselves decide order.

Throughout, $\Sigma = \{-1,+1\}^{\mathbb{Z}^d}$, and $M_1(\Sigma)$, the probability measures on Σ with the topology of convergence of cylinder probabilities, is compact and metrizable. We write $\sigma_e = \sigma_x\sigma_y$ for $e = xy$; translations act by $(\tau_a J)_{xy} = J_{x-a,y-a}$, $(\tau_a\sigma)_x = \sigma_{x-a}$, and on measures by push-forward; $\mathcal{G}_J = \mathcal{G}_J(\beta)$ is the set of Gibbs states for J (Chapter 1). We write ν also for the product law $\nu^{\otimes E_d}$ of J, as in "ν-a.e. J".

3.1 Finite volumes and their limits

Let $\mu_{L,J}^{\mathrm{per}}$ be the Gibbs measure on $\Lambda_L = \{-L,\dots,L\}^d$ with periodic boundary condition (Chapter 1). By compactness, for each fixed J the sequence $(\mu_{L,J}^{\mathrm{per}})_L$ has convergent subsequences, and every limit point is a Gibbs state for J. The question is whether a limit exists, and if not, how the limit points are visited. For the two-dimensional Ising ferromagnet at $\beta > \beta_c$ the periodic measures converge to $\frac{1}{2}\mu^+ + \frac{1}{2}\mu^-$, by flip symmetry and because $\mu^\pm$ are the only pure states [171, p. 54], and plus boundary conditions converge to μ^+: monotonicity selects the extreme states. A spin glass has no such structure. With many competing pure states one expects the state in Λ_L to favour those that "fit" the boundary condition best, and which these are to change erratically

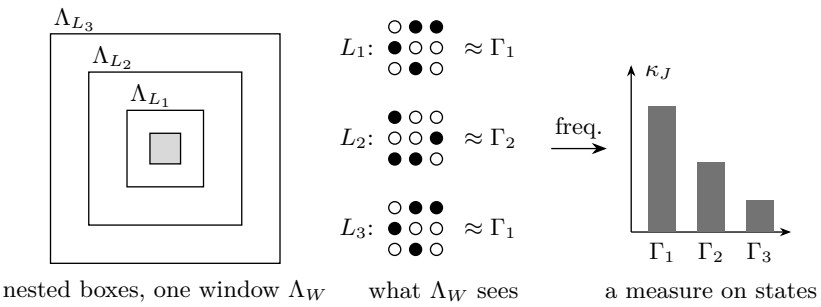


Figure 3.1: Chaotic size dependence and the metastate. Along deterministic periodic boxes the state seen in a fixed window may jump between Gibbs states $\Gamma_1, \Gamma_2, \dots$ in a J-dependent way; what can converge is the frequency with which each state is seen, the metastate κ_J of Section 3.2.

with L: *chaotic size dependence* [172]. For typical J, different J-independent subsequences of volumes have different limits. In Fig. 3.1 the state seen in a fixed window does not settle down, but the *frequency* with which each state is seen may.

Without any picture, the limit, when it exists, does not see certain changes of the boundary condition. Two boundary conditions on Λ_L are *flip-related* if one is obtained from the other by reversing the couplings on a set B_L of boundary edges: periodic and antiperiodic ones are, and so are any two fixed ones, whereas periodic and fixed ones are not [183, Sec. 5.2.4].

Theorem 3.1 (Newman–Stein [172]; see [183, Thm. 2])**.** *Let ν be symmetric. Suppose that along a deterministic sequence of volumes Λ_{L_k}, with coupling-independent boundary conditions, the finite-volume Gibbs measures converge for ν-a.e. J to a state ρ_J. If in every volume (or in any J-independent subset of the volumes) the boundary condition is replaced by a flip-related one, then the new sequence also converges for ν-a.e. J, and its limit is ρ_J.*

Proof of the identification of the limit. We prove that if both sequences converge, the limits agree; the existence of the second limit is proved in [172]. Fix a finite set A and let $X^1_L(J)$, $X^2_L(J)$ be the expectations of $\sigma_A = \prod_{x\in A}\sigma_x$ in the two finite-volume measures, with limits X^1, X^2. Let $\mathbb{E}_r$ denote conditional expectation given the couplings in Λ_r. Flip-relatedness means $X^2_L(J) = X^1_L(J^{(B_L)})$, where $J^{(B_L)}$ reverses the couplings on B_L. For $L > r$ the set B_L lies outside Λ_r, and since ν is symmetric, $J^{(B_L)}$ and J have the same conditional law given the couplings in Λ_r. Hence $\mathbb{E}_r[X^1_L] = \mathbb{E}_r[X^2_L]$, and by bounded convergence $\mathbb{E}_r[X^1 - X^2] = 0$ for every r. By the martingale convergence theorem $X^1 - X^2 = \lim_r \mathbb{E}_r[X^1 - X^2] = 0$ a.s. Countably many A determine the limit states. □

This is the rigorous core of the heuristic that many incongruent states cause chaotic size dependence: if switching from periodic to antiperiodic boundary conditions changed the state deep inside, the periodic sequence could not converge [183, Sec. 5.2.4]. Nor can one simply "choose an ordered state". For finite $S \subset \mathbb{Z}^d$ let $g_S = -1$ on S and $+1$ off S, acting by $(g_S J)_{xy} = g_S(x)g_S(y)J_{xy}$ and $(g_S\sigma)_x = g_S(x)\sigma_x$: the gauge transformation of Chapter 1, which reverses the couplings on the edge boundary of S and so preserves the iid symmetric law ν.

Lemma 3.2 (No covariant ordered selection; a standard gauge argument)**.** *Let ν be invariant under every g_S, S finite (for instance iid symmetric couplings). Let $J \mapsto \Gamma_J$ be a measurable map into $M_1(\Sigma)$ such that, for every finite S, $\Gamma_{g_S J} = (g_S)_*\Gamma_J$ for ν-a.e. J. Then Γ_J is invariant under the global spin flip for ν-a.e. J.*

Proof. Let A be finite with $|A|$ odd, and $m_A(J) = \Gamma_J(\sigma_A)$. Let h be a bounded function of the couplings on a finite edge set F, and choose a finite S containing A and every endpoint of every edge of F. Then $h(g_S J) = h(J)$, because both endpoints of each edge of F lie in S, while $m_A(g_S J) = \Gamma_J(\prod_{x\in A} g_S(x)\sigma_x) = -m_A(J)$. Invariance of ν gives $\mathbb{E}[h\, m_A] = \mathbb{E}[(h\, m_A)\circ g_S] = -\mathbb{E}[h\, m_A]$, so $\mathbb{E}[h\, m_A] = 0$. Cylinder functions h are dense in $L^2(\nu)$, hence $m_A = 0$ a.s. There are countably many odd A, and a measure on Σ all of whose odd moments vanish is flip-invariant. □

Corollary 3.3. *Let ν be iid symmetric. If the finite-volume Gibbs measures with a fixed, coupling-independent boundary condition (plus, say) converge for ν-a.e. J along a deterministic sequence of volumes, the limit is flip-symmetric.*

Proof. For S inside the box, the finite-volume measure transforms exactly as $\mu_{L,g_S J} = (g_S)_*\mu_{L,J}$, because the boundary spins are untouched. Intersect the full-measure convergence set with its image under g_S, pass to the limit using continuity of $(g_S)_*$, and apply Lemma 3.2. □

So for symmetric couplings, plus boundary conditions either fail to converge or converge to a symmetric state; in a two-state picture one expects the first, the plus measures alternating between α and $\bar\alpha$ [171, p. 59]. Order cannot appear as a magnetization of a covariant state; it must be read off from even correlations at large distance, pure-state decompositions, two replicas, or the way a state is selected.

3.2 Metastates

The way out, proposed by Aizenman and Wehr [8] and, for spin glasses, by Newman and Stein [174, 175], is to track the *distribution* of finite-volume states, "moving the discussion one rung up" [8, p. 504]. We first record what one wants of the result [190, Defs. 3.1–3.2]. For finite $B \subset \mathbb{Z}^d$ and $J_B \in \mathbb{R}^{E(B)}$ define the local reweighting $L_{J_B}\colon M_1(\Sigma) \to M_1(\Sigma)$ by

$$(L_{J_B}\Gamma)(f) = \frac{\Gamma\big(f\, e^{\beta\sum_{e\in E(B)} J_{B,e}\sigma_e}\big)}{\Gamma\big(e^{\beta\sum_{e\in E(B)} J_{B,e}\sigma_e}\big)}. \tag{3.1}$$

If $\Gamma \in \mathcal{G}_J$ then $L_{J_B}\Gamma \in \mathcal{G}_{J+J_B}$, where $J + J_B$ adds J_B to the couplings on $E(B)$: this is the effect of a finite change of couplings on a Gibbs state.

Definition 3.4 (Metastate)**.** A *metastate* is a measurable map $J \mapsto \kappa_J$ from coupling configurations to probability measures on $M_1(\Sigma)$ such that, for ν-a.e. J:

(i) (support on Gibbs states) $\kappa_J(\mathcal{G}_J) = 1$;

(ii) (coupling covariance) for every finite B and $J_B \in \mathbb{R}^{E(B)}$, $\kappa_{J+J_B} = (L_{J_B})_*\kappa_J$.

It is *translation covariant* if in addition $\kappa_{\tau_a J} = (\tau_a)_*\kappa_J$ for every $a \in \mathbb{Z}^d$. The *metastate-average state* (barycenter) is $\rho_J = \int \Gamma\, \kappa_J(d\Gamma)$; it is again a Gibbs state for J.

Coupling covariance makes the dependence on any finite set of couplings explicit, so that coupling derivatives can be taken without boundary terms (Chapter 4). A gauge transformation is a special finite change of couplings, $g_S J = J + \Delta J$ with ΔJ supported on the edge boundary of S, and $L_{\Delta J}\Gamma = (g_S)_*\Gamma$ for every $\Gamma \in \mathcal{G}_J$ (condition on the spins outside S and flip all spins of S in the finite sum; the normalizing factor equals 1). So a metastate should satisfy $\kappa_{g_S J} = (g_S)_*\kappa_J$, and then its barycenter satisfies the hypothesis of Lemma 3.2 (Corollary 3.6 below).

Theorem 3.5 (Aizenman–Wehr [8, Prop. 5.1 and App. I]; in the form of [171, Thm. 4.8])**.** *Let ν be iid, $\beta < \infty$, and consider periodic boundary conditions. There is a deterministic subsequence L_n such that the joint law of $(J, \mu^{\mathrm{per}}_{L_n,J})$ converges to a probability measure $K^\dagger$ on $\mathbb{R}^{E_d} \times M_1(\Sigma)$ whose J-marginal is ν. Its conditional distribution κ_J given J satisfies $\kappa_J(\mathcal{G}_J) = 1$ for ν-a.e. J.*

Convergence is in the product topology: of all finite-dimensional distributions of the couplings together with the cylinder probabilities of the state. The resulting κ_J, the *periodic* metastate, is coupling covariant and, because the torus is translation invariant, translation covariant [8, Prop. 5.1(iii)–(iv)], [171, p. 55], [190, p. 9]. It is the law of the state seen near the origin given the couplings near the origin, averaged over the couplings far away.

Proof sketch. This sketch omits the measure-theoretic details, which are in Section A.5 together with the full proof. The J-marginal is fixed and $M_1(\Sigma)$ is compact, so the laws are tight and Prokhorov's theorem gives a convergent deterministic subsequence; disintegrating the limit defines κ_J. The DLR equations of a finite Λ, tested inside a larger finite Λ', are closed conditions on (J, Γ) that the torus measures satisfy exactly once Λ' lies well inside the torus, so they survive the weak limit by the portmanteau theorem. □

Whether the subsequence is needed at all is open (Open problem 3.8 below).

Corollary 3.6. *For iid symmetric couplings, the barycenter ρ_J of the periodic metastate (Theorem 3.5) is flip-symmetric for ν-a.e. J.*

Proof. For S inside the torus, $\mu^{\mathrm{per}}_{L,g_S J} = (g_S)_* \mu^{\mathrm{per}}_{L,J}$, and $g_S J$ has the law of J. So the law of $(J, \mu^{\mathrm{per}}_{L,J})$ is invariant under the continuous map $(J, \Gamma) \mapsto (g_S J, (g_S)_* \Gamma)$, and so is its limit $K^\dagger$. Disintegrating gives $\kappa_{g_S J} = (g_S)_* \kappa_J$ a.s., hence $\rho_{g_S J} = (g_S)_* \rho_J$ a.s.; apply Lemma 3.2. □

For the periodic metastate more is true: the torus Hamiltonian is even, so every $\mu^{\mathrm{per}}_{L,J}$ is flip-invariant, and since flip invariance is a closed condition, for ν-a.e. J, κ_J-almost every Γ is flip-invariant. The gauge argument matters for constructions whose finite-volume states are not symmetric, such as fixed boundary conditions (Corollary 3.3), where only covariant averages are forced to be symmetric. The second construction keeps J fixed and treats the volume as the random variable.

Theorem 3.7 (Newman–Stein [174, 178]; as stated in [171, Thm. 4.11])**.** *Along a further deterministic subsequence L_{n_k} of the volumes of Theorem 3.5 and a deterministic subsequence K_m, for ν-a.e. J,*

$$\frac{1}{K_m} \sum_{k \le K_m} \delta_{\mu^{\mathrm{per}}_{L_{n_k},J}} \longrightarrow \kappa_J \qquad (m \to \infty),$$

where κ_J is the metastate of Theorem 3.5.

So the Aizenman–Wehr metastate is also the empirical distribution of states along a sparse sequence of volumes at fixed J, the precise form of Fig. 3.1; both subsequences are chosen independently of J. In mean-field models of random-field and Hopfield type the almost-sure analogue fails, convergence holding only in distribution [171, p. 60]; for the EA model the question is open.

Open problem 3.8 (Newman [171, Research Problems 4.7 and 4.10])**.** Prove that the joint law of $(J, \mu^{\mathrm{per}}_{L,J})$ converges as $L \to \infty$ without passing to a subsequence, or that the empirical measures of Theorem 3.7 converge almost surely along the full sequence of volumes.

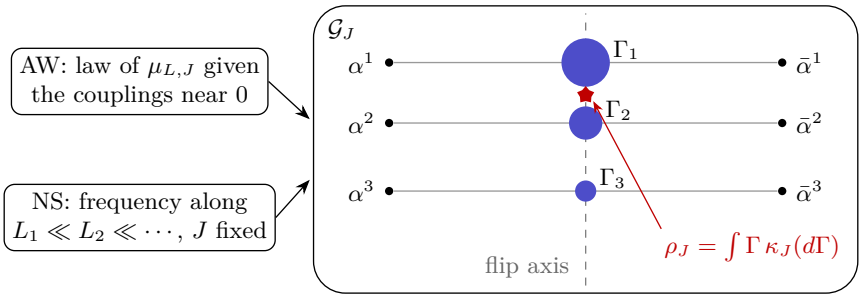


Figure 3.2: A metastate as a measure on states, drawn for a chaotic-pairs scenario. Pure states come in flip pairs $(\alpha^k, \bar\alpha^k)$; each Gibbs state $\Gamma_k = \frac{1}{2}(\alpha^k + \bar\alpha^k)$ lies on the flip axis, and κ_J gives it the weight shown by the disc; both constructions (AW: Aizenman–Wehr; NS: Newman–Stein) give the same κ_J along suitable deterministic volumes. The barycenter ρ_J (star) is flip-symmetric (Corollary 3.6) and forgets which pair a volume selects.

The joint-limit construction gives the covariance that the proofs use; the empirical one gives the interpretation of $\kappa_J(\{\Gamma\})$ as the fraction of large volumes in which Γ is seen. The metastate is insensitive to the remaining freedom in the boundary condition.

Theorem 3.9 (Newman–Stein [177]; see [183, Thm. 3])**.** *Let ν be symmetric. Two metastates constructed along the same deterministic sequence of volumes, using two sequences of flip-related, coupling-independent boundary conditions, coincide for ν-a.e. J.*

Proof sketch. This sketch omits the passage to a common subsequence. There the metastate is the Aizenman–Wehr conditional distribution, and the symmetry argument of Theorem 3.1, applied to bounded continuous functions of the finite-volume state, identifies the two joint limits [183, pp. 39–40]. □

In particular the periodic and antiperiodic metastates coincide, as do all metastates built from fixed boundary conditions chosen independently of J [190, p. 10]; from now on we work with the periodic metastate unless stated otherwise.

3.3 What translation covariance forbids

Translation covariance turns natural functions of the couplings into translation-invariant ones, which for iid couplings are constant. Applied to the right objects, this gives the first rigorous constraints on the pictures of Chapter 2.

Lemma 3.10 (Ergodicity of iid couplings)**.** *Let ν be an iid product law on $\mathbb{R}^{E_d}$. If f is measurable and $f \circ \tau_a = f$ ν-a.s. for every $a \in \mathbb{Z}^d$, then f is ν-a.s. constant.*

This is the standard mixing argument: approximate an invariant event by a cylinder event, which is independent of its far translates (Section A.6).

Proposition 3.11 (Newman [171, Prop. 4.4])**.** *For iid couplings and every β, the number $N(J,\beta) \in \{1, 2, \ldots, \infty\}$ of pure states is measurable and ν-a.s. equal to a constant $N(d, \nu, \beta)$.*

Measurability is the substantive part [171, pp. 48–49]; τ_a maps the pure states for J bijectively onto those for $\tau_a J$, so Lemma 3.10 gives the rest. Determining $N(d, \nu, \beta)$ for any $d \geq 2$ and large β is open [171, Research Problem 4.5].

Let $J \mapsto \rho_J$ be a measurable, translation-covariant choice of Gibbs state, for instance the barycenter of the periodic metastate. The joint law $\nu(dJ)\,\rho_J(d\sigma)\,\rho_J(d\sigma')$ of the couplings and two replicas is then invariant under $(J, \sigma, \sigma') \mapsto (\tau_a J, \tau_a \sigma, \tau_a \sigma')$, so by the multiparameter ergodic theorem the site overlap

$$q(\sigma, \sigma') = \lim_{L \to \infty} \frac{1}{|\Lambda_L|} \sum_{x \in \Lambda_L} \sigma_x \sigma'_x \tag{3.2}$$

exists almost surely, and by Fubini's theorem, for ν-a.e. J it exists $\rho_J \otimes \rho_J$-a.s.; its law P_J is the *overlap distribution* of ρ_J. A translation-covariant metastate has three such objects [190, p. 14]: the law $P_{J,\Gamma}$ of q under $\Gamma \otimes \Gamma$ for Γ drawn from κ_J (the weighted law of the overlaps of pairs of pure states of Γ), its average $P_J = \int P_{J,\Gamma}\,\kappa_J(d\Gamma)$, and the law $P_{J,\rho}$ of q under $\rho_J \otimes \rho_J$.

Theorem 3.12 (Newman–Stein [174, 175]; [171, Prop. 4.6])**.** *Let ν be iid, $d \geq 1$, $\beta < \infty$.*

(a) If $J \mapsto \rho_J$ is a measurable, translation-covariant Gibbs state, its overlap distribution P_J equals a deterministic probability measure P on $[-1, 1]$ for ν-a.e. J.

(b) For a translation-covariant metastate, $P_J = \int P_{J,\Gamma}\,\kappa_J(d\Gamma)$ and $P_{J,\rho}$ are ν-a.s. deterministic.

Proof. (a) If σ, σ' are independent with law ρ_J, then $\tau_a\sigma, \tau_a\sigma'$ are independent with law $\rho_{\tau_a J}$, and $q(\tau_a\sigma, \tau_a\sigma') = q(\sigma, \sigma')$ because shifting the boxes in (3.2) changes the average by $O(|\Lambda_L \triangle (\Lambda_L + a)|/|\Lambda_L|) = O(1/L)$. Hence $P_{\tau_a J} = P_J$ a.s. For each $k \geq 1$ the moment $J \mapsto \int q^k\,P_J(dq)$ is a bounded measurable translation-invariant function, so it is a.s. constant by Lemma 3.10. A probability measure on $[-1, 1]$ is determined by its moments; intersect the countably many full-measure events. (b) The same argument applies to $\int q^k P_J(dq) = \int \kappa_J(d\Gamma) \int q^k\,d(\Gamma \otimes \Gamma)$, which is translation invariant by translation covariance of κ, and to the barycenter, which is a translation-covariant state. □

The theorem is a no-go result, and it matters exactly what it excludes. In Newman's *standard SK picture* [171, pp. 53–54] (Chapter 2) the periodic states converge in some natural way to a J-dependent state $\rho_J = \sum_\alpha W_\alpha\,\alpha$ with countably many pure states α, whose overlap distribution P_J has the four features of the Parisi solution: (1) it depends on J; (2) it is a countable sum of point masses; (3) its J-average is continuous apart from two atoms at $\pm q_{\mathrm{EA}}$; (4) the overlaps are ultrametric.

Route warning 3.13 (Self-averaging does not exclude many states). Theorem 3.12 excludes feature (1) for every measurable translation-covariant state, hence the standard SK picture; since $P = \int P_J\,d\nu$, a deterministic P_J cannot satisfy both (2) and (3). It does *not* exclude a deterministic overlap distribution with many atoms, ultrametric or not, nor many pure states, and says nothing about how $P_{J,\Gamma}$ depends on the state Γ drawn from the metastate. Non-self-averaging is permitted in short range only for $P_{J,\Gamma}$, and only if the metastate is nontrivial [190, p. 15]: the opening used by the metastate interpretation of replica symmetry breaking (Section 3.6). Strength: "self-averaging, hence no replica symmetry breaking" is false as an implication.

Drawing two replicas from one state Γ or from the barycenter is not a technicality: the difference can carry the whole order-parameter signal.

Observation 3.14 (Same state versus averaged state; cf. [187, 188, 221]). Fix J, a probability measure κ on Gibbs states and its barycenter ρ. For a finite window W let $Q_W = |W|^{-1}\sum_{x\in W}\sigma_x\sigma'_x$. Then

$$\int \kappa(d\Gamma)\,(\Gamma\otimes\Gamma)(Q_W^2) - (\rho\otimes\rho)(Q_W^2) = \frac{1}{|W|^2}\sum_{x,y\in W}\mathrm{Var}_\kappa\big[\Gamma(\sigma_x\sigma_y)\big] \;\geq 0.$$

Indeed, $(\Gamma\otimes\Gamma)(\sigma_x\sigma'_x\sigma_y\sigma'_y) = \Gamma(\sigma_x\sigma_y)^2$, while for the barycenter it is $\rho(\sigma_x\sigma_y)^2$, the square of the κ-average of $\Gamma(\sigma_x\sigma_y)$. In a chaotic-pairs scenario each Γ is $\frac{1}{2}(\alpha+\bar\alpha)$ for one flip pair, so $(\Gamma\otimes\Gamma)(Q_W^2)$ is close to q_{EA}^2 for large W, while the barycenter averages $\alpha(\sigma_x)\alpha(\sigma_y)$ over many incongruent pairs and its replica overlap can be much smaller. A statement of "EA order" must say from which object the replicas are drawn (Chapter 1).

3.4 Interfaces are space-filling

The *interface* between two spin configurations σ,σ' is the set of edges satisfied in one and not in the other, $I(\sigma,\sigma') = \{e : \sigma_e \neq \sigma'_e\}$; on the dual lattice it is a union of domain walls. For two states Γ,Γ' and $\varepsilon>0$ put $I_\varepsilon(\Gamma,\Gamma') = \{e : |\Gamma(\sigma_e) - \Gamma'(\sigma_e)| > \varepsilon\}$. Following [14, Def. 1.1], Γ and Γ' are *incongruent* if for some $\varepsilon>0$ the set I_ε has positive lower density among the edges of Λ_L; distinct states whose interfaces have zero density, like the Dobrushin states of the ferromagnet, are *regionally congruent*. In covariant constructions only the first can occur.

Lemma 3.15 (Stationary sets are empty or dense)**.** *Let I be a random subset of E_d whose law is invariant under translations, and let $\delta_L = |I\cap E(\Lambda_L)|/|E(\Lambda_L)|$. Then $I=\varnothing$ almost surely on the event $\{\delta_L\to 0\}$.*

Proof. The event $Z=\{\delta_L\to 0\}$ is translation invariant, since shifting I changes δ_L by $O(1/L)$. Each of the d edge directions contributes the same number of edges to $E(\Lambda_L)$, so by invariance $\mathbb{E}[\delta_L\mathbf{1}_Z] = \frac{1}{d}\sum_{i=1}^d \mathbb{P}(e_i\in I,\,Z)$, where e_i is the edge from 0 to the i-th unit vector. The left side tends to 0 by bounded convergence and does not depend on L, so $\mathbb{P}(e\in I, Z)=0$ for every edge e (by invariance again). A countable union over edges gives $\mathbb{P}(I\neq\varnothing, Z)=0$. □

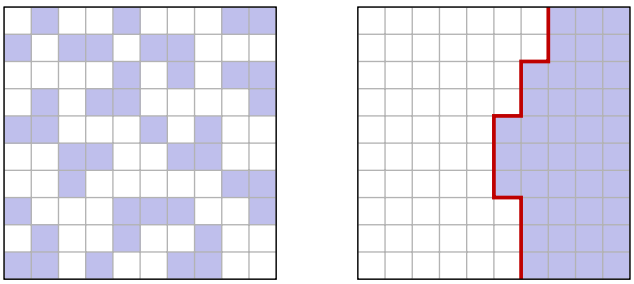
Since δ_L converges a.s. by the ergodic theorem, this is a dichotomy: a.s. either I is empty or its density is positive.

Theorem 3.16 (Newman–Stein [179, 182]; see [183, Thm. 4])**.** *Let (J,σ,σ') be random with a translation-invariant joint law; for instance $\sigma\sim\kappa_J$, $\sigma'\sim\kappa'_J$ independently, where κ,κ' are translation-covariant zero-temperature metastates (measures on ground states). Then almost surely either $\sigma'=\pm\sigma$, or the interface $I(\sigma,\sigma')$ has positive density.*

Proof. $I(\sigma,\sigma')$ is a translation-covariant function of (σ,σ'), so its law is translation invariant. By Lemma 3.15 and the ergodic theorem, almost surely either $I=\varnothing$ or its density is positive. If $I=\varnothing$ then $\sigma_x\sigma_y=\sigma'_x\sigma'_y$ on every edge, and since $\mathbb{Z}^d$ is connected, $\sigma'=\pm\sigma$. □

The proof uses nothing about ground states: covariance forces the dichotomy, so distinct ground states built from coupling-independent boundary conditions are incongruent, with $d_s=d$. At positive temperature, applied to $I_\varepsilon(\alpha,\alpha')$, it gives the following.

Theorem 3.17 (Newman–Stein [187, Thm. 5.2])**.** *Let κ be a periodic metastate at $\beta<\infty$, and let α,α' be pure states sampled from the pure-state decompositions of states drawn from κ_J.*



space-filling: $d_s = d$ one wall: $d_s = d - 1$

Figure 3.3: Two ways for σ' to differ from σ. Shaded cells mark where $\sigma' = -\sigma$; the interface is their boundary. Left: positive density at every scale. Right: a single wall, whose density in Λ_L tends to zero; this cannot occur in a translation-covariant construction (Theorem 3.16).

Almost surely, either α and α' are incongruent, or $\alpha(\sigma_e) = \alpha'(\sigma_e)$ for every edge e; in the latter case all even correlation functions of α and α' coincide. If in addition α and α' have the same EA order parameter, then in the latter case $\alpha' \in \{\alpha, \bar{\alpha}\}$.

Proof sketch. The sketch omits measurability and the analytic details. The dichotomy is Lemma 3.15 applied to I_ε for rational ε. Equal edge correlations stay equal under finite coupling changes [187, Lemma 4.1], and differentiating in the couplings propagates the equality to all even correlations; with equal EA order parameters, clustering then gives $\alpha(\sigma_x) = \pm\alpha'(\sigma_x)$ with a common sign [187, proof of Thm. 5.2]. □

Route warning 3.18 (No new states from thin excitations). A tempting route to many states is to extract a new thermodynamic state from a low-energy system-size excitation with a fractal interface of dimension $d_s < d$. In any translation-covariant construction this fails outright (strength: proved false): such an interface has density zero, so by Lemma 3.15 it is absent from every fixed window in the limit. Such excitations can matter for finite-size overlap distributions, as in the trivial–nontrivial picture of Chapter 2, but produce no new state; regionally congruent states, if they exist, can only come from boundary conditions chosen as functions of the couplings [189, Remark after Thm. 6.2].

3.5 Inside a Gibbs state: weights and structure

A Gibbs state $\Gamma \in \mathcal{G}_J$ has a unique pure-state decomposition $\Gamma = \int \alpha \, w_\Gamma(d\alpha)$ [106, Thm. 7.26], written $\sum_\alpha W_\alpha \, \alpha$ when countable; coupling covariance acts on it simply.

Lemma 3.19 (Reweighting under a finite coupling change)**.** *Let B be finite, $J_B \in \mathbb{R}^{E(B)}$, $X = \exp(\beta \sum_{e \in E(B)} J_{B,e}\sigma_e)$ and $r_\alpha = \alpha(X)$.*

(a) If α is a pure state for J, then $L_{J_B}\alpha$ is a pure state for $J + J_B$.

(b) $L_{J_B}\Gamma = \int L_{J_B}\alpha \, w'(d\alpha)$ with $w'(d\alpha) = r_\alpha \, w_\Gamma(d\alpha) / \int r \, dw_\Gamma$; for a countable decomposition, $W'_\alpha = r_\alpha W_\alpha / \sum_\gamma r_\gamma W_\gamma$.

Proof. (b) is a one-line computation: $\Gamma(fX) = \int \alpha(fX) \, w_\Gamma(d\alpha) = \int r_\alpha \, (L_{J_B}\alpha)(f) \, w_\Gamma(d\alpha)$; divide by $\Gamma(X) = \int r \, dw_\Gamma$. For (a), $L_{J_B}\alpha$ is a Gibbs state for $J + J_B$, and pure states are exactly the

Gibbs states that are trivial on the tail σ-field [106, Thm. 7.7]. For a tail event T, $\alpha(T) \in \{0,1\}$ gives $\alpha(\mathbf{1}_T X) \in \{0, \alpha(X)\}$, so $(L_{J_B}\alpha)(T) \in \{0,1\}$. □

So a finite coupling change keeps the pure states (up to $\alpha \mapsto L_{J_B}\alpha$) and reweights them by r_α; flip-related pure states have $r_\alpha = r_{\bar{\alpha}}$, so the weights $\frac{1}{2}, \frac{1}{2}$ of a flip pair never move. Unequal weights move, which is incompatible with translation covariance.

Theorem 3.20 (Newman–Stein [185, Thm. 1])**.** *The periodic metastate cannot assign positive probability to Gibbs states $\Gamma = \sum_\alpha W_\alpha \alpha$ with not all $W_\alpha = \frac{1}{2}$, if over all these Γ only countably many pure states appear.*

Proof sketch. This follows [185, pp. 9–10] and omits measurability. With countably many pure states, the ordered weights in ρ_J, and their distribution within states drawn from κ_J, are translation invariant, hence constant, hence unchanged by finite coupling changes; but for distinct pure states α, δ with nonzero weights a small change J_B with $r_\alpha \neq r_\delta$ multiplies W_α/W_δ by $r_\alpha/r_\delta \neq 1$ (Lemma 3.19), changing that distribution. □

Theorem 3.21 (Newman–Stein; as stated in [187, Thm. 6.1])**.** *Let κ be a periodic metastate at fixed $\beta < \infty$. Every Gibbs state Γ in the support of κ_J consists of (a) a single pure state, or (b) a single pair of spin-reversed pure states, or (c) infinitely many (countably or uncountably many) incongruent pure-state pairs.*

Theorem 3.22 (Newman–Stein; as stated in [187, Thm. 6.2])**.** *Suppose that the periodic metastate κ_J is supported on nontrivial mixed states, that is, on Gibbs states containing more than one pure-state pair. Then the barycenter ρ_J is a single mixed Gibbs state supported on uncountably many pure states, and its pure-state decomposition has no atoms.*

So the only finite structures inside a periodic metastate are a single state and a single flip pair; anything richer is infinite, and once the states drawn from κ_J are nontrivial mixtures the barycenter is atomless.

Remark 3.23 (Equal self-overlaps). Equal EA order parameters for the pure states of one Gibbs state (Parisi's replica equivalence) are proved by Newman, Read and Stein, in the form that every translation- and coupling-invariant observable of a pure state is the same for $w_\Gamma \otimes w_\Gamma$-a.e. pair, for ν-a.e. J and κ_J-a.e. Γ, in short-range mixed p-spin models whose coupling laws have no isolated points and nonzero variance for every interaction set (every even set in the flip-invariant case) [191, Thm. 1]. These hypotheses exclude the nearest-neighbour EA model [191, remark (f)], so here they remain a hypothesis (Theorem 3.17).

3.6 The four classes

The structure theorems suggest two binary distinctions [190, pp. 11–14]: a metastate is *trivial* if κ_J is a point mass for ν-a.e. J, and *nontrivial* (dispersed) otherwise; a Gibbs state is *trivial* if its pure-state decomposition is a single orbit of the symmetry group (in zero field, a single pure state or a single flip pair), and *nontrivial* otherwise. If almost every state drawn from the metastate has the same character, there are four classes.

1) *Trivial metastate, trivial Gibbs states.* The droplet/scaling picture and the trivial–nontrivial picture of Chapter 2 (a single flip pair), and the high-temperature phase.

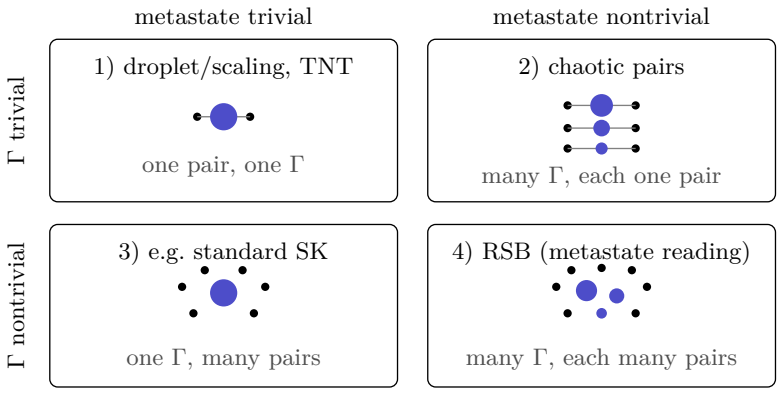


Figure 3.4: The Newman–Read–Stein classes. Dots are pure states, discs are Gibbs states drawn from κ_J (area $\approx$ weight). The standard SK picture is excluded by Theorem 3.12; a class-3 scenario with a self-averaging overlap distribution is not. At $T = 0$ only the top row survives; all four classes are otherwise open in every dimension [190, p. 13].

2) *Nontrivial metastate, trivial Gibbs states.* Chaotic pairs: each large volume sees essentially one flip pair, which changes chaotically with the volume [183, Sec. 8.1].

3) *Trivial metastate, nontrivial Gibbs states.* No standard physical proponent; Newman and Stein's standard SK picture is of this type.

4) *Nontrivial metastate, nontrivial Gibbs states.* Replica symmetry breaking in its metastate reading [221]; Newman and Stein's nonstandard SK picture.

For continuous coupling laws, ground states drawn from a zero-temperature metastate are single flip pairs, so at $T = 0$ only classes 1) and 2) occur [190, p. 14] (Fig. 3.4).

In every class P_J and $P_{J,\rho}$ are deterministic (Theorem 3.12), so non-self-averaging can only be a property of $P_{J,\Gamma}$ over a nontrivial metastate. In class 3 the nontrivial Gibbs state must have an atomless decomposition, and in class 4 the barycenter is atomless; if, as replica theory suggests, each Γ in class 4 has countably many pure states, the metastate itself must be atomless [190, p. 15] (Theorems 3.20 and 3.22). In class 2 a metastate supported on finitely many flip pairs is not excluded by these results (Open problem 3.24). Interfaces between distinct pure states drawn from the metastate are space-filling in every class (Theorems 3.16 and 3.17). In the CMR colouring of two periodic replicas, the number of infinite blue clusters is constrained by the kind of pure pair the replicas sampled: if two clusters never occur, classes 3 and 4 are excluded, and a count that is not almost surely constant excludes class 1; nothing finer is known (Section 8.7).

Replica theory predicts class 4 (Chapter 2): in Read's derivation [221] the metastate is dispersed over nontrivial mixtures, the barycenter has power-law correlations, and the metastate is nontrivial whenever $q(0) \neq \int_0^1 q(x)\,dx$. Newman and Stein argued heuristically against it: the invariance of Theorem 3.9 is hard to reconcile with a nontrivial distribution of weights over many pure states [183, Sec. 7.6], [185, p. 10]; their later joint review with Read records all four classes as open [190, p. 13].

Open problem 3.24 (Newman [171, Research Problem 4.14])**.** Determine whether for some d, ν and β the periodic metastate is, for ν-a.e. J, supported on (a) many Gibbs states, (b) Gibbs

states that decompose into more than two pure states, or (c) Gibbs states with the further properties of the nonstandard SK picture. Equally open: whether a chaotic-pairs metastate must be supported on infinitely many flip pairs [190, p. 15].

3.7 Overlap identities

In the Sherrington–Kirkpatrick model the decisive structural input is the Ghirlanda–Guerra identities for the overlaps of replicas [107], with stochastic stability [3]; in distributional form they imply ultrametricity, and with the Parisi formula they determine the overlap law. Their derivation uses only Gaussian integration by parts and energy concentration, so it works in short range; we give it because it shows what the identities can and cannot say about order.

Let $G=(V,E)$ be a finite graph with $N=|E|$ edges (for instance Λ_L with free boundary condition) and iid standard Gaussian $(J_e)_{e\in E}$, $H(\sigma)=-\sum_e J_e\sigma_e$, $\langle\cdot\rangle$ the product of Gibbs measures $\propto e^{-\beta H}$ for independent replicas $\sigma^1,\sigma^2,\dots$ at fixed J, and $\mathbb{E}\langle\cdot\rangle$ its disorder average. The *link overlap* of two replicas is

$$R_{ll'}=\frac{1}{N}\sum_{e\in E}\sigma_e^l\sigma_e^{l'}\in[-1,1],\qquad R_{ll}=1,$$

which for $G=\Lambda_L$ is the edge overlap q_Λ^e of Chapter 1, natural here as the covariance of the energy: $\mathbb{E}[H(\sigma^1)H(\sigma^2)]=N\,R_{12}$.

Gaussian integration by parts, $\mathbb{E}[gF(g)]=\mathbb{E}[F'(g)]$ for standard Gaussian g and C^1 functions F with F,F' of polynomial growth (as in Chapter 1), applied to the couplings one at a time computes the correlation of the energy with any observable.

Lemma 3.25 (Energy–overlap formula)**.** *Let f be a bounded function of $\sigma^1,\dots,\sigma^n$ that does not depend on J. Then*

$$\mathbb{E}\big\langle f\cdot(-H(\sigma^1))\big\rangle=\beta N\Big(\sum_{l=1}^{n}\mathbb{E}\langle fR_{1l}\rangle-n\,\mathbb{E}\langle fR_{1,n+1}\rangle\Big).$$

In particular $\mathbb{E}\langle -H\rangle=\beta N(1-\mathbb{E}\langle R_{12}\rangle)\in[0,2\beta N]$.

Proof. For an n-replica function g, differentiating the numerator and the n normalizations of $\langle g\rangle$ gives $\partial_{J_e}\langle g\rangle=\beta\sum_{l\le n}\langle g\,\sigma_e^l\rangle-n\beta\langle g\,\sigma_e^{n+1}\rangle$. Take $g=f\sigma_e^1$ and use $\sigma_e^1\sigma_e^1=1$: $\partial_{J_e}\langle f\sigma_e^1\rangle=\beta\sum_{l\le n}\langle f\sigma_e^1\sigma_e^l\rangle-n\beta\langle f\sigma_e^1\sigma_e^{n+1}\rangle$. Now $\mathbb{E}\langle f(-H(\sigma^1))\rangle=\sum_e\mathbb{E}[J_e\langle f\sigma_e^1\rangle]=\sum_e\mathbb{E}[\partial_{J_e}\langle f\sigma_e^1\rangle]$ by Gaussian integration by parts, and $\sum_e\sigma_e^l\sigma_e^{l'}=NR_{ll'}$. □

Proposition 3.26 (Stochastic stability; Contucci–Giardinà [60, Thms. 1–2])**.** *For f as in Lemma 3.25,*

$$\frac{d}{d\beta}\mathbb{E}\langle f\rangle=\beta N\,\Delta_n(f),\qquad \Delta_n(f)=\sum_{\substack{l,l'\le n\\ l\neq l'}}\mathbb{E}\langle fR_{ll'}\rangle-2n\sum_{l\le n}\mathbb{E}\langle fR_{l,n+1}\rangle+n(n+1)\,\mathbb{E}\langle fR_{n+1,n+2}\rangle.$$

Consequently, for $0\le\beta_1<\beta_2$, $\big|\int_{\beta_1}^{\beta_2}\Delta_n(f)\,\beta\,d\beta\big|\le 2\|f\|_\infty/N$.

The proof differentiates $\langle f\rangle$ in β and applies Lemma 3.25 to each replica (Section A.7).

The name: perturbing $-\beta H$ by $\sqrt{\lambda}K$ (K independent Gaussian, covariance R) shifts β^2 by λ/N, so the λ-derivative $\Delta_n(f)/2$ of the quenched state vanishes on average over temperature [60, Sect. 2]; the Ghirlanda–Guerra identities hold term by term.

Proposition 3.27 (Ghirlanda–Guerra identities for the link overlap)**.** *Let G_N be finite graphs with $N \to \infty$ edges and iid standard Gaussian couplings. For $n \geq 1$ and a measurable $f\colon [-1,1]^{n(n-1)/2} \to [-1,1]$ of the overlaps $(R_{ll'})_{l<l'\leq n}$, let*

$$\delta_N(\beta) = \mathbb{E}\langle f R_{1,n+1}\rangle - \frac{1}{n}\mathbb{E}\langle f\rangle\,\mathbb{E}\langle R_{12}\rangle - \frac{1}{n}\sum_{l=2}^{n}\mathbb{E}\langle f R_{1l}\rangle.$$

Then $\int_{\beta_1}^{\beta_2} |\delta_N(\beta)|\,d\beta \to 0$ for all $0 < \beta_1 < \beta_2$.

This is Contucci and Giardinà's form for short-range Gaussian models, including EA [61, Thm. 1]; Chatterjee proved the identities without averaging wherever the limiting pressure is differentiable, for a class including EA with the bond overlap [47, Thm. 1.2] (both use site normalization).

Proof. Apply Lemma 3.25 to f and to $f \equiv 1$ and subtract: a two-line computation gives

$$-n\,\delta_N(\beta) = \frac{1}{\beta N}\,\mathbb{E}\big\langle f\cdot(-H(\sigma^1) + \mathbb{E}\langle H\rangle)\big\rangle, \qquad \text{so} \qquad |\delta_N(\beta)| \leq \frac{1}{n\beta_1 N}\,\mathbb{E}\big\langle |H(\sigma^1) - \mathbb{E}\langle H\rangle|\big\rangle.$$

The right side is the thermal and disorder fluctuation of the energy per edge, whose integral over $[\beta_1,\beta_2]$ tends to zero by Lemma 3.28 below (a routine estimate, Section A.8). □

Lemma 3.28 (Concentration of the energy)**.** *In the setting of Proposition 3.27, for $0 < \beta_1 < \beta_2$,*

$$\int_{\beta_1}^{\beta_2} \frac{1}{N}\mathbb{E}\big\langle |H(\sigma) - \mathbb{E}\langle H\rangle|\big\rangle\,d\beta \leq \Big(\frac{2\beta_2(\beta_2-\beta_1)}{N}\Big)^{1/2} + C(\beta_2)\Big(\frac{\beta_2-\beta_1+1}{N^{1/4}}\Big),$$

where $C(\beta_2)$ depends only on β_2 (the proof in Section A.8 gives $C(\beta_2) = 8(\beta_2+1)$).

Route warning 3.29 (Overlap identities do not decide order). It is tempting to argue that the Ghirlanda–Guerra identities and stochastic stability force ultrametricity, hence a Parisi-type nontrivial overlap structure. Each step fails without a new input.

(i) The proof of Proposition 3.27 used nothing about the graph, so the identities hold on a path, where (free boundary condition) the edge variables are independent with means $\tanh(\beta J_e)$ and R_{12} concentrates at $\mathbb{E}\tanh^2(\beta J)$; an overlap law concentrated at one point satisfies every identity (both sides equal that point times $\mathbb{E}\langle f\rangle$). So the identities are compatible with a unique, paramagnetic state (strength: proved false as an implication).

(ii) They concern the link overlap; the site overlap needs overlap equivalence, not proved for the EA model and predicted to fail by the trivial–nontrivial picture (Chapter 2).

(iii) More is known only for an enlarged model: adding to the EA Hamiltonian independent Gaussian fields with covariances R^p, $p \geq 2$, and generic positive coefficients (not nearest-neighbour), Contucci, Mingione and Starr prove stochastic stability for every power and the Ghirlanda–Guerra identities in distribution, hence, by Panchenko's theorem, ultrametricity of the asymptotic link-overlap array [65, Thm. 3.1], which still allows equilateral triangles [65, Sec. 5]; for nearest-neighbour EA only the linear identities are proved.

(iv) They concern finite-volume replica measures; relating them to pure states or metastates requires exchanging limits, and the way replicas are drawn matters (Observation 3.14).

Arguin, Newman and Stein derive metastate analogues from Lemma 3.19, as evidence for Poisson–Dirichlet weights in countable mixtures [15].

3.8 Summary and outlook

In every dimension, covariance rules out the standard SK picture for iid couplings at every positive temperature, and a covariant ordered state for symmetric iid couplings, though not many states; it makes covariant interfaces empty or space-filling and restricts finite structures in the periodic metastate; the Ghirlanda–Guerra identities, proved for Gaussian couplings, are blind to order. None of this decides whether $N(d, \nu, \beta) > 1$: all four classes remain open (Open problem 3.24).

Chapter 4 turns coupling covariance into variance bounds, sharpest in $d = 2$; zero temperature is Chapter 5; graphical characterizations of the overlap are in Chapters 7 and 8.

Chapter 4

Fluctuations and the two-dimensional programme

Can the randomness of the couplings itself forbid many states? For the random-field Ising model the answer is the Imry–Ma argument [123]: flipping a region of linear size ℓ gains a random field energy of order $\ell^{d/2}$ and costs a surface energy of order ℓ^{d-1}, so the field wins for $d < 2$, the case $d = 2$ is marginal, and long-range order is predicted to be destroyed in $d \leq 2$. Aizenman and Wehr turned this into a theorem by comparing the fluctuations of a free-energy difference, of order $|\Lambda|^{1/2}$, with a deterministic surface bound, of order $|\partial_e \Lambda|$. For spin glasses the same comparison can be set up, with the coupling J_e playing the role of the field and the edge correlation $\Gamma(\sigma_e)$ playing the role of the magnetization. This chapter carries it out and explains what it gives.

The answer has three parts. In every dimension, two states drawn from metastates whose mean edge correlations differ have a free-energy difference whose variance grows like the volume; we prove this (Arguin, Newman, Stein and Wehr) in Section 4.2. In two dimensions, where the square root of the volume and the surface have the same size, the variance bound turns into a contradiction: for continuous coupling laws with $\mathbb{E}J^4 < \infty$ (so not for $\pm J$), all metastates have the same mean edge correlations, so covariant selections from one metastate cannot be told apart by their edge means (Section 4.3). The same comparison controls disorder chaos at zero temperature and, combined with a surface-order upper bound, would act in every dimension (Section 4.4). What the method cannot do is equally precise: it cannot separate a flip pair into its two pure states, so it cannot show that the two-dimensional model has no order, and in $d \geq 3$ the surface beats the square root of the volume. These limits are recorded as route warnings in Section 4.6, and the chapter closes with a table that sets the rigorous constraints of Chapter 3 and this chapter against the physical pictures of Chapter 2.

The two-dimensional target is the following conjecture, supported by all numerical evidence (Chapter 2).

Conjecture 4.1 (Uniqueness in two dimensions)**.** *For the EA model on $\mathbb{Z}^2$ with a symmetric coupling law, for every $\beta < \infty$ and ν-a.e. J, there is a unique infinite-volume Gibbs state.*

Newman's Research Problem 4.2 [171] asks for a proof of this, or of the weaker statement that there is no $\beta < \infty$ at which the Gibbs state is a.s. non-unique (Open problem 4.20).

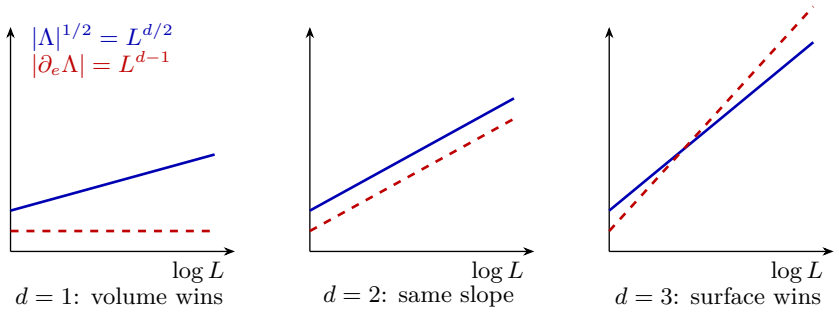


Figure 4.1: Volume against surface on a logarithmic scale. Solid: the size $|\Lambda|^{1/2}$ of the random fluctuations of a free-energy difference; dashed: the deterministic surface bound. In $d = 1$ the fluctuations eventually exceed the bound, so a variance bound alone gives a contradiction. In $d = 2$ the two have the same slope, and a Gaussian tail is needed (Section 4.3). In $d \geq 3$ the surface wins and there is no contradiction (Section 4.6).

4.1 Imry–Ma and Aizenman–Wehr

Consider a disordered system whose Hamiltonian contains a term $-\varepsilon \sum_x \eta_x \, g(\tau_x \sigma)$ with iid η_x of mean 0, variance 1 and all exponential moments finite, coupled to a local bounded observable g. In the random-field Ising model $g(\sigma) = \sigma_0$; in the EA model the couplings themselves are the disorder and $g(\sigma) = \sigma_0 \sigma_{e_i}$. The rounding theorem says that in low dimension the thermal average of the observable coupled to the disorder cannot jump.

Theorem 4.2 (Aizenman–Wehr [8, Thm. 4.1, Cors. 4.2–4.3])**.** *Let $d \leq 2$, and consider a system of bounded spins of the type of [8, Sect. 3], with disorder strength $\varepsilon \neq 0$. If the law of η has no isolated atoms, then for every $T > 0$, for almost every η, any two Gibbs states ρ, ρ' satisfy*

$$\lim_{L \to \infty} \frac{1}{|\Lambda_L|} \sum_{x \in \Lambda_L} \left[\rho(g \circ \tau_x) - \rho'(g \circ \tau_x) \right] = 0;$$

at $T = 0$ the same holds for ground states if the law of η is absolutely continuous. For the ferromagnetic random-field Ising model in $d \leq 2$ it follows that, at every $T > 0$ and field strength $\varepsilon \neq 0$, the Gibbs state is unique for almost every field.

The mechanism, in Bovier's summary [33, Sect. 2.4.2], is a dilemma. The free-energy difference between plus and minus boundary conditions in Λ is bounded by a constant times $|\partial_e \Lambda|$. Its random part has fluctuations of order $C(\beta)|\Lambda|^{1/2}$, where $C(\beta)$ is tied to the jump of the order parameter. If both hold, $C(\beta)|\Lambda|^{1/2} \leq \text{const}\, |\partial_e \Lambda|$, and in $d \leq 2$ this forces $C(\beta) = 0$. To make this work, Aizenman and Wehr had to show that the fluctuations are Gaussian on the scale $|\Lambda|^{1/2}$, by martingale methods, and to work with infinite-volume objects; this is where the metastate first appeared [8, Sect. 5]. The passage from "no jump" to uniqueness uses the FKG inequality: in the random-field Ising model the plus and minus states are the extreme states, and equal magnetizations force them to coincide. Chapter 10 returns to the random-field model in full.

For the EA model with a symmetric coupling law, the observable conjugate to J_e is σ_e, and the theorem concerns the spatial average of edge correlations. Aizenman and Wehr note that for symmetric laws their general theorem "yields no new information even for $d = 2$ dimensions, since its conclusion holds by the model's 'gauge symmetry' " [8, Remark after Cor. 4.3]. Two

obstacles stand between the rounding argument and spin glasses, as Arguin, Newman, Stein and Wehr explain [14, Sect. 2]. First, there is no known way to connect spin-glass Gibbs states to boundary conditions "in a translation-invariant, measurable way", so one cannot follow a given state as the couplings change. Second, there is a cancellation problem: unlike in a ferromagnet, the sign of the energy difference across an interface varies with the couplings, so contributions of order one per edge can cancel. Metastates and coupling covariance (Chapter 3) solve the first; an assumption of asymmetry on average, below, addresses the second, and it buys a statement for each edge separately instead of a spatial average.

4.2 A volume-order variance bound

Fix $d \geq 1$, $\beta < \infty$, and a coupling law ν that is continuous (has no atoms) with $\mathbb{E}J^4 < \infty$. Let κ, κ' be two metastates in the sense of Chapter 3, coupling covariant and translation covariant, and let

$$M = \nu(dJ)\,\kappa_J(d\Gamma)\,\kappa'_J(d\Gamma')$$

be the joint law of the couplings and of two states drawn from them. For a box $\Lambda = \Lambda_L$ and $H_{\Lambda,J}(\sigma) = -\sum_{e \in E(\Lambda)} J_e \sigma_e$, the *free-energy difference* of the two states in Λ is

$$F_\Lambda(J, \Gamma, \Gamma') = \log \Gamma\big(e^{\beta H_{\Lambda,J}}\big) - \log \Gamma'\big(e^{\beta H_{\Lambda,J}}\big) \tag{4.1}$$

[14, Eq. (3)]. By the DLR equations, $\Gamma(e^{\beta H_{\Lambda,J}})$ is, up to a factor $e^{\pm 2\beta \sum_{e \in \partial_e \Lambda} |J_e|}$, the inverse partition function of Λ with boundary condition drawn from Γ, so F_Λ is β times the difference of the free energies of Λ with boundary conditions drawn from Γ and from Γ', up to boundary terms. Here, as in Chapter 1, $\partial_e \Lambda$ is the set of edges with exactly one endpoint in Λ.

Definition 4.3 (Asymmetry on average; [14, Assumption 4.1])**.** The pair (κ, κ') is *asymmetric on average* if for some edge e

$$\nu\big\{J : \ \kappa_J(\Gamma(\sigma_e)) \neq \kappa'_J(\Gamma'(\sigma_e))\big\} > 0,$$

that is, if the barycenters ρ_J, ρ'_J differ in some nearest-neighbour correlation with positive probability.

By translation covariance the condition then holds for every translate of e. It implies that $\kappa_J \otimes \kappa'_J$ samples incongruent pairs. As Arguin, Newman and Stein put it later, the bound below needs "an inherent asymmetry" between the two states "on average" [17, Sect. 1.3]: the difference $\Gamma(\sigma_e) - \Gamma'(\sigma_e)$ must not only fluctuate but have nonzero mean.

Theorem 4.4 (Arguin–Newman–Stein–Wehr [14, Thm. 4.2])**.** *If (κ, κ') is asymmetric on average, there are $c > 0$ and L_0 such that $\operatorname{Var}_M(F_\Lambda) \geq c\,|\Lambda|$ for every box $\Lambda = \Lambda_L$ with $L \geq L_0$.*

The proof has three ingredients: a lower bound for the variance of a function of independent blocks, a formula for the derivative of the conditional mean of F_Λ with respect to a coupling, and the asymmetry assumption to make each block contribute. Figure 4.2(a) shows the structure.

Lemma 4.5 (Wehr–Aizenman [245, Prop. 3.1])**.** *Let X be a square-integrable function of independent random elements $Y_1, \dots, Y_N, Z$. Then $\operatorname{Var} X \geq \sum_{k=1}^{N} \operatorname{Var}(\mathbb{E}[X \mid Y_k])$.*

Proof. Let $X_k = \mathbb{E}[X \mid Y_1, \dots, Y_k]$, with $X_0 = \mathbb{E}X$. The increments $X_k - X_{k-1}$ are orthogonal, so $\operatorname{Var} X \geq \operatorname{Var} X_N = \sum_k \mathbb{E}[(X_k - X_{k-1})^2]$. Since X_{k-1} is independent of Y_k, $\mathbb{E}[X_k - X_{k-1} \mid Y_k] = \mathbb{E}[X \mid Y_k] - \mathbb{E}X$, and Jensen's inequality gives $\mathbb{E}[(X_k - X_{k-1})^2] \geq \operatorname{Var}(\mathbb{E}[X \mid Y_k])$. □

For a finite box $B \subset \Lambda$ write J_B for the couplings on $E(B)$, $J^{B,0}$ for J with J_B set to zero, and

$$M_{B^c} = \nu(dJ_{B^c})\, \kappa_{J^{B,0}}(d\Gamma)\, \kappa'_{J^{B,0}}(d\Gamma').$$

Let $\delta_e(\Gamma, \Gamma') = \Gamma(\sigma_e) - \Gamma'(\sigma_e)$.

Lemma 4.6 (Smooth conditional mean; [14, Lemmas 5.3–5.4])**.** *Let $B \subset \Lambda$ be a finite box.*

(a) If $\Phi(J, \Gamma, \Gamma')$ is integrable, then the function

$$J_B \longmapsto M_{B^c}[\Phi(J, L_{J_B}\Gamma, L_{J_B}\Gamma')],$$

defined for every $J_B \in \mathbb{R}^{E(B)}$, is a version of $M(\Phi \mid J_B)$.

(b) The version $G_B(J_B)$ of $M(F_\Lambda \mid J_B)$ given by (a) is continuously differentiable, with $\partial_{J_e} G_B = -\beta\, m_e$ for $e \in E(B)$, where $m_e(J_B) = M_{B^c}[\delta_e(L_{J_B}\Gamma, L_{J_B}\Gamma')]$ is the version of $M(\delta_e \mid J_B)$ given by (a).

Proof. (a) Coupling covariance gives $\kappa_J = (L_{J_B})_* \kappa_{J^{B,0}}$, and J_B is independent of J_{B^c}, so for every bounded $g(J_B)$ one has $M(\Phi\, g) = \int \nu(dJ_B)\, g(J_B)\, M_{B^c}[\Phi(J, L_{J_B}\Gamma, L_{J_B}\Gamma')]$, which is the claim. (b) Put $X_B = \exp(\beta \sum_{e \in E(B)} J_e \sigma_e)$. Since $E(B) \subset E(\Lambda)$,

$$\log(L_{J_B}\Gamma)(e^{\beta H_{\Lambda, J}}) = \log \Gamma(e^{\beta H_{\Lambda, J^{B,0}}}) - \log \Gamma(X_B),$$

and the first term does not depend on J_B. Differentiating the second gives $-\beta(L_{J_B}\Gamma)(\sigma_e)$. These derivatives are bounded by β and continuous in J_B, and $|F_\Lambda| \leq 4\beta \sum_{e \in \partial_e \Lambda} |J_e|$ is integrable (Lemma 4.7 below), so differentiation under M_{B^c} is justified by dominated convergence. □

The sign of the derivative depends on the orientation chosen in (4.1) and plays no role.

Proof of Theorem 4.4. Step 0. $\operatorname{Var}_M(F_\Lambda) \geq \operatorname{Var}(G)$ with $G = M(F_\Lambda \mid J_\Lambda)$.

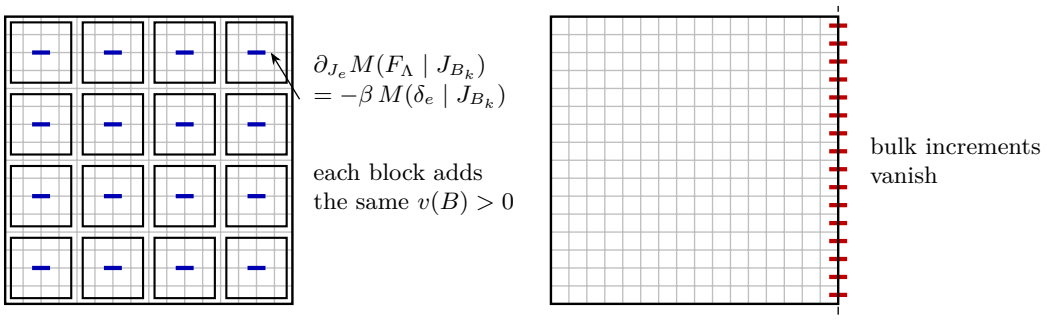
Step 1 (blocks). Fix a cube B, to be chosen in Step 3, and let $B_1, \dots, B_N$ be disjoint translates of B inside Λ, with $N \geq c_B|\Lambda|$ once $L \geq L_0$, the side length of B. G is a function of the independent families $J_{B_1}, \dots, J_{B_N}$ and the remaining couplings in Λ, so Lemma 4.5 and the tower property give $\operatorname{Var}(G) \geq \sum_{k=1}^N \operatorname{Var}(M(F_\Lambda \mid J_{B_k}))$.

Step 2 (every block contributes the same). For a block B_k, let Z', Z'' be independent copies of J_{B_k}. Then $\operatorname{Var}(M(F_\Lambda \mid J_{B_k})) = \frac{1}{2}\mathbb{E}[(G_{B_k}(Z'') - G_{B_k}(Z'))^2]$, and by Lemma 4.6(b)

$$G_{B_k}(Z'') - G_{B_k}(Z') = -\beta \int_0^1 \sum_{e \in E(B_k)} m_e(Z' + s(Z'' - Z'))(Z''_e - Z'_e)\, ds.$$

The right side involves only the law of $(J_{B_k^c}, \Gamma, \Gamma')$ under $M_{B_k^c}$, which does not depend on Λ and, by translation covariance of the metastates and invariance of ν, does not depend on k. So $\operatorname{Var}(M(F_\Lambda \mid J_{B_k})) = v(B)$ for all k and all $\Lambda \supset B_k$.

Step 3 (each block contributes something). Let e be an edge for which Definition 4.3 holds, and let B range over cubes with central edge e. As B increases to $\mathbb{Z}^d$, the martingale convergence theorem gives $M(\delta_e \mid J_B) \to M(\delta_e \mid J) = \kappa_J(\Gamma(\sigma_e)) - \kappa'_J(\Gamma'(\sigma_e))$ a.s., and the limit is nonzero



(a) blocks $B_k \subset \Lambda$, $N \asymp |\Lambda|$ (b) seam S of the torus, $|S| \asymp L^{d-1}$

Figure 4.2: Two martingale arguments. (a) Lower bound (Theorem 4.4): reveal the couplings block by block; by translation covariance every block contributes the same amount to the variance, and asymmetry on average makes this amount positive, so the variance is at least of order $|\Lambda|$. (b) Upper bound (Chapter 2): periodic and antiperiodic boundary conditions differ by the signs of the seam couplings; revealing all other couplings first gives conditional mean zero, so only the seam contributes, and the variance is at most of order $|S|$.

with positive probability by asymmetry on average. Fix B so large that m_e is not a.s. zero. If $v(B) = 0$, then G_B is a.s. constant, hence constant on the support $(\operatorname{supp}\nu)^{E(B)}$, because it is continuous. Since ν has no atoms, its support has no isolated points, so every point of the support is a limit, along each coordinate direction, of other points of the support; hence $\nabla G_B = 0$ on the support, and $m_e = 0$ a.s., a contradiction. So $v(B) > 0$.

Conclusion. For $L \geq L_0$, $\operatorname{Var}_M(F_\Lambda) \geq N\, v(B) \geq c_B v(B)\, |\Lambda|$. □

The hypotheses enter at identifiable places: coupling covariance in Lemma 4.6, translation covariance in Step 2, continuity of ν and asymmetry on average in Step 3. Inspecting the proof, the product form of M enters only through these two covariances of the conditional law of (Γ, Γ') given J and through the two J-measurable means $\kappa_J(\Gamma(\sigma_e))$ and $\kappa'_J(\Gamma'(\sigma_e))$.

4.3 Two dimensions

The free-energy difference has a deterministic bound by the boundary couplings.

Lemma 4.7 (Surface bound; [14, Prop. 6.1, Lemma 6.2])**.** *If Γ and Γ' are Gibbs states for J, then*

$$|F_\Lambda(J, \Gamma, \Gamma')| \leq 4\beta \sum_{e \in \partial_e \Lambda} |J_e|, \qquad \textit{and consequently} \qquad |M(F_\Lambda \mid J_\Lambda)| \leq 4\beta\, |\partial_e \Lambda|\, \mathbb{E}|J|.$$

Proof. Let $K = \sum_{e \in \partial_e \Lambda} |J_e|$ and let μ_Λ be the Gibbs measure on Λ with free boundary condition. For $f \geq 0$ depending only on σ_Λ, the DLR equations write $\Gamma(f)$ as the Γ-average of a ratio in which the boundary interaction $e^{-\beta H_{\partial_e \Lambda}}$, a number between $e^{-\beta K}$ and $e^{\beta K}$, appears in the numerator and the denominator; hence $e^{-2\beta K} \mu_\Lambda(f) \leq \Gamma(f) \leq e^{2\beta K} \mu_\Lambda(f)$. Apply this with $f = e^{\beta H_{\Lambda, J}}$ to Γ and to Γ'. The boundary couplings are independent of J_Λ, which gives the second bound. □

Compare this with Theorem 4.4, which gives $\operatorname{Var} M(F_\Lambda \mid J_\Lambda) \geq c|\Lambda|$ for large Λ (Steps 0–3 bound the variance of the conditional mean). The conditional mean is bounded by $C|\partial_e\Lambda|$, so its variance is at most $C^2|\partial_e\Lambda|^2$. For $d = 1$ this already contradicts $c|\Lambda|$. For $d = 2$ both are of order L^2 (Fig. 4.1), and one needs more than the variance: that the conditional mean, rescaled by $|\Lambda|^{1/2}$, has at least a Gaussian tail. This follows from a martingale central limit argument going back to Aizenman and Wehr.

Proposition 4.8 (Gaussian lower bound; [14, Prop. 6.3 and Sect. 6.2], after [8, Prop. A.2.1])**.** *Let $d = 2$ and suppose (κ, κ') is asymmetric on average. Write $G_\Lambda = M(F_\Lambda \mid J_\Lambda)$. There is $c > 0$ such that for every $t \in \mathbb{R}$,*

$$\liminf_{L\to\infty} \mathbb{E} \exp\Bigl(t\,\frac{G_\Lambda - \mathbb{E}G_\Lambda}{|\Lambda|^{1/2}}\Bigr) \geq e^{ct^2}.$$

Proof sketch. This sketch omits the uniformity estimates, which are in Section A.9. Reveal the couplings of Λ one edge at a time in lexicographic order. The martingale increments are bounded by $2\beta(|J_e| + \mathbb{E}|J|)$, so the fourth-moment assumption gives the Lindeberg condition. By translation covariance the conditional variances of the increments are, up to boundary effects, translates of one function, so by the L^2 ergodic theorem their average converges to a constant σ^2, which is positive because it is the limit of $\operatorname{Var} M(F_\Lambda \mid J_\Lambda)/|E(\Lambda)|$ and Theorem 4.4 bounds this from below. A martingale central limit bound for the moment generating function then gives the claim. □

Theorem 4.9 (Arguin–Newman–Stein–Wehr [14, Cor. 4.3])**.** *Let $d = 2$, $\beta < \infty$, and let ν be continuous with $\mathbb{E}J^4 < \infty$. For every pair of metastates κ, κ' and every edge e,*

$$\kappa_J(\Gamma(\sigma_e)) = \kappa'_J(\Gamma'(\sigma_e)) \qquad \text{for } \nu\text{-a.e. } J;$$

that is, all metastates have barycenters with the same nearest-neighbour correlations.

Proof. If not, the pair is asymmetric on average. For $\Lambda = \Lambda_L \subset \mathbb{Z}^2$ one has $|\partial_e\Lambda| = 4|\Lambda|^{1/2}$, so by Proposition 4.8, for $t > 0$, $\liminf_L \mathbb{E}\exp(t\,(G_\Lambda - \mathbb{E}G_\Lambda)/|\partial_e\Lambda|) \geq e^{ct^2/16}$, while by Lemma 4.7, $|G_\Lambda - \mathbb{E}G_\Lambda| \leq 8\beta\,\mathbb{E}|J|\,|\partial_e\Lambda|$, so the left side is at most $e^{8\beta\mathbb{E}|J|\,t}$. These are incompatible for large t. □

Since the theorem holds for every pair of metastates in the sense of Chapter 3, it applies in particular to two metastates obtained from one metastate by a covariant selection of its states: in two dimensions, covariant selections from one metastate cannot be told apart by their mean edge correlations. Weights give such selections: the set of weights of the atoms of a metastate is invariant under translations and under finite coupling changes, so it is a.s. constant, and the atoms carrying a given weight form a covariant family. Arguin, Newman, Stein and Wehr use these tags to derive restrictions on metastates with atomic support from Theorem 4.9 [14, Cor. 4.4 and Sect. 7].

Two caveats keep the result in proportion. It concerns metastates, that is, covariant objects such as those built from coupling-independent boundary conditions; and, as its authors stress, such results are "far from the conjecture that in $d = 2$ there exists a single Gibbs state" [14, remark after Cor. 4.4].

4.3.1 Zero temperature

At zero temperature the role of the free energy is played by the energy difference $\Delta H_\Lambda(J,\sigma,\sigma') = H_{\Lambda,J}(\sigma) - H_{\Lambda,J}(\sigma')$ of two ground states drawn from zero-temperature metastates. To differentiate with respect to a coupling one needs to know how a ground state responds to a change of couplings inside a box, and this information is carried by the *excitation metastate* [16, Def. 5.2], which records, besides the ground state, the energy costs of all local rearrangements in each finite box.

For excitation metastates κ, κ' built from periodic boundary conditions (more generally, any two with the covariance properties of [16, Def. 5.2]), Arguin, Newman, Stein and Wehr prove the analogue of Theorem 4.4: if $\nu\{J : \kappa_J(\sigma_e) \neq \kappa'_J(\sigma_e)\} > 0$ for some edge e, then $\mathrm{Var}_M(\Delta H_\Lambda) \geq c|\Lambda|$ [16, Thm. 5.5]. In two dimensions this gives the zero-temperature form of Theorem 4.9.

Theorem 4.10 (Arguin–Newman–Stein–Wehr; from [16, Thm. 5.5 and Sect. 6])**.** *Let $d = 2$, $T = 0$, and let ν be continuous and symmetric with $\mathbb{E}J^4 < \infty$. Any two translation- and coupling-covariant zero-temperature kernels (excitation metastates in the sense of [16, Def. 5.2]) κ, κ' have the same mean edge correlations: $\kappa_J(\sigma_e) = \kappa'_J(\sigma_e)$ for every edge e, ν-a.s.*

Proof sketch. This sketch omits the construction of excitation metastates and the uniformity estimates. The argument, given in outline in [16, Sect. 6], follows the proofs of Theorems 4.4 and 4.9 with F_Λ replaced by ΔH_Λ. The derivative of the conditional mean is now $-M(\sigma_e - \sigma'_e \mid J_B)$ [16, Lemma 5.3]; the surface bound is $|\Delta H_\Lambda| \leq 4\sum_{e\in\partial_e\Lambda} |J_e|$, and the moment generating function has a Gaussian lower bound as in Proposition 4.8. □

At zero temperature edge variables take the values ± 1, which makes one consequence immediate.

Remark 4.11 (after [16, Sect. 6])*.* In the setting of Theorem 4.10, suppose a zero-temperature metastate is atomic and some weight value is carried by exactly one ground-state pair α. Then the metastate is supported on α alone. Indeed, translations and finite coupling changes act on the atoms of the metastate by bijections that preserve their weights, so the atom of the given weight is selected covariantly and δ_α is a covariant kernel (this is the tagging by weights of [16, Sect. 6]); so by the theorem $\alpha_e = \sum_\gamma W_\gamma \gamma_e$ for every edge e, where $\gamma_e = \gamma_x\gamma_y \in \{\pm 1\}$ and the sum runs over the pairs γ of the metastate with their weights W_γ. A convex combination of numbers in $\{\pm 1\}$ equals ± 1 only if all of them do, so every pair of positive weight has the same edge variables as α, hence equals α. Arguin, Newman, Stein and Wehr derive further restrictions on atomic zero-temperature metastates in [16, Sect. 6].

Two further zero-temperature results in two dimensions do not use variance bounds and are treated in Chapter 5. For continuous symmetric coupling laws of unbounded support, two ground-state pairs drawn from periodic metastates are equal or differ by a single domain wall of positive density [181, Thm. 1], and in the half-plane with free boundary condition on the boundary line the metastate is supported on a single pair [13, Thm. 1.2].

What these results leave open at zero temperature is the existence of metastates on several pairs at all.

Open problem 4.12 (Several ground-state pairs in two dimensions; compare [16, Sect. 6])**.** Show that in $d = 2$ no zero-temperature periodic metastate is supported on $M \geq 2$ incongruent ground-state pairs; for instance, that none gives weight $\frac{1}{2}$ to each of two incongruent pairs. The analogous questions at positive temperature, and for atomless metastates, are open as well.

Newman and Stein have developed restricted metastates, which group the pure states of a metastate by their edge overlap with a randomly chosen reference state, together with bounds on the fluctuations of free-energy and energy differences between incongruent states, as part of a programme toward uniqueness in two dimensions [187–189]. Among their results, for continuous symmetric coupling laws with finite moments, two zero-temperature metastates that are each supported on a single ground-state pair are supported on the same pair, for periodic, antiperiodic, free and fixed boundary conditions [189, Sect. 6.1].

4.4 Upper bounds and disorder chaos

If one had an upper bound of surface order for the *variance* of the same free-energy difference, the comparison with Theorem 4.4 would act in every dimension: $c\,L^d \le C\,L^{d-1}$ is false for large L whatever d is. Such a bound is known for finite volumes with gauge-related boundary conditions: the surface variance bound of Aizenman–Fisher and Newman–Stein, proved in Chapter 2 as stated in [227, Thm. 1] and [17, Thm. 5.1], gives $\mathrm{Var}(\log Z_{\mathrm{P}} - \log Z_{\mathrm{AP}}) \le 4\beta^2 \mathbb{E}[J^2]\, L^{d-1}$ for periodic and antiperiodic boundary conditions on a box of side L and a symmetric coupling law. Antiperiodic boundary conditions are periodic ones with the couplings on one seam reversed; revealing the couplings off the seam first gives conditional mean zero, so only the L^{d-1} seam couplings contribute (Fig. 4.2(b)).

In the language of the stiffness exponent of Chapter 2, this is the rigorous form of the bound $\theta \le (d-1)/2$ [17, Sect. 5.1]. The obstacle to combining it with Theorem 4.4 is that the two concern different quantities: a finite-volume periodic/antiperiodic difference, and the restriction to a box of a free-energy difference between two infinite-volume states. As Arguin, Newman, Stein and Wehr note, the lower bound "is for a quantity that differs somewhat from that for which an upper bound has been obtained, which prevents us from applying of our result to dimensions higher than two" [14, Sect. 1].

At zero temperature Arguin, Newman and Stein connected the variance of ground-state energy differences with disorder chaos. Let ν be standard Gaussian, let $\kappa^{(1)}, \kappa^{(2)}$ be zero-temperature metastates built from coupling-independent boundary conditions, and let

$$P = \nu(dJ)\,\nu(dJ')\,\kappa_J^{(1)}(d\sigma^1)\,\kappa_J^{(2)}(d\sigma^2).$$

Inside Λ interpolate $J_e(t) = e^{-t}J_e + \sqrt{1-e^{-2t}}\,J'_e$, and let $\sigma^i(t)$ be the corresponding ground-state paths [17, Sect. 2]. Let q^e_Λ be the edge overlap over the edges $E(\Lambda)$ of Λ (Chapter 1), and I the event that $\limsup_\Lambda q^e_\Lambda(\sigma^1,\sigma^2) < 1$ (incongruence).

Theorem 4.13 (Arguin–Newman–Stein [17, Thm. 1.2, Def. 1.3, Cor. 1.4])**.** *For all $t > 0$,*

$$\mathrm{Var}\big(H_{\Lambda,J}(\sigma^1) - H_{\Lambda,J}(\sigma^2)\big) \ge 2|E(\Lambda)| \int_0^t \Big\{ \mathbb{E}[1 - q^e_\Lambda(\sigma^1,\sigma^2)] - \sum_{i=1,2} \big(2\,\mathbb{E}[1 - q^e_\Lambda(\sigma^i,\sigma^i(s))]\big)^{1/2} \Big\} e^{-s}\,ds.$$

Say that there is absence of disorder chaos at scale $a \in [0,1]$ *if for every $\varepsilon > 0$ there are events A_ε with $P(A_\varepsilon) > 1-\varepsilon$ and $C(\varepsilon) > 0$ such that $q^e_\Lambda(\sigma^i,\sigma^i(t)) > 1-\varepsilon$ on A_ε for $i = 1,2$, all $t \le C(\varepsilon)|\Lambda|^{-a}$ and all large Λ. If $P(I) > 0$ and there is absence of disorder chaos at scale a, then $\mathrm{Var}(H_{\Lambda,J}(\sigma^1) - H_{\Lambda,J}(\sigma^2)) \ge C|\Lambda|^{1-a}$ for some $C > 0$.*

So incongruence without chaos forces large fluctuations. The authors also show that the variance is bounded below by a constant whenever $P(I) > 0$, and that a bound $|\partial D_b| \le C|\Lambda|^{\gamma}$ on the critical droplets of single couplings gives absence of chaos at every scale $a > 2\gamma$ [17, Thm. 1.5, Cor. 1.7]. To turn this into a statement about the existence of incongruent states one needs an upper bound of the type of the surface variance bound for the same quantity.

Conjecture 4.14 (Arguin–Newman–Stein [17, Conj. 5.3])**.** *For σ^1, σ^2 as above,* $\mathrm{Var}(H_{\Lambda,J}(\sigma^1) - H_{\Lambda,J}(\sigma^2)) \le A\, L^{d-1}$ *with A independent of $\Lambda = \Lambda_L$.*

Corollary 4.15 (Conditional; [17, Claim 5.4])**.** *Assume Conjecture 4.14. If there is absence of disorder chaos at some scale $a < 1/d$, then $P(I) = 0$: zero-temperature metastates built from coupling-independent boundary conditions are not supported on incongruent ground states.*

Proof. If $P(I) > 0$, Theorem 4.13 and the conjecture give $CL^{d(1-a)} \le AL^{d-1}$ for all large L, so $a \ge 1/d$. □

This is the precise sense of "chaos ⇔ incongruence": incongruent states can exist only if the ground states are chaotic under perturbations of size $|\Lambda|^{-1/d}$ inside Λ, a relation the authors connect to the droplet theory of Fisher and Huse [17, Sect. 5]. Disorder chaos itself is the subject of Chapter 6.

4.5 Bond observables in two dimensions

Dario, Harel and Peled quantified the Aizenman–Wehr argument for general disordered spin systems in two dimensions [68]. For the EA model their setting is $\mathbb{Z}^2$ with independent Gaussian couplings $J_e = \lambda\eta_e + h$ (η_e standard Gaussian, $\lambda > 0$, $h \in \mathbb{R}$), and the observable coupled to the disorder is the vector $f_v(\sigma) = (\sigma_v\sigma_{v+e_1}, \sigma_v\sigma_{v+e_2})$. Their results control this observable for *all* boundary conditions, including ones chosen as functions of the couplings, and do not use coupling covariance.

Theorem 4.16 (Dario–Harel–Peled [68, Thms. 1 and 3])**.** *Let $d = 2$, $\beta > 0$, $\lambda > 0$.*

(a) There are $C, c > 0$ depending only on λ (and on the model, but not on β) such that for every $L \ge 3$,

$$\mathbb{P}\Big(\sup_{\xi_1,\xi_2}\Big|\frac{1}{|\Lambda_L|}\sum_{v\in\Lambda_L}\big(\langle f_v\rangle^{\xi_1}_{\Lambda_L} - \langle f_v\rangle^{\xi_2}_{\Lambda_L}\big)\Big| > \frac{C}{(\ln\ln L)^{1/4}}\Big) \le e^{-cL},$$

the supremum being over all boundary conditions.

(b) For any two translation-covariant Gibbs states $\eta \mapsto \mu_1^\eta, \mu_2^\eta$, $\langle f_0\rangle_{\mu_1^\eta} = \langle f_0\rangle_{\mu_2^\eta}$ almost surely.

Part (b) is the analogue of Theorem 4.9 for translation-covariant Gibbs states that need not come from a metastate: all of them have the same nearest-neighbour correlations. The authors conjecture the pointwise version of (a), that the boundary condition has no effect on $\langle f_0\rangle$ in the limit, uniformly over boundary conditions [68, Conj. 2.1], and they record that the zero-temperature versions of their results hold; in particular, if a translation-covariant ground state supported on flip pairs exists, then all translation-covariant ground states are supported on such pairs [68, p. 10].

4.6 Route warnings

Route warning 4.17 (Volume against surface works only in $d \leq 2$). The route "volume-order fluctuations of a free-energy difference contradict its surface bound, hence no incongruent states" needs $|\Lambda|^{1/2}$ to be at least of the order of $|\partial_e \Lambda|$. In $d \geq 3$, $L^{d/2} \ll L^{d-1}$, and the argument gives nothing even if every variance bound of Section 4.2 is granted (strength: fails even under optimistic assumptions). The way around is to compare variances with variances, which needs a surface-order variance bound for the same quantity (Conjecture 4.14); the known surface variance bound of Chapter 2 concerns a different one. Such a bound would exclude incongruent states in every dimension, so it is not an easier lemma than the question it is meant to settle.

Route warning 4.18 (A unique metastate does not give $q_{\mathrm{EA}} = 0$). The route "all two-dimensional metastates agree, or the periodic metastate is supported on one pair, hence the two-dimensional EA model has no order at positive temperature" fails outright (strength: proved false as an implication). For the two-dimensional Ising ferromagnet below its critical temperature, the periodic measures converge to $\frac{1}{2}(\mu^+ + \mu^-)$, so the periodic metastate is a single point mass, yet $q_{\mathrm{EA}} > 0$ (Chapter 3). In the same way, a metastate supported on one flip pair $\Gamma = \frac{1}{2}(\alpha + \bar{\alpha})$ with $q_{\mathrm{EA}} > 0$ is consistent with every result of this chapter. The variance method compares states only through the observables conjugate to the couplings, the edge correlations, and a flip pair has identical even correlations: $\alpha(\sigma_A) = \bar{\alpha}(\sigma_A)$ for every finite set A of even size. So no comparison of this kind can exclude a decomposition $\Gamma = \frac{1}{2}(\alpha + \bar{\alpha})$ with $\alpha \neq \bar{\alpha}$; this also applies to Theorem 4.16 and to its conjectured pointwise form. A proof of Conjecture 4.1 needs an input that detects long-range order directly, such as decay of $\langle \sigma_0 \sigma_x \rangle_J^2$ at large $|x|$, or an order parameter conjugate to some randomness together with monotonicity, which for the random-field Ising model is supplied by the FKG inequality.

Route warning 4.19 (Variance bounds see a metastate only through its edge means). We record an observation about the method class of this chapter. The bounds of Sections 4.2 and 4.3 are proved through the J-conditional mean $M(F_\Lambda \mid J_\Lambda)$ of a free-energy or energy difference, whose derivatives with respect to the couplings are differences of mean edge correlations (Lemma 4.6). Consider a metastate giving weight $\frac{1}{2}$ to each of two incongruent flip pairs, and a pair law for (Γ, Γ') that is symmetric under exchanging the two pairs; a pair law that is not symmetric would itself single out one of the pairs covariantly. Since $F_\Lambda(J, \Gamma', \Gamma) = -F_\Lambda(J, \Gamma, \Gamma')$, the J-conditional mean of F_Λ vanishes identically (the conditional law of F_Λ itself need not be trivial), exactly as it does for a metastate supported on a single pair. So bounds of this kind cannot tell the two situations apart (strength: proved, for this method class). Excluding such metastates needs either a covariant labelling of the pairs, which would already give asymmetry on average, or information beyond J-conditional means.

4.7 The pictures against the constraints

Table 4.1 collects the rigorous constraints of Chapter 3 and this chapter and records how each bears on the pictures of Chapter 2. Only the standard SK picture is excluded outright, in every dimension. In two dimensions the remaining many-state classes survive only in forms in which all covariant selections of states have the same mean edge correlations. In $d \geq 3$ no rigorous result excludes chaotic pairs or the metastate reading of replica symmetry breaking. Table 4.2 places the numerical evidence beside this; it is heuristic and is discussed in Chapter 2.

Constraint (where)	D/S	TNT	CP	RSB	class 3
Covariant overlap laws self-average (Chapter 3)	✓	✓	✓	✓[a]	✓[b]
Metastate invariant under flip-related b.c. (Chapter 3)	✓	✓	✓	✓	✓
Interfaces empty or space-filling (Chapter 3)	✓	✓	○	○	○
State, pair or infinitely many pairs; no countable unequal weights; atomless barycenter (Chapter 3)	✓	✓	✓	○	○
Linear Ghirlanda–Guerra identities, link overlap (Chapter 3)	✓	✓	✓	✓	✓
Volume-order variance under asymmetry, any d (Theorem 4.4)	–	–	○	○	○
$d = 2$, $T > 0$: all metastates have equal mean edge correlations (Theorem 4.9)	✓	✓	○	○	○
$d = 2$, $T = 0$: the same for zero-temperature kernels (Theorem 4.10)	✓	✓	○	n/a	n/a
Incongruence needs chaos at scale $\geq 1/d$, given Conjecture 4.14 (Corollary 4.15)	✓	✓	○	○	n/a

Table 4.1: Rigorous constraints against the pictures: droplet/scaling (D/S), trivial–nontrivial (TNT), chaotic pairs (CP), replica symmetry breaking in the metastate reading (RSB, class 4) and class 3 of Chapter 3. ✓: consistent with the picture as usually stated; ○: consistent, but the constraint restricts the form the picture can take; –: no bearing; n/a: the picture concerns only $T > 0$ (at $T = 0$ only classes 1 and 2 occur). [a]Non-self-averaging only over the states drawn from the metastate. [b]The overlap distribution must be deterministic. The standard SK picture is excluded by the first row.

Picture	Evidence cited for	Evidence cited against
RSB	simulations in $d = 3$ down to $0.64\,T_c$ at $L \leq 32$ fit RSB at experimental length scales [125]; a numerical Aizenman–Wehr metastate in $d = 3$ appears dispersed [29]	domain-wall fractal dimension below d for $d \leq 5$ (approximate ground-state methods) [243]
Droplet/scaling	$d_s < d$ for interfaces [243]	the $d = 3$ data of [125]
TNT	bulk excitations whose energy exponent is $\theta' \approx 0$, with $d - d_s \approx 0.42$, in $d = 3$ [205]	the authors note a possible crossover at larger sizes [205]
Chaotic pairs	no direct test; needs space-filling interfaces	as for RSB [243]

Table 4.2: Numerical evidence for the pictures (heuristic; finite sizes, see Chapter 2). In two dimensions the numerics point to a single state at every $T > 0$, in line with Conjecture 4.1.

4.8 Summary and outlook

Coupling covariance turns the free-energy difference between two covariant families of states into a martingale whose increments are differences of mean edge correlations. When those means differ, the variance grows like the volume (Theorem 4.4); in two dimensions, against the surface bound, this forces all metastates to agree on their mean edge correlations for continuous coupling laws with $\mathbb{E}J^4 < \infty$, at positive and (for symmetric such laws) at zero temperature (Theorems 4.9 and 4.10), which restricts atomic metastates. The method sees a metastate only through its edge means, cannot separate a flip pair into its two pure states, and does not reach $d \geq 3$ without a surface-order variance bound for the same quantity (Route warnings 4.17, 4.18 and 4.19). Three

questions remain open: metastates on several pairs in two dimensions (Open problem 4.12), the upper bound of Conjecture 4.14, and Conjecture 4.1 itself, which needs an input that detects long-range order directly.

Open problem 4.20 (Newman [171, Research Problem 4.2])**.** Determine, for all symmetric coupling laws, and in particular for $\pm J$ and Gaussian couplings, whether the two-dimensional EA model has a unique Gibbs state at every $\beta < \infty$, or at least whether there is no $\beta < \infty$ at which the Gibbs state is a.s. non-unique.

Newman's weaker form, for at least some law with $\nu(\mathbb{R} \setminus \{0\}) > p_c(2) = \frac{1}{2}$, is answered by the diluted laws of Chapter 13, slightly above the percolation threshold, by an exact cancellation that is not known to extend to undiluted laws.

The zero-temperature side of the two-dimensional problem, domain walls and half-plane uniqueness, is developed in Chapter 5; Chatterjee's chaos theorems, which complement Theorem 4.13, in Chapter 6; the random-field model, where monotonicity separates the plus and minus states and the programme succeeds, in Chapter 10; and a comparison route to two-dimensional non-order, which targets the decay of two-point functions directly, in Chapter 11.

Part III

Zero temperature

Chapter 5

Ground states and their stability

At zero temperature the Gibbs measure of a finite Edwards–Anderson system sits on its ground states, and the question "how many low-temperature states are there?" becomes "how many ground states are there?". For a continuous coupling law a finite box has, almost surely, a single ground-state pair. The infinite lattice is different. There a ground state is a configuration that no *finite* flip can improve, and nothing obvious forbids many of them. This chapter asks how many ground-state pairs the EA model on $\mathbb{Z}^d$ has, and how stable they are under changes of the couplings and under the natural zero-temperature dynamics.

The question matters because the pictures of Chapter 2 already disagree at $T = 0$. Heuristically, the droplet and scaling picture predicts a single ground-state pair in every dimension for states built from coupling-independent boundary conditions, while replica-symmetry-breaking and chaotic-pairs pictures predict infinitely many mutually incongruent pairs in high enough dimension. A theorem either way in some $d \geq 2$ would be the zero-temperature version of the order question of this review.

The rigorous record is short. In $d = 1$ the answer is one pair (Example 5.4). In $d = 2$, for continuous symmetric laws of unbounded support, two ground-state pairs drawn from periodic metastates differ by at most one domain wall (Theorem 5.9) and on the half-plane the metastate is a single pair (Theorem 5.12); for continuous laws with support $\mathbb{R}$ the full set of half-plane ground states has two or infinitely many elements (Theorem 5.13); for Gaussian couplings the unsatisfied dual edges of a translation-invariant ground state have only finite clusters (Theorem 5.16). In $d \geq 3$ nothing beyond the general facts of Section 5.1 is known. Almost every proof below runs on one engine: *translation invariance*, used through the ergodic theorem and through Burton–Keane counting of ends, together with *local coupling modifications* that keep events of positive probability. By the end of the chapter the reader can prove the basic structure of infinite-volume ground states, the matching description of planar ground states, and fixation of zero-temperature dynamics for continuous laws with finite mean, and can state the two-dimensional theorems exactly.

5.1 Ground states on infinite graphs

Let $G = (V, E)$ be an infinite, connected, locally finite graph, usually $\mathbb{Z}^d$ with nearest-neighbour edges E_d, and let $J = (J_e)_{e \in E}$ be couplings with $J_e \neq 0$. For an edge $e = xy$ write $\sigma_e := \sigma_x \sigma_y$; the edge is *satisfied* in σ when $J_e \sigma_e > 0$. For a finite set $A \subset V$ let ∂A be the set of edges with exactly one endpoint in A, and let σ^A be σ with the spins in A reversed. Only the edges of ∂A

change status, so

$$H_J(\sigma^A) - H_J(\sigma) = 2 \sum_{e \in \partial A} J_e \sigma_e, \tag{5.1}$$

a finite sum even though H_J itself is not defined on an infinite graph.

Definition 5.1 (Ground states)**.** A configuration $\sigma \in \{-1,1\}^V$ is a *ground state* for J if $\sum_{e \in \partial A} J_e \sigma_e \geq 0$ for every finite $A \subset V$, and a *strict* ground state if the inequality is strict for every nonempty finite A. Ground states come in pairs $\pm\sigma$; a *ground-state pair* (GSP) is determined by $(\sigma_e)_{e \in E}$. We write $\mathrm{GS}(J)$ for the set of ground states. In a finite box Λ with a boundary condition (fixed spins outside Λ, free, or periodic), a finite-volume ground state minimises the energy among configurations obeying the boundary condition; equivalently (5.1) is nonnegative for every $A \subseteq \Lambda$.

This is the convention of the notation table. Two easy facts make the definition usable.

Lemma 5.2 (Existence)**.** *Let all $J_e \neq 0$, let $\Lambda_n \uparrow \mathbb{Z}^d$ be boxes with arbitrary boundary conditions, and let $\sigma^{(n)}$ be finite-volume ground states. Every limit point of $(\sigma^{(n)})$ in the product topology is a ground state. In particular* $\mathrm{GS}(J) \neq \emptyset$.

Proof. The space $\{-1,1\}^{\mathbb{Z}^d}$ is compact. Fix a finite A; for n large, A and the endpoints of ∂A lie in Λ_n, so $\sum_{\partial A} J_e \sigma_e^{(n)} \geq 0$. This sum depends on finitely many spins and passes to the limit along the subsequence. □

Lemma 5.3 (No ties for continuous laws)**.** *Let the couplings be iid with an atomless law ν. Then ν-a.s., for every finite nonempty set F of edges and every sign vector $s \in \{\pm 1\}^F$, $\sum_{e \in F} s_e J_e \neq 0$ (in a periodic box F may contain wrap-around edges, whose couplings are also iid with law ν). Consequently, almost surely every ground state is strict, and every finite box with a fixed boundary condition has a unique ground state (a unique pair for free or periodic boundary conditions).*

Proof. A signed sum of independent atomless variables is atomless, and there are countably many pairs (F, s). If two finite-volume ground states, not related by a global flip, differ on A, then (5.1), with ∂A taken in the box graph, gives a vanishing signed sum over the nonempty edge set ∂A. □

Example 5.4 (One dimension)**.** On $\mathbb{Z}$ with iid atomless couplings, $\mathrm{GS}(J) = \{\pm\bar{\sigma}\}$ almost surely, where $\bar{\sigma}_x \bar{\sigma}_{x+1} = \operatorname{sgn} J_{x,x+1}$. Indeed $\bar{\sigma}$ satisfies every edge. If a ground state σ left an edge e unsatisfied, pick an edge f with $|J_f| < |J_e|$ (there are infinitely many on each side, almost surely) and let A be the interval between e and f. Then $\partial A = \{e, f\}$ and the flip cost (5.1) is $2(J_e\sigma_e + J_f\sigma_f) \leq 2(|J_f| - |J_e|) < 0$. The argument uses that every edge of $\mathbb{Z}$ separates two half-lines; on $\mathbb{Z}^2$ every frustrated plaquette (Chapter 1) has an unsatisfied edge in every configuration.

Definition 5.5 (Interfaces and incongruence)**.** For configurations σ, τ the *interface* is $\sigma\Delta\tau := \{e : \sigma_e \neq \tau_e\}$, the edges satisfied in exactly one of them. It is the edge boundary of the disagreement set $\{x : \sigma_x \neq \tau_x\}$ and is empty exactly when $\tau = \pm\sigma$. In $d = 2$ we draw $\sigma\Delta\tau$ on the dual lattice; its connected components are *domain walls*. Two GSPs are *incongruent* if $\sigma\Delta\tau$ has a positive density of edges, and *regionally congruent* if it is nonempty with density zero [181]. The *common-satisfaction graph* is $D_0(\sigma, \tau) := (V, \{e : J_e\sigma_e > 0,\ J_e\tau_e > 0\})$.

The first structural fact says that two strict ground states cannot differ, or agree, on a finite island. Part (a) is the mechanism by which Newman and Stein exclude closed domain walls [181, Proposition 1]; part (b) is the same computation applied to the edges both states satisfy.

Proposition 5.6 (No finite droplets)**.** *Let σ, τ be strict ground states for the same J on an infinite connected locally finite graph. Then*

(a) every connected component of $\{x : \sigma_x = \tau_x\}$ and of $\{x : \sigma_x \neq \tau_x\}$ is infinite;

(b) every connected component of $D_0(\sigma, \tau)$ is infinite.

By Lemma 5.3, for iid atomless couplings both statements hold almost surely for all pairs of ground states at once.

Proof. (a) Let A be a finite component of one of the two sets. Each edge of ∂A joins agreement to disagreement, so $\tau_e = -\sigma_e$ on ∂A, and the two strict inequalities $\sum_{\partial A} J_e\sigma_e > 0$ and $\sum_{\partial A} J_e\tau_e > 0$ contradict each other. (b) Let A be the vertex set of a finite component of D_0. No edge of ∂A is satisfied in both states, so with $a_e = \operatorname{sgn}(J_e\sigma_e)$ and $b_e = \operatorname{sgn}(J_e\tau_e)$ we have $a_e + b_e \leq 0$, and $\sum_{\partial A} J_e\sigma_e + \sum_{\partial A} J_e\tau_e = \sum_{\partial A} |J_e|(a_e + b_e) \leq 0$, which contradicts strictness. In both cases $\partial A \neq \emptyset$ because G is infinite and connected. □

At $\beta = \infty$ the doubly satisfied edges are exactly the blue edges of the two-replica representation of Chapter 7, so (b) says that two strict ground states have only infinite blue clusters; Chapter 8 asks the much harder positive-temperature question. Strictness is essential: for $\pm J$ couplings a finite droplet can have zero flip cost, and flipping it produces a second ground state with a finite island.

Next comes the only local tool of the chapter, a coupling so strong that no ground state can afford to violate it.

Lemma 5.7 (Super-satisfied edges)**.** *Let $e = xy$ satisfy*

$$|J_{xy}| > \min\Big\{\textstyle\sum_{z\sim x,\, z\neq y} |J_{xz}|,\ \sum_{z\sim y,\, z\neq x} |J_{yz}|\Big\}. \tag{5.2}$$

Then $J_e\sigma_e > 0$ in every ground state σ, so e lies in no interface between ground states. Moreover, if σ is a ground state and e is satisfied in σ, then σ remains a ground state when $|J_e|$ is increased with its sign kept.

Proof. If $J_e\sigma_e < 0$, flipping $\{x\}$ (or $\{y\}$, whichever attains the minimum) costs at most $2(-|J_{xy}| + \sum_{z\neq y} |J_{xz}|) < 0$ by (5.1). For the second claim, raising $|J_e|$ adds $|J'_e| - |J_e| \geq 0$ to every flip cost with $e \in \partial A$ and changes no other. □

Edges satisfying (5.2) are called super-satisfied [13, Proposition 2.6] and appear in the proof of Proposition 3 of [181]. The second half of the lemma is how they are used (Fig. 5.1): raising couplings already satisfied in both ground states of interest changes neither state, keeps an event of positive probability positive when the coupling law has unbounded support, and creates barriers that domain walls must avoid. They reappear in Chapter 6 as the reason the edge overlap is not chaotic.

The global tool is translation invariance. A *zero-temperature metastate* is a measurable assignment $J \mapsto \kappa_J$ of a probability measure on GSPs, obtained (by compactness, along subsequences) as a limit of the laws of finite-volume GSPs in boxes Λ_{L_j} with a coupling-independent boundary condition, along a J-independent sequence of sizes (Chapter 3; for $T = 0$

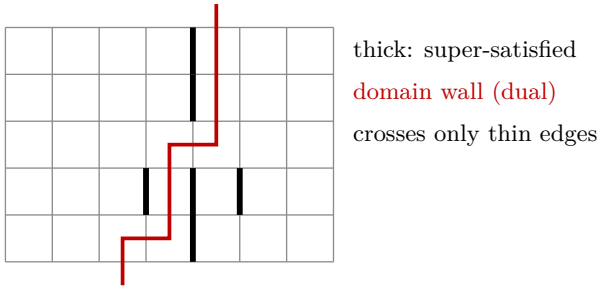


Figure 5.1: A super-satisfied edge (thick) is satisfied in every ground state, so no domain wall can cross it. The column of thick vertical edges at the centre leaves a single opening, and the wall between two ground states (red, on the dual lattice) must pass through it. Raising couplings that are already satisfied in both ground states creates such barriers without changing either state.

see [181, Section 2]). With periodic boundary conditions the law $\nu(dJ)\,\kappa_J(d\alpha)\,\kappa'_J(d\alpha')$ of a pair of GSPs drawn from two such metastates is translation invariant.

Lemma 5.8 (Zero or positive density)**.** *Let (J, α, α') have a translation-invariant law P on $\mathbb{Z}^d$. Then P-a.s. the interface $\alpha\Delta\alpha'$ is either empty or has positive density* $\lim_{L\to\infty} |E_d(\Lambda_L)|^{-1}|\alpha\Delta\alpha' \cap E_d(\Lambda_L)|$.

The density exists by the multidimensional ergodic theorem, and the lemma is the stationarity lemma of Chapter 3 applied to the stationary edge set $\alpha\Delta\alpha'$. So distinct GSPs drawn from translation-covariant metastates are always incongruent [181, Proposition 1]; regionally congruent pairs can arise only from boundary conditions chosen with knowledge of J. The same ergodic argument shows that $|\mathrm{GS}(J)|$ is almost surely a constant [172], [11, Section 2].

5.2 Two dimensions: at most one domain wall

In $d = 2$ the interface between two GSPs is a set of dual edges, and planarity turns the general facts above into strong geometric constraints. By Proposition 5.6(a) a domain wall has no closed loops and no dangling ends. It cannot branch either: branch points would occur with positive density by Lemma 5.8, and a tree-like wall with positive density of branch points would send of order L^2 disjoint branches across the boundary of the $L \times L$ square, which has only $O(L)$ edges. Hence domain walls are doubly infinite self-avoiding dual paths, and together they cut the plane into strips and at most two half-planes [181, Proposition 1]. The theorem below removes the possibility of several walls.

Theorem 5.9 (Newman–Stein [181])**.** *Let $d = 2$ and let the couplings be iid with a symmetric, continuous law of unbounded support. For ν-a.e. J, let α and α' be GSPs chosen independently from periodic-boundary-condition metastates κ_J and κ'_J, built along J-independent sequences of sizes. Then, with probability one, either $\alpha = \alpha'$ or $\alpha\Delta\alpha'$ is a single domain wall with strictly positive density.*

Proof sketch. We omit the excitation metastate, which records how the GSPs change when finitely many couplings vary, and the continuity of the critical values of single couplings. Suppose two adjacent walls W_1, W_2 occur with positive probability. A *rung* is a dual path R from W_1 to W_2 with no other wall vertex; its edges are off the interface, so its energy $E_R = \sum_{e\in R} J_e\sigma_e$ is the

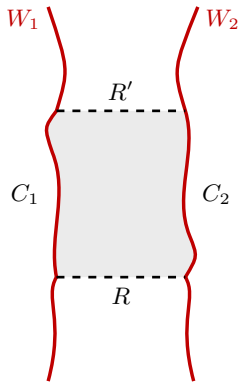


Figure 5.2: Two adjacent domain walls W_1, W_2 (red, dual) and two rungs R, R' (dashed). Flipping the shaded region costs $E_R + E_{R'} + E(C_1, C_2)$ in one ground-state pair and $E_R + E_{R'} - E(C_1, C_2)$ in the other, because the wall edges are satisfied in exactly one of them. Cheap rungs therefore make both pairs nearly unstable, and the proof of Theorem 5.9 shows that rung energies can neither accumulate at zero nor stay bounded away from it.

same in α and α'. Two rungs R, R' and the wall segments C_1, C_2 between them bound a finite region (Fig. 5.2). Its flip costs in α and α' are $E_R + E_{R'} \pm E(C_1, C_2)$, where $E(C_1, C_2)$ is the energy of the wall segments in α; the sign flips because the wall edges change status. Both costs are positive, so $E_R + E_{R'} > 0$. A property of rungs that occurs at all occurs on both sides of the origin by translation invariance, so no rung has negative energy, and zero energy has probability zero. Let $E' = \inf_R E_R$. If $E' = 0$, translation invariance supplies arbitrarily cheap short rungs on both sides of the origin, so every wall bond can be moved across its critical value by an arbitrarily small change, which contradicts the almost sure positivity of these "flexibilities" [181, Propositions 2 and 4]. If $E' > 0$, one super-satisfies the bonds touching W_1 near a nearly optimal rung (Lemma 5.7; this is where unbounded support enters) and then lowers one rung coupling slightly. Neither α nor α' changes, but a rung below the infimum appears, an event of positive probability that contradicts translation invariance [181, Proposition 3 and Lemma 1]. □

The authors explain how the restriction to periodic boundary conditions can be relaxed to other coupling-independent ones by averaging over translates of the box [181, Section 2.3], and "suspect that such domain walls do not in fact occur" [181, p. 4]. Numerically the case for a single pair in $d = 2$ is strong.

Physics prediction 5.10 (Two-dimensional domain walls)**.** For Gaussian couplings on $\mathbb{Z}^2$ the domain-wall energy decays with scale, with stiffness exponent $\theta < 0$ and walls of fractal dimension $d_f \approx 1.27$; the best estimates of [139] are $\theta = -0.2793(3)$ and $d_f = 1.27319(9)$. A negative θ is read as $T_c = 0$ and as a single ground-state pair in coupling-independent constructions.

Remark 5.11 (The two-dimensional periodic metastate). Fluctuation arguments also constrain how the metastate itself distributes its weight over ground-state pairs [16, Section 6], [189]; Chapter 4 gives the account. Coupling-dependent boundary conditions can produce other limits, as in Loebl's construction of regionally incongruent ground states [149].

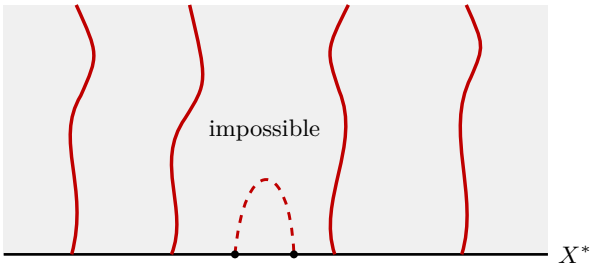


Figure 5.3: Tethered domain walls on the half-plane. A dual path inside the interface with both ends on the dual axis X^* (dashed) would need energies of opposite signs in the two ground states, both positive. So each wall meets X^* once, and a positive density of tether points gives infinitely many walls, which contradict the one-wall theorem.

5.3 The half-plane

On the half-plane $H = \mathbb{Z} \times \mathbb{N}$ the free bottom boundary gives domain walls something to hold on to, and the one-wall theorem then finishes the argument. Let $\Lambda_n = [-n, n] \times [0, 2n]$ with periodic boundary conditions on the left and right sides and free ones on the top and bottom, and let K_n be the joint law of the couplings and the GSP of Λ_n. Let X^* be the dual x-axis, the row of dual vertices just below H. A ground state of H satisfies the flip condition for finite subsets of H; equivalently, the energy $\sum J_e \alpha_e$ is positive along every dual circuit and along every dual path that starts and ends at distinct points of X^* [13, equations (3)–(4)].

Theorem 5.12 (Arguin–Damron–Newman–Stein [13])**.** *Let the couplings be iid with a symmetric, continuous law of unbounded support. The measures K_n converge as $n \to \infty$, without passing to subsequences, to a measure K supported on infinite-volume GSPs of H, and for ν-a.e. J the conditional law $K(\cdot \mid J)$ is supported on a single GSP. Equivalently, for ν-a.e. J there is a unique metastate for these boundary conditions, and it is supported on a single GSP.*

Proof sketch. We omit the excitation-metastate formalism and the coupling modification that pushes an interface to the boundary [13, Corollary 2.9]. Draw α, α' independently from two metastates and suppose $\alpha \Delta \alpha' \neq \emptyset$ with positive probability. A domain wall that meets X^* is *tethered. Each tethered wall meets X^* only once*: a dual path inside $\alpha \Delta \alpha'$ joining two points of X^* would have energies of opposite signs in α and α', while both must be positive. Super-satisfying edges along a path (Lemma 5.7) shows that the interface reaches X^* with positive probability, and horizontal translation invariance with the ergodic theorem then gives a positive density of tether points, each on its own wall. Hence there are infinitely many tethered walls, and the expected number crossing $[-n, n]$ at height k is at least cn uniformly in k [13, Propositions 3.3–3.4]. Shifting the half-plane down by k and averaging over k produces a translation-invariant law on pairs of full-plane GSPs whose interface has at least two walls with positive probability. The proof of Theorem 5.9 uses only translation invariance and the excitation structure, and so it applies to this limit [13, Theorem 2.10]. This is a contradiction. □

Theorem 5.12 concerns the GSP seen by one family of boundary conditions. The full set of ground states is a different object.

Theorem 5.13 (Arguin–Damron [11]). *Let the couplings be iid with a continuous law whose support is $\mathbb{R}$. For the EA model on the half-plane $\mathbb{Z}\times\mathbb{N}$, the number of ground states $|\mathrm{GS}(J)|$ is either 2 with ν-probability one or ∞ with ν-probability one.*

A metastate may miss ground states that no coupling-independent boundary condition selects [11, p. 4]. The device of Arguin and Damron is the uniform measure on $\mathrm{GS}(J)$ when this set is finite, which is measurable and covariant, and to which the interface method can be applied. The dichotomy is the spin-glass analogue of Wehr's theorem that an Ising ferromagnet with random couplings has, in every $d\geq 2$, either two or infinitely many ground states [244].

5.4 Frustration, matching and unsatisfied edges in the plane

The unsatisfied edges of a planar ground state have a combinatorial description that makes planar ground states computable and that underlies the last two-dimensional theorem.

Lemma 5.14 (Parity of unsatisfied edges). *For any configuration σ on a 2-connected plane graph (so that every face P is bounded by a cycle ∂P) and any face P, the number of unsatisfied edges of P is odd if and only if $\prod_{e\in\partial P}\operatorname{sgn} J_e=-1$, that is, if P is frustrated.*

Proof. $\prod_{e\in\partial P}\operatorname{sgn}(J_e\sigma_e)=\prod_{e\in\partial P}\operatorname{sgn} J_e$, since every vertex of the cycle ∂P appears in two of its edges. □

Thus the set $U^*(\sigma)$ of unsatisfied dual edges has odd degree exactly at the frustrated faces: it pairs them up by dual paths (Fig. 5.4). Writing $U(\sigma)$ for the set of unsatisfied primal edges,

$$H_J(\sigma)=-\sum_e |J_e|+2\sum_{e\in U(\sigma)}|J_e|. \tag{5.3}$$

Proposition 5.15 (Planar ground states by matching). *Let G be a finite 2-connected plane graph (for example a box of $\mathbb{Z}^2$) with free boundary conditions and nonzero couplings. The map $\pm\sigma\mapsto U(\sigma)$ is a bijection from configurations modulo a global flip onto the edge sets U such that every bounded face P has $|U\cap\partial P|$ odd exactly when P is frustrated. Ground-state pairs correspond to such sets of minimal weight $\sum_{e\in U}|J_e|$.*

Proof. Injectivity: $U(\sigma)$ determines $\sigma_e=\operatorname{sgn} J_e\,(-1)^{\mathbf{1}\{e\in U\}}$ for every edge. Surjectivity: given U with the parity property, the signs $s_e:=\operatorname{sgn} J_e\,(-1)^{\mathbf{1}\{e\in U\}}$ have product $+1$ around every bounded face. Boundaries of bounded faces generate the cycle space of a planar graph over $\mathbb{Z}/2$, so the product of s_e around every cycle is $+1$, and since G is connected there is σ, unique up to a global flip, with $\sigma_e=s_e$. Then $U(\sigma)=U$. The last claim is (5.3). □

Finding such a set is a minimum-weight T-join problem in the dual; it reduces to minimum-weight perfect matching of the frustrated faces, with dual shortest-path distances, and is solvable in polynomial time. The method goes back to Bieche, Uhry, Maynard and Rammal [28]. On general graphs the ground-state problem is a weighted maximum cut, which is NP-hard [21], [49, p. 3]. Planar sampling at positive temperature is discussed in Chapter 14.

In infinite volume the same lemma shows that the unsatisfied dual edges of a ground state form a forest: a dual cycle of unsatisfied edges encloses a finite set whose flip satisfies all of them and lowers the energy. Berger and Tessler proved that this forest has only finite trees, and that the statement survives at low positive temperature. The objects are Gaussian couplings and an

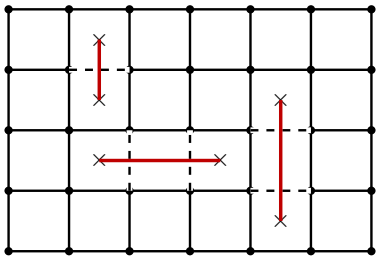


Figure 5.4: Frustration and matching. Dashed edges carry negative couplings, and crosses mark the frustrated plaquettes. In the all-plus configuration the unsatisfied edges are the dashed ones, and their dual edges (red) pair the frustrated plaquettes by dual paths, as Lemma 5.14 requires. A ground state is a configuration whose pairing has the least total $|J|$, which is a matching problem.

EA_β spin glass distribution: a joint law of (J, σ) on $\mathbb{Z}^2$ whose J-marginal is iid standard Gaussian and whose conditional law of σ in every finite region, given the outside, is the finite-volume Gibbs measure at inverse temperature β. At $\beta = \infty$ this means that σ is almost surely a ground state [26, Definitions 1.1 and 1.3]. Translation-invariant examples come from periodic boxes by compactness.

Theorem 5.16 (Berger–Tessler [26])**.** *There is a finite β^* such that for every $\beta \in (\beta^*, \infty]$ and every translation-invariant EA_β spin glass distribution on $\mathbb{Z}^2$ with Gaussian couplings, almost surely every connected component of the set of unsatisfied dual edges is finite. At $\beta = \infty$ the unsatisfied dual edges of a translation-invariant ground state form almost surely a forest all of whose components are finite [26, Theorem 1.7].*

Since every frustrated plaquette carries an odd number of them, the unsatisfied dual edges have positive density [26, Remark 1.6]; the theorem says that they nevertheless do not percolate.

Proof sketch. We omit the Peierls-type estimate that bounds unsatisfied cycles at large finite β, the measurable extraction of a spanning forest, and the large-deviation input. At $\beta = \infty$ there are no unsatisfied dual cycles; at large β there are only finitely many through each vertex, and a translation-invariant spanning forest with the same components can be extracted [26, Lemmas 2.8–2.9]. By the Burton–Keane counting of encounter points [38], each infinite tree of a translation-invariant forest has one or two ends [26, Lemma 2.14]. In the two-ended case the bi-infinite paths have positive density and cut a box $[N]^2$ into regions that can be two-coloured. Flipping the regions of one colour satisfies all path edges, a gain of order N^2 (a large-deviation bound keeps the total $|J|$ along long connected sets proportional to their length), against a boundary cost of order N. In the one-ended case the mass-transport principle gives an expected $N \log N$ forest edges on paths that cross the box [26, Lemma 2.15], and the same comparison with the boundary applies. Either way some finite flip lowers the energy, a contradiction. □

The authors ask whether the ground state is unique up to a global flip, and whether the theorem has an analogue on $\mathbb{Z}^d$, $d > 2$ [26, Section 3].

5.5 Zero-temperature dynamics and fixation

A different way to look for ground states is to let the system relax. In the zero-temperature Glauber dynamics on $\{-1,1\}^{\mathbb{Z}^d}$ [103, 180], the initial configuration σ^0 is iid symmetric, each site carries an independent rate-one Poisson clock, and when the clock at x rings the spin σ_x flips with probability 1, 1/2 or 0 according as

$$\Delta H_x(\sigma) := 2\sum_{y\sim x} J_{xy}\sigma_x\sigma_y \tag{5.4}$$

is negative, zero or positive. The process exists by the usual graphical construction. A site is an *F-site* if it flips finitely often and an *I-site* otherwise. Their densities ρ_F, ρ_I are almost sure constants by translation ergodicity, and the pair (d,ν) is of *type F*, *type I* or *type M* according as $\rho_F = 1$, $\rho_I = 1$ or $0 < \rho_F < 1$ [103]. Type F means that the configuration converges, to a state stable under single flips that need not be a ground state.

The basic theorem is an energy argument, stated as in [103, Theorem 2].

Theorem 5.17 (Nanda–Newman–Stein [168])**.** *For any d and any coupling law ν with finite mean, $\mathbb{E}|J_e| < \infty$, almost surely there are only finitely many flips with $\Delta H_x(\sigma) < 0$ at each site x.*

Proof. We follow the energy argument sketched in [180, Section 4], with the bookkeeping made explicit. Let $\varepsilon_x(t) := -\frac{1}{2}\sum_{y\sim x} J_{xy}\sigma_x(t)\sigma_y(t)$ be the energy at x at time t. A flip at x changes ε_x by $\Delta H_x/2$ and ε_y, for each neighbour y, by $J_{xy}\sigma_x\sigma_y$ (values before the flip). By translation invariance of the joint law of couplings, initial spins and clocks, the expected change of ε_0 caused by flips at the neighbours equals the expected value of $\sum_{y\sim 0} J_{0y}\sigma_0\sigma_y = \Delta H_0/2$ summed over flips at the origin (Section A.10). Hence

$$\mathbb{E}\varepsilon_0(t) - \mathbb{E}\varepsilon_0(0) = \mathbb{E}\sum_{\text{flips at } 0 \text{ in } (0,t]} \Delta H_0. \tag{5.5}$$

Every flip has $\Delta H_0 \le 0$, $\mathbb{E}\varepsilon_0(0) = 0$ because σ^0 is symmetric and independent of J, and $\varepsilon_0(t) \ge -\frac{1}{2}\sum_{y\sim 0}|J_{0y}|$. So $\mathbb{E}\sum|\Delta H_0| \le d\,\mathbb{E}|J_e|$ uniformly in t, and by monotone convergence the sum of $|\Delta H_0|$ over all flips at 0 is almost surely finite. Every strictly energy-lowering flip at 0 has $|\Delta H_0| \ge m_0(J)$, the smallest nonzero value of $2|\sum_{y\sim 0} s_y J_{0y}|$ over $s \in \{\pm 1\}^{2d}$, and $m_0(J) > 0$ because it is a minimum of finitely many positive numbers (if there are none, no such flip can occur). So there are finitely many such flips at 0, and by translation invariance at every site. □

Corollary 5.18 (Continuous couplings fixate)**.** *If ν is continuous with finite mean, then (d,ν) is of type F for every d.*

Proof. For an atomless law every sign pattern gives $\sum_y s_y J_{0y} \neq 0$ almost surely, so $\Delta H_x \neq 0$ at all times and every flip is strictly energy-lowering. □

When ν has atoms, zero-energy flips are possible and happen at rate 1/2 forever unless something blocks them. For the biased $\pm J$ law $\nu_p = p\delta_{+1} + (1-p)\delta_{-1}$ of the notation table, the known results are these.

Theorem 5.19 (Gandolfi–Newman–Stein [103])**.** *(a) For $d = 2$ and every $0 < p < 1$, the law ν_p is of type M.*

(b) For every $d \geq 2$, a law ν that is neither continuous nor of the form $\alpha\delta_J + (1-\alpha)\delta_{-J}$ with $J \geq 0$ and $0 \leq \alpha \leq 1$ is of type M.

The proof of (b) is local [103, Section II]: cubes with strong couplings of one sign inside, weak couplings to the outside and a matching initial pattern never flip ($\rho_F > 0$), and an atom of ν lets one surround a site by frozen cubes so that it always sees $\Delta H = 0$ and flips forever ($\rho_I > 0$). The proof of (a) uses the contour representation of planar $\pm J$ models and shows only that some site in a finite region is of each type.

Route warning 5.20 (The energy argument says nothing about zero-energy flips). Theorem 5.17 controls only flips that lower the energy, so for $\pm J$ couplings it cannot decide the type. In $d = 2$ the planar contour representation does decide it; in $d \geq 3$ no current method controls zero-energy flips [180, Section 6.2].

Open problem 5.21 (Newman–Stein [180])**.** Determine the type (F, I or M) of $\pm J$ models on $\mathbb{Z}^d$ for $d \geq 3$, and of continuous laws with infinite mean for $d \geq 2$ (problems 3 and 2 of [180, Section 6.2]; for the second class the authors conjecture type F).

5.6 Where the two-dimensional methods stop

Route warning 5.22 (Domain-wall arguments need one-dimensional interfaces). The proofs of Theorems 5.9, 5.12 and 5.16 use that interfaces are dual *paths*: they can be adjacent, be joined by short rungs, be tethered to a line, and be counted by their ends with Burton–Keane and mass-transport arguments. In $d \geq 3$ an interface is a $(d-1)$-dimensional surface, and none of these notions applies. (In the disordered monomer–dimer model the symmetric difference of two ground states is a union of paths in every dimension, and uniqueness is proved in every d [144].) For the EA model no substitute for these path arguments is known in $d \geq 3$; this is our assessment, and there is no counterexample.

Route warning 5.23 (Metastate uniqueness is not uniqueness of ground states). A complete two-dimensional result for metastates would concern only GSPs selected by coupling-independent boundary conditions: on the half-plane the metastate is a single pair, yet Theorem 5.13 leaves open whether $|\mathrm{GS}(J)|$ is 2 or ∞.

To summarise: ground states exist, continuous laws exclude finite islands, super-satisfied edges are frozen, and covariant constructions give only incongruent distinct pairs. In two dimensions interfaces are dual paths: at most one wall between periodic metastate pairs, a single pair on the half-plane, two or infinitely many ground states there, and, for Gaussian couplings, finite unsatisfied clusters. Continuous laws with finite mean fixate in every dimension. Chapter 6 asks how far the ground state moves when the couplings move, and the common-satisfaction clusters of Proposition 5.6 return as blue clusters in Chapter 8.

Open problem 5.24. (a) Show that two GSPs drawn from periodic metastates on $\mathbb{Z}^2$ coincide (Conjecture 1 of [181]), or equivalently exclude the single positive-density wall of Theorem 5.9.

(b) Determine $|\mathrm{GS}(J)|$ on the half-plane and on $\mathbb{Z}^2$ for continuous couplings.

(c) Find any $d \geq 3$ and any continuous symmetric law for which the number of GSPs in a coupling-independent metastate is determined. Chatterjee states the question for $|\mathrm{GS}(J)|$: "The main question is whether this number is greater than two in some dimension and for some symmetric and continuous distribution of the disorder" [49, p. 11].

Chapter 6

Chaos and superconcentration

How much does the ground state move when the couplings move a little? Fisher and Huse [86] and Bray and Moore [36] predicted that in a spin glass it moves a lot: an arbitrarily small change of the disorder reorganises the state on large scales. This *disorder chaos* is a signature of glassiness in the physics literature, and since 2009 it has become one of the few places where short-range spin glasses have theorems.

The theorems split along the two overlaps of Chapter 1. The *edge* overlap is not chaotic at any temperature: it stays of order one under small perturbations (Chatterjee [46], Arguin and Hanson [12]). The *site* overlap of zero-temperature ground states is chaotic: on a box of side L in any dimension it is lost once the perturbation is much larger than $1/L$ [49]. The second result rests on the observation at the centre of this chapter: as a function of the couplings, the product $\sigma_i\sigma_j$ of ground-state spins has no Hermite components of low degree, because a gauge flip along any cut that the index set does not touch reverses its sign. Ornstein–Uhlenbeck smoothing damps high-degree components, so a small perturbation destroys the correlation of distant spins. With a few variational inequalities the same identity gives Chatterjee's results on large and rough droplets, low-energy macroscopic excitations and large critical droplets; a two-step gauge flip gives his slow decay of boundary influence.

One caution shapes the reading of the whole chapter: every chaos statement for the site overlap, and most of the others, already holds on the one-dimensional chain, whose ground state is unique and unfrustrated (Example 6.2 and Route warning 6.24). By the end of the chapter the reader can prove the variance identity behind "superconcentration $\iff$ chaos", the absence of edge chaos, Chatterjee's site-chaos theorems, the existence of macroscopic low-energy excitations, and the $1/L$ lower bound on boundary sensitivity.

6.1 Perturbing the couplings

Throughout, $G = (V, E)$ is a finite connected graph with a (possibly empty) set $B \subseteq V$ of boundary vertices, the interior $V^\circ := V \setminus B$ is connected and nonempty, and $|E| \geq 2$. A boundary condition $\eta \in \{\pm 1\}^B$ is fixed in advance and does not depend on J. The ground state $\sigma = \sigma^J$ minimises H_J among configurations equal to η on B; for an atomless coupling law it is almost surely unique, and for $B = \emptyset$ (free boundary condition) it is a unique pair. Let $d(\cdot, \cdot)$ be the graph distance, $d(i, B) := \min_{k \in B} d(i, k)$ (infinite if $B = \emptyset$), and $D(i, j) := \min\{d(i, j),\ d(i, B) + d(j, B)\}$.

Definition 6.1 (Two perturbations)**.** Let J' be an independent copy of J, let $p \in (0, 1)$, and write $p = 1 - e^{-t}$ with $t > 0$. For iid standard Gaussian couplings, the *Ornstein–Uhlenbeck*

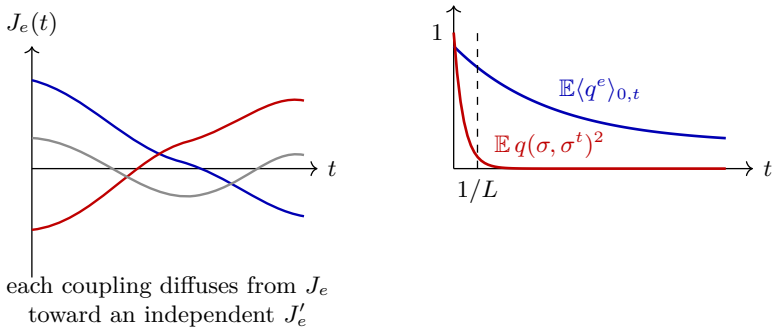


Figure 6.1: The Ornstein–Uhlenbeck interpolation between two disorder samples (left) and the two overlaps along it (right, schematic). The edge overlap stays of order one for $t = O(1)$, at a rate controlled by the free-energy variance (Corollary 6.6); the site overlap of ground states is lost once $t \gg 1/L$ (Theorem 6.14).

(OU) perturbation is $J(t) := e^{-t}J + \sqrt{1-e^{-2t}}\,J' = (1-p)J + \sqrt{2p-p^2}\,J'$. The *resampling perturbation*, available for any law, replaces each J_e by J'_e independently with probability p. In both cases $J(t)$ has the law of J, and we write σ^t for the ground state of $J(t)$.

On G the site and edge overlaps of the notation table are $q(\sigma,\tau) := |V^\circ|^{-1}\sum_{i\in V^\circ}\sigma_i\tau_i$ and $q^e(\sigma,\tau) := |E|^{-1}\sum_{e\in E}\sigma_e\tau_e$, with $\sigma_e := \sigma_x\sigma_y$ for $e = xy$; for free boundary conditions only q^2 is meaningful. *Site chaos* means that $\mathbb{E}\, q(\sigma,\sigma^t)^2 \to 0$ for perturbations that vanish with the system size; *edge chaos* means the same for q^e.

Example 6.2 (The chain)**.** Let G be the path $\{0,\dots,L\}$ with free boundary conditions. Then $\sigma_i\sigma_j = \prod_{e\in[i,j]}\operatorname{sgn} J_e$, the product over the edges between i and j. Under resampling, $\mathbb{E}[\operatorname{sgn} J_e\,\operatorname{sgn} J_e(t)] = 1-p$, so

$$\mathbb{E}[\sigma_i\sigma_j\sigma_i^t\sigma_j^t] = (1-p)^{|i-j|}, \qquad \mathbb{E}\, q(\sigma,\sigma^t)^2 = \frac{1}{(L+1)^2}\sum_{i,j}(1-p)^{|i-j|} \le \frac{2}{(L+1)p},$$

while $\mathbb{E}\, q^e(\sigma,\sigma^t) = 1-p$. (Under the OU perturbation $1-p$ is replaced by $\frac{2}{\pi}\arcsin e^{-t}$.) The site overlap is lost once $p \gg 1/L$ and the edge overlap is not, although the chain has a unique, unfrustrated ground-state pair. Keep this example in mind: the general theorems below show the same two behaviours in every dimension.

6.2 The Hermite basis and the Ornstein–Uhlenbeck semigroup

The tool for everything that follows is harmonic analysis on Gaussian space. Let $h_0 = 1, h_1(x) = x, h_2(x) = (x^2-1)/\sqrt{2},\dots$ be the normalised Hermite polynomials, an orthonormal basis of L^2 of the standard Gaussian law, with $h_k(-x) = (-1)^k h_k(x)$. For a multi-index $n \in \mathbb{N}^E$ put $h_n(J) := \prod_{e\in E} h_{n_e}(J_e)$, $|n| := \sum_e n_e$, $E(n) := \{e : n_e > 0\}$ and $\delta(n) := |E(n)|$, and let $G(n)$ be the graph formed by the edges $E(n)$ and their endpoints $V(n)$. Every $f \in L^2(J)$ expands as $f = \sum_n \hat f(n) h_n$ with $\hat f(n) = \mathbb{E}[f h_n]$, and $\mathbb{E} f^2 = \sum_n \hat f(n)^2$. The basis diagonalises both perturbations [49, Section 2.1]:

$$\mathbb{E}[h_n(J(t)) \mid J] = e^{-|n|t}h_n(J) \ \ \text{(OU)}, \qquad \mathbb{E}[h_n(J(t)) \mid J] = e^{-\delta(n)t}h_n(J) \ \ \text{(resampling)}. \tag{6.1}$$

The first is the eigenfunction property of the OU semigroup [48, Chapters 2 and 6]; the second holds because a resampled coordinate kills every term with $n_e \geq 1$ (Section A.11). Hence, for $f, g \in L^2$,

$$\mathbb{E}[f(J)\, g(J(t))] = \sum_n e^{-|n|t} \hat{f}(n)\hat{g}(n), \qquad \mathbb{E}\big[\mathbb{E}[f(J(t)) \mid J]^2\big] = \sum_n e^{-2|n|t} \hat{f}(n)^2, \tag{6.2}$$

with $\delta(n)$ in place of $|n|$ for resampling. So a perturbation of size t keeps the components of degree $\ll 1/t$ and damps those of degree $\gg 1/t$. Chaos questions become questions about where the Hermite mass of an observable sits.

6.3 Superconcentration and the edge overlap

Let G be a finite graph with free boundary conditions and iid standard Gaussian couplings, let $\beta \in (0, \infty]$, and let

$$F_\beta(J) := -\frac{1}{\beta} \log \sum_\sigma e^{-\beta H_J(\sigma)} \quad (\beta < \infty), \qquad F_\infty(J) := \min_\sigma H_J(\sigma).$$

Write $\langle\cdot\rangle_{0,t}$ for the product of the Gibbs measures at β for J and for $J(t)$ (OU), each replica drawn independently given (J, J'); at $\beta = \infty$ the Gibbs measure is the uniform law on the ground states. The free energy and the edge overlap are tied by an exact identity.

Theorem 6.3 (Chatterjee [46])**.** *For every finite graph and every $\beta \in (0, \infty]$,*

$$\operatorname{Var} F_\beta = |E| \int_0^\infty e^{-t}\, \psi(t)\, dt, \qquad \psi(t) := \mathbb{E}\langle q^e(\sigma, \tau)\rangle_{0,t}. \tag{6.3}$$

Moreover $\psi(t) = \sum_{k \geq 0} \mu_k e^{-kt}$ with $\mu_k \geq 0$ and $\sum_k \mu_k = \psi(0) \leq 1$; in particular ψ is nonnegative, nonincreasing and completely monotone.

This is [46, Theorem 3.8] for the EA covariance $\operatorname{Cov}(-H_J(\sigma), -H_J(\tau)) = |E|\, q^e(\sigma, \tau)$, with complete monotonicity from [46, Lemma 3.3]. Chatterjee proves it by Gaussian integration by parts; we give the equivalent computation in the Hermite basis.

Proof. For f smooth with bounded derivatives, or Lipschitz, Gaussian integration by parts gives $\widehat{\partial_e f}(m) = \sqrt{m_e + 1}\, \hat{f}(m + \mathbf{1}_e)$ (Section A.11). By (6.2),

$$\sum_e \mathbb{E}[\partial_e f(J)\, \partial_e f(J(t))] = \sum_e \sum_m e^{-|m|t} (m_e + 1) \hat{f}(m + \mathbf{1}_e)^2 = e^t \sum_n |n|\, e^{-|n|t} \hat{f}(n)^2.$$

Multiplying by e^{-t} and integrating over $t \in (0, \infty)$ gives $\sum_{n \neq 0} \hat{f}(n)^2 = \operatorname{Var} f$ (all terms are nonnegative). For $f = F_\beta$ with $\beta < \infty$, $\partial_e F_\beta = -\langle \sigma_e \rangle_J$, so $\sum_e \partial_e F_\beta(J)\, \partial_e F_\beta(J(t)) = |E|\, \langle q^e \rangle_{0,t}$, which is (6.3). At $\beta = \infty$, F_∞ is Lipschitz with $\partial_e F_\infty = -\sigma_e$ almost everywhere. Finally, $\psi(t) = |E|^{-1} \sum_e \sum_m e^{-|m|t} \widehat{\partial_e F}(m)^2$, which gives the series with $\mu_k = |E|^{-1} \sum_e \sum_{|m|=k} \widehat{\partial_e F}(m)^2$. □

Corollary 6.4 (Superconcentration $\iff$ edge chaos)**.** *Let G_N be finite graphs with $|E_N| \to \infty$, $\beta \in (0, \infty]$ fixed, and ψ_N as above. Then $\operatorname{Var} F_\beta^{(N)} = o(|E_N|)$ if and only if $\psi_N(t) \to 0$ for every $t > 0$.*

Proof. Since ψ_N is nonincreasing, $|E_N|^{-1}\operatorname{Var} F_\beta^{(N)} \geq \int_0^s e^{-t}\psi_N(t)\,dt \geq (1-e^{-s})\psi_N(s)$. Conversely, if $\psi_N \to 0$ pointwise then the integral tends to zero by dominated convergence, as $0 \leq \psi_N \leq 1$. □

A variance much smaller than the number of couplings is called *superconcentration*; it happens exactly when the edge overlap is chaotic. For the EA model on bounded-degree graphs the variance is of volume order.

Theorem 6.5 (Chatterjee [46, Theorem 1.8])**.** *Let G be a finite graph of maximum degree Δ with iid standard Gaussian couplings. For every $\beta \in (0,\infty]$, $\operatorname{Var} F_\beta \geq \frac{9|E|}{32}\min\{\beta^2, 1/(4\Delta^2)\}$.*

This generalises the square-lattice bound of Wehr and Aizenman [245] (Chapter 4); with Theorem 6.3 it excludes edge chaos.

Corollary 6.6 (No edge chaos; Chatterjee [46])**.** *In the setting of Theorem 6.5, let $v := \frac{9}{32}\min\{\beta^2, 1/(4\Delta^2)\}$. Then $\mathbb{E}\langle q^e(\sigma,\tau)\rangle_{0,t} \geq \frac{v}{2}\,e^{-2t/v}$ for all $t \geq 0$.*

Proof. This is [46, Theorem 1.6] with explicit constants, by a discrete form of [46, Lemma 3.9]. Let $w := \operatorname{Var} F_\beta/|E|$, so $w \geq v$ and, by (6.3), $w = \sum_k \mu_k/(k+1) \leq \sum_k \mu_k \leq 1$. Let $K := \lfloor 2/w \rfloor - 1 \geq 1$. Then $\sum_{k>K}\mu_k/(k+1) \leq 1/(K+2) < w/2$, so $\sum_{k\leq K}\mu_k \geq \sum_{k\leq K}\mu_k/(k+1) > w/2$, and $\psi(t) \geq \sum_{k\leq K}\mu_k e^{-kt} \geq \frac{w}{2}e^{-Kt} \geq \frac{w}{2}e^{-2t/w}$. The last expression increases with w. □

The variance identity sees only the edge overlap, because the covariance of the Gaussian field $-H_J$ is $|E|\,q^e$; the site overlap "can well be chaotic" [46, p. 8]. For general laws a direct argument shows why the edge overlap is stable: a positive density of edges is frozen by their neighbourhood.

Theorem 6.7 (Arguin–Hanson [12])**.** *Let $\Lambda \subset \mathbb{Z}^d$ be a box with edge set Λ^*, and let H_J include the edges from Λ to its exterior vertex boundary, with any boundary condition there. Let the couplings be iid with a continuous law, let $\varepsilon, \varepsilon'$ be independent iid fields of continuous mean-zero variables, independent of J, and let $\sigma(\cdot)$ denote the ground state. For every $\delta > 0$ there is $t_0(\delta) > 0$, independent of Λ, such that for all $t < t_0(\delta)$*

$$\mathbb{E}\Big[\frac{1}{|\Lambda^*|}\sum_{e\in\Lambda^*}\sigma_e(J+t\varepsilon)\,\sigma_e(J+t\varepsilon')\Big] > (1-\delta)\,\mathbb{P}(A),$$

where A is an explicit event with $\mathbb{P}(A) > 0$ independent of Λ.

Proof sketch for a symmetric law with 0 in its support. We omit the continuity estimate; the proof is in Section A.12, following [12, Sections 1.2 and 2.1]. On the open event $A_e = \{J_{xy} > \sum_{z\sim x, z\neq y}|J_{xz}|\}$, which has the same positive probability $\mathbb{P}(A)$ for every edge $e = xy$, the edge is super-satisfied (Lemma 5.7) and stays so after small perturbations with high probability. Off A_e the two perturbed values of σ_e are conditionally independent given J, so their product has nonnegative expectation. □

Without the support assumption Arguin and Hanson freeze a block whose couplings lie away from 0 [12, Section 2.2]. Both proofs give stability on a set of edges of small positive density; whether the edge overlap stays close to one is open (Open problem 6.26).

Remark 6.8 (Edge chaos and incongruence). For standard Gaussian couplings, the conditional corollary of Arguin, Newman and Stein in Chapter 4 reads: assume their conjectured surface-order bound on the variance of ground-state energy differences; then absence of disorder chaos in the edge overlap at some scale $\alpha < 1/d$, in the sense defined there, excludes incongruent ground states in zero-temperature metastates built with coupling-independent boundary conditions [17].

6.4 Site chaos: the gauge identity

We return to the setting of Section 6.1. To cover general symmetric laws and positive temperature at once, we use the representation of Chen, Kim and Sen [56]: the couplings are $\rho(J_e)$ with J_e iid standard Gaussian and ρ an odd function, so that every symmetric law arises (Chatterjee's setting is $\rho(x) = x$). For $\beta \in (0, \infty]$ and $i, j \in V^\circ$ let $\phi_{ij}(J) := \langle \sigma_i \sigma_j \rangle^{\eta}_{\beta, \rho(J)}$ be the Gibbs average with boundary condition η, where at $\beta = \infty$ the average is uniform over the ground states (for an atomless law, simply $\sigma_i \sigma_j$ for the ground state). For $a \in \{\pm 1\}^V$ write $(a \circ J)_{xy} := a_x a_y J_{xy}$; this is the gauge transformation of Chapter 1.

Lemma 6.9 (Gauge identity in the Hermite basis)**.** *Let $a \in \{\pm 1\}^V$ with $a \equiv 1$ on B. Then $\phi_{ij}(a \circ J) = a_i a_j \, \phi_{ij}(J)$, and for every $n \in \mathbb{N}^E$*

$$\hat{\phi}_{ij}(n) = a_i a_j \prod_{e \in E} a_e^{n_e} \; \hat{\phi}_{ij}(n), \qquad a_e := a_x a_y \ (e = xy).$$

Proof. Since ρ is odd, $\rho((a \circ J)_e) = a_e \rho(J_e)$, so $H_{\rho(a \circ J)}(a\sigma) = H_{\rho(J)}(\sigma)$, where $(a\sigma)_x = a_x \sigma_x$. The map $\sigma \mapsto a\sigma$ preserves the boundary condition because $a \equiv 1$ on B, so it carries the Gibbs measure (or the set of ground states) for J onto that for $a \circ J$, and $\phi_{ij}(a \circ J) = a_i a_j \phi_{ij}(J)$. The law of J is invariant under $J \mapsto a \circ J$ and $h_n(a \circ J) = \prod_e a_e^{n_e} h_n(J)$ by the parity of h_k. Therefore $\hat{\phi}_{ij}(n) = \mathbb{E}[\phi_{ij}(a \circ J) h_n(a \circ J)] = a_i a_j \prod_e a_e^{n_e} \hat{\phi}_{ij}(n)$. □

For $\beta = \infty$, $\rho(x) = x$ and $a = 1 - 2\mathbf{1}_A$ this is Chatterjee's Lemma 1 [49, Section 2.1]; for free boundary conditions and all β it is [56, Lemma 2.2]. The lemma kills every coefficient for which some admissible a has $a_i a_j \prod_e a_e^{n_e} = -1$.

Corollary 6.10 (Chatterjee [49])**.** *If $i \neq j$ and $\hat{\phi}_{ij}(n) \neq 0$, then $i, j \in V(n)$, and the connected components of $G(n)$ containing i and j either coincide or both meet B. In particular $\delta(n) = |E(n)| \geq D(i, j)$.*

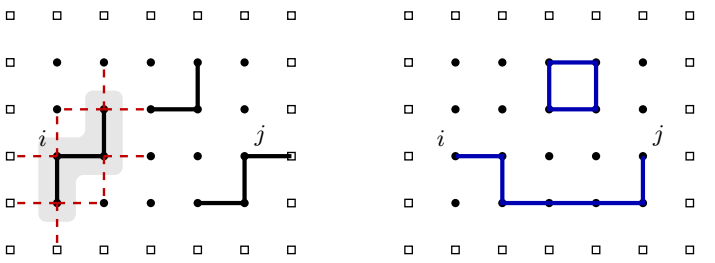
Proof. Let A be the component of i in $G(n)$, or $A = \{i\}$ if $i \notin V(n)$, and suppose $A \cap (B \cup \{j\}) = \emptyset$. Take $a = -1$ on A and $a = +1$ elsewhere. Every edge of $E(n)$ has both endpoints in A or both outside, so $a_e = 1$ on $E(n)$, while $a_i a_j = -1$; Lemma 6.9 gives $\hat{\phi}_{ij}(n) = 0$. Exchange i and j. A path in $E(n)$ from i to j, or two paths from i and from j to B, has at least $D(i, j)$ edges. □

Remark 6.11 (The parity form). Taking $a = -1$ at a single vertex $v \notin B$ in Lemma 6.9 gives more: for distinct $i \neq j$, if $\hat{\phi}_{ij}(n) \neq 0$, then every vertex $v \notin B$ has an even number of incident edges with n_e odd, except i and j, which have an odd number. In homological terms, $E_{\mathrm{odd}}(n) := \{e : n_e \text{ odd}\}$ is a mod-2 chain with boundary $\{i, j\}$ relative to B (Fig. 6.2b). By the handshake lemma, the component of i in $E_{\mathrm{odd}}(n)$ contains j or a vertex of B, so $|E_{\mathrm{odd}}(n)| \geq D(i, j)$. This is our reformulation, and it refines Corollary 6.10; the converse fails [56, Remark 2.3]. In the chain it forces $E_{\mathrm{odd}}(n) = [i, j]$.

The damping estimate is now immediate.

Lemma 6.12 (Damping; Chatterjee [49])**.** *For either perturbation, every $\beta \in (0, \infty]$ and all distinct $i, j \in V^\circ$,*

$$\mathbb{E}\big[\mathbb{E}[\phi_{ij}(J(t)) \mid J]^2\big] \leq e^{-2tD(i,j)} \, \mathbb{E}\phi_{ij}^2 \leq e^{-2tD(i,j)}.$$



(a) a component A of $G(n)$ (b) an admissible odd support

Figure 6.2: The gauge identity. (a) Thick edges are $E(n)$, squares are boundary vertices. The component A of i (shaded) meets neither j nor B; reversing the spins on A and the couplings on ∂A (dashed) reverses $\sigma_i\sigma_j$ and leaves h_n unchanged, so $\hat\phi_{ij}(n) = 0$. (b) The edges with odd n_e form a mod-2 chain with boundary $\{i, j\}$ up to boundary vertices, such as a path plus a cycle (Remark 6.11).

Proof. By (6.2) the left side is $\sum_n e^{-2|n|t}\hat\phi_{ij}(n)^2$ for OU and $\sum_n e^{-2\delta(n)t}\hat\phi_{ij}(n)^2$ for resampling. Since $\delta(n) \le |n|$, both are at most $\sum_n e^{-2\delta(n)t}\hat\phi_{ij}(n)^2$, and by Corollary 6.10 only $\delta(n) \ge D(i,j)$ contributes. Parseval gives $\sum_n \hat\phi_{ij}(n)^2 = \mathbb{E}\phi_{ij}^2 \le 1$. □

Theorem 6.13 (Local chaos; Chatterjee [49])**.** *Let the couplings be iid standard Gaussian, take $p \in (0,1)$, and let σ, σ' be the ground states before and after a perturbation of either kind with parameter p. Then $|\mathbb{E}(\sigma_i\sigma_j\sigma'_i\sigma'_j)| \le (1-p)^{D(i,j)}$ for all $i, j \in V^\circ$.*

Proof. With $p = 1 - e^{-t}$, $|\mathbb{E}[\sigma_i\sigma_j\,\mathbb{E}(\sigma'_i\sigma'_j \mid J)]| \le \mathbb{E}|\mathbb{E}(\sigma'_i\sigma'_j \mid J)| \le \mathbb{E}[\mathbb{E}(\sigma'_i\sigma'_j \mid J)^2]^{1/2} \le e^{-tD(i,j)}$ by Lemma 6.12. □

Since $\sigma_i\sigma_j$ and $\sigma'_i\sigma'_j$ are each uniform on $\{\pm1\}$, the theorem says that the relative orientation of two spins becomes independent of its perturbed version once $p \gg 1/D(i,j)$. Summing over pairs gives global chaos; we rename the source's growth constants $\alpha, \beta, \gamma, \delta$ to a, m, a', m'.

Theorem 6.14 (Site chaos; Chatterjee [49])**.** *Let the couplings be iid standard Gaussian. Suppose that for every $i \in V^\circ$ and $r \ge 1$ the number of j with $d(i,j) \le r$ is at most ar^m, and the number of j with $d(j,B) \le r$ is at most $a'|B|r^{m'}$. Then for both perturbations and every $p \in (0,1)$,*

$$\mathbb{E}\,q(\sigma,\sigma')^2 \le \frac{C(|V^\circ|p^{-m} + |B|^2p^{-2m'})}{|V^\circ|^2},$$

with C depending only on a, m, a', m'. For the box $\{0,\dots,L\}^d$ with the usual boundary and any coupling-independent boundary condition, $\mathbb{E}\,q^2 \le C(d)L^{-1}p^{-1}$ if $d = 1$ and $\mathbb{E}\,q^2 \le C(d)L^{-2}p^{-2}$ if $d \ge 2$; with free or periodic boundary conditions, $\mathbb{E}\,q^2 \le C(d)(Lp)^{-d}$ for all $d \ge 1$.

Proof. With $p = 1 - e^{-t}$, expand the square and use Theorem 6.13:

$$\mathbb{E}\,q(\sigma,\sigma')^2 = \frac{1}{|V^\circ|^2}\sum_{i,j\in V^\circ}\mathbb{E}(\sigma_i\sigma_j\sigma'_i\sigma'_j) \le \frac{1}{|V^\circ|} + \frac{1}{|V^\circ|^2}\sum_{i\ne j} e^{-tD(i,j)}.$$

Group the pairs by $k \le tD(i,j) < k+1$. There are at most $a|V^\circ|(k+1)^m t^{-m}$ pairs with $d(i,j) \le (k+1)/t$, and at most $(a'|B|)^2(k+1)^{2m'}t^{-2m'}$ pairs with $d(i,B) + d(j,B) \le (k+1)/t$;

the growth bounds hold trivially when $(k+1)/t < 1$. Summing against e^{-k} bounds the double sum by $C(|V^\circ|t^{-m} + |B|^2 t^{-2m'})$. Finally $t = -\log(1-p) \geq p$, and the diagonal term $1/|V^\circ|$ is absorbed. For boxes, $|V^\circ| \asymp L^d$, $|B| \asymp L^{d-1}$, $m = d$ and $m' = 1$. □

So the ground states before and after a perturbation of size $p \gg L^{-1}$ disagree on about half of the sites (the threshold "may not be optimal" [49, p. 4]). With couplings $\rho(J_e)$, Lemma 6.12 holds for every symmetric law and every β, so Theorems 6.13 and 6.14 extend to that generality, with both perturbations applied to the underlying Gaussian variables J_e and with ϕ_{ij} in place of $\sigma_i\sigma_j$ whenever $\beta < \infty$ or the coupling law has atoms. For free boundary conditions this is the theorem of Chen, Kim and Sen.

Theorem 6.15 (Chen–Kim–Sen [56])**.** *Let $G = ([N], E)$ be a finite graph with free boundary conditions, couplings $\rho(J_e)$ with J_e iid standard Gaussian and ρ odd, and $\beta \in [0, \infty]$. Let σ be drawn from the Gibbs measure for J and τ from that for $J(t)$, independently given (J, J'), for either perturbation, and let $B_r(i)$ be the ball of radius r in the graph distance. Then $\mathbb{E}\langle q(\sigma,\tau)^2\rangle_{0,t} \leq N^{-1}\max_i |B_r(i)| + e^{-tr}$ for every integer $r \geq 1$.*

Proof sketch. This is [56, Theorem 1.3] for pair interactions. By (6.2) and Corollary 6.10, $\mathbb{E}\langle\sigma_i\sigma_j\tau_i\tau_j\rangle_{0,t} = \mathbb{E}[\phi_{ij}(J)\phi_{ij}(J(t))] \leq e^{-td(i,j)}\mathbb{E}\phi_{ij}^2$. Split $\sum_{i,j}$ into pairs with $d(i,j) \leq r$ and the rest. □

6.5 Consequences: valleys, droplets and critical droplets

Chaos produces low-energy excitations. Let $\sigma(p)$ be the ground state after an OU perturbation with parameter p, and let $A := \{i \in V^\circ : \sigma_i(p) \neq \sigma_i\}$ be the *droplet* of overturned spins (for free or periodic boundary conditions, A is defined up to complement and ∂A is well defined). Write $\Delta(A) := H_J(\sigma^A) - H_J(\sigma)$ for its energy cost in the unperturbed system.

Lemma 6.16 (A variational inequality; Chatterjee [49])**.** *With $J(p) = (1-p)J + \sqrt{2p-p^2}\,J'$,*

$$\frac{\Delta(A)}{|\partial A|} \leq \frac{2\sqrt{2p-p^2}}{1-p}\max_{e\in E}|J'_e|.$$

Proof. On edges of ∂A the product σ_e changes sign, elsewhere it does not, so $\Delta(A) = H_J(\sigma(p)) - H_J(\sigma) = 2\sum_{\partial A} J_e\sigma_e$. Minimality of $\sigma(p)$ for $J(p)$ gives

$$0 \leq H_{J(p)}(\sigma) - H_{J(p)}(\sigma(p)) = -2\sum_{\partial A} J_e(p)\sigma_e = -(1-p)\Delta(A) - 2\sqrt{2p-p^2}\sum_{\partial A} J'_e\sigma_e,$$

and the last sum is at most $|\partial A|\max|J'_e|$ in absolute value. □

When the perturbation is chaotic the droplet is macroscopic, and the lemma gives a macroscopic region whose energy is small compared with its boundary. Let

$$F := \min\Big\{\frac{\Delta(A)}{|\partial A|} : A \subseteq V^\circ,\ \tfrac{1}{4}|V^\circ| \leq |A| \leq \tfrac{3}{4}|V^\circ|\Big\}.$$

Theorem 6.17 (Multiple valleys; Chatterjee [49])**.** *In the setting of Theorem 6.14 there are a universal constant C_1 and a constant C_2 depending only on a, m, a', m' such that for every $p \in (0, \frac{1}{2})$*

$$\mathbb{P}\big(F \geq C_1\sqrt{p\log|E|}\big) \leq \frac{1}{|E|} + \frac{C_2(|V^\circ|p^{-m} + |B|^2p^{-2m'})}{|V^\circ|^2}.$$

For the box $\{0,\dots,L\}^d$, $d \ge 2$, $L \ge 2$, with any coupling-independent boundary condition and every $K \ge 1$, $\mathbb{P}(F \ge C_1(d)L^{-1/2}\sqrt{K \log L}) \le C_2(d)K^{-2}$, with K^{-d} in place of K^{-2} for free or periodic boundary conditions and $d \ge 3$; in $d = 1$ the same argument gives C_2K^{-1}. Here $C_1(d)$ and $C_2(d)$ depend only on d.

Proof. Since $|A| = \frac{1}{2}|V^\circ|(1 - q(\sigma, \sigma(p)))$, the droplet violates the size constraint only if $|q| > \frac{1}{2}$, which has probability at most $4\mathbb{E}q^2$ by Chebyshev. On the complementary event, Lemma 6.16 and $p < \frac{1}{2}$ give $F \le 4\sqrt{2p}\max_e |J'_e|$. The Gaussian tail bound $\mathbb{P}(\max_e |J'_e| \ge x) \le 2|E|e^{-x^2/2}$ with $x = 2\sqrt{\log(2|E|)}$ makes this maximum exceed x with probability at most $1/|E|$. Apply Theorem 6.14. For boxes with $K < L/2$ take $p = K/L$ and use $\log|E| \le C(d)\log L$; for $K \ge L/2$ use $F \le 2\max_e |J_e|$ and the Gaussian tail, after enlarging $C_1(d)$. □

In a ferromagnet every macroscopic flip costs energy proportional to its boundary, so F is bounded below; here F is at most $L^{-1/2}$ up to logarithms, which bears on the debate of Chapter 2 without settling it [49, p. 12]. The good regions are random, not fixed in advance:

Theorem 6.18 (Chatterjee [49])**.** *In the setting of Theorem 6.14 there are $C_1, C_2, C_3 > 0$ depending only on the maximum degree of G such that for every $A \subset V^\circ$, $\mathbb{P}(\Delta(A)/|\partial A| < C_1) \le C_2e^{-C_3|\partial A|}$.*

The proof uses *special edges*, with $J_e > S_e + 2$ where S_e is the sum of $|J_f|$ over the edges f sharing an endpoint with e [49, Section 2.5]: stabilisers as in Lemma 5.7 with a margin, each adding at least 2 to $\Delta(A)$ when it lies in ∂A, and linearly many of them lie in a well-separated subset of ∂A outside an exponentially small event.

The droplet itself has a large boundary.

Theorem 6.19 (Rough droplets; Chatterjee [49])**.** *Consider $\{0,\dots,L\}^d$ with the usual boundary and some boundary condition, or with free or periodic boundary conditions, iid standard Gaussian couplings, and an OU perturbation with parameter $p \in (0,1)$. Then*

$$\mathbb{P}\Big(|\partial A| \ge C_1(1-p)\sqrt{p}\,L^d(\log L)^{-1/2}\Big) \ge 1 - 3e^{-C_2pL^d} - C_3L^{-C_4},$$

with $C_1, \dots, C_4 > 0$ depending only on d.

Proof sketch. We omit the concentration estimates. Gaussian concentration and a lower bound of order L^d on the ground-state energy show that $\Delta(A) \ge cpL^d$ with the stated probability. Then Lemma 6.16 with $\max_e |J'_e| = O(\sqrt{\log L})$ converts this into the bound on $|\partial A|$ [49, Section 2.4]. □

For $p = L^{-\alpha}$ the boundary is at least $L^{d-\alpha/2}$ up to logarithms, consistent with a fractal dimension strictly greater than $d-1$ [49, abstract]. The droplet need not be connected [49, p. 5], and a large total boundary may come from many small pieces (Fig. 6.3).

Finally, chaos forces large critical droplets. For an edge $e = \{i,j\}$ let σ^1 and σ^2 minimise the energy under the constraints $\sigma_i = \sigma_j$ and $\sigma_i = -\sigma_j$. Neither depends on J_e, and for every value of J_e the ground state is one of them. The *critical droplet* $\mathcal{D}(e)$ is the set where they disagree (for free or periodic boundary conditions, its size is $\min\{|\mathcal{D}(e)|, |V \setminus \mathcal{D}(e)|\}$); it is connected. This is the finite-volume form of the critical droplets of Newman and Stein discussed in Chapter 4.

Theorem 6.20 (Critical droplets; Chatterjee [49])**.** *In the setting of Theorem 6.14,*

$$\frac{1}{|E|}\sum_{e \in E}\mathbb{E}|\mathcal{D}(e)| \ge \frac{C|V^\circ|}{|E|\max\{|V^\circ|^{-1/m}, (|B|/|V^\circ|)^{1/m'}\}},$$

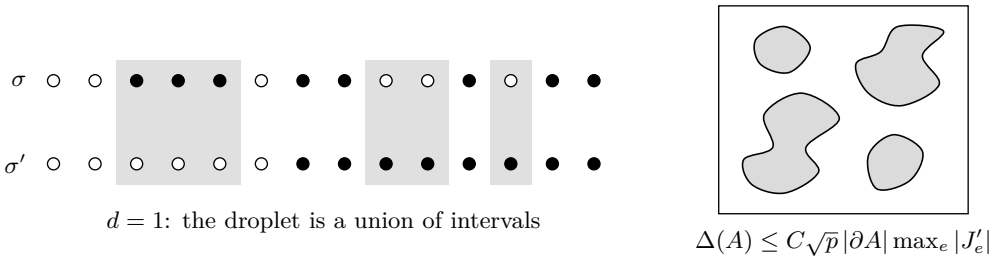


Figure 6.3: Perturbation droplets. Left: in $d = 1$ the two ground states disagree on a union of intervals (shaded) ending at the edges where $\operatorname{sgn} J_e$ changed, about $\sqrt{p}\,L$ of them under the OU perturbation. Right: in higher dimensions the droplet A may have many pieces; Theorem 6.19 bounds the total boundary $|\partial A|$, not the roughness of a single interface.

with $C > 0$ depending only on a, m, a', m'. For $\{0, \dots, L\}^d$ with any coupling-independent boundary condition, or with free or periodic boundary conditions, the average of $\mathbb{E}|\mathcal{D}(e)|$ over edges is at least $C(d)L$ for all $d \geq 1$. With periodic boundary conditions and $d \geq 2$, $\mathbb{E}|\partial\mathcal{D}(e)|^{d/(d-1)} \geq C(d)L$ for every edge e.

Proof sketch. We omit the bookkeeping [49, Section 2.6]. Resampling a Binomial$(|E|, p)$ number of uniformly chosen edges one at a time realises a resampling perturbation, and each step changes $\sigma_i\sigma_j$ with the probability that one uniform resampling does, so $\mathbb{P}(\sigma_i\sigma_j \neq \sigma'_i\sigma'_j) \leq |E|p\,\mathbb{P}(\sigma_i\sigma_j \neq \tilde{\sigma}_i\tilde{\sigma}_j)$, with $\tilde{\sigma}$ the ground state after one uniform resampling. The left side, summed over pairs, is of order $|V^\circ|^2$ by Theorem 6.14 once p is large enough. The right side counts pairs separated by the critical droplet of the resampled edge. The torus bound follows from an edge-isoperimetric inequality. □

6.6 Boundary sensitivity

The two-dimensional random-field ground state forgets its boundary condition exponentially fast [75]. In the EA model a single spin always depends on the boundary, since reversing every boundary spin reverses the ground state; the meaningful quantity is the product of two neighbouring spins.

Theorem 6.21 (Chatterjee [49])**.** *Let $d \geq 2$, $L \geq 2$, let $V \subseteq \mathbb{Z}^d$ be the box of side length L centred at the origin with vertex boundary B, and let the couplings be iid standard Gaussian. For $\eta \in \{\pm 1\}^B$ let σ^η be the ground state with boundary condition η. Fix a neighbour j of the origin and let $\mathcal{E}$ be the event that $\sigma_0^\eta \sigma_j^\eta$ is not the same for all η. Then $\mathbb{P}(\mathcal{E}) \geq 1/(32L)$.*

Proof sketch. We omit the bookkeeping, which is in Section A.13 following [49, Section 2.7]. On a torus of side $100L$, reversing the couplings on the edge boundary ∂U of $U = \{-L, \dots, L\}^d$ is the gauge transformation $a = 1 - 2\mathbf{1}_U$, so it reverses $\sigma_0\sigma_k$ for $k = 4Le_1$. Do it in two steps, first on the upper half S_1 of ∂U and then on the lower half S_2 (Fig. 6.4); each reversed coupling field has the law of J. In step r, $\sigma_0\sigma_k$ changes only if $\sigma_u\sigma_v$ changes for some edge $\{u, v\}$ of a path P_r from 0 to k, of length at most $16L$, along which the boxes $u + V$ contain no edge of S_r. The couplings inside $u + V$ do not change, so the box sees only a change of boundary condition, which has probability at most $\mathbb{P}(\mathcal{E})$ per edge. Hence $1 \leq 2 \cdot 16L\,\mathbb{P}(\mathcal{E})$. □

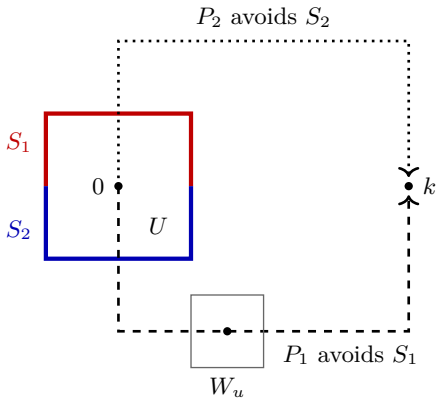


Figure 6.4: The geometry behind Theorem 6.21. Reversing the couplings on ∂U reverses $\sigma_0\sigma_k$. Done in two steps, on S_1 (red) and then on S_2 (blue), each step changes $\sigma_0\sigma_k$ only if an edge changes on a path that avoids the reversed half, which is boundary sensitivity in a small box W_u.

The argument needs only a symmetric atomless law, and $d \geq 2$ for the two paths (Section A.14). In $d = 1$ the conclusion holds because $\sigma_0^\eta \sigma_j^\eta$ depends on η exactly when $|J_{0j}|$ is the smallest coupling in V, an event of probability $1/|E(V)| \geq 1/(32L)$.

Route warning 6.22 (No exponential decay of boundary influence). A route to two-dimensional uniqueness modelled on the random-field Ising model would need boundary influence on local observables to decay exponentially, as it does there [75]. For the EA model this is false: by Theorem 6.21 the influence on a nearest-neighbour product decays at most like $1/L$. What separates the two models is the FKG inequality [49, p. 9] (Chapter 10). The bound is compatible with uniqueness, since $1/L \to 0$; it rules out a mechanism, not a conclusion.

6.7 The spectral sample

Squared Hermite coefficients form a probability measure: at $\beta = \infty$ with an atomless law, $\mu_{ij}(n) := \hat{\phi}_{ij}(n)^2$ sums to $\mathbb{E}(\sigma_i\sigma_j)^2 = 1$. A sample from μ_{ij} is the *spectral sample* of $\sigma_i\sigma_j$, and by (6.2) $\mathbb{E}[\sigma_i\sigma_j\sigma_i^t\sigma_j^t] = \sum_n \mu_{ij}(n)e^{-|n|t}$ for the OU perturbation. The size of the spectral sample therefore sets the chaos scale, and Corollary 6.10 says that its support connects i to j or both to the boundary, an analogy with percolation [22, p. 5]. Basu Roy Chowdhury and Ganguly show that it is typically strictly longer than the shortest connection. Let $V = \{1, \dots, n\}^2$ with free boundary conditions and Gaussian couplings, n even, $u = (1, n/2)$, $v = (n, n/2)$, let L_{uv} be the straight line of edges from u to v, and write $\mu(S) := \sum_{n' : E(n') = S} \mu_{uv}(n')$ for edge sets S.

Theorem 6.23 (Basu Roy Chowdhury–Ganguly [22])**.** *There are $c, \varepsilon > 0$, independent of n, such that $\mu(L_{uv}) \leq e^{-cn}$ and $\sum_{|S| \leq (1+\varepsilon)n} \mu(S) \leq e^{-cn}$.*

Proof sketch. We omit the extension from the line to sets of size $(1 + \varepsilon)n$. Conditioning on the couplings in a set S gives $\mathbb{E}[\mathbb{E}(\sigma_u\sigma_v \mid J_S)^2] = \sum_{E(n') \subseteq S} \mu_{uv}(n')$, so it suffices to show that $|\mathbb{E}(\sigma_u\sigma_v \mid J_{L_{uv}})| \leq e^{-cn}$ with high probability. A vertical cut through an edge e^* of the line would reverse $\sigma_u\sigma_v$, but J_{e^*} is frozen by the conditioning. A *barrier* around e^*, a ring of weak couplings ($|J| \leq 1$) with strong spokes ($|J| \geq 100$), makes $\sigma_u\sigma_v$ independent of the sign of J_{e^*}, so the cut can route around it [22, Lemma 3.1]. Barriers depend only on coupling magnitudes,

occur with positive probability, and have linearly many disjoint chances along the line, so all of them fail only with probability e^{-cn}. □

Aligned endpoints avoid exponentially many shortest paths [22, Remark 2.4]; numerics of Bray and Moore suggest, heuristically, a fractal spectral sample of size of order $d^{1+\kappa}$ at distance d. The spectral view follows noise sensitivity [25, 105], [49, p. 13]. The superconcentration–chaos equivalence has since been proved for first-passage percolation [1] and for supercritical chemical distance [192]; see also [72] for a disordered ferromagnet, and [50] for stability of optimisers under perturbations of bounded size.

6.8 What the chaos theorems do and do not show

Route warning 6.24 (Site chaos is not evidence of many ground states). It is tempting to read site chaos, macroscopic low-energy excitations or large critical droplets as evidence for many ground states (Chapter 5). The chain refutes this: Theorems 6.14 and 6.20 cover $d = 1$ [49, pp. 4, 8], in $d = 1$ the quantity F of Theorem 6.17 is at most of order L^{-1} [49, p. 7], and the droplet boundary has order $\sqrt{p}\,L$, matching Theorem 6.19 up to a logarithm, yet the chain has a unique, unfrustrated ground-state pair (Example 5.4). These signatures come from sign disorder and gauge symmetry; one system-spanning wall already gives site overlap near zero with edge overlap near one. Multiplicity would show in the edge overlap at small perturbation, which is open (Open problem 6.26). This reading of $d = 1$ is ours; Chatterjee's own caveats follow.

Route warning 6.25 (Finite volume). The results are finite-volume, zero-temperature statements, which "do not settle the infinite volume question of whether the Edwards–Anderson model has incongruent ground states" [49, p. 2]. Translating a finite-volume chaos bound into a statement about metastates needs a limit that none of these papers takes.

To summarise: the free-energy variance is of volume order, so the edge overlap is not chaotic; the gauge identity puts the Hermite spectrum of $\sigma_i\sigma_j$ at degree at least $D(i, j)$, which gives site chaos, low-energy excitations and large droplets; a two-step gauge flip gives the $1/L$ bound on boundary sensitivity.

Open problem 6.26 (Chaos in the EA model)**.** (a) *Edge chaos in the strong sense.* "is it true that under the periodic boundary condition on $\{-L, \dots, L\}^d$, for fixed neighboring sites i and j, $\lim_{p\to 0} \lim_{L\to\infty} |\mathbb{E}(\sigma_i\sigma_j\sigma'_i\sigma'_j)| < 1$?" [49, p. 10]. Arguin and Hanson conjecture that in $d = 2$ the edge overlap stays close to one on an event of probability close to one [12, Conjecture 1.2].

(b) *The size of F.* Its correct order for $d \geq 2$, and the existence of macroscopic regions with small $\Delta(A)/\sqrt{|\partial A|}$ [49, p. 7].

(c) *The chaos threshold.* Do perturbations smaller than L^{-1} produce site chaos [49, p. 4]?

(d) *Long-range order.* Does the EA model, like the RFIM, order in $d \geq 3$ [49, p. 10]?

(e) *Self-averaging of the edge overlap.* Show that $\mathrm{Var}\langle q^e(\sigma, \tau)\rangle_{0,0} \to 0$ as the graph grows [46, Section 1.7, problem (8)].

(f) *Fractality of the spectral sample.* Replace the factor $1 + \varepsilon$ in Theorem 6.23 by a sequence $\varepsilon_n \to \infty$, and study the Fourier spectrum at positive temperature [22, Remark 2.5].

Problem (a) meets Chapter 5 through Remark 6.8; the random-field comparison is in Chapter 10, and the programme as a whole in Chapter 16.

Part IV

Graphical representations

Chapter 7

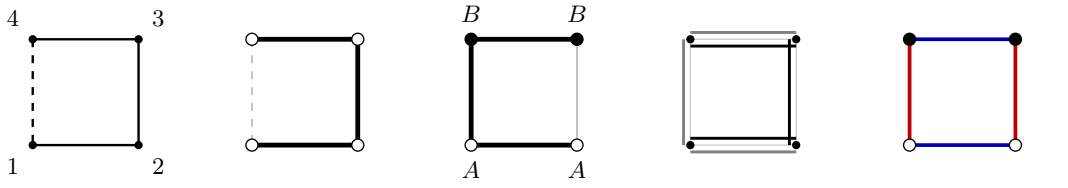
Cluster representations of frustrated Ising systems

In an Ising ferromagnet, order has a geometric face. The Fortuin–Kasteleyn (FK) random-cluster representation [90] writes the two-point function $\langle \sigma_x \sigma_y \rangle$ as the probability that x and y are joined by a path of open bonds, so long-range order and percolation are the same statement, and the whole toolbox of percolation (monotonicity, comparison, uniqueness of the infinite cluster) becomes available for the spins. This chapter asks how much of that dictionary survives frustration. The question matters for the central question of the review in a direct way: if spin-glass order on $\mathbb{Z}^d$ had a geometric signature that one could prove, the ordering problem would become a percolation problem.

The answer has two halves. Every representation of this chapter is exact, with explicit marginals and conditional laws, and some classical percolation arguments go through once the couplings are averaged. But in each representation percolation by itself is weaker than order, because frustration lets the sign information carried by a cluster cancel. By the end of the chapter the reader will be able to write the frustrated Edwards–Sokal coupling and its two-replica relatives (double FK, the Chayes–Machta–Redner colouring, signed FK); prove that averaging one coupling restores insertion tolerance, so that there is at most one infinite FK cluster (Gandolfi–Keane–Newman), and that on the Nishimori line the averaged FK bonds are independent (Gandolfi, Mazza); explain with De Santis and Gandolfi's frustrated Peierls argument why, if two-dimensional uniqueness holds as conjectured, FK percolation cannot be the order signal there; and see all the two-replica objects as functions of one pair of FK configurations. Chapter 8 then studies the CMR blue clusters; the formulas of every representation are collected in Appendix D.

7.1 One plaquette, many pictures

Take the frustrated square of Chapter 1, with three positive couplings and one negative one (Fig. 7.1(a)): every configuration violates an odd number of its edges. Each representation records this in its own way. FK bonds open only on satisfied edges, so the frustrated cycle is never fully open (panel (b)); signed FK labels the sites of a cluster by the coupling sign along open paths (panel (c)); two-replica representations colour the edges satisfied in both replicas, in one, or in neither (panels (d) and (e)). The chapter explains these pictures one at a time and then shows that they are all functions of one probability space.



(a) couplings J (b) FK of σ (c) FK of τ (d) ω^1 and ω^2 (e) CMR colouring

Figure 7.1: One frustrated plaquette (positive couplings solid, the negative one dashed) under the representations of this chapter; open circles are $+1$, filled circles -1. (b) With $\sigma \equiv +1$, FK bonds open only on the three satisfied edges. (c) With $\tau = (+,+,-,-)$ the open bonds are $12, 34, 41$, and the coupling sign along open paths splits the cluster into companions A and B, on which τ is constant. (d) ω^1 of σ (inner strands) and ω^2 of τ (outer). (e) Their union coloured by satisfaction, drawn on the overlap $q = \sigma\tau$: blue edges (satisfied in both) join equal overlaps, red edges (satisfied in one) opposite ones, and the grey cycle has an even number of red edges.

7.2 The Edwards–Sokal coupling with frustration

Throughout this section $G = (V, E)$ is a finite graph, $J \in \mathbb{R}^E$ a fixed coupling configuration and $\beta \geq 0$. The Gibbs measure with free boundary condition is $\mu^{\beta}_{G,J}(\sigma) \propto \exp(\beta \sum_{xy \in E} J_{xy}\sigma_x\sigma_y)$, with expectation $\langle\cdot\rangle_J$, and $p_e = 1 - e^{-2\beta|J_e|}$. An edge $e = xy$ is satisfied by σ if $J_e\sigma_x\sigma_y > 0$. We identify $\omega \in \{0,1\}^E$ with its set of open edges, and write $k(\omega)$ for the number of connected components of the graph (V, ω), isolated vertices included.

Definition 7.1 (Frustration-free configurations and relative parity)**.** A configuration $\omega \subseteq E$ is *J-frustration-free* if $J_e \neq 0$ for every $e \in \omega$ and every cycle $C \subseteq \omega$ has $\prod_{e \in C} \operatorname{sgn} J_e = +1$. Write FF_J for the set of such ω. For $\omega \in \mathrm{FF}_J$ and $x, y \in V$, the *relative parity* $\operatorname{sgn}_J(x, y; \omega)$ is the product of $\operatorname{sgn} J_e$ along any open path from x to y, and 0 if there is no such path. Frustration-freeness is exactly what makes the relative parity independent of the path.

Chapter 1 set up the Edwards–Sokal (ES) coupling for couplings of both signs [80, 132, 170, 230] and used it to prove Newman's high-temperature uniqueness theorem. We recall it for a finite graph with free boundary condition: an open bond certifies that its edge is satisfied, and given the bonds each cluster carries one free global sign. The ES coupling is

$$\Phi_{G,J}(\sigma, \omega) = \frac{1}{\widetilde{Z}_{G,J}} \prod_{e=xy \in E} \Big[(1-p_e)\mathbf{1}\{\omega_e = 0\} + p_e \mathbf{1}\{\omega_e = 1\}\mathbf{1}\{J_e\sigma_x\sigma_y > 0\}\Big], \tag{7.1}$$

with $\widetilde{Z}_{G,J} = e^{-\beta\sum_e |J_e|} Z^{\beta}_{G,J}$, where $Z^{\beta}_{G,J}$ is the Ising partition function; summing an edge factor over ω_e gives $e^{-\beta|J_e|}e^{\beta J_e\sigma_x\sigma_y}$. Its spin marginal is $\mu^{\beta}_{G,J}$, and its bond marginal, the *frustrated FK measure*, is

$$\varphi_{G,J}(\omega) = \frac{1}{\widetilde{Z}_{G,J}} \prod_{e \in E} p_e^{\omega_e}(1-p_e)^{1-\omega_e}\, 2^{k(\omega)}\, \mathbf{1}\{\omega \in \mathrm{FF}_J\}, \tag{7.2}$$

because a frustration-free ω is compatible with exactly $2^{k(\omega)}$ spin configurations and any other ω with none. Given σ, the bonds are independent, Bernoulli(p_e) on satisfied edges and closed elsewhere. Given $\omega \in \mathrm{FF}_J$, choose a root x_C in each cluster C; then

$$\sigma_y = \sigma_{x_C} \operatorname{sgn}_J(x_C, y; \omega) \quad (y \in C), \qquad (\sigma_{x_C})_C \text{ independent fair signs.} \tag{7.3}$$

These facts are proved in Chapter 1 and [170, §1]. Taking expectations in (7.3),

$$\langle\sigma_x\sigma_y\rangle_J = \mathbb{E}_{\varphi_{G,J}}[\operatorname{sgn}_J(x,y;\omega)] = \varphi_{G,J}(x\leftrightarrow y,\ \operatorname{sgn}_J=+1) - \varphi_{G,J}(x\leftrightarrow y,\ \operatorname{sgn}_J=-1), \quad (7.4)$$

since a pair in one cluster has a determined product and a pair in two clusters a product of independent fair signs. In particular $|\langle\sigma_x\sigma_y\rangle_J| \le \varphi_{G,J}(x\leftrightarrow y)$, with equality when $J \ge 0$ [170, (1.8) and (4.10)]. Formula (7.4) contains the whole difficulty of the subject: the connectivity $\varphi(x\leftrightarrow y)$ can stay large while the two signed terms cancel.

Remark 7.2 (The conditioned form). Comparing (7.2) with the ferromagnetic random-cluster measure $\varphi_{G,|J|}$, which has the same edge parameters, gives $\varphi_{G,J} = \varphi_{G,|J|}(\cdot \mid \mathrm{FF}_J)$, $Z^{\beta}_{G,J} = Z^{\beta}_{G,|J|}\,\varphi_{G,|J|}(\mathrm{FF}_J)$ and $\langle\sigma_x\sigma_y\rangle_J = \mathbb{E}_{\varphi_{G,|J|}}[\operatorname{sgn}_J(x,y;\omega) \mid \mathrm{FF}_J]$, Aizenman's "frustration-adjusted FK" representation [2, Thm. 7.2]. Since FF_J is decreasing, positive association of $\varphi_{G,|J|}$ gives the domination $\varphi_{G,J} \preccurlyeq \varphi_{G,|J|}$ of Chapter 1; nothing comparable holds in the other direction, one way to see why FK methods give non-order more easily than order.

In infinite volume one works with weak limits of these measures, or with ES and FK Gibbs measures defined by their finite-volume conditional laws; for fixed ±1 couplings their theory, including the one-to-one correspondence of spin and ES Gibbs states, is in [104, §3].

7.3 Averaging over the couplings restores insertion tolerance

For a ferromagnet, every bond can be opened with conditional probability bounded below, whatever the other bonds are; this *insertion tolerance* drives the Burton–Keane uniqueness argument [38]. At a fixed frustrated J it fails: if the endpoints of e are already connected by an open path whose coupling product has the wrong sign, opening e would close a frustrated cycle, so $\varphi_{G,J}(\omega_e = 1 \mid \omega_{E\setminus e}) = 0$. Gandolfi, Keane and Newman observed that averaging one coupling repairs this: if the sign of J_e is unknown, at least one sign makes opening e legal, and changing one coupling changes the partition function by a bounded factor.

Fix $p \in (0,1)$ and put $\beta = -\frac{1}{2}\log(1-p)$, so that $p = 1 - e^{-2\beta}$. Let $J = (J_e)$ be iid fair signs and consider the joint law of couplings and bonds on a finite graph,

$$\mathbb{P}_G(J,\omega) = 2^{-|E|}\,\varphi_{G,J}(\omega), \qquad J \in \{-1,1\}^E,\ \omega \in \{0,1\}^E. \quad (7.5)$$

Lemma 7.3 (Insertion tolerance after averaging; Gandolfi–Keane–Newman [101, Thm. 3])**.** *For every finite graph G, every $e \in E$ and every $\omega' \in \{0,1\}^{E\setminus e}$ with $\mathbb{P}_G(\omega_{E\setminus e} = \omega') > 0$,*

$$\mathbb{P}_G(\omega_e = 1 \mid \omega_{E\setminus e} = \omega') \ \ge\ \frac{p}{2}.$$

Proof. We follow [101, pp. 524–526]. Write $w(\omega) = p^{|\omega|}(1-p)^{|E|-|\omega|}2^{k(\omega)}$, so $\varphi_{G,J}(\omega) = w(\omega)\mathbf{1}\{\omega \in \mathrm{FF}_J\}/\widetilde{Z}_J$. Fix the couplings J' off e; it suffices to show that for every J',

$$\sum_{J_e=\pm1} \varphi_{G,(J',J_e)}(\omega' \cup e) \ \ge\ \frac{p}{2-p} \sum_{J_e=\pm1} \varphi_{G,(J',J_e)}(\omega'), \quad (7.6)$$

because summing over J' then bounds the conditional odds of $\{\omega_e = 1\}$ below by $p/(2-p)$, which is the claimed bound $p/2$ on the probability. Let x, y be the endpoints of e.

Case 1: x and y are not connected in ω'. Adding e creates no cycle, so $\omega' \cup e$ is J-frustration-free exactly when ω' is, and k drops by one. Hence $\varphi_J(\omega' \cup e) = \frac{p}{2(1-p)}\varphi_J(\omega')$ for both values of J_e, and (7.6) holds because $\frac{p}{2(1-p)} \ge \frac{p}{2-p}$.

Case 2: x and y are connected in ω'. If ω' is J'-frustrated, both sides of (7.6) vanish. Otherwise let $s = \operatorname{sgn}_{J'}(x, y; \omega')$. Now $\omega' \cup e$ is frustration-free if and only if $J_e = s$, and k does not change. Writing $Z_\pm$ for $\widetilde{Z}_J$ with $J_e = \pm s$, the left side of (7.6) is $\frac{p}{1-p} w(\omega')/Z_+$ and the right side is $\frac{p}{2-p} w(\omega')(1/Z_+ + 1/Z_-)$. Changing the single coupling J_e multiplies every Ising weight by a factor in $[e^{-2\beta}, e^{2\beta}]$, so $Z_+/Z_- \le e^{2\beta} = 1/(1-p)$, and

$$\frac{p/(1-p)}{1 + Z_+/Z_-} \;\ge\; \frac{p/(1-p)}{1 + 1/(1-p)} = \frac{p}{2-p}. \qquad \square$$

The couplings are integrated out. This *annealed* insertion tolerance makes Burton–Keane work for frustrated FK, and variants of it reappear for the two-replica colourings in Chapter 8.

Now let $\mathbb{T}_L^d$ be the discrete torus and let $\mathbb{P}$ be any weak limit, along some $L_n \to \infty$, of the joint laws $\mathbb{P}_{\mathbb{T}^d_{L_n}}$ of (7.5), viewed as laws on $\{-1,1\}^{E_d} \times \{0,1\}^{E_d}$. Write $\mathbb{P}^J$ for its conditional law of ω given J. Translation invariance of the tori makes $\mathbb{P}$ translation invariant.

Theorem 7.4 (Uniqueness and existence of the infinite FK cluster; Gandolfi–Keane–Newman [101, Thm. 3, Cor. 1–2])**.** *Let $d \ge 1$, let the couplings be iid fair ± 1, let $p = 1 - e^{-2\beta} \in (0,1)$, and let $\mathbb{P}$ be a weak limit of periodic joint laws as above.*

(i) For every edge e, $\mathbb{P}(\omega_e = 1 \mid \omega_{E_d \setminus e}) \ge p/2$ almost surely, and the bond marginal of $\mathbb{P}$ stochastically dominates Bernoulli($p/2$) bond percolation.

(ii) For almost every J, $\mathbb{P}^J$-almost surely there is at most one infinite cluster.

(iii) If $p/2 > p_c^{\mathrm{bond}}(\mathbb{Z}^d)$, then for almost every J, $\mathbb{P}^J$-almost surely there is exactly one infinite cluster, and $\mathbb{P}(0 \leftrightarrow \infty) \ge \theta^{\mathrm{bond}}(p/2) > 0$.

Part (iii) needs $p_c^{\mathrm{bond}}(\mathbb{Z}^d) < \frac{1}{2}$, which holds for $d \ge 3$ and fails for $d = 2$ [101, p. 526].

Proof. (i) On each torus, Lemma 7.3 gives $\mathbb{E}[\mathbf{1}\{\omega_e = 1\} f] \ge \frac{p}{2}\mathbb{E}[f]$ for every nonnegative function f of the bonds other than e. For local f this inequality passes to weak limits, and a monotone-class argument extends it to all bounded measurable $f \ge 0$; this is the conditional bound. A conditional lower bound $p/2$ given all other bonds implies, in finite volume, stochastic domination of Bernoulli($p/2$) (open the edges one at a time with a monotone coupling), and domination is preserved under weak limits. (ii) The bond marginal of $\mathbb{P}$ is translation invariant and, by (i), has positive finite energy for insertions. The Burton–Keane argument, in the form [101, Thm. 1] that needs only these two properties, gives at most one infinite cluster almost surely; Fubini's theorem transfers the statement to $\mathbb{P}^J$ for almost every J. (iii) If $p/2 > p_c^{\mathrm{bond}}$, Bernoulli($p/2$) has an infinite cluster almost surely and the origin lies in it with probability $\theta^{\mathrm{bond}}(p/2) > 0$; both events are increasing, so (i) transfers them to $\mathbb{P}$, and Fubini and (ii) give exactly one infinite cluster for almost every J. $\square$

The finite-volume details of (i) are in Section A.15. Uniqueness holds in every dimension and at every temperature, while existence needs a margin ($p/2 > p_c$) that two dimensions cannot provide; and neither says anything about $\langle \sigma_0 \sigma_x \rangle_J$, because of the signs in (7.4). As Newman put it, with a unique infinite cluster $\varphi(0 \leftrightarrow x)$ does not decay, but the two signed terms may still cancel asymptotically [170, p. 258].

7.4 Frustrated FK clusters percolate in the plane

In two dimensions $p/2 \le \frac{1}{2} = p_c^{\mathrm{bond}}(\mathbb{Z}^2)$, so averaging says nothing. De Santis and Gandolfi found a route that works for *every* sign configuration: a Peierls argument with one combinatorial lemma, that a closed dual circuit can always be "reopened" in two complementary halves without creating frustration.

Let γ be the set of edges of $\mathbb{Z}^2$ crossed by a self-avoiding circuit of the dual lattice, inside a box Λ. Given bonds ω off γ and a subset $A \subseteq \gamma$, call A *conditionally unfrustrated* if opening all edges of A, keeping $\gamma \setminus A$ closed, creates no frustrated cycle.

Lemma 7.5 (Splitting a dual circuit; De Santis–Gandolfi [71, Lemma 2.1])**.** *For every box Λ, every $J \in \{-1,1\}^{E(\Lambda)}$, every dual circuit image $\gamma \subseteq E(\Lambda)$ and every J-frustration-free configuration of $E(\Lambda) \setminus \gamma$, there is $A \subseteq \gamma$ such that both A and $\gamma \setminus A$ are conditionally unfrustrated.*

Proof sketch. We omit the bookkeeping of the cycle decomposition; see [71, pp. 1786–1787]. Build A greedily, adding an edge of γ whenever this keeps A conditionally unfrustrated. If $\gamma \setminus A$ contained a frustrated self-avoiding cycle π, then π would cross γ an even number $k \ge 2$ of times (it must enter and leave the interior), and each rejected crossing edge b_j comes with an open frustrated cycle through b_j and A. Replacing each b_j in π by the rest of its witnessing cycle yields a cycle inside the open edges of $A \cup (\Lambda \setminus \gamma)$ with sign $(-1)^{k+1} = -1$, contradicting the construction of A. □

Theorem 7.6 (FK percolation for every $\pm\mathbf{1}$ configuration; De Santis–Gandolfi [71, Cor. 2.3])**.** *On $\mathbb{Z}^2$, for every $J \in \{-1,1\}^{E_2}$ and every $p > 16/17$, that is $\beta > \frac{1}{2}\log 17$, every weak limit ν of the frustrated FK measures $\varphi_{\Lambda_n,J}$ along boxes $\Lambda_n \uparrow \mathbb{Z}^2$ has an infinite open cluster ν-almost surely.*

Proof sketch. We give the Peierls estimate and a short form of the passage from "finitely many closed circuits" to an infinite cluster; the details are in [71, pp. 1787–1790]. Fix a dual circuit image γ, the bonds ω off γ, and A from Lemma 7.5, with $|A| = m$. Opening any $S \subseteq A$, with $\gamma \setminus A$ closed, keeps the configuration frustration-free, multiplies $p^{|\omega|}(1-p)^{|E|-|\omega|}$ by $(p/(1-p))^{|S|}$ and lowers $k(\omega)$ by at most $|S|$; so its weight relative to $\{\gamma \equiv 0\}$ is at least $(p/(2(1-p)))^{|S|}$. Summing over $S \subseteq A$ and over $S \subseteq \gamma \setminus A$ (the empty set is counted twice),

$$\varphi_{\Lambda,J}(\gamma \equiv 0 \mid \omega_{\Lambda\setminus\gamma}) \le \frac{1}{x^m + x^{|\gamma|-m} - 1} \le \frac{2}{x^m + x^{|\gamma|-m}} \le x^{-|\gamma|/2}, \qquad x = \frac{2-p}{2(1-p)},$$

using $x^m + x^{|\gamma|-m} \ge 2$ and the arithmetic–geometric mean inequality. The bound is uniform in the outside configuration and passes to every weak limit. At most $k3^k$ dual circuits of length k surround the origin, and $\sum_k k3^k x^{-k/2} < \infty$ exactly when $9 \cdot 2(1-p) < 2-p$, that is $p > 16/17$. By Borel–Cantelli only finitely many dual circuits around the origin are closed, so for n large no closed dual circuit surrounds the box $\Lambda'_n := \{-n,\dots,n\}^2$. Then some open cluster meeting Λ'_n is infinite: otherwise Λ'_n together with the open clusters that meet it would be a finite connected set, and its outer edge boundary, all of whose edges are closed, would contain a closed dual circuit surrounding Λ'_n. □

The same argument works with Peierls contours in $\mathbb{Z}^d$, $d \ge 3$ [71, Lemma 2.5], and its two-replica version (red-bond parity in place of frustration) gives an infinite grey cluster in the CMR colouring of Section 7.6 [215]. De Santis and Gandolfi's conclusion is that in two dimensions "percolation gives no clear indication on whether a phase transition could take place in the EA model" [71, p. 1782] (Route warning 7.9).

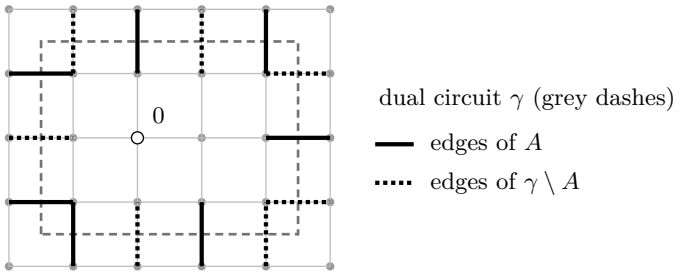


Figure 7.2: The frustrated Peierls argument. The primal edges crossed by a dual circuit γ around the origin split into two classes, each of which can be opened while the other stays closed without closing a frustrated cycle (Lemma 7.5), so whatever the couplings a closed circuit costs a factor $x^{-|\gamma|/2}$, $x = (2-p)/(2(1-p))$, which beats the $k3^k$ circuits of length k once $p > 16/17$. The split shown is illustrative: A depends on the couplings and on the bonds off γ.

7.5 The Nishimori line: averaged FK bonds are independent

Let the couplings be iid with the biased law $\nu_p = p\delta_{+1} + (1-p)\delta_{-1}$, $p \in [\frac{1}{2}, 1)$, and let $\beta = \beta_{\mathrm{N}}(p)$ be the Nishimori temperature, $\tanh \beta_{\mathrm{N}} = 2p - 1$. Gauge symmetry (Chapter 14) then makes the disorder average of $\varphi_{G,J}$ a product measure.

Theorem 7.7 (Gandolfi [99, Thm. 5]; Mazza [158, §2])**.** *Let $G = (V, E)$ be a finite graph, let J be iid with law ν_p, $p \in [\frac{1}{2}, 1)$, and let $\beta = \beta_{\mathrm{N}}(p)$. Then for every $\omega \subseteq E$,*

$$\mathbb{E}\big[\varphi_{G,J}(\omega)\big] = (\tanh\beta)^{|\omega|}(1-\tanh\beta)^{|E|-|\omega|}, \qquad \tanh\beta = 2p-1.$$

That is, the disorder-averaged FK configuration is Bernoulli$(2p-1)$ bond percolation.

Gandolfi proves the identity for every coupling-independent boundary condition, wired ones included [99, Thm. 5 and Cor. 6]; we give the free-boundary argument, following the gauge computation of [158, §2].

Proof. On the Nishimori line $p = e^{\beta}/(2\cosh\beta)$ and $1-p = e^{-\beta}/(2\cosh\beta)$, so the disorder law is $\mathbb{P}(J) = (2\cosh\beta)^{-|E|}\exp(\beta\sum_e J_e)$. For $\tau \in \{-1,1\}^V$ let $J^\tau_{xy} = J_{xy}\tau_x\tau_y$. If g is gauge invariant, $g(J^\tau) = g(J)$ for all τ, then averaging the change of variables $J \mapsto J^\tau$ over τ gives

$$\mathbb{E}g(J) = 2^{-|V|}\sum_J g(J)\sum_\tau \mathbb{P}(J^\tau) = 2^{-|V|}(2\cosh\beta)^{-|E|}\sum_J g(J)\, Z^\beta_{G,J}, \tag{7.7}$$

because $\sum_\tau \exp(\beta\sum_{xy} J_{xy}\tau_x\tau_y) = Z^\beta_{G,J}$. By (7.1) and (7.2), with $|J_e| = 1$, $\varphi_{G,J}(\omega) = e^{\beta|E|}p_e^{|\omega|}(1-p_e)^{|E|-|\omega|}N_J(\omega)/Z^\beta_{G,J}$, where $N_J(\omega)$ is the number of spin configurations satisfying every edge of ω. Both $N_J(\omega)$ and $Z^\beta_{G,J}$ are gauge invariant, so (7.7) applies and the partition function cancels:

$$\mathbb{E}\,\varphi_{G,J}(\omega) = 2^{-|V|}(2\cosh\beta)^{-|E|}e^{\beta|E|}p_e^{|\omega|}(1-p_e)^{|E|-|\omega|}\sum_J N_J(\omega).$$

Exchanging the sums, $\sum_J N_J(\omega) = \sum_\sigma \#\{J : J_e\sigma_x\sigma_y = 1\ \forall e \in \omega\} = 2^{|V|}2^{|E|-|\omega|}$. With $p_e = 1 - e^{-2\beta}$ one has $e^{\beta}p_e/(2\cosh\beta) = \tanh\beta$ and $2e^{\beta}(1-p_e)/(2\cosh\beta) = 1-\tanh\beta$, which is the claim. □

The cancellation of $Z_{G,J}$ against the disorder weight, itself a Gibbs weight, is the hallmark of the Nishimori line. An immediate consequence bounds correlations by independent percolation.

Corollary 7.8 (cf. Gandolfi [99], Mazza [158])**.** *In the setting of Theorem 7.7,* $\mathbb{E}|\langle\sigma_x\sigma_y\rangle_J| \le \mathbb{P}_{2p-1}(x \leftrightarrow y)$*, where* $\mathbb{P}_{2p-1}$ *is Bernoulli*$(2p-1)$ *bond percolation on* G*. In particular, for free boundary conditions on boxes* $\Lambda_L \subset \mathbb{Z}^d$ *with* $2p-1 < p_c^{\mathrm{bond}}(\mathbb{Z}^d)$*,* $\sup_L \mathbb{E}|\langle\sigma_0\sigma_x\rangle_{\Lambda_L,J}| \to 0$ *as* $|x| \to \infty$*.*

Proof. Average $|\langle\sigma_x\sigma_y\rangle_J| \le \varphi_{G,J}(x \leftrightarrow y)$ (see (7.4)) over J and use Theorem 7.7. On Λ_L, $\mathbb{P}_{2p-1}(0 \leftrightarrow x)$ is at most the probability that 0 is joined to distance $|x|$ in Bernoulli$(2p-1)$ percolation on $\mathbb{Z}^d$, which tends to zero below p_c. □

On $\mathbb{Z}^2$ the corollary covers $p < 3/4$ on the Nishimori line, and Bernoulli$(2p-1)$ percolates for $p > 3/4$ [99, p. 430]. On the binary tree the Nishimori-line ferromagnetic transition is at $\tanh\beta = 1/\sqrt{2}$, while the averaged FK bonds percolate from $\tanh\beta = 1/2$ on [99, Cor. 7]; in between, FK clusters percolate and the magnetization vanishes, so FK percolation plays a "marginal role in determining phase transitions" [99, abstract].

Route warning 7.9 (FK percolation is not the order signal)*. Tempting step:* prove spin-glass order by proving percolation of frustrated FK clusters, as one would for a ferromagnet. *Status: fails, with proofs.* (i) In $\mathbb{Z}^2$, FK clusters percolate at $\beta > \frac{1}{2}\log 17$ for *every* ±1 configuration (Theorem 7.6), whereas two-dimensional uniqueness at all positive temperature is the conjectured picture (Chapter 4). (ii) On the binary tree on the Nishimori line, the averaged FK bonds percolate on an interval of temperatures where there is no magnetic order [99, Cor. 7]. (iii) In general only $|\langle\sigma_x\sigma_y\rangle_J| \le \varphi(x \leftrightarrow y)$ holds, so percolation is necessary for long-range correlation but carries no sign; with a unique infinite cluster (Theorem 7.4) the signed terms of (7.4) may still cancel [170, p. 258]. Numerically (heuristic), FK clusters of the three-dimensional $\pm J$ model percolate well above the spin-glass transition [151, p. 3], and in two dimensions at a temperature that stays positive as the system grows [166].

7.6 Two replicas: double FK and the CMR colouring

Spin-glass order is a statement about two replicas (Chapter 1), whose natural local variable is the overlap sign $q_x = \sigma_x\tau_x$, invariant under the joint global flip and under gauge transformations. A representation adapted to order should therefore live on pairs of configurations. Throughout this section G is finite, $J \in \mathbb{R}^E$ is fixed, and σ, τ are independent with law $\mu^{\beta}_{G,J}$.

Definition 7.10 (Double FK)**.** Given (σ,τ), let ω^1 and ω^2 be independent, with ω^1 distributed as the ES bonds of σ and ω^2 as those of τ (given the spins, as after (7.2)). Then (σ,ω^1) and (τ,ω^2) are independent ES pairs. The *double FK* (or two-replica FK) configuration is $\omega^1 \cap \omega^2$, the edges occupied in both copies.

It was introduced for spin glasses by Newman and Stein [185, §4.1] and studied by Machta, Newman and Stein [151, 152]. A doubly occupied edge is satisfied in both replicas and, given the spins, occupied with probability p_e^2. The second representation occupies the same edges with a larger probability and also keeps the edges satisfied in exactly one replica.

Definition 7.11 (CMR colouring)**.** Given (σ,τ), colour the edges independently: an edge satisfied by both σ and τ is *blue* with probability $1 - e^{-4\beta|J_e|}$ and closed otherwise; an edge

satisfied by exactly one of them is *red* with probability $p_e = 1 - e^{-2\beta|J_e|}$ and closed otherwise; an edge satisfied by neither is closed. *Grey* edges are those that are blue or red. Blue, red and grey clusters are the connected components of these edge sets.

This is the two-replica representation of Chayes, Machta and Redner [55], introduced for ferromagnets in a field; its spin-glass form is [151, (7)–(13)] (Fig. 7.1(e)). The colouring encodes the overlap directly.

Proposition 7.12 (Structure of the CMR colouring)**.** *Let G be finite, $J \in \mathbb{R}^E$, $\beta \geq 0$, and let (σ, τ, c) be distributed as above, with c the colouring; write $\mathbf{P}_J$ for this law.*

(a) $q_x = q_y$ across every blue edge and $q_x = -q_y$ across every red edge. Hence every grey cycle has an even number of red edges, every blue cluster lies in one overlap sector $\{q = +1\}$ or $\{q = -1\}$, and every all-blue cycle C has $\prod_{e \in C} \operatorname{sgn} J_e = +1$.

(b) Summing out the colours gives the product measure $\mu^{\beta}_{G,J} \otimes \mu^{\beta}_{G,J}$.

(c) Fix a colouring c and a grey cluster K. Flipping τ on K is a bijection of $\{(\sigma, \tau) : (\sigma, \tau, c)$ has positive weight$\}$ that preserves the joint weight.

(d) For all $x, y \in V$,

$$\langle \sigma_x \sigma_y \rangle_J^2 = \mathbf{P}_J(x \overset{\text{grey}}{\leftrightarrow} y, \text{ even red parity}) - \mathbf{P}_J(x \overset{\text{grey}}{\leftrightarrow} y, \text{ odd red parity}). \tag{7.8}$$

Proof. (a) On a blue edge both replicas satisfy e, so $\sigma_x\sigma_y = \tau_x\tau_y = \operatorname{sgn} J_e$ and $q_xq_y = 1$; on a red edge exactly one does, so $\sigma_x\sigma_y = -\tau_x\tau_y$ and $q_xq_y = -1$. Multiplying q_xq_y around a cycle gives the parity statement, and multiplying $\sigma_x\sigma_y = \operatorname{sgn} J_e$ around an all-blue cycle gives the last one. (b) Write $r_e = e^{-2\beta|J_e|}$. Relative to $e^{2\beta|J_e|}$ the two-replica Boltzmann weight of an edge is 1, r_e or r_e^2 according as it is satisfied in two, one or no replicas. The colour kernel splits these as $1 = (1 - r_e^2) + r_e^2$ (blue or closed), $r_e = r_e(1 - r_e) + r_e^2$ (red or closed) and r_e^2 (closed); summing over the colour returns the Boltzmann weight. (c) The flip leaves $\sigma_x\sigma_y$ and $\tau_x\tau_y$ unchanged for edges inside K or outside K. Each edge with exactly one endpoint in K is closed, since K is a grey cluster; the flip changes its satisfaction pattern (two↔one, or one↔none), and in every case its closed weight is r_e^2. (d) If x and y lie in different grey clusters, the flip of (c) on the cluster of x changes q_x but not q_y, so $\mathbb{E}[q_xq_y \mid c] = 0$. If they lie in one grey cluster, $q_xq_y = (-1)^{\#\text{red}}$ along any grey path, by (a). Finally $\mathbb{E}_J[q_xq_y] = \langle\sigma_x\sigma_y\rangle_J\langle\tau_x\tau_y\rangle_J = \langle\sigma_x\sigma_y\rangle_J^2$. □

Identity (7.8) is [151, (16)]. Compared with (7.4), the sign has moved from the couplings to the red edges, where the two replicas disagree about an edge: the sign problem has become a question of two-replica geometry. Blue edges are the natural carriers of overlap order, and the conjectured order signal is *two* infinite blue clusters, one in each overlap sector, of unequal densities (Chapter 8). The two representations are realised on one space as follows.

Observation 7.13 (Union coupling). In the setting of Definition 7.10, colour e blue if $e \in \omega^1 \cup \omega^2$ and both replicas satisfy e, red if $e \in \omega^1 \cup \omega^2$ and exactly one replica satisfies e, and closed otherwise. Then (J, σ, τ, c) has the law of the CMR colouring. Consequently

$$\text{grey} = \omega^1 \cup \omega^2, \qquad \omega^1 \cap \omega^2 \subseteq \text{blue}.$$

Proof. Given (σ,τ), the edges are independent in both constructions. On a doubly satisfied edge, $\mathbb{P}(e\in\omega^1\cup\omega^2)=1-(1-p_e)^2=1-e^{-4\beta|J_e|}$. On an edge satisfied by exactly one replica, only that replica's bond can be open, with probability p_e. On an edge satisfied by neither, both bonds are closed. □

The occupation probabilities are compared in [151, p. 6], and [55] contrasts the colouring with two independent Swendsen–Wang chains. Two consequences are immediate: double-FK clusters refine CMR blue clusters, so double-FK percolation implies blue percolation; and every blue cluster lies inside a connected component of $\{q=+1\}$ or of $\{q=-1\}$, a cluster of Houdayer's algorithm (Section 7.10).

On the complete graph the blue clusters are understood. In the $\pm J$ SK model with couplings $\pm N^{-1/2}$, let $D_k(N,\beta)$ be the density of the k-th largest blue cluster.

Theorem 7.14 (Machta–Newman–Stein [152, Thms. 1 and 4])**.** *For every $0<\beta<\infty$, in probability as $N\to\infty$,*

$$D_1-\tfrac12(1+|Q|)\to 0,\qquad D_2-\tfrac12(1-|Q|)\to 0,\qquad D_3\to 0,$$

where Q is the site overlap of the two replicas. Consequently, for $0<\beta\le 1$ there are exactly two giant blue clusters, each of asymptotic density $\frac12$. For $\beta>1$ there are one or two giant blue clusters whose densities add to one, and under the properties of the SK overlap distribution labelled P1 and P2 in [152] there is a unique one of maximal density in $(\frac12,1]$ and a second one of positive smaller density.

The paramagnetic half is the decisive caution for the blue route: two equal giant blue clusters exist throughout $0<\beta\le 1$, where $Q\to 0$. Blue percolation is a geometric transition far above the spin-glass transition at $\beta=1$; what signals order is the *difference* of the two densities, a structure Machta, Newman and Stein also found numerically in three dimensions.

Physics prediction 7.15 (Percolation signature of the spin-glass transition; Machta–Newman–Stein [151])**.** For the short-range EA model in dimensions where it orders, blue clusters percolate already above the spin-glass temperature, as two infinite clusters of equal density, one in each overlap sector; the spin-glass transition is the temperature below which their densities differ.

Monte Carlo studies support this in three dimensions [151, 167]; in two dimensions the blue-cluster thresholds move toward zero temperature with system size [166], as expected if $T_c=0$. These are heuristic findings.

7.7 Signed FK percolation

Gandolfi, Newman and Stein refine FK connectivity for ±1 couplings differently [104]: they keep the sign information that it throws away. In the notation of Definition 7.1, two sites of one FK cluster are *positively* connected if $\operatorname{sgn}_J(x,y;\omega)=+1$ and *negatively* connected if it is -1. Positive connection is an equivalence relation on each cluster, with at most two classes, and two negative connections in series give a positive one.

Definition 7.16 (Signed companions; Gandolfi–Newman–Stein [104, §§1, 4])**.** Let $J\in\{-1,1\}^{E_d}$ and let ω be J-frustration-free with at most one infinite cluster C_∞. The *signed ferromagnetic companions* C_∞^A, C_∞^B are the two classes of positive connection in C_∞ (one may be

empty). For two replicas (ω^1, ω^2), sites $x, y \in C_\infty(1) \cap C_\infty(2)$ are *coherently* connected if $\operatorname{sgn}_J(x, y; \omega^1) \operatorname{sgn}_J(x, y; \omega^2) = +1$; the two coherence classes of $C_\infty(1) \cap C_\infty(2)$ are the *paired coherent companions*. For a van Hove sequence $\Gamma = (\Lambda_n)$ write $d_\Gamma(C) = \lim_n |C \cap \Lambda_n|/|\Lambda_n|$ when the limit exists.

By (7.3), the spins of a replica are constant on each companion and opposite on the two, while finite clusters get independent fair signs; so a magnetization must be a density difference of the two companions, and an overlap one of the two coherent companions. The theorems make this exact at fixed couplings, for extremal states.

Theorem 7.17 (Gandolfi–Newman–Stein [104, Thm. 5.7])**.** *Let $d \geq 1$, $\beta \geq 0$, $p = 1 - e^{-2\beta}$, and fix $J \in \{-1, 1\}^{E_d}$. Let μ be an extremal Gibbs state, π its Edwards–Sokal lift and ν its FK marginal, and assume $\pi(\text{at most one infinite FK cluster}) = 1$. Fix a van Hove sequence Γ and assume that $d_\Gamma(C_\infty^A)$ and $d_\Gamma(C_\infty^B)$ exist ν-almost surely when the infinite cluster exists (both set to zero otherwise). Then $m_\Gamma = \lim_n |\Lambda_n|^{-1} \sum_{x \in \Lambda_n} \mu(\sigma_x)$ exists and*

$$|m_\Gamma| = \mathbb{E}_\nu |d_\Gamma(C_\infty^A) - d_\Gamma(C_\infty^B)|.$$

Theorem 7.18 (Gandolfi–Newman–Stein [104, Thm. 8.5])**.** *In the setting of Theorem 7.17, let $\mu^{(1)}, \mu^{(2)}$ be extremal Gibbs states whose Edwards–Sokal lifts each have at most one infinite FK cluster almost surely, and let $\nu^{\otimes 2}$ be the product of their FK marginals. Fix Γ and assume that the densities of the two paired coherent companions exist $\nu^{\otimes 2}$-almost surely; let $D_\Gamma \geq 0$ be their difference (larger minus smaller), with $D_\Gamma = 0$ when there is no infinite paired cluster. Then $q_\Gamma = \lim_n (\mu^{(1)} \otimes \mu^{(2)})(|\Lambda_n|^{-1} \sum_{x \in \Lambda_n} \sigma_x \tau_x)$ exists, and $q_\Gamma = \varepsilon_J \, \mathbb{E}_{\nu^{\otimes 2}}[D_\Gamma]$ for some $\varepsilon_J \in \{-1, 1\}$. In particular $q_\Gamma \neq 0$ if and only if $\nu^{\otimes 2}(D_\Gamma > 0) > 0$.*

For stationary joint laws of ± 1 couplings and two-replica Edwards–Sokal Gibbs states in which each replica has at most one infinite FK cluster, $\mathbb{E}|Q|$, with Q the spatial site overlap of Chapter 1, equals the expected coherent-companion density difference [104, Cor. 8.9]. These exact recodings turn "order" into "the companions are distinguishable by density", the right geometric statement but no easier to prove. Through Observation 7.13 the companions can be compared with the CMR colouring.

Observation 7.19 (Coherent companions inside the grey cluster). In the union coupling, let x, y be joined in both ω^1 and ω^2. Then

$$\operatorname{sgn}_J(x, y; \omega^1) \operatorname{sgn}_J(x, y; \omega^2) = q_x q_y = (-1)^{\#\text{red edges on any grey path from } x \text{ to } y}.$$

Hence $C_\infty(1) \cap C_\infty(2)$ lies in one grey cluster, and its two paired coherent companions are its intersections with the two red-parity classes of that grey cluster.

Proof. If x, y are joined in ω^1, then $\sigma_x \sigma_y = \operatorname{sgn}_J(x, y; \omega^1)$ by (7.3), and similarly for τ; so the product of the two parities is $q_x q_y$. Paths in ω^1 are grey paths, since grey $= \omega^1 \cup \omega^2$, and $q_x q_y$ equals the red parity of any grey path by Proposition 7.12(a). □

So double FK, CMR and signed FK read the same pair (ω^1, ω^2) with the spins through its intersection, its coloured union and its parities; none supplies an inequality that the others lack.

7.8 Two copies and disagreement percolation

Van den Berg's disagreement percolation says that if two independent samples from two Gibbs measures disagree only on finite clusters, the measures coincide (as recalled in [100, §4.1]). Gandolfi refined it [100]: conditioning two independent copies with the same boundary condition on the sum of their spins fixes the overlap region $\{q = +1\}$ and leaves a symmetric Gibbs field on $\{q = -1\}$, whose random-cluster representation has "active" bonds. The covariance of two local events is then bounded by the averaged probability that active bonds join them [100, Thm. 3.1], and such decay gives extremality of a limit state.

For the $\pm J$ EA model, the active bonds in the non-overlap region are CMR blue bonds between sites with $q = -1$ [100, §4.4]. For a boundary condition η on the complement of a box Λ, let $\Pi^{\eta}_{\Lambda,J}$ be the joint law of two copies of $\mu^{\eta}_{\Lambda,J}$ and their CMR blue bonds, and let $\mathcal{A}$ be the event that the blue bonds inside $\{q = -1\}$ have an infinite cluster.

Theorem 7.20 (Gandolfi [100, Cor. 3.3–3.4 and Thm. 4.3])**.** *Let $d \geq 1$, let J be ± 1-valued, $\beta > 0$. Assume that for almost every J and for every deterministic sequence of boundary conditions $(\eta_{\Lambda^c})_{\Lambda}$, every subsequential limit Π_J of $\Pi^{\eta_{\Lambda^c}}_{\Lambda,J}$ satisfies $\Pi_J(\mathcal{A}) = 0$. Then:*

(i) for almost every J, every limit point of $\mu^{\eta_{\Lambda^c}}_{\Lambda,J}$ along a deterministic sequence of boundary conditions is an extremal Gibbs state;

(ii) if, in addition, every Gibbs state for J is such a limit point, the Gibbs state for J is unique.

Part (ii) holds because a convex set all of whose elements are extremal is a single point. The theorem is a criterion for *non*-order, attractive only where blue percolation in the disagreement region is expected to fail, as Gandolfi suggested for two dimensions [100, §1]; in three dimensions blue clusters percolate above the spin-glass temperature in simulations [151, 167], and in high dimensions blue percolation in both overlap sectors is a theorem in periodic limits (Chapter 8). An infinite-volume form for two arbitrary Gibbs states, with uniqueness stated through pairs of extremal states, is in Observations 8.19 and 8.20.

Open problem 7.21 (Two-dimensional disagreement percolation; after Gandolfi [100, §1])**.** For the $\pm J$ EA model on $\mathbb{Z}^2$ and every $\beta < \infty$, show that, for almost every J and every deterministic sequence of boundary conditions, CMR blue bonds do not percolate inside the disagreement region of two copies.

7.9 Frustration-adjusted random currents

The random-current representation, the other great geometric tool for ferromagnets, also has a frustrated form. Write a current as $n : E \to \mathbb{Z}_+$, with sources ∂n the vertices of odd total current and weight $w(n) = \prod_e (\beta|J_e|)^{n_e}/n_e!$ for the ferromagnet $|J|$. Call n frustration-free if every cycle supported on $\{n_e \neq 0\}$ contains an even number of negative couplings.

Theorem 7.22 (Aizenman [2, Thm. 7.5])**.** *For a finite Ising system with couplings $J \in \mathbb{R}^E$ of either sign,*

$$\frac{Z^{\beta}_{G,J}}{Z^{\beta}_{G,|J|}} = \mathbb{P}^{\varnothing,\varnothing}\big(n_1 + n_2 \text{ is frustration-free}\big),$$

where n_1, n_2 are independent sourceless random currents of the ferromagnet $|J|$.

As in Remark 7.2, a frustrated system is a ferromagnet conditioned on a global, nonmonotone event. The identity is exact and positive, useful for systems with a little frustration; for the symmetric EA model no comparison inequality is known that would exploit it.

7.10 Cluster moves for simulation

Most two-replica cluster definitions were invented for Monte Carlo simulation, and their efficiency claims are empirical. *Swendsen–Wang* alternates the two conditional laws of the ES coupling (Section 7.2); it is inefficient for spin glasses because FK clusters percolate well above the glass transition and then contain most sites [151, p. 3]. *Replica Monte Carlo* [229] builds clusters on $\sigma_x\tau_x$ for replicas at different temperatures, the ancestor of all two-replica moves [166, p. 7]. *Houdayer moves* [120] flip a connected component of $\{q = -1\}$ in both replicas, preserving $H_J(\sigma)+H_J(\tau)$ and the overlap field; both overlap sectors percolate "as soon as the site percolation threshold is less than 0.5", for instance in $d = 3$ [120, p. 3], and flipping a system-spanning cluster only exchanges the replicas; Zhu, Ochoa and Katzgraber use these isoenergetic moves at low temperatures in any dimension, where frustration keeps the clusters from spanning [247]. *Jörg's algorithm* [128] occupies doubly satisfied edges with probability $1 - e^{-4\beta|J_e|}$ and flips a resulting cluster in both replicas, so its clusters are CMR blue clusters [166, (13) and p. 7]; adding red bonds gives the full CMR move [151, 223]. Keiser and Machta report gains from two-replica cluster moves with population annealing for the frustrated, non-random two-dimensional ANNNI model [136].

Flipping both replicas on a blue cluster preserves the two-replica CMR weight: the edges leaving a blue cluster are not blue, and the flip exchanges "satisfied in both" with "satisfied in neither" and the two "satisfied in one" states, all of which carry the same weight off blue. This flip is the heart of the crossing lemma of Chapter 8. Swendsen–Wang conserves nothing, Houdayer the overlap field and the total energy of the pair, Jörg the overlap field, the full CMR move the overlap up to grey-cluster flips. A move that conserves the overlap cannot by itself sample overlap fluctuations, which is why these moves are always paired with others.

7.11 How the representations fit together

The representations of this chapter can be compared in two ways: by the recipe that builds each of them on a single edge, and by realizing all of them on one probability space. The first explains why there are so many and what each one buys; the second shows what their percolation statements give.

7.11.1 One recipe: allowance matrices

On an edge e, list the local configurations c of the spins at its ends, with their Boltzmann weights $W(c)$ normalized so that a satisfied edge has weight 1. With $r_e = e^{-2\beta|J_e|}$, one replica has $c \in \{\mathrm{sat}, \mathrm{unsat}\}$ with weights $1, r_e$; two replicas have the satisfaction pattern $c \in \{++, +-, -+, --\}$ (first sign for σ, second for τ) with weights $1, r_e, r_e, r_e^2$. A graphical representation is an *admissible pair* (A, λ): a 0/1 *allowance matrix* $A[b, c]$, whose rows are bond types b, and weights $\lambda_b \geq 0$ that do not depend on the spins, such that

$$W(c) = \sum_b \lambda_b \, A[b, c] \qquad \text{for every column } c. \tag{7.9}$$

(a) frustrated FK, one replica

	sat	unsat	λ
open	1	0	$1-r$
closed	1	1	r

(b) double FK

	$++$	$+-$	$-+$	$--$	λ
$(1,1)$	1	0	0	0	$(1-r)^2$
$(1,0)$	1	1	0	0	$r(1-r)$
$(0,1)$	1	0	1	0	$r(1-r)$
$(0,0)$	1	1	1	1	r^2

(c) CMR

	$++$	$+-$	$-+$	$--$	λ
blue	1	0	0	0	$1-r^2$
red	0	1	1	0	$r(1-r)$
closed	1	1	1	1	r^2

(d) Houdayer: one row per pattern

	$++$	$+-$	$-+$	$--$	λ
$++$	1	0	0	0	1
$+-$	0	1	0	0	r
$-+$	0	0	1	0	r
$--$	0	0	0	1	r^2

Table 7.1: Allowance matrices on one edge, with $r = e^{-2\beta|J_e|}$: rows are bond types, columns the satisfaction patterns, and every column sums to its Boltzmann weight 1, r, r or r^2. In (b) the row (i,j) means "open in the first replica iff $i=1$, in the second iff $j=1$". Signed FK uses the rows of (a) with a sign label, Gandolfi's active bonds and Jörg's clusters the blue bonds of (c) on $\{q=-1\}$.

In probabilistic terms, an admissible pair is a *coupling*: one probability space that carries the spins and a bond configuration, whose spin marginal is the Gibbs measure (this is (7.9)), and whose bonds are independent given the spins, each column's weight being spread over the bond types it allows. The edge representations of this chapter (all but the random currents of Section 7.9) differ only in the coupling they choose; Section 7.11.2 realizes them on one space. The joint weight of spins and bonds is $\prod_e \lambda_{b_e} A[b_e, c_e]$, and three facts follow from (7.9) alone. Given the spins, the bonds are independent, with $\mathbb{P}(b \mid c) = \lambda_b A[b,c]/W(c)$. Given the bonds, the spins are uniform over the configurations that every edge allows. The bond marginal is $\prod_e \lambda_{b_e}$ times the number of such configurations. So *frustration enters only through that count*, as consistency conditions around cycles, and which conditions appear is decided by the zeros of A. This is the general framework of cluster algorithms of Kandel and Domany [130], introduced for frustrated models in [131]; the term *allowance matrix*, and rows that freeze satisfied edges, are from Coddington and Han [58], and Chayes and Machta develop the joint spin–bond measures it defines [53, 54]. Table 7.1 lists the matrices behind this chapter.

What a representation can do is decided by which moves of finite clusters are free. Let g act on the local configuration at a site: reverse one replica, reverse both, or exchange the two replicas.

Observation 7.23 (Free cluster moves). Fix the bonds and a configuration of positive weight, and let C be a finite set of sites. Applying g at every site of C, with the bonds kept, preserves the joint weight if and only if every edge stays allowed: $A[b_e, c_e'] = 1$ for every edge e, where c_e' is its pattern after the move. The move then exchanges configurations of equal weight.

Proof. The joint weight is $\prod_e \lambda_{b_e} A[b_e, c_e]$ and the move keeps every b_e, so the weight changes only through the factors $A[b_e, c_e] \in \{0,1\}$. □

Only edges with an endpoint in C change pattern, so the condition is on interior and border rows. For FK, open edges inside a cluster stay satisfied when the cluster is reversed, and the border rows are closed, which is allowed in both columns: this is the Swendsen–Wang move. For

CMR, reversing both replicas on a blue cluster keeps its blue edges in $++$ and sends a border pattern $+-$ to $-+$ and $++$ to $--$, which closed and red rows allow; this is the blue move of the crossing lemma and of Observation 8.19. Reversing τ alone is free on a grey cluster but not on a blue cluster with a red border edge, which would become $++$ or $--$ (Proposition 7.12(c)). Houdayer's matrix has one row per pattern, so its bonds carry no information beyond the spins and the observation applies with the rows relabelled: reversing both replicas on a component of $\{q = -1\}$ keeps the Boltzmann weight because its border edges are mixed, the move exchanges $+-$ and $-+$, and the two mixed columns have equal weight r. This is the isoenergetic move of Section 7.10.

So the zeros of A decide which finite-cluster symmetries exist, those symmetries decide which quantities cancel over finite clusters, and that decides what infinite clusters can tell apart. Rows with few allowed columns are *stiff*: they give rigid structure (labels, parity classes, at most two infinite blue clusters) but fewer free moves. Rows with many are *soft*: they give more free moves and exact cancellation, but less structure. Double FK lets each replica reverse its own finite clusters, a full $\mathbb{Z}_2 \times \mathbb{Z}_2$ symmetry, so finite clusters cancel exactly and the overlap equals a companion density difference (Theorem 7.18); it has no overlap-labelled cluster structure. CMR is double FK relabelled column by column (Observation 7.13): in the $++$ column the rows $(1,1)$, $(1,0)$ and $(0,1)$ become blue, with weight $(1-r)^2 + 2r(1-r) = 1 - r^2$, and in the mixed columns the replica's own open row becomes red. Red welds blue clusters into grey clusters with rigid overlap parity, which is behind the at-most-two theorem of Chapter 8, and blue clusters keep a free swap. The cost is that inside an infinite grey cluster a finite blue cluster can only be reversed in both replicas, which keeps q, so its overlap does not cancel; this is why the spatial overlap and $D_+ - D_-$ can differ (Chapter 8). No edge-local recipe does better, in the following sense. Call a two-replica row *connecting* if it is allowed only on columns among $++$ and $--$ (so its clusters have constant overlap), *neutral* if it is allowed on all four columns (so it is invariant under reversing either replica at one endpoint), and a *defect* row otherwise; the defect mass is $\sum_c \sum_{b \text{ defect}} \lambda_b A[b,c]$.

Observation 7.24 (CMR is defect-optimal on one edge). For every admissible pair on one edge of two replicas, the defect mass is at least $2r(1-r)$. Equality holds exactly when the connecting mass is $1 - r^2$ on $++$ and 0 on $--$, the neutral mass is r^2, and the defect mass is $r(1-r)$ on each mixed column: the column masses of blue, closed and red in CMR.

Proof. Let n be the total neutral weight, which is the same in every column. The columns $+-$ and $-+$ carry no connecting rows, so their defect masses are $r - n$ each; the column $--$ gives $n \le r^2$. Hence the defect mass is at least $2(r-n) \ge 2r(1-r)$, with equality only if $n = r^2$, which leaves no connecting or defect mass on $--$, and no defect mass on $++$, whose remaining weight $1 - r^2$ is then connecting. □

Blue and red cannot be merged into one row, because their weights differ. Admissible pairs are far from unique: the rows "$++$ or $--$" with weight r^2, "$+-$ or $-+$" with weight r and "$++$" with weight $1 - r^2$ also satisfy (7.9), but the first row joins doubly unsatisfied edges, so its clusters say nothing about satisfaction. We know of no choice of A that creates monotonicity: in every example above the compatibility count is a non-monotone constraint.

The recipe also works on blocks larger than an edge; this is a tangent, taken up in Section 7.12.

Table 7.2 sorts the representations of this chapter by these two questions.

Representation	What it tells apart	Where the frustration is hidden
frustrated FK (a)	uniqueness, sign-blind, by domination by the $\|J\|$-ferromagnet (Section 7.2)	the compatibility count of the bond marginal
signed FK (a) with labels	magnetization, as a companion density difference (Theorem 7.17)	sign labels computed from J along open paths
double FK (b)	the overlap, exactly, by finite-cluster cancellation (Theorem 7.18)	the compatibility count, in each replica
CMR (c)	overlap sectors, and a criterion for two states to coincide (Proposition 7.12, Observation 8.19)	the compatibility count, through blue and red cycle conditions
Houdayer (d)	disagreement clusters; free isoenergetic moves (Section 7.10)	nowhere: the bonds are the patterns themselves
FK on the Nishimori line	Bernoulli percolation bounds, since the averaged bonds are independent (Theorem 7.7)	the disorder average, which makes the averaged bonds Bernoulli($\tanh\beta$)
GKN averaging	at most one infinite cluster (Theorem 7.4)	the disorder average, which restores insertion tolerance
KBD plaquettes, fully frustrated lattice	cluster moves for fully frustrated models (Section 7.12)	inside the weights of each black plaquette of a checkerboard

Table 7.2: Choosing a representation means answering two questions: what should the clusters tell apart, and where should the frustration be hidden.

Table 7.3: What percolation, or its absence, gives in each representation.

Representation	What percolation or its absence gives	What it does not give
frustrated FK	no percolation in the dominating ferromagnet $\Rightarrow$ uniqueness (Newman, Chapter 1); at most one infinite cluster (Theorem 7.4)	order: FK percolates at low temperature in $\mathbb{Z}^2$ for every J (Theorem 7.6), and on trees without magnetic order (Route warning 7.9)
double FK	contained in blue (Observation 7.13)	order without a density difference: in SK there is no doubly occupied giant for $\beta \leq 1$, and for $\beta > 1$, under properties of the SK overlap, two of different densities [152, Thm. 5]
CMR blue	unequal densities $\Rightarrow$ distinct spin-flip-related Gibbs states (Chapter 8); no blue percolation in the disagreement region $\Rightarrow$ extremality (Theorem 7.20)	order from percolation alone: two equal giants in the SK paramagnet (Theorem 7.14); percolation with finite χ_{SG} on $\mathbb{Z}^d$ (Chapter 8)
signed FK	companion density difference $\Leftrightarrow$ magnetization or overlap (Theorems 7.17 and 7.18)	a way to prove the difference: the criteria are equivalences

7.11.2 One probability space

Figure 7.3 collects the relations of this chapter over one probability space (the couplings, two replicas and their ES bonds); its top arrow has a random-current analogue (Section 7.9). Table 7.3 records what percolation proves in each representation. Throughout, *absence* of percolation, suitably uniform, implies uniqueness or decay, while *presence* needs an extra asymmetry (a sign that does not cancel, or a density difference) before it says anything about order.

Route warning 7.25 (Two-replica percolation is not order either). *Tempting step:* prove order by proving percolation of blue or doubly occupied clusters. *Status: fails, with proofs, in the*

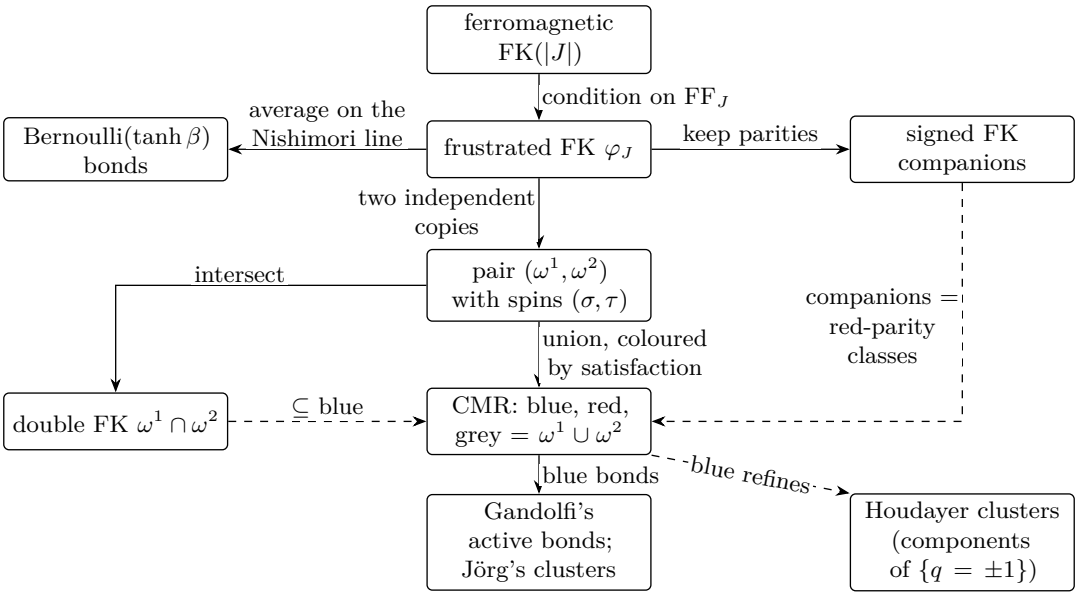


Figure 7.3: The cluster representations of this chapter as functions of one probability space. Solid arrows are constructions, dashed arrows inclusions or identifications. All two-replica objects are read off one pair of FK configurations, so their percolation statements can be compared on one space; the relations are bookkeeping, not inequalities.

mean-field bridge and on $\mathbb{Z}^d$. In the SK model two equal-density giant blue clusters exist for every $0 < \beta \leq 1$, where the overlap vanishes (Theorem 7.14); on $\mathbb{Z}^d$ in high dimension, both overlap sectors contain infinite blue clusters of exactly equal density while χ_{SG} is finite (Chapter 8). What must be proved is a density imbalance, and the criteria that recode order as an imbalance (Theorem 7.18, (7.8)) are equivalences that carry no monotonicity or comparison inequality. *Credit:* the equal-giants phenomenon and the imbalance signature are Machta, Newman and Stein's [151, 152].

7.12 A tangent: cluster representations beyond edges

This section is a side quest. Nothing in it bears on order or non-order; it records what the recipe of Section 7.11.1 gives when the blocks are larger than an edge, which matters mainly for cluster algorithms.

Any decomposition of the Hamiltonian into local blocks can replace the edges, with the local configurations of a block as columns. Call a row *freeze* F, for a set F of edges of the block, if it is allowed exactly on the configurations that satisfy every edge of F. A matrix of freeze rows always gives a cluster algorithm: in finite volume, draw the bonds given the spins, then reverse each cluster of frozen edges independently with probability one half. Reversing a union of clusters keeps every frozen edge satisfied, so the move is free (Observation 7.23), and the two steps alternate the two conditional laws of the coupling, which gives detailed balance; a positive weight on the empty row of every block lets one step reach every configuration. What matters is how large the clusters are. Following Kandel, Ben-Av and Domany, one makes the weight of the empty row as large as possible; it is at most $\min_c W(c)$, and one asks for $\lambda_\emptyset = \min_c W(c)$.

Plaquettes. Kandel, Ben-Av and Domany use plaquettes [131]. Colour the faces of the square lattice as a checkerboard, so that every edge lies in exactly one black plaquette. In the fully frustrated model (unit magnitudes, every plaquette frustrated) the sum S_p of the four terms $J_{xy}\sigma_x\sigma_y$ of a black plaquette p takes only the values ± 2, since the number of unsatisfied edges is odd, and

$$e^{\beta S_p} = e^{-2\beta} + \big(e^{2\beta} - e^{-2\beta}\big)\big(I_{p,H} + I_{p,V}\big),$$

where $I_{p,H}$ and $I_{p,V}$ indicate that both horizontal, respectively both vertical, edges of p are satisfied. So a black plaquette has three freeze rows: empty, "freeze the horizontal pair" and "freeze the vertical pair". On a plaquette with three satisfied edges exactly one pair is satisfied, and it is frozen with probability $1 - e^{-4\beta}$; on a plaquette with one satisfied edge the plaquette is empty [94]. The frustration of each plaquette is absorbed into its own weights, and a frozen pair of opposite edges never closes a cycle inside the plaquette.

Cubes and hypercubes. The checkerboard has an analogue in every dimension.

Observation 7.26 (Edge-disjoint unit cubes). For $d \geq 2$, the unit cubes $w + [0,1]^d$ with $w \in \mathbb{Z}^d$ and $w_1 \equiv \cdots \equiv w_d \pmod 2$ contain every edge of $\mathbb{Z}^d$ exactly once and every vertex exactly twice.

Proof. A vertex v lies in $w + [0,1]^d$ iff $w_i \in \{v_i - 1, v_i\}$ for every i; for each of the two parities exactly one such w has all its coordinates of that parity. An edge $\{v, v + e_j\}$ lies in $w + [0,1]^d$ iff $w_j = v_j$ and $w_i \in \{v_i - 1, v_i\}$ for $i \neq j$; the parity is then that of v_j, and every other w_i is determined. □

For $d = 2$ these are the black plaquettes, and for $d = 3$ the decomposition that Hamze et al. use to plant ground states [114].

Observation 7.27 (Freeze matrices on the fully frustrated cube). Let every plaquette of the cubic lattice be frustrated, with unit magnitudes, so that all six faces of each cube of Observation 7.26 are frustrated. A cube then has 128 configurations up to a global flip, with between 3 and 9 unsatisfied edges, so $\min_c W(c) = r^9$.

(a) Admissible freeze matrices with $\lambda_\emptyset = r^9$ exist. One applies the plaquette rule on each face at coupling $J/2$. Another, whose frozen sets are forests, minimizes among all admissible freeze matrices with $\lambda_\emptyset = r^9$ the expected number of frozen edges of an isolated cube at every temperature; the minimum is $6(1-r)(1+r^2)/(1+r^3)$.

(b) In every admissible freeze matrix, every row allowed only on ground-state configurations of the cube freezes at least six edges. On an even torus every ground state is a ground state of each cube, so as $T \to 0$, on a ground state and with probability tending to one, every cube freezes at least six edges, at least $3/2$ per site, against $9/4$ for Swendsen–Wang and 1 for the plaquette rule in two dimensions.

(c) $\lambda_\emptyset = \min_c W(c)$ is attainable for each of the six classes of $\pm J$ cubes (by the number and position of frustrated faces).

Each statement is checked by an exact computation in rational arithmetic (polynomial identities in r with rational coefficients, and an exact dual certificate for the minimum in (a)) and, except for the face rule in (a), again by a second program written independently from the definitions; both programs are in [220].

The face rule of (a) works in every dimension, but for $d = 5$ its empty-row weight falls below $\min_c W(c)$.

Observation 7.28 (The face rule on hypercubes). On the cubes of Observation 7.26, split each coupling equally among the $d - 1$ two-dimensional faces that contain its edge, and apply the plaquette rule on every frustrated face (any admissible freeze matrix on the others). This gives an admissible freeze matrix in every d. If every two-dimensional face of the cube is frustrated, it attains $\lambda_\emptyset = \min_c W(c)$ exactly when some configuration has exactly one satisfied edge on every two-dimensional face; such a configuration exists for $d = 2, 3, 4$ and not for $d = 5$.

Proof sketch. A product of admissible block matrices is admissible for the sum of the block Hamiltonians. A frustrated face has at least one satisfied edge, and each edge lies in $d - 1$ faces, so a configuration has at most $3 \cdot \#\text{faces}/(d-1) = \frac{3}{4}|E|$ unsatisfied edges, while the face rule has $\lambda_\emptyset = r^{3|E|/4}$. The cases $d \le 5$ are settled by exhaustive search [220]. □

How much smaller the clusters are is a separate, empirical question. In short Monte Carlo runs on the cubic lattice (no error bars), the largest cluster of the forest rule of Observation 7.27(a) at $r = \frac{1}{2}$ holds about a tenth of an $L = 12$ torus, against about four fifths for Swendsen–Wang; at $r = 0.4$ and at $r \le 0.1$ the largest cluster of both rules holds most of the torus. These runs are not part of [220]. Coddington and Han report that their generalized cluster algorithm for the cubic fully frustrated model produces a very large cluster at the critical point and does no better than the Metropolis algorithm [58].

Remark 7.29 (Frustration is moved, not removed). For two replicas, blocks larger than an edge lower the defect of Observation 7.24 in the cases tested, without removing it. In two-replica linear programs of the same kind on the plaquettes of a checkerboard and on pairs of adjacent plaquettes (double precision; unpublished, and not part of [220]), at $r = 10^{-3}$ the optimal defect is about 10^{-5} of the weight of an unfrustrated plaquette and about one half of the weight of a frustrated one, against three quarters for CMR; on two adjacent plaquettes it is about 0 when both are frustrated and about one third when only one is. This suggests that the cost sits on frustration that a block cannot pair inside itself, so that a fixed translation-covariant tiling keeps a positive density of it. On one edge, allowing negative weights removes the defect but turns it into a sign problem. In every form tested, frustration is moved, into compatibility constraints, sign labels, the disorder average or uncancelled cluster borders, but not removed. This is heuristic. The single-replica plaquette rule above is unaffected, since it needs no overlap cancellation.

7.13 Summary

Every cluster representation has an exact frustrated version, with two-point identities (7.4) and (7.8) that express correlations through signed connectivity. Averaging a coupling makes the infinite FK cluster unique when it exists (Theorem 7.4), on the Nishimori line the averaged FK bonds are independent (Theorem 7.7), and the two-replica objects are functions of one pair of FK configurations (Observations 7.13 and 7.19). None supplies an inequality that forces the signs to add up (Route warnings 7.9 and 7.25), and Newman asks whether this can be seen on the lattice itself. Chapter 8 continues on $\mathbb{Z}^d$, and Chapter 14 on the Nishimori line.

Open problem 7.30 (FK percolation with uniqueness; Newman [170, p. 258])**.** Prove that for some $d \ge 2$, some symmetric coupling law and some β, the frustrated FK clusters percolate while the Gibbs state is unique for almost every J.

Chapter 8

Blue percolation and imbalance

Chapter 7 ended with a picture and a warning. The picture, due to Machta, Newman and Stein, is that spin-glass order shows up in the Chayes–Machta–Redner (CMR) colouring of two replicas as two infinite blue clusters, one in each overlap sector, of *unequal* density (Physics prediction 7.15). The warning is that percolation alone is not order: in the SK paramagnet the two giant blue clusters exist and have equal densities. This chapter follows the picture on $\mathbb{Z}^d$ as far as present proofs go. It splits the ordering statement into four parts:

(O1) both overlap sectors contain infinite blue clusters;
(O2) each sector contains at most one;
(O3) the two infinite-cluster densities differ (*imbalance*);
(O4) imbalance gives distinct Gibbs states related by the spin flip, and overlap order in a local channel.

Figure 8.1 shows how they fit and what is known. Part (O1) in high dimension, part (O2) and the implication (O3)⇒(O4) are theorems; (O3) is open, and it is the whole difficulty. We prove a crossing lemma that turns correlation into blue connection, and, by a two-replica cavity argument with no computer input, that both sectors percolate in high dimension; we then explain why the same regime has percolation *without* order, and why the technique cannot reach low dimensions even granting percolation there. A further section reads the blue count as a statement about the pair of pure states that two periodic replicas sample, which connects the route to the metastate classes of Chapter 3. This is our own route; its results carry their current status.

8.1 The setting: periodic joint limits and the two sectors

Fix a symmetric coupling law ν on $\mathbb{R}$, independent of the dimension, and write $J_e = \varepsilon_e W_e$ with fair signs ε_e independent of iid magnitudes $W_e \geq 0$. Assume

$$0 < m := \mathbb{E}W < \infty. \tag{8.1}$$

On a finite graph $G = (V, E)$ the *two-replica blue law* $\mathbf{P}_{G,\beta}$ is the joint law of (J, σ, τ, B): the couplings are iid with law ν; given J, the replicas σ, τ are independent with law $\mu^{\beta}_{G,J}$; and given (J, σ, τ) the blue indicators B_e are independent with

$$\mathbf{P}_{G,\beta}(B_e = 1 \mid J, \sigma, \tau) = (1 - e^{-4\beta|J_e|})\, \mathbf{1}\{J_e\sigma_x\sigma_y > 0,\ J_e\tau_x\tau_y > 0\}. \tag{8.2}$$

This is the blue part of the CMR colouring of Definition 7.11. The disorder is quenched: each replica is normalized by $Z_{G,J}$ before the couplings are averaged, so the joint law carries the factor

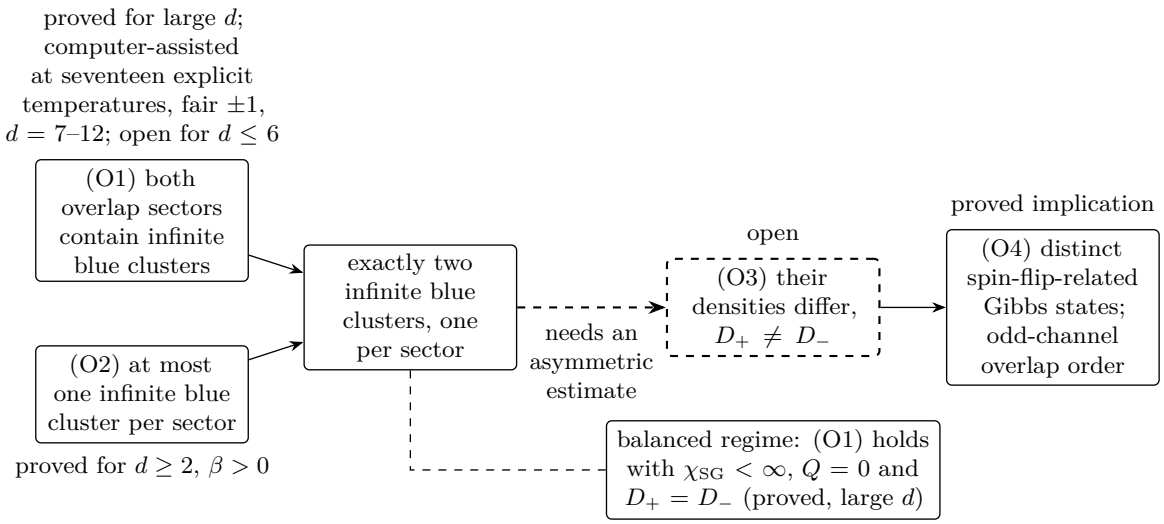
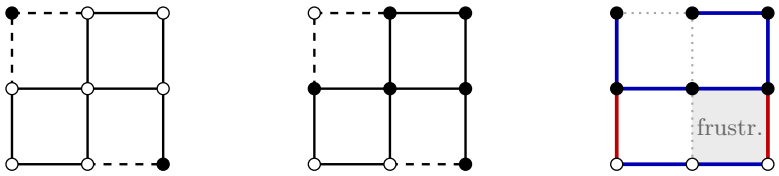


Figure 8.1: The blue-cluster route to order. Solid boxes are theorems, solid arrows proved implications, the dashed box the open step. The balanced regime shows that the left half can hold in a disordered system, so all the ordering content sits in the dashed arrow: an estimate that treats the two overlap sectors differently.



(a) couplings and σ (b) couplings and τ (c) CMR colouring on $q = \sigma\tau$

Figure 8.2: Two replicas on a 3×3 patch (couplings: positive solid, negative dashed; spins: open $+1$, filled -1). In (c) sites show $q = \sigma\tau$; activated edges are blue if satisfied in both replicas, red if in exactly one; dotted edges are closed. The two blue clusters lie in opposite sectors (bottom row $q = +1$) inside one grey cluster, joined by red edges: the configuration the at-most-two theorem allows for infinite clusters. The replicas violate different edges of the shaded frustrated plaquette.

$Z_{G,J}^{-2}$. We realize $B_e = \mathbf{1}\{U_e \le 1 - e^{-4\beta|J_e|}\}\mathbf{1}\{e \text{ doubly satisfied}\}$ with iid uniform activation variables U_e. The overlap sign $q_x = \sigma_x\tau_x$ splits the vertices into the two *overlap sectors* $\{q = \pm1\}$; a blue edge joins sites of equal overlap, so each blue cluster lies in one sector (Proposition 7.12; example in Fig. 8.2).

Infinite-volume statements are made for limits of tori $\mathbb{T}_L^d = (\mathbb{Z}/L\mathbb{Z})^d$, $L \ge 3$. Let $\mathcal{L}_{d,\beta}$ be the set of local accumulation points of $\mathbf{P}_{\mathbb{T}_L^d,\beta}$ as $L \to \infty$; it is nonempty (the coupling marginals are fixed and the other local spaces are finite) and its elements, the *selected periodic joint limits*, are translation invariant. They need not be products of two copies of one Gibbs state, but their spin marginal given J is Gibbs for two uncoupled replicas [214, §6.1]. Write $0 \overset{\mathrm{b}}{\leftrightarrow} \infty$ for the event that the blue cluster of the origin is infinite and $0 \overset{\mathrm{b}}{\leftrightarrow} \partial B_n$ for a blue connection from 0 to graph distance n inside the ball B_n. For $s = \pm1$, let D_s be the density of vertices of overlap s that lie in infinite blue clusters; it exists almost surely by stationarity. Reversing one replica, $\tau \mapsto -\tau$,

preserves every torus law, keeps the blue edges and exchanges the sectors, so it preserves every selected limit and exchanges D_+ and D_-. In particular $\mathbb{E}[D_+ - D_-] = 0$ always; *imbalance* means $\mathbb{E}(D_+ - D_-)^2 > 0$.

8.2 The crossing lemma: correlation forces blue connection

Before percolation, a short and useful fact: whatever the couplings, two strongly correlated spins are likely to be joined by blue edges. The proof is one flip: off blue connection, reversing both replicas on the blue cluster of v exchanges the aligned and anti-aligned configurations of (σ_v, τ_v) relative to (σ_u, τ_u) at equal weight.

Lemma 8.1 (Crossing lemma)**.** *Let G be a finite graph, $J \in \mathbb{R}^E$ and $\beta \geq 0$, with free boundary condition. Let σ, τ be independent with law $\mu^{\beta}_{G,J}$, let B be the blue indicators (8.2), and write $\mathbf{P}_J$ for the joint law of (σ, τ, B). For all $u \neq v$,*

$$\mathbf{P}_J(u \overset{\mathrm{b}}{\leftrightarrow} v) \;\geq\; |\langle \sigma_u \sigma_v \rangle_J|.$$

Proof. Let $w(\sigma) = \exp(\beta \sum_e J_e \sigma_x \sigma_y)$ and, for $s = \pm 1$, $Z(s) = \sum_{\sigma : \sigma_u = 1, \sigma_v = s} w(\sigma)$, so that $\langle \sigma_u \sigma_v \rangle_J = (Z(+) - Z(-))/(Z(+) + Z(-))$. The joint weight of (σ, τ, B) is $w(\sigma)w(\tau)$ times the product of the blue kernels. Let $\mathcal{D} = \{u \overset{\mathrm{b}}{\not\leftrightarrow} v\}$ and let N_s be the weight of $\mathcal{D} \cap \{\sigma_u = \tau_u = 1,\ \sigma_v = \tau_v = s\}$.

On $\mathcal{D}$, reverse both σ and τ on the blue cluster C of v, keeping B. Edges inside C and edges away from C keep their states. An edge with one endpoint in C is not blue; relative to $e^{2\beta|J_e|}$, its weight times the probability of being non-blue is $e^{-4\beta|J_e|}$ if it is satisfied in both replicas or in neither, and $e^{-2\beta|J_e|}$ if it is satisfied in exactly one. The flip exchanges "both" with "neither" and the two "exactly one" states, so it preserves the weight. It is an involution of $\mathcal{D}$ (the blue cluster of v does not change, and it does not contain u) that maps $\{\sigma_v = \tau_v = s\}$ to $\{\sigma_v = \tau_v = -s\}$ and fixes the spins at u. Hence $N_+ = N_-$, and since $N_s \leq Z(s)^2$, both are at most $\min(Z(+), Z(-))^2$.

By the global flip of each replica it suffices to count configurations with $\sigma_u = \tau_u = 1$. If $\sigma_v \neq \tau_v$ then $q_u \neq q_v$ and u, v are disconnected, which contributes $2Z(+)Z(-)$. Writing $\pi_\pm = Z(\pm)/(Z(+) + Z(-))$,

$$\mathbf{P}_J(\mathcal{D}) = \frac{2Z(+)Z(-) + N_+ + N_-}{(Z(+) + Z(-))^2} \leq 2\pi_+\pi_- + 2\min(\pi_+, \pi_-)^2 = 2\min(\pi_+, \pi_-) = 1 - |\langle \sigma_u \sigma_v \rangle_J|.$$

□

The flip is the "blue move" of Chayes, Machta and Redner's cluster algorithm [55]. On finite graphs the lemma gives a converse to the imbalance question: strong order, measured against a reference configuration, forces a giant blue cluster and an imbalance between the sectors.

Corollary 8.2 (Reference alignment forces a giant blue cluster and imbalance)**.** *Let G be a finite graph with random couplings J of any law, $\beta \geq 0$, free boundary condition, and the two-replica blue law. Fix $a \in V$ and a configuration $\eta = \eta(J) \in \{-1, 1\}^V$ with $\eta_a = 1$, and suppose that for some $\delta \in [0, \frac{1}{2})$,*

$$\mathbb{E}\, \mu^{\beta}_{G,J}(\sigma_v \neq \eta_v \sigma_a) \leq \delta \qquad \text{for every } v \in V.$$

Let $C_\pm$ be the largest blue clusters in the two sectors (size 0 if a sector is empty), $I = (|C_+| - |C_-|)/|V|$, and $Q = |V|^{-1} \sum_v q_v$. Then

$$\frac{\mathbb{E}|C_{\mathrm{b}}(a)|}{|V|} \geq 1 - 2\delta, \qquad \mathbb{E}I^2 \geq (1 - 4\delta)_+^2, \qquad \mathbb{E}Q^2 \geq (1 - 2\delta)^4.$$

Proof. Let $e_v(J) = \mu^{\beta}_{G,J}(\sigma_v \neq \eta_v\sigma_a)$; then $\eta_v\langle\sigma_a\sigma_v\rangle_J = 1 - 2e_v$, so $|\langle\sigma_a\sigma_v\rangle_J| \geq 1 - 2e_v$. By Lemma 8.1, $\mathbb{E}|C_{\mathrm{b}}(a)| = \sum_v \mathbf{P}(a \overset{\mathrm{b}}{\leftrightarrow} v) \geq \sum_v(1 - 2\mathbb{E}e_v)$, which gives the first bound. Put $h = |C_{\mathrm{b}}(a)|/|V|$. The cluster of a lies in the sector q_a, and the opposite sector has at most $|V| - |C_{\mathrm{b}}(a)|$ vertices, so $q_aI \geq 2h - 1$ pathwise; hence $I^2 \geq (2h-1)^2_+$, and Jensen's inequality for the convex map $x \mapsto (2x-1)^2_+$ gives $\mathbb{E}I^2 \geq (2\mathbb{E}h - 1)^2_+ \geq (1 - 4\delta)^2_+$. For the overlap, $\mathbb{E}[q_aq_v] = \mathbb{E}\langle\sigma_a\sigma_v\rangle_J^2 \geq \mathbb{E}(1 - 2e_v)^2 \geq (1 - 2\delta)^2$, so $\mathbb{E}[q_aQ] \geq (1 - 2\delta)^2$ and, by Cauchy–Schwarz, $\mathbb{E}Q^2 \geq (\mathbb{E}[q_aQ])^2$. □

The last bound uses only the spins; the first two are the CMR content. On hierarchical lattices a reference with small δ is available at low temperature, and the corollary then yields a giant blue cluster and imbalance (Chapter 12). On $\mathbb{Z}^d$ no such reference is known, and supplying one is the ordering problem itself.

8.3 At most two infinite blue clusters

The second link (O2) is a uniqueness statement in the style of Burton and Keane. It holds at every positive temperature and says nothing about existence. Here the coupling law satisfies (D1)–(D2) of [215]: iid couplings with $\nu(0,\infty) > 0$, $\nu(-\infty, 0) > 0$ and no atom at zero. An *admissible joint measure* is a subsequential local limit of the torus laws of (J, σ, τ, c), with c the full CMR colouring.

Theorem 8.3 ([215, Thm. 1.9])**.** *Let $d \geq 2$, let the couplings satisfy (D1)–(D2), and let $\mathbf{P}_\beta$ be an admissible joint measure.*

(a) There are $0 < \beta_{\mathrm{sub}} \leq \beta_{\mathrm{grey}} < \infty$, depending on d and the coupling law, such that almost surely there is no infinite grey cluster for $\beta < \beta_{\mathrm{sub}}$ and exactly one for $\beta > \beta_{\mathrm{grey}}$.

(b) For every $\beta > 0$, almost surely there are at most two infinite blue clusters, at most one in each overlap sector. If there are two, they lie in the same grey cluster and every grey path between them has an odd number of red edges.

Proof sketch. We omit the resampling case analysis and the ergodic reduction [215, §§2–4]. Blue edges are not insertion tolerant even after averaging (an edge is blue only if both replicas satisfy it), so two substitutes replace Lemma 7.3. First, given everything outside a finite set of edges, the law of the couplings there is equivalent to the iid prior [215, Prop. 2.1]: grey edges have finite energy, so Burton–Keane gives a unique infinite grey cluster, and a Peierls argument on red-edge parity (as in Theorem 7.6) gives one for large β. Second, resampling couplings, spins and colours in a box joins, with positive conditional probability, all infinite blue arms of one overlap sign that meet it (*label-class coalescence*). With finitely many, but at least two, infinite blue clusters in a sector, such a merger would change their number with positive probability, against ergodicity; infinitely many would yield a component with three ends, impossible for a translation-invariant random subgraph of $\mathbb{Z}^d$ [113, Prop. 3.3]. □

Corollary 8.4 ([215, Cor. 4.7])**.** *If, in addition, $\mathbf{P}_\beta$ is ergodic and invariant under $\tau \mapsto -\tau$, the number of infinite blue clusters is 0 or 2 almost surely, and the densities of the infinite blue clusters of the two sectors are equal.*

So imbalance is delicate: any ergodic component that keeps the replica-flip symmetry is balanced, and since the selected limits are replica-flip invariant as a whole, imbalance can only appear as a spontaneous breaking of this symmetry across ergodic components.

8.4 Both sectors percolate in high dimension

We now prove (O1) in high dimension. The theorem holds under the finite-mean assumption (8.1); we prove it for laws with an exponential moment (including $\pm J$ and Gaussian) and refer to [214, §3] for the extension, which skips targets incident to rare large couplings.

Theorem 8.5 ([214, Thm. 1.1])**.** *Let the coupling law be symmetric and satisfy* (8.1)*.*

(i) For each fixed $c > 1$ and all sufficiently large d, every $\mathbf{P} \in \mathcal{L}_{d,c/(2dm)}$ satisfies $\mathbf{P}(0 \overset{\mathrm{b}}{\leftrightarrow} \infty,\ q_0 = s) > 0$ for $s = \pm 1$, and $\mathbf{P}$-almost surely each overlap sector contains an infinite blue cluster.

(ii) For each fixed $0 < c < 1$ there are $d_0 < \infty$ and $\rho < 1$ such that for $d \geq d_0$, $0 \leq \beta \leq c/(2dm)$, every $\mathbf{P} \in \mathcal{L}_{d,\beta}$ and every $n \geq 1$, $\mathbf{P}(0 \overset{\mathrm{b}}{\leftrightarrow} \partial B_n) \leq 2\rho^n$.

In particular the first blue onset $\beta_{\mathrm{b}}(d) = \inf\{\beta : \text{some } \mathbf{P} \in \mathcal{L}_{d,\beta} \text{ has } \mathbf{P}(0 \overset{\mathrm{b}}{\leftrightarrow} \infty) > 0\}$ satisfies $2dm\,\beta_{\mathrm{b}}(d) \to 1$.

For couplings with no atom at zero, Theorem 8.3 upgrades (i) to *exactly* two infinite blue clusters, one per sector. The scale $\beta \sim 1/(2dm)$ is natural: at high temperature an edge is blue with probability of order $m\beta$, and a site has $2d$ edges. The difficulty is dependence: the averaged law of the blue edges is not a product, and an exploration reveals failures as well as successes, so the proof needs a lower bound on the probability of the next blue edge that survives *every* history.

8.4.1 The leaf identity and the cost of restoring an edge

The idea is to make the next target a leaf, whose spins can be summed out so that the blue probability of its edge is exactly computable, and then to restore the other edges at v one at a time; because the sign of a restored edge is fair, each restoration costs a factor quadratic, not linear, in β.

Lemma 8.6 (Two-replica leaf identity [214, Lemma 2.1])**.** *Let G be finite and let v be a vertex whose only edge is $e = uv$. Under $\mathbf{P}_{G,\beta}$, conditionally on the couplings off e, on W_e, on all spins $(\sigma_x, \tau_x)_{x \neq v}$ and on all blue indicators $(B_f)_{f \neq e}$, the probability that e is blue equals $\tanh(\beta W_e)$. Moreover the joint law of the couplings off e, the spins off v and the blue indicators off e does not depend on J_e.*

Proof. Let $a = \beta W_e$ and $t = \tanh a$. Given everything off v, each replica satisfies e with probability $e^a/(e^a + e^{-a}) = (1+t)/2$, independently of the other, and a doubly satisfied edge is activated with probability $1 - e^{-4a} = 4t/(1+t)^2$. The product is t. If $W_e = 0$ both sides vanish. For the last claim, summing the spin at v gives the factor $2\cosh a$ in each replica's Boltzmann weight and in its partition function, whatever the sign of J_e; the blue indicators off e are functions of the spins off v, their couplings and their activation variables. □

Lemma 8.7 (Sign-averaged insertion [214, Lemma 2.2])**.** *Let G be finite, $f = xy \in E$, and let $G \setminus f$ be G with f removed (vertices kept). Fix the couplings off f and the magnitude W_f, and write $\mathbb{E}_J^G$ for expectation under $\mu_{G,J}^{\beta} \otimes \mu_{G,J}^{\beta}$ and the activation variables. For every nonnegative function F of (σ, τ) and of the couplings and activation variables of edges other than f,*

$$\operatorname{sech}^2(\beta W_f)\,\mathbb{E}_J^{G\setminus f}[F] \;\leq\; \frac{1}{2}\sum_{\varepsilon_f = \pm 1} \mathbb{E}_J^{G}[F] \;\leq\; \cosh(4\beta W_f)\,\mathbb{E}_J^{G\setminus f}[F].$$

Proof. Let $a = \beta W_f$, $t = \tanh a$, $z = \langle \sigma_x \sigma_y \rangle_{G\setminus f,J}$, $r = \sigma_x\sigma_y$ and $s = \tau_x\tau_y$. Since $e^{a\varepsilon r} = \cosh a\,(1+t\varepsilon r)$, the density of the law on G with sign ε relative to the law on $G\setminus f$ is $(1+t\varepsilon r)(1+t\varepsilon s)/(1+t\varepsilon z)^2$; the squared denominator is the two quenched normalizations. Averaging over ε gives the density

$$\begin{aligned} K_t(r,s;z) &= \frac{1}{2}\,\frac{(1+tr)(1+ts)}{(1+tz)^2} + \frac{1}{2}\,\frac{(1-tr)(1-ts)}{(1-tz)^2} \\ &= \frac{\cosh^2 h}{\cosh^2 a}\cosh\bigl(a(r+s)-2h\bigr), \qquad h = \operatorname{artanh}(tz), \end{aligned} \tag{8.3}$$

where the second form uses $1 \pm tr = e^{\pm ar}/\cosh a$ for $r = \pm 1$ and $1 \pm tz = e^{\pm h}/\cosh h$. As $|z| \le 1$, $|h| \le a$, so the prefactor lies in $[\operatorname{sech}^2 a, 1]$ and the argument of the last cosh has absolute value at most $4a$. This bounds K_t between $\operatorname{sech}^2 a$ and $\cosh 4a$. The other edges' couplings and activation variables are unaffected, so the bounds hold for every such F. □

Expanding (8.3) in t shows the point of averaging the sign: the terms of order t in the two summands are opposite, and $K_t = 1 + t^2(rs - 2z(r+s) + 3z^2) + O(t^4)$. With a fixed sign the density would be $1 + O(t)$, and $2d$ restorations at $\beta \sim 1/(2dm)$ would cost a factor bounded away from one; with fair signs they cost $1 + O(1/d)$ in total. Averaging also over W gives the constants

$$b = \mathbb{E}\tanh(\beta W), \qquad \ell = \mathbb{E}\operatorname{sech}^2(\beta W), \qquad M = \mathbb{E}\cosh(4\beta W), \qquad \varrho = \ell/M \le 1.$$

For a vertex v write $\operatorname{star}(v)$ for its set of edges and let $\mathcal{F}_v^{\mathrm{ext}}$ be the sigma-field generated by the spins $(\sigma_x, \tau_x)_{x \ne v}$ and the blue indicators of edges not in $\operatorname{star}(v)$.

Corollary 8.8 (Conditional opening bounds [214, Cor. 2.3])**.** *Let G be finite with maximum degree Δ and suppose $M < \infty$. For every edge $e = uv$ and every $H \in \mathcal{F}_v^{\mathrm{ext}}$ with $\mathbf{P}_{G,\beta}(H) > 0$,*

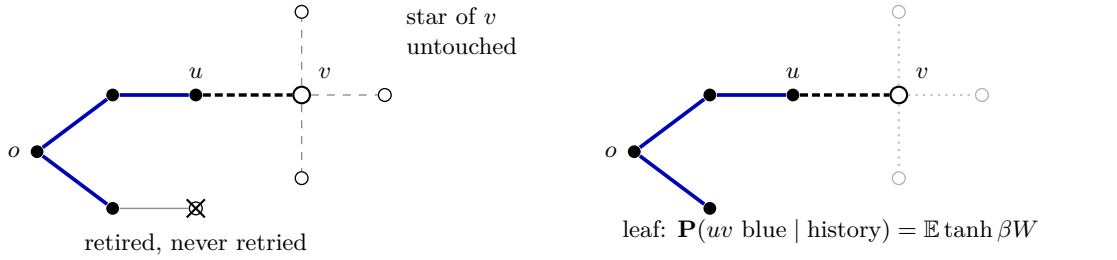
$$b\,\varrho^{\Delta-1} \le \mathbf{P}_{G,\beta}(B_e = 1 \mid H) \le b\,\varrho^{-(\Delta-1)}, \qquad \mathbf{P}_{G,\beta}(q_v = s \mid H) \ge \tfrac{1}{2}\varrho^{\Delta} \quad (s = \pm 1).$$

Proof. Let G' be G with the edges of $\operatorname{star}(v) \setminus \{e\}$ removed, so that v is a leaf. The event H and the variable $\mathbf{1}_H B_e$ are the same functions of the retained coordinates in G and in G'. In G', Lemma 8.6 gives $\mathbb{E}_{G'}[\mathbf{1}_H B_e] = \mathbb{E}_{G'}[\mathbf{1}_H \tanh(\beta W_e)] = b\,\mathbb{E}_{G'}[\mathbf{1}_H]$, because H does not involve J_e and the law of the exterior does not depend on J_e. Now restore the removed edges one at a time, with fresh couplings drawn from the law. By Lemma 8.7, averaged over the restored magnitude and then over all other couplings, each restoration multiplies both $\mathbb{E}[\mathbf{1}_H B_e]$ and $\mathbb{E}[\mathbf{1}_H]$ by factors in $[\ell, M]$; neither observable involves a restored edge. After at most $\Delta - 1$ restorations, $\mathbb{E}_G[\mathbf{1}_H B_e] \ge \ell^{\Delta-1} b\,\mathbb{E}_{G'}[\mathbf{1}_H]$ and $\mathbb{E}_G[\mathbf{1}_H] \le M^{\Delta-1}\mathbb{E}_{G'}[\mathbf{1}_H]$; dividing gives the lower bound, and exchanging the roles of ℓ and M gives the upper bound. For the overlap bound remove the whole star; an isolated vertex has q_v fair and independent of the rest, and the Δ restorations cost at most ϱ^{Δ}. □

The allowed histories may contain closed as well as open blue observations and the spins outside v, but nothing about the star of the target; the exploration below respects this.

8.4.2 One attempt per target, and thinning

The second ingredient turns a conditional lower bound into an embedded independent percolation cluster, by two classical devices: never test a target twice, so that its star is untouched when it is tested, and discard some successes with an extra coin, so that every target is accepted with probability exactly p (the thinning step of Grimmett and Stacey [110, proof of Lemma 5]).



(a) exploration: test uv once

(b) cavity graph; each restored edge costs $1 + O(\beta^2)$

Figure 8.3: The two-replica cavity argument. (a) The exploration has reached some vertices through blue edges and retired a failed target; the next target v was never examined, so its star is untouched. (b) Deleting the other star edges makes v a leaf, blue with probability exactly $\mathbb{E}\tanh(\beta W)$ whatever the history; restoring each with a fair sign costs a factor in $[\ell, M]$, which is $1 + O(\beta^2)$ because the linear term cancels. Thinning each success to probability exactly $p = b(\ell/M)^{2d-1}$ embeds an iid site cluster in the blue cluster.

Proposition 8.9 (One-attempt exploration [214, Prop. 2.4])**.** *Let $G = (V, E)$ be countable and locally finite, $o \in V$, and let $\omega \in \{0,1\}^E$ have any law P. Call an event an* off-star history at v *if it specifies finitely many edge variables, none in* $\mathrm{star}(v)$*. Suppose that for every oriented edge uv and every off-star history H at v of positive probability, $P(\omega_{uv} = 1 \mid H) \geq p$. Then there is a coupling of ω with independent Bernoulli(p) site variables $(X_v)_{v \neq o}$, $X_o = 1$, such that the occupied site cluster $C_X(o)$ is contained in the open cluster $C_\omega(o)$. The same holds on any prescribed induced subgraph.*

The proof is a breadth-first exploration that tests each target once, through the edge by which it is first reached, and thins each success by an independent coin to probability exactly p; it is written out in Section A.16.

Figure 8.3 shows the two steps together.

8.4.3 Proof of the high-dimensional theorem

The last input is Kesten's asymptotics of the percolation thresholds of $\mathbb{Z}^d$.

Theorem 8.10 (Kesten [138, Thm. 1])**.** *For all d,*

$$\frac{1}{2d-1} \leq p_c^{\mathrm{bond}}(\mathbb{Z}^d) \leq p_c^{\mathrm{site}}(\mathbb{Z}^d) \leq \frac{1}{2d} + O\Big(\frac{(\log\log d)^2}{d \log d}\Big).$$

Proof of Theorem 8.5 for laws with an exponential moment. Assume $\mathbb{E}e^{\alpha W} < \infty$ for some $\alpha > 0$ and let $\beta \leq \alpha/8$, so that $M < \infty$. Put $\Delta = 2d$ and

$$p = b\,\varrho^{2d-1}, \qquad u = b\,\varrho^{-(2d-1)}.$$

Constants. From $\tanh x \geq x - x^3/3$, $\operatorname{sech}^2 x \geq 1 - x^2$ and $\cosh y \leq 1 + \frac{y^2}{2}\cosh y$,

$$b \geq m\beta - \tfrac{1}{3}\beta^3 \mathbb{E}W^3, \qquad \ell \geq 1 - \beta^2 \mathbb{E}W^2, \qquad M \leq 1 + 8\beta^2 \mathbb{E}\big[W^2 \cosh(4\beta W)\big],$$

so $\varrho \geq 1 - C\beta^2$ for a constant C depending only on the law, and $b \leq m\beta$. At $\beta = c/(2dm)$ this gives $2d\,p \to c$ and $(2d-1)u \to c$ as $d \to \infty$: the $2d-1$ restorations cost $(1 - C\beta^2)^{2d-1} = 1 - O(1/d)$, while the leaf gain is of order $1/d$.

Lower bound, part (i). Fix $c > 1$ and $\beta = c/(2dm)$. By Theorem 8.10, $2d\,p_c^{\mathrm{site}}(\mathbb{Z}^d) \to 1$, so $p > p_c^{\mathrm{site}}(\mathbb{Z}^d)$ for all large d. Fix R and let $L > 2R+1$, so that the ball B_R of $\mathbb{T}_L^d$ is isomorphic to the ball of $\mathbb{Z}^d$. Run the exploration of Proposition 8.9 inside B_R on the blue indicators. Its histories are events about blue indicators of edges outside the target's star, so they lie in $\mathcal{F}_v^{\mathrm{ext}}$ and Corollary 8.8 (with $\Delta = 2d$) gives the hypothesis with parameter p. Hence

$$\mathbf{P}_{\mathbb{T}_L^d,\beta}(0 \overset{\mathrm{b}}{\leftrightarrow} \partial B_R) \;\geq\; P_p^{\mathrm{site}}(0 \leftrightarrow \partial B_R \mid X_0 = 1) \;\geq\; \theta^{\mathrm{site}}(p)/p > 0,$$

uniformly in R and L. The event on the left is local, so the bound holds for every $\mathbf{P} \in \mathcal{L}_{d,\beta}$; letting $R \to \infty$ gives $\mathbf{P}(0 \overset{\mathrm{b}}{\leftrightarrow} \infty) \geq \theta^{\mathrm{site}}(p)/p$. Reversing one replica keeps the blue edges and exchanges the two values of q_0, so this probability splits equally between $s = \pm 1$. That each sector then contains an infinite blue cluster almost surely follows by conditioning on the exterior of a target (Section A.17).

Upper bound, part (ii). Fix $0 < c < 1$ and $\rho = (1+c)/2$; for large d, $(2d-1)u \leq \rho$ at $\beta = c/(2dm)$, and also for all smaller β, since u is nondecreasing in β. A blue connection from 0 to ∂B_n inside B_n contains a self-avoiding blue path of length n from 0. Expose such a path $x_0, \dots, x_n$ edge by edge: the event that the first $k-1$ edges are blue involves no edge of $\mathrm{star}(x_k)$, because the path is self-avoiding, so by the upper bound of Corollary 8.8 each new edge is blue with conditional probability at most u. There are at most $2d(2d-1)^{n-1}$ such paths, so on every large torus $\mathbf{P}(0 \overset{\mathrm{b}}{\leftrightarrow} \partial B_n) \leq \frac{2d}{2d-1}((2d-1)u)^n \leq 2\rho^n$, and the local event passes to every selected limit.

Onset. Part (ii) excludes percolation below $(1-\varepsilon)/(2dm)$ and part (i) gives it at $(1+\varepsilon)/(2dm)$, for every fixed ε and large d; no monotonicity in β is needed for the squeeze. □

Symmetry of the coupling law is used once, in Lemma 8.7, to cancel the linear term (a mean-zero law that is not symmetric does not give this cancellation); the quenched normalization produces the squared denominators; and the dimension enters only through Kesten's asymptotic threshold. Machta, Newman and Stein's comparison of the posterior law of the couplings, given the spins, with tilted product laws [152, Lemmas 4.1–4.2] gives a second proof in high dimension, through an independent mixed site–bond model [214, App. F].

8.4.4 Explicit dimensions

For unit couplings the asymptotic statement can be made explicit, at the price of a computer-assisted proof.

Theorem 8.11 (computer-assisted)**.** *Let the couplings be iid uniform on $\{-1, 1\}$, $t = \tanh\beta$, and let (d, t) be one of the seventeen pairs*

$$\begin{aligned}
&d = 12: \;\; t \in \{\tfrac{3}{25}, \tfrac{11}{100}, \tfrac{23}{200}, \tfrac{1}{8}, \tfrac{13}{100}\}, \qquad d = 11: \;\; t \in \{\tfrac{13}{100}, \tfrac{3}{25}\}, \qquad d = 10: \;\; t \in \{\tfrac{27}{200}, \tfrac{13}{100}\},\\
&d = 9: \;\; t \in \{\tfrac{7}{50}, \tfrac{29}{200}, \tfrac{3}{20}, \tfrac{31}{200}, \tfrac{4}{25}\}, \qquad d = 8: \;\; t = \tfrac{3}{20}, \qquad d = 7: \;\; t \in \{\tfrac{31}{200}, \tfrac{3}{20}\}.
\end{aligned}$$

Then there is an explicit $\theta_*(d,t) > 0$ such that every $\mathbf{P} \in \mathcal{L}_{d,\beta}$ satisfies $\mathbf{P}(0 \overset{\mathrm{b}}{\leftrightarrow} \infty,\; q_0 = s) \geq \theta_*$ for $s = \pm 1$, and almost surely each overlap sector contains an infinite blue cluster.

The seventeen cases are from a revised version of [214] that is not yet public; the public version proves the case $d = 22$, $t = 11/125$ [214, Cor. 4.3]. The certificates are described in Appendix B.

Proof sketch. We omit all estimates and certificates. Locally, the exploration first reveals the whole overlap field and then queries blue edges only between sites of overlap $+1$; in gauge variables the remaining law of a target's star is an explicit tilt of a product law, a bond floor of about $0.9 \tanh 2\beta$ survives any number of failed queries, and the plus sites dominate an Ising field with a small ferromagnetic coupling and a negative field (a one-sided Holley argument). Globally, Kesten's threshold is replaced by a weighted second moment over oriented paths ($d = 9$ to 12) or paths with short lateral excursions ($d = 7$ to 9), whose pair interaction is bounded by an explicit finite criterion. The computer decides finitely many inequalities between explicit sums of rationals exactly and bounds the second-moment criteria rigorously (rational enclosures, or directed rounding followed by an exact rational comparison); two independently written programs check each pair. □

At the listed points $2d\beta$ lies between about 2 and 3; they are points of percolation, not estimates of the onset, and nothing is claimed in dimension six or below.

8.5 Percolation without order

The same regime shows that (O1) carries no ordering information by itself. The tool is a sign-averaged mean-square bound on correlations.

Proposition 8.12 (Centred susceptibility bound [214, Prop. 5.1])**.** *Let G be finite of maximum degree Δ, with zero field and symmetric iid couplings. Let $F(u) = \mathbb{E}\cosh(uW)$, assume $F(4\beta) < \infty$, and put*

$$A_\beta = \tfrac{1}{2}(F(4\beta) + F(2\beta)), \qquad B_\beta = \tfrac{1}{2}(F(4\beta) - F(2\beta)), \qquad \kappa_\Delta = \Delta B_\beta A_\beta^{\Delta-1}.$$

If $\kappa_\Delta < 1$, then $\sup_x \sum_y \mathbb{E}\langle\sigma_x\sigma_y\rangle_J^2 \le 1/(1-\kappa_\Delta)$. For $\Delta = 2d$, every $\mathbf{P} \in \mathcal{L}_{d,\beta}$ has $\chi_{\mathrm{SG}} = \sum_x \mathbb{E}[q_0 q_x] \le 1/(1-\kappa_{2d})$ and spatial overlap $Q = 0$ almost surely.

The mechanism is the mean-square, sign-averaged one that also drives the cavity L^2 lemma of Chapter 9. As $\beta \to 0$, $A_\beta = 1 + 5s_2\beta^2 + O(\beta^4)$ and $B_\beta = 3s_2\beta^2 + O(\beta^4)$ with $s_2 = \mathbb{E}W^2$, so at the blue onset scale $\beta = c/(2dm)$ one has $\kappa_{2d} = 3s_2c^2/(2dm^2) + O(d^{-2}) \to 0$.

Theorem 8.13 (Percolation with finite spin-glass susceptibility [214, Cor. 5.2 and Cor. 5.4])**.** *Let the coupling law be symmetric with an exponential moment, and fix $c > 1$. For all sufficiently large d and $\beta = c/(2dm)$, every $\mathbf{P} \in \mathcal{L}_{d,\beta}$ has infinite blue clusters in both overlap sectors almost surely, while*

$$\chi_{\mathrm{SG}} = 1 + O_c(d^{-1}), \qquad Q = 0 \quad \text{and} \quad D_+ = D_- \quad \text{almost surely.}$$

More generally, $D_+ = D_-$ almost surely in every selected limit whenever $\kappa_{2d} < 1$.

Proof sketch. Percolation is Theorem 8.5(i), and the susceptibility bound is Proposition 8.12 with the expansion above. We omit the proof of exact balance [214, Cor. 5.4]: by Proposition 8.16 below, a nonzero $D_+ - D_-$ would force a nonzero overlap in some finite odd local channel, and the bound of Proposition 8.12 extends to every such channel. □

So the system is a paramagnet in every sense the review uses, yet both sectors percolate. The regime extends well above the onset: $\chi_{\mathrm{SG}} \le 2$ up to β of order $d^{-1/2}$ [214, Cor. 5.2], while blue

percolation sets in at order d^{-1}, so any susceptibility transition in these dimensions lies strictly below the first blue onset in temperature [214, Rem. 5.3]. At the explicit points of Theorem 8.11 an exact-arithmetic certificate also gives $\chi_{\mathrm{SG}} < 7/2$ and $D_+ = D_-$; its criterion and reduction belong to the same revised version and are not reproduced here (Appendix B).

Route warning 8.14 (Percolation without order). *Tempting step:* prove blue percolation in both sectors in $d = 3$ and conclude order. *Status: fails even granting the percolation, with proofs in the analogous regimes.* Both sectors percolate with $\chi_{\mathrm{SG}} < \infty$ and exactly balanced densities on $\mathbb{Z}^d$ in high dimension (Theorem 8.13); this is the lattice case of Route warning 7.25, whose mean-field half and credit (Machta, Newman and Stein) are given there. Order requires (O3), and any replica-flip-symmetric ergodic component is balanced (Corollary 8.4).

8.6 Imbalance would give order

The right half of Fig. 8.1 is a proved implication. Machta, Newman and Stein observed that a percolating blue cluster of strictly larger density than all others breaks spin-flip symmetry, through "agree" or "disagree" boundary conditions on it [151, pp. 4–5]; in the periodic ensemble no statement about finite blue clusters is needed.

Proposition 8.15 (Imbalance gives spin-flip-related Gibbs states [214, Prop. 6.1])**.** *Fix d and $\beta > 0$. Let a stationary law of (J, σ, τ, B) have, for almost every J, a spin marginal that is Gibbs for two uncoupled replicas, and the activation kernel (8.2) given the spins; every selected periodic joint limit qualifies. If $\mathbf{P}(D_+ \neq D_-) > 0$, then for iid couplings, for almost every J there exist two distinct Gibbs states for J that are images of each other under the global spin flip.*

Proof sketch. We omit the measurability arguments. For fixed J, $D = D_+ - D_-$ is unchanged by finite changes of the activation variables, hence (zero-one law) a function $d_J(\sigma, \tau)$ of the spins, and by finite changes of the spins; so d_J is a tail function of the two replicas, odd under reversal of the first. Conditioning on $\{d_J > 0\}$ preserves the product Gibbs specification, as tail conditioning does, and reversing the first replica gives a second Gibbs law with the opposite sign of d_J; disintegrating both over the common law of the second replica yields distinct, spin-flip-related Gibbs states of the first. Ergodicity of the iid couplings upgrades positive probability in J to almost every J. □

Imbalance also has a consequence for the overlap, although only in a local channel that need not be a single site. For a finite set A write $q_{A,x} = \prod_{y \in A} q_{x+y}$ and Q_A for its spatial average.

Proposition 8.16 (Imbalance forces an odd local order channel [214, Prop. 6.2])**.** *Let $\mathbf{P} \in \mathcal{L}_{d,\beta}$ and suppose $\mathbf{P}(D_+ \neq D_-) > 0$. Then for some finite set A of odd cardinality, $\mathbb{E}Q_A^2 > 0$, and the susceptibility $\sum_x \mathbb{E}[q_{A,0} q_{A,x}]$ of that channel is infinite.*

The proof approximates the signed infinite-blue density by finite-radius signed blue-arm events, expands them in spin monomials of the two replicas, and uses their oddness under one replica flip to keep only odd channels [214, App. D]; the same source derives a cusp of the two-replica pressure with the exponent $h \sum_x q_{A,x}$, by the mechanism of van Enter and Griffiths [237].

Remark 8.17 (From an odd channel to the site overlap). A bond-interpolation comparison of odd local channels with the site overlap upgrades Proposition 8.16, for symmetric iid couplings, to ordinary overlap order: imbalance implies $\mathbb{E}Q^2 > 0$, and $\chi_{\mathrm{SG}} = \infty$ for ± 1 or bounded-support

couplings; equivalently, $Q = 0$ almost surely forces $D_+ = D_-$. This upgrade is in a revised version of [214], which is not yet public.

The converse, order $\Rightarrow$ imbalance, is open on $\mathbb{Z}^d$: finite *grey* clusters carry fair overlap signs and cancel (Proposition 7.12(c)), but finite blue clusters inside an infinite grey cluster need not, so the spatial overlap and $D_+ - D_-$ can differ; on finite graphs Corollary 8.2 gives it from a good reference configuration.

8.7 What the blue count says about the periodic metastate

The at-most-two theorem is an upper bound, and an upper bound alone cannot tell the metastate classes of Section 3.6 apart: it holds for every admissible joint measure, whatever the states behind it. It becomes informative once it is paired with a lower bound, a statement that distinct states force blue percolation. This section supplies one and reads the blue count as a description of the pair of pure states that two periodic replicas sample. The results are ours and unpublished; they are stated as observations, with proofs or sketches. Throughout, $K_{J,\sigma,\tau}$ is the product colouring kernel of Definition 7.11, N_s is the number of infinite blue clusters in the sector $\{q = s\}$, and $N_b = N_+ + N_-$.

8.7.1 Two periodic replicas sample one state twice

Recall the periodic metastate κ_J of Theorem 3.5, the pure-state decomposition $\Gamma = \int \alpha \, w_\Gamma(d\alpha)$ of a Gibbs state (Section 3.5) and the global flip $\alpha \mapsto \bar{\alpha}$.

Observation 8.18 (Admissible measures are same-state replica laws). Let the couplings be iid, $\beta < \infty$, and let $\mathbf{P}$ be an admissible joint measure along volumes L_k. Along a further subsequence the periodic states converge as in Theorem 3.5, and the resulting periodic metastate satisfies

$$\mathbf{P}(dJ, d\sigma, d\tau, dc) = \nu(dJ) \int \kappa_J(d\Gamma)\, \Gamma(d\sigma)\, \Gamma(d\tau)\, K_{J,\sigma,\tau}(dc). \tag{8.4}$$

Conversely, every periodic metastate defines an admissible joint measure by (8.4). Moreover, for ν-a.e. J:

(a) κ_J-a.e. Γ is invariant under the global spin flip;

(b) $\mathbf{P}(\cdot \mid J) = \int \Pi_J(d\alpha\, d\alpha')\, \alpha \otimes \alpha' \otimes K_J$, where $\Pi_J = \int \kappa_J(d\Gamma)\, w_\Gamma \otimes w_\Gamma$ is the law of the *pure pair* sampled by the two replicas.

Proof sketch. We omit the measurability of $J \mapsto \Pi_J$. The torus law is the integral of $\mu^{\mathrm{per}}_{L,J} \otimes \mu^{\mathrm{per}}_{L,J} \otimes K_J$ against the law of $(J, \mu^{\mathrm{per}}_{L,J})$. For a local function f of (J, σ, τ, c) that is continuous in J, the map $(J, \Gamma) \mapsto (\Gamma \otimes \Gamma \otimes K_J)(f)$ is bounded and continuous: cylinder probabilities are continuous in Γ, and the colour probabilities are continuous in J_e, including at $J_e = 0$, where both vanish. Prokhorov's theorem gives a further subsequence along which $(J, \mu^{\mathrm{per}}_{L,J})$ converges, and (8.4) follows; the converse is the same computation. For (a), the torus Hamiltonian is even, so every $\mu^{\mathrm{per}}_{L,J}$ is flip-invariant, and flip invariance is a closed condition. For (b), insert $\Gamma \otimes \Gamma = \iint w_\Gamma(d\alpha)\, w_\Gamma(d\alpha')\, \alpha \otimes \alpha'$. □

Two *independently* drawn states $\Gamma, \Gamma' \sim \kappa_J$ would give the spin law $\rho_J \otimes \rho_J$ instead. That law is admissible only if δ_{ρ_J} is itself a periodic metastate: if $\int \kappa'_J(d\Gamma)\, \Gamma \otimes \Gamma = \rho_J \otimes \rho_J$, the barycenter

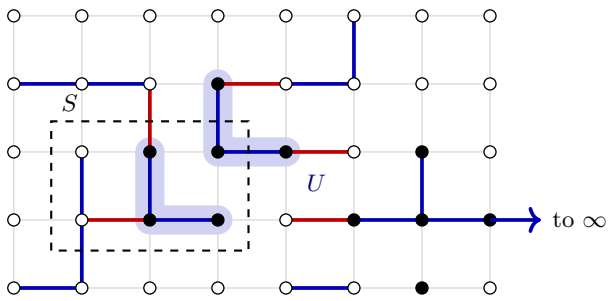


Figure 8.4: The exchange behind Observation 8.19. Sites show $q = \sigma\tau$ (open $+1$, filled -1), with blue and red edges as in Fig. 8.2. The disagreement sites of the window S lie in two finite blue clusters, whose union U is shaded. Reversing both replicas on U exchanges σ and τ on S and changes nothing in the picture: edges leaving U are red or closed, and their weights are permuted among equal entries. The infinite blue cluster of $\{q = -1\}$ on the right does not meet S; one that did would be the event A_S.

of κ'_J is ρ_J, and testing on $f \otimes f$ shows that $\Gamma(f)$ has zero variance under κ'_J for every local f. So in a nontrivial metastate, pairs of pure states taken from *different* states Γ, such as the incongruent cross pairs of the chaotic-pairs picture, never enter the CMR colouring of periodic replicas; the pairs that do enter come from inside one Γ.

8.7.2 Distinct states force blue percolation where they disagree

The lower bound is the flip of the crossing lemma, made on the disagreement side. Where $q = -1$ the replicas are opposite, so reversing both on a set where $q = -1$ *exchanges* them there. If every disagreement site in a window lies in a finite blue cluster, the exchange can be made on those clusters at no cost, and the two replicas then have the same law in the window (Fig. 8.4).

Observation 8.19 (Blue disagreement percolation). Let $G = (V, E)$ be a countable locally finite graph, $J \in \mathbb{R}^E$ arbitrary and $\beta < \infty$. Let α, α' be any two Gibbs states for J on G, and let $\mathbf{Q} = \alpha \otimes \alpha' \otimes K_J$ be the law of (σ, τ, c) with $\sigma \sim \alpha$ and $\tau \sim \alpha'$ independent and c the CMR colouring. For finite $S \subset V$ and every function f of σ_S,

$$|\alpha(f) - \alpha'(f)| \le 2\|f\|_\infty \, \mathbf{Q}(A_S),$$

where A_S is the event that some $x \in S$ with $q_x = -1$ lies in an infinite blue cluster. In particular, if $\mathbf{Q}$-almost surely no infinite blue cluster lies in $\{q = -1\}$, then $\alpha = \alpha'$; if none lies in $\{q = +1\}$, then $\alpha' = \bar{\alpha}$.

Proof. On A_S^c let U be the union of the blue clusters of the sites $x \in S$ with $q_x = -1$: a finite set, on which $q \equiv -1$ (Proposition 7.12(a)). Let Φ reverse both σ and τ on U on the event A_S^c, keep the colouring, and be the identity on A_S. Since $\sigma = -\tau$ on U and $\sigma = \tau$ on $S \setminus U$, Φ exchanges σ_S and τ_S.

The colouring stays admissible and U is unchanged. Edges inside U have both endpoints reversed and keep their satisfaction pattern. An edge with one endpoint in U is not blue, and reversing that endpoint in both replicas exchanges "satisfied by both" with "by neither" and the two "by exactly one" states, so a red edge stays red. Since q and the colouring are unchanged, so are U and A_S, and Φ is an involution.

Φ preserves $\mathbf{Q}$. Condition on $\{U = u\}$ for a finite set u, and on the spins off u and the colours off the set $E(u)$ of edges touching u. The event $\{U = u\}$ is determined by (σ_u, τ_u), $c_{E(u)}$ and the spins off u, and Φ leaves it invariant; the colours off $E(u)$ depend only on the spins off u, so they do not change the conditional law of (σ_u, τ_u). By the DLR equations for α and for α', the conditional weight of $(\sigma_u, \tau_u, c_{E(u)})$ is the product over $e \in E(u)$ of the Boltzmann factors of both replicas and the colour probability. Relative to $e^{2\beta|J_e|}$, and with $r_e = e^{-2\beta|J_e|}$, this factor is $1 - r_e^2$ for a blue edge, $r_e(1 - r_e)$ for a red edge and r_e^2 for a closed edge, for every satisfaction pattern that allows the colour. Φ leaves the factors of edges inside u unchanged and permutes those of edges leaving u among equal entries. Summing over the countably many u gives $\mathbf{Q} \circ \Phi^{-1} = \mathbf{Q}$.

Hence $\mathbf{E}_{\mathbf{Q}}[(f(\sigma) - f(\tau))\mathbf{1}_{A_S^c}]$ equals its own negative and vanishes, and $|\alpha(f) - \alpha'(f)| = |\mathbf{E}_{\mathbf{Q}}[(f(\sigma) - f(\tau))\mathbf{1}_{A_S}]| \le 2\|f\|_\infty \mathbf{Q}(A_S)$. For the last statement apply this to the pair $(\alpha, \bar{\alpha}')$: reversing τ keeps the colouring and exchanges the sectors. □

The exchange is the one Chayes, Machta and Redner used for ferromagnets with plus and minus boundary conditions [55], and Gandolfi used the same symmetry for two copies with a common boundary condition (Theorem 7.20). Blue clusters in $\{q = -1\}$ lie inside the clusters of the disagreement set $\{\sigma \ne \tau\}$, so the observation contains the disagreement-percolation criterion for independent copies [235, 236]. No monotonicity is used. Percolation does not conversely give order: in the regime of Theorem 8.13 both sectors percolate while $Q = 0$.

Observation 8.20 (Extremality and uniqueness criteria). In the setting of Observation 8.19, say that a pair (α, α') has *no disagreement percolation* if $\alpha \otimes \alpha' \otimes K_J$-almost surely no infinite blue cluster lies in $\{q = -1\}$.

(a) If (μ, μ) has no disagreement percolation, then μ is extremal.

(b) If every pair of extremal Gibbs states has no disagreement percolation, the Gibbs state for J is unique.

Proof. (a) Write $\mu \otimes \mu = \iint w_\mu(d\alpha)\, w_\mu(d\alpha')\, \alpha \otimes \alpha'$. The hypothesis holds for $w_\mu \otimes w_\mu$-a.e. pair, so Observation 8.19 gives $\alpha = \alpha'$ for almost every pair, and a probability measure whose square is carried by the diagonal is a point mass. (b) Observation 8.19 makes any two extremal states equal, and every Gibbs state is a mixture of extremal ones. □

Part (a) is the infinite-volume form of Theorem 7.20(i). Part (b) is stated through pairs of extremal states, so it needs no assumption on how Gibbs states arise as boundary limits. Copies with a common boundary condition test a pair (α, α') only when both occur in the decomposition of one limit state. Two periodic replicas are such copies, and inside one periodic state Γ they test $w_\Gamma \otimes w_\Gamma$-almost every pair of its pure states, which is what the next two observations use.

8.7.3 The blue count reads the pair of pure states

For extremal α, α' the CMR law $\alpha \otimes \alpha' \otimes K_J$ is trivial on the tail σ-field: extremal states are asymptotically decoupled [106, Prop. 7.9], the decoupling passes to products, and the colours are independent given the spins. The events $\{N_s \ge 1\}$ and the limits $D_\pm$ and Q do not change under finite modifications of (σ, τ, c), so each is a constant under the CMR law of almost every pure pair. With at most one infinite cluster per sector, $N_s \in \{0, 1\}$ is then a property of the pair, and Observation 8.19 says which. Write $q_{\alpha\alpha'} = \lim_n |\Lambda_n|^{-1} \sum_{x \in \Lambda_n} \alpha(\sigma_x)\, \alpha'(\sigma_x)$ for the overlap of two pure states.

Observation 8.21 (Pure-pair dictionary). Let $d \geq 2$, let the couplings satisfy (D1)–(D2), let $\beta > 0$, and let $\mathbf{P}$ be an admissible joint measure with pure-pair law Π_J (Observation 8.18). For ν-a.e. J and Π_J-a.e. (α, α'), almost surely under $\alpha \otimes \alpha' \otimes K_J$: if $\alpha' \neq \alpha$ then $N_- = 1$, and if $\alpha' \neq \bar{\alpha}$ then $N_+ = 1$. Consequently the pair type determines the blue count as in Table 8.1. In particular, for Π_J-a.e. pair of distinct pure states that are not flips of each other there are exactly two infinite blue clusters, one in each sector, lying in one grey cluster, and every grey path between them has an odd number of red edges.

pure pair (α, α')	N_+	N_-	Q
$\alpha' = \alpha = \bar{\alpha}$ (one flip-invariant state)	$N_+ = N_- \in \{0, 1\}$		0
$\alpha' = \alpha \neq \bar{\alpha}$ (same state, broken flip symmetry)	1	0 or 1	$q_{\alpha\alpha} \geq 0$
$\alpha' = \bar{\alpha} \neq \alpha$ (a flip pair)	0 or 1	1	$-q_{\alpha\alpha} \leq 0$
$\alpha' \notin \{\alpha, \bar{\alpha}\}$ (not flip-related)	1	1	$q_{\alpha\alpha'}$

Table 8.1: The blue count of a pure pair, almost surely, for almost every pair sampled by two periodic replicas (Observation 8.21). Read backwards: $N_b = 0$ forces one flip-invariant state; $N_b = 1$ forces a flip-related pair with broken flip symmetry, the percolating sector being $+$ exactly when $\alpha' = \alpha$; and $Q \neq 0$ forces percolation in the sector of sign $\operatorname{sgn} Q$.

Proof sketch. We omit the measurability details. Under the tail-trivial law of a pure pair, $\{N_- \geq 1\}$ has probability 0 or 1; if 0, Observation 8.19 gives $\alpha = \alpha'$, and similarly $\{N_+ \geq 1\}$ with $\bar{\alpha}'$ in place of α'. Theorem 8.3 holds $\mathbf{P}$-almost surely, hence, by Observation 8.18(b), for almost every pure pair, and it gives $N_s \leq 1$ and the grey statement. In the first row, reversing τ maps the pair law to itself ($\alpha' = \bar{\alpha}'$) and exchanges the sectors, so the two constants agree. The constant value of Q is its mean, $q_{\alpha\alpha'}$ by bounded convergence, which vanishes when α is flip-invariant, since then every magnetization vanishes. □

8.7.4 Consequences for the four classes

Let $W_k(\Gamma)$ be the weights of the atoms of w_Γ, grouped into flip orbits (a flip pair $\{\alpha, \bar{\alpha}\}$ or a single flip-invariant pure state), and let

$$s = \mathbb{E} \int \kappa_J(d\Gamma) \sum_k W_k(\Gamma)^2$$

be the probability that the two periodic replicas fall in the same flip orbit. Then $s = 1$ exactly when almost every Γ is trivial in the sense of Section 3.6.

Observation 8.22 (Blue counts and the metastate classes). Let $d \geq 2$, let the couplings satisfy (D1)–(D2), let $\beta > 0$, and let $\mathbf{P}$ be an admissible joint measure, with periodic metastate κ as in Observation 8.18.

(a) If $\mathbf{P}(N_- \geq 1) = 0$ or $\mathbf{P}(N_+ \geq 1) = 0$, then for ν-a.e. J, κ_J-a.e. Γ is a single flip-invariant pure state, and $N_b = 0$ and $Q = 0$ almost surely.

(b) $\mathbf{P}(N_b = 2) \geq 1 - s$. Hence if two infinite blue clusters never occur, almost every Γ is a single pure state or a single flip pair, which excludes type (c) of Theorem 3.21 and classes 3 and 4. In classes 3 and 4, $\mathbf{P}(N_b = 2) > 0$.

(c) On $\{Q \neq 0\}$ the sector of sign $\operatorname{sgn} Q$ contains an infinite blue cluster. So $\mathbb{E}Q^2 > 0$ implies $\mathbf{P}(N_b \geq 1) > 0$.

(d) If for ν-a.e. J and κ_J-a.e. Γ almost every pure state of Γ is flip-invariant (for instance under Gibbs uniqueness), then $N_+ = N_-$, $D_+ = D_-$ and $Q = 0$ almost surely. Conversely, if $\mathbf{P}(N_b = 1) > 0$, $\mathbf{P}(D_+ \neq D_-) > 0$ or $\mathbf{P}(Q \neq 0) > 0$, then for ν-a.e. J, with positive κ_J-probability Γ gives positive weight to pure states $\alpha \neq \bar{\alpha}$; in particular there are at least two extremal Gibbs states.

(e) In class 1, N_b is almost surely constant: equal to 1 or 2 if the trivial state is a flip pair, and to 0 or 2 if it is a single flip-invariant pure state.

Proof sketch. We use Table 8.1 and omit the zero-one arguments in J, which use the ergodicity of iid couplings. (a) If $N_- = 0$ almost surely, then $\alpha' = \alpha$ for almost every pure pair, so $w_\Gamma \otimes w_\Gamma$ is carried by the diagonal and w_Γ is a point mass; the pure state is flip-invariant by Observation 8.18(a), and the first row of the table gives $N_+ = N_- = 0$ and $Q = 0$. If $N_+ = 0$ almost surely, then $\alpha' = \bar{\alpha}$ for almost every pair, which again forces $w_\Gamma = \delta_\alpha$ with $\alpha = \bar{\alpha}$. (b) By the table, $N_b \leq 1$ requires $\alpha' \in \{\alpha, \bar{\alpha}\}$, which has probability $\sum_k W_k(\Gamma)^2$ given Γ, since only atoms are hit twice. If $\mathbf{P}(N_b = 2) = 0$ then $s = 1$; a nontrivial Γ has $\sum_k W_k^2 < 1$. (c) is the last column of the table. (d) If $\alpha = \bar{\alpha}$ or $\alpha' = \bar{\alpha}'$, reversing the corresponding replica maps the pair law to itself and exchanges the sectors, so the constants are symmetric; the converse is the contrapositive. (e) If $\Gamma_J = \frac{1}{2}(\alpha + \bar{\alpha})$, the four pure pairs are exchanged by the joint flip, which changes no observable, and by the reversal of one replica, which exchanges the sectors; so N_b is a function of J, translation invariant and hence constant, and its value is read off the table. □

Part (d) complements Corollary 8.4: that corollary needs an ergodic, replica-flip-invariant joint measure, while (d) needs flip invariance only of the pure states, which allows nontrivial metastates of flip-invariant pure states. Its converse is a metastate form of Proposition 8.15. Parts (c) and (d) put order and imbalance on the two sides of blue percolation: order forces percolation in the sector of its sign, imbalance forces spin-flip-related states, and percolation in both sectors forces neither order nor imbalance (Theorem 8.13).

Route warning 8.23 (Two replicas see one pair of pure states at a time). *Tempting step:* many pure states should produce many infinite blue clusters, so the at-most-two theorem should exclude replica symmetry breaking, or at least separate chaotic pairs from the droplet picture. *Status: fails; the at-most-two theorem holds for the admissible measure of every periodic metastate, whatever its class (Observation 8.18).* Two periodic replicas sample one pair of pure states, and almost every sampled pair of distinct pure states that are not flips of each other gives exactly two clusters, however many states there are (Observation 8.21). What the count does see is in Observation 8.22: never two clusters excludes classes 3 and 4, and a count that is not almost surely constant excludes class 1. These results do not separate class 1 from class 2 beyond that constancy, or class 3 from class 4, or count pure states beyond one flip orbit versus several; this is a limit of the present tools, not a proof that the classes cannot be told apart. The converse of Observation 8.22(c) fails, since both sectors percolate with $Q = 0$ in the regime of Theorem 8.13; the other converses are open.

Remark 8.24 (Other boundary conditions). Observation 8.19 needs no boundary condition, and the proof of Theorem 8.3 extends to translation-invariant joint laws with the same comparison of coupling laws, finite-box non-nullness of the two-replica specification and the same local colour kernel [215, Discussion]. So the dictionary should extend to translation-covariant metastates

built from fixed, coupling-independent boundary conditions, with boxes centred at a random point to restore translation covariance; we have not checked this in detail. The conclusion of Observation 8.22(a) and the values in (e) use the flip invariance of periodic states and fail for a pure state $\Gamma = \alpha \neq \bar{\alpha}$, which has $N_+ = 1$. With plus boundary conditions in a two-state picture one expects only the pairs (α, α) and $(\bar{\alpha}, \bar{\alpha})$, and then $N_+ = 1$ always.

Part (a) also suggests a periodic analogue of Open problem 7.21, with a weaker conclusion: for $\pm J$ couplings on $\mathbb{Z}^2$ and every β, show that $\mathbf{P}(N_- \geq 1) = 0$ for every admissible joint measure. This would make the periodic states single flip-invariant pure states and $Q = 0$ at every temperature, so no EA order in periodic limits; uniqueness of all Gibbs states would still need the pair hypothesis of Observation 8.20(b).

Remark 8.25 (Two dimensions). In $d = 2$ the dictionary meets a half-plane uniqueness result that needs neither FKG nor finite energy: for a translation-invariant ergodic edge percolation on $\mathbb{Z}^2$ with finitely many infinite clusters, the restriction to the upper half-plane has almost surely at most one infinite cluster [142, Lemma 2.5], and the authors remark that their methods extend to half-planes bounded by lines of rational slope. Applied to the blue edges of an ergodic component of an admissible joint measure, with Theorem 8.3, this shows that almost surely two coexisting infinite blue clusters do not both have infinite pieces in the upper half-plane, nor, granting the extension, in any rational half-plane. By Observation 8.21, almost every incongruent pair of pure states inside one periodic state would therefore come with two infinite blue clusters of opposite overlap, no rational half-plane containing infinite pieces of both. If such coexistence could be excluded for the admissible joint measures, every periodic state would be a single pure state or a single flip pair (Observation 8.22(b)). Excluding it also for the independent-state law $\rho_J \otimes \rho_J$, which contains the incongruent pairs drawn from different states and to which the at-most-two argument should extend (Remark 8.24), would make the periodic metastate a point mass (class 1) at every β. That would exclude the exchangeable metastates of Route warning 4.19, which the variance method cannot reach, but it would not by itself give non-order. These consequences are unaudited.

8.8 Why the technique stops short of three dimensions

The proof of Theorem 8.5 needs an opening estimate that beats the percolation threshold, and Kesten's asymptotics; the explicit proofs push both down to $d = 7$. Going further would not, by itself, settle anything.

Route warning 8.26 (Imbalance needs an estimate that distinguishes the sectors). *Tempting step:* push the opening estimates down to $d = 3$, then deduce order. *Status: fails even under optimistic assumptions.* Grant both-sector blue percolation in $d = 3$. Every estimate of the kind used here is invariant under the reversal of one replica, which exchanges the sectors, so it gives the same lower bound in both and none on $|D_+ - D_-|$; the balanced regime of Theorem 8.13 shows that such estimates are compatible with exact balance. What is needed is a sector-asymmetric estimate, such as a boundary signal surviving to large scales. *Credit:* the limitation is stated in [214, §6.3].

Route warning 8.27 (The Machta–Newman–Stein auxiliary model is subcritical in dimension six). *Tempting step:* reach low dimensions with the independent mixed site–bond model, allowing repeated attempts at a target [214, App. F], to which Machta, Newman and Stein's comparison of the posterior law of the couplings with tilted product laws leads [152, Lemmas 4.1–4.2]. *Status:*

proved to fail at $d = 6$ for unit couplings. With $t = \tanh\beta$ and $\Delta = 2d$, the auxiliary model's site-times-bond parameter satisfies $p_\Delta \le t/(1 + 3\Delta t^2) \le 1/(2\sqrt{3\Delta})$, which at $d = 6$ is at most $1/12 < 1/11 \le p_c^{\text{site}}(\mathbb{Z}^6)$ (Theorem 8.10), so the auxiliary model is subcritical by path counting [214, App. G]. A model-specific input is needed below this point.

Nor does the engine of Theorem 8.11 certify $d = 6$ (floating-point screens of its pair-interaction criterion exceed the threshold for every path family tried): numerical evidence about one engine, neither a lower bound nor a statement about percolation in $d = 6$.

8.9 Summary

On $\mathbb{Z}^d$ the number of infinite CMR blue clusters is settled: at most one per sector, for $d \ge 2$ and couplings without an atom at zero (Theorem 8.3). Existence in both sectors is proved in high dimension, at $\beta = c/(2dm)$ for each fixed $c > 1$ (Theorem 8.5), and, computer-assisted, at seventeen explicit temperatures for fair ± 1 couplings in $d = 7$–12 (Theorem 8.11); it is open for $d \le 6$. Imbalance would give distinct spin-flip-related Gibbs states and local overlap order (Propositions 8.15 and 8.16), but percolation is compatible with a paramagnet (Theorem 8.13) and is proved by sector-symmetric estimates (Route warning 8.26). The open step follows; by Theorem 8.13 its β must lie outside the region $\kappa_{2d} < 1$, which in high dimension contains an interval at scale $\beta \asymp d^{-1/2}$, well above the blue onset. Paired with the fact that distinct states force blue percolation where they disagree (Observation 8.19), the count reads the pair of pure states sampled by two periodic replicas (Observation 8.21): never seeing two clusters would exclude metastate classes 3 and 4, but the count tells chaotic pairs from droplets only through its constancy (Route warning 8.23). On hierarchical lattices the analogous statement is proved (Chapter 12); the comparison with $\mathbb{Z}^d$ resumes in Chapter 15.

Open problem 8.28 (Persistent imbalance [214, Hyp. 6.3])**.** For iid uniform ± 1 couplings, find a finite d, a finite $\beta > 0$ and $\mathbf{P} \in \mathcal{L}_{d,\beta}$ with $\mathbb{E}(D_+ - D_-)^2 > 0$.

Part V

Toward non-order in low dimension

Chapter 9

High temperature and its edge

At high enough temperature every spin system forgets its boundary condition, and for the Edwards–Anderson model this is easy to prove. But the easy proofs stop well above the temperature at which simulations see a transition (or, in $d = 2$, see none), and between the two lies the whole region in which the order question of this review is decided. This chapter asks how far down in temperature the absence of order can be proved, in which sense, and what stops each method; a proof of non-order down to $T = 0$ in $d = 2$ would settle half of the central question of Chapter 1.

Two notions of "no order" are in play (Definition 9.1): uniqueness, which concerns every Gibbs state at once, and L^2 non-order, which gives rates and is what averaging over the coupling signs makes accessible. Sign-blind comparisons with the ferromagnet of couplings $|J|$, sharpened by De Santis, and the multiscale expansion of Fröhlich and Imbrie for unbounded couplings give uniqueness (Sections 9.2 and 9.3). Integrating out one spin and averaging over the signs of its couplings gives L^2 non-order uniformly in the dimension and, with computer certificates, down to $T \geq 1.90926$ on $\mathbb{Z}^2$ and $T \geq 2.73214$ on $\mathbb{Z}^3$ for $\pm J$ couplings, for finite subgraphs with free boundary conditions and tori of side at least 11 (Sections 9.4 and 9.5); Fig. 9.3 sets these against the numerical critical temperatures. The Kac limit closes the chapter (Section 9.7). By the end the reader can prove the cavity L^2 inequality and the dimension-uniform bound on χ_{SG}, and say why none of these methods can reach the transition.

9.1 Two notions of a high-temperature phase

Definition 9.1 (Uniqueness and L^2 non-order)**.** Fix d, a coupling law ν and $\beta \geq 0$.

(a) *Uniqueness* holds if for ν-a.e. J the set $\mathcal{G}_J(\beta)$ of infinite-volume Gibbs states has one element.

(b) Let $\mathcal{F}$ be a family of finite graphs with boundary prescriptions (free boxes, tori, boxes with coupling-independent boundary conditions). L^2 *non-order* holds on $\mathcal{F}$ with susceptibility bound $\bar{\chi}$ if $\sup_u \sum_v \mathbb{E}\langle \sigma_u \sigma_v \rangle_G^2 \leq \bar{\chi}$ for every $G \in \mathcal{F}$, and *exponential* L^2 *decay* holds if $\mathbb{E}\langle \sigma_u \sigma_v \rangle_G^2 \leq C e^{-m \operatorname{dist}_G(u,v)}$ with $C < \infty$ and $m > 0$ independent of $G \in \mathcal{F}$ and of u, v.

Uniqueness forces the unique state to be flip invariant, but gives no rate of decay and no bound on χ_{SG}; L^2 non-order gives both, but only for the prescriptions in $\mathcal{F}$, and it rules out the finite-volume and stationary order targets of Chapter 1 for them (Corollary 9.9). The simplest

example shows the quantity that governs L^2 decay. On a tree with free boundary conditions the gauge removes every sign (Chapter 1), and

$$\langle \sigma_u \sigma_v \rangle_J = \prod_{e \in \gamma(u,v)} \tanh(\beta J_e), \tag{9.1}$$

where $\gamma(u, v)$ is the path from u to v.

Example 9.2 (Trees and the Bethe value)**.** Let $T_{\Delta,R}$ be the ball of radius R around the root of the tree in which every vertex has degree Δ, with free boundary conditions and iid couplings, and let $m_2 = \mathbb{E}\tanh^2(\beta J_e)$. By (9.1) and independence along the path, $\mathbb{E}\langle \sigma_u \sigma_v \rangle_J^2 = m_2^{\mathrm{dist}(u,v)}$, so

$$\sum_v \mathbb{E}\langle \sigma_{\mathrm{root}} \sigma_v \rangle_J^2 = 1 + \sum_{n=1}^{R} \Delta(\Delta-1)^{n-1} m_2^n,$$

which is bounded uniformly in R if and only if $(\Delta - 1)\, m_2 < 1$. For $\Delta = 2d$ the borderline $(2d-1)\mathbb{E}\tanh^2(\beta J) = 1$ is the *Bethe value*; at large d it is $\beta \approx ((2d-1)\mathbb{E}J^2)^{-1/2}$. As an estimate of the transition on $\mathbb{Z}^d$ it is a mean-field heuristic; on trees the computation is exact (such analyses go back to [52]). On $\mathbb{Z}^d$ cycles carry frustration, (9.1) fails, and the methods below recover a version of the tree recursion at the cost of a constant that keeps them away from the Bethe value.

9.2 The high-temperature expansion and Griffiths singularities

For a finite graph $G = (V, E)$ with free boundary conditions, write $e^{\beta J_e \sigma_x \sigma_y} = \cosh(\beta J_e)\,(1 + \sigma_x \sigma_y \tanh \beta J_e)$ and expand the product. Summing over σ kills every term with a vertex of odd degree, so

$$Z_{G,J} = 2^{|V|} \prod_{e \in E} \cosh(\beta J_e) \sum_{F \subseteq E:\, \partial F = \emptyset} w(F), \qquad \langle \sigma_A \rangle_G = \frac{\sum_{\partial F = A} w(F)}{\sum_{\partial F = \emptyset} w(F)}, \tag{9.2}$$

with $w(F) = \prod_{e \in F} \tanh(\beta J_e)$, where ∂F is the set of vertices of odd degree in (V, F). An even subgraph splits into vertex-disjoint connected components, the *polymers*, over which w factorizes: the partition function is that of a gas of non-overlapping polymers with activities $w(\gamma)$ (Fig. 9.1). For bounded couplings and small β all activities are small and the standard cluster expansion converges, giving uniqueness and exponential decay, on a domain that deteriorates as the range of values of the couplings grows [96, Sect. 8].

For unbounded couplings this fails because of rare regions: arbitrarily large cubes on which all couplings are large occur somewhere, correlation lengths near them are large, and the radius of convergence of the ordinary expansion is presumably zero; even for bounded support with nonzero variance it is expected to diverge inside the single-phase region [96, p. 146]. The prototype is the dilute ferromagnet (bond density p): its free energy has an essential singularity in the field at $h = 0$ for all $p \in (0, 1)$ and $T < T_c(1)$, the undiluted critical temperature [108], which makes the high-temperature expansion diverge well above $T_c(p)$ [51, p. 403]: rare large clusters behave like pieces of an ordered ferromagnet even below the percolation threshold. For the EA model the same mechanism is expected, with rare regions of large or unfrustrated couplings; we are not aware of a proof that Griffiths singularities occur for symmetric couplings on $\mathbb{Z}^d$.

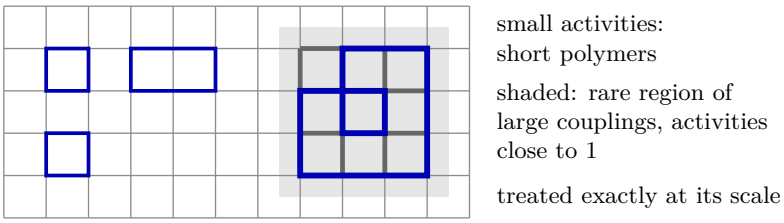


Figure 9.1: The high-temperature expansion (9.2) as a gas of non-overlapping polymers (blue) with activities $\prod_{e\in\gamma}\tanh(\beta J_e)$. Where couplings are bounded, activities are small and long polymers rare. In a rare region of large couplings (shaded, thick edges) activities are close to 1, so polymers of every size there carry weight of order one and the expansion cannot converge; such regions occur somewhere with probability one. Fröhlich and Imbrie absorb them, scale by scale, into the unperturbed system.

The theorem of Fröhlich and Imbrie beats this obstacle (the same paper treats the random-field Ising model at large field variance and a predominantly ferromagnetic glass, Chapter 10); its boundary spins are arbitrary [96, Sect. 8, p. 174]. Connected correlations $\langle\sigma_{i_1};\dots;\sigma_{i_n}\rangle$ have *tree decay with rate* M if, for fixed sites $x_1,\dots,x_n$ and their dilations $a_k=\theta x_k$, $\theta=1,2,\dots$, they are bounded by $C(J,h;x_1,\dots,x_n)\exp(-M|T(a_1,\dots,a_n)|)$, where $|T|$ is the Euclidean length of the shortest tree with endpoints $a_1,\dots,a_n$ [96, (1.14)].

Theorem 9.3 (Fröhlich–Imbrie [96, Thm. 8.1])**.** *Let $d\geq 1$, let the couplings be iid Gaussian with any mean and any variance, and let the fields be iid with any law, independent of the couplings. For all sufficiently small $\beta>0$, with probability one, for every finite family of sites $i_1,\dots,i_n$ the connected correlation $\langle\sigma_{i_1};\dots;\sigma_{i_n}\rangle_{\beta,\Lambda}$ in finite volume converges as $\Lambda\uparrow\mathbb{Z}^d$, exponentially fast, to a limit that does not depend on the boundary condition, and the limits have tree decay with rate $M\geq c\,|\log\beta|>0$.*

Proof sketch. We omit all estimates. Small couplings ($|J_e|\leq 1$ after rescaling) are expanded as in (9.2); around large ones one expands about a local energy minimizer [96, (8.3)], merging singular components close on increasing length scales into the unperturbed system step by step. Large singular regions have probability exponentially small in their volume, which beats their entropy, so $\log Z$ has an almost surely absolutely convergent local expansion [96, Sect. 1 and 8]. □

9.3 Comparison with the ferromagnet, and a strict improvement

Newman's theorem (Chapter 1) compares the frustrated random-cluster measure of J with the ferromagnetic one of $|J|$: if the ferromagnet with couplings $|J|$ has a unique Gibbs state at β, so does J. For iid couplings this gives uniqueness whenever $\mathbb{E}(1-e^{-2\beta|J_e|})<p_c(d)$, and for $\pm J$ couplings whenever $T>T_c^{\mathrm{Ising}}(d)$. No argument of the same kind can do better:

Observation 9.4 (Sign-blind criteria stop at the ferromagnet). Let $\mathcal{K}$ be a set of coupling magnitudes $K=(|J_e|)_e$ such that for every $K\in\mathcal{K}$ and every choice of signs the EA model with couplings $(\pm K_e)_e$ has a unique Gibbs state at β. Then the ferromagnet with couplings K has a unique Gibbs state at β for every $K\in\mathcal{K}$, because it is one of these sign choices.

So uniqueness for $\pm J$ couplings below $T_c^{\mathrm{Ising}}(d)$ must use the signs, and can only hold for almost every sign configuration. De Santis used the almost sure presence of frustrated plaquettes; we state his result in the interval form in which we use it.

Theorem 9.5 (De Santis [70, Thm. 1]). *Let $d \geq 2$ and let the couplings be iid with law $\nu = \nu_{1-\tilde{p}}$, that is, $\mathbb{P}(J_e = -1) = \tilde{p} \in (0,1)$ and $\mathbb{P}(J_e = +1) = 1 - \tilde{p}$. There is $\beta^* > \beta_c^{\mathrm{Ising}}(d)$, depending only on d and $\tilde{p}$, such that for ν-a.e. J the Gibbs state is unique for every $\beta < \beta^*$.*

Proof sketch. We omit the Markov-chain coupling and the strict inequality for percolation thresholds. The frustrated random-cluster measure is the ferromagnetic one conditioned on no open cycle being frustrated (Chapter 1). On each block containing a frustrated plaquette its block-update kernel is dominated by the ferromagnetic one with the block's couplings lowered to some $J_2 < 1$ [70, Sect. 3 and pp. 22–23], so the frustrated measure is dominated by a ferromagnet with couplings 1 or J_2, the latter on a positive density of blocks [70, (65)–(66)], whose critical point is strictly above $\beta_c^{\mathrm{Ising}}(d)$ by a strict inequality of Aizenman–Grimmett type [70, Thm. 4 and (67)]; the comparison of Chapter 1 turns absence of percolation into uniqueness [70, Thm. 3 and (68)]. □

This answers an open problem of Newman [70, p. 2]. The gain is strict but not explicit, and the comparison is from above only, so it proves uniqueness and never order. Going well below T_c^{Ising} requires the law of the signs; the rest of the chapter uses it, replacing uniqueness by the L^2 notion.

9.4 The cavity L^2 inequality

We prove a susceptibility bound in every dimension, with a constant that does not degrade as $d \to \infty$, from one exact computation: integrate out one spin and average the squared correlation over the signs of its couplings. It reproduces the tree recursion of Example 9.2 with one extra factor, which accounts for frustration.

Setting. Let $G = (V, E)$ be a finite multigraph without loops (parallel edges are allowed; the torus of side 2 has them). A symmetric law is the law of $\varepsilon|J|$ with ε a fair sign independent of $|J|$, so we write $J_e = \varepsilon_e |J_e|$ with independent fair signs, and put $K_e = \beta |J_e|$, $t_e = \tanh K_e$. All Gibbs measures in this section have free boundary conditions. For $S \subseteq V$ let $\sigma_S = \prod_{x \in S} \sigma_x$ and

$$\Xi_G(S) := \mathbb{E}\,\langle \sigma_S \rangle_G^2,$$

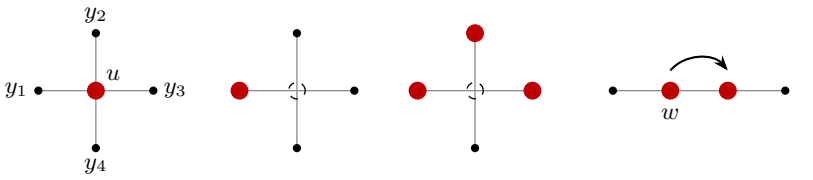
the expectation being over all couplings of G. Thus $\Xi_G(\emptyset) = 1$, $\Xi_G(S) = 0$ for $|S|$ odd, and $\Xi_G(\{u,v\}) = \mathbb{E}\langle \sigma_u \sigma_v \rangle_G^2$.

Lemma 9.6 (Cavity L^2 inequality). *Let $u \in V$ have star edges $e_i = uy_i$, $i = 1, \dots, D$ (the y_i need not be distinct), let $H = G - u$ and $B \subseteq V \setminus \{u\}$. Fix the couplings of H and the star magnitudes $K_i = K_{e_i}$, and average only over the star signs $\varepsilon = (\varepsilon_i)$. Then*

$$\mathbb{E}_\varepsilon \langle \sigma_u \sigma_B \rangle_G^2 \;\leq\; L_u \sum_{A \subseteq [D],\ |A| \text{ odd}} \Big(\prod_{i \in A} t_i^2 \Big)\, \langle \sigma_{B \triangle \partial A} \rangle_H^2, \qquad L_u := \frac{\prod_{i=1}^D \cosh^2 K_i}{\min_{s \in \{\pm 1\}^D} \cosh^2\big(\sum_i s_i K_i\big)}, \tag{9.3}$$

where ∂A is the set of y_i, $i \in A$, counted modulo 2. Consequently, for iid symmetric couplings,

$$\Xi_G(\{u\} \cup B) \leq \sum_{A \subseteq [D],\ |A| \text{ odd}} \omega_u(A)\, \Xi_H(B \triangle \partial A), \qquad \omega_u(A) := \mathbb{E}\Big[L_u \prod_{i \in A} t_i^2 \Big]. \tag{9.4}$$



(a) particle at u (b) $A = \{1\}$: weight t_1^2 $A = \{1,2,3\}$: weight $t_1^2 t_2^2 t_3^2$ (c) a child on a particle: both vanish

Figure 9.2: One step of the cavity expansion. (a) The observable σ_S is a set of particles (red); one particle u is expanded. (b) Integrating out σ_u and averaging over its star signs replaces u by an odd set ∂A of neighbours, with weight $\prod_{i\in A} t_i^2$ times the loss L_u, and deletes u (dashed). (c) A child landing on a particle cancels it. Each step deletes a vertex, so every coupling is used at most once and the disorder average of a history factorizes over its steps.

Proof. Let μ be the Gibbs measure of H and $h(\sigma) = \sum_i K_i \varepsilon_i \sigma_{y_i}$. Summing out σ_u gives $\langle \sigma_u \sigma_B \rangle_G = \mu(\sigma_B \sinh h)/\mu(\cosh h)$. Since $\varepsilon_i \sigma_{y_i} = \pm 1$, $e^{\pm h} = \prod_i (\cosh K_i \pm \varepsilon_i \sigma_{y_i} \sinh K_i)$, and the odd part is

$$\sinh h = \sum_{|A| \text{ odd}} s_A \, \varepsilon^A \sigma_{\partial A}, \qquad s_A = \prod_{i\in A} \sinh K_i \prod_{i\notin A} \cosh K_i, \quad \varepsilon^A = \prod_{i\in A} \varepsilon_i.$$

Hence $\mu(\sigma_B \sinh h) = \sum_{A \text{ odd}} s_A \varepsilon^A \mu(\sigma_{B \triangle \partial A})$. The characters ε^A are orthonormal under the fair product law, and μ does not depend on the star signs, so $\mathbb{E}_\varepsilon \mu(\sigma_B \sinh h)^2 = \sum_{A \text{ odd}} s_A^2 \, \mu(\sigma_{B\triangle\partial A})^2$ with $s_A^2 = \prod_i \cosh^2 K_i \prod_{i\in A} t_i^2$. For the denominator,

$$\mu(\cosh h) \geq \min_\sigma \cosh h(\sigma) \geq \min_{s \in \{\pm 1\}^D} \cosh\Big(\sum_i s_i K_i\Big);$$

the last minimum does not involve the signs, and it remains a lower bound when some y_i coincide. Dividing gives (9.3). The star magnitudes, the star signs and the couplings of H are independent, and $L_u \prod_{i\in A} t_i^2$ depends only on the star magnitudes, so averaging (9.3) gives (9.4). □

If $D = 1$, then $L_u = 1$ and (9.3) is an equality: a leaf passes its correlation on with the factor t^2, and iterating along a path recovers Example 9.2. For $D \geq 2$ the factor $L_u \geq 1$ is the price of frustration, the worst case in which the neighbours of u are fully magnetized against each other; no argument that, like this proof, holds for every flip-invariant measure in place of the Gibbs measure of H can lower it (Route warning 9.13).

Iterating Lemma 9.6 produces a branching and annihilating system of particles (Fig. 9.2). The weights of one step depend only on the magnitudes of the star, which are fresh at each step, so for iid couplings we may write $\rho_D := \sum_{A \subseteq [D],\, |A| \text{ odd}} \omega(A)$ for the total weight of a vertex with D remaining edges, and $\bar\rho_D := \max_{D' \leq D} \rho_{D'}$. Two bounds make these explicit. Since $\min_s \cosh(\cdot) \geq 1$, $L_u \leq \prod_i \cosh^2 K_i$ and $\omega(A) \leq \prod_{i\in A} \mathbb{E} \sinh^2 K \prod_{i\notin A} \mathbb{E} \cosh^2 K$; summing over odd A with $\mathbb{E}\cosh^2 K - \mathbb{E}\sinh^2 K = 1$ gives

$$\bar\rho_D \leq \tfrac{1}{2}\big[(\mathbb{E}\cosh 2\beta|J|)^D - 1\big] \quad \text{(every symmetric law)}, \qquad \rho_D = \frac{(\cosh 2\beta)^D - 1}{2\cosh^{2[D \text{ odd}]}\beta} \quad (\pm J), \quad (9.5)$$

the second because for $\pm J$ the minimum in L_u is 1 for even D and $\cosh\beta$ for odd D.

To sum over the target before expanding, we need row sums of multi-point functions. For an induced subgraph $H = G - W$ with $W \neq \emptyset$ and $T \subseteq V(H)$ with $|T|$ odd, put

$$\chi_H(T) := \sum_{v \in V(H)} \Xi_H(T \triangle \{v\}),$$

and call (H, T) *non-fresh* if every vertex of T is adjacent in G to a vertex of W. For $|T| = 1$, $\chi_H(T)$ is a row sum of squared two-point functions; branching creates the larger T.

Lemma 9.7 (Summed row recursion)**.** *For every such (H, T) and every $w \in T$,*

$$\chi_H(T) \leq \Xi_H(T \setminus \{w\}) + \sum_{A \text{ odd}} \omega_w(A)\, \chi_{H-w}((T \setminus \{w\}) \triangle \partial A),$$

the sum running over the odd subsets of the star of w in H. If (H, T) is non-fresh, so is every pair on the right, and $\sum_A \omega_w(A) \leq \bar{\rho}_{\Delta-1}$ when G has maximum degree at most Δ.

The proof is bookkeeping with (9.4); it is in Section A.18.

Theorem 9.8 (Spin-glass susceptibility at high temperature)**.** *Let the coupling law be symmetric with $\mathbb{E}\cosh(2\beta|J|) < \infty$, let $\Delta \geq 2$, and let G be a finite multigraph of maximum degree at most Δ, with free boundary conditions.*

(a) If $\bar{\rho}_{\Delta-1} < 1$, then $\sup_{u \in V} \sum_{v \in V} \mathbb{E}\langle \sigma_u \sigma_v \rangle_G^2 \leq 1 + \dfrac{\bar{\rho}_\Delta}{1 - \bar{\rho}_{\Delta-1}}$.

(b) If the law is strictly sub-Gaussian, $\mathbb{E}\cosh(\lambda J) \leq e^{\lambda^2 \mathbb{E}J^2/2}$ for all $\lambda \in \mathbb{R}$ ($\pm J$, Gaussian and uniform laws qualify), and $a := (\Delta - 1)\beta^2 \mathbb{E}J^2 < \frac{1}{2}\ln 3$, then

$$\sup_{u \in V} \sum_{v \in V} \mathbb{E}\langle \sigma_u \sigma_v \rangle_G^2 \leq 1 + \frac{e^{4a} - 1}{3 - e^{2a}}. \tag{9.6}$$

Proof. (a) Let M be the maximum of $\chi_H(T)$ over the finitely many non-fresh pairs (H, T) with H an induced subgraph of G. By Lemma 9.7 with $\Xi \leq 1$, $M \leq 1 + \bar{\rho}_{\Delta-1} M$, hence $M \leq 1/(1 - \bar{\rho}_{\Delta-1})$. For the root, the term $v = u$ is 1, and for $v \neq u$, (9.4) at u with $B = \{v\}$ gives $\sum_{v \neq u} \Xi_G(\{u, v\}) \leq \sum_A \omega_u(A) \chi_{G-u}(\partial A) \leq \bar{\rho}_\Delta M$, since $(G - u, \partial A)$ is non-fresh.

(b) Put $n = \Delta - 1$ and $s = \mathbb{E}\sinh^2 K = \frac{1}{2}(\mathbb{E}\cosh 2\beta|J| - 1)$. Strict sub-Gaussianity at $\lambda = 2\beta$ gives $1 + 2s \leq e^{2\beta^2 \mathbb{E}J^2} = e^{2a/n}$, so by (9.5), $\bar{\rho}_n \leq \frac{1}{2}(e^{2a} - 1) < 1$ and $\bar{\rho}_\Delta \leq \frac{1}{2}(e^{2a(n+1)/n} - 1) \leq \frac{1}{2}(e^{4a} - 1)$. Substituting into (a) gives (9.6). □

The cavity mechanism goes back to Chayes, Chayes, Sethna and Thouless [52] and to Fröhlich and Zegarlinski [97]. Comparable high-temperature bounds for diluted spin glasses on sparse random graphs are due to Guerra and Toninelli [112] and to Efthymiou and Zampetakis [83, (1.4)]; they are not uniform over graphs with short cycles.

On $\mathbb{Z}^d$, where $\Delta = 2d$, the right side of (9.6) depends on β and d only through $a = (2d-1)\beta^2 \mathbb{E}J^2$, so the bound is uniform in the dimension. It holds for $\beta < 0.7411\,((2d-1)\mathbb{E}J^2)^{-1/2}$, against the Bethe value $\beta \approx ((2d-1)\mathbb{E}J^2)^{-1/2}$ of Example 9.2; for instance $\chi_{\mathrm{SG}} \leq 2$ in every dimension when $a \leq 0.22284$. The centred susceptibility bound of Chapter 8 restores the edges of a star one at a time; averaging over the whole star at once, by Parseval, moves the constant from about 0.39 to 0.74 of the mean-field scale at large d. Disorder-averaged squared correlations were

bounded before by Berretti, with exponential decay for almost every J for Gaussian couplings at small β times their standard deviation, via Cauchy–Schwarz in the couplings in a two-replica graph expansion [27] (see [96, pp. 150–151]; for slowly decaying tails the terms need not be integrable), and by Fröhlich and Zegarlinski, with weak uniqueness and decay of $\mathbb{E}\langle\sigma_A;\sigma_B\rangle^2$ for long-range glasses at sufficiently high temperature [97].

Corollary 9.9 (No overlap at high temperature)**.** *Let $d \geq 1$ and let β and the law satisfy (a) or (b) of Theorem 9.8 with $\Delta = 2d$, with bound $\bar\chi$.*

(i) For boxes Λ_L with free boundary conditions and for tori $(\mathbb{Z}/L\mathbb{Z})^d$, $L \geq 2$: $\chi_{\mathrm{SG}}(\Lambda) \leq \bar\chi$ and $\mathbb{E}\langle q_\Lambda^2\rangle \leq \bar\chi/|\Lambda|$, where q_Λ is the overlap of two independent replicas.

(ii) Let $\mathbf{P}$ be any subsequential limit of the laws of (J,σ,τ) on the tori, with σ,τ independent samples of the torus Gibbs measure and everything extended periodically. Then $\sum_x \mathbf{E}[\sigma_0\tau_0\sigma_x\tau_x] \leq \bar\chi$, and the spatial overlap $Q = \lim_n q_{\Lambda_n}$ vanishes $\mathbf{P}$-almost surely.

Part (i) is Theorem 9.8 with $\mathbb{E}\langle q_\Lambda^2\rangle = \chi_{\mathrm{SG}}(\Lambda)/|\Lambda|$ (Chapter 1); part (ii) passes to the limit with the ergodic theorem and Fatou's lemma (Section A.18).

The same expansion, with a tilt that rewards long histories, gives exponential decay at fixed d, also for boundary conditions that do not depend on the couplings. For a box Λ and a boundary condition η outside it, let $\mu_{\Lambda,\eta}$ be the Gibbs measure in Λ and $\partial^{\mathrm{in}}\Lambda$ the set of sites of Λ with a neighbour outside.

Theorem 9.10 (Exponential $\boldsymbol{L^2}$ decay in fixed dimension)**.** *Let the coupling law be symmetric with $\mathbb{E}\cosh(2\beta|J|) < \infty$, let $d \geq 1$, and suppose $(\mathbb{E}\cosh 2\beta|J|)^{2d-1} < 3$; for $\pm J$ couplings it suffices that $c^{2d-1} - c < 2$ with $c = \cosh 2\beta$. There are $r > 1$ and $C < \infty$, depending only on d, β and the law, such that:*

(i) for every finite subgraph G of $\mathbb{Z}^d$ with free boundary conditions and every torus $G = (\mathbb{Z}/L\mathbb{Z})^d$, $L \geq 2$, $\mathbb{E}\langle\sigma_x\sigma_y\rangle_G^2 \leq C^2 r^{-\mathrm{dist}_G(x,y)}$;

(ii) for every box Λ and every boundary condition η independent of the couplings of the edges with an endpoint in Λ, with $\delta_x = \mathrm{dist}(x,\partial^{\mathrm{in}}\Lambda)$,

$$\mathbb{E}\,\mu_{\Lambda,\eta}(\sigma_x)^2 \leq C\,r^{-1-\delta_x}, \qquad \mathbb{E}\,\mu_{\Lambda,\eta}(\sigma_x\sigma_y)^2 \leq C^2 r^{-\min\{\mathrm{dist}(x,y),\,\delta_x+\delta_y+2\}}.$$

The $\pm J$ condition reads $\beta < 0.49065$ ($T > 2.0381$) for $d = 2$ and $\beta < 0.35781$ ($T > 2.7948$) for $d = 3$; for Gaussian couplings the general condition is $(2d-1)\beta^2 < \frac{1}{2}\ln 3$.

The proof, encoding the boundary condition as one extra vertex, is in Section A.19. By (ii), a metastate κ_J built from coupling-independent boundary conditions along boxes exhausting $\mathbb{Z}^d$ has $\mathbb{E}\int \Gamma(\sigma_0)^2\,\kappa_J(d\Gamma) = 0$, as $\Gamma \mapsto \Gamma(\sigma_0)^2$ is bounded and continuous. The condition of the theorem implies $\bar\rho_{2d-1} < 1$, and for $\pm J$ its form $c^{2d-1} - c < 2$ is equivalent to it (the exact ρ_D of (9.5) is nondecreasing in D); this is the hypothesis of Theorem 9.8(a), so in fixed d the susceptibility bound holds in the same range, which for $\pm J$ in $d = 2$ allows $(2d-1)\beta^2$ up to 0.722, while part (b) trades this range (0.549) for d-independent constants.

9.5 Certified windows in two and three dimensions

In low dimension the loss L_u, which assumes the neighbours of u fully magnetized, is the weak point. Two refinements recover part of it, at the cost of a computer certificate; the first takes the worst case over the boundary of a window around the star.

Lemma 9.11 (Layer cavity inequality)**.** *In the setting of Lemma 9.6, let $R \subseteq V(H)$ contain the y_i, and for a configuration η on the vertex boundary of R in H let μ_R^η be the Gibbs measure of H on R with boundary condition η. Then (9.3) holds with L_u replaced by*

$$L_R := \frac{\prod_i \cosh^2 K_i}{D_{\min}^2}, \qquad D_{\min} := \min_{\eta,\ \varepsilon}\ \mu_R^\eta(\cosh h).$$

If the y_i are distinct and $R = \{y_1, \dots, y_D\}$ is an independent set of H (as in every subgraph of $\mathbb{Z}^d$), then $L_R \le 1/\min_s x_s^2$ with $x_s = \frac{1}{2}\big[\prod_i(1 + s_i \tanh\alpha_i) + \prod_i(1 - s_i \tanh\alpha_i)\big]$, where $\tanh\alpha_i = t_i T_i$ and $T_i = \tanh\big(\sum_{f \ni y_i,\, f \neq e_i} K_f\big)$.

The proof, by the Markov property and multilinearity, is in Section A.20. Special cases of the closed form are standard local-energy identities (compare [200, eq. (7)]).

Since $\alpha_i < K_i$ whenever $T_i < 1$, the gain comes from the neighbours of u not being rigidly magnetized. The second refinement tracks the past: a particle created by an expansion sits next to deleted vertices, unavailable to its children and unmagnetized. A particle's *type* records its last k lattice moves; every type has finitely many children, and a *potential* $\varphi \ge 1$ on types with

$$\sum_{A \text{ odd} \subseteq \mathrm{avail}(\pi)} L(\pi)\, t^{2|A|} \prod_{m \in A} \varphi(\mathrm{child}(\pi, m)) \le \lambda\, \varphi(\pi) \qquad \text{for every type } \pi, \tag{9.7}$$

for some $\lambda < 1$ and type-dependent bounds $L(\pi)$ on the loss, makes the total weight of histories of length n at most $\lambda^n \varphi(\mathrm{root})$. A history joining x to y has at least $\mathrm{dist}(x, y)/(2d)$ steps.

Theorem 9.12 (Certified $\boldsymbol{L^2}$ windows for $\boldsymbol{\pm J}$ couplings)**.** *Let the couplings be iid $\pm J$. The following hold, by computer-assisted proofs.*

(a) Let $d = 2$ and $\beta \le \mathrm{artanh}(0.4806) = 0.523764$ ($T \ge 1.90926$), or $d = 3$ and $\beta \le \mathrm{artanh}(0.3505) = 0.366014$ ($T \ge 2.73214$). For every finite subgraph G of $\mathbb{Z}^d$ with free boundary conditions and every torus of side at least 11, $\mathbb{E}\langle\sigma_x\sigma_y\rangle_G^2 \le C\lambda^{\mathrm{dist}_G(x,y)/(2d)}$ with $\lambda = 0.9995$, $C \le 1.02 \cdot 10^4$ for $d = 2$ and $C \le 3.8 \cdot 10^3$ for $d = 3$.

(b) Let $d = 2$ and $\beta \le \mathrm{artanh}(0.4762) = 0.518058$ ($T \ge 1.93028$). For every rectangular box $\Lambda \subset \mathbb{Z}^2$ whose sides both have at least 8 sites and every boundary condition η independent of the couplings, with $\lambda = 0.9995$ and $\varphi_{\max} = 4.1487331$, writing $\delta'_x = \mathrm{dist}(x, \mathbb{Z}^2 \setminus \Lambda)$,

$$\mathbb{E}\,\mu_{\Lambda,\eta}(\sigma_x)^2 \le \frac{\varphi_{\max}}{1-\lambda}\,\lambda^{\delta'_x/4}, \qquad \mathbb{E}\,\mu_{\Lambda,\eta}(\sigma_x\sigma_y)^2 \le \frac{\varphi_{\max}^2}{1-\lambda}\,\lambda^{\min\{\mathrm{dist}(x,y),\,\delta'_x+\delta'_y\}/4}.$$

What the certificates check. With windows of radius 2 ($d = 2$) or 1 ($d = 3$ and (b)) in Lemma 9.11 and types of memory 5 ($d = 2$) or 3 ($d = 3$), the computer evaluates the losses $L(\pi)$ exactly, rounds a floating-point potential φ up to multiples of 10^{-9}, and verifies (9.7) for every type in exact rational arithmetic at $t_0 = 4806/10000$ ($d = 2$) and $3505/10000$ ($d = 3$), with largest ratio 0.99930010. The reduction of every subgraph, coupling configuration and $\beta \le \beta_0$ to this check is proved in Section A.20 (see also Appendix B).

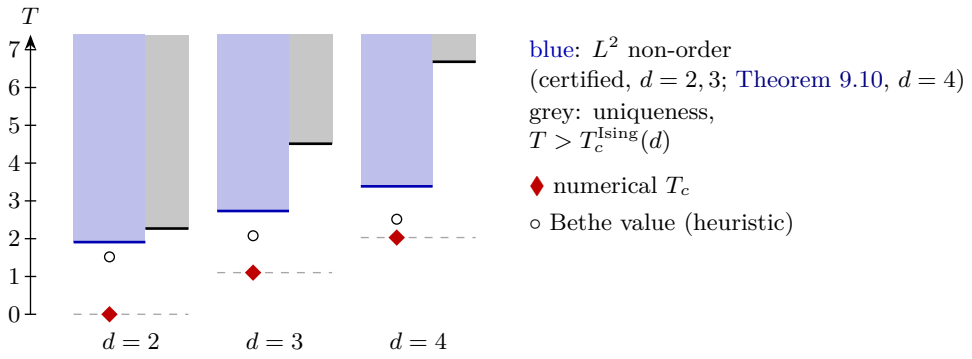


Figure 9.3: Proved high-temperature regions for $\pm J$ couplings. Grey: uniqueness, $T > T_c^{\mathrm{Ising}}(d)$ (exact in $d = 2$, numerical in $d \geq 3$), enlarged by De Santis by an unquantified amount; blue: L^2 non-order; diamonds: numerical T_c (0, 1.10, 2.03 [127, 155]); circles: the Bethe value (Example 9.2). Between diamonds and blue the numerics see a paramagnet but no current method proves the absence of order: in $d = 2$, everything below $T \approx 1.91$.

The windows lie strictly inside the interval that Observation 9.4 leaves to sign-using methods, between Newman's $T > 2.269$ ($d = 2$), $T > 4.51$ ($d = 3$) and the numerical $T_c = 0$, $T_c = 1.1019(29)$ [116, 127, 233] (Fig. 9.3). For Gaussian couplings the same method certifies $\beta \leq 0.4374$ on $\mathbb{Z}^2$ and $\beta \leq 0.3338$ on $\mathbb{Z}^3$; in $d = 2$ this is *weaker* than Newman's explicit sufficient condition $\mathbb{E}(1 - e^{-2\beta|J|}) < \frac{1}{2}$, which holds up to $\beta \approx 0.5438$, so there the sign-using method does not reach beyond the comparison with the ferromagnet. The windows stop where they do because the layer loss is a worst case over the window signs, the loss of the gauge ferromagnet, while particles move with fair-sign weights t^2; for $\pm J$ in $d = 2$ that ferromagnet orders at $\beta > 0.4407$, so the loss saturates and floating-point thresholds level off near $\beta \approx 0.53$ (numerical, not a theorem).

9.6 Where the cavity method stops

Route warning 9.13 (The cavity method stops at $\frac{1}{2}\ln 3$). *Tempting step:* sharpen the losses until the susceptibility bound reaches the Bethe value $(2d-1)\mathbb{E}\tanh^2(\beta J) = 1$, or the numerical T_c. *Status: fails for every bound of the form* (9.3) *with a constant smaller than* L_u *that holds for all flip-invariant cavity measures (proved).* Since the proof of Lemma 9.6 uses the Gibbs measure μ of H only through $\langle \sigma_u \sigma_B \rangle_G = \mu(\sigma_B \sinh h)/\mu(\cosh h)$ and its independence of the star signs, (9.3) in this form holds for every probability measure μ on the spins of H (*cavity measures*). In every degree D, for all $K_i > 0$, distinct y_i and $B = \{b\}$ with $b \notin \{y_1, \ldots, y_D\}$, the supremum over flip-invariant ones of the ratio of the left side of (9.3) to its sum over A (without L_u) equals L_u (Lemma A.8): no inequality of the form (9.3) with a constant smaller than L_u holds for all of them. With the branching count this caps the method, as $d \to \infty$ at fixed $a = (2d-1)\beta^2\mathbb{E}J^2$, at $e^a \sinh a = 1$, that is $a = \frac{1}{2}\ln 3$: the denominator contributes e^a, and the odd-subset branching $\sinh a$ where self-avoiding paths would give a (removing only the first gives $a < 0.881$, only the second $a < 0.567$, both $a < 1$, the Bethe value). Reaching it needs an inequality that averages the cavity couplings inside the denominator and absorbs the four-point functions created by $|A| \geq 3$; none is known.

Remark 9.14 (L^2 non-order is not uniqueness). The results of Sections 9.4 and 9.5 average over all fair signs, including those on edges to the boundary, so they need coupling-independent

boundary conditions (one chosen from nearby couplings, as a pure state would be, breaks the average). They say nothing about uniqueness below the reach of Theorem 9.5, nor give the L^2 boundary independence of even observables that Fröhlich and Zegarlinski obtain at sufficiently high temperature for their couplings [97].

9.7 The Kac limit

Another approach from above makes the interaction long but finite. In the Kac model of Franz and Toninelli [92] each spin interacts with about γ^{-d} others through couplings of variance of order γ^d, so on the scale γ^{-1} the model looks like the Sherrington–Kirkpatrick (SK) model $H^{\mathrm{SK}} = -(2|\Lambda|)^{-1/2}\sum_{i,j\in\Lambda} J_{ij}\sigma_i\sigma_j - h\sum_i \sigma_i$. Precisely, on $\mathbb{Z}^d$,

$$H_\Lambda^{(\gamma)}(\sigma) = -\sum_{i,j\in\Lambda}\sqrt{\frac{w(i-j;\gamma)}{2W(\gamma)}}\,J_{ij}\sigma_i\sigma_j - h\sum_{i\in\Lambda}\sigma_i, \qquad w(r;\gamma)=\gamma^d\phi(\gamma r), \quad W(\gamma)=\sum_{i\in\mathbb{Z}^d} w(i;\gamma),$$

summed over ordered pairs, with $\phi \geq 0$ smooth and fast decaying and J_{ij} iid $N(0,1)$. Let $f^{(\gamma)}(\beta,h) = -\lim_{L\to\infty}(\beta|\Lambda|)^{-1}\mathbb{E}\ln Z_\Lambda^{(\gamma)}$, which exists, and let f^{SK} be the SK free energy.

Theorem 9.15 (Franz–Toninelli [92, Thm. 1])**.** *Let $\sum_{i\in\mathbb{Z}^d}\phi(i) < \infty$. Then for all β and h, $\limsup_{\gamma\to 0} f^{(\gamma)}(\beta,h) \leq f^{\mathrm{SK}}(\beta,h)$. If in addition the Fourier transform of ϕ is nonnegative, then $f^{(\gamma)} \geq f^{\mathrm{SK}}$ for every γ, a bound due to Guerra and Toninelli (as recalled in [92]), and therefore $\lim_{\gamma\to 0} f^{(\gamma)}(\beta,h) = f^{\mathrm{SK}}(\beta,h)$.*

The first inequality is a Guerra-type interpolation with SK blocks of side $1 \ll \ell \ll \gamma^{-1}$ [92, (13)]. With Talagrand's proof of the Parisi formula, it follows that the Parisi theory gives the correct free energy in the Kac limit [92, p. 2]; with small random perturbations added, the local overlap distribution on the scale γ^{-1} converges to that of the perturbed SK model for almost every value of the perturbation parameters [91, Thm. 2].

Route warning 9.16 (From the Kac limit to long-range order). *Tempting step:* use local mean-field behaviour at scale γ^{-1}, for small fixed γ, to prove spin-glass order at large scales. *Status: fails without a new input, even granting the Kac-limit theorems.* As the authors note, local closeness to mean field has, as in non-disordered models, no implications for long-range order, and nothing can be said a priori for the unperturbed model ($\varepsilon \equiv 0$ in their notation) [91, p. 8]. For ferromagnetic Kac models the passage from local to global order is a contour argument around the mean-field minimizer; the spin-glass analogue, a large-deviation theory for overlap profiles, is outlined only informally [91, Sect. V]; we know of no rigorous version.

9.8 What is established, and what is open

Sign-blind comparison gives uniqueness for every coupling configuration above the ferromagnetic critical temperature of $|J|$ and no further, up to De Santis's strict gain and Fröhlich and Imbrie's extension to unbounded Gaussian couplings. Sign averaging gives L^2 non-order much further down, but the cavity bound in its present form cannot pass $\frac{1}{2}\ln 3$ at large d, and the Kac limit says nothing about large scales. The threads resume in Chapters 8, 10, 11 and 13.

Open problem 9.17 (The Bethe value in high dimension)**.** Show that for every $\varepsilon > 0$ there is $d_0(\varepsilon)$ such that for $d \geq d_0$ and iid symmetric couplings with an exponential moment, $(2d -$

$1)\mathbb{E}\tanh^2(\beta J) \le 1-\varepsilon$ implies $\sup_L \chi_{\mathrm{SG}}(\Lambda_L) < \infty$ with free boundary conditions. A sufficient statement would be a one-vertex inequality of the form of Lemma 9.6 with L_u replaced by $1+\delta_\Delta$, $\delta_\Delta \to 0$, for the two-point functions alone.

Open problem 9.18 (All temperatures in two dimensions)**.** For iid $\pm J$ or Gaussian couplings on $\mathbb{Z}^2$, prove $\mathbb{E}\langle\sigma_x\sigma_y\rangle_\Lambda^2 \to 0$ as $|x-y| \to \infty$, uniformly in boxes with free boundary conditions, at every $T > 0$. This is the L^2 form of Newman's Research Problem 4.2 [171, p. 48]; current methods stop at $T \approx 1.91$.

Open problem 9.19 (Griffiths singularities for spin glasses)**.** For the EA model on $\mathbb{Z}^d$ with Gaussian couplings, prove or disprove that the quenched free energy is not analytic in β (or in a uniform field h at $h = 0$) at some β in the uniqueness region.

Chapter 10

The random-field Ising model as a mirror

The random-field Ising model is the disordered spin system whose order question has been answered. In two dimensions an arbitrarily weak random field destroys the order of the Ising ferromagnet at every temperature; in three and more dimensions a weak field leaves it intact at low temperature. Both halves are theorems, and the recent proofs are short enough to teach. This chapter asks which of their steps would carry over to the Edwards–Anderson model, and why the ones that matter do not.

The answer has a clean shape. The random-field proofs combine *probabilistic modules* (Gaussian concentration and chaining for free-energy differences, counting of lattice animals, and the Aizenman–Wehr comparison of volume fluctuations with surface terms) with *ferromagnetic seeds* (a deterministic Hamiltonian with known ordered phases and a surface tension, extremal plus and minus states, the FKG inequality and monotone couplings). The modules transfer to the EA model and have already produced EA theorems: the finiteness of unsatisfied dual clusters (Chapter 5), the two-dimensional results on bond observables (Chapter 4) and the chaos theorems (Chapter 6). The seeds do not, and for the central one the failure is an identity. The joint flip of spins and disorder that drives the modern Peierls proofs for the random field is, for the EA model, exactly a gauge transformation, which leaves the joint law of couplings and spins invariant (Proposition 10.8).

By the end of the chapter the reader can prove the joint-flip identities in both models, state the random-field theorems with the idea behind each, and read off from one table (Table 10.1) which step fails for the EA model, and how strongly.

10.1 The model and the Imry–Ma picture

Definition 10.1 (Random-field Ising model)**.** Let $d \geq 1$, $\Lambda \subset \mathbb{Z}^d$ finite, $J > 0$, a uniform field $h \in \mathbb{R}$, a disorder strength $\varepsilon \geq 0$, and iid random fields $(\eta_x)_{x \in \mathbb{Z}^d}$ (standard Gaussian unless stated). With plus or minus boundary conditions,

$$H^{\pm}_{\Lambda,\eta}(\sigma) = -J \sum_{xy \in E(\Lambda)} \sigma_x \sigma_y \;\mp\; J \sum_{\substack{xy \in \partial_e \Lambda \\ x \in \Lambda}} \sigma_x - \sum_{x \in \Lambda} (h + \varepsilon \eta_x)\, \sigma_x,$$

and $\mu^{\pm}_{\Lambda,\eta} \propto e^{-\beta H^{\pm}_{\Lambda,\eta}}$, with $T = 1/\beta$. The *boundary influence* at the origin is $m_{\Lambda,\eta} = \mu^{+}_{\Lambda,\eta}(\sigma_0 = 1) - \mu^{-}_{\Lambda,\eta}(\sigma_0 = 1)$, and $m(L) = \mathbb{E}\, m_{\Lambda(L),\eta}$ for the graph-distance ball $\Lambda(L)$ of radius L, so that $m(L) = \frac{1}{2}(\mathbb{E}\langle\sigma_0\rangle^{+}_{\Lambda(L)} - \mathbb{E}\langle\sigma_0\rangle^{-}_{\Lambda(L)})$ [10, (1.5)].

The interaction is ferromagnetic for every field configuration, so the FKG inequality holds for every η, the measures $\mu^{\pm}_{\Lambda,\eta}$ are monotone in the boundary condition, and their infinite-volume limits $\mu^{\pm}_{\eta}$ exist and are the largest and smallest Gibbs states [37, Sect. 2], [8, Sect. 4b]. Uniqueness of the Gibbs state is therefore equivalent to $\mu^{+}_{\eta} = \mu^{-}_{\eta}$, and to $\lim_L m(L) = 0$.

The heuristic is that of Imry and Ma [123]. "In a region of linear dimension L, the typical fluctuation of magnetic field energy should be of the order of $L^{d/2}$, while the energy of a domain wall surrounding the region is of the order of L^{d-1}" [122, p. 146]. In $d \le 2$ the field wins on large scales, and in $d \ge 3$ the wall does; Chapter 4 draws the comparison. One identity, used repeatedly below, says which observable the disorder couples to.

Lemma 10.2 (Conjugate observables)**.** *For every finite Λ and boundary condition b,*

$$\frac{\partial}{\partial \eta_x} \log Z^{b}_{\Lambda,\eta} = \beta\varepsilon\, \langle\sigma_x\rangle^{b}_{\Lambda,\eta} \quad \textit{(random field)}, \qquad \frac{\partial}{\partial J_e} \log Z^{b}_{\Lambda,J} = \beta\, \langle\sigma_x\sigma_y\rangle^{b}_{\Lambda,J} \quad \textit{(EA, } e = xy\textit{)}.$$

In particular $\partial_{\eta_x} \log(Z^{+}_{\Lambda,\eta}/Z^{-}_{\Lambda,\eta}) = \beta\varepsilon(\langle\sigma_x\rangle^{+} - \langle\sigma_x\rangle^{-}) \ge 0$ by FKG. The observable σ_x is odd under the global flip; the observable $\sigma_x\sigma_y$ is even.

Proof. Differentiate $\log \sum_\sigma e^{-\beta H(\sigma)}$; the sign is the monotonicity of $\mu^{\pm}$ in the boundary condition. □

In the random-field model the free-energy difference between the two boundary conditions is a monotone function of the field whose derivative is the disagreement of the magnetizations. This is what lets the Aizenman–Wehr argument conclude something about σ_x, and through FKG about the Gibbs states. In the EA model the derivative has no sign, which is the cancellation problem of Chapter 4, and the observable it controls cannot tell a state from its flip.

10.2 Two dimensions: no order, and how fast it is lost

The general rounding theorem of Aizenman and Wehr is stated and discussed in Chapter 4; for the random field it gives uniqueness.

Theorem 10.3 (Aizenman–Wehr [8, Cor. 4.3 with Thm. 4.1])**.** *Let $d \le 2$, and let the fields be iid with mean 0, variance 1 and $\mathbb{E}e^{s\eta_x} < \infty$ for all real s. For every $T > 0$, $h \in \mathbb{R}$ and $\varepsilon \neq 0$, the nearest-neighbour random-field Ising model has a unique Gibbs state for almost every field configuration.*

At $T = 0$ the same holds for ground states when the field law is absolutely continuous [8, Thm. 4.1]. The proof has two parts. The first, valid for any disordered system and proved in Chapter 4, says that in $d \le 2$ the translation averages of the observable conjugate to the disorder agree in all Gibbs states. The second is specific: "In RFIM the Fortuin–Kasteleyn–Ginibre (FKG) inequality is satisfied for an arbitrary field configuration", so the plus and minus states are the highest and lowest states, the differences of their magnetizations are nonnegative, and equal averages force them to coincide [8, p. 500]. The second part is exactly what the EA model lacks.

The rate at which order is lost was settled in three steps.

Theorem 10.4 (Aizenman–Peled; Aizenman–Harel–Peled; Ding–Xia)**.** *Let $d = 2$, let the fields be iid standard Gaussian, $T \geq 0$, $h \in \mathbb{R}$ and $\varepsilon > 0$.*

(a) For a finite-range interaction there are $C, \gamma > 0$, depending on the interaction, T and ε, with $m(L) \leq C/L^{\gamma}$ for all large L [6, Thm. 1.1].

(b) For the nearest-neighbour interaction of strength J, $m(L) \leq C(J/\varepsilon)\, e^{-c(J/\varepsilon)L}$ with $c, C > 0$ that do not depend on T, h or L [10, Thm. 1.1].

(c) For $h = 0$ and every $T \in [0, \infty)$ there is $c = c(\varepsilon, T) > 0$ with $\mathbb{E}(\langle\sigma_0\rangle^{+}_{\Lambda_N} - \langle\sigma_0\rangle^{-}_{\Lambda_N}) \leq c^{-1}e^{-cN}$ for all $N \geq 1$ [75, Thm. 1.1].

Proof sketch. We omit the tortuosity estimates and the bootstrap. The common starting point is monotonicity: by FKG, $\langle\sigma_0\rangle^{+} - \langle\sigma_0\rangle^{-} \geq 0$ for every field [6, (1.8)], and the difference is a disagreement probability. At $T = 0$ the two ground states satisfy $\sigma^{\Lambda,+} \geq \sigma^{\Lambda,-}$ [75, (3)]; at $T > 0$ Ding and Xia use a monotone coupling of the two measures, and Aizenman, Harel and Peled a continuous extension of the spins to the edges, with two independent samples, in which the origin must be joined to the boundary by a path of disagreement [10, Sect. 1.2], [75, p. 2]. The field in a box is then used "to fight the influence of the boundary condition" [75, p. 2]: by Lemma 10.2 the free-energy difference of the two boundary conditions grows with the field at a rate given by the disagreement, and its size is limited by the boundary. Aizenman and Peled quantify the Aizenman–Wehr dichotomy to get a power law; the two exponential bounds add the observation that disagreement paths are tortuous, in the sense of Aizenman and Burchard, and a bootstrap from fast power-law decay to exponential decay [10, Sect. 1.2], [75, p. 2]. □

At weak disorder the length scale on which the two-dimensional order is lost is very large. For $T = 0$ and fixed $\theta \in (0, 1)$, the correlation length $\min\{N : \mathbb{E}\, m_{\Lambda_N,\eta} \leq \theta\}$ scales as $e^{\varepsilon^{-4/3+o(1)}}$, the upper bound holding for all $T \geq 0$ [74], as summarized in [76, p. 2], and a matching lower bound holds for small $T > 0$ [76, Thm. 1.2]. The exponent $4/3$ comes from the greedy lattice animal normalized by its boundary size, which grows like a power of $\log N$ in $d = 2$ and stays bounded in $d \geq 3$ [76, p. 4]. That quantity returns in the next section.

10.3 Three and more dimensions: order by a Peierls map

Theorem 10.5 (Imbrie; Bricmont–Kupiainen)**.** *Let the fields be iid centred Gaussian with variance ε^2 in the normalization of the sources, and $h = 0$.*

(a) ($d = 3$, $T = 0$.) There is $C > 0$ such that for ε small, every site i and every rectangular box Λ centred at the origin, the plus ground state satisfies $\mathbb{P}(\sigma_i^{\min}(\Lambda^{+}) = -1) \leq \exp(-C/\varepsilon^2)$ [122, Thm. 1.1].

(b) ($d \geq 3$, $T > 0$.) For β large and ε small, $\langle\sigma_0\rangle^{+} \geq 1 - \exp(-O(\beta))$ with probability at least $1 - \exp(-O(\varepsilon^{-2}))$, and connected correlations decay exponentially [37, Thm. 2.1].

Imbrie's proof is a multiscale contour induction that uses coarse-grained contours, following Chalker and Fisher, Fröhlich and Spencer [44, 89]: "one obtains that contours of a certain size are rare, assuming only that smaller contours are rare and that they can be neglected" [122, p. 146]. Bricmont and Kupiainen use an exact renormalization group in the space of contour models and field distributions, under which "the system flows to the zero temperature, zero

field fixed point" [37, p. 542]. They also note that the free energy is C^∞ but analyticity "is not expected to hold, because of Griffiths' singularities" [37, Remark 2.2], the phenomenon of Chapter 9. Ding and Zhuang found a short proof.

Theorem 10.6 (Ding–Zhuang [76, Thm. 1.1], reproving Imbrie and Bricmont–Kupiainen)**.** *Let $d \geq 3$, $J = 1$, $h = 0$ and iid standard Gaussian fields. There is $c > 0$ such that for all $0 \leq T \leq c$, $0 \leq \varepsilon \leq c$ and $N \geq 1$, with probability at least $1 - e^{-c/T} - e^{-c/\varepsilon^2}$, $\mu^+_{\Lambda_N,\eta}(\sigma_0 = 1) \geq 1 - e^{-c/T} - e^{-c/\varepsilon^2}$.*

Ding, Liu and Xia extended this to the whole low-temperature phase of the pure model: for $d \geq 3$, $T < T_c$ and $\gamma > 0$ there is $\varepsilon_0(d, T, \gamma) > 0$ such that $\mathbb{P}(m_{\Lambda_N,\eta} \geq \varepsilon_0) \geq 1 - \gamma$ for all N when $\varepsilon \leq \varepsilon_0$ [77, Thm. 1.1].

The engine of Theorem 10.6 is a Peierls map on the joint space of fields and spins. Its proof has three ingredients, the first two of which we prove. Let $\mathbf{P}^+$ be the joint law of (η, σ) in which η is iid standard Gaussian with density ϕ and, given η, $\sigma \sim \mu^+_{\Lambda,\eta}$; its density is $f^+(\eta,\sigma) = \prod_{u\in\Lambda} \phi(\eta_u)\, e^{-\beta H^+_{\Lambda,\eta}(\sigma)}/Z^+_{\Lambda,\eta}$. For $A \subseteq \Lambda$ let η^A and σ^A be η and σ reversed on A, and let ∂A be the set of edges with exactly one endpoint in A.

Proposition 10.7 (The joint flip in the random-field model [76, (16) and Lemma 3.1])**.**

(i) If $h = 0$ and $\sigma_u\sigma_v = -1$ for every edge $uv \in \partial A$ (with $\sigma \equiv +1$ outside Λ), then

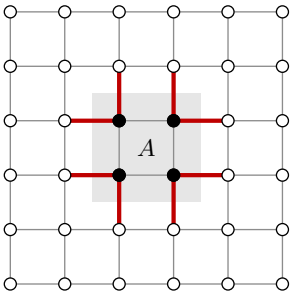
$$\frac{f^+(\eta,\sigma)}{f^+(\eta^A,\sigma^A)} = e^{-2\beta J|\partial A|}\, \frac{Z^+_{\Lambda,\eta^A}}{Z^+_{\Lambda,\eta}}.$$

(ii) The free-energy difference $\Delta_A(\eta) = T\log Z^+_{\Lambda,\eta^A} - T\log Z^+_{\Lambda,\eta}$ satisfies $\mathbb{E}[\Delta_A \mid \eta_{A^c}] = 0$ and $\mathbb{P}(|\Delta_A| \geq \lambda \mid \eta_{A^c}) \leq 2\exp(-\lambda^2/(8\varepsilon^2|A|))$.

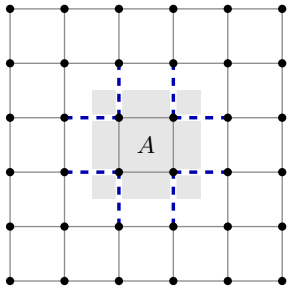
Proof. (i) Since $h = 0$, the field term is invariant: $(\varepsilon\eta^A_u)\sigma^A_u = \varepsilon\eta_u\sigma_u$ for every u. Edges inside A or outside A keep $\sigma_u\sigma_v$, and each edge of ∂A goes from -1 to $+1$, so $H^+_{\Lambda,\eta^A}(\sigma^A) = H^+_{\Lambda,\eta}(\sigma) - 2J|\partial A|$; the Gaussian density is even. (ii) Given η_{A^c}, the vectors η and η^A have the same law and Δ_A changes sign under $\eta \leftrightarrow \eta^A$, so its conditional mean is 0. By Lemma 10.2, $\partial_{\eta_v}\Delta_A = -\varepsilon(\langle\sigma_v\rangle_{\eta^A} + \langle\sigma_v\rangle_\eta)$ for $v \in A$, so Δ_A is $2\varepsilon\sqrt{|A|}$-Lipschitz in η_A, and Gaussian concentration gives the tail. □

Proof sketch of Theorem 10.6. We omit the coarse-graining estimate. Under the map of Proposition 10.7, the spin part gains the deterministic factor $e^{-2\beta|\partial A|}$, the field term is unchanged, and the price is the ratio of partition functions, $e^{\beta\Delta_A}$. By (ii), Δ_A has fluctuations of order $\varepsilon|A|^{1/2}$, and by the isoperimetric inequality $|A| \leq c_d|\partial A|^{d/(d-1)}$ this is $o(|\partial A|)$ for large A exactly when $d \geq 3$: the Imry–Ma count. The supremum over the exponentially many simply connected $A \ni 0$ is controlled by chaining and coarse graining, as for the greedy lattice animal normalized by its boundary: $\mathbb{P}(\sup_A |\Delta_A|/|\partial A| \geq 1) \leq e^{-c/\varepsilon^2}$ for $d \geq 3$ and $\varepsilon \leq c$ [76, (17)], adapting [44, 89]. On the complementary event each contour around the origin costs at least $e^{-\beta|\partial A|}$, and summing over contours with the count $|\{A : |\partial A| = n\}| \leq (2dn)^d(16d^3)^{2n}$ [76, (15)] gives the theorem for small T. □

The decisive line is (i): the ferromagnetic bonds crossing ∂A supply a gain $2J|\partial A|$ that does not depend on the disorder, and the disorder enters only through a free-energy difference with



(a) random field: flip σ and η on A; the 8 red bonds become satisfied: gain $2J|\partial A|$



(b) EA: flip σ on A and J on ∂A (blue); every edge keeps its status: gain 0

Figure 10.1: The joint flip in the two models. (a) In the random-field model the spins and the fields on A are reversed; the field energy is unchanged and the ferromagnetic bonds across ∂A, unsatisfied before the flip, are satisfied after it, which is the Peierls gain. (b) In the EA model the analogous map reverses the spins on A and the couplings on ∂A. Every term $J_{xy}\sigma_x\sigma_y$ is unchanged, the partition function is unchanged, and for a symmetric law the coupling distribution is unchanged: there is nothing to gain and nothing to pay.

mean zero, which the probabilistic modules control. Ding, Liu and Xia keep this map and replace the deterministic gain, which is too weak just below T_c, by a comparison along the map with the pure Ising model: "It seems the only plausible route here is to use $\mu_0^+(\sigma_o = 1) > 1/2$ for $T < T_c$ (which is essentially the definition of T_c) and to make comparisons between what happens along the Peierls mapping for the Ising model without and with external field" [77, p. 4], with FK-Ising clusters, coarse graining and spatial mixing doing the rest.

10.4 The mirror image: the joint flip is a gauge transformation

For the EA model the natural joint map flips the spins on A and the couplings on ∂A. It is the gauge transformation of Chapter 1 with $\tau = -1$ on A (Fig. 10.1).

Proposition 10.8 (The joint flip is a gauge transformation)**.** *Let the coupling law ν be symmetric, let $\Lambda \subset \mathbb{Z}^d$ be finite, let b be a boundary condition and $\beta \geq 0$, and let $\mathbf{P}^b$ be the joint law of (J, σ) in which the couplings of the edges with an endpoint in Λ are iid with law ν and, given J, $\sigma \sim \mu_{\Lambda,J}^{\beta,b}$. For $A \subseteq \Lambda$ let $J^{\partial A}$ be J with the signs of the couplings on ∂A reversed (including the edges from A to $\mathbb{Z}^d \setminus \Lambda$) and σ^A be σ reversed on A. Then:*

(i) $H^b_{\Lambda,J^{\partial A}}(\sigma^A) = H^b_{\Lambda,J}(\sigma)$ *for every σ, and* $Z^b_{\Lambda,J^{\partial A}} = Z^b_{\Lambda,J}$*;*

(ii) the map $(J, \sigma) \mapsto (J^{\partial A}, \sigma^A)$ preserves $\mathbf{P}^b$; equivalently, the ratio of the joint weights of (J, σ) and of its image is identically 1, and the analogue of Δ_A, $\log Z^b_{\Lambda,J^{\partial A}} - \log Z^b_{\Lambda,J}$, vanishes identically;

(iii) $\mathbf{P}^b(\sigma_x = -1) = \frac{1}{2}$ *for every $x \in \Lambda$ and every boundary condition b, the plus boundary condition included.*

Proof. Let $\tau = -1$ on A and $\tau = +1$ elsewhere. Then $J^\tau_{xy} = J_{xy}\tau_x\tau_y$ reverses exactly the couplings with one endpoint in A, so $J^\tau = J^{\partial A}$, and $\sigma\tau = \sigma^A$. The gauge lemma of Chapter 1

gives $H^b_{\Lambda,J^\tau}(\sigma\tau) = H^b_{\Lambda,J}(\sigma)$ (since $\tau = 1$ off Λ), and summing over the bijection $\sigma \mapsto \sigma\tau$ gives the equality of partition functions. Reversing the signs of independent symmetric couplings preserves their product law, so $\mathbf{P}^b(dJ, \sigma) = \nu^{\otimes}(dJ)\, \mu^{\beta,b}_{\Lambda,J}(\sigma)$ is mapped to itself. For (iii) take $A = \{x\}$: the map exchanges $\{\sigma_x = -1\}$ and $\{\sigma_x = 1\}$. □

So the event whose probability the joint Peierls argument bounds by a small number, $\sigma_0 = -1$ under plus boundary conditions, has annealed probability exactly $\frac{1}{2}$ in the EA model. This is the statement $\mathbb{E}\langle\sigma_x\rangle = 0$ of Chapter 1 seen through the Peierls map. The quantity that could signal EA order is gauge invariant, an overlap or a squared correlation, and the joint map acts trivially on it.

Remark 10.9 (Biased laws). For a law that is not symmetric the map still preserves the Hamiltonian and the partition function, but it changes the disorder density by $\prod_{e\in\partial A} \nu(-J_e)/\nu(J_e)$. For the biased law $\nu_p = p\delta_{+1} + (1-p)\delta_{-1}$ and $e^{-2\beta_{\mathrm{N}}} = (1-p)/p$ this factor is $\exp(-2\beta_{\mathrm{N}} \sum_{e\in\partial A} J_e)$: a surface term made of the couplings on ∂A, at the Nishimori inverse temperature. It is the same sum $\sum_{e\in\partial A} J_e$ that controls Peierls arguments with random contours for biased disorder (Chapter 14).

The other ferromagnetic seeds fail for reasons that are almost as short. First, there is no unperturbed Hamiltonian.

Observation 10.10 (No distinguished configuration). Let ν be symmetric with $\mathbb{E}|J_e| < \infty$ and Λ finite with free boundary conditions. Then $\mathbb{E}H_J(\sigma) = 0$ for every σ, and for every $\tau \in \{\pm1\}^\Lambda$ the random functions $\sigma \mapsto H_J(\sigma\tau)$ and $\sigma \mapsto H_J(\sigma)$ have the same law.

Proof. $\mathbb{E}J_e = 0$; and $H_J(\sigma\tau) = H_{J^\tau}(\sigma)$ with J^τ equal in law to J (Chapter 1). □

In the random-field model $H = H^{(0)} + \varepsilon H^{(1)}$, where $H^{(0)}$ is the Ising ferromagnet, with two ground states and a surface tension, and $\varepsilon H^{(1)}$ is a small random perturbation; Imbrie, Bricmont–Kupiainen and Ding–Zhuang all expand around $H^{(0)}$. For symmetric EA couplings the whole Hamiltonian is disorder, and the energy landscape looks statistically the same from every configuration. A deterministic part appears only for biased laws, and then it is a ferromagnet. Fröhlich and Imbrie's low-temperature theorem is of this kind: for $d \geq 2$, iid Gaussian couplings of mean 1 and variance less than $(\kappa\beta)^{-1}$, β large and plus boundary conditions, there is, with probability one, nonzero spontaneous magnetization, and connected correlations have tree decay with rate at least a constant times β [96, Thm. 8.2]. The order it produces is ferromagnetic.

Second, there is no monotone coupling. Holley's inequality and the monotone couplings behind Theorems 10.3 and 10.4 need a partial order on configurations in which the Gibbs measures are monotone. For a frustrated EA system no order obtained from the coordinatewise one by reversing some coordinates will do; we give a short argument.

Lemma 10.11 (No monotone order under frustration)**.** *Let G be a finite simple graph, J couplings with $J_e \neq 0$, $\beta > 0$ and μ the Gibbs measure with free or fixed boundary conditions (or any single-site fields). For $\tau \in \{\pm1\}^V$ order configurations by $\sigma \leq_\tau \sigma'$ if $\tau_x\sigma_x \leq \tau_x\sigma'_x$ for all x. If μ satisfies the FKG lattice condition for $\leq_\tau$, then $J_{xy}\tau_x\tau_y > 0$ for every edge; hence if some cycle of G is frustrated, μ satisfies the FKG lattice condition for no $\leq_\tau$.*

Proof. Write $s = \tau\sigma$, so that $J_{xy}\sigma_x\sigma_y = J^\tau_{xy}s_xs_y$ and fields stay single-site. Fix an edge xy and two configurations that agree off $\{x, y\}$ and have $(s_x, s_y) = (+,-)$ and $(-,+)$; their maximum and minimum for $\leq_\tau$ have $(+,+)$ and $(-,-)$. In the ratio $\mu(\sigma\vee\sigma')\mu(\sigma\wedge\sigma')/(\mu(\sigma)\mu(\sigma'))$ every term involving x or y but not both cancels, and the edge xy contributes $e^{4\beta J^\tau_{xy}}$. The lattice

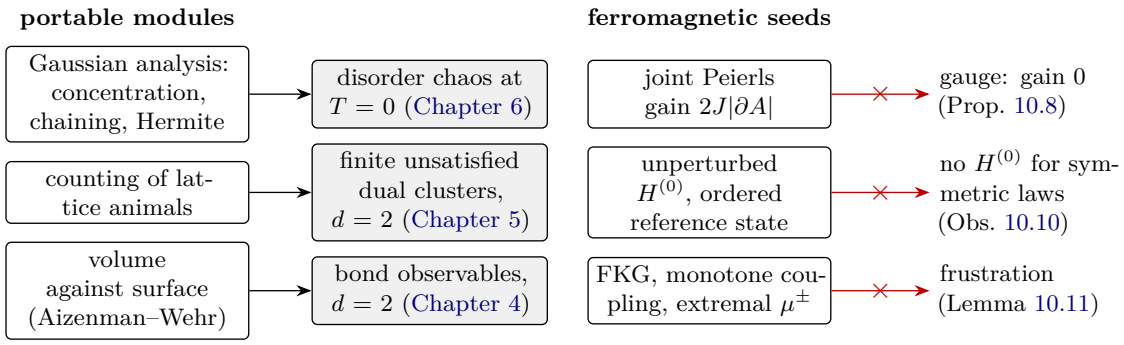


Figure 10.2: The two kinds of ingredient in the random-field proofs. Left: probabilistic modules, which use only Gaussian or product structure of the disorder and have already given EA theorems. Right: ferromagnetic seeds, which supply the order-producing comparison; each fails for the EA model, by an identity, by the absence of a deterministic part, or by frustration. Each random-field proof of order, or of exponential decay at weak disorder, discussed in this chapter uses at least one seed.

condition requires this to be at least 1, so $J^\tau_{xy} \geq 0$, and $J^\tau_{xy} \neq 0$. If $J^\tau > 0$ on every edge, then around any cycle $\prod \operatorname{sgn} J_e = \prod \operatorname{sgn} J^\tau_e = 1$, so no cycle is frustrated. The same computation with one site shows that $\mu(\tau_x \sigma_x = 1 \mid \sigma_{V\setminus\{x\}})$ is nondecreasing for $\leq_\tau$ only if $J^\tau_{xz} \geq 0$ for every neighbour z of x; so under frustration no $\leq_\tau$ makes the Gibbs measures monotone in the boundary condition either, already for single-site volumes. □

Unordered couplings of two boundary conditions exist, as in the disagreement percolation of van den Berg and Maes [236], but they give uniqueness only when a disagreement process is subcritical, a condition of high-temperature type. Third, the obvious graphical substitutes for the ordered Ising phase do not signal order: the frustrated random-cluster measure of Chapter 7 percolates in $d = 2$ at low temperature for every $\pm J$ configuration [71], where no order is expected, and the averaged two-replica laws lack the positivity that coarse graining uses (Chapter 1).

10.5 The transfer table

Table 10.1 collects the steps of the random-field proofs, what each uses, and its status for the EA model. Strengths follow the route-warning convention of the review: *false* (an identity, a counterexample or a theorem), *fails without a new input* (even under optimistic assumptions), *portable* (transfers as stated).

The portable rows deserve one comment each. The concentration bound of Proposition 10.7(ii) and the chaining lemma behind it hold for any Lipschitz function of iid Gaussian variables; Ding and Zhuang's version reads: if (X_t) is a mean-zero Gaussian process and (Y_t) a mean-zero process with $\mathbb{P}(|Y_s - Y_t| \geq \lambda) \leq 2e^{-\lambda^2/(2d(s,t)^2)}$, where $d(s,t)^2 = \mathbb{E}(X_s - X_t)^2$, then $\mathbb{E} \sup_t Y_t \leq C\, \mathbb{E} \sup_t X_t$ for a universal C [76, Lemma 2.1]. EA free energies are Lipschitz functions of the couplings, so this applies to them unchanged. The counting of connected sets, combined with large deviations, is what Berger and Tessler use to keep the total $|J|$ along long connected sets proportional to their length [26, Lemmas 2.4–2.5]. What is missing for the EA model is not these estimates but the deterministic gain they are compared with, and Proposition 10.8 says that for the joint map that gain is exactly zero.

Random-field step	What it uses	For the EA model	Strength
Joint Peierls map [76, 77]	flipping σ and η on A keeps the field term and gains $2J\lvert\partial A\rvert$	it is a gauge transformation: gain 0, law preserved, $\mathbf{P}^b(\sigma_0 = -1) = \frac{1}{2}$ (Proposition 10.8)	false (identity)
Weak disorder around $H^{(0)}$ [37, 122]	a deterministic ferromagnet with surface tension	no deterministic part for symmetric laws (Observation 10.10); biased laws give ferromagnetic order	false (structural)
Ordered reference state [77]	$\mu_0^+(\sigma_0 = 1) > \frac{1}{2}$ below T_c	an ordered EA reference state is what one wants to prove; its two-replica form likewise	fails without a new input
FK coarse graining, spatial mixing [77]	FK-Ising percolation, positive association	FK percolation without order in $d = 2$ [71]; no positivity for two-replica laws (Chapter 1)	false as an order signal
Monotone coupling, disagreement percolation [6, 8, 10, 75]	FKG, $\mu^- \le \mu^+$	no coordinatewise monotone order under frustration (Lemma 10.11); unordered versions only at high temperature [236]	fails without a new input
Exponential decay of boundary influence, $d = 2$ [10, 75]	monotone coupling	false: the influence on $\sigma_0\sigma_j$, j a neighbour of 0, is at least $1/(32L)$ [49] (Chapter 6)	false (theorem)
Volume against surface [8]	conjugate observable σ_x, odd under the flip	transfers for bond observables [68] (Chapter 4); the conjugate $\sigma_x\sigma_y$ is even (Lemma 10.2)	portable, blind to the flip
Concentration, chaining, lattice animals [67, 74, 76]	Lipschitz functions of iid Gaussians, counting connected sets	apply verbatim; used by Berger–Tessler [26] (Chapter 5)	portable, comparator missing

Table 10.1: Which step of the random-field proofs transfers to the EA model. The probabilistic modules in the last two rows transfer; every step that produces order or exponential decay fails, and the first and sixth fail by an identity or a theorem.

10.6 Route warnings

Route warning 10.12 (Joint spin–disorder Peierls for EA order). *Tempting step:* flip spins and couplings together, as Ding and Zhuang flip spins and fields, and bound the cost of a contour by a free-energy difference controlled by chaining. *Status: false, by an identity.* For a symmetric law the joint map is a gauge transformation, so the gain and the free-energy difference both vanish identically and the bounded event has probability $\frac{1}{2}$ (Proposition 10.8). A Peierls argument for EA order must act on a gauge-invariant object, such as two replicas or their overlap, for which no analogue of the gain $2J\lvert\partial A\rvert$ is known.

Route warning 10.13 (Perturbation around a ferromagnet). *Tempting step:* treat the EA Hamiltonian as a small perturbation of a model with known ordered phases, as in the low-temperature theorems for weakly disordered ferromagnets. *Status: false for symmetric laws (structural).* There is no deterministic part (Observation 10.10). For biased laws there is one, it is a ferromagnet, and what the perturbative theorems produce is ferromagnetic order [96, Thm. 8.2].

Route warning 10.14 (Two-dimensional exponential decay, then uniqueness). *Tempting step:* transfer the two-dimensional exponential decay of boundary influence (Theorem 10.4) to the EA model and conclude uniqueness. *Status: false for the decay, and fails without a new input*

for the method. Chatterjee's theorem in Chapter 6 shows that the boundary influence on a nearest-neighbour product is at least $1/(32L)$ at zero temperature, so the decay is at most polynomial; and the monotone coupling on which the random-field proofs rest does not exist under frustration for any of the orders $\leq_\tau$, since none makes the Gibbs measures monotone in the boundary condition (Lemma 10.11 and its proof). In Chatterjee's words, "a major technical advantage of the RFIM, in comparison with the EA model, is the FKG inequality" [49, p. 9].

10.7 What the mirror shows

The random-field model is solved because its Hamiltonian has a deterministic ferromagnetic part and its measures are monotone. Remove either and the proofs stop: the probabilistic estimates that control the disorder survive, but there is nothing for them to be compared with. For the EA model the cleanest instance is the joint flip, which for symmetric couplings is a symmetry. The constructive lesson is that an EA order proof has to work with gauge-invariant objects from the start, and that the disorder must be coupled to something that sees the global flip.

Open problem 10.15 (Long-range order in $\boldsymbol{d \geq 3}$; Chatterjee [49, p. 10])**.** "In dimensions $d \geq 3$, the RFIM exhibits long-range order; it remains open whether an analogous statement holds for the EA model."

Open problem 10.16 (A disorder conjugate to the overlap)**.** The observables conjugate to the EA couplings are even under the global flip (Lemma 10.2). Consider two replicas with an additional iid random replica coupling $\sum_x \lambda_x \sigma_x \tau_x$ of small symmetric strength, whose conjugate observable $\sigma_x \tau_x$ is odd under the flip of one replica. Develop the Aizenman–Wehr and Dario–Harel–Peled arguments for this two-replica system in $d = 2$, and determine what they imply for the overlap of the unperturbed model as the strength of the replica coupling tends to 0.

The threads resume in three places. Biased couplings, where the joint map produces a surface term at the Nishimori temperature (Remark 10.9), are the subject of Chapter 14. Comparison arguments that use the exact fair average of the signs, rather than a ferromagnetic seed, are developed in Chapter 11. The lessons of this chapter return in the list of barriers of Chapter 15: the fair average must be used exactly, and there is no positivity to lean on.

Chapter 11

Comparison with forests and electrical networks

In two dimensions the EA model is expected to have no spin-glass order at any positive temperature, but for $\pm J$ or Gaussian couplings no proof is known (for diluted laws just above the percolation threshold, see Chapter 13). This chapter asks whether non-order can be obtained by *comparison*: is the squared two-point function of the glass, averaged over the signs of the couplings, bounded by the connection probability of a simpler model whose behaviour in $d = 2$ is understood? The Edwards–Sokal route of the ferromagnet does not work as it stands, since those clusters percolate where no order is expected (Chapter 7); the comparison has to see the signs.

We follow one candidate as far as present proofs go: the *arboreal gas*, the weighted random spanning forest with weight $\sinh^2 K_e$ on an edge of magnitude K_e. The comparison it suggests (the forest-domination conjecture FD) is open, and what is proved is a chain around it:

$$\mathrm{FD} \xRightarrow{\text{forest/conductance bound}} \mathrm{RD} \xRightarrow{\text{Nash-Williams}} \mathrm{NWB} \xRightarrow{\mathbb{Z}^2} \mathbb{E}\langle\sigma_u\sigma_v\rangle^2 \le \frac{8\,\bar{y}(\beta)}{\ln(2r+1)},$$

where $\bar{y}(\beta) = \mathbb{E}\sinh^2(\beta|J_e|)$ and r is the distance. The first arrow is a theorem about forests, the last a short argument with annular cuts, and the middle statements are conjectures proved on bundles, series–parallel networks and a few small graphs (Fig. 11.3). This is one of the two routes of our own programme (the other is Chapter 8); the forest bound is in a preprint [216], and the rest is proved or sketched here with its current status. The reader will be able to prove a single-edge inequality and from it the bundle, cut and series–parallel bounds, the forest bound, and the conditional non-order theorem on $\mathbb{Z}^2$; the last section shows, through relatives of the arboreal gas, what any comparison model must contain.

11.1 Fixed magnitudes, fair signs, and electrical networks

For symmetric couplings the sign of each coupling is a fair coin independent of its magnitude. All statements of this chapter are inequalities on a finite graph with fixed magnitudes and averaged signs; the coupling law enters at the end, through one number.

Definition 11.1 (The fair-sign glass on a finite graph)**.** Let $G = (V, E)$ be a finite multigraph without loops, with *magnitudes* $K = (K_e)_{e\in E} \in [0,\infty)^E$, and let $\varepsilon = (\varepsilon_e)_{e\in E}$ be independent fair signs. Write $\langle\cdot\rangle_{G,\varepsilon}$ for the Ising measure on $\{-1,+1\}^V$ with weight $\exp(\sum_{e=xy}\varepsilon_e K_e\sigma_x\sigma_y)$

(free boundary condition), and for $u \neq v$ put $c = c_{uv}(\varepsilon) = \langle \sigma_u \sigma_v \rangle_{G,\varepsilon}$ and $\mathcal{Q}_G(u,v) = \mathbb{E}_\varepsilon\, c^2$, where $\mathbb{E}_\varepsilon$ averages the signs only. The *activity* and *conductance* of an edge are $t_e = \tanh K_e$ and $y_e = \sinh^2 K_e$, so that $t_e^2 = y_e/(1+y_e)$. For a function F of the signs, $D_e F = \frac{1}{2}(F|_{\varepsilon_e=+1} - F|_{\varepsilon_e=-1})$ is the *sign-flip derivative* of e.

We write $\mathcal{Q}$ rather than Q to keep Q for overlaps. For a single edge $c = \varepsilon t$, so $\mathcal{Q} = t^2$ and

$$\frac{\mathcal{Q}}{1-\mathcal{Q}} = \sinh^2 K = y : \tag{11.1}$$

on one edge the odds are the conductance, which is the reason for the choice. If J has a symmetric law, the signs $\operatorname{sgn} J_e$, completed by an independent fair coin where $J_e = 0$, are independent fair signs independent of the magnitudes $|J_e|$; hence, on a finite Λ with free boundary conditions,

$$\mathbb{E}\,\langle \sigma_u \sigma_v \rangle_{\Lambda,J}^2 = \mathbb{E}_K\, \mathcal{Q}_\Lambda(u,v), \qquad K_e = \beta |J_e|, \tag{11.2}$$

where $\mathbb{E}_K$ averages the magnitudes. The coupling law will enter only through the mean conductance $\bar{y}(\beta) = \mathbb{E} \sinh^2(\beta |J_e|)$, which is $\sinh^2 \beta$ for $\pm J$ couplings and $\frac{1}{2}(e^{2\beta^2} - 1)$ for Gaussian ones, finite at every β for both.

11.1.1 Electrical tools

We use a few classical facts about electrical networks [150, Ch. 2]. Give each edge e the conductance $y_e \geq 0$. The *effective conductance* between $u \neq v$ is the Dirichlet minimum

$$C_{\mathrm{eff}}^G(u,v) = \min\Big\{ \sum_{e=xy \in E} y_e (f(x) - f(y))^2 : \ f \colon V \to \mathbb{R},\ f(u) = 1,\ f(v) = 0 \Big\}, \tag{11.3}$$

and $R_{\mathrm{eff}} = 1/C_{\mathrm{eff}} \in (0, \infty]$. A minimizer f is the *potential*; it is harmonic off $\{u, v\}$ and takes values in $[0,1]$.

Lemma 11.2 (Electrical facts)**.**

(i) Concavity and Rayleigh monotonicity. *$C_{\mathrm{eff}}^G(u,v)$ is a concave, nondecreasing and 1-homogeneous function of y. Deleting edges, or passing to an induced subgraph that contains u and v, does not increase it.*

(ii) Series and parallel. *Effective conductances add for networks joined in parallel at their terminals, and effective resistances add for networks joined in series.*

(iii) Nash-Williams' inequality *[169]. If $\Pi_1, \dots, \Pi_k$ are pairwise disjoint edge sets each of which meets every u–v path, then $R_{\mathrm{eff}}^G(u,v) \geq \sum_j \big(\sum_{e \in \Pi_j} y_e\big)^{-1}$.*

Proof. (i) is read off from (11.3): a minimum of linear functions of y with nonnegative coefficients is concave, nondecreasing and 1-homogeneous, deleting an edge removes a nonnegative term, and the minimizer restricts to an admissible function on $G[S]$; (ii) is classical. For (iii), by Thomson's principle R_{eff} is the minimal energy $\sum_e \theta_e^2 / y_e$ of a unit u–v flow, which carries total flow at least one across each Π_j, so by Cauchy–Schwarz the energy on Π_j is at least $1/y(\Pi_j)$, with $y(\Pi) = \sum_{e \in \Pi} y_e$. □

On $\mathbb{Z}^2$ the inequality (iii) produces the logarithm that drives everything below. For $u \in \mathbb{Z}^2$ and $j \geq 0$ let $\Pi_j(u)$ be the set of edges joining $\{z : |z-u|_\infty = j\}$ to $\{z : |z-u|_\infty = j+1\}$. These sets are disjoint, $|\Pi_j(u)| = 4(2j+1)$, and $\Pi_j(u)$ meets every path from u to a point v with $|v-u|_\infty > j$, because $|z-u|_\infty$ changes by at most one along an edge (Fig. 11.1(b)). With unit conductances, for $r = |u-v|_\infty \geq 1$,

$$R_{\mathrm{eff}}^{\mathbb{Z}^2,1}(u,v) \;\geq\; \sum_{j=0}^{r-1} \frac{1}{4(2j+1)} \;\geq\; \frac{1}{4}\int_0^r \frac{dx}{2x+1} \;=\; \frac{\ln(2r+1)}{8}, \tag{11.4}$$

and the same bound holds in every finite subgraph of $\mathbb{Z}^2$ containing u and v, where the sets $\Pi_j(u) \cap E(\Lambda)$ still separate them. The two-dimensional lattice is *recurrent*: its point-to-point resistance diverges.

11.2 One edge, one bundle, one cut

Everything proved in this chapter about the glass rests on one inequality for a single edge in the presence of an external coupling b between its endpoints.

Lemma 11.3 (Single edge)**.** *Let $K \geq 0$, $b \in \mathbb{R}$, $y = \sinh^2 K$, and let ε be a fair sign. Then*

$$\operatorname{Var}_\varepsilon \tanh(b+\varepsilon K) = \Big(\tfrac{1}{2}\big[\tanh(b+K) - \tanh(b-K)\big]\Big)^2 \;\leq\; y\,\mathbb{E}_\varepsilon \operatorname{sech}^2(b+\varepsilon K),$$

with equality if and only if $b = 0$ or $K = 0$.

Proof. For a function of one fair sign the variance is the square of half its increment. Using $\tanh\alpha - \tanh\gamma = \sinh(\alpha-\gamma)/(\cosh\alpha\cosh\gamma)$ and $\cosh(b+K)\cosh(b-K) = \cosh^2 b + \sinh^2 K$,

$$\tfrac{1}{2}\big[\tanh(b+K) - \tanh(b-K)\big] = \frac{\sinh K \cosh K}{\cosh^2 b + y}.$$

Since $\cosh^2 K = 1+y \leq \cosh^2 b + y$, its square is at most $y/(\cosh^2 b + y) = y\operatorname{sech}(b+K)\operatorname{sech}(b-K)$, and by the arithmetic–geometric mean inequality this is at most $\frac{y}{2}[\operatorname{sech}^2(b+K) + \operatorname{sech}^2(b-K)]$. The first step is strict when $b \neq 0$ and $K > 0$; the second is an equality exactly when $b = 0$ or $K = 0$, and at $b = 0$ both sides equal $\tanh^2 K$. □

At $b = 0$ the lemma is (11.1); the external coupling makes it useful, since in a graph the rest of the network acts on an edge as an external coupling between its endpoints. The second tool is the Efron–Stein inequality [82] (used in Chapter 1) on the sign cube: for independent fair signs and any F,

$$\operatorname{Var} F \;\leq\; \sum_e \mathbb{E}\,(D_e F)^2, \tag{11.5}$$

an equality for one sign, and in general by conditioning on one coordinate at a time.

The simplest graphs with cycles are bundles of parallel edges between u and v, where the forest model and the electrical network are computed by hand and the comparison becomes a scalar inequality.

Theorem 11.4 (Fair signs on a bundle)**.** *For every $k \geq 1$ and $K_1, \dots, K_k \geq 0$, with $\varepsilon_1, \dots, \varepsilon_k$ independent fair signs, $X = \sum_i \varepsilon_i K_i$ and $y_i = \sinh^2 K_i$,*

$$\mathbb{E}\operatorname{sech}^2 X \;\geq\; \frac{1}{1+\sum_i y_i},$$

with equality if and only if at most one K_i is nonzero. Equivalently, on the multigraph with k parallel edges between u and v, $\mathcal{Q}(u,v) \leq y(E)/(1+y(E))$ with $y(E) = \sum_i y_i$.

Proof. On the bundle $c = \tanh X$, so $\mathcal{Q} = 1 - \mathbb{E}\operatorname{sech}^2 X$. The function $F = \tanh X$ is odd under $\varepsilon \mapsto -\varepsilon$, so $\mathbb{E}F = 0$ and $\mathcal{Q} = \operatorname{Var} F$. Write $X_{-i} = X - \varepsilon_i K_i$. Then $D_i F = \frac{1}{2}[\tanh(X_{-i} + K_i) - \tanh(X_{-i} - K_i)]$ does not depend on ε_i, and Lemma 11.3 with $b = X_{-i}$ gives $(D_i F)^2 \leq y_i\, \mathbb{E}[\operatorname{sech}^2 X \mid \varepsilon_j,\, j \neq i]$. By (11.5),

$$1 - \mathbb{E}\operatorname{sech}^2 X = \operatorname{Var} F \leq \sum_i y_i\, \mathbb{E}\operatorname{sech}^2 X,$$

which rearranges to the claim. If $K_i > 0$ and $K_j > 0$ for some $j \neq i$, then $X_{-i} \neq 0$ with positive probability, and the step for i is strict by the equality case of Lemma 11.3. If at most one K_i is nonzero, both sides equal $1/(1+y_i)$. □

As $C_{\text{eff}} = y(E)$ on the bundle, Theorem 11.4 reads $\mathcal{Q}/(1-\mathcal{Q}) \leq C_{\text{eff}}(u,v)$: the odds of the bundle are at most those of its electrical network, with equality only for a single edge.

Two contrasts show that the inequality depends on the exact fair average. First, with *fixed* equal signs on two parallel edges of magnitude K, $c = \tanh 2K$, and since $1 + 2y = \cosh 2K$, $\tanh^2 2K = 1 - \operatorname{sech}^2 2K > 1 - \operatorname{sech} 2K = 2y/(1+2y)$: the bound fails for every $K > 0$. Second, any bias of the sign law toward unfrustrated configurations breaks it:

Proposition 11.5 (Tilted sign laws violate the bundle bound)**.** *For $\alpha > 0$ let $\mathbb{P}_\alpha(\varepsilon) \propto Z_\varepsilon^\alpha$ on $\{\pm 1\}^E$, where $Z_\varepsilon = \sum_\sigma \exp(\sum_e \varepsilon_e K_e \sigma_x \sigma_y)$. On three parallel edges of equal magnitude K, for every $\alpha > 0$ there is $K_0(\alpha)$ such that $\mathbb{E}_{\mathbb{P}_\alpha} c^2 > 3y/(1+3y)$ for all $K > K_0(\alpha)$.*

The proof is an explicit computation on the eight sign patterns (Section A.22).

The tilt $\alpha = 1$ is the gauge-symmetrized Nishimori law of Chapter 14 (the "planted" law), and $\alpha = 2$ is the sign marginal of the annealed two-replica measure; both favour unfrustrated cycles and make the glass more correlated than the forest (Route warning 11.23).

On a general graph, fair signs give a first estimate through a symmetry that has no ferromagnetic analogue: flipping the spins on one side of a cut (Fig. 11.1(a)).

Lemma 11.6 (Random cuts)**.** *Let $u \neq v$ and let C be a random set of edges, independent of ε, every realization of which meets every u–v path. Then*

$$\mathcal{Q}_G(u,v) \;\leq\; \sum_{e \in E} \mathbb{P}(e \in C)\, \mathbb{E}_\varepsilon (D_e c)^2.$$

In particular $\mathcal{Q}_G(u,v) \leq \sum_e |f(x_e) - f(y_e)|\, \mathbb{E}(D_e c)^2$ for every $f\colon V \to [0,1]$ with $f(u) = 1$, $f(v) = 0$, for instance the potential of (11.3); here $e = x_e y_e$.

Proof. Fix a realization of C and let S be the vertex set of the component of u in $(V, E \setminus C)$, so that $v \notin S$ and the edge boundary $\delta(S)$ is contained in C. The map $\sigma \mapsto \sigma'$ that flips every spin in S carries the Ising measure with signs ε to the one with the signs on $\delta(S)$ reversed, and

it changes $\sigma_u\sigma_v$ into $-\sigma_u\sigma_v$. So c changes sign when every sign on $\delta(S)$ is reversed. Given the signs off $\delta(S)$, the signs on $\delta(S)$ are fair and their law is invariant under that reversal, hence $\mathbb{E}[c \mid \varepsilon_{E\setminus\delta(S)}] = 0$ and, by (11.5) applied to the signs on $\delta(S)$,

$$\mathcal{Q} = \mathbb{E}\,\operatorname{Var}\big(c \mid \varepsilon_{E\setminus\delta(S)}\big) \le \sum_{e\in\delta(S)} \mathbb{E}(D_e c)^2 \le \sum_{e\in C} \mathbb{E}(D_e c)^2.$$

Averaging over C gives the first claim. For the second, let U be uniform on $(0,1)$ and independent of ε, and let C be the set of edges between $\{f > U\}$ and $\{f \le U\}$. It separates u (where $f = 1$) from v (where $f = 0$), and e lies in C exactly when U falls between the values of f at its endpoints. □

The influence $\mathbb{E}(D_e c)^2$ of an edge is controlled by Lemma 11.3, with the rest of the graph as the external coupling.

Lemma 11.7 (Restoring one edge)**.** *For every edge e and every choice of the signs off e,*

$$(D_e c)^2 \;\le\; y_e\,\mathbb{E}_{\varepsilon_e}(1 - c^2).$$

The proof conditions on $\sigma_a\sigma_b$, $e = ab$, which writes $c = \alpha_e + B_e \tanh(b_e + \varepsilon_e K_e)$ with α_e, B_e, b_e independent of ε_e, and applies Lemma 11.3; it is in Section A.23.

Theorem 11.8 (Cut inequality)**.** *For every finite multigraph, all magnitudes, $u \ne v$ and every set C of edges meeting every u–v path,*

$$\frac{\mathcal{Q}_G(u,v)}{1 - \mathcal{Q}_G(u,v)} \;\le\; y(C) = \sum_{e\in C} y_e.$$

Proof. By Lemma 11.6 with the deterministic cut C and Lemma 11.7, $\mathcal{Q} \le \sum_{e\in C} \mathbb{E}(D_e c)^2 \le \sum_{e\in C} y_e(1 - \mathcal{Q})$. □

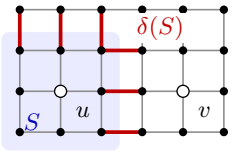
For a single edge this is (11.1), and on a bundle it is Theorem 11.4 up to the strictness. But it is local: on $\mathbb{Z}^2$ the minimal cut between two points has four edges whatever their distance, so alone it gives no decay. To gain from distance one must combine many cuts, and that is where the conjectures begin.

The single-edge step of Theorem 11.4 is the local-energy inequality of Kitatani, Nishimori and Aoki [141, Thm. 1], and the cut mechanism of Theorem 11.8 also appears in Chatterjee's zero-temperature analysis [49, Sect. 1.7 and Cor. 5]. On trees, the analogue of the electrical comparison developed below is due to Evans, Kenyon, Peres and Schulman [84, Thm. 1.2].

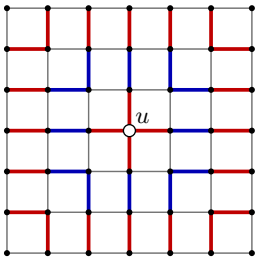
11.3 The arboreal gas and its electrical bound

Which model should the glass be compared with? On a tree the answer is forced: the gauge removes the signs (Chapter 1), $c = \prod_e \varepsilon_e t_e$ along the path from u to v, and $\mathcal{Q} = \prod_e t_e^2$ is the connection probability of Bernoulli(t_e^2) percolation. With cycles Bernoulli percolation overcounts, because a frustrated cycle cancels; the candidate that keeps the tree values and forbids cycles is Bernoulli(t_e^2) percolation conditioned to be acyclic.

Definition 11.9 (Arboreal gas)**.** For a finite multigraph G with weights $y \in [0,\infty)^E$, the *arboreal gas* is the probability measure on forests (acyclic edge sets) $F \subseteq E$ given by $\mathbb{P}^{\mathrm{AG}}_{G,y}(F) \propto \prod_{e\in F} y_e$. We write $u \leftrightarrow v$ when u and v lie in the same tree of F.



(a) flip all spins in S: only the signs on $\delta(S)$ change, and $\sigma_u\sigma_v \mapsto -\sigma_u\sigma_v$



(b) disjoint cuts Π_0, Π_1, Π_2 around u (red, blue, red), of sizes $4, 12, 20$

Figure 11.1: Cuts. (a) Flipping the spins in $S \ni u$ reverses exactly the coupling signs on $\delta(S)$ (red), so for fair signs $c = \langle\sigma_u\sigma_v\rangle$ has mean zero given the other signs (Lemma 11.6); one cut gives a bound that does not improve with distance. (b) The annular cuts $\Pi_j(u)$ on $\mathbb{Z}^2$ have $4(2j+1)$ edges; Nash-Williams' inequality turns them into the logarithmic resistance (11.4), which a comparison with an electrical network converts into decay.

With $y_e = \sinh^2 K_e$ this is Bernoulli(t_e^2) percolation conditioned to be acyclic, since $t^2/(1-t^2) = y$; on a tree $\mathbb{P}^{\mathrm{AG}}(u \leftrightarrow v) = \prod t_e^2 = \mathcal{Q}(u,v)$. With constant weights the arboreal gas is the $q \to 0$ limit of the random-cluster model with p/q fixed. In two dimensions, Bauerschmidt, Crawford, Helmuth and Swan proved through a supersymmetric representation that on every finite subgraph of $\mathbb{Z}^2$ with constant weight its connection probabilities decay at least like a power of the distance [23, Theorem 1.3], and Halberstam and Hutchcroft developed the theory of its infinite-volume Gibbs measures [113]. The comparison needs a forest-level understanding, supplied by an exact identity.

For a multigraph H with weights y, let $T(H)$ be the weighted sum of its spanning trees, and for $u \neq v$ let $T_2(H; u|v)$ be the weighted sum of its spanning forests with exactly two trees, one containing u and the other v.

Lemma 11.10 (Kirchhoff)**.** *For every finite multigraph H, weights $y \geq 0$ and $u \neq v$, $T(H) = C_{\mathrm{eff}}^{H}(u,v)\, T_2(H; u|v)$.*

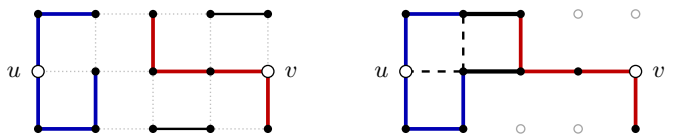
This is Kirchhoff's formula [150, Ch. 4]: with L the weighted Laplacian and $R = V(H)\setminus\{u,v\}$, the matrix-tree theorem gives $T(H)$ as the minor of L without v and, after identifying u with v, $T_2(H; u|v) = \det L_{RR}$; if $T_2 > 0$ that minor is $\det L_{RR}$ times the Schur complement $L_{uu} - L_{uR}L_{RR}^{-1}L_{Ru}$, which is $C_{\mathrm{eff}}^{H}(u,v)$ by minimizing (11.3) over the values on R; and $T_2 = 0$ forces $T = 0$, as a spanning tree minus an edge of its u–v path is a separating two-forest.

For a forest F and a vertex x let $S_x(F)$ be the vertex set of the tree of F containing x, and write $G[S]$ for the subgraph induced by S, with all edges of G between vertices of S.

Theorem 11.11 (Forest odds are an average conductance [216, Theorem 1.1])**.** *For every finite multigraph G, weights $y \geq 0$ and $u \neq v$,*

$$\frac{\mathbb{P}^{\mathrm{AG}}(u \leftrightarrow v)}{\mathbb{P}^{\mathrm{AG}}(u \not\leftrightarrow v)} = \mathbb{E}^{\mathrm{AG}}\Big[C_{\mathrm{eff}}^{G[S_u \cup S_v]}(u,v) \;\Big|\; u \not\leftrightarrow v\Big] \;\leq\; C_{\mathrm{eff}}^{G}(u,v), \tag{11.6}$$

where the conductances are y. Equivalently $\mathbb{P}^{\mathrm{AG}}_{G,y}(u \leftrightarrow v) \leq 1/(1 + R_{\mathrm{eff}}(u,v))$. Equality holds if and only if $C_{\mathrm{eff}}^{G-E(u,v)}(u,v) = 0$, where $E(u,v)$ is the set of edges joining u to v.



(a) a forest F with $u \not\leftrightarrow v$ (b) the induced subgraph $G[S_u \cup S_v]$

Figure 11.2: The cluster identity of Theorem 11.11. (a) A forest separating u from v: the tree of u (blue), the tree of v (red) and the rest (black). (b) The subgraph induced by the two vertex sets: it contains the two trees, the edges of G inside each vertex set (dashed) and the edges joining them (thick). Kirchhoff's formula turns the forests that connect u to v inside this region into its effective conductance, so the forest odds are the average conductance of the explored region; Rayleigh monotonicity then enlarges the region to G.

Proof. Write Z_{c} and Z_{s} for the weighted sums of the forests with $u \leftrightarrow v$ and $u \not\leftrightarrow v$. A forest with $u \leftrightarrow v$ whose tree at u has vertex set S is a spanning tree of $G[S]$ together with a forest of $G[V \setminus S]$; a forest with $u \not\leftrightarrow v$ and $S_u \cup S_v = S$ is a spanning two-forest of $G[S]$ separating u from v, together with a forest of $G[V \setminus S]$. Both decompositions are bijective, so with $Z(\cdot)$ the forest sum,

$$Z_{\mathrm{c}} = \sum_{S \ni u,v} T(G[S])\, Z(G[V \setminus S]), \qquad Z_{\mathrm{s}} = \sum_{S \ni u,v} T_2(G[S]; u|v)\, Z(G[V \setminus S]).$$

By Lemma 11.10, $T(G[S]) = C_{\mathrm{eff}}^{G[S]}(u,v)\, T_2(G[S]; u|v)$; so Z_{c} is the Z_{s}-weighted sum of $C_{\mathrm{eff}}^{G[S]}$, which is the identity. Rayleigh monotonicity (Lemma 11.2(i)) gives $C_{\mathrm{eff}}^{G[S]} \le C_{\mathrm{eff}}^{G}$. For equality: the empty forest is separating, has positive weight and $S = \{u,v\}$, where $C_{\mathrm{eff}}^{G[\{u,v\}]} = y(E(u,v))$; since the direct edges are in parallel with $G - E(u,v)$, equality forces $C_{\mathrm{eff}}^{G-E(u,v)} = 0$. Conversely, if that conductance vanishes then $y(E(u,v)) \le C_{\mathrm{eff}}^{G[S]} \le C_{\mathrm{eff}}^{G} = y(E(u,v))$ for every $S \ni u,v$. □

Figure 11.2 shows the objects in (11.6). The identity says more than the inequality: the forest odds are the conductance of the region the forest has explored around u and v, and the bound only forgets that the region is random. On a fixed graph at large weights the conditional law concentrates on spanning two-forests ($S_u \cup S_v = V$), so the bound is asymptotically sharp [216, Proposition 1.2]. For a ghost vertex joined to every vertex of a torus with constant weights, (11.6) for the pair (vertex, ghost) is the Mermin–Wagner bound of Bauerschmidt, Crawford, Helmuth and Swan [23, Theorem 4.2].

The two-dimensional consequence is immediate from (11.4): in every finite subgraph H of $\mathbb{Z}^2$ with weights $y_e \le Y$,

$$\mathbb{P}^{\mathrm{AG}}_{H,y}(u \leftrightarrow v) \;\le\; \frac{1}{1 + \ln(2r+1)/(8Y)} \;\le\; \frac{8Y}{\ln(2r+1)}, \qquad r = |u - v|_\infty. \tag{11.7}$$

The arboreal-gas note proves this also for boundary conditions on vertices at distance at least r from u [216, Theorem 1.3], and with the resampling theorem of Halberstam and Hutchcroft it gives, without supersymmetry, that translation-invariant arboreal-gas Gibbs measures on $\mathbb{Z}^2$ have no infinite tree [216, Corollary 1.5]; the supersymmetric method gives the stronger power law. For random weights with $\mathbb{E}y_e \le \bar{y}$, concavity of C_{eff} (Lemma 11.2(i)) and Jensen's inequality give $\mathbb{E}_y \mathbb{P}^{\mathrm{AG}}_{H,y}(u \leftrightarrow v) \le \mathbb{E}_y C_{\mathrm{eff}}(u,v) \le \bar{y}\, C_{\mathrm{eff}}^{H,1}(u,v) \le 8\bar{y}/\ln(2r+1)$, with only a first moment.

The cut inequality also has a forest counterpart: deleting the first edge of C on the u–v path of a forest with $u \leftrightarrow v$ is injective for each choice of that edge and divides the weight by its y_e, so $\mathbb{P}^{\mathrm{AG}}(u \leftrightarrow v)/\mathbb{P}^{\mathrm{AG}}(u \not\leftrightarrow v) \le y(C)$ for every set C meeting every u–v path.

11.4 The comparison conjectures and what they would give

Conjecture 11.12 (Forest domination, FD)**.** *For every finite multigraph G, all magnitudes and all $u \neq v$,*

$$\mathcal{Q}_G(u,v) \;\le\; \mathbb{P}^{\mathrm{AG}}_{G,y}(u \leftrightarrow v), \qquad y_e = \sinh^2 K_e.$$

It is an equality on trees, and it is proved on bundles (Theorem 11.4: the arboreal gas there uses at most one edge and gives $y(E)/(1+y(E))$), on networks built from single edges by parallel composition and by series composition with a single edge (including cycles, fans and ladders, at suitable terminal pairs), and, by exact computer-assisted certificates, on the Wheatstone bridge and K_4 at every pair and all magnitudes. The proofs for the series–parallel classes are not reproduced here; the bridge and K_4 certificates, each verified by two independent implementations, are described in Appendix B. The conjecture is closed under minors (a magnitude 0 deletes an edge, a magnitude ∞ contracts it after a gauge change, on both sides), so FD on the finite subgraphs of $\mathbb{Z}^2$ covers every finite planar graph, and a proof for all magnitudes cannot use any structure of the square lattice beyond planarity. Exact and numerical searches on many small graphs found no counterexample; this is evidence, not proof.

Two weaker conjectures sit below it. For disjoint edge sets $\Pi_1, \dots, \Pi_k$, each meeting every u–v path, put $\mathrm{NW}(\Pi_1, \dots, \Pi_k) = \big(\sum_j 1/y(\Pi_j)\big)^{-1}$ (with $1/0 = \infty$).

Conjecture 11.13 (Resistance domination, RD, and the Nash-Williams bound, NWB)**.** *For every finite multigraph, all magnitudes and all $u \neq v$:*

(RD) $\quad \mathcal{Q}_G(u,v)/(1-\mathcal{Q}_G(u,v)) \le C^G_{\mathrm{eff}}(u,v)$*;*

(NWB) $\quad \mathcal{Q}_G(u,v)/(1-\mathcal{Q}_G(u,v)) \le \mathrm{NW}(\Pi_1, \dots, \Pi_k)$ *for all such disjoint* $\Pi_1, \dots, \Pi_k$*.*

Proposition 11.14 (The ladder)**.** *(a) FD on G implies RD on G. (b) RD on G implies NWB on G. (c) NWB with $k = 1$ is Theorem 11.8, which holds on every graph. (d) On a bundle the three statements coincide with Theorem 11.4; hence every tilted law of Proposition 11.5 violates all three.*

Proof. (a) If $\mathcal{Q} \le \mathbb{P}^{\mathrm{AG}}(u \leftrightarrow v)$, then $\mathcal{Q}/(1-\mathcal{Q})$ is at most the forest odds, which are at most C_{eff} by Theorem 11.11. (b) $C_{\mathrm{eff}} = 1/R_{\mathrm{eff}} \le \mathrm{NW}(\Pi_1, \dots, \Pi_k)$ by Nash-Williams' inequality. (c) With one set, $\mathrm{NW} = y(\Pi_1)$. (d) On a bundle the forest probability is $y(E)/(1+y(E))$, $C_{\mathrm{eff}} = y(E)$, and the only minimal separating set is the bundle itself, of weight $y(E)$. □

In the form (11.6), FD says $\mathcal{Q}/(1-\mathcal{Q}) \le \mathbb{E}^{\mathrm{AG}}[C^{G[S_u \cup S_v]}_{\mathrm{eff}} \mid u \not\leftrightarrow v]$: it is RD on the random region explored by the forest, and since conductances add in parallel and are exact on bundles, RD relaxes FD only at series steps. The bottom of the ladder already gives two-dimensional non-order.

Theorem 11.15 (Conditional non-order in two dimensions)**.** *Assume NWB holds on every finite subgraph of $\mathbb{Z}^2$. Let the coupling law be symmetric and $\beta > 0$ with $\bar{y}(\beta) = \mathbb{E}\sinh^2(\beta|J_e|) < \infty$.*

Then for every finite subgraph Λ of $\mathbb{Z}^2$ and all $u, v \in \Lambda$ with $r = |u - v|_\infty \geq 1$, with free boundary conditions,

$$\mathbb{E}\,\langle \sigma_u \sigma_v \rangle_{\Lambda,J}^2 \;\leq\; \frac{8\,\bar{y}(\beta)}{\ln(2r+1)}.$$

In particular this holds if FD or RD holds on every finite subgraph of $\mathbb{Z}^2$.

Proof. Let $\Pi_j = \Pi_j(u) \cap E(\Lambda)$ for $0 \leq j < r$; these are disjoint and separate u from v in Λ. By (11.2) and NWB on Λ, $\mathbb{E}\langle \sigma_u \sigma_v \rangle_{\Lambda,J}^2 = \mathbb{E}_K \mathcal{Q}_\Lambda(u,v) \leq \mathbb{E}_K \,\mathrm{NW}(\Pi_0, \dots, \Pi_{r-1})$, where $y_e = \sinh^2(\beta |J_e|)$. The map $(a_0, \dots, a_{r-1}) \mapsto (\sum_j 1/a_j)^{-1}$ is concave and nondecreasing on $[0,\infty)^r$, and each $a_j = y(\Pi_j)$ is linear in y with $\mathbb{E}\, y(\Pi_j) = \bar{y}\,|\Pi_j| \leq \bar{y}\, 4(2j+1)$. By Jensen's inequality and (11.4),

$$\mathbb{E}_K \,\mathrm{NW}(\Pi_0, \dots, \Pi_{r-1}) \;\leq\; \Big(\sum_{j<r} \frac{1}{4(2j+1)\,\bar{y}} \Big)^{-1} \;\leq\; \frac{8\,\bar{y}}{\ln(2r+1)}.$$

The last statement follows from Proposition 11.14. □

The bound is uniform in the volume and holds at every temperature with $\bar{y}(\beta) < \infty$ (every temperature for $\pm J$ and Gaussian couplings). It is logarithmic and weak at low temperature: for $\pm J$, $\bar{y} = \sinh^2 \beta$, so it says something only when $\ln(2r+1) > 8 \sinh^2 \beta$. On any graph whose finite subgraphs satisfy RD the same argument gives $\mathbb{E}\langle \sigma_u \sigma_v \rangle^2 \leq \bar{y}\, C_{\mathrm{eff}}^{G,1}(u,v)$ with unit conductances, which tends to 0 as v moves away when the graph is recurrent, and stays bounded below on the transient $\mathbb{Z}^d$, $d \geq 3$. The comparison sorts graphs by recurrence and is silent in $d \geq 3$, as it must be if the physics of Chapter 2 is right.

FD itself gives more, since the arboreal gas decays faster than its electrical bound.

Corollary 11.16 (A power law from FD)**.** *Assume FD holds on every finite subgraph of $\mathbb{Z}^2$. For $\pm J$ couplings and every $\beta > 0$ there is $c(\beta) > 0$ such that for every finite subgraph Λ of $\mathbb{Z}^2$ and $0, x \in \Lambda$ with $x \neq 0$, with free boundary conditions, $\mathbb{E}\langle \sigma_0 \sigma_x \rangle_{\Lambda,J}^2 \leq |x|_\infty^{-c(\beta)}$.*

Proof. All magnitudes equal β, so $\mathcal{Q}_\Lambda(0,x) \leq \mathbb{P}^{\mathrm{AG}}_{\Lambda,y}(0 \leftrightarrow x)$ with the constant weight $y = \sinh^2 \beta$, and the power-law bound for the arboreal gas with constant weight [23, Theorem 1.3] applies; any of the usual norms on $\mathbb{Z}^2$ dominates $|x|_\infty$. □

For general symmetric laws with $\bar{y} < \infty$, a weight-inhomogeneous form of [23, Lemma 4.1] gives the exponent $1/(144(\bar{y}+3))$ under FD; we omit the argument.

11.5 Series–parallel networks

The middle of the ladder, RD, is proved on every series–parallel network, by an induction that shows both what composes and what does not.

A *two-terminal network* (N, u, v) is a finite multigraph with two marked vertices. *Parallel composition* identifies the two pairs of terminals of two networks; *series composition* identifies the second terminal of one with the first terminal of the other. A network is *series–parallel* if it can be built from single edges by these operations, like the diamond hierarchical lattices of Chapter 12 between their poles; the Wheatstone bridge, K_4 minus an edge with the two degree-two vertices as terminals, is the simplest network that is not (Fig. 11.4).

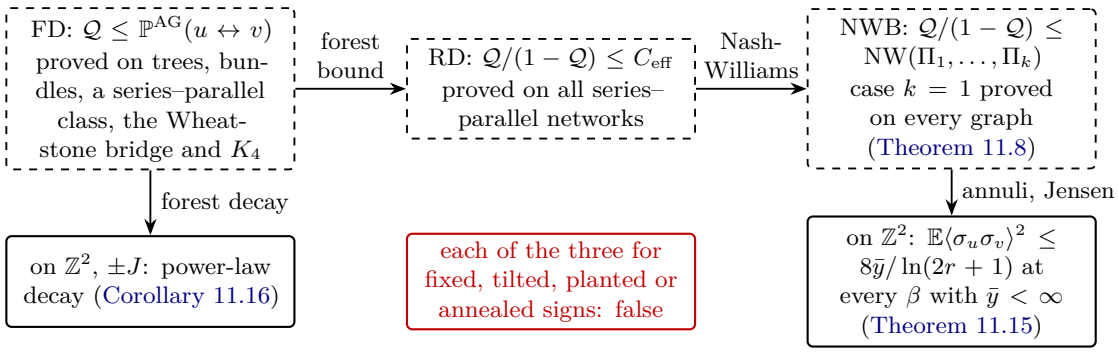


Figure 11.3: The comparison ladder. Each arrow is a proved implication; every arrow out of a dashed box is conditional on that box. The bottom of the ladder, NWB on the finite subgraphs of $\mathbb{Z}^2$, already gives non-order at every temperature (Theorem 11.15); only its one-cut case is proved in general. The red box records that all three inequalities fail for fixed signs and for every tilted law $\mathbb{P}_\alpha$, $\alpha > 0$ (Propositions 11.5 and 11.14), so a proof must use the fair average exactly.

Theorem 11.17 (RD on series–parallel networks)**.** *For every two-terminal series–parallel network (N, u, v) and all magnitudes,*

$$\mathcal{Q}_N(u,v) \;\le\; \frac{C_{\mathrm{eff}}(u,v)}{1 + C_{\mathrm{eff}}(u,v)}, \qquad y_e = \sinh^2 K_e.$$

Correlations with an external coupling. The inductive statement must allow an external coupling between the terminals, since in a composition each piece sees the other as one. Define the effective coupling K_N of (N, u, v) by $\tanh K_N = c$ (it depends on the signs); adding a u–v edge of deterministic coupling b changes the correlation to $\tanh(b + K_N)$. Put

$$I_e(b) = \mathbb{E}_\varepsilon\big(D_e \tanh(b + K_N)\big)^2, \qquad S(b) = \mathbb{E}_\varepsilon \operatorname{sech}^2(b + K_N),$$

so that $I_e(0) = \mathbb{E}(D_e c)^2$ and $S(0) = 1 - \mathcal{Q}$. The composition rules are exact: summing out the interior of each piece leaves a weight $\exp(K_i \sigma_u \sigma_v)$ on the terminals, so in parallel $K_{N_1 \| N_2} = K_1 + K_2$; in series, conditioning on the spin at the junction gives $c_{N_1 - N_2} = c_1 c_2$.

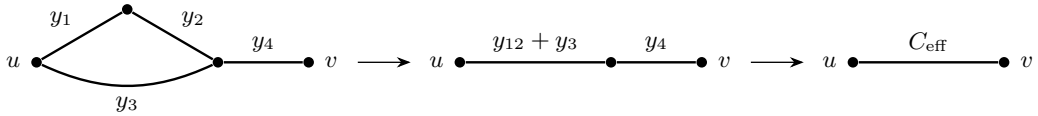
Definition 11.18 (Tame networks)**.** Let $\rho \ge 0$ and let $w_e = \mathbb{P}(e \in C)$ be the marginals of a random edge set C, independent of the signs, every realization of which meets every u–v path. The network is *tame with* (ρ, w) if

$$\sum_e w_e \, I_e(b) \;\le\; \rho \, S(b) \qquad \text{for every } b \in \mathbb{R}.$$

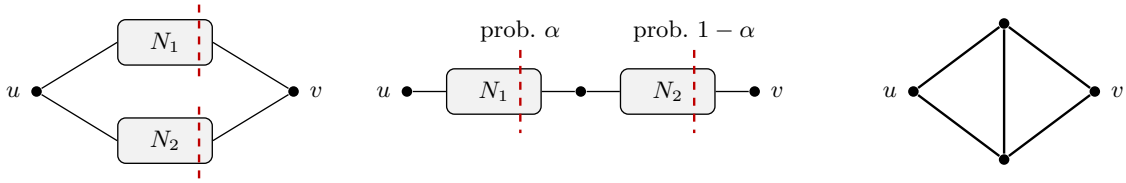
By Lemma 11.3, a single edge is tame with $(y_e, 1)$, using $C = \{e\}$.

Theorem 11.19 (Closure)**.** *Let N_1, N_2 be two-terminal networks, tame with (ρ_1, w^1) and (ρ_2, w^2), with disjoint sign vectors.*

(a) $N_1 \parallel N_2$ is tame with $(\rho_1 + \rho_2,\ w^1 \oplus w^2)$.



(a) a series–parallel network reduces to one edge; ρ in Theorem 11.19 follows the same rules



(b) parallel: $C = C_1 \cup C_2$, $\rho = \rho_1 + \rho_2$

(c) series: $\rho = \frac{\rho_1\rho_2}{\rho_1+\rho_2}$, $\alpha = \frac{\rho_2}{\rho_1+\rho_2}$

(d) the Wheatstone bridge W

Figure 11.4: Series–parallel networks. (a) Conductances combine by the parallel and series rules until one edge remains ($y_{12} = y_1y_2/(y_1+y_2)$, $C_{\text{eff}} = (y_{12}+y_3)y_4/(y_{12}+y_3+y_4)$); the induction of Theorem 11.17 carries the inequality through the same steps. (b),(c) The random separating sets of the closure theorem: in parallel the union of the two cuts, in series one of them, chosen with probability α so that the halves carry equal weight. (d) The Wheatstone bridge, the simplest network that cannot be reduced this way, where the composition argument stops.

(b) $N_1 - N_2$ is tame with $(\rho_1\rho_2/(\rho_1+\rho_2),\ \alpha w^1 \oplus (1-\alpha)w^2)$, where $\alpha = \rho_2/(\rho_1+\rho_2)$ (and $\rho = 0$ with any α if $\rho_1 = \rho_2 = 0$).

The parameter ρ composes exactly as a conductance (Lemma 11.2(ii)), and nothing in the theorem requires the pieces to be series–parallel.

Part (a) follows by averaging first over the signs of one piece, at the shifted coupling the other piece induces; both parts are proved in Section A.24.

The series step is the heart of the matter: N_1 sees the external coupling b reduced by the correlation $\tau_2 = c_2$ across N_2, to $\operatorname{artanh}(\tanh b \cdot \tau_2)$, and the two pieces share the budget $1 - \tanh^2(b + K_N)$. It reduces to the *series-slack inequality* for the Möbius maps $\varphi_a(x) = (a+x)/(1+ax)$, $\lambda_a(t) = t(1-a^2)/(1-a^2t^2)$:

$$\lambda_a(z)^2\big(1-\varphi_{az}(x)^2\big) + \lambda_a(x)^2\big(1-\varphi_{ax}(z)^2\big) \ \le\ 1-\varphi_a(xz)^2 \qquad (a,x,z\in(-1,1)). \tag{11.8}$$

splitting the budget between the two halves, each at the reduced coupling it sees, never exceeds the budget of the whole. The random cut picks N_1 with probability α, so that both halves carry the same weight ρ, which is what (11.8) absorbs.

Proof of Theorem 11.17. By induction on the construction, using Lemma 11.3 for single edges and Theorem 11.19 for the two compositions, every series–parallel network N is tame with $(C_{\text{eff}}(N), w)$ for the marginals w of the composed random set; the parameter equals $C_{\text{eff}}(N)$ because it composes by the conductance rules. Lemma 11.6 for that random set and tameness at $b = 0$ give $\mathcal{Q} \le \sum_e w_e I_e(0) \le C_{\text{eff}}(N)(1-\mathcal{Q})$. □

As a check, on the diamond lattices of Chapter 12, where much more is known for particular laws, the method covers only dimension below two but every symmetric law at once.

Corollary 11.20 (Diamond hierarchical lattices)**.** *Let D_n be the level-n diamond lattice in which each edge is replaced by s parallel chains of b edges, with poles A, B (dimension $1 + \log s/\log b$).*

For every symmetric coupling law and every β with $\bar{y}(\beta) < \infty$, with free boundary conditions,

$$\mathbb{E}\,\langle \sigma_A \sigma_B \rangle_{D_n,J}^2 \;\le\; \bar{y}(\beta)\,(s/b)^n.$$

In particular, when $s < b$ there is no spin-glass order between the poles at any β with $\bar{y}(\beta) < \infty$.

Proof. By Theorem 11.17, $\mathcal{Q} \le C_{\mathrm{eff}}(y)$, and by concavity and Jensen's inequality $\mathbb{E}_K C_{\mathrm{eff}}(y) \le C_{\mathrm{eff}}(\mathbb{E}y) = \bar{y}\, C_{\mathrm{eff}}^{D_n,1} = \bar{y}(s/b)^n$, since D_n is s parallel chains of b copies of D_{n-1}. □

Beyond series–parallel networks, Lemma 11.6 with the level-set cut of the potential, and the closure induction, suggest an inequality that would give RD on every graph.

Conjecture 11.21 (Flow-weighted influence bound)**.** *For every finite multigraph, all magnitudes and $u \ne v$, with f the potential of* (11.3)*,*

$$\sum_{e=xy} |f(x) - f(y)|\, \mathbb{E}_\varepsilon (D_e c)^2 \;\le\; C_{\mathrm{eff}}(u,v)\,(1 - \mathcal{Q}_G(u,v)).$$

By Lemma 11.6 the left side is at least $\mathcal{Q}$, so Conjecture 11.21 implies RD, and even a version with a universal constant factor would give Theorem 11.15 with a constant. On series–parallel networks the composed random set has marginals $|f(x) - f(y)|$, so the conjecture holds there, with the external coupling, at every b, and on the Wheatstone bridge for all magnitudes by an exact computer-assisted certificate (a polynomial in the five variables $\tanh^2 K_e$, nonnegative on the unit cube by Bernstein expansions on five charts; two independent certificates), a finite computation that does not compose.

Observation 11.22 (The per-path form is false)*.* On series–parallel networks the same induction gives $\sum_{e\in P} I_e(b)/y_e \le S(b)$ for every simple u–v path P, which would imply Conjecture 11.21 on every graph through the flow decomposition. It fails on the 3×3 grid between opposite corners: for suitable magnitudes in $[0.5, 8]$, $\sum_{e\in P} I_e(0)/y_e = 1.4154\, S(0)$ along a monotone path through the centre, while at these magnitudes the ratio of the two sides of Conjecture 11.21 is 0.0086 (exact rational arithmetic, two independent implementations; not yet independently reviewed): the path crosses weak edges with almost no potential drop inside a strongly coupled block, so a proof must use the flow weights.

What remains for $\mathbb{Z}^2$ is composition across an annulus, along separators with many vertices; every local device we have tried for it fails, and other approaches to two-dimensional non-order stall there too (Chapters 4 and 15).

11.6 What a comparison model must contain

Each failure of a close relative of the arboreal gas teaches a property a comparison must have.

Route warning 11.23 (The fair average must be used exactly)*. Attractive:* prove FD, or RD, by an argument robust to small changes of the sign law. *Proved false:* fixed signs violate the inequality on two parallel edges, and every positive tilt $\mathbb{P}_\alpha \propto Z_\varepsilon^\alpha$, including the planted and annealed laws, on three parallel edges at large coupling (Propositions 11.5 and 11.14(d)), so no argument that also applies to a slightly biased law proves the fair statement. The proof of Theorem 11.4 uses $\mathbb{E}\tanh X = 0$ and the variance of one fair sign.

The arboreal gas is the $q = 0$ member of the family $\varphi_q(\omega) \propto \prod_{e\in\omega} y_e\, q^{\gamma(\omega)}$ on edge sets $\omega \subseteq E$, $\gamma(\omega)$ the number of independent cycles; $q = 1$ is Bernoulli(t_e^2) percolation, and $q > 0$ the random-cluster measure with edge parameter $qy_e/(1 + qy_e)$ and cluster weight q. Let C_q be the comparison $\mathcal{Q} \le \varphi_q(u \leftrightarrow v)$.

Proposition 11.24 (The arboreal gas is the edge of the family)**.** *(a) For every $q < 0$, C_q fails on two parallel edges, for activities at which the total weight is positive. (b) For $q \in (0, 1]$ and constant weight y with $qy > 1$, on every finite graph φ_q stochastically dominates Bernoulli percolation with parameter $qy/(1 + qy) > \frac{1}{2}$, which is supercritical on $\mathbb{Z}^2$.*

Part (a) is a computation on two parallel edges, and (b) a comparison with Bernoulli percolation [109, Theorem 3.21], where $p_c(\mathbb{Z}^2) = \frac{1}{2}$ [137] (Section A.25). So the cycle weight must be exactly zero: negative weights overshoot on the smallest cycle, and positive ones reintroduce percolation at low temperature, where in case (b) the right side of C_q stays bounded below on $\mathbb{Z}^2$. The forest identity (11.6) extends to $q > 0$ only with an electrical factor, unbounded on subgraphs of $\mathbb{Z}^2$ (Observation A.12; derivations and numerics not yet independently reviewed); the forest proof works as it is exactly 1 at $q = 0$.

Observation 11.25 (A Gaussian sibling fails). Replace the Ising spins by a Gaussian field whose precision matrix is the identity plus a signed graph Laplacian (each edge $e = xy$ adds y_e to the diagonal at x, y and $-\varepsilon_e t_e/(1 - t_e^2)$ off it), so that the normalized correlation $\rho_\varepsilon(u, v) = G_{uv}/(G_{uu}G_{vv})^{1/2}$ is εt on one edge and $\mathbb{E}\rho^2 = \prod t_e^2$ on trees. The analogue $\mathbb{E}_\varepsilon \rho_\varepsilon(u, v)^2 \le \mathbb{P}^{\mathrm{AG}}(u \leftrightarrow v)$ of FD holds on every bundle, the triangle and the four-cycle, but fails on the Wheatstone bridge between the terminals at homogeneous $t = 1/10$, by the ratio 1.00028, where the Ising glass satisfies FD (exact rational arithmetic, two independent implementations; not yet independently reviewed): the comparison must use that the Ising spin is compact.

The last relative is a continuous-spin glass in which nothing cancels exactly: angles θ_x, couplings $K'_e \cos(\theta_x - \theta_y - \phi_e)$ with independent uniform fluxes ϕ_e, and $I_1(K'_e)/I_0(K'_e) = \tanh K_e$, so that on one edge $\mathbb{E}_\phi |\langle e^{i(\theta_u - \theta_v)}\rangle|^2 = t_e^2$ and on trees the two sides agree. The comparison of $\mathcal{Q}$ with this quantity on the finite subgraphs of $\mathbb{Z}^2$ would give a power law at every temperature for $\pm J$ couplings, since the XY side is controlled uniformly in the fluxes by a complex shift of the angles [159] (we omit the details).

Route warning 11.26 (The reference must reproduce exact $\pm J$ cancellations). *Attractive:* compare with the uniform-flux XY glass, whose two-dimensional decay is accessible. *False in general (high-precision numerics, not certified):* the comparison fails on three parallel edges at $t = 0.9$ ($0.85735 > 0.83748$) and, with unequal activities, on a subgraph of $\mathbb{Z}^2$ (three arcs of lengths $3, 5, 5$); these values come from high-precision numerical evaluation by three independent methods, agreeing to 10^{-15}, not from an interval certificate. Heuristically, a frustrated $\pm J$ square of equal magnitudes cancels exactly (the engine of Chapter 13), while a uniform flux never cancels a cycle; Observation A.13 records a numerical failure at homogeneous activity on a ring of squares in $\mathbb{Z}^2$.

11.7 Where the comparison stands

The forest side is understood, the arboreal gas being bounded by its electrical network on every graph through an exact identity, and the glass side one cut at a time and on series–parallel networks, for every symmetric law. The step that would finish the two-dimensional problem, NWB or RD on the finite subgraphs of $\mathbb{Z}^2$, is an open finite-graph inequality with no free

parameter, whose proof must use the fair sign law exactly, the compactness of the Ising spin and the exact cancellation of frustrated cycles. By the proof of Theorem 11.15, the first alternative of Open problem 11.27 gives $\mathbb{E}\langle\sigma_u\sigma_v\rangle^2 \le 8M\bar{y}/\ln(2r+1)$ on $\mathbb{Z}^2$ at every temperature. On the decorated lattices of Chapter 13 exact cancellations on cycles decide the phase, and forest signs alone cannot tell order from disorder. The thread resumes in Chapters 15 and 16.

Open problem 11.27 (The planar comparison)**.** Prove that there is a constant $M < \infty$ such that $\mathcal{Q}_G(u,v) \le M\, C^G_{\mathrm{eff}}(u,v)$ for every finite subgraph G of $\mathbb{Z}^2$, all magnitudes and all $u \ne v$; or prove Conjecture 11.21 on one infinite family of graphs that are not series–parallel and exhaust a recurrent graph.

Part VI

Testbeds and the Nishimori line

Chapter 12

Hierarchical and long-range proxies

On $\mathbb{Z}^d$ with $\pm J$ or Gaussian couplings, neither the existence of spin-glass order at low temperature nor its absence at every positive temperature is proved in any dimension $d \geq 2$. This chapter asks on which graphs, close to $\mathbb{Z}^d$ in some respect, the same model with the same kind of disorder can be decided, and what the proofs there teach about $\mathbb{Z}^d$. Three families answer with theorems. On *Migdal–Kadanoff (MK) hierarchical lattices* the renormalization group is exact and ordering is a question about one distributional recursion: for Gaussian couplings, order at low temperature is proved on every MK lattice of fractal dimension at least three and excluded below 2.58, while for $\pm J$ couplings an exact cancellation excludes order at every temperature on a family whose dimension approaches three from below. On *Dyson's hierarchical lattice* order reduces exactly to a sum of sign losses over scales, controlled in a perturbed model only under an open hypothesis. On *one-dimensional chains* with power-law couplings, coupling-dependent boundary conditions produce uncountably many Gibbs states at every temperature, and on the Nishimori line a phase transition is proved for decay exponents between 1 and 3/2.

Each proof uses a structure that $\mathbb{Z}^d$ lacks: closure of the renormalization on pair couplings (Proposition 12.9), a fractal dimension that does not behave like d (Route warning 12.11), fluctuation bounds that stop at the Imry–Ma threshold (Route warning 12.18). The reader will be able to derive the exact MK recursion, prove the dilution criterion and the four-arm-star obstruction, and state the MK order theorems and the exotic-state theorem with their hypotheses.

12.1 Migdal–Kadanoff lattices and their exact recursion

The Migdal–Kadanoff approximation replaces $\mathbb{Z}^d$ by a graph on which the bond-moving renormalization is exact. The graph is built by repeatedly replacing every edge by a small two-terminal network.

Definition 12.1 (MK lattice)**.** Fix integers $b \geq 2$ (the *scale factor*) and $s \geq 2$ (the *branching number*). The lattice D_0 is one edge between two vertices A and B, the *poles*. For $n \geq 1$, D_n consists of s parallel *branches* between A and B, each a chain of b copies of D_{n-1} glued in series. The integer n is the *level*. Equivalently, D_n is D_{n-1} with every edge replaced by a copy of D_1. The graph distance between the poles is b^n, the number of edges is $(bs)^n$, and the *fractal dimension* is $d_{\mathrm{MK}} = \log(bs)/\log b = 1 + \log s/\log b$. The case $b = 2$ is the *diamond lattice* with s branches; $s = 2^{d-1}$ is the MK approximation of $\mathbb{Z}^d$, so $(b, s) = (2, 4)$ is the three-dimensional diamond lattice.

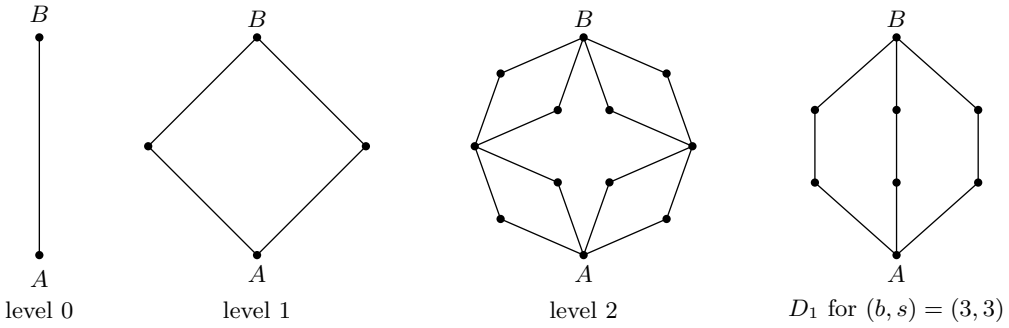


Figure 12.1: The diamond lattice $(b,s)=(2,2)$ at levels 0, 1, 2, and the unit D_1 for $(b,s)=(3,3)$. Each level replaces every edge by a copy of D_1; the poles have degree s^n, and every other vertex sits in a nested chain of sub-lattices, which makes the renormalization exact.

Figure 12.1 shows the first levels. Collet and Eckmann and [219] fix $b=2$ and write n for the number of branches; we always write (b,s), as Okuyama and Ohzeki do.

Couplings $J=(J_e)$ are iid with law ν, and the Gibbs measure on D_n has *free poles*: all spins, including σ_A and σ_B, are summed. Every copy D of D_k inside D_n is a *sub-lattice of level k*, with its own poles a_D, b_D. For $\eta,\eta'\in\{-1,1\}$ let $Z_D(\eta,\eta')$ be the Boltzmann sum over the spins of D other than its poles, with pole spins (η,η') and only the couplings of D. By the global flip, $Z_D(+,+)=Z_D(-,-)$ and $Z_D(+,-)=Z_D(-,+)$, so

$$Z_D(\eta,\eta')=C_D\,e^{\xi_D\eta\eta'},\qquad \xi_D:=\tfrac12\log\frac{Z_D(+,+)}{Z_D(+,-)}. \tag{12.1}$$

The number ξ_D is the *effective coupling* of D; it is dimensionless (it includes β). Write ξ_k for a random variable with the law of ξ_D for D of level k, and put

$$a\star a':=\operatorname{artanh}(\tanh a\tanh a'),\qquad a_1\star\cdots\star a_b:=\operatorname{artanh}\Big(\prod_{j=1}^{b}\tanh a_j\Big).$$

Lemma 12.2 (Exact recursion)**.** *(a) If D is a chain of sub-lattices with effective couplings $\eta_1,\dots,\eta_b$ glued in series, then $\xi_D=\eta_1\star\cdots\star\eta_b$; if D consists of branches with effective couplings $\eta_1,\dots,\eta_s$ in parallel, then $\xi_D=\eta_1+\cdots+\eta_s$. Consequently*

$$\xi_n\overset{d}{=}\sum_{i=1}^{s}\ \bigstar_{j=1}^{b}\ \xi_{n-1}^{(i,j)},\qquad \xi_0=\beta J_e, \tag{12.2}$$

with bs independent copies of ξ_{n-1} on the right.

(b) With free poles, $\langle\sigma_A\sigma_B\rangle_J=\tanh\xi_{D_n}$; hence $\mathbb{E}\langle\sigma_A\sigma_B\rangle_J^2=\mathbb{E}\tanh^2\xi_n$, and pole order, $\mathbb{E}\langle\sigma_A\sigma_B\rangle_J^2\to1$, holds if and only if $|\xi_n|\to\infty$ in probability.

(c) One has $\operatorname{sgn}(a\star a')=\operatorname{sgn}(aa')$ and $|a\star a'|\le\min(|a|,|a'|)$. As $\beta\to\infty$, $\beta^{-1}(\beta x\star\beta y)\to\operatorname{sgn}(xy)\min(|x|,|y|)$: at zero temperature a chain transmits its weakest bond, with the product of the signs.

Proof. (a) For two bonds in series with middle spin τ, $\sum_{\tau=\pm1}e^{a\eta\tau+a'\tau\eta'}=2\cosh(a\eta+a'\eta')=2\cosh a\cosh a'\,(1+\tanh a\tanh a'\,\eta\eta')$, and $1+t\eta\eta'=e^{c\eta\eta'}/\cosh c$ when $\tanh c=t$. Induction

on b gives the chain; parallel branches share only the poles, so their pole weights multiply and the exponents add. For (12.2), sum first over the interior spins of each of the bs copies of D_{n-1}; these interiors are disjoint and carry disjoint (hence independent) couplings, and what remains is D_1 with effective couplings $\xi_{n-1}^{(i,j)}$. (b) With free poles, $\mu(\sigma_A\sigma_B = \pm 1) \propto Z_{D_n}(+,\pm)$, so $\langle\sigma_A\sigma_B\rangle_J = \tanh \xi_{D_n}$. (c) $|\tanh a \tanh a'| \le \min(|\tanh a|, |\tanh a'|)$, and artanh is odd and increasing. From $1 \pm \tanh u \tanh v = \cosh(u \pm v)/(\cosh u \cosh v)$ one gets $u \star v = \frac{1}{2}\log\frac{\cosh(u+v)}{\cosh(u-v)}$, so $\beta^{-1}(\beta x \star \beta y) \to \frac{1}{2}(|x+y| - |x-y|) = \operatorname{sgn}(xy)\min(|x|,|y|)$. □

The recursion is in [59, §III] for $b = 2$, and in [202, (2)–(5)] for $J_n(\beta) = \xi_n/\beta$. It reduces *pole* order to one sequence of laws. Bulk statements (spins inside the lattice, the site overlap) need one more idea, used in every proof below: a vertex of level j is the midpoint of a branch of a sub-lattice of level j, which sits in a unique chain of nested sub-lattices of levels $j, \dots, n$ (its *ancestor chain*). Given the pole spins of the sub-lattice of level k, the next spin down feels only the two effective couplings of the halves of its branch, so a union bound along the chain transfers coupling growth to the bulk [59, §VII].

12.2 No order by exact dilution

The simplest MK theorem is a non-order theorem with a one-page proof. For $\pm J$ couplings the branch contributions in (12.2) after one step all have the same magnitude, so with an even number of branches they cancel exactly with positive probability; cancelled bonds are absent, and absence of a path from A to B forces $\langle\sigma_A\sigma_B\rangle_J = 0$. What remains is a percolation question one level down.

Theorem 12.3 (Okuyama–Ohzeki [201, Thms. 1–3])**.** *Let $b \ge 2$, let $s \ge 2$ be even, let the couplings be iid $\pm J$ with $\nu = \frac{1}{2}(\delta_{+1} + \delta_{-1})$, and put $\zeta_s = \binom{s}{s/2}2^{-s}$. Suppose*

$$\zeta_s < \left[1 - (1-\zeta_s)^b\right]^s. \tag{12.3}$$

Then there are $C, c > 0$ such that for every $n \ge 1$ and every $\beta \in (0,\infty)$, with free poles,

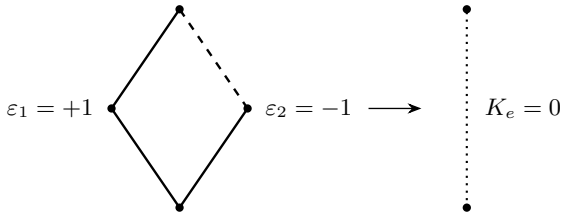
$$\mathbb{E}\langle\sigma_A\sigma_B\rangle_J^2 \le Ce^{-cb^n}, \qquad \mathbb{P}(\xi_{D_n} = 0) \ge 1 - Ce^{-cb^n}.$$

The bounds are uniform in β, so they also hold for the limit $\beta \to \infty$. In particular there is no pole order at any temperature, including zero temperature, and the free-energy difference between the pole conditions $(+,+)$ and $(+,-)$ vanishes with probability tending to one (no stiffness).

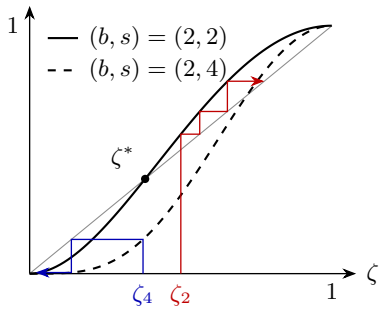
Proof. Step 1: one exact step creates holes. View D_n as D_{n-1} with each edge e replaced by a unit U_e, a copy of D_1. Sum over the $s(b-1)$ interior spins of every unit. By Lemma 12.2(a) the unit becomes one bond with coupling

$$K_e = \sum_{i=1}^{s} \bigstar_{j=1}^{b} \beta J_{ij} = \operatorname{artanh}\big((\tanh\beta)^b\big) \sum_{i=1}^{s} \varepsilon_i, \qquad \varepsilon_i = \prod_{j=1}^{b} J_{ij},$$

where J_{ij} are the couplings of U_e. The ε_i are iid fair signs, so $K_e = 0$ exactly when half of them are $+1$, with probability ζ_s, independently of β and b. Distinct units carry disjoint couplings, so the events $\{K_e = 0\}$ are independent. The marginal of the Gibbs measure on the vertices



(a) $s = 2$, $b = 2$: the two branch contributions cancel exactly



(b) vacancy map $\zeta \mapsto [1 - (1 - \zeta)^b]^s$

Figure 12.2: The dilution mechanism. (a) With $\pm J$ couplings every branch transmits the magnitude $\operatorname{artanh}(\tanh^b \beta)$ with an independent fair sign, so an even number of branches cancels with probability ζ_s at every temperature. (b) The vacancy probability evolves by the hierarchical percolation map: for $(2, 2)$, $\zeta_2 = \frac{1}{2}$ lies above the unstable fixed point $(3 - \sqrt{5})/2$ and flows to 1 (red, no order); for the three-dimensional diamond $(2, 4)$, $\zeta_4 = \frac{3}{8}$ flows to 0 (blue) and the criterion says nothing.

of D_{n-1} is the Ising model on D_{n-1} with couplings (K_e), and the bonds with $K_e = 0$ can be deleted.

Step 2: disconnection kills correlation and stiffness. Let $\omega = \{e : K_e \neq 0\}$. If A and B lie in different components of $(V(D_{n-1}), \omega)$, the measure factorizes over components, and each factor is invariant under the global flip of its component (free poles, no field). Hence $\langle \sigma_A \sigma_B \rangle_J = \langle \sigma_A \rangle_J \langle \sigma_B \rangle_J = 0$, and flipping the component of B gives $Z_{D_n}(+, +) = Z_{D_n}(+, -)$, that is $\xi_{D_n} = 0$. Since $|\langle \sigma_A \sigma_B \rangle_J| \leq 1$, we get $\mathbb{E}\langle \sigma_A \sigma_B \rangle_J^2 \leq x_{n-1}$ and $\mathbb{P}(\xi_{D_n} \neq 0) \leq x_{n-1}$, where x_k is the probability that A and B are connected in Bernoulli bond percolation of density $1 - \zeta_s$ on D_k.

Step 3: the hierarchical percolation recursion. A and B are connected in D_k if and only if some branch has all b of its sub-lattices connected, and the bs sub-lattices are independent, so $x_k = f(x_{k-1})$ with $f(x) = 1 - (1 - x^b)^s$ and $x_0 = 1 - \zeta_s$. Condition (12.3) says exactly $f(x_0) < x_0$. The map f is increasing with $f(0) = 0$, $f(1) = 1$, and $f(x) \leq s x^b < x$ near 0; Okuyama and Ohzeki show that f has exactly one fixed point $x^* \in (0, 1)$, by writing the fixed-point equation as $\log(1 - x)/\log(1 - x^b) = s$ and checking that the left side decreases strictly from ∞ to 1 [201, (17)–(19)]. Hence $f(x) < x$ on $(0, x^*)$ and $f(x) > x$ on $(x^*, 1)$, so $x_0 < x^*$, the sequence x_k decreases, and its limit is a fixed point below x^*, namely 0. For the rate, the union bound $f(x) \leq s x^b$ gives, for $y_k = s^{1/(b-1)} x_k$, the recursion $y_{k+1} \leq y_k^b$. Choose k_0 with $\lambda := y_{k_0} < 1$; then $x_{k_0+m} \leq s^{-1/(b-1)} \lambda^{b^m}$, which is the claimed $C e^{-c b^n}$. The bounds of Step 2 do not depend on β (at $\beta = \infty$, Lemma 12.2(c) gives $K_e = \sum_i \varepsilon_i$ in energy units, which vanishes with the same probability). □

Figure 12.2 shows the two ingredients. The consequences are all in [201]. Condition (12.3) holds for $s = b$ ($d_{\mathrm{MK}} = 2$) and every even s [201, p. 3], so the MK approximation of the square-lattice $\pm J$ glass has no order at any temperature; it holds for $(b, s) = (5, 6)$ and $(6, 8)$ ($d_{\mathrm{MK}} \approx 2.11, 2.16$), and fails for $(3, 6)$, $(4, 6)$, $(3, 4)$ and every diamond lattice $(2, s)$ with $s \geq 4$. For large even s and b of order $\sqrt{\pi s/2}\,(\log s - \log\log s + C)$ with $C > \log 2$, it holds and $d_{\mathrm{MK}} = 3 - (4 \log\log s + O(1))/\log s$ [201, (9)]: there are MK lattices of dimension arbitrarily

close to three from below with no $\pm J$ order at any temperature. Two exact steps instead of one give a vacancy probability above the threshold on the self-dual hierarchical lattice of dimension $\log 5/\log 2$ (the diamond with a bond between the two midpoints) [201, Thm. 4]. The mechanism needs an atom of the effective coupling at 0, so it gives nothing for odd s or for continuous laws such as the Gaussian [201, p. 4]. On $\mathbb{Z}^d$ it is applied to diluted couplings in Chapter 13.

12.3 Order by coupling growth

Order on an MK lattice means that $|\xi_n|$ grows. The difficulty is visible already in (12.2): a sum of s independent symmetric terms grows like $\sqrt{s}$ times its typical size, while the series map shrinks. Collet and Eckmann made this precise for many branches; the obstruction to doing it for few branches is anti-concentration, that is, excluding that ξ_n stays near 0 with large probability.

Theorem 12.4 (Collet–Eckmann [59, Thm. 4.1 and Thm. 7.1])**.** *Let $b = 2$, let the couplings be iid with an even law of variance 1 and a bounded density, and let $Q = \mathbb{E}[\min(J^2, J'^2)]/8$ with J, J' independent copies. There is $N_0 = N_0(Q)$ such that for every even $s > N_0$ and every sufficiently large β, with the pole spins fixed, every vertex x of D_n other than the poles satisfies $\mathbb{E}\langle\sigma_x\rangle_J = 0$ and $\mathbb{E}\langle\sigma_x\rangle_J^2 > \frac{1}{2}$, for every level n, and $\mathbb{E}\langle\sigma_x\rangle_J^2 \to 1$ as $\beta \to \infty$.*

The proof bounds the variance of ξ_k by a Berry–Esseen estimate in s and its density by a local limit theorem, then runs the ancestor-chain argument; $N_0 = \max(1600/Q^3, (40 \cdot 320^3)^2) > 1.7 \times 10^{18}$ [59, p. 386].

For Gaussian couplings the question is now settled for every (b, s) except possibly one marginal case. Let G_n be the zero-temperature effective coupling, obtained from (12.2) with the series map replaced by its limit in Lemma 12.2(c) and $G_0 \sim N(0,1)$, let g_n be its density and φ the standard normal density, and put

$$\alpha_n = \frac{\varphi(0)}{g_n(0)}, \qquad \Lambda(b,s) = \lim_{n\to\infty} \alpha_n^{1/n} = \sup_{n\geq 1} \alpha_n^{1/n}.$$

The limit exists because α_n is supermultiplicative [202, §III.D]. A large Λ means that the central density of the zero-temperature coupling decays fast, so G_n escapes from 0. Write $q_n = |V(D_n)|^{-1} \sum_x \sigma_x \tau_x$ for the site overlap of two replicas with free poles.

Theorem 12.5 (Okuyama–Ohzeki [202, Thms. 1–3])**.** *Let $b, s \geq 2$ and let the couplings be iid $N(0,1)$.*

(a) If $\Lambda(b,s) \leq 1$, then $\lim_n \mathbb{E}\langle q_n^2\rangle = 0$ for every $\beta \in (0,\infty)$. This holds in particular if $6s < (b+1)(b+2)$, or $\pi s < b^2$, or $24s \leq 7(b+1)(b+2)$ with $b \geq 6$; and $\Lambda(b,s) < 1$ whenever $d_{\mathrm{MK}} < 1 + \log 17/\log 6 \approx 2.581$.

(b) If $\Lambda(b,s) > 1$, there is $\beta_{\mathrm{SG}} < \infty$ such that $\liminf_n \mathbb{E}\langle q_n^2\rangle > 0$ for every $\beta > \beta_{\mathrm{SG}}$. This holds in particular if $\frac{b}{\sqrt{s}}\big(\frac{b\sqrt{s(s-1)}}{bs-1}\big)^{s-1} < 1$, which is the case for every (b,s) with $d_{\mathrm{MK}} \geq 3$.

(c) With $C(b,s) = 4\lceil \log_2(bs)\rceil$: if $\alpha_n > 1$ for some n then $\Lambda > 1$, and if $C(b,s) n \alpha_n < 1$ for some n then $\Lambda < 1$. If $\Lambda(b,s) \neq 1$, one of the two holds for every n with $e^{n|\log\Lambda|} > C(b,s)n$.

So Λ classifies the Gaussian MK lattices, part (c) decides the sign of $\Lambda - 1$ by a finite computation whenever $\Lambda \neq 1$, and part (a) covers $\Lambda = 1$, whose occurrence is open (Open problem 12.20).

Proof sketch. We omit the analysis of Λ (part (c) and the explicit conditions). *Absence.* At zero temperature both maps preserve symmetric unimodality and the peakedness order, and a Gaussian scale-mixture comparison controls $g_n(0)$; positive temperature only shrinks couplings further, and small effective couplings make $\mathbb{E}\langle q_n^2\rangle$ small [202, §IV]. *Existence.* Build a reference $\eta = \eta(J,\beta)$ from the top down: $\eta_A = +1$, $\eta_B = \operatorname{sgn}\xi_{D_n}$, and inside every sub-lattice the interior skeleton spins that minimize its effective Hamiltonian given its pole spins (on a branch whose effective bonds cannot all be satisfied, exactly the weakest one is left unsatisfied). A Peierls estimate in one cell bounds the probability of a thermal deviation by the sup norm $D_k(\beta)$ of the density of ξ_k/β, and a union bound over the ancestor chain gives, uniformly over vertices v and levels n,

$$\mathbb{E}\,\mu_n(\sigma_v \neq \eta_v\sigma_A) \le \delta_n(\beta) := \frac{\log 2}{\beta}D_n(\beta) + \frac{sb(b-1)}{\beta}\sum_{k=0}^{n-1} D_k(\beta) \tag{12.4}$$

[202, Lemma 19, (82)]. If $\Lambda > 1$, the densities decay geometrically and $\sup_n \delta_n(\beta) \le C_\delta/\beta$ for $\beta \ge \beta_0$ [202, Lemmas 20–23]; then $\mathbb{E}\langle q_n^2\rangle \ge (1-4\delta_n)_+^2$. □

The proof stays in the exact-recursion line of Collet and Eckmann: control of the density at zero, through the peakedness order, takes the place of their local limit theorem, and the existence half is a reference configuration with a one-cell Peierls estimate, the same structure as in [219].

The estimate (12.4) is the hypothesis of the reference-alignment corollary of Chapter 8, which turns order against a reference into a giant blue cluster and an imbalance between the overlap sectors of the CMR colouring; so every Gaussian MK lattice that orders has blue imbalance.

Corollary 12.6 (Blue imbalance on Gaussian MK lattices)**.** *Let $b, s \ge 2$ with $\Lambda(b,s) > 1$, iid $N(0,1)$ couplings and free poles, and let β_0, C_δ be as in (12.4). For $\beta > \max(\beta_0, 4C_\delta)$ and every level n, in the CMR colouring of two replicas,*

$$\frac{\mathbb{E}|C_{\mathrm{b}}(A)|}{|V(D_n)|} \ge 1 - \frac{2C_\delta}{\beta}, \qquad \mathbb{E}I_n^2 \ge \Big(1 - \frac{4C_\delta}{\beta}\Big)^2, \qquad \mathbb{E}\langle q_n^2\rangle \ge \Big(1 - \frac{2C_\delta}{\beta}\Big)^4,$$

where $C_{\mathrm{b}}(A)$ is the blue cluster of A and $I_n = (|C_+| - |C_-|)/|V(D_n)|$ the imbalance of the largest blue clusters in the two overlap sectors.

Proof. The reference η of (12.4) depends on (J,β) only and has $\eta_A = 1$; it satisfies the hypothesis of the reference-alignment corollary of Chapter 8 with $a = A$ and $\delta = C_\delta/\beta < \frac{1}{4}$, whose three conclusions are the three bounds. □

For the diamond lattices $b = 2$ a second route covers every law with a bounded density; its input is a one-step anti-concentration estimate in place of Collet and Eckmann's local limit theorem.

Theorem 12.7 ([219, Thm. A])**.** *Let $b = 2$, $s \ge 4$, and let the couplings be iid with a density bounded by $\varphi_{\max} < \infty$ (no symmetry or moment condition). There are $\beta_0 = \beta_0(s, \varphi_{\max}) < \infty$ and $C = C(s, \varphi_{\max}) < \infty$ such that for every $\beta \ge \beta_0$ and every level $n \ge 1$:*

(i) with free poles, $\mathbb{E}\langle\sigma_A\sigma_B\rangle_J^2 \ge 1 - C/\beta$, and $\mathbb{E}\langle\sigma_A\sigma_B\rangle_J^2 \to 1$ as $n \to \infty$;

(ii) for every choice of fixed pole spins and every vertex x other than the poles, $\mathbb{E}\langle\sigma_x\rangle_J^2 \ge 1 - C/\beta$;

(iii) with free poles, $\mathbb{E}\langle q_n^2\rangle \ge 1 - C/\beta$;

(iv) with free poles, in the CMR colouring, $\mathbb{P}(B \in C_{\mathrm{b}}(A),\ |C_{\mathrm{b}}(A)| > (1-\theta)|V(D_n)|) \geq 1 - C/(\theta\beta)$ *for* $\theta \in (0,1)$*; the largest blue cluster not containing* A *has density tending to* 0 *in probability and in* L^1 *as* $n \to \infty$*; and* $\mathbb{E}I_n^2 \geq 1 - C/\beta$.

For standard Gaussian couplings and $s = 4$*, explicit versions of these bounds are positive for* $\beta \geq 520$*, verified in exact rational arithmetic.*

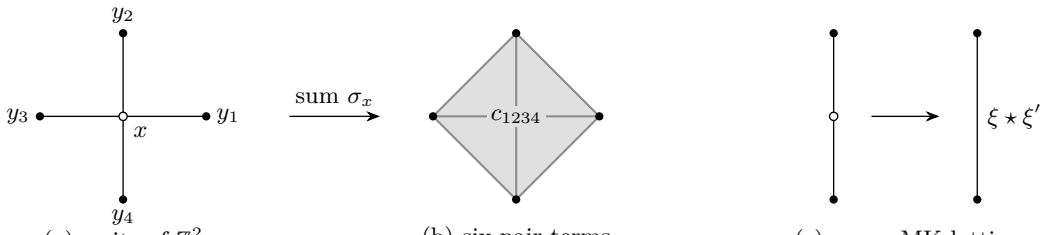
Proof sketch. We describe the coupling-growth step and omit the transfer and the blue-cluster bookkeeping, which follow the ancestor-chain argument above and the crossing lemma of Chapter 8. Track per level a bound r_k on the density of ξ_k restricted to an event E_k determined by one sub-lattice of level k, and a bound w_k on $\mathbb{P}(E_k^c)$. A branch is *good* if both halves lie in their events and their moduli are well separated. Given the larger half, the smaller one is confined to a range of probability W with density at most $r_k/(\mathbb{P}(E_k)W)$ there; at zero temperature the branch transmits the smaller half up to sign (Lemma 12.2(c)), and averaging over the larger half cancels $1/W$ against the probability of the range. So a good branch has density at most $r_k\, p_2(r_k\,\cdot)$, $p_2(x) = 2(1-2|x|)_+$, and Rogozin's inequality (the density of a sum of independent variables is at most that of a sum of uniform variables with the same density bounds) bounds the density of the sum over g good branches by $r_k\kappa_g$, where $\kappa_g = p_2^{*g}(0)$: $\kappa_1 = 2$, $\kappa_2 = \frac{4}{3}$, $\kappa_3 = \frac{11}{10}$, $\kappa_4 = \frac{302}{315} < 1$. Bad branches only shift the sum. For $s \geq 4$ and $r_0 = \varphi_{\max}/\beta$ small, r_k and w_k decay geometrically, which forces $|\xi_k| \to \infty$ and supplies the small-ball bounds the transfer needs [219, §§4–5]. At zero temperature $r_k\kappa_3$ is attained, so one-step density bounds cannot treat $s = 3$. □

For Gaussian couplings near the critical temperature, the same transfer runs on a computer-assisted growth input.

Theorem 12.8 (computer-assisted; [219, Thm. B and Cor. 2.1])**.** *Let* $b = 2$*, let the couplings be iid* $N(0,1)$*, and let* (s, T^*) *be one of* $(3, \frac{4}{15})$, $(4, \frac{6}{7})$, $(5, \frac{6}{5})$, $(7, \frac{7}{4})$, $(8, 2)$*. For every* $T = 1/\beta \leq T^*$*, with free poles: for each* $u \geq 0$, $\mathbb{P}(|\xi_n| \leq u)$ *decays geometrically in* n*;* $\mathbb{E}\langle\sigma_A\sigma_B\rangle_J^2 \to 1$*;* $\liminf_n \mathbb{E}\langle q_n^2\rangle \geq \kappa_Q^2$*;* $\liminf_n \mathbb{E}|C_{\mathrm{b}}(A)|/|V(D_n)| \geq \kappa_C$ *while the largest other blue cluster has expected density tending to* 0*; and* $\liminf_n \mathbb{E}I_n^2 \geq \kappa_C^2$*, with explicit constants* $\kappa_Q, \kappa_C > 0$*. The same conclusions hold on the diamond lattice with any* $s' \geq s$ *branches at the same temperatures; in particular for* $s' = 6$ *at* $T \leq \frac{6}{5}$ *and for every* $s' \geq 9$ *at* $T \leq 2$.

The certificate is a finite list of inequalities between explicit laws, checked in exact rational and interval arithmetic. For symmetric unimodal inputs the one-step map preserves symmetric unimodality and is monotone in the peakedness order (X is more peaked than Y if $\mathbb{P}(|X| \leq t) \geq \mathbb{P}(|Y| \leq t)$ for all t); the computer checks that explicit staircase laws, each more peaked than the true law at its level, reach a family that the map sends into itself with growth factor $\lambda > 1$, itself a finite check [219, §§6–7]. The only temperature-dependent condition is that the first law be more peaked than that of $\beta|J_e|$, so a certificate at T^* covers every lower temperature. The $T^*(s)$ are 0.89 to 0.98 times population-dynamics estimates of $T_c(s)$ (a heuristic comparison), the constants are astronomically small ($\kappa_C \geq 10^{-28.2}$ at $s = 4$), and for $s = 3$, $d_{\mathrm{MK}} = \log_2 6 \approx 2.585$ is just above the threshold 2.581 of Theorem 12.5(a).

On D_n the imbalance is degenerate, one giant blue cluster containing both poles and dust [219, §1.5], while on $\mathbb{Z}^d$ both sectors are expected to percolate in an ordered phase and imbalance is open (Chapter 8). For $\pm J$ couplings, which have no density, order on the diamond lattices is open for every $s \geq 3$ [219, §1.1], and Theorem 12.3 excludes it only for $s = 2$.



(a) a site of $\mathbb{Z}^2$ and its star

(b) six pair terms and a four-spin term

(c) on an MK lattice: one bond

Figure 12.3: Why the MK recursion does not close on $\mathbb{Z}^d$. (a)–(b) Summing out one spin of $\mathbb{Z}^2$ couples its four neighbours by six pair terms and a four-spin term, strictly negative for equal bonds (Proposition 12.9). (c) On an MK lattice a summed spin has two neighbours, and the result is again one pair coupling.

12.4 What does not transfer to the cubic lattice

The MK proofs are about one scalar recursion with independent inputs. Two features make that possible, and both fail on $\mathbb{Z}^d$. The first is closure: on an MK lattice, summing out the interior of a sub-lattice produces a pair coupling between its two poles and nothing else, independent of those of disjoint sub-lattices. On $\mathbb{Z}^d$ summing out even one spin couples all of its $2d$ neighbours.

Proposition 12.9 (The four-arm star)**.** *Let x be a vertex whose only neighbours are $y_1, \dots, y_4$, joined to x by bonds with dimensionless couplings $K_1, \dots, K_4$. Then, for $\eta \in \{-1, 1\}^4$ the spins at $y_1, \dots, y_4$,*

$$\sum_{\sigma_x = \pm 1} e^{\sigma_x \sum_i K_i \eta_i} = \exp\Big(c_\emptyset + \sum_{i<j} c_{ij} \eta_i \eta_j + c_{1234}\, \eta_1 \eta_2 \eta_3 \eta_4 \Big),$$

$$c_{1234} = \frac{1}{16} \sum_{\eta \in \{\pm 1\}^4} \eta_1 \eta_2 \eta_3 \eta_4 \log\Big(2 \cosh \sum_i K_i \eta_i \Big).$$

If $K_1 = \cdots = K_4 = K \neq 0$, then $c_{1234} = \frac{1}{8}[\log\cosh 4K - 4 \log\cosh 2K] < 0$. Consequently, for iid Gaussian couplings and every $\beta > 0$, $c_{1234} \neq 0$ with positive probability; the same holds for the coefficient of any product of four arm spins when x has $2d \geq 4$ neighbours, as on $\mathbb{Z}^d$.

Proof. The function $\eta \mapsto \log 2\cosh(\sum_i K_i \eta_i)$ is even under $\eta \mapsto -\eta$, so its expansion in the characters $\eta_S = \prod_{i \in S} \eta_i$ contains only even $|S|$, and orthogonality gives the formula for c_{1234}. With equal couplings, $\sum_i \eta_i \in \{\pm 4, \pm 2, 0\}$ occurs for 2, 8 and 6 configurations with $\eta_1\eta_2\eta_3\eta_4 = +1$, -1 and $+1$, so $c_{1234} = \frac{1}{16}[2 \log 2\cosh 4K - 8 \log 2\cosh 2K + 6 \log 2] = \frac{1}{8}[\log\cosh 4K - 4 \log\cosh 2K]$. With $c = \cosh 2K > 1$, $\cosh^4 2K - \cosh 4K = c^4 - 2c^2 + 1 = (c^2 - 1)^2 > 0$, so $c_{1234} < 0$. The coefficient is continuous in $(K_1, \dots, K_4)$, hence nonzero on a neighbourhood of (K, K, K, K), which has positive Gaussian probability. With $2d$ arms the coefficient of $\eta_1\eta_2\eta_3\eta_4$ is continuous in all $2d$ couplings and reduces to c_{1234} when the other arms vanish. □

Route warning 12.10 (Lifting the scalar MK recursion to $\mathbb{Z}^d$). *Tempting step:* prove order on $\mathbb{Z}^3$ by a renormalization of independent pair couplings, as on the diamond lattice with $s = 4$. *Status: the needed closure is false.* Exact elimination of spins on $\mathbb{Z}^d$ produces multi-spin interactions with positive probability (Proposition 12.9, Fig. 12.3), and the pair couplings it produces are

correlated through shared bonds. This excludes an exact scalar recursion; it does not exclude order on $\mathbb{Z}^d$, a renormalization that tracks boundary interactions, or approximate comparisons.

The second feature is the geometry. The poles of D_n have degree s^n and high-level vertices have unbounded degree, while every vertex of $\mathbb{Z}^d$ has degree $2d$; the fractal dimension, a count of edges against distance, does not control the ordering question.

Route warning 12.11 (The MK dimension is not a proxy for d). *Tempting step:* read an MK theorem at $d_{\mathrm{MK}} = 3$ (or < 3) as evidence for order (or its absence) on $\mathbb{Z}^3$. *Status: proved misleading within the MK family.* For $\pm J$ couplings there are MK lattices with $d_{\mathrm{MK}} \to 3^-$ and no order at any temperature (Theorem 12.3 and the remarks after it). For Gaussian couplings, order holds for every (b, s) with $d_{\mathrm{MK}} \geq 3$, but for every fixed $d_{\mathrm{MK}} < 3$ the condition $\pi s < b^2$ of Theorem 12.5(a) holds once b is large, while the diamond with $s = 3$ orders at $d_{\mathrm{MK}} \approx 2.585$ (Theorem 12.8). So the MK lower critical dimension depends on the construction and not on d_{MK} alone, as Okuyama and Ohzeki also stress [201, pp. 3, 5].

12.5 Dyson's hierarchical spin glass

Dyson's hierarchical lattice keeps the block structure of a renormalization but connects all pairs inside a block. Castellana, Barra and Guerra define the hierarchical Edwards–Anderson model (HEA) on 2^n sites $\{0, \dots, 2^n - 1\}$ by the recursion

$$H_{k+1}(\sigma) = H_k(\sigma^{(1)}) + H_k(\sigma^{(2)}) - 2^{-(k+1)\varrho} \sum_{i<j\leq 2^{k+1}} J_{ij}^{(k+1)} \sigma_i \sigma_j, \qquad H_0 = 0, \tag{12.5}$$

where $\sigma^{(1)}, \sigma^{(2)}$ are the two halves, the $J_{ij}^{(k)}$ are iid $N(0,1)$, and ϱ is a decay exponent [42, Def. 1] (written σ there). For $\varrho > \frac{1}{2}$ the quenched free energy, with an independent Gaussian random field allowed, has a thermodynamic limit [42, Thm. 1]; the interesting range is $\frac{1}{2} < \varrho < 1$. No order theorem is known, only an exact reduction of order to a quantity on each scale.

Call the dyadic intervals of length 2^k the *blocks of level* k; each block B of level $k \geq 1$ has two halves $B^{(1)}, B^{(2)}$ of level $k-1$. For two configurations σ, τ let $q_i = \sigma_i \tau_i$, $m(B) = |B|^{-1} \sum_{i\in B} q_i$, and $r_1 = m(B^{(1)})$, $r_2 = m(B^{(2)})$. The *sign loss* of B is

$$\ell_B = \min(|r_1|, |r_2|)\, \mathbf{1}\{r_1 r_2 < 0\}, \qquad \bar{\ell}_k = 2^{-(n-k)} \sum_{B \text{ of level } k} \ell_B.$$

Lemma 12.12 (Magnitude telescoping)**.** *For every pair of configurations on 2^n sites, $|q| = 1 - \sum_{k=1}^n \bar{\ell}_k$, where $q = m(\{0, \dots, 2^n - 1\})$, and $0 \leq \ell_B \leq \sqrt{(r_1 r_2)_-}$. Consequently, for every law of (σ, τ), $\mathbb{E}|q| = 1 - \sum_{k=1}^n \mathbb{E}\bar{\ell}_k$, and for a sequence of laws indexed by n, $\liminf_n \mathbb{E}q^2 > 0$ if and only if $\limsup_n \sum_{k=1}^n \mathbb{E}\bar{\ell}_k^{(n)} < 1$.*

Proof. Since $m(B) = (r_1 + r_2)/2$, $|m(B)| = (|r_1| + |r_2|)/2$ if $r_1 r_2 \geq 0$ and $|m(B)| = (|r_1| + |r_2|)/2 - \min(|r_1|, |r_2|)$ otherwise; that is, $|m(B)| = (|r_1| + |r_2|)/2 - \ell_B$. Let A_k be the average of $|m(B)|$ over the 2^{n-k} blocks B of level k. Every block of level $k-1$ is a half of exactly one block of level k, so averaging gives $A_k = A_{k-1} - \bar{\ell}_k$. Since $A_0 = 1$ and $A_n = |q|$, the identity follows. The bound uses $\min(|r_1|, |r_2|) \leq \sqrt{|r_1 r_2|}$. The last statement follows from $(\mathbb{E}|q|)^2 \leq \mathbb{E}q^2 \leq \mathbb{E}|q|$. □

So the overlap is lost scale by scale, a sign disagreement between the halves of a block costing only the overlap of the minority half, and any multiscale proof of order on this hierarchy has

to bound $\sum_k \mathbb{E}\bar{\ell}_k$ below 1, uniformly in the size. The bound $\ell_B \le \sqrt{(r_1 r_2)_-}$ ties the loss to the negative part of the *tensor overlap* $r_1 r_2 = z_B(\sigma)\cdot z_B(\tau)$, where $z_B(\sigma) = 2^{-(k-1)}\sigma_{B^{(1)}} \otimes \sigma_{B^{(2)}}$ is a unit vector. Talagrand's positivity principle, in Panchenko's form, says that a suitable random perturbation g of *any* measure π on a unit sphere makes the scalar product of two independent draws typically nonnegative: for every $\varepsilon > 0$ there is v with $\mathbb{E}\pi_g^{\otimes 2}(z^1 \cdot z^2 \le -\varepsilon) \le \varepsilon$, uniformly in π and in the dimension, for $g(z) = v\sum_{p\ge 1} 2^{-p} x_p g_p(z)$ with iid uniform x_p and iid Gaussian p-spin interactions g_p [206, Thm. 1].

Theorem 12.13 (Conditional order in a perturbed HEA)**.** *Let $\frac{1}{2} < \varrho < 1$ and consider the HEA (12.5), or any base weight, with Panchenko's perturbation of strength $v_k = \beta 2^{k(1-\varrho)}$ added on the tensor vector z_B of every block B of level k. Let $\mathcal{J}(v)$ be the supremum, over dimensions and base measures, of $\mathbb{E}\pi_g^{\otimes 2}[((z^1\cdot z^2)_-)^{1/2}]$ at strength v, and $\mathcal{J}^*(v) = \sup_{v'\ge v}\mathcal{J}(v')$. Then for every n and $\beta > 0$,*

$$\mathbb{E}|q| \ge 1 - \sum_{k=1}^{n} \mathcal{J}(v_k) \ge 1 - \frac{1}{(1-\varrho)\ln 2}\int_\beta^\infty \mathcal{J}^*(v)\,\frac{dv}{v}.$$

Hence, if (SP-int) $\int_1^\infty \mathcal{J}^*(v)\,dv/v < \infty$*, there is $\beta_0(\varrho)$ such that $\inf_n \mathbb{E}q^2 \ge \frac{1}{4}$ for all $\beta \ge \beta_0$.*

Proof sketch. This conditional theorem is ours; we give the deduction. By Lemma 12.12, $\mathbb{E}|q| = 1 - \sum_k \mathbb{E}\bar{\ell}_k$ and $\mathbb{E}\ell_B \le \mathbb{E}\sqrt{(r_1 r_2)_-}$. Conditionally on everything except the perturbation of B, the law of z_B is some measure on the sphere tilted by its own Panchenko perturbation, so $\mathbb{E}\ell_B \le \mathcal{J}(v_k)$ uniformly. Since v_k grows geometrically and $\mathcal{J}^*$ is nonincreasing, $\mathcal{J}^*(v_k) \le ((1-\varrho)\ln 2)^{-1}\int_{v_{k-1}}^{v_k}\mathcal{J}^*(v)\,dv/v$ with $v_0 = \beta$. Under (SP-int) the tail integral is at most $\frac{1}{2}$ for large β, so $\mathbb{E}|q| \ge \frac{1}{2}$ and $\mathbb{E}q^2 \ge (\mathbb{E}|q|)^2$. □

The hypothesis (SP-int) is model-free: it asks for a rate in the positivity principle, uniform over base measures, for instance $\mathcal{J}(v) = O((\log v)^{-1-\delta})$. Only $\mathcal{J}(v) \to 0$ is known; tracking the constants of the published proof gives, heuristically, rates of order $(\log\log v/\log v)^{1/2}$ or slower, which are not integrable against dv/v. The theorem says nothing directly about the unperturbed HEA or about $\mathbb{Z}^d$, where the coupling between adjacent cubes has the surface link overlap as its covariance and does not see the signs of the volume overlaps.

Open problem 12.14 (A rate for the positivity principle)**.** Prove (SP-int), or any rate $\mathcal{J}(v)$ integrable against dv/v, for Panchenko's perturbation uniformly over measures on spheres of all dimensions (it suffices to treat measures supported on normalized tensors $x\otimes y/(|x||y|)$ of sign vectors). Alternatively, prove $\limsup_n \sum_k \mathbb{E}\bar{\ell}_k < 1$ at low temperature for the unperturbed HEA.

12.6 One-dimensional long-range chains

A second way to tune an effective dimension is to keep the lattice $\mathbb{Z}$ and let the couplings decay slowly. Let

$$J_{xy} = c_{|x-y|}Z_{xy}, \qquad c_r = r^{-\alpha/2}, \tag{12.6}$$

with Z_{xy} iid, mean 0, variance 1, so that $\operatorname{Var} J_{xy} = |x-y|^{-\alpha}$ (other papers write $|J_{xy}| \sim r^{-a}$ or $r^{-\sigma}$, $a = \sigma = \alpha/2$). For $\alpha > 1$ the couplings at a site are square summable, so for a configuration chosen independently of the couplings the field $\sum_y J_{xy}\sigma_y$ converges almost surely; for $\alpha \le 2$ they are not absolutely summable, so a boundary condition chosen with knowledge of the couplings can exert an infinite field on a spin.

For coupling-independent boundary conditions two statements are known. For $\alpha > 2$ and Z_{xy} with exponential moments, at every temperature the finite-volume Gibbs states converge almost surely to a pure state that does not depend on the boundary condition [40, Thms. 3.4, 3.5 and 3.7]. For every $\alpha > 1$ (in any dimension, with decay exponent $\alpha d/2$) and $|\mathbb{E}Z^p| \le \gamma^p p!$, at sufficiently high temperature the infinite-volume state is weakly unique and the averaged squared connected correlation decays like $|x-y|^{-\alpha}$ [97, Thm. 1 and (1.19)]. Physicists expect a transition for $1 < \alpha < 2$, of mean-field type for $\alpha < 4/3$ [143], as summarized in [203, p. 2]; no proof exists.

12.6.1 Exotic states

Gandolfi, Newman and Stein showed that with *coupling-dependent* boundary conditions the non-summability already produces uncountably many Gibbs states at every temperature. Their setting is $\mathbb{Z}^d$ with $J_{xy} = c_{y-x} Z_{xy}$, where the $Z_{xy} = Z_{yx}$ are iid with mean 0 and variance 1 and c is deterministic with $c_0 = 0$, and

$$\sum_y |c_y| = \infty, \qquad \sum_y c_y^2 < \infty; \tag{12.7}$$

for (12.6) this is exactly $1 < \alpha \le 2$. Because $\sum_y |J_{xy}| = \infty$, sums over y are taken along the fixed boxes $\Lambda_m = \{-m, \dots, m\}^d$. For a configuration σ and a finite set A, the energy cost of flipping A is

$$\Delta H_A(\sigma) = \lim_{m \to \infty} 2 \sum_{x \in A} \sum_{y \in \Lambda_m \setminus A} J_{xy} \sigma_x \sigma_y \qquad \text{(the limit may be } +\infty\text{)}.$$

A *ground state* has $\Delta H_A \ge 0$ for every finite A; it is *rigid* if $\Delta H_A = +\infty$ for every nonempty A, and *pliable* if $\Delta H_A < \infty$ for every A. For a pliable ground state, finite perturbations are allowed boundary conditions, and Gibbs states are defined by the DLR equations with the same order of summation.

Theorem 12.15 (Gandolfi–Newman–Stein [102, Thms. 2.1–2.3])**.** *Assume (12.7), and let $S \subset \mathbb{Z}^d$ be infinite with $\sum_{y \in S} |c_{y-x}| < \infty$ for every x. With probability one:*

(a) for every $\hat\sigma \in \{-1,1\}^S$ there is a rigid ground state equal to $\hat\sigma$ on S;

(b) for every $\hat\sigma \in \{-1,1\}^S$ and every family (θ_A) of nonnegative numbers indexed by finite sets, there is a pliable ground state $\tilde\sigma$ equal to $\hat\sigma$ on S with $\Delta H_A(\tilde\sigma) \ge \theta_A$ for every A;

(c) if $\tilde\sigma$ is a pliable ground state with $\sum_A e^{-\beta \Delta H_A(\tilde\sigma)} < \infty$, then $\mu(\{\tilde\sigma^A\}) \propto e^{-\beta \Delta H_A(\tilde\sigma)}$, where $\tilde\sigma^A$ is $\tilde\sigma$ flipped on A, is an infinite-volume Gibbs state at inverse temperature β, supported on finite perturbations of $\tilde\sigma$.

In particular, choosing θ_A that grow faster than the logarithm of an enumeration of the finite sets, at every $\beta > 0$ there are uncountably many mutually singular Gibbs states.

The construction behind (a) is a forest (Fig. 12.4): each site x receives an infinite set U_x of children with $\sum_{y \in U_x} |c_{y-x}| = \infty$, and the trees are rooted at the points of S. On each tree the spins are forced from the root down, the child agreeing with the sign of its coupling to the parent, so each site receives a divergent field from its tree neighbours. Every other term $J_{xy}\sigma_x\sigma_y$ multiplies a sign determined by tree couplings only, so, conditionally on the tree couplings, these terms are independent and centred with summable variances, and their sum converges. The proof, with the pliable variant (b) and the Gibbs states (c), is in Section A.28.

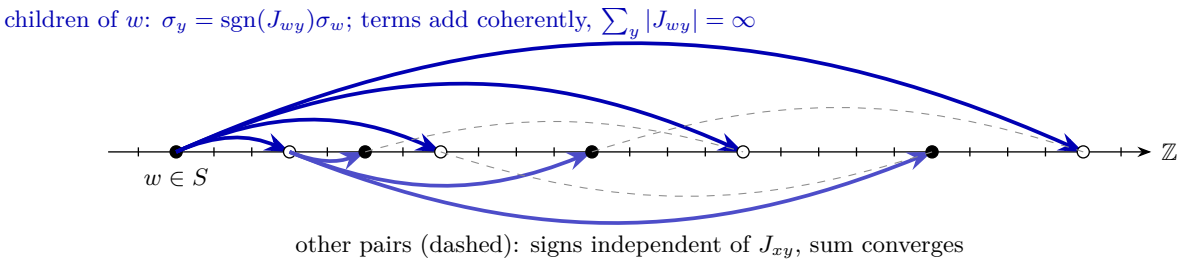


Figure 12.4: Rigid ground states of a long-range chain (Theorem 12.15). Along tree edges the spin follows the sign of the coupling, so the field from the tree neighbours diverges; the remaining couplings (dashed) multiply independent signs and give a convergent sum. Flipping any finite set costs infinite energy.

Route warning 12.16 (Gibbs non-uniqueness is not the order criterion here). *Tempting step:* take the number of Gibbs states as the order parameter of the chain (12.6). *Status: proved inadequate.* For $1 < \alpha \le 2$ it is uncountable at every temperature, including the high-temperature region where coupling-independent boundary conditions give a unique limit (Theorem 12.15, [97]). The meaningful statements are weak uniqueness for coupling-independent boundary conditions, overlap order, or metastate statements (Chapter 3). That weak uniqueness fails at low temperature for some $1 < \alpha \le 2$ is open [102, p. 4].

12.6.2 A transition on the Nishimori line

On the Nishimori line of a Gaussian chain, Okuyama and Ohzeki prove a phase transition. The couplings are independent, $J_{xy} \sim N(4\beta|x-y|^{-\alpha}, 4|x-y|^{-\alpha})$, with the mean tied to the variance by the factor β (the Gaussian Nishimori condition of Chapter 14, on which long-range order is simultaneously ferromagnetic and spin-glass order). The order parameter is the block-spin parameter $m^2 = \lim_p \lim_N 2^{-2p}\mathbb{E}\langle S_{p,r}^2\rangle$, $S_{p,r}$ the sum of the spins in a block of 2^p sites, with free boundary conditions [203, (15)–(22)].

Theorem 12.17 (Okuyama–Ohzeki [203, Thms. 1–3, Cor. 4])**.** *For the one-dimensional Gaussian chain on the Nishimori line with $1 < \alpha < \frac{3}{2}$: the long-range order parameter is positive at sufficiently low temperature, and it vanishes when $32\beta^2 \sum_{i\ge 1} i^{-\alpha} < 1$. Hence the chain has a phase transition at a finite temperature. The positive lower bound is proved first for the Dyson hierarchical version of the chain (couplings $N(\beta b_N 2^{-2N}, b_N 2^{-2N})$ on all pairs of a block of level N, $b_N = 2^{(2-\alpha)N}$), whose order parameter is at most that of the chain for every $\alpha > 1$ and β.*

Proof sketch. We omit the estimates; the architecture is Dyson's block-spin method for the ferromagnet [79] adapted to the Nishimori line. (1) The Nishimori-line Griffiths inequalities (Chapter 14) make $f_N(p) = 2^{-2p}\mathbb{E}\langle S_{p,r}^2\rangle_N$ nondecreasing in the size N, and compare the chain with its hierarchical version, whose couplings the chain's dominate in mean and variance pair by pair. (2) The Gibbs–Bogoliubov inequality on the line turns the loss of block order from one level to the next into a difference of quenched pressures [203, Lemma 5]. (3) A Guerra-type interpolation and Gaussian concentration bound that difference: $f_N(N) \ge f_{N-1}(N-1) - O(\beta^{-1}2^{-(3/2-\alpha)N}$ log-factors) [203, Lemmas 6–8, (28)]. The errors are summable exactly when $\alpha < \frac{3}{2}$; then $\liminf_N f_N(N)$ is at least $f_1(1) - O(1/\beta)$, positive for large β. The high-temperature

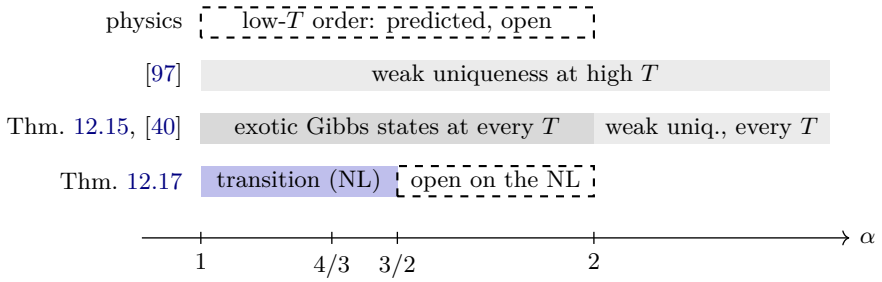


Figure 12.5: The chain with $\operatorname{Var} J_{xy} = |x-y|^{-\alpha}$; shaded bars are theorems, dashed boxes open statements, and NL marks statements on the Nishimori line. Off the Nishimori line nothing is proved at low temperature for $1 < \alpha \le 2$ (the exotic states say nothing about order); on the line the proved range ends at the Imry–Ma threshold $\alpha = \frac{3}{2}$.

part combines the same Griffiths inequality with Dyson's argument for the ferromagnet [203, p. 8]. □

Route warning 12.18 (Dyson's method stops at the Imry–Ma threshold). *Tempting step:* push Theorem 12.17 to $\frac{3}{2} \le \alpha < 2$, or off the Nishimori line. *Status: fails under the method's own assumptions.* The concentration step produces an error of order $2^{N/2}$ at level N against a coupling scale $2^{(2-\alpha)N}$, and for $\alpha \ge \frac{3}{2}$ these errors are not summable. Even optimal sub-Gaussian concentration gives an error of this order, the Imry–Ma count of fluctuations in a block; for the random-field Ising model on the same hierarchy the method proves order exactly for $1 < \alpha < \frac{3}{2}$, and there no order is expected for $\alpha > \frac{3}{2}$. So the method cannot distinguish spin-glass from random-field fluctuation costs. Off the line the Griffiths and Gibbs–Bogoliubov inputs are absent; for binary couplings the interpolation, the Gibbs–Bogoliubov inequality and the reproduction property of the Gaussian law all fail [203, p. 20].

Open problem 12.19 (Long-range chains)**.** (a) For the Gaussian chain on the Nishimori line, decide whether there is long-range order at low temperature for $\frac{3}{2} \le \alpha \le 2$ [203, p. 1]. (b) For the symmetric chain (12.6) with some $1 < \alpha \le 2$, prove that weak uniqueness fails at low temperature, or that it holds at every temperature.

12.7 Summary

On MK lattices ordering is one recursion (Lemma 12.2), decided for Gaussian couplings by the sign of $\Lambda(b,s) - 1$, with blue imbalance wherever there is order; order holds on diamond lattices for every bounded-density law once $s \ge 4$, and is excluded for $\pm J$ couplings on lattices of dimension arbitrarily close to three. None of this transfers directly (Proposition 12.9 and Route warning 12.11). On Dyson's lattice, order in the perturbed model follows from a rate in the positivity principle that is not proved; on long-range chains a transition is proved on the Nishimori line up to the Imry–Ma threshold.

Open problem 12.20 (MK lattices)**.** (a) Decide whether $\Lambda(b,s) = 1$ for some integers $b, s \ge 2$ [202, Rem. 5]. (b) Decide whether $\pm J$ couplings order at low temperature on the diamond lattice with $s \ge 3$ branches; more generally treat odd s and even s failing (12.3) [201, p. 5]. (c) Prove

Collet and Eckmann's expectation that order holds at every temperature below the critical one [59, Rem. 7.3].

The dilution mechanism returns on $\mathbb{Z}^d$ in Chapter 13, the Nishimori line in Chapter 14, and why scale-uniform control, which the hierarchies supply through exact recursions and unbounded degrees, is the wall on $\mathbb{Z}^d$ in Chapter 15.

Chapter 13

Diluted and decorated lattices

Is there *any* symmetric coupling law on $\mathbb{Z}^2$ whose nonzero couplings percolate and for which one can prove that the Gibbs state is unique at every temperature? Newman asked exactly this, as the "for some ν" form of his Research Problem 4.2: determine, "either for all ν or at least for some ν with $\nu(\mathbb{R} \setminus \{0\}) > p_c(2) = 1/2$", whether $\beta_c(2, \nu) = \infty$ [171, p. 48] (the full statement is in Chapter 1). Below the percolation threshold the answer is classical: Newman's comparison with the ferromagnet of couplings $|J_e|$ gives uniqueness at every β when the nonzero couplings do not percolate (Chapter 1). Above it the $|J|$-ferromagnet orders at low temperature, the comparison says nothing, and any proof of uniqueness must use the signs, uniformly down to zero temperature.

This chapter shows one mechanism that does it. For discrete couplings, a frustrated cycle of a particular shape can *cancel exactly*: summing out the spins inside it leaves a zero effective coupling at every temperature. If such cancellations remove enough bonds, what remains is subcritical percolation, and uniqueness follows at all temperatures at once, with no expansion. The same mechanism appears in three places: in the physics of diluted spin glasses (Jörg and Ricci-Tersenghi), on hierarchical lattices (Okuyama and Ohzeki, Chapter 12), and on $\mathbb{Z}^d$ itself for diluted $\pm J$ couplings slightly above the percolation threshold. On decorated lattices it gives a second lesson: two models on the same graph, one ordered and one not, can have the same law of signs on every forest, so a proof that uses only forest information about the disorder cannot decide order.

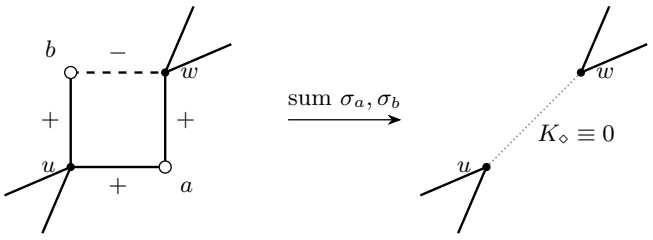
By the end of the chapter the reader can prove the cancellation identity and the reduction from finite effective clusters to uniqueness, the no-transition theorem on decorated lattices and the forest-information barrier, and can state the uniqueness theorem for diluted $\pm J$ couplings on $\mathbb{Z}^d$ with the structure of its proof.

13.1 Exact cancellation and finite effective clusters

The basic computation sums out a spin of degree two. Throughout, couplings multiply β as in Chapter 1.

Lemma 13.1 (Series and parallel decimation)**.**

(i) For a spin s joined only to u and w, by couplings x and y, $\sum_{s=\pm 1} e^{\beta s(x\sigma_u + y\sigma_w)} = C\, e^{k\sigma_u\sigma_w}$ with $C > 0$ and $\tanh k = \tanh(\beta x)\tanh(\beta y)$.



(a) path signs $\epsilon_1 = +1$ (via a) and $\epsilon_2 = -1$ (via b)

(b) $\sum_{\sigma_a,\sigma_b}(\cdots) = 4\cosh 2\beta$ for all σ_u, σ_w

Figure 13.1: The diamond cancellation (Lemma 13.1(iii)). Solid edges are $+1$ couplings, the dashed edge is -1; the middle vertices a, b have no other couplings. The path through a prefers $\sigma_u = \sigma_w$ and the path through b prefers $\sigma_u = -\sigma_w$, with exactly the same strength, so after summing out σ_a, σ_b the terminals do not interact at any temperature. The frustration of the square has become a missing bond.

(ii) Effective couplings between the same two spins add. Hence a diamond*, two paths $u\,a\,w$ and $u\,b\,w$ with nonzero couplings whose middle vertices a, b have no other couplings, has effective coupling*

$$K_\diamond = \operatorname{artanh}\big(\tanh(\beta J_{ua})\tanh(\beta J_{aw})\big) + \operatorname{artanh}\big(\tanh(\beta J_{ub})\tanh(\beta J_{bw})\big).$$

(iii) For unit magnitudes, $K_\diamond = \frac{1}{2}(\epsilon_1 + \epsilon_2)\log\cosh 2\beta$, where $\epsilon_1 = J_{ua}J_{aw}$ and $\epsilon_2 = J_{ub}J_{bw}$ are the path signs. So $K_\diamond = 0$ for every $\beta > 0$ if and only if the diamond is frustrated, $\epsilon_1\epsilon_2 = -1$.

(iv) In general $K_\diamond(\beta) = 0$ for every $\beta > 0$ if and only if the two path signs are opposite and the magnitude pairs $\{|J_{ua}|, |J_{aw}|\}$ and $\{|J_{ub}|, |J_{bw}|\}$ agree as multisets.

Proof. (i) $2\cosh(\beta(x\sigma_u + y\sigma_w))$ depends only on $\sigma_u\sigma_w$, so it equals $Ce^{k\sigma_u\sigma_w}$ with $e^{2k} = \cosh\beta(x+y)/\cosh\beta(x-y)$, and $\tanh k = \frac{\cosh\beta(x+y)-\cosh\beta(x-y)}{\cosh\beta(x+y)+\cosh\beta(x-y)} = \tanh(\beta x)\tanh(\beta y)$. (ii) The exponent is a sum of a term in s_a and a term in s_b, so the double sum factorizes and the two effective couplings add. (iii) If $|x| = |y| = 1$, $2\cosh(\beta(x\sigma_u + y\sigma_w))$ is $2\cosh 2\beta$ when $xy\sigma_u\sigma_w = 1$ and 2 otherwise, so $k = \frac{1}{2}xy\log\cosh 2\beta$. (iv) artanh is odd and injective, so $K_\diamond \equiv 0$ if and only if the path signs are opposite and $\beta \mapsto \tanh(\beta x)\tanh(\beta y)$ coincide for the two magnitude pairs; expanding, $\tanh(\beta x)\tanh(\beta y) = xy\beta^2 - \frac{1}{3}xy(x^2+y^2)\beta^4 + O(\beta^6)$, so the products and the sums of squares agree, hence the multisets. □

Figure 13.1 shows the cancellation. Two parallel two-bond paths of opposite sign are an old device: they already appear in the decorated lattices of Kasai, Miyazima and Syozi, where each pair of neighbouring spins is joined through a ferromagnetic and an antiferromagnetic decorated bond, with annealed occupations [133, §2]. What is used here is that, for quenched random signs, the frustrated diamond transmits *nothing* at every temperature, and which diamonds cancel does not depend on β.

The second ingredient turns finite effective clusters into uniqueness. It is stated for any locally finite graph, so that it applies to both settings of this chapter.

Lemma 13.2 (Cancelling blocks)**.** *Let $G = (V, E)$ be locally finite, with couplings J and $\beta \geq 0$. Suppose $V = B \sqcup \bigsqcup_k I_k$ with finite, pairwise disjoint* blocks *I_k, and for each k two vertices*

$u_k \neq w_k$ of B such that every edge with nonzero coupling that meets I_k has both endpoints in $I_k \cup \{u_k, w_k\}$. Summing over the spins of I_k the Boltzmann weight of the edges that meet I_k, with $\sigma_{u_k}, \sigma_{w_k}$ fixed, gives $C_k e^{K_k \sigma_{u_k} \sigma_{w_k}}$ for some $C_k > 0$ and $K_k \in \mathbb{R}$. Let G^{eff} be the multigraph on B whose edges are the edges of G inside B with nonzero coupling and the pairs $\{u_k, w_k\}$ with $K_k \neq 0$. If every component of G^{eff} is finite, then the Gibbs measure for J at β is unique. Under it the spins on different components of G^{eff} are independent, and given the spins on B the spins of different blocks are independent, each with a law depending only on $\sigma_{u_k}, \sigma_{w_k}$.

Proof. The summed weight is a positive function of $(\sigma_{u_k}, \sigma_{w_k})$, invariant under the global flip, hence of the stated form. Let μ be a Gibbs measure and ν its marginal on B.

The marginal is an effective Gibbs measure. Let $\Lambda \subset B$ be finite and let Λ^+ be Λ together with the blocks joined to Λ by an edge of nonzero coupling; local finiteness and disjointness make Λ^+ finite. A block with a terminal in Λ that is not in Λ^+ has nonzero couplings only inside it and to its other terminal if that lies outside Λ, so its summed weight is a flip-invariant function of at most one spin, hence constant, and $K_k = 0$; such blocks give no effective edge at Λ. An edge with nonzero coupling that meets Λ^+ either lies in B and meets Λ, or belongs to one of the blocks in Λ^+ (an edge from Λ into a block puts that block in Λ^+ and makes its endpoint a terminal). In the DLR equation for Λ^+ the sum over the block spins factorizes and produces $\prod_k C_k e^{K_k \sigma_{u_k} \sigma_{w_k}}$, so the conditional law of σ_Λ given the spins off Λ^+ is the finite-volume Gibbs law of the effective model on G^{eff} given $\sigma_{B \setminus \Lambda}$. This is measurable with respect to the spins on $B \setminus \Lambda$, which lie off Λ^+, so by the tower property ν satisfies the DLR equations of the effective model.

Finite components. If D is a finite component of G^{eff}, no effective edge leaves D, so the DLR equation for $\Lambda = D$ says that σ_D is independent of the spins outside, with the finite Gibbs law π_D of the effective model on D. By induction over finitely many components, ν is the product of the π_D on cylinder events, hence $\nu = \bigotimes_D \pi_D$.

Lifting. For finitely many blocks $I_1, \dots, I_n$, the only nonzero couplings touching $I_1 \cup \cdots \cup I_n$ are inside the blocks or to their terminals, so the DLR equation for this set gives a product of kernels, each depending on its two terminal spins. Hence μ is determined by ν on cylinder events, so μ is unique; a Gibbs measure exists by compactness, as in [95, Theorem 6.26]. □

Remark 13.3 (One null set for all temperatures). In both applications below, the set of blocks with $K_k \neq 0$ is the same for every $\beta > 0$, by Lemma 13.1(iii)–(iv). The hypothesis of Lemma 13.2 is then an event about J alone, and on it the Gibbs state is unique simultaneously for every $\beta \in [0, \infty)$. This is the strong quantifier order in Newman's definition of β_c [171, (3.58)]: one null set of couplings serves all temperatures.

13.2 Decorated lattices

The simplest setting puts the cancelling diamonds in by hand. Fix $d \geq 2$ and $r \geq 1$, and let $G_{d,r}$ be the graph obtained from $\mathbb{Z}^d$ by replacing every nearest-neighbour edge $\{x, x'\}$ by a chain of r diamonds: new vertices $x = z_0, z_1, \dots, z_r = x'$, and for each j two paths $z_{j-1} a_j z_j$ and $z_{j-1} b_j z_j$ (Fig. 13.2). The $3r - 1$ new vertices of an edge form its *block*; the vertices of $\mathbb{Z}^d$ are the *backbone*. The graph is periodic, has maximal degree $4d$ and is quasi-isometric to $\mathbb{Z}^d$. Every edge of $G_{d,r}$ carries a coupling of magnitude 1. Summing out a block, by Lemma 13.1, gives the backbone

coupling

$$\tanh K_e = \prod_{j=1}^{r} \tanh K_{\diamond,j}, \qquad K_{\diamond,j} \in \{0, \pm L(\beta)\}, \quad L(\beta) = \log\cosh 2\beta, \tag{13.1}$$

which vanishes as soon as one diamond of the chain is frustrated. With all couplings $+1$ it equals $K_r(\beta) = \operatorname{artanh}\big((\tanh L(\beta))^r\big)$, which tends to ∞ as $\beta \to \infty$.

Theorem 13.4 (No transition on a decorated lattice)**.** *Let $d \geq 2$ and $r \geq 1$ with $(2d-1)2^{-r} < 1$, and put $\lambda = (2d-1)2^{-r}$.*

(a) Put independent fair ± 1 couplings on all edges of $G_{d,r}$. There is an event of probability one on which the Gibbs measure is unique for every $\beta \in [0,\infty)$, and for backbone vertices x, y,

$$\sup_{\beta \geq 0} \mathbb{E}\, \langle \sigma_x \sigma_y \rangle_{\beta,J}^2 \;\leq\; \frac{2d}{(2d-1)(1-\lambda)}\, \lambda^{|x-y|_1}.$$

(b) Put all couplings equal to $+1$. For every β with $K_r(\beta) > \beta_c^{\mathrm{Ising}}(d)$, the critical inverse temperature of the Ising model on $\mathbb{Z}^d$, the Gibbs measure is not unique.

Proof. (a) The blocks satisfy the hypothesis of Lemma 13.2, with the backbone as B and the endpoints of the original edge as terminals; G^{eff} consists of the backbone edges whose chain contains no frustrated diamond. Each diamond is frustrated with probability $\frac{1}{2}$, independently, so the effective edges form Bernoulli bond percolation on $\mathbb{Z}^d$ with parameter 2^{-r}, the same for every $\beta > 0$. A cluster containing x and y contains a self-avoiding path from x of length $n \geq |x-y|_1$, and there are at most $2d(2d-1)^{n-1}$ such paths, so

$$\mathbb{P}(x \leftrightarrow y \text{ in } G^{\mathrm{eff}}) \;\leq\; \sum_{n \geq |x-y|_1} 2d(2d-1)^{n-1} 2^{-rn} = \frac{2d}{(2d-1)(1-\lambda)}\, \lambda^{|x-y|_1}.$$

In particular all effective clusters are finite almost surely, and Lemma 13.2 and Remark 13.3 give uniqueness for every β on one event. Under the unique measure the spins of different effective clusters are independent and, by flip invariance, have mean zero, so $\langle \sigma_x \sigma_y \rangle = 0$ unless $x \leftrightarrow y$ in G^{eff}; this gives the bound.

(b) Let Λ_n be a box of the backbone, and consider the decorated graph made of Λ_n and the blocks of all edges with an endpoint in Λ_n, with plus boundary condition on the backbone vertices outside Λ_n. Summing out the blocks turns it exactly into the Ising model on Λ_n with coupling $K_r(\beta)$ on every edge and plus boundary condition. So the magnetization of a backbone vertex under the decorated plus measure is that of the Ising model at inverse temperature $K_r(\beta)$, which is at least the spontaneous magnetization $m^*(K_r(\beta))$; this is positive when $K_r(\beta) > \beta_c^{\mathrm{Ising}}(d)$ [95, Theorem 3.25 and Definition 3.32], and $\beta_c^{\mathrm{Ising}}(d) < \infty$ for $d \geq 2$ by the same theorem. Subsequential limits of the decorated plus and minus measures are Gibbs measures with opposite, nonzero backbone magnetizations. □

By (a), with two independent replicas from the unique state, for boxes Λ of the backbone

$$\mathbb{E}\langle q_\Lambda^2 \rangle = |\Lambda|^{-2} \sum_{x,y \in \Lambda} \mathbb{E}\langle \sigma_x \sigma_y \rangle^2 \leq \frac{2d}{(2d-1)(1-\lambda)} \Big(\frac{1+\lambda}{1-\lambda}\Big)^d \frac{1}{|\Lambda|},$$

so the stationary target of Chapter 1 fails for every metastate, all of which are concentrated on the unique state. The thermodynamic Edwards–Anderson response also vanishes: the dependency colouring of the finite clusters used in [217] adapts, and we omit it. The mechanism is classical decoration–iteration; the theorem is stated for its lessons.



(a) one edge of $\mathbb{Z}^d$ replaced by $r = 3$ diamonds; the middle diamond is frustrated, so $K_e = 0$

(b) a spanning forest of the block (blue): on it, iid signs and gauged plus signs have the same law

Figure 13.2: A decorated edge of $G_{d,r}$. (a) Summing out the block multiplies the diamond tanh's, so one frustrated diamond (shaded) disconnects x from x'; for iid fair signs a block survives with probability 2^{-r}, and when $(2d-1)2^{-r} < 1$ the surviving backbone edges are subcritical percolation. (b) On any forest, such as the blue one, iid fair signs and the gauge image $\eta_x\eta_y$ of the all-plus model are both iid fair. The two models differ only in the products of signs around the dotted closing edges, that is, around cycles: they are frustrated with probability $\frac{1}{2}$ in one model and never in the other.

13.2.1 Forest information cannot decide order

Part (b) of Theorem 13.4 describes a ferromagnet. Its gauge transform looks like a spin glass.

Lemma 13.5 (Gauged laws are fair on forests)**.** *Let $a \in \{\pm1\}^E$ be any sign pattern on a graph and let η be independent fair signs on the vertices. Then for every finite forest $F \subseteq E$ the signs $(\eta_x a_{xy} \eta_y)_{xy\in F}$ are independent and fair.*

Proof. The map $\eta \mapsto (\eta_x\eta_y)_{xy\in F}$ from $\{\pm1\}^{V(F)}$ to $\{\pm1\}^F$ is a group homomorphism, and it is surjective: root each tree of F and choose η along it. The image of the uniform measure under a surjective homomorphism of finite groups is uniform; multiplying by the fixed a_F preserves uniformity. □

Theorem 13.6 (Forest information does not decide order)**.** *On $G_{d,r}$ with $(2d-1)2^{-r} < 1$, let J be independent fair ±1 couplings and let $\tilde{J}_{xy} = \eta_x\eta_y$ with independent fair vertex signs η. For every finite forest F, the families $(J_e)_{e\in F}$ and $(\tilde{J}_e)_{e\in F}$ have the same law. Under J, almost surely the Gibbs state is unique at every β. Under $\tilde{J}$, for every β with $K_r(\beta) > \beta_c^{\mathrm{Ising}}(d)$, almost surely there are at least two Gibbs states, and in the gauge image of the plus state $\langle\sigma_x\sigma_y\rangle^2 \geq m^*(K_r(\beta))^4 > 0$ for all backbone vertices x, y.*

Proof. Equality on forests is Lemma 13.5 with $a \equiv +1$. Uniqueness under J is Theorem 13.4(a). The map $\sigma \mapsto \eta\sigma$ carries the Gibbs measures of the all-plus model bijectively to those of $\tilde{J}$ (Chapter 1), so non-uniqueness follows from Theorem 13.4(b). In the plus state of the all-plus model the backbone marginal is the plus state of the Ising model at $K_r(\beta)$, so for backbone vertices $\langle\sigma_x\sigma_y\rangle^+ \geq \langle\sigma_x\rangle^+\langle\sigma_y\rangle^+ = m^*(K_r(\beta))^2$ by the second Griffiths inequality; the gauge multiplies $\langle\sigma_x\sigma_y\rangle$ by $\eta_x\eta_y$. □

The theorem extends to adaptive queries: an algorithm that queries coupling signs one at a time, choosing each query from the answers so far, sees exactly the same law of transcripts under J and $\tilde{J}$ as long as the union of its queries stays a forest (we omit the proof, an induction on the transcript using Lemma 13.5). So any argument that uses the disorder only through its law on forests, adaptively or not, would give the same answer for an ordered and a disordered model on the same periodic, bounded-degree graph with unit magnitudes. Only the law of the signs around cycles distinguishes them: under J each diamond is frustrated with probability $\frac{1}{2}$, under $\tilde{J}$ none

is. The same phenomenon occurs on $\mathbb{Z}^d$ itself for the diluted laws of the next section, where the Mattis couplings $\eta_x\eta_y|J_{xy}|$, with η iid fair vertex signs independent of J, have the same forest statistics as the iid couplings and order at low temperature [217, Remark 8.1]; and on the square lattice with unit magnitudes there are two centred, full-support, finite-range sign laws, neither iid, that agree on every forest, one with zero and one with positive Edwards–Anderson response at a fixed large β (Remark 15.4). The comparison conjectures of Chapter 11 respect this lesson: they are statements about the exact iid law, and the forest model enters only as a bound.

13.3 Diluted $\pm J$ couplings on $\mathbb{Z}^d$

On the cubic lattice nothing is decorated by hand; the cancelling diamonds have to be found in the disorder. Take the diluted law

$$\mu_p = (1-p)\,\delta_0 + \tfrac{p}{2}(\delta_{+1} + \delta_{-1})$$

of Chapter 1. The nonzero couplings form Bernoulli bond percolation $\omega(J)$ with parameter p; let $p_c = p_c(d)$ be its threshold, with $p_c(2) = \frac{1}{2}$ [137]. For $p \le p_c$ the Gibbs state is unique at every β by Newman's comparison (Chapter 1). For $p > p_c$ the ferromagnet with couplings $|J_e|$ has distinct plus and minus states for every β with $p\tanh\beta > p_c$ [171, Theorem 3.10], [9], and frustrated plaquettes occupy a positive density of the infinite cluster [217, Theorem 1.1(4)]. The comparison with the $|J|$-ferromagnet then fails at low temperature.

Definition 13.7 (Isolated diamonds)**.** A plaquette of $\mathbb{Z}^d$ is an *isolated diamond* of ω if its four edges are open and exactly one of its two diagonal pairs consists of vertices of ω-degree two; that pair is *internal*, the other *terminal*. The *diminished configuration* $\hat\omega(J)$ is $\omega(J)$ with the four edges of every frustrated isolated diamond removed.

An internal vertex has exactly two open edges, both in its diamond, so it belongs to no other isolated diamond; distinct isolated diamonds share no edge, and distinct internal vertices are never joined by an open edge [217, Lemma 2.3]. So the internal pairs of the frustrated isolated diamonds are blocks in the sense of Lemma 13.2, each with $K_k = 0$ by Lemma 13.1(iii), and the effective graph is $\hat\omega(J)$ minus the internal vertices (Fig. 13.3). Uniqueness at every temperature therefore follows from finiteness of the clusters of $\hat\omega(J)$, an event that does not depend on β. Its probability is a percolation question.

Lemma 13.8 (The diminished model [217, Lemma 3.4])**.** *Let ω be Bernoulli(p) bond percolation and, independently, let each plaquette carry a Bernoulli(s) mark. Delete the four edges of every isolated diamond of ω that is marked. For J with law $\mu_p^{\otimes E_d}$, the pair $(\omega(J), \hat\omega(J))$ has the law of the resulting pair at $s = \frac{1}{2}$.*

Proof. Given $\omega(J)$, the signs of the open edges are independent and fair. The frustration indicators of the isolated diamonds are products of signs over pairwise disjoint sets of four edges, hence independent fair coins given ω. □

The deletion rule is local, but it is not monotone in ω: opening an edge can create an isolated diamond or destroy one. It can only close edges, so the diminished model has no infinite cluster for $p < p_c$; the point is a strict inequality.

Proposition 13.9 (Subcriticality slightly above p_c [217, Proposition 3.5])**.** *For every $d \geq 2$ there is an explicit $g_d > 0$ such that for every $s \in (0, \frac{1}{2}]$ and every $p < p_c(d) + g_d s$ the diminished model has no infinite cluster, and the probability that the cluster of the origin reaches distance L decays exponentially in L.*

Proof sketch. We omit the surgery construction and the counting, which occupy most of the source. The strategy is that of Aizenman and Grimmett for essential enhancements [5], in the form used by Grimmett and Stacey for a non-monotone diminishment of bond percolation [110]. Write $\theta_L(p, s)$ for the probability that the origin is connected to distance L in the diminished model. A Russo-type formula, valid because the deletion rule has finite range, bounds $\partial_p \theta_L$ by a sum of pivotal probabilities over edges and expresses $-\partial_s \theta_L$ as a sum of pivotal probabilities over plaquettes. A local surgery, built for this rule, converts an edge that is pivotal for the crossing into a nearby plaquette whose deletion is pivotal: inside a box around the edge the crossing is rerouted through a planted isolated diamond that becomes its only connection. Counting the surgery's cost gives $-\partial_s \theta_L \geq g_d \, \partial_p \theta_L$, and integrating along the curve $s \mapsto (p_0 + g_d s, s)$ from a subcritical $p_0 < p_c$ gives $\theta_L(p_0 + g_d s, s) \leq \theta_L(p_0, 0)$. The surgery replaces the general combinatorial step of [5], in the spirit of the treatment of essential enhancements of bond percolation on $\mathbb{Z}^d$ by Balister, Bollobás and Riordan [20]. □

The hand-made constant is tiny ($g_2 = 2^{-988}/1130$); sharper surgeries, and in $d = 2$ an exact computer-verified comparison of local pivotal probabilities, enlarge it. Putting the pieces together:

Theorem 13.10 (Diluted $\pm J$ uniqueness above p_c [217, Theorem 1.1(1)–(2), Corollary 1.2, Theorem 1.4])**.** *Let $d \geq 2$, let J have law $\mu_p^{\otimes E_d}$, and let $p < p_c(d) + g_d/2$ with g_d as in Proposition 13.9.*

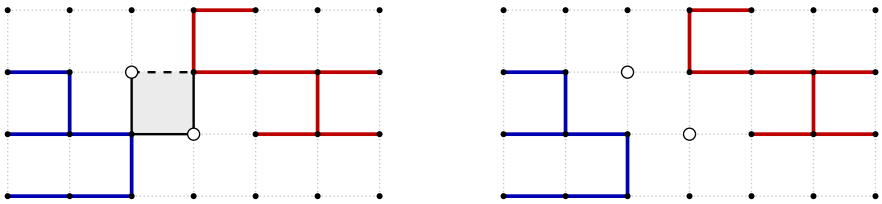
(a) There is an event of probability one on which the Gibbs measure is unique for every $\beta \in [0, \infty)$.

(b) There are $C, c > 0$, depending only on d and p, with $\sup_{\beta \geq 0} \mathbb{E}\langle \sigma_0 \sigma_x \rangle_{\beta,J}^2 \leq C e^{-c|x|}$.

In $d = 2$ the conclusions hold for every $p < \frac{1}{2} + g^*$, with an explicit $g^* > 1.1692 \cdot 10^{-5}$. Hence $\beta_c(2, \mu_p) = \bar{\beta}_c(2, \mu_p) = \infty$ for $p \in (\frac{1}{2}, \frac{1}{2} + g^*)$, where $\mu_p(\mathbb{R} \setminus \{0\}) = p > \frac{1}{2}$.

Proof sketch. We omit the correlation bound (b), which follows from the explicit structure of the unique state given by Lemma 13.2, and the explicit windows. For (a): by Lemma 13.8 and Proposition 13.9 with $s = \frac{1}{2}$, all clusters of $\hat{\omega}(J)$ are finite almost surely; by Lemma 13.2 and Remark 13.3, the Gibbs state is unique for every β on that event. The value g^* in $d = 2$ is computer-assisted: a sparser deletion rule dominates the diminished model, and by planarity the comparisons needed along a two-path curve reduce to finitely many exact inequalities between counts of configurations in a box, one for each non-crossing partition of its boundary, checked exhaustively in exact arithmetic by two programs that share no code. □

So the "for some ν" form of Newman's Research Problem 4.2 has an affirmative answer, through the diluted laws μ_p with p slightly above $\frac{1}{2}$; the "for all ν" form, and in particular the $\pm J$ and Gaussian laws, is open. In the same window the source also proves that the thermodynamic Edwards–Anderson order parameter vanishes and, for $p > p_c$, that the $|J|$-ferromagnet orders: a percolating $|J|$-graph, an ordered $|J|$-ferromagnet and a positive density of frustration do not imply spin-glass order at any temperature on $\mathbb{Z}^d$ [217, Theorem 1.1(3)–(4)]. The argument needs



(a) $\omega(J)$: a frustrated isolated diamond (shaded) is the only link between two clusters

(b) $\hat{\omega}(J)$: after exact decimation the two effective clusters are finite and independent

Figure 13.3: Decimation becomes dilution. (a) Open edges of $\omega(J)$ (thick; dotted edges have $J_e = 0$); the shaded plaquette is an isolated diamond whose internal vertices (white) have degree two, and its path signs are opposite. (b) Summing out the two internal spins removes the diamond at every temperature, and the blue and red clusters no longer interact. The diminished configuration is Bernoulli percolation with each isolated diamond deleted with probability $\frac{1}{2}$; slightly above p_c this extra deletion is enough to make all effective clusters finite.

two atoms of the coupling law: an atom at 0, to sit near the percolation threshold, and an atom of $|J_e|$, to make exact cancellation possible (Lemma 13.1(iv)). For diluted symmetric laws whose magnitude law has an atom, the same conclusions hold with $\frac{1}{2}$ replaced by the cancellation probability of an isolated diamond [217, Theorem 9.1].

13.4 One mechanism in four settings

The mechanism, exact finite decimation producing zero effective bonds and then subcritical percolation, was proposed in physics, then made rigorous on hierarchical lattices and on $\mathbb{Z}^d$ (Table 13.1).

Jörg and Ricci-Tersenghi studied diluted $\pm J$ models at very low temperature. On the hierarchical lattice they argued that "if the original model had a bond density slightly larger than the percolating density p_c ... these exact cancellations would produce a renormalized lattice which is below the percolating point, preventing any long-range correlation from arising", so that $p_c < p_{\mathrm{SG}}$ "for any frustrated model with discrete couplings on the hierarchical lattice" [129, p. 3]. Their Monte Carlo study of the three-dimensional diluted model indicates a percolating phase without long-range order, again $p_c < p_{\mathrm{SG}}$ [129, p. 3].

Okuyama and Ohzeki made the hierarchical version rigorous [201]. On a Migdal–Kadanoff lattice in which each edge is replaced by s parallel chains of b bonds, with s even and undiluted $\pm J$ couplings, one exact renormalization step gives each renormalized bond a zero coupling with probability $\binom{s}{s/2}2^{-s}$, the probability that the s equal-magnitude branch contributions cancel; if this vacancy probability exceeds the hierarchical percolation threshold, $\mathbb{E}\langle\sigma_A\sigma_B\rangle^2$ between the poles decays exponentially at every temperature, including zero [201, Theorems 1–2]. The criterion and its proof are in Chapter 12. On the hierarchical lattice the induced dilution is independent over the renormalized bonds, and a percolation recursion suffices. On $\mathbb{Z}^d$ the cancelling diamonds form a dependent, non-monotone diminishment, and a strict inequality is needed. On decorated lattices the cancellation is built into the geometry, and path counting suffices.

Setting	Exact step and zero bonds	Percolation input	Status
Hierarchical lattice, diluted $\pm J$ [129]	first decimation steps; exact cancellations	renormalized lattice below threshold	physics argument; numerics in $d = 3$
MK lattice, $\pm J$, s even [201]	one RG step; the s branches cancel with probability $\binom{s}{s/2}2^{-s}$	hierarchical percolation recursion	theorem (Chapter 12)
Decorated $\mathbb{Z}^d$, iid ± 1 (Theorem 13.4)	sum out each block; zero if one of r diamonds is frustrated	Bernoulli(2^{-r}), path counting	theorem
$\mathbb{Z}^d$, diluted $\pm J$ (Theorem 13.10)	sum out internal pairs; frustrated isolated diamonds, probability $\frac{1}{2}$	strict inequality for a diminishment	theorem, $p < p_c + g_d/2$

Table 13.1: One mechanism in four settings: an exact finite decimation produces zero effective bonds, and the remaining bond process is subcritical. The settings differ in how independent the zero bonds are, and hence in how much percolation theory the last step needs.

13.5 Route warnings

Route warning 13.11 (Exact cancellation needs atoms). *Attractive:* run the same argument for undiluted $\pm J$ or Gaussian couplings on $\mathbb{Z}^2$. *Needed:* a finite exact decimation that produces zero effective bonds with probability above a percolation threshold. *Fails for continuous laws:* by Lemma 13.1(iv) exact cancellation needs equal magnitudes, an event of probability zero when the magnitude law has no atom; Okuyama and Ohzeki note the same limitation on hierarchical lattices [201, p. 4]. *No known decimation for undiluted $\pm J$ on $\mathbb{Z}^d$:* every vertex has degree $2d$, so there are no isolated diamonds, and eliminating spins on $\mathbb{Z}^d$ produces multi-spin interactions with positive probability rather than an effective bond (Chapter 12); whether some other finite decimation creates enough exact zeros is not known.

Route warning 13.12 (Coarse hypotheses do not imply order). *Attractive:* prove spin-glass order at low temperature from coarse properties of the model: dimension, bounded degree, periodicity, iid symmetric signs, unit magnitudes, a percolating $|J|$-graph with an ordered $|J|$-ferromagnet, a positive density of frustration. *Proved false:* the decorated lattices of Theorem 13.4 have all the geometric properties and no transition, and the diluted laws of Theorem 13.10 have all the disorder properties on $\mathbb{Z}^d$ itself and no transition. An ordering proof must use a quantitative input, such as non-atomic magnitudes or strong supercriticality, and must use the law of the signs on cycles (Theorem 13.6).

13.6 What the testbeds show

In this chapter exact cancellation gave uniqueness at every temperature for symmetric short-range models whose nonzero couplings percolate, on $\mathbb{Z}^d$ and on graphs quasi-isometric to it. The proofs are elementary once the cancellation is found, except for the strict percolation inequality, and they show plainly what an ordering proof must avoid: the decorated lattices and the diluted laws satisfy every coarse hypothesis one might hope to use, and the decorated pair shows that the law of the signs on forests does not determine the answer, so a proof must use the signs around cycles. They say nothing about the undiluted $\pm J$ or Gaussian models, for which no exact zero is available.

Open problem 13.13 (Diluted couplings in two dimensions)**.** For the diluted law μ_p on $\mathbb{Z}^2$, prove that the Gibbs state is almost surely unique at every $\beta < \infty$ for every $p < 1$; or, more

modestly, for all p in an interval $(\frac{1}{2}, \frac{1}{2} + \delta)$ with δ of order one.

The comparison route of Chapter 11 is designed to handle the undiluted case, where cancellations are not exact but still decide the answer. The barriers met here, forest information and coarse hypotheses, reappear with the others in Chapter 15; the Nishimori line, where gauge symmetry gives exact identities of a different kind, is Chapter 14.

Chapter 14

Gauge symmetry and the Nishimori line

For symmetric couplings the gauge symmetry of Chapter 1 is mostly an obstacle: it makes every averaged correlation vanish and forces order to be detected through overlaps. For biased couplings it becomes a tool. On one curve of the plane of bias and temperature, the *Nishimori line*, the law of the couplings is itself a Boltzmann weight, the quenched normalization cancels against it, and quantities out of reach elsewhere become exact: the internal energy, the averaged random-cluster law (Chapter 7), and the identity between the laws of magnetization and overlap. It is the one place in the phase diagram of a short-range Ising glass where both order and its absence are proved in low dimension, and it carries the only proved transition of a one-dimensional long-range glass (Chapter 12).

The chapter derives the gauge identities from one lemma, with Nishimori's inequality and the shape of the ferromagnetic phase, the magnetization–overlap identity and the Griffiths inequalities on the line; proves, near full bias on $\mathbb{Z}^2$, ferromagnetic order at the Nishimori point and overlap order with Gibbs non-uniqueness up to twice the Nishimori temperature, by a Peierls argument with random contours, and states the best certified bound for overlap order; and turns to what is open or heuristic: reentrance below the multicritical point, the duality conjecture, the inference reading and planar free fermions. One theme runs through it: on the Nishimori line spin-glass and ferromagnetic order coincide, so the line bounds the zero-bias glass but never reaches it.

14.1 Biased couplings and the gauge average

Recall the gauge transformation of Chapter 1: for $\tau \in \{-1,1\}^V$, $J^\tau_{xy} = J_{xy}\tau_x\tau_y$ and $(\sigma\tau)_x = \sigma_x\tau_x$, so that $H_{J^\tau}(\sigma\tau) = H_J(\sigma)$. A symmetric coupling law is invariant under $J \mapsto J^\tau$. The laws below are not, but their failure to be invariant has a special form.

Definition 14.1 (Nishimori-type laws and the Nishimori line)**.** A law ν on $\mathbb{R}$ is of *Nishimori type* with parameter $\beta_{\mathrm{N}} \geq 0$ if $\nu(dJ) = c^{-1}e^{\beta_{\mathrm{N}}J}\nu_s(dJ)$ for a symmetric measure ν_s with $c = \int e^{\beta_{\mathrm{N}}J}\nu_s(dJ) < \infty$. The main examples are the biased $\pm J$ law $\nu_p = p\delta_{+1} + (1-p)\delta_{-1}$, $p \in [\frac{1}{2}, 1)$, with

$$\nu_p(J) = \frac{e^{\beta_{\mathrm{N}}J}}{2\cosh\beta_{\mathrm{N}}}, \qquad e^{-2\beta_{\mathrm{N}}} = \frac{1-p}{p}, \qquad \tanh\beta_{\mathrm{N}} = 2p-1,$$

and the Gaussian law $N(J_0, \Delta^2)$ with $\beta_{\mathrm{N}} = J_0/\Delta^2$. The *Nishimori line* is $\beta = \beta_{\mathrm{N}}$; for $\pm J$ it is the curve $T = 1/\beta_{\mathrm{N}}(p) = 2/\log\frac{p}{1-p}$ in the (p, T) plane.

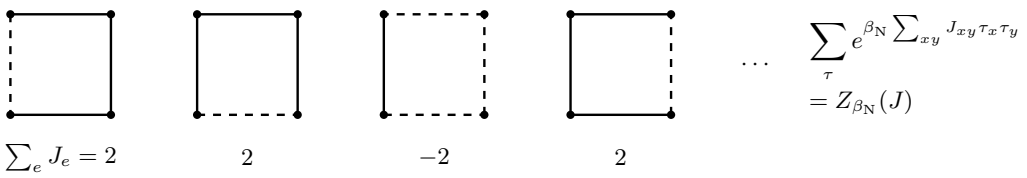


Figure 14.1: Part of the gauge orbit of a frustrated plaquette (solid +1, dashed −1). Gauge-invariant quantities are constant on the orbit, but the biased law weights each member by $e^{\beta_{\mathrm{N}} \sum_e J_e}$; summed over the orbit these are Boltzmann weights of the gauge spins τ, so the orbit's probability is proportional to $Z_{\beta_{\mathrm{N}}}(J)$, which on the Nishimori line cancels the normalization of the Gibbs average.

We work on a finite graph $G = (V, E)$ with iid couplings, with either the free boundary condition or the plus boundary condition on a set $\partial \subset V$ of boundary vertices whose spins are fixed to +1. The gauge group $\mathcal{T}$ is $\{-1, 1\}^V$ in the first case and $\{\tau : \tau = 1 \text{ on } \partial\}$ in the second; in both cases μ_{β,J^τ} is the image of $\mu_{\beta,J}$ under $\sigma \mapsto \sigma\tau$ for $\tau \in \mathcal{T}$. A function $F : \mathbb{R}^E \times \{-1, 1\}^V \to \mathbb{R}$ is *gauge covariant* if $F(J^\tau, \sigma\tau) = F(J, \sigma)$ for all $\tau \in \mathcal{T}$. Write $+$ for the all-plus configuration.

Lemma 14.2 (Nishimori identity)**.** *Let the couplings be iid of Nishimori type with parameter β_{N}, and let F be gauge covariant with $\mathbb{E} \max_\sigma |F(J, \sigma)| < \infty$. Then*

$$\mathbb{E}\langle F(J, \sigma)\rangle_{\beta_{\mathrm{N}},J} = \mathbb{E}F(J, +).$$

Both sides equal $|\mathcal{T}|^{-1} c^{-|E|} \int \nu_s^{\otimes E}(dJ)\, Z_{\beta_{\mathrm{N}}}(J)\, \langle F(J, \cdot)\rangle_{\beta_{\mathrm{N}},J}$.

Proof. Write $\mathbb{E}F(J, +) = c^{-|E|} \int \nu_s^{\otimes E}(dJ)\, e^{\beta_{\mathrm{N}} \sum_e J_e} F(J, +)$. For $\tau \in \mathcal{T}$ the map $J \mapsto J^\tau$ changes the sign of the couplings on the edges with $\tau_x\tau_y = -1$ and preserves $\nu_s^{\otimes E}$. Substitute it, and use $\sum_e J_e^\tau = \sum_{xy} J_{xy}\tau_x\tau_y$ and $F(J^\tau, +) = F(J, \tau)$ (covariance with $\sigma = \tau$):

$$\mathbb{E}F(J, +) = c^{-|E|} \int \nu_s^{\otimes E}(dJ)\, e^{\beta_{\mathrm{N}} \sum_{xy} J_{xy}\tau_x\tau_y} F(J, \tau).$$

Averaging over $\tau \in \mathcal{T}$ turns the integrand into $|\mathcal{T}|^{-1} Z_{\beta_{\mathrm{N}}}(J) \langle F(J, \cdot)\rangle_{\beta_{\mathrm{N}},J}$. The function $G(J) = \langle F(J, \cdot)\rangle_{\beta_{\mathrm{N}},J}$ is gauge invariant (substitute $\sigma = \sigma'\tau$ in the Gibbs sum), so the same computation with G in place of F gives the same expression for $\mathbb{E}G(J)$. □

Averaged over a gauge orbit, the disorder weight $e^{\beta_{\mathrm{N}} \sum_e J_e}$ becomes the partition function at β_{N}, the normalization of the Gibbs average (Fig. 14.1); seen through any covariant observable, a Gibbs sample at β_{N} is indistinguishable from the planted configuration $+$ (Section 14.7). Nishimori's classical identities follow by choosing F.

Corollary 14.3 (Nishimori [193, (4.15), (4.45), (4.50)])**.** *In the setting of Lemma 14.2:*

(a) if $\mathbb{E}|J_e| < \infty$, the internal energy on the line is exact: $\mathbb{E}\langle H_J\rangle_{\beta_{\mathrm{N}},J} = -\sum_e \mathbb{E}J_e$, which for $\pm J$ is $-|E| \tanh \beta_{\mathrm{N}}$;

(b) for every $A \subset V$ and every $\beta \geq 0$,

$$\mathbb{E}\langle \sigma_A\rangle_{\beta,J} = \mathbb{E}\big[\langle \sigma_A\rangle_{\beta_{\mathrm{N}},J} \langle \sigma_A\rangle_{\beta,J}\big];$$

in particular $\mathbb{E}\langle \sigma_A\rangle_{\beta_{\mathrm{N}},J} = \mathbb{E}\langle \sigma_A\rangle^2_{\beta_{\mathrm{N}},J} \geq 0$;

(c) (Nishimori's inequality) for every $A \subset V$ and every $\beta \geq 0$, $|\mathbb{E}\langle\sigma_A\rangle_{\beta,J}| \leq \mathbb{E}|\langle\sigma_A\rangle_{\beta_{\mathrm{N}},J}|$.

Proof. (a) Take $F(J,\sigma) = H_J(\sigma)$, which is covariant with $\max_\sigma |H_J(\sigma)| \leq \sum_e |J_e|$, integrable by assumption, and note $H_J(+) = -\sum_e J_e$ and $\mathbb{E}J_e = \tanh\beta_{\mathrm{N}}$ for $\pm J$. (b) Take $F(J,\sigma) = \sigma_A\langle\sigma_A\rangle_{\beta,J}$; it is covariant because $\langle\sigma_A\rangle_{\beta,J^\tau} = \tau_A\langle\sigma_A\rangle_{\beta,J}$. (c) Take absolute values in (b) and use $|\langle\sigma_A\rangle_{\beta,J}| \leq 1$. □

Part (b) at $\beta = \beta_{\mathrm{N}}$ equates a ferromagnetic quantity, the averaged correlation, with a spin-glass quantity, the averaged *squared* correlation. Part (c), uniform in the volume, combined with the averaged random-cluster law of Chapter 7, gives an explicit region without ferromagnetic order.

Proposition 14.4 (No ferromagnetic order to the left of the line's transition)**.** *Let the couplings on $\mathbb{Z}^d$ be iid ν_p, and consider boxes Λ_L with free boundary conditions.*

(a) For every $\beta \geq 0$, L and x, $|\mathbb{E}\langle\sigma_0\sigma_x\rangle_{\Lambda_L,\beta}| \leq \mathbb{E}|\langle\sigma_0\sigma_x\rangle_{\Lambda_L,\beta_{\mathrm{N}}(p)}|$. Hence, if the right side tends to 0 as $|x| \to \infty$ uniformly in L, then $\sup_{\beta\geq 0}\sup_L |\mathbb{E}\langle\sigma_0\sigma_x\rangle_{\Lambda_L,\beta}| \to 0$.

(b) The hypothesis of (a) holds whenever $2p-1 < p_c^{\mathrm{bond}}(\mathbb{Z}^d)$; on $\mathbb{Z}^2$, whenever $p < \frac{3}{4}$.

Proof. (a) is Corollary 14.3(c) with $A = \{0, x\}$. (b) is the percolation bound on the Nishimori line of Chapter 7, from Gandolfi's averaged FK law: $\mathbb{E}|\langle\sigma_0\sigma_x\rangle_{\Lambda_L,\beta_{\mathrm{N}}}| \leq \mathbb{P}_{2p-1}(0 \leftrightarrow x \text{ in } \Lambda_L)$, which tends to 0 uniformly in L below p_c^{bond}; on $\mathbb{Z}^2$, $p_c^{\mathrm{bond}} = \frac{1}{2}$ [99, p. 430]. □

Part (a) is Nishimori's inequality [193], and a weaker form of (b), with threshold $1 - 1/\sqrt{2}$ for the antiferromagnetic fraction, is due to Mazza [158]. In words: at fixed bias, ferromagnetic order at any temperature forces order at the Nishimori temperature. The ferromagnetic phase therefore lies to the right of the point M where the Nishimori line leaves it; below M its boundary is vertical or bends to the right (reentrance), never to the left [193, pp. 57–58] (Fig. 14.2).

14.2 Magnetization and overlap

The correlation identity has a distributional form. Let $m = |V|^{-1}\sum_x \sigma_x$, and for two replicas σ, τ drawn independently from $\mu_{\beta_1,J}$ and $\mu_{\beta_2,J}$ (same couplings, possibly different temperatures) let $q = |V|^{-1}\sum_x \sigma_x\tau_x$.

Theorem 14.5 (Nishimori–Sherrington [196, (11)])**.** *On a finite graph with free boundary conditions and iid couplings of Nishimori type, for every $\beta \geq 0$ the law of m under $\mathbb{E}\langle\cdot\rangle_{\beta,J}$ equals the law of q under $\mathbb{E}\langle\cdot\rangle_{\beta,J} \otimes \langle\cdot\rangle_{\beta_{\mathrm{N}},J}$. In particular, on the line ($\beta = \beta_{\mathrm{N}}$) the averaged magnetization law and the averaged overlap law coincide.*

Nishimori and Sherrington state the identity for $\pm J$ couplings; the proof below applies to every Nishimori-type law.

Proof. For a bounded function ϕ let $F(J,\sigma) = \sum_{\sigma'} \mu_{\beta,J}(\sigma')\,\phi(|V|^{-1}\sum_x \sigma_x\sigma'_x)$. Substituting $\sigma' \mapsto \sigma'\tau$ in the sum shows that F is covariant. By Lemma 14.2, $\mathbb{E}\langle F\rangle_{\beta_{\mathrm{N}},J}$, which is $\mathbb{E}\phi(q)$ for replicas at β_{N} and β, equals $\mathbb{E}F(J,+) = \mathbb{E}\langle\phi(m)\rangle_{\beta,J}$. □

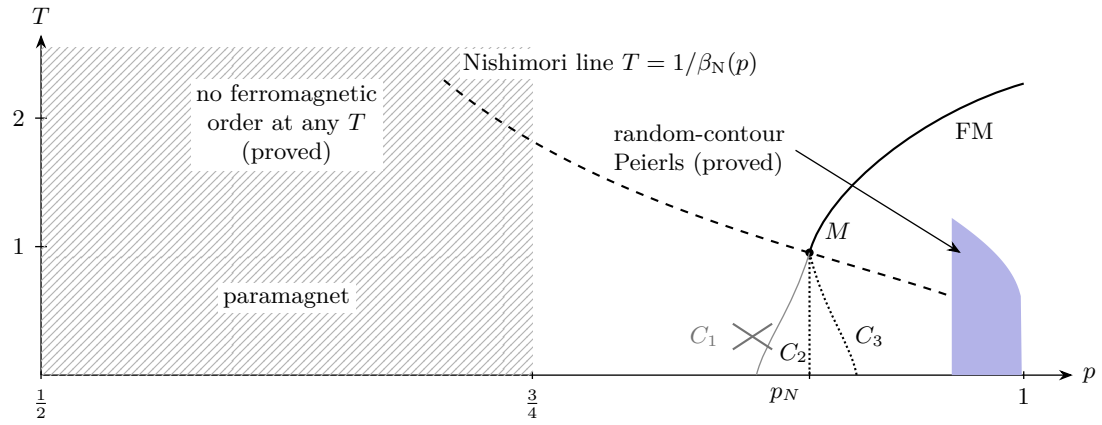


Figure 14.2: Schematic phase diagram of the biased $\pm J$ model on $\mathbb{Z}^2$ (p the fraction of positive bonds; FM the ferromagnetic phase). There is no ferromagnetic order at any temperature to the left of the point $M \approx (0.8908, 0.953)$ where the Nishimori line meets the ferromagnetic boundary (Proposition 14.4(a)); the hatched region $p < \frac{3}{4}$ is proved by the averaged FK law. Below M the boundary may be vertical (C_2) or reentrant (C_3), not of type C_1; which holds is open (Open problem 14.12). In the shaded strip $p > 0.9634$, $T \le 2/\beta_{\mathrm{N}}(p)$, random contours prove order (Theorem 14.8).

On the line, therefore, $\mathbb{E}\langle q^2 \rangle = \mathbb{E}\langle m^2 \rangle$ in every finite volume: there is overlap order exactly when there is ferromagnetic order, and the line never passes through a phase with $q > 0 = m$ [193, p. 56]. For the shape of the overlap law Nishimori and Sherrington draw a further conclusion, which rests on an input about the magnetization law.

Physics prediction 14.6 (No replica symmetry breaking on the line [196])**.** If in the thermodynamic limit the averaged magnetization law on the Nishimori line consists of two symmetric atoms, as for a ferromagnet, then so does the averaged overlap law; a phase with a nontrivial overlap distribution, if present, lies off the line, below it.

The identity is exact at every point; the two-atom form of the magnetization law is expected but not proved in finite dimension.

14.3 Griffiths inequalities on the line

For ferromagnets the second Griffiths inequality says that correlations increase with every coupling. It fails for spin glasses in general (Chapter 1), but a version holds on the Nishimori line if every bond is moved along the line, bias and temperature together.

Theorem 14.7 (Kitatani [140, (17)–(18)]; Morita–Nishimori–Contucci for Gaussian couplings [165])**.** *Let G be finite with free boundary conditions. Let the coupling of each edge e be independent with a law of Nishimori type $\nu_{b_e}(dJ) \propto e^{b_e J}\nu_{s,e}(dJ)$ (fixed symmetric part, parameter b_e), and let bond e carry its own inverse temperature $\beta_e = b_e$. Then for every edge B and every $C \subset V$, $\mathbb{E}\log Z$ and $\mathbb{E}\langle \sigma_C \rangle$ are nondecreasing functions of b_B, when $\beta_B = b_B$ moves with it.*

Proof sketch. We omit the computation. Differentiate in b_B: the derivative falls on the Boltzmann weight and on the law of J_B, for which $\partial_b \nu_b(dJ) = (J - \mathbb{E}_b J)\nu_b(dJ)$. Kitatani rewrites both terms

using two consequences of the Nishimori form, $\mathbb{E}f(J_B) = \mathbb{E}[\tanh(b_B J_B) f(J_B)]$ for odd f and a correlation inequality for even increasing f, and the gauge identities of Corollary 14.3, which turn the derivative into an average of squares [140, §3]. For Gaussian couplings Morita, Nishimori and Contucci use Gaussian integration by parts instead [165]; Okuyama and Ohzeki extend the inequalities, for Gaussian disorder, to the XY and $\mathbb{Z}_q$ gauge glasses and the gauge-invariant Potts model [204]. □

Two consequences are used repeatedly. Taking b_B from 0 (bond absent) to β_{N} shows that on the line adding a bond never decreases an averaged correlation, so free-boundary correlations on the line are nondecreasing in the volume and have infinite-volume limits. Moving all bonds together shows that $\mathbb{E}\langle\sigma_C\rangle$ at $(p, \beta_{\mathrm{N}}(p))$ is nondecreasing in p: the magnetized part of the Nishimori line is a half-line, and the line crosses the ferromagnetic boundary at most once. These are the monotonicity statements that Dyson-type comparisons need, and they drive the one-dimensional theorem of Chapter 12. Off the line no such monotonicity holds.

14.4 Order near full bias: random contours

Order on the Nishimori line for p close to 1 is a ferromagnetic statement, and the Peierls argument proves it, with the twist that the energy of a contour is random. Flipping the spins inside a contour γ changes the energy by twice the sum S_γ of the couplings it crosses; for biased couplings S_γ is typically of order $|\gamma|$, and a large-deviation bound replaces the deterministic Peierls estimate. Avron, Roepstorff and Schulman observed this for weakly diluted and weakly frustrated ferromagnets [19], as stated in [33, p. 43], where the argument is given for dilute ferromagnets [33, Thm. 2.6–Cor. 2.8]; Horiguchi and Morita work it out on $\mathbb{Z}^2$ for several laws, including $\pm J$ [119, §§2 and 4]. For $\pm J$ the exponential moment of $-S_\gamma$ is minimized exactly at the Nishimori coupling:

$$\min_t \mathbb{E}e^{-tJ_e} = \min_t \bigl(pe^{-t} + (1-p)e^{t}\bigr) = 2\sqrt{p(1-p)} =: a_p, \qquad \text{attained at } t = \beta_{\mathrm{N}}(p). \tag{14.1}$$

Horiguchi and Morita's convergence condition for the averaged contour sum is, for $\pm J$, $3a_p < 1$ at $2\beta \geq \beta_{\mathrm{N}}$ [119, (2.15)]. The form below replaces 3^n by the connective constant μ of the square lattice and uses a favourable finite patch so that convergence alone suffices; the same criterion $2\mu\sqrt{p(1-p)} < 1$ appears in the analysis of optimal decoding of the toric code [73, (65)–(67)].

Theorem 14.8 (Random-contour Peierls bound)**.** *Let the couplings on $\mathbb{Z}^2$ be iid ν_p with $p \geq 0.96$ and $\mu a_p < 1$, where μ is the connective constant of self-avoiding walks on $\mathbb{Z}^2$. Consider square boxes Λ with plus boundary condition and write $m_{\Lambda,x} = \langle\sigma_x\rangle^{+}_{\Lambda,J,\beta}$.*

(a) For every $\beta \geq \beta_{\mathrm{N}}(p)/2$ there is $c = c(p,\beta) > 0$ such that $\mathbb{E}m^2_{\Lambda,x} \geq c$ for every square box Λ and every $x \in \Lambda$. Hence, for two replicas with the same boundary condition, $\liminf_\Lambda \mathbb{E}\langle q_\Lambda\rangle \geq c$ and $\liminf_\Lambda \mathbb{E}\langle q^2_\Lambda\rangle \geq c^2$.

(b) At $\beta = \beta_{\mathrm{N}}(p)$, $\liminf_\Lambda \mathbb{E}\langle |\Lambda|^{-1}\sum_x \sigma_x\rangle^{+}_{\Lambda} \geq c$: the Nishimori point is ferromagnetic.

(c) For each $\beta \geq \beta_{\mathrm{N}}(p)/2$, almost every J has at least two infinite-volume Gibbs states.

With Couronné's bound $\mu \leq 2.662343$ [66], the hypothesis holds for $p > 0.963389$; with $\mu \leq 3$ it holds for $p > (1+\sqrt{8/9})/2 \approx 0.9714$.

Proof. Contours. Extend the spins by $+1$ outside Λ. For a simple dual polygon γ surrounding x, let S_γ be the sum of the couplings on the primal edges it crosses and E_γ the event that all these edges have disagreeing endpoints. Flipping the spins inside γ maps E_γ injectively into the disjoint event that all of them agree, and multiplies the Boltzmann weight by $e^{2\beta S_\gamma}$. Hence $\mu^+_{\Lambda,J,\beta}(E_\gamma) \le f_\beta(S_\gamma) := (1+e^{2\beta S_\gamma})^{-1}$. The disagreeing edges form an even subgraph of the dual lattice, which is a union of simple dual polygons; if $\sigma_x = -1$, a primal path from x to the boundary crosses it an odd number of times, so one of these polygons surrounds x. Thus

$$m_{\Lambda,x} \ge 1 - 2 \sum_{\gamma \text{ around } x,\ \gamma\subset\Lambda^*} f_\beta(S_\gamma). \tag{14.2}$$

Averaging. For $\beta \ge \beta_{\mathrm N}/2$ one has $f_\beta(S) \le e^{-\beta_{\mathrm N} S}$ for every real S (for $S \ge 0$ use $2\beta \ge \beta_{\mathrm N}$, for $S < 0$ the right side exceeds 1). Let G_R be the event that all couplings within distance R of x are $+1$. If γ has n edges, k of them in that ball, independence and (14.1) give $\mathbb{E}[e^{-\beta_{\mathrm N} S_\gamma} \mid G_R] = e^{-\beta_{\mathrm N} k} a_p^{n-k} \le a_p^n$, since $e^{-\beta_{\mathrm N}}/a_p = 1/(2p) \le 1$. The number N_n of polygons of length n around x is at most $\frac{n}{2} c_{n-1}$, where c_k counts k-step self-avoiding walks: the polygon crosses the horizontal half-line to the right of x at one of at most $n/2$ edges, and through a given dual edge there are at most c_{n-1} polygons. Since $c_k \le C_\rho \rho^k$ for every $\rho > \mu$, $\mu a_p < 1$ makes $\sum_n N_n a_p^n$ finite. As $R \to \infty$, every fixed polygon eventually lies in the ball, where $f_\beta(S_\gamma) = (1+e^{2\beta n})^{-1}$ on G_R; by dominated convergence

$$T_R := \sum_{\gamma \text{ around } x} \mathbb{E}[f_\beta(S_\gamma) \mid G_R] \longrightarrow \sum_n \frac{N_n}{1+e^{2\beta n}} \le \sum_{n \ge 4 \text{ even}} 2n\, 3^{n-2} e^{-2\beta n},$$

using $c_{n-1} \le 4 \cdot 3^{n-2}$. For $\beta \ge \beta_{\mathrm N}/2$ and $p \ge 0.96$, $e^{-2\beta} \le ((1-p)/p)^{1/2} \le 0.205$ and the right side is below 0.3. Fix R with $T_R < \frac12$ and put $d = 1 - 2T_R > 0$. By (14.2), $\mathbb{E}[m_{\Lambda,x} \mid G_R] \ge d$ for every box $\Lambda \ni x$, and by Jensen's inequality $\mathbb{E} m^2_{\Lambda,x} \ge \mathbb{P}(G_R)\, \mathbb{E}[m_{\Lambda,x} \mid G_R]^2 \ge p^{|B_R|} d^2 =: c$, where $|B_R|$ is the number of edges in the ball.

Overlap and magnetization. Since $\langle \sigma_x \tau_x \rangle = m^2_{\Lambda,x}$, averaging over x gives $\mathbb{E}\langle q_\Lambda \rangle \ge c$, and $\mathbb{E}\langle q_\Lambda^2 \rangle \ge (\mathbb{E}\langle q_\Lambda \rangle)^2$. At $\beta = \beta_{\mathrm N}$, Corollary 14.3(b) with the plus boundary condition gives $\mathbb{E} m_{\Lambda,x} = \mathbb{E} m^2_{\Lambda,x} \ge c$, which is (b).

Gibbs states. Take joint subsequential limits of the law of $(J, \mu^+_{\Lambda,J})$; they are supported on pairs with μ a Gibbs state for J, and the bound $\mathbb{E}\mu(\sigma_0)^2 \ge c$ survives because $\mu \mapsto \mu(\sigma_0)^2$ is continuous. So with positive probability there is a Gibbs state with $\mu(\sigma_0) \ne 0$, whose spin flip is a different Gibbs state. The event of non-uniqueness is translation invariant, and the disorder is ergodic, so it has probability one, as in [33, Cor. 2.8]. □

We give the form with μ and the favourable patch ourselves; its ingredients are classical. A finer local analysis of contours, certified by computer, lowers the bias further for the overlap at very low temperature.

Theorem 14.9 (computer-assisted)**.** *Let the couplings on $\mathbb{Z}^2$ be iid ν_p with $p \ge 0.9537$, and let $\beta \ge 10^6$. Then the thermodynamic Edwards–Anderson response $-\partial_{\lambda+} f^{(2)}(\beta, 0)$ is positive, where $f^{(2)}(\beta,\lambda)$ is the quenched free energy of two replicas coupled by $-\lambda \sum_x \sigma_x \tau_x$. For the disorder-dependent boundary condition $\eta(J)$, where $\eta_x = -1$ at the sites x of one checkerboard class at which at least three of the four incident couplings are negative and $\eta_x = +1$ elsewhere, square boxes have $\liminf_\Lambda \mathbb{E}\langle q_\Lambda \rangle \ge c_0 > 0$; and almost every J has more than one infinite-volume Gibbs state.*

The proof is the contour argument of Theorem 14.8 with three refinements: a gauge transformation on one checkerboard class, which makes the transformed four-bond stars independent; local estimates that count several low-cost images of a contour at its corners and short caps; and a finite-memory walk matrix that keeps the exact correlations inside transformed stars. The computer checks one inequality, $\mathcal{M}v \le \frac{9999}{10000}\, v$ entrywise for an explicit positive integer vector v and an explicit nonnegative rational matrix $\mathcal{M}$ with 453,233,520 rows encoding the weights of the refined contour sum, confirmed in exact arithmetic by the generating program and a separately written checker; everything else is proved by hand (Section A.29; full details in [212]). A closely related computer-assisted method, failure paths counted by finite-memory transfer matrices with an integer potential, gives a lower bound 0.0465 on the threshold of minimum-weight decoding of the surface code under independent bit flips [234]; that is a zero-temperature decoding statement, and neither result implies the other. The conclusion does not distinguish a ferromagnet from a phase with zero magnetization and positive overlap, and the method is close to its limit: the local estimates level off near $p \approx 0.951$ (a heuristic extrapolation), against the numerical zero-temperature threshold $p \approx 0.897$.

Route warning 14.10 (Biased order does not reach the fair glass). *Tempting step:* approach the symmetric model from the biased side, by proving order on or near the Nishimori line for smaller and smaller bias. *Status: the methods of this chapter cannot cross* $p = \frac{1}{2}$. By Theorem 14.5 every order theorem on the line is a ferromagnetic theorem, and the Peierls bounds above use the positive mean of S_γ. As $p \downarrow \frac{1}{2}$, $\beta_{\mathrm{N}}(p) \to 0$: the line runs to infinite temperature, and the gauge identities reduce to the statement that averaged correlations vanish (Chapter 1). The zero-bias glass sits at the end of the temperature axis that the line never visits.

14.5 Below the multicritical point: vertical or reentrant?

Below M the question left open by Proposition 14.4 is whether cooling at fixed p can destroy order, so that the boundary bends to the right (*reentrance*). The two-temperature form of Theorem 14.5 turns this into a statement about overlaps between temperatures.

Corollary 14.11 (Reentrance and two-temperature overlaps)**.** *On a finite graph with free boundary conditions and iid couplings of Nishimori type, for every* β, $\mathbb{E}\langle m^2\rangle_\beta = \mathbb{E}\langle q^2\rangle_{\beta\otimes\beta_{\mathrm{N}}}$, *where the right side is the overlap of independent replicas at* β *and at* β_{N}. *Hence, along a sequence of boxes, the point* (p,β) *is ferromagnetic in the sense* $\liminf \mathbb{E}\langle m^2\rangle_\beta > 0$ *if and only if the two-temperature overlap between* β *and* $\beta_{\mathrm{N}}(p)$ *is nontrivial in the same sense. In particular, if for every* $\beta \ge \beta_{\mathrm{N}}(p)$ *the overlap between* β *and* $\beta_{\mathrm{N}}(p)$ *is nontrivial, there is no reentrance at* p*; and if there is reentrance at* p*, the overlap between the cold temperature and the Nishimori temperature vanishes.*

Proof. Apply Theorem 14.5 with $\phi(x) = x^2$. □

Nishimori, Ohzeki and Okuyama derive these identities in a model that interpolates the disorder, and relate reentrance to temperature chaos [198, (14)–(15)]: reentrance makes the configurations at the cold and the Nishimori temperature decorrelate completely. The theory does not decide whether either phenomenon occurs [198, Rem. 2].

Heuristic arguments point both ways. Nishimori's argument from the singularity of the entropy of frustration suggests a vertical boundary ("not a rigorous proof" [193, p. 63]), and a modified $\pm J$ model with correlated disorder, which provably has a spin-glass phase, supports

the same picture under an unproved assumption [193, §4.8]. Le Doussal and Harris argue from gauge invariance that the multicritical point lies on the Nishimori line [148]. Numerical analyses support a slightly reentrant boundary on $\mathbb{Z}^2$, with the zero-temperature threshold above the multicritical bias by about 0.006 [117, p. 4]; the question is regarded as "still unsettled" [241, p. 38].

Open problem 14.12 (Verticality)**.** For iid $\pm J$ couplings on $\mathbb{Z}^2$ (or $\mathbb{Z}^d$, $d \geq 2$), decide whether the boundary of the ferromagnetic phase below the multicritical point is vertical, that is, whether the zero-temperature critical bias equals the bias p_N at which the Nishimori line leaves the ferromagnetic phase, or strictly reentrant. Nishimori's inequality gives that the zero-temperature threshold is at least p_N; neither equality nor strict inequality is proved.

14.6 The duality conjecture

The location of M on $\mathbb{Z}^2$ has a conjectured closed form, from a duality argument for the replicated model. It is a heuristic, and it is numerically not exact, but the derivation is short and shows where the self-duality of the square lattice enters.

Physics prediction 14.13 (Nishimori–Nemoto [195]; Nishimori [194, Conj. 2])**.** The multicritical point of the $\pm J$ model on $\mathbb{Z}^2$ lies on the Nishimori line at the bias p_{dual} with

$$-p_{\mathrm{dual}} \log p_{\mathrm{dual}} - (1 - p_{\mathrm{dual}}) \log(1 - p_{\mathrm{dual}}) = \tfrac{1}{2} \log 2, \qquad p_{\mathrm{dual}} = 0.889972\ldots.$$

Heuristic derivation. This is not a proof: it uses the replica limit and an identification of the multicritical point with a self-dual point, neither of which is justified. Average Z^n over the couplings. On each edge the averaged weight depends only on the number k of replicas in which the two endpoints disagree, $x_k = pe^{(n-2k)\beta} + (1-p)e^{-(n-2k)\beta}$. The Wu–Wang duality of the $\mathbb{Z}_2^n$ model maps the square lattice to itself and the weights to their discrete Fourier transform; the principal weight becomes $x_0^* = 2^{-n/2} \sum_k \binom{n}{k} x_k = 2^{-n/2} (2 \cosh \beta)^n$ [194, (2)–(5)]. On the line, $e^{2\beta} = p/(1-p)$, the ansatz $x_0 = x_0^*$ reads

$$p^{1+n} + (1-p)^{1+n} = 2^{-n/2}.$$

Both sides equal 1 at $n = 0$; equating the coefficients of n gives $p \log p + (1-p) \log(1-p) = -\frac{1}{2} \log 2$. Nishimori shows that for each integer n the replicated system is self-dual with a fixed point at $x_0 = x_0^*$, and that if a certain function in the duality relation is analytic and the replicated free energy has a unique singularity, the singularity is at that fixed point [194, Thm. 2, Cor. 2]. Analyticity is not proved, nor are the replica limit and the existence of the multicritical point on the line [194, p. 7]. □

Route warning 14.14 (Duality does not give the exact multicritical point). *Tempting step:* take the duality value as exact, or use duality to prove absence of a transition on self-dual lattices. *Status: numerically not exact; heuristic.* Monte Carlo finite-size scaling locates the multicritical point at $p^* = 0.89081(7)$, $T^* = 0.9528(4)$, so that $p^* - p_{\mathrm{dual}} = 0.00084(7)$ [117, pp. 1, 5]; on three pairs of mutually dual hierarchical lattices the predicted sum rule $H(p_1) + H(p_2) = 1$ (in bits) comes out as 1.0172, 0.9829, 0.9911 [194, p. 8]. A refined duality with larger clusters gives 0.890813 on the square lattice [199], in agreement with the numerics. For any probabilistic use, the dual of the random-bond model has weights that are not positive [241, p. 38].

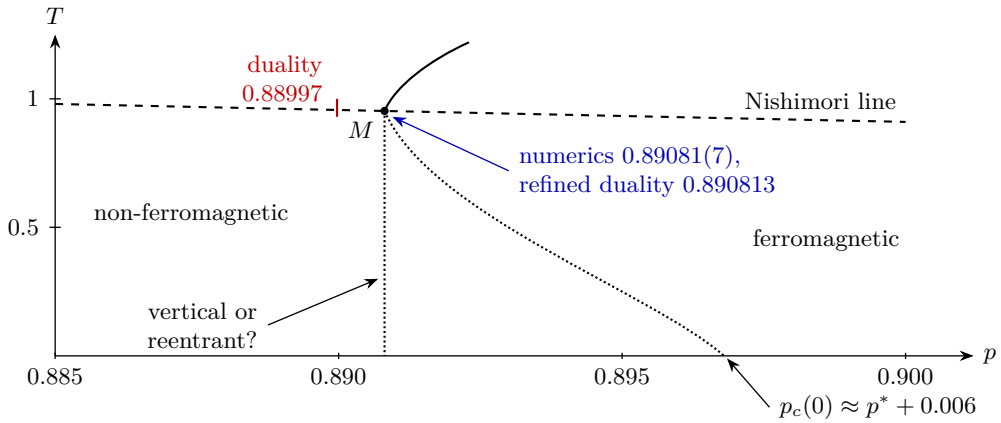


Figure 14.3: Near the multicritical point M on $\mathbb{Z}^2$ (axes stretched; all numbers are numerical estimates or conjectures). Duality places M at $p = 0.88997$, left of the Monte Carlo value 0.89081(7), which a refined duality reproduces. Below M numerics favour a small reentrance, with zero-temperature threshold $p_c(0)$; neither alternative is proved (Open problem 14.12).

14.7 The inference reading

The Nishimori identity has a direct statistical meaning. Plant a uniformly random configuration $\theta \in \{-1,1\}^V$ and observe each edge through a binary symmetric channel: $Y_{xy} = \theta_x\theta_y\xi_{xy}$ with iid $\xi_{xy} = \pm 1$, $\mathbb{P}(\xi = +1) = p$. By Bayes' rule the posterior law of θ given Y is $\mathbb{P}(\theta \mid Y) \propto \prod_{xy} \nu_p(Y_{xy}\theta_x\theta_y) \propto e^{\beta_{\mathrm{N}} \sum_{xy} Y_{xy}\theta_x\theta_y}$, the Gibbs measure of the $\pm J$ model with couplings Y at the Nishimori temperature [197].

Observation 14.15 (One identity in three languages). Let J be iid ν_p and let τ be uniform on $\{-1,1\}^V$, independent of J. Then (τ, J^τ) has the law of (θ, Y) in the planted model, and the Nishimori identity of Lemma 14.2 is the statement that for every function f,

$$\mathbb{E}\big[\langle f(\sigma, Y)\rangle_{\mathrm{post}}\big] = \mathbb{E} f(\theta, Y),$$

the posterior average of f has the same expectation as f evaluated at the truth (the "Nishimori property" of Bayes-optimal inference). The Nishimori–Sherrington identity at $\beta = \beta_{\mathrm{N}}$ is the special case $f(\sigma, Y) = \langle \phi(q(\sigma,\sigma'))\rangle_{\mathrm{post},\sigma'}$, with $q(\sigma,\sigma') = |V|^{-1}\sum_x \sigma_x\sigma'_x$ and σ' an independent posterior sample: the overlap of a posterior sample with the truth has the law of the overlap of two independent posterior samples.

Proof. $\mathbb{P}(\tau, J^\tau = y) = 2^{-|V|}\prod_{xy}\nu_p(y_{xy}\tau_x\tau_y) = \mathbb{P}(\theta = \tau, Y = y)$. Given Y, a posterior sample σ and the truth θ have the same conditional law, so $\mathbb{E}\langle f(\sigma, Y)\rangle_{\mathrm{post}} = \mathbb{E} f(\theta, Y)$. To see that this is Lemma 14.2, put $\bar{F}(J,\sigma) = 2^{-|V|}\sum_\tau f(\sigma\tau, J^\tau)$, which is gauge covariant. Then $\mathbb{E}\bar{F}(J,+) = \mathbb{E} f(\theta, Y)$, and substituting $\sigma \mapsto \sigma\tau$ in the Gibbs average gives $\mathbb{E}\langle \bar{F}(J,\sigma)\rangle_{\beta_{\mathrm{N}},J} = \mathbb{E}\langle f(\sigma, Y)\rangle_{\mathrm{post}}$. Conversely, for covariant F put $f(\sigma, Y) = F(Y,\sigma)$; then $\bar{F} = F$, so the two statements are equivalent. □

This identification is our own way of organizing well-known facts, each of which is in [193, 197]. It explains why the line appears in coding and synchronization problems: for the toric code with perfect syndrome measurement, optimal decoding is posterior inference of this kind, and Dennis,

Kitaev, Landahl and Preskill identify the accuracy threshold of the code with the Nishimori-line critical point, numerically $1 - p_N \approx 0.109$ [73, §IV.F].

14.8 Planar models and free fermions

On planar graphs the Ising partition function at fixed couplings is a determinant, for couplings of either sign.

Theorem 14.16 (Kac–Ward–Feynman–Sherman, signed weights; Aizenman–Warzel [7, Thm. 4.1 and p. 12])**.** *Let $G = (V, E)$ be a finite planar graph drawn with straight edges, and $J \in \mathbb{R}^E$ arbitrary. Then*

$$\frac{Z_{G,J}(\beta)}{2^{|V|}\prod_e \cosh(\beta J_e)} = \sqrt{\det(1 - KW)},$$

where W is the diagonal matrix on oriented edges with entries $\tanh(\beta J_e)$, K is the non-backtracking matrix on oriented edges with the phase $e^{i\vartheta/2}$ for a turn by the angle ϑ, and the square root is the principal value.

Aizenman and Warzel prove the identity for small weights through a graph zeta function and extend it to all weights by analyticity, which covers signed couplings; so the planar Edwards–Anderson model at fixed J is a free-fermion model. This does not decide two-dimensional ordering: the disorder average of $\log\det(1 - KW)$ is a random-matrix problem in symmetry class D, studied by network models and field theory [222], and bounds that depend only on the moduli $|\tanh(\beta J_e)|$ cannot see frustration, so they cannot exceed the high-temperature region of the ferromagnet with couplings $|J_e|$. One exact fact about symmetric laws is elementary.

Proposition 14.17 (Read–Ludwig [222, §II, (5)–(9)])**.** *Let G be finite, the couplings iid with a symmetric law, and γ a set of edges (for planar G, the edges crossed by a dual path between two faces α and α'). Let J^γ be J with the signs reversed on γ, and $X = Z_{G,J^\gamma}/Z_{G,J} = \langle \mu_\alpha \mu_{\alpha'} \rangle_J$, the Kadanoff–Ceva disorder correlator. Then X and $1/X$ have the same law, $\log X$ is symmetric, and $\mathbb{E}X^m \geq 1$ for every integer m.*

Proof. For symmetric ν the map $J \mapsto J^\gamma$ preserves the law of J and exchanges $Z_{G,J}$ and Z_{G,J^γ}, so $X \overset{d}{=} 1/X$. Hence $\mathbb{E}X^m = \frac{1}{2}\mathbb{E}(X^m + X^{-m}) \geq 1$ by $a + a^{-1} \geq 2$. □

Writing $X = e^{-\beta \Delta F_\gamma}$, the free-energy cost ΔF_γ of the defect line γ is symmetric around 0 for every symmetric law. Read and Ludwig use this, with $|\langle \sigma\sigma \rangle| \leq 1$, to exclude a metallic phase of class D in the random-bond Ising model with symmetric couplings [222, p. 3]: a rigorous inequality applied to a field-theory prediction. No probabilistic statement about two-dimensional order has yet been derived from the fermionic representation.

14.9 Summary

On the Nishimori line one identity (Lemma 14.2) gives the exact energy, Nishimori's inequality, the equality of the magnetization and overlap laws and, with the averaged FK law, absence of ferromagnetic order for $p < \frac{3}{4}$ on $\mathbb{Z}^2$; Griffiths inequalities hold along the line; random contours prove ferromagnetic order at the Nishimori point for $p > 0.9634$, and a computer-assisted refinement proves overlap order and Gibbs non-uniqueness for $p \geq 0.9537$ at $\beta \geq 10^6$. Open are

the shape of the boundary below the multicritical point (Open problem 14.12), its location, and every statement at zero bias, which the line does not reach (Route warning 14.10).

Open problem 14.18 (Order on the Nishimori line of $\mathbb{Z}^2$)**.** Prove ferromagnetic order at the Nishimori point for every p above the line's transition point p_N, or at least narrow the range between the proved $p > 0.9634$ and the numerical $p_N \approx 0.891$. On the other side, prove that the line is disordered for some $p > \frac{3}{4}$, so that the averaged FK threshold of Proposition 14.4 is not the magnetic one.

The long-range chain on the Nishimori line is in Chapter 12; the conditional reductions that would connect biased and symmetric order, and why they fail, are in Chapter 15 and Chapter 16.

Part VII

Where proofs break, and what to aim at

Chapter 15

Barriers

Fifty years of attempts have produced no proof that the short-range Edwards–Anderson model orders in any dimension, and, for $\pm J$ or Gaussian couplings, no proof that it fails to order at every positive temperature in $d = 2$ (for diluted laws just above the percolation threshold, Chapter 13 gives one). Some of the reasons are theorems: a natural class of arguments, from the routes of the earlier chapters, cannot reach the target. Heuristically they share one feature (Remark 15.13): every argument that works on a tree, a series–parallel network or a hierarchical lattice uses a resource that the cubic lattice with fair iid couplings lacks. Provably, two cheaper substitutes for the exact fair average over coupling signs, fixed signs and every positive tilt of the sign law, already violate the comparison inequalities on bundles (Proposition 15.2). By the end the reader can state nine such barriers with proofs or proof ideas, and test a new candidate inequality on small graphs and laws (Section 15.11).

Figure 15.1 shows the routes of this book as a siege: each line of attack, what it has taken, and the wall that has held it.

15.1 What a barrier is

A barrier speaks about all arguments of a given kind at once.

Definition 15.1 (Method class and barrier)**.** A *method class* is specified by the information its arguments may use about the coupling law and the Gibbs measures (a family of marginals, identities or inequalities) and by a *target*: one of the order-parameter targets (T1)–(T3) of Chapter 1, their negations, or a stated comparison inequality. A *barrier* for the class is one of the following:

(i) two models that agree on all the information the class uses and disagree on the target (*indistinguishability*);
(ii) a model satisfying the premises of the class for which an inequality the class needs is false (*false universal inequality*);
(iii) a proof that the estimates the class produces cannot reach the target at any scale (*quantitative ceiling*).

A barrier blocks an argument only when the hypotheses match; it says nothing about logical independence or computational hardness. Most results here are the author's; they are proved here or in the chapter cited, we say when one rests on a hypothesis, and computations not yet independently checked are labelled as observations.

15.2 The fair average must be used exactly

Question. Can a proof replace the exact average over iid fair coupling signs by something more robust: a law with a small bias, the planted law of the Nishimori line, the annealed law, or a bound that holds for every fixed sign pattern?

The sharpest test is the arboreal-gas comparison of Chapter 11, forest domination (FD). On a finite multigraph $G = (V, E)$ with magnitudes $K_e = \beta|J_e| \geq 0$ and iid fair signs ε_e, write $\langle\cdot\rangle_{G,\varepsilon K}$ for the free Gibbs measure with couplings $\varepsilon_e K_e/\beta$, put $y_e = \sinh^2 K_e$, and let $P^{\mathrm{AG}}_{G,y}$ be the arboreal gas, the probability measure on spanning forests $F \subset E$ with weight $\prod_{e\in F} y_e$. The comparison conjectured in Chapter 11 is

$$\mathbb{E}_\varepsilon \langle\sigma_u\sigma_v\rangle^2_{G,\varepsilon K} \;\leq\; P^{\mathrm{AG}}_{G,y}(u \leftrightarrow v) \qquad (u, v \in V). \tag{15.1}$$

On a bundle of k parallel edges between u and v, $\langle\sigma_u\sigma_v\rangle_J = \tanh(\sum_i \varepsilon_i K_i)$ and $P^{\mathrm{AG}}(u \leftrightarrow v) = r/(1+r)$ with $r = \sum_i y_i$; for the fair law (15.1) holds on every bundle (the bundle theorem of Chapter 11). It holds there with no room to spare in the direction of the law.

Proposition 15.2 (Edge of validity on bundles)**.** *(a)* Fixed signs. *On two parallel edges with the same deterministic sign and magnitude $K > 0$, $\tanh^2(2K) > P^{\mathrm{AG}}(u \leftrightarrow v) = 2y/(1+2y)$.*

(b) Tilted laws. *For $\alpha > 0$ let $P_\alpha(\varepsilon) \propto Z_\varepsilon^\alpha$ on $\{\pm1\}^E$, where Z_ε is the partition function with couplings $\varepsilon_e K_e$; $\alpha = 1$ is the planted (Nishimori) law and $\alpha = 2$ the annealed law. On three parallel edges of equal magnitude K, for every $\alpha > 0$ there is $K_0(\alpha)$ such that for $K > K_0(\alpha)$,*

$$\mathbb{E}_{P_\alpha}\langle\sigma_u\sigma_v\rangle_J^2 > P^{\mathrm{AG}}(u \leftrightarrow v).$$

Proof. (a) With $t = \tanh K$, $\tanh^2(2K) = 4t^2/(1+t^2)^2$ and $2y/(1+2y) = 2t^2/(1+t^2)$; their ratio is $2/(1+t^2) > 1$. (b) is the tilt proposition of Chapter 11. With $S = \varepsilon_1 + \varepsilon_2 + \varepsilon_3$ one has $Z_\varepsilon \propto \cosh(KS)$, so the tilt raises the weight of the two frustration-free patterns $|S| = 3$ against the six others by the factor $(\cosh 3K/\cosh K)^\alpha \to \infty$. This drives $(1 - \mathbb{E}_{P_\alpha}\langle\sigma_u\sigma_v\rangle_J^2)(1+3\sinh^2 K)$ to 0, while $(1 - P^{\mathrm{AG}}(u \leftrightarrow v))(1 + 3\sinh^2 K) = 1$. □

Figure 15.2 shows the three regimes. The proof of (b) only uses that the tilt reweights the frustration-free patterns by a factor that grows with K. For a product bias the failure is already on two parallel edges: if the signs are iid with mean $m \neq 0$, then $\mathbb{E}\langle\sigma_u\sigma_v\rangle_J^2 = \frac{1+m^2}{2}\tanh^2 2K$, and its ratio to $P^{\mathrm{AG}}(u \leftrightarrow v) = 2y/(1+2y)$ is $(1+m^2)/(1+\tanh^2 K)$, which exceeds 1 whenever $\tanh^2 K < m^2$. A method that is robust to a small bias of the sign law, or that only uses an inequality valid for the planted or annealed law, would therefore also prove a false statement, and (a) shows that no pointwise-in-signs version of (15.1) exists.

Lesson. A proof of a comparison of the type (15.1) must use that the signs are exactly fair and independent. *Escape.* Identities that hold only at the fair point, such as the Efron–Stein decomposition over the sign cube with exact one-sign derivatives used for bundles in Chapter 11, or the gauge identities of Section 15.9 combined with independence.

15.3 Forest information cannot decide order

Question. Is partial information about the coupling law, such as its restriction to forests or its marginals on small sets of bonds, enough to decide whether the model orders?

No. Let $A = (A_e)$ be any law of signs on the edges of a graph, η an independent iid fair field of vertex signs, and $J_{xy} = \eta_x A_{xy} \eta_y$. By the forest lemma of Chapter 13, applied for each value of A, the signs $(J_e)_{e \in F}$ on every finite forest F are iid fair, even conditionally on A, and the same holds for adaptive queries whose union stays a forest. On cycles the law of J remembers the frustration $\prod_{e \in C} A_e$ of A, and two backgrounds with opposite behaviour give a barrier of type (i).

Theorem 15.3 (Forest-information barrier)**.** *Let D be the decorated square lattice in which each edge of $\mathbb{Z}^2$ is replaced by two internally disjoint paths of length two, with unit coupling magnitudes. Let $A^+ \equiv +1$, and let A^- give, on every such pair of paths, sign product $+1$ to one path and -1 to the other. Put $J^\pm = \eta A^\pm \eta$ as above. Then J^+ and J^- have the same law on every forest of D and the same adaptive forest-query transcripts, while:*

(a) if $K(\beta) = \log\cosh 2\beta > \frac{1}{2}\log(1+\sqrt{2})$, the model J^+ has at least two Gibbs states, and in the gauge image of the plus state $\langle \sigma_x \sigma_y \rangle_J^2 \geq m^*(K(\beta))^4 > 0$ for all vertices x, y of $\mathbb{Z}^2$, where $m^*(K)$ is the spontaneous magnetization of the square-lattice Ising model at coupling K;

(b) for every $\beta < \infty$ the model J^- has a unique Gibbs state, and $\langle \sigma_x \sigma_y \rangle_J = 0$ whenever x and y are at graph distance greater than 2; in particular $\mathbb{E}\langle q_\Lambda^2 \rangle = O(|\Lambda|^{-1})$ for boxes Λ.

Proof sketch. We give the mechanism; the infinite-volume steps are those of Chapter 13. Summing the two midpoint spins of a decorated edge with end spins s, t gives weight $4\cosh^2[\beta(s+t)] = 4\cosh(2\beta)\, e^{K(\beta)st}$ when both paths are positive, and $4\cosh[\beta(s+t)]\cosh[\beta(s-t)] = 4\cosh 2\beta$, independent of s, t, when the paths have opposite signs. So A^+ decimates exactly to the square-lattice ferromagnet at coupling $K(\beta)$, and (a) follows as in the no-transition theorem for decorated lattices of Chapter 13 (with one diamond per edge), together with the second Griffiths inequality on the backbone. Under A^- every decorated edge is a frustrated diamond, so in the cancelling-blocks lemma of Chapter 13 the effective graph has no edges: for every β the Gibbs state is unique, the backbone spins are independent fair signs, and given them the midpoint spins are independent, each with mean $\tanh(\beta(a_1 s + a_2 t)) = \frac{1}{2}(a_1 s + a_2 t)\tanh 2\beta$, where s, t are its two backbone neighbours and a_1, a_2 the signs of its two edges. Hence the correlation of two spins vanishes unless their backbone neighbourhoods meet, which forces graph distance at most 2. Gauge transformations do not change squared correlations. □

Remark 15.4. Neither J^+ nor J^- is iid. The same indistinguishability holds on the ordinary square lattice with unit magnitudes and full-support, finite-range, centred sign laws. Take independent blocks of black-checkerboard plaquettes with flux -1 with probability $1-\delta$ and $+1$ with probability δ, with uniform sign vectors given the fluxes. For fixed β and $0 < \delta \leq 24^{-4/\varepsilon}$, where $\varepsilon = \ell/(4(\log 3 + \ell))$ and $\ell = -\log(1 - e^{-4\beta})$, every Gibbs state has zero one-spin means and every doubled state zero local overlap, almost surely; yet at fixed $\beta \geq 10^6$ this law has the same joint sign law on every forest as a randomly gauged biased law, which orders there (Chapter 14). (Idea: the rare unfrustrated plaquettes are isolated, a Peierls count shows that a frustrated background of plaquettes destroys one-spin means, and forest marginals see only the uniform sign vectors. We omit the details.) Relatedly, a proof that selects a subforest of $\mathbb{Z}^d$ using the whole disorder, as a scaffold, gains nothing: for iid couplings with $|J_e| < \infty$, almost surely the Ising model on every nearest-neighbour subforest, with any signs and at every β, has a unique Gibbs state, and CMR blue clusters sampled on it are finite. (Sketch: on a tree the root magnetization in the plus state is at most $\sum_{e \in \Pi} \prod_{e' \leq e} \tanh(\beta |J_{e'}|)$ for every cutset Π; a subtree

of $\mathbb{Z}^d$ has at most $(2n+1)^d$ vertices at depth n; and a first-moment count over self-avoiding paths shows that almost surely every path of length $n \geq n_0$ has $\prod_e \tanh(\beta|J_e|) \leq e^{-cn}$.)

Lesson. A proof of order or non-order must use the cycle statistics of the actual iid law. *Escape.* Arguments that see the law of frustrations on cycles, such as the exact sign average of Chapter 11, or that use the independence of the couplings, which the gauged laws lack.

15.4 Percolation is not order

Question. Does percolation of a graphical representation decide spatial order?

No, in either representation. Frustrated Fortuin–Kasteleyn clusters percolate at low temperature for every ± 1 configuration on $\mathbb{Z}^2$ [71] (Chapter 7), where no order is expected. Both overlap sectors of the CMR colouring contain infinite blue clusters in high dimension where $\chi_{\mathrm{SG}} < \infty$ and the sectors have exactly equal densities [214] (Chapter 8), and the SK paramagnet has two giant blue clusters of equal density [152]. Percolation is even: it is invariant under the reversal of one replica, which exchanges the two overlap sectors, while order is detected by an odd quantity; in the recodings of Gandolfi, Newman and Stein [104] (Chapter 7) magnetization and overlap are density imbalances between signed companions of the infinite cluster.

Lesson. A cluster signature needs a cancellation statement: an imbalance between the sectors, or a bound on the mass of finite clusters inside the infinite one. *Escape.* An estimate that treats the sectors differently, the open step of Chapter 8.

15.5 The qualitative hypotheses do not imply order

Question. Do the four qualitative features used by heuristic arguments for order (symmetric iid couplings, a percolating set of bonds, an ordered ferromagnet with the same magnitudes, frustrated plaquettes of positive density) imply order?

Corollary 15.5 (A counterexample on $\mathbb{Z}^d$; from the diluted uniqueness theorem of Chapter 13)**.** *Let $d \geq 2$. There is $g_d > 0$ such that for the diluted law $\mu_p = (1-p)\delta_0 + \frac{p}{2}(\delta_{+1} + \delta_{-1})$ with $p_c^{\mathrm{bond}}(d) < p < p_c^{\mathrm{bond}}(d) + g_d/2$:*

(i) μ_p is symmetric and the couplings are iid;

(ii) the bonds with $J_e \neq 0$ percolate, and the ferromagnet with couplings $|J_e|$ on them orders at large β;

(iii) frustrated plaquettes have positive density;

and yet, for μ_p-a.e. J, the EA Gibbs measure is unique simultaneously for every $\beta \in [0, \infty)$.

Proof. (i) is the definition, and (iii) holds because a unit square has all four bonds present with an odd number of negative signs with probability $p^4/2 > 0$, so frustrated plaquettes have density $p^4/2$ by the ergodic theorem. The first half of (ii) is $p > p_c$; the second is recorded with the uniqueness statement in [217], whose main theorem gives the conclusion (Chapter 13). □

The mechanism is an exact cancellation (Chapter 13): a frustrated isolated diamond transmits zero effective coupling at every temperature, and removing such diamonds leaves a subcritical effective bond process.

Lesson. An ordering proof must use a quantitative input: supercriticality well above threshold, magnitudes without atoms, or the geometry of $\mathbb{Z}^d$. *Escape.* Hypotheses that fail for μ_p near p_c and for decorated graphs, such as the bounded coupling density of the hierarchical proofs of Chapter 12.

15.6 No positivity survives the quenched average

Question. Do the positivity properties behind ferromagnetic proofs (Griffiths inequalities, FKG, reflection positivity, positive random-current weights) survive averaging over symmetric couplings?

Four of them already fail in Chapter 1, on graphs with one or two plaquettes: the second Griffiths inequality (on a 2×3 box near $\beta \approx 0.695$ [64]), monotonicity of $\mathbb{E}\langle\sigma_x\sigma_y\rangle_J^2$ under adding a fair bond, quenched positivity with a ferromagnetic reference (asked about in [63]), and the FKG lattice condition for the averaged overlap law of the free square. This section adds positive double-current weights and bond reflection positivity. All the witnesses have at most eight vertices and one source: the two-replica law contains the factor Z_J^{-2}, which depends on the signs.

Lemma 15.6 (Loop normalization)**.** *On the cycle C_n with iid fair signs J_e, let $F = \prod_e J_e$, $t = \tanh\beta$ and $c = \cosh\beta$. Then $Z_J = 2^n c^n(1 + Ft^n)$, so $\mathbb{E}F = 0$ but*

$$\mathbb{E}\Big[\frac{F}{Z_J^2}\Big] = -\frac{2t^n}{2^{2n}c^{2n}(1-t^{2n})^2} < 0.$$

In particular, in the two-replica current expansion of $\mathbb{E}\langle\sigma_x\sigma_y\rangle_J^2$ the term pairing one x–y arc with the complementary arc has weight proportional to $\mathbb{E}[F/Z_J^2]$, which is negative.

Proof. The high-temperature expansion of Z_J on a cycle has two even subgraphs, the empty one and the cycle. Averaging over $F = \pm 1$ gives $\frac{1}{2}[(1+t^n)^{-2} - (1-t^n)^{-2}] = -2t^n/(1-t^{2n})^2$. The current statement follows because the two arcs together carry every edge once, with weight $\beta^n F$. □

Frustrated cycles have smaller partition functions, hence larger weight Z_J^{-2} in the quenched two-replica law; this negative correction breaks reflection positivity on an even cycle at every temperature (Fig. 15.3).

Proposition 15.7 (Reflection positivity fails on the six-cycle)**.** *On the cycle with vertices $0, \dots, 5$ and edges $\{i, i+1\}$ (indices mod 6), let the couplings be iid fair signs, $\beta > 0$, and let P be the law of $q = (\sigma_i\tau_i)_i$ under $\mathbb{E}[\mu_J \otimes \mu_J]$. Let $\theta(i) = 5 - i$, a reflection exchanging $\{0,1,2\}$ and $\{3,4,5\}$ and fixing the edges $\{5,0\}$ and $\{2,3\}$. With $C = \cosh 2\beta$, define $f(q_0, q_1, q_2)$ to be 1 at $(+,+,+)$, $-C^2$ at $(+,-,+)$ and 0 otherwise, and $\theta f(q) = f(q_5, q_4, q_3)$. Then*

$$\mathbb{E}_P[f\,\theta f] = -\frac{kC^6}{\mathcal{Z}} < 0, \qquad k = \frac{2t^6v^6}{1+t^{12}}, \quad \mathcal{Z} = (C+1)^6 + (C-1)^6 - 2kC^6,$$

where $t = \tanh\beta$ and $v = \tanh 2\beta$. So P is not bond reflection positive.

Proof. Given q, the second replica is $\tau = q\sigma$, and the two-replica weight is $\exp(\beta\sum_i J_i\sigma_i\sigma_{i+1}(1 + q_iq_{i+1}))$: edges with $q_iq_{i+1} = -1$ drop out, and the $m(q)$ others carry the coupling $2\beta J_i$. If q is not constant these edges form a forest and the sum over σ is $2^6C^{m(q)}$; if q is constant it is $2^6C^6(1 + Fv^6)$. Dividing by Z_J^2 from Lemma 15.6 and averaging over the fair sign F with $a = t^6$, using $\mathbb{E}_F(1+Fa)^{-2} = (1+a^2)/(1-a^2)^2$ and $\mathbb{E}_F F(1+Fa)^{-2} = -2a/(1-a^2)^2$, gives

$$P(q) = \mathcal{Z}^{-1}C^{m(q)}\big(1 - k\,\mathbf{1}\{q \text{ constant}\}\big),$$

and $\mathcal{Z}$ is the stated normalization because $\sum_q C^{m(q)} = (C+1)^6 + (C-1)^6$ (a 2×2 transfer matrix). Without the correction, $C^{m(q)}$ is a nearest-neighbour ferromagnetic weight. On the

support of $f\,\theta f$ both halves have $q = +$ at their ends, so both fixed edges contribute C^2 and the uncorrected form factorizes as $C^2\big(\sum_{q_0 q_1 q_2} f\, C^{m_{\text{left}}}\big)^2$; the sum is $1 \cdot C^2 + (-C^2) \cdot 1 = 0$. The correction acts only on the two constant configurations, and of these only $(+, \dots, +)$ is in the support, with $f\,\theta f = 1$ and weight $-kC^6/\mathcal{Z}$. □

The construction works on every even cycle of length at least six, and the obstruction survives on the smallest full rectangle with a bond reflection and more than one row.

Proposition 15.8 (Reflection positivity fails on the $\mathbf{2 \times 4}$ box; computer-assisted)**.** *Let $V = \{0,1\} \times \{0,1,2,3\}$ with all ten nearest-neighbour edges, free boundary conditions, iid fair signs and $\beta = \log 2$, and let P be the averaged overlap law as above. For the reflection $\theta(r,c) = (r, 3-c)$ write the left half as $a = q_{(0,0)}$, $a' = q_{(1,0)}$ (outer column) and $b = q_{(0,1)}$, $b' = q_{(1,1)}$ (inner column), and let $f = 0$ if $b \neq b'$ and $f = s\,(5 - 9 \cdot \mathbf{1}\{a = a' = s\})$ if $b = b' = s$. Then $\mathbb{E}_P[f\,\theta f] \approx -0.0033682 < 0$, an explicit rational number whose exact value is given in Appendix B.*

Proof sketch. We omit the enumeration (Appendix B): at $\beta = \log 2$ every Boltzmann weight is a power of 4 up to a common factor, so the value is an exact rational sum over 1024 disorders (six distinct partition functions) and 256 spin configurations. By continuity the strict inequality persists for β near $\log 2$ and for symmetric coupling laws close to the fair ± 1 law. □

Corollary 15.9 (No local reflection-positive lift)**.** *In the setting of Proposition 15.8, there is no reflection-positive probability measure on an enlarged configuration space with a map to overlap configurations that commutes with the reflections, makes the left-half overlaps measurable with respect to the enlarged left-half variables, and has P as its image.*

Proof. Composing f with the left-half overlaps would give an admissible half-space function with a negative reflection form. □

Lesson. The quenched average destroys positivity already on one or two plaquettes, through the sign dependence of Z_J^{-2}. The annealed ensemble removes this factor, but it changes the model, and by Proposition 15.2 it is on the wrong side of the comparison inequalities. *Escape.* Inequalities specific to one observable (such as the comparisons of Chapter 11), a disorder-averaged infrared bound that does not pass through reflection positivity, or a positive representation that is not local in the sense of Corollary 15.9. No lattice-specific positivity statement on infinite cubic boxes is ruled out.

15.7 Scale-uniform control is the wall

Question. Where do local and one-step arguments stop, and why does the hierarchical ordering proof not transfer to $\mathbb{Z}^d$?

The ordering proofs on Migdal–Kadanoff lattices [59, 219] (Chapter 12) have a simple skeleton: nested blocks, a failure event at each level, and a union bound over the levels. On a general graph: let G be finite with couplings J and zero field, and let two replicas be drawn independently, each with the Gibbs law with fixed spins on a root set T_N (the root spins may be random). A *separator chain* for a vertex x is a decreasing sequence of vertex sets $D_N \supseteq \cdots \supseteq D_j$ with terminal sets $T_k \subset D_k$, such that every edge meeting $D_k \setminus T_k$ lies in D_k, the interiors $D_k \setminus T_k$

decrease, and $T_{k-1} \subset T_k \cup Y_k$ with $Y_k = T_{k-1} \setminus T_k$ inside $D_k \setminus T_k$, and $Y_j = \{x\}$. For terminal spins s on T_k let $\mu_k(\cdot \mid s)$ be the Gibbs law on $D_k \setminus T_k$, and

$$\phi_k(s) = 1 - \sum_{\eta \in \{\pm 1\}^{Y_k}} \mu_k(\sigma_{Y_k} = \eta \mid s)^2,$$

the probability that two independent samples disagree somewhere on the new terminals.

Lemma 15.10 (Separator chains)**.** *Let* Cons_k *be the event that the overlap* $q = \sigma\tau$ *is constant on* T_k*, and* $t_0 \in T_N$*. For every* J*,*

$$\mathbb{P}(q_x \neq q_{t_0},\ \mathrm{Cons}_N \mid J) \le \sum_{k=j}^{N} \mathbb{E}\big[\mathbf{1}_{\mathrm{Cons}_k}\, \phi_k(\sigma_{T_k}) \mid J\big] \le \sum_{k=j}^{N} \max_s \phi_k(s).$$

Proof. If Cons_k holds with value c and $q \equiv c$ on Y_k, then Cons_{k-1} holds with value c; so the event on the left lies in the union over k of $\mathrm{Cons}_k \cap \{q \not\equiv c \text{ on } Y_k\}$. Given both replicas on T_k, each has law $\mu_k(\cdot \mid \text{its terminal spins})$ on $D_k \setminus T_k$ (Markov property); on Cons_k with value c, $\tau = c\sigma$ on T_k, so by spin-flip symmetry $c\tau_{Y_k}$ is an independent copy of σ_{Y_k}, and $q \not\equiv c$ on Y_k with conditional probability $\phi_k(\sigma_{T_k})$. □

On the diamond lattice the terminals of a block are its two poles and Y_k is the midpoint of the branch containing x, with field $h = \xi_1 s_a + \xi_2 s_b$ given the poles (ξ_1, ξ_2 the effective couplings of the two halves); so $\phi_k = \frac{1}{2}\operatorname{sech}^2 h \le \frac{1}{2}\operatorname{sech}^2||\xi_1| - |\xi_2||$ for every pole configuration, and growing effective couplings make these bounds summable. On bounded-degree graphs with bounded couplings this worst-case form cannot work (Fig. 15.4).

Proposition 15.11 (A scale-independent floor)**.** *Let* D *be a finite block with terminal set* T*, couplings* J*, zero field, and for* $t \in T$ *let* m_t *be the sum of* $|J_{tv}|$ *over the neighbours* $v \in D \setminus T$ *of* t*. For every* $x \in D \setminus T$ *there is a terminal configuration* s^* *with*

$$2p(1-p) \ge \tfrac{1}{2}\operatorname{sech}^2(\beta \max_{t \in T} m_t), \qquad p = \mu(\sigma_x = +1 \mid s^*).$$

Consequently, on a graph of maximal degree Δ *with* $|J_e| \le M$*, every separator chain has* $\max_s \phi_k(s) \ge \frac{1}{2}\operatorname{sech}^2(\beta\Delta M)$ *at every level with* $Y_k \neq \emptyset$*, and the worst-case bound of Lemma 15.10 is not summable over unboundedly many levels, at any* β*.*

Proof. By spin-flip symmetry, $p(-s) = 1 - p(s)$. Start from a configuration s with $p(s) \ge \frac{1}{2}$ and flip the terminals one at a time until reaching $-s$. Flipping a terminal t changes the log-weight of every interior configuration by at most $2\beta m_t$, so the odds $O = p/(1-p)$ change by a factor in $[e^{-4\beta m_t}, e^{4\beta m_t}]$. Since $\log O$ goes from ≥ 0 to ≤ 0, at some flip of a terminal t one of the two configurations s^* has $|\log O(s^*)| \le 2\beta m_t$. Finally $2p(1-p) = 2O/(1+O)^2 = \frac{1}{2}\operatorname{sech}^2(\frac{1}{2}\log O)$. For the consequence, $\phi_k(s)$ is at least the disagreement probability $2p(1-p)$ at any single $y \in Y_k$, and $\max_t m_t \le \Delta M$. □

A proof on $\mathbb{Z}^d$ must therefore control per-level failures in a Gibbs-typical rather than worst-case sense. The cleanest form of what is needed is a criterion on b-adic blocks. On the torus $(\mathbb{Z}/b^N\mathbb{Z})^d$, call the scale-$k$ blocks the cubes of side b^k in the b-adic hierarchy; for a block B let its *sign* be the sign of $\sum_{x \in B} q_x$ (with a fixed tie-breaking site), and call B *incoherent* if some child block has a different sign. Let ε_k be the average over scale-k blocks of the probability that the block is incoherent.

Proposition 15.12 (Multiscale coherence criterion)**.** *For every law of two configurations* (σ, τ) *on the torus, with* $Q = b^{-Nd} \sum_x q_x$*,*

$$\mathbb{E}Q^2 \geq \Big(1 - 2\sum_{k=1}^{N} \varepsilon_k\Big)_+^2.$$

In particular, for the EA model, if $\sum_k \varepsilon_k^{(N)} \leq \eta < \frac{1}{2}$ *uniformly in* N*, then* $\liminf_N \mathbb{E}Q^2 > 0$*.*

Proof. If every block containing x at scales $1, \dots, N$ is coherent, the sign of the torus equals the sign of the site block $\{x\}$, which is q_x. So the fraction F of sites with q_x different from the torus sign satisfies $\mathbb{E}F \leq \sum_k \varepsilon_k$, by averaging over x. Since the torus sign times Q equals $1 - 2F$, $\mathbb{E}|Q| \geq 1 - 2\sum_k \varepsilon_k$, and $\mathbb{E}Q^2 \geq (\mathbb{E}|Q|)^2$. □

The criterion is an obligation without a supplier. Heuristically it is compatible with the droplet picture, where ε_k decays like a power of the block size, and not with the trivial–nontrivial picture of Chapter 2; on binary hierarchies it has an exact telescoping form (Chapter 12).

Remark 15.13 (Contours and suppliers)*.* Two further results of the author point the same way; we indicate only the arguments. For unit couplings on $\mathbb{Z}^d$, $d \geq 2$, local cut certificates for interfaces, which interpolate one cut at a time, give a divergent contour sum at every temperature: they certify a cost of at most $1/(2(2d-1))$ per cut edge, while the number of contours around the origin grows exponentially in their size at a larger rate. Assuming the droplet scaling hypothesis for large walls (heuristic), every fixed-block contour supplier fails at every block size at which the hypothesis holds, since the origin is then enclosed by a wall of linear size R with probability at least of order $R^{-\theta}$, while a supplier needs bounds exponentially small in the wall's area. Exact elimination of spins on $\mathbb{Z}^d$ produces multi-spin interactions with positive probability (Chapter 12), which excludes an exact scalar recursion for pair couplings. Heuristically, every proof we know of order or of summable per-scale bounds in a random environment uses a small parameter (giving only decay), a positive mean surface cost, a planted structure, a linear structure with multiscale analysis, or an exact recursion with unbounded degree; the symmetric iid model on $\mathbb{Z}^d$, $d \geq 3$, at low temperature has none of the last four.

The non-ordering side meets the same wall: the comparison inequalities of Chapter 11 are proved on every series–parallel network by composing two-terminal blocks, the boxes of $\mathbb{Z}^2$ are not series–parallel, and the cut inequality that holds on every graph controls one cut at a time, while the minimal cut between two far points stays bounded.

Lesson. Bounds uniform over boundary conditions at each scale have a scale-independent floor on bounded-degree graphs; summable control must be Gibbs-typical. *Escape.* None is known. The obligation is the hypothesis of Proposition 15.12 (on the non-ordering side, the comparison on boxes); on hierarchical lattices unbounded degree supplies it.

15.8 What a comparison reference must contain

Question. Comparison with a reference model has reached furthest on the non-ordering side (Chapter 11). Which properties must the reference have? Table 15.1 collects the requirements: the first is Proposition 15.2, the others come from the warnings and observations on comparison models of Chapter 11 and from computations by the author that are not reproduced here.

Lesson. On this evidence (the proved rows and the observations of the table), a comparison reaching all temperatures in $d = 2$ should use compact Ising spins, exact $\pm J$ cancellations and

Requirement	What fails without it	Status
The exact fair sign average	every positive tilt, the planted and the annealed law violate (15.1) on the triple edge	proved (Proposition 15.2)
Compact spins	the Gaussian (connection-Laplacian) analogue of the glass violates the forest comparison on the Wheatstone bridge	observation (exact)
Spins that are units of $\mathbb{R}$, $\mathbb{C}$ or $\mathbb{H}$	tree-calibrated gauge glasses with $SO(3)$, $O(2)$, $SO(4)$, $O(4)$ or $U(2)$ connections violate it on two parallel edges; at fourth order the margin is positive exactly for $\mathbb{Z}_2$, $U(1)$, $SU(2)$	observation
Exact $\pm J$ cancellations	a uniform-flux XY reference fails on the triple edge (three independent high-precision methods) and, at equal activities, on a ring of twelve squares in $\mathbb{Z}^2$ closed by one edge, where each $\pm J$ square cancels with probability $\frac{1}{2}$ and a uniform flux never cancels	high-precision numerics (triple edge); observation (Monte Carlo, ring)
Cluster weight $q = 0$	random-cluster references with $q < 0$ fail; for $q \in (0, 1]$ they cannot give two-dimensional non-order once $\sinh^2 \beta > 1/q$	proved (Chapter 11)
A large-scale harmonic mechanism	tree-like references give non-order only above a finite temperature	heuristic

Table 15.1: Requirements on a reference model for comparison inequalities of the type (15.1). "Observation" marks exact or numerical computations by the author that are not reproduced here and have not yet been independently checked.

a reference with a large-scale harmonic mechanism. *Escape.* The arboreal gas meets every requirement (it matches the glass on trees, has cluster weight zero and is controlled by effective conductances); what is missing is the comparison on boxes.

15.9 Disorder-conjugacy tools are blind to order

Question. The random-field model was solved with identities from differentiating in the disorder: Gaussian integration by parts, the Aizenman–Wehr martingale [8], the stability estimates of Dario, Harel and Peled [68] (Chapters 4 and 10). Can their EA analogues separate order from non-order?

The first-order identities follow from gauge invariance alone. For an edge $e = xy$ write $s_e = \sigma_x \sigma_y$ and $C_{ef}(J) = \langle s_e s_f \rangle_J - \langle s_e \rangle_J \langle s_f \rangle_J$.

Proposition 15.14 (Gauge diagonalization)**.** *Let G be a finite simple graph with zero field and free or periodic boundary conditions, and let the law of J be invariant under all gauge transformations (not necessarily a product law). Then $\mathbb{E} C_{ef} = 0$ for $e \neq f$, and for every gauge-invariant $w : \mathbb{R}^E \to \mathbb{R}^E$ with $\mathbb{E}|w|^2 < \infty$,*

$$\mathbb{E}\Big[\sum_{e,f} w_e w_f C_{ef}\Big] = \mathbb{E}\Big[\sum_{e} w_e^2 C_{ee}\Big] \leq \mathbb{E}|w|^2.$$

For Gaussian couplings, $\mathbb{E}[J_f \langle s_e \rangle_J] = \beta\, \mathbb{E} C_{ef}$, which therefore vanishes for $e \neq f$. All of this holds for the Mattis law [157] $J_{xy} = \eta_x \eta_y K$, η iid fair, which on $\mathbb{Z}^d$, $d \geq 2$, is ordered when βK is large: the gauge image of the symmetric mixture of the plus and minus states of the ferromagnet has $\langle \sigma_0 \sigma_x \rangle_J^2 \geq m^*(\beta K)^4 > 0$ for all x, where $m^*(\beta K)$ is the spontaneous magnetization of the Ising model on $\mathbb{Z}^d$ at coupling βK.

Proof. For $X \subset V$ let J^X be J with the signs flipped on the edges with exactly one endpoint in X. Then $\langle g(\sigma) \rangle_{J^X} = \langle g(\sigma^X) \rangle_J$, where σ^X flips the spins in X, so $C_{ef}(J^X) = \pm C_{ef}(J)$ with the

sign -1 exactly when exactly one of e, f crosses the boundary of X. If $e \neq f$, choose a vertex x of e that is not a vertex of f (possible because G is simple) and $X = \{x\}$: then $C_{ef}(J^X) = -C_{ef}(J)$, and gauge invariance gives $\mathbb{E}C_{ef} = -\mathbb{E}C_{ef}$. The same argument applies with the gauge-invariant weights $w_e w_f$, and $C_{ee} \leq 1$. The Gaussian identity is integration by parts. The Mattis law is gauge invariant, gauge transformations preserve squared correlations, and in the symmetric plus–minus mixture of the ferromagnet $\langle\sigma_0\sigma_x\rangle \geq \langle\sigma_0\rangle_+\langle\sigma_x\rangle_+ = m^*(\beta K)^2$ by FKG. □

So the first-order identities behind the random-field arguments hold equally in an ordered model. What separates the EA and Mattis laws is the product structure of the couplings, while the gauge-covariant directions the identities need are not product directions; at second order integration by parts gives another collective-mode variance, not a closed bound (compare the warning on edge means in Chapter 4).

Lesson. Disorder-conjugacy identities follow from gauge invariance and cannot by themselves decide order. *Escape.* An essential use of the independence of the couplings, or an observable conjugate to some randomness together with a monotonicity, which for the random-field model is FKG (Chapter 10).

15.10 Biased order does not transfer to fair disorder

Question. Order is proved for strongly biased couplings and on the Nishimori line (Chapter 14). Can it be transported to fair couplings by a change of measure?

A generic transfer turns a large expectation of an observable $0 \leq F \leq 1$ under the biased product law P_p on n bonds ($P_p(J_e = +1) = p$) into a statement under the fair law $P_{1/2}$.

Proposition 15.15 (Exponential cost of transfer)**.** *For $p \neq \frac{1}{2}$ and $c \in (0,1)$ let $B_n(c) = \inf\{\mathbb{E}_{1/2}F : \ 0 \leq F \leq 1, \ \mathbb{E}_p F \geq c\}$. Then*

$$\lim_{n\to\infty} \frac{1}{n} \log B_n(c) = -D(p\|\tfrac{1}{2}), \qquad D(p\|\tfrac{1}{2}) = p\log(2p) + (1-p)\log(2(1-p)).$$

Proof. Let $L_n = \mathrm{d}P_p/\mathrm{d}P_{1/2}$; under P_p, $n^{-1}\log L_n \to D$ almost surely. For $\varepsilon > 0$, $\mathbb{E}_p F = \mathbb{E}_{1/2}[FL_n] \leq e^{n(D+\varepsilon)}\mathbb{E}_{1/2}F + P_p(\log L_n > n(D+\varepsilon))$, and the last term tends to 0, so $\mathbb{E}_{1/2}F \geq (c - o(1))e^{-n(D+\varepsilon)}$. Conversely $F_n = \mathbf{1}\{\log L_n \geq n(D-\varepsilon)\}$ has $\mathbb{E}_p F_n \to 1$ and $\mathbb{E}_{1/2}F_n = \mathbb{E}_p[L_n^{-1}F_n] \leq e^{-n(D-\varepsilon)}$. □

Gauge-invariant observables do not escape this. The frustrations $\Phi_i = \prod_{e\in\partial P_i} J_e$ of n vertex-disjoint plaquettes are iid with mean $(2p-1)^4$ under the biased law and fair under $P_{1/2}$, and every gauge-invariant function of the couplings on these plaquettes is a function of (Φ_i). So Proposition 15.15, applied to the Φ_i, gives the same exponential cost with rate $D(p'\|\frac{1}{2})$ per plaquette, $p' = \frac{1}{2}(1+(2p-1)^4)$; for instance $F_n = \mathbf{1}\{n^{-1}\sum_i \Phi_i \geq \frac{1}{2}(2p-1)^4\}$ has $\mathbb{E}_p F_n \to 1$ and $\mathbb{E}_{1/2}F_n \leq \exp(-n(2p-1)^8/8)$ by Hoeffding's inequality. In infinite volume the two product laws are mutually singular.

Lesson. Order proved for biased or planted couplings says nothing about fair couplings through a change of measure; the proofs of Chapter 14 use the bias (a positive mean surface cost, or the Nishimori identities) in every step. *Escape.* A structure shared exactly by the biased and fair models, such as the gauge identities, with an argument that does not degrade with the bias.

15.11 A kill-test suite

Before trying to prove a new candidate inequality, evaluate it on the small graphs and laws of Fig. 15.5: it should pass where the fair glass satisfies (15.1) and fail where the barriers say any valid argument must. Its author should also say which cycles of the disorder it retains and which state or metastate supplies its replicas.

(K1) *Triple edge and fixed-sign double edge.* The fair glass satisfies (15.1) on every bundle (Chapter 11); fixed signs violate it on two edges (Proposition 15.2(a)).

(K2) *Ten or more parallel edges at high temperature.* The uniform-flux XY comparison fails on $k \geq 10$ parallel edges at the small coupling tested and holds there for $k \leq 9$ (observation, with a rigorous error bound on the difference); as $\mathbb{Z}^2$ has at most four disjoint paths between two vertices, this separates multigraph from planar statements.

(K3) *A positive tilt and the planted law.* A candidate, or an argument, that survives unchanged under $P_\alpha \propto Z^\alpha$, $\alpha > 0$, cannot imply (15.1), by Proposition 15.2(b).

(K4) *The Wheatstone bridge and K_4,* the smallest graphs that are not series–parallel. The fair glass satisfies (15.1) on both, for every pair and all activities (exact certificates, Appendix B); its Gaussian sibling violates the forest comparison on the bridge (observation).

(K5) *Corners of the 3×3 grid.* At a degree-two terminal, $\mathbb{E}\langle\sigma_u\sigma_v\rangle_J^2$ is the sum of the influences of its two edges (observation), so global Poincaré-type bounds are exact reformulations of (15.1) there; per-path bounds along monotone paths fail between opposite corners (observation, exact values). The fair glass satisfies (15.1) for every pair at every homogeneous activity on the free 2×3, 2×4, 2×5, 3×3, 3×4 grids and the 3×3 torus (certificates, Appendix B).

(K6) *A strong edge in parallel with a long chain of squares.* A reference without exact $\pm J$ cancellations fails here at equal couplings, on a subgraph of $\mathbb{Z}^2$ (observation; Table 15.1).

(K7) *Decorated lattices* (Theorem 15.3, Chapter 13): a criterion that sees only forest information, percolation, or the qualitative hypotheses of Corollary 15.5 cannot be sound here.

15.12 Summary

Apart from those marked as observations or heuristics, the barriers of this chapter are proved statements about classes of arguments, not about the EA model, and none excludes order or non-order. A proof must use the exact fair average over independent signs, the law of frustrations on cycles, a quantitative input and an odd cluster signature; it must do without positivity of the quenched law and use more than first-order disorder identities; and above all it needs control that is uniform across scales and typical rather than worst-case (Section 15.7), as on trees, series–parallel networks and hierarchical lattices. Chapter 16 collects the open statements that would close the gap.

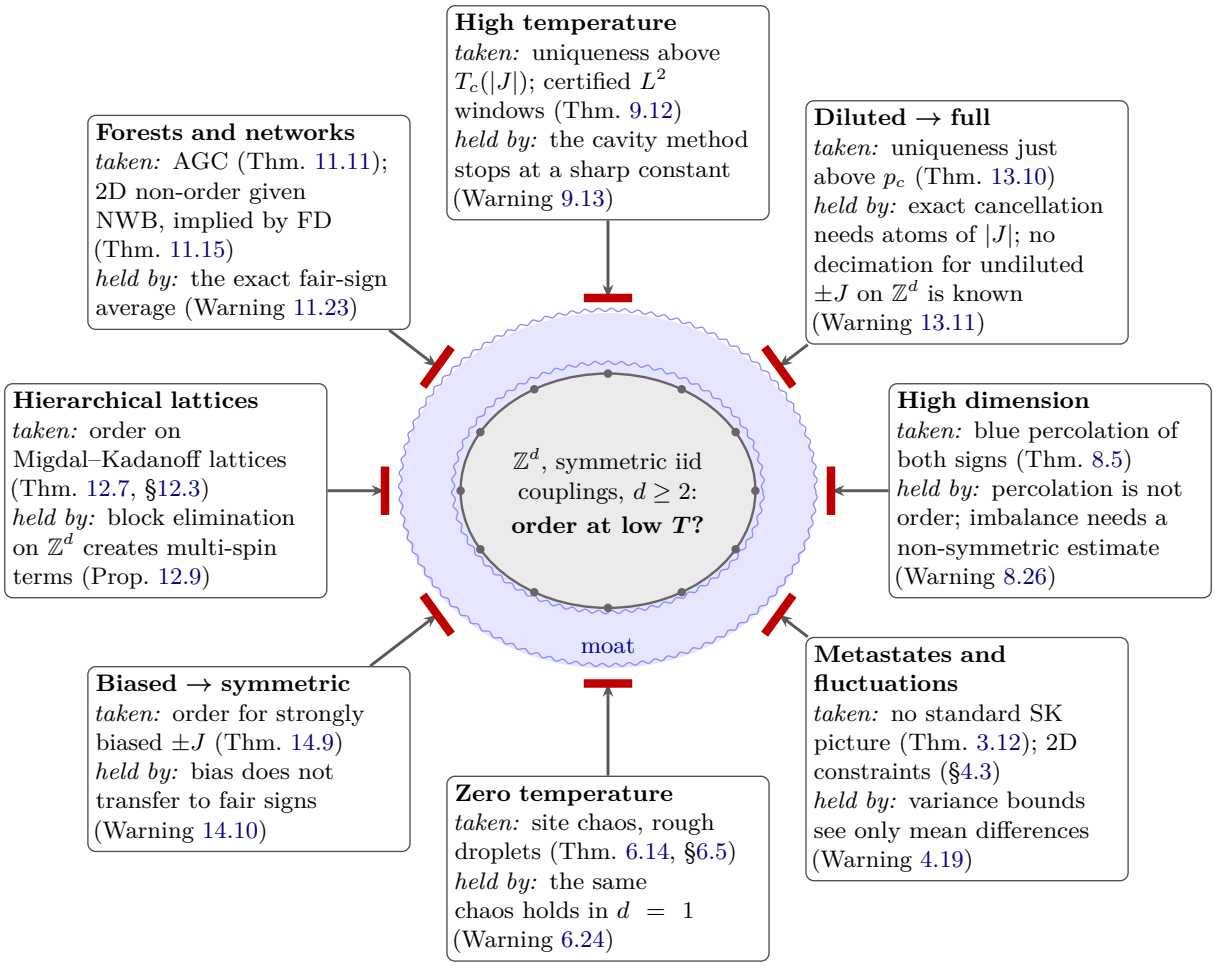


Figure 15.1: The siege of the short-range spin glass. Each box is a line of attack, with what it has taken and the bottleneck that has held it so far (red wall); the numbers point to the theorem and to the warning or barrier where each is explained. Temperature runs vertically: down from high temperature and up from zero temperature. Dimension and geometry run horizontally: down from high dimension and across from hierarchical lattices. The diagonals come from nearby models: diluted and biased couplings, forests and electrical networks, metastates. The walls are the subject of this chapter.

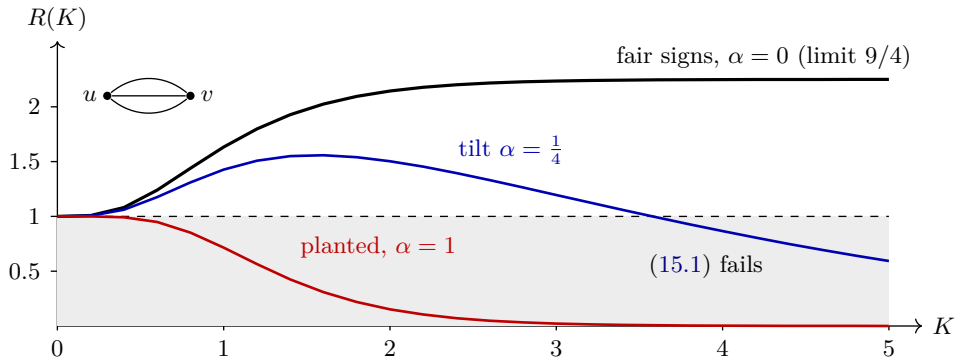


Figure 15.2: The triple edge: R is the ratio of $1-\mathbb{E}\langle\sigma_u\sigma_v\rangle_J^2$ to $1-P^{\mathrm{AG}}(u\leftrightarrow v)$, and (15.1) holds exactly when $R\geq 1$. For fair signs R increases to $9/4$; every positive tilt eventually drops below 1, the planted law already at small couplings. An argument that applies to a tilted law cannot prove the comparison.

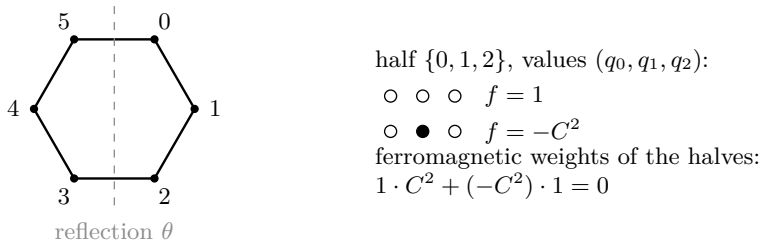


Figure 15.3: The six-cycle witness of Proposition 15.7 (open circles: overlap +1). The two half-configurations cancel exactly in the ferromagnetic part of the overlap law, so the reflection form sees only the correction from frustrated cycles, which is negative because they carry the larger weight Z_J^{-2}.

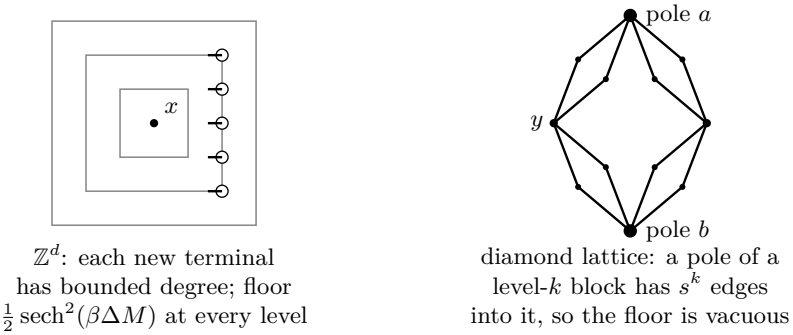


Figure 15.4: Why the hierarchical ordering proof does not transfer. On $\mathbb{Z}^d$ (left) terminals of nested blocks have bounded degree, and Proposition 15.11 bounds the worst-case failure of every level below, independently of the scale. On the diamond lattice (right) the poles have unboundedly many edges into the block and their effective couplings grow.

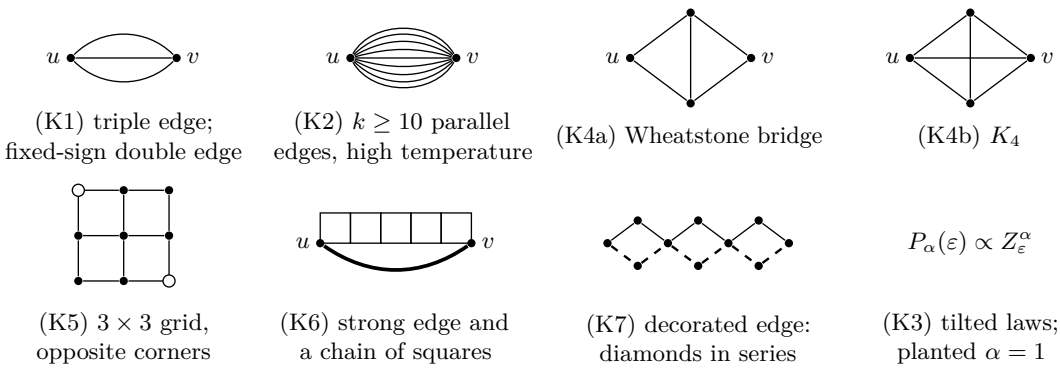


Figure 15.5: The kill-test suite: bundles (K1, K2) and tilted laws (K3) test the use of the exact fair average; (K4) are the smallest graphs that are not series–parallel; (K5) has degree-two terminals of a box; (K6) tests exact $\pm J$ cancellations, (K7) the use of cycle information.

Chapter 16

Open problems and programmes

What should one try to prove next? The central question of Chapter 1, order in $d \geq 3$ and, for undiluted laws such as $\pm J$ and Gaussian couplings, its absence at every positive temperature in $d = 2$, is far from any present method, and Chapter 15 explains several reasons why. Between it and what is known there is a layer of precise statements that are open, that would change the picture if proved, and for which an attack point can be named. This chapter ranks them (Section 16.1), draws the conditional reductions that connect them to the central question (Section 16.2), and describes the programmes that work on them, with their frontiers and the open problems they state themselves, quoted with locators (Section 16.3). It ends with the status of Newman's research problems of 1997 (Section 16.4).

The ranking is a judgment about what might be done within a year with the tools of this review. It weighs three things: how close the nearest proved statement is, whether the barriers of Chapter 15 leave room, and how much the result would change. It is not a ranking of importance; the central question itself is ranked last because no route to it is in sight.

16.1 Ranked targets

Table 16.1 lists the targets. The first four are stated below as open problems; the others are stated in the chapters cited.

Two dimensions at zero temperature. In $d = 2$ the fluctuation programme of Arguin, Newman, Stein and Wehr proves, for continuous coupling laws with $\mathbb{E}J^4 < \infty$ (so not for $\pm J$), that all metastates have the same mean edge correlations, at positive temperature [14, Cor. 4.3] and, for zero-temperature metastates and symmetric such laws, in the analogous form of [16, Thm. 5.5 and Sect. 6] (both stated in Chapter 4, the first with a proof and the second with a proof sketch). These restrict atomic metastates, since weights tag covariant families of ground-state pairs. Newman and Stein have developed restricted metastates and fluctuation bounds between incongruent states as part of a programme toward uniqueness in two dimensions [187–189]; among their results, two zero-temperature metastates that are each supported on a single ground-state pair are supported on the same pair [189, Sect. 6.1]. The natural next step of the programme is the following.

Open problem 16.1 (Several ground-state pairs in two dimensions)**.** Let $d = 2$ and let ν be continuous and symmetric with $\mathbb{E}J^4 < \infty$. Show that no zero-temperature periodic metastate is

	Target	Nearest proved statement and attack point	Where
1	$d = 2$, $T = 0$: no periodic metastate on $M \geq 2$ incongruent ground-state pairs	for continuous symmetric laws with $\mathbb{E}J^4 < \infty$, all zero-temperature metastates have the same mean edge correlations; needs a pathwise input (Observation 16.2)	Open problem 16.1, Chapters 4 and 5
2	forest or resistance domination (FD, RD) on the boxes of $\mathbb{Z}^2$, or on one infinite family that is not series–parallel	FD proved on bundles, a series–parallel class, the Wheatstone bridge and K_4; RD on all series–parallel networks; either on the boxes of $\mathbb{Z}^2$ would give two-dimensional non-order at every temperature	Open problem 16.3, Chapter 11
3	$\pm J$ couplings on Migdal–Kadanoff lattices left open by the dilution criterion; whether $\Lambda(b, s) = 1$ occurs for Gaussian couplings	exact dilution after one or several decimation steps; Gaussian classification	Open problem 16.4, Chapter 12
4	one-dimensional long-range chains on the Nishimori line for $\frac{3}{2} \leq \alpha \leq 2$	transition for $1 < \alpha < \frac{3}{2}$ by Dyson's method	Open problem 16.5, Chapters 12 and 14
5	an integrable rate in the positivity principle	conditional order on the perturbed Dyson hierarchical model	Chapter 12
6	both-sector blue percolation in some $d \leq 6$	proved for all large d, and at seventeen explicit temperatures for fair ± 1 couplings in $d = 7$–12 (computer-assisted; revised preprint not yet public); no present engine reaches $d \leq 6$, and for unit couplings the Machta–Newman–Stein auxiliary model is subcritical at $d = 6$	Chapter 8
7	edge chaos in the strong sense; a surface-order variance bound for ground-state energy differences	the edge overlap is not chaotic in the averaged sense; the surface bound for periodic against antiperiodic boundary conditions	Chapters 4 and 6
8	FK percolation together with a unique Gibbs state on a lattice; power-law decay under single FK percolation	FK percolation on $\mathbb{Z}^2$ for every ± 1 configuration at $\beta > \frac{1}{2} \log 17$; at most one infinite FK cluster for iid fair ± 1 couplings in periodic limits	Chapter 7
9	persistent blue imbalance; a Gibbs-typical supplier of multiscale coherence on $\mathbb{Z}^d$	both reduce order to one estimate (Fig. 16.1)	Chapters 8 and 15
10	the central question	—	Chapter 1

Table 16.1: Targets ranked by how reachable they look with present tools. The column "Where" points to the statement and to the chapters with the background.

supported, for ν-a.e. J, on $M \geq 2$ incongruent ground-state pairs; already the case of two pairs with weight $\frac{1}{2}$ each is open.

What makes this case hard is a property of the method, not of the model.

Observation 16.2 (What the variance method can see). The variance bounds of Chapter 4 control the J-conditional mean of an energy difference between two covariant kernels, whose derivatives with respect to the couplings are differences of mean edge correlations. For a metastate giving

weight $\frac{1}{2}$ to each of two incongruent pairs, with an exchangeable pair law, this conditional mean vanishes identically, as it does for a metastate on one pair. So no argument that uses only J-conditional means of difference functionals, edge means and the covariance axioms can exclude such metastates. An exclusion needs an input of another kind: pathwise minimality of ground states under finite flips, the single-wall geometry of two ground-state pairs drawn from periodic metastates [181] (Chapter 5), or the structure of critical droplets (Chapters 5 and 6).

The observation is the route warning on edge means of Chapter 4, read as a guide: of the three inputs, the single-wall theorem already says, for continuous symmetric laws of unbounded support, that two distinct pairs from periodic metastates differ by one doubly infinite domain wall of positive density, so the remaining task is to show that such a wall cannot be part of a metastate. Positive temperature is further away. Even a full solution of Open problem 16.1 leaves the positive-temperature statement of Newman's Research Problem 4.2, since a metastate on one ordered flip pair is consistent with every fluctuation bound (Chapter 4).

Two dimensions at positive temperature: comparison. The comparison route of Chapter 11 targets the decay of $\mathbb{E}\langle\sigma_u\sigma_v\rangle_J^2$ directly. For magnitudes K_e and conductances $y_e = \sinh^2 K_e$, resistance domination (RD) asks that $\mathbb{E}_\varepsilon\langle\sigma_u\sigma_v\rangle^2 \le C_{\mathrm{eff}}(u,v)/(1+C_{\mathrm{eff}}(u,v))$, and forest domination (FD), the comparison with the arboreal gas restated in Chapter 15, implies it (Chapter 11).

Open problem 16.3 (Comparison on boxes)**.** Prove RD, or FD, for iid fair signs on the free boxes of $\mathbb{Z}^2$ at all activities; or, as a first step, on one infinite family of graphs that is not series–parallel and exhausts a recurrent lattice.

By the implications proved in Chapter 11, the Nash-Williams bound (NWB) on the finite subgraphs of $\mathbb{Z}^2$, and hence RD or FD there, gives $\mathbb{E}\langle\sigma_u\sigma_v\rangle_J^2 \le 8\bar{y}(\beta)/\ln(2r+1)$ with $\bar{y}(\beta) = \mathbb{E}\sinh^2(\beta|J_e|)$ and $r = |u-v|_\infty$, at every β with $\bar{y}(\beta) < \infty$; and for $\pm J$ couplings a power law is proved conditional on FD there. A candidate proof must pass the kill tests of Chapter 15; in particular it must fail for every positive tilt of the sign law, and it cannot rely on series–parallel composition alone. The most specific missing piece is an influence-flow bound beyond series–parallel networks; it implies RD. The bound holds on series–parallel networks and, by a computer-assisted certificate (Appendix B), on the Wheatstone bridge between its terminals; its per-path form fails on the 3×3 grid in exact computations not yet independently reviewed (Chapter 11).

Hierarchical lattices with $\pm J$ couplings. Okuyama and Ohzeki prove absence of order for symmetric $\pm J$ couplings on Migdal–Kadanoff lattices with even branching number when one exact decimation step creates enough vacancies, and classify the Gaussian case completely: there is order at low temperature if and only if a parameter $\Lambda(b,s)$ exceeds 1 (Chapter 12); open is only whether $\Lambda(b,s) = 1$ occurs, the case in which their finite-step criterion does not terminate. In their notation s is the number of parallel branches and b the number of bonds in series in each branch. They write: "Several important cases remain open, including the case of an odd branching number, continuous coupling distributions, and even-branching MK lattices for which the sufficient condition is not satisfied" [201, p. 5]; the continuous case has since been settled for Gaussian couplings [202].

Open problem 16.4 (Okuyama–Ohzeki)**.** Decide spin-glass order for symmetric $\pm J$ couplings on Migdal–Kadanoff lattices with odd branching number, and with even branching number where

the one-step dilution condition fails. "Whether $\Lambda(b, s) = 1$ can occur for some integers $b, s \geq 2$ remains an open question" [202, Remark 5].

An exact cancellation count over more than one decimation step, of the kind used for the self-dual lattice in [201], combined with a computer-assisted comparison with percolation, looks within reach for small lattices. On the ordering side, the proof of order on diamond lattices with $s \geq 4$ branches [219] needs a coupling law with a bounded density, so it does not cover $\pm J$ couplings, for which order at scale factor two is open for every $s \geq 3$.

Long-range chains on the Nishimori line. Okuyama and Ohzeki prove a phase transition for the one-dimensional Gaussian chain with variance decaying like $r^{-\alpha}$ on the Nishimori line for $1 < \alpha < \frac{3}{2}$, by Dyson's hierarchical method, and state: "For $3/2 \leq \alpha \leq 2$, the existence of a phase transition remains an open problem" [203, abstract].

Open problem 16.5 (Okuyama–Ohzeki)**.** Prove or disprove a phase transition on the Nishimori line for the one-dimensional long-range Gaussian spin glass with $\frac{3}{2} \leq \alpha \leq 2$. Off the Nishimori line, prove that weak uniqueness fails at low temperature for some $1 < \alpha \leq 2$: in the words of Gandolfi, Newman and Stein, "there is as yet no proof that weak uniqueness breaks down at low temperatures" [102, p. 4].

The barrier at $\alpha = \frac{3}{2}$ is a per-level error of the order of the square root of the block volume, the Imry–Ma count of Chapter 10. On the Nishimori line magnetization and overlap have the same law (Chapter 14), so a second-moment bound for block free energies sharper than sub-Gaussian concentration might pass it; this is a suggestion, not a computation.

The remaining targets. Targets 5 to 9 are stated where their background is. The rate in the positivity principle (Chapter 12) is a question about mean-field Gibbs measures with no reference to finite dimension, and a positive answer would give order in the perturbed Dyson hierarchical model. Blue percolation in $d \leq 6$ (Chapter 8) would not by itself give order, but it would show that the cavity technique is not confined to high dimension. The chaos questions are listed at the end of Chapter 6; the surface-order variance bound is the conjecture of Arguin, Newman and Stein in Chapter 4, which, together with absence of disorder chaos at some scale $a < 1/d$, would show, for Gaussian couplings in every dimension, that zero-temperature metastates built from coupling-independent boundary conditions are not supported on incongruent ground states. Machta, Newman and Stein ask whether "the presence of 'single' [...] FK percolation would lead to a slower, e.g. power-law, decay of correlation functions even though the Gibbs state is unique" [153, p. 3]; on $\mathbb{Z}^2$, where uniqueness is conjectured at every temperature, the FK clusters percolate at $\beta > \frac{1}{2} \log 17$ for every ± 1 configuration (Chapter 7), which makes it the natural test case. Target 9 restates the ordering problem in the form to which the two routes of Chapters 8 and 15 reduce it.

16.2 A map of conditional reductions

Figure 16.1 collects the implications proved in this review between open statements and conclusions about order. Each dashed box is an open statement (a conjecture or hypothesis); each solid arrow is a proved implication, stated in the chapter indicated. The figure is a guide to what a single new estimate would buy.

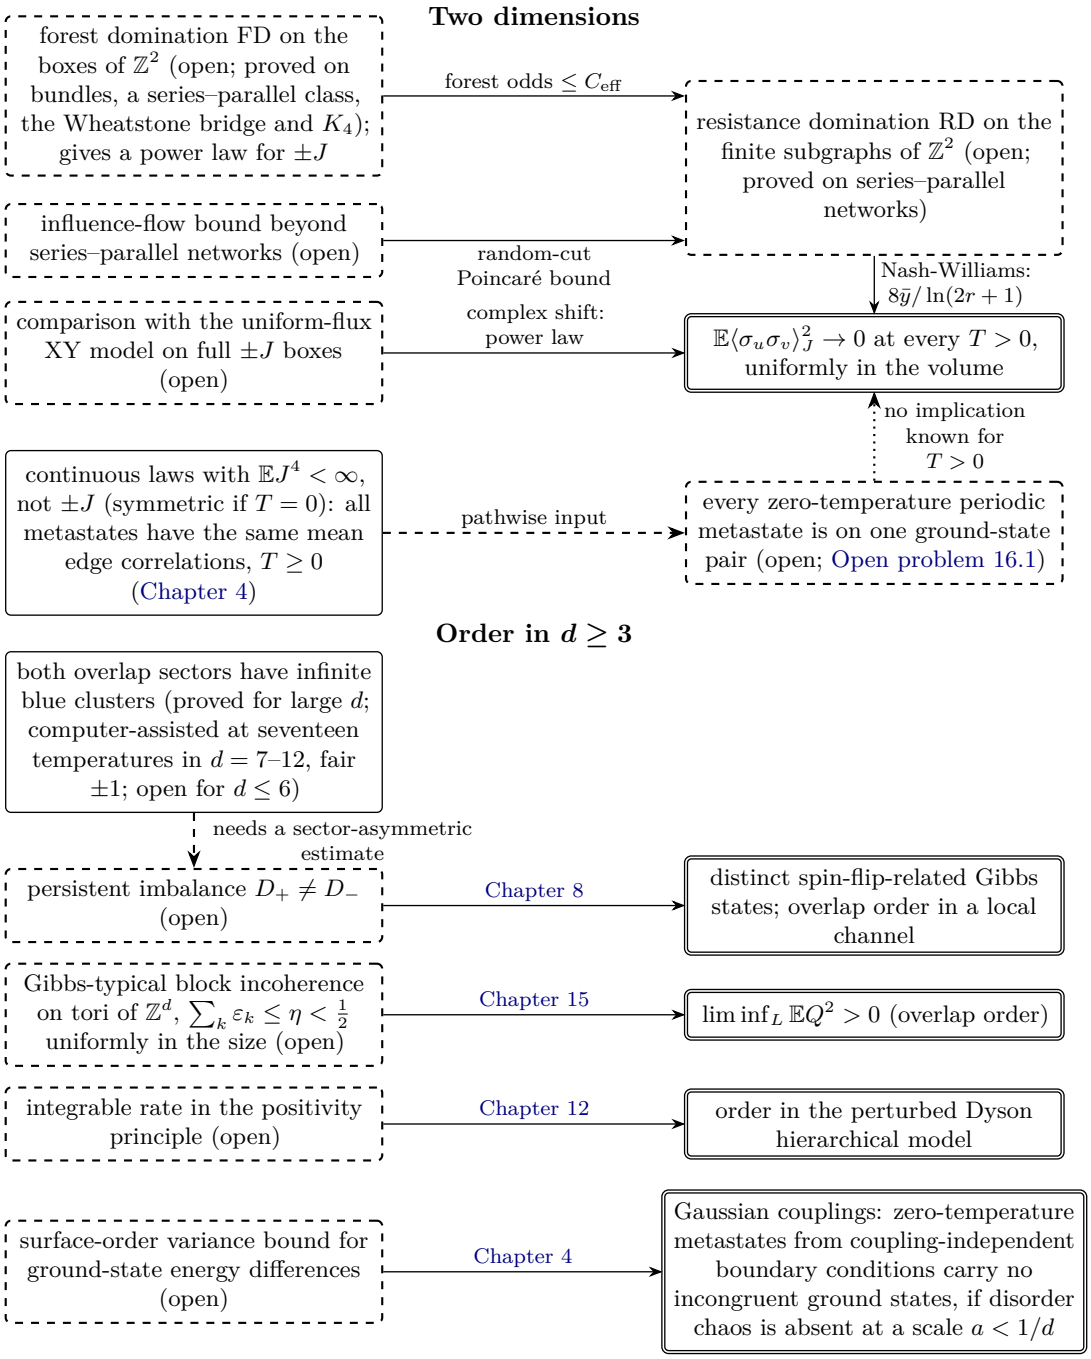


Figure 16.1: Conditional reductions. Dashed boxes are open statements, solid boxes theorems (with the range where they are proved), double boxes conclusions; solid arrows are proved implications, labelled by the mechanism or the chapter. The dashed arrows mark steps that are not proved: the one that symmetric estimates cannot make (Chapter 8) and the pathwise input of Observation 16.2; the dotted arrow marks the absence of any known passage from zero to positive temperature in $d=2$. One solid arrow, the complex shift from the XY comparison, is a proved implication whose details are omitted in Chapter 11.

Two features of the map are worth stating. First, on every route that reaches a conclusion through solid arrows the open statement is a single estimate and the other steps are proved; the zero-temperature route in two dimensions is the exception, since it needs a pathwise input and has no known passage to positive temperature. Second, the open statements on the two sides have the same shape: each asks for control that is uniform across scales and uses the exact average over fair signs, which is the wall of Chapter 15.

16.3 Programmes, frontiers and attack points

16.3.1 Newman, Stein and collaborators: metastates, fluctuations, $d = 2$

The programme of Newman and Stein, with Aizenman, Wehr, Arguin, Damron, Read, Gandolfi and Machta, built the language of the review: metastates and their structure (Chapter 3), the self-averaging no-go for the standard SK picture, space-filling interfaces, the single-wall theorem and half-plane uniqueness at zero temperature (Chapter 5), volume-order fluctuation bounds and the two-dimensional constraints (Chapter 4), the relation between chaos and incongruence, and signed FK representations [104] (Chapter 7). Its own open-problem lists remain the reference. Newman and Stein's list of 1998 asks, among other questions: "1) Prove (or disprove) the widely held belief that the model has a thermodynamic phase transition above some lower critical dimension d_c [...] 3) For a given d and T, find the number of pure state pairs [...] 5) If multiple Gibbs states are shown to exist, determine whether the low-temperature phase persists in the presence of a small external magnetic field" [176]. Arguin, Newman, Stein and Wehr close their two-dimensional theorems with the programme's own aim: "Our hope is that this method of proof using fluctuation bounds could be extended to further restrict the possible structures of the set of Gibbs states in dimension two and higher" [14, p. 8]. Newman and Stein's list of 2004 adds: "Open Question 3. How many ground state pairs is the $T = 0$ periodic boundary condition metastate supported on, as a function of d?" [184, p. 8]. The review of Newman, Read and Stein records the status of the four metastate classes: "all four classes remain open as possibilities" [190, arXiv v1, p. 13], and for chaotic pairs, whether the metastate can be supported on finitely many Gibbs states: "as of now the question remains open" [190, arXiv v1, p. 15]. In the summary of Camia and Stein, which of the four pictures holds "remains a fundamental open problem in the mathematics and physics of the statistical mechanics of disordered systems" [39, p. 4].

Frontier and attack points. In $d = 2$, Open problem 16.1 and the observation following it. In every dimension, the surface-order variance bound for the same quantity that the volume-order bound controls (Chapter 4), which would turn the two-dimensional argument into a variance-against-variance argument in all d. A third, more speculative, point is the exclusion of the class with a trivial metastate and nontrivial Gibbs states, where the weights of an atomless decomposition transform covariantly under coupling changes and an ergodic argument in the space of decompositions might apply.

16.3.2 Nishimori, Ohzeki and Okuyama: gauge identities, hierarchical lattices, long range

This programme proves exact statements where the gauge structure is strongest: the Nishimori identities and correlation inequalities on the Nishimori line, the equality in law of magnetization and overlap there, Griffiths-type inequalities on the line for Gaussian couplings (Morita, Nishimori and Contucci; the $\pm J$ case is Kitatani's) and their extension to other gauge glasses (Chapter 14);

absence of order on Migdal–Kadanoff lattices by exact dilution [201]; the classification of Gaussian Migdal–Kadanoff lattices [202]; and the Nishimori-line transition in one-dimensional long-range chains [203] (Chapter 12). The authors state the limits of their methods plainly: "Whether this approach can be extended to spin-glass models on conventional d-dimensional lattices remains unclear" [202, p. 20], and "An important further problem is whether the present method can be extended to spin-glass models in two dimensions" [201, p. 5]. On reentrance of the ferromagnetic phase boundary below the Nishimori point, Nishimori, Ohzeki and Okuyama relate reentrance to temperature chaos and note that "the present theory does not determine the existence or absence of either phenomenon itself" [198, p. 10] (Chapter 14).

Frontier and attack points. Open problems 16.4 and 16.5; the vertical or reentrant shape of the phase boundary below the Nishimori point on $\mathbb{Z}^2$ (Chapter 14). The scale barrier of Chapter 15 explains why the hierarchical results do not transfer: exact block elimination on $\mathbb{Z}^d$ creates multi-spin terms, and the Migdal–Kadanoff proofs use unbounded degree.

16.3.3 Chatterjee and zero-temperature chaos

Chatterjee's variance identity links superconcentration and edge chaos, and with a volume-order variance bound it shows that the edge overlap is not chaotic. For Gaussian couplings, a gauge identity for the Hermite expansion of ground-state correlations shows that the site overlap is chaotic at zero temperature in every dimension and, with variational inequalities, gives rough droplet boundaries; a two-step gauge flip gives a polynomial lower bound on boundary sensitivity (Chapter 6). Chen, Kim and Sen extend site chaos to general symmetric disorder and every temperature, with free boundary conditions [56], and Basu Roy Chowdhury and Ganguly study the spectral sample [22]. The open questions are Chatterjee's own: "is it true that under the periodic boundary condition on $\{-L,\dots,L\}^d$, for fixed neighboring sites i and j, $\lim_{p\to 0}\lim_{L\to\infty}|\mathbb{E}(\sigma_i\sigma_j\sigma_i'\sigma_j')|<1$?" and "In dimensions $d\geq 3$, the RFIM exhibits long-range order [...]; it remains open whether an analogous statement holds for the EA model" [49, p. 10]; Arguin and Hanson conjecture that in $d=2$ the edge overlap stays close to one [12, Conj. 1.2].

Frontier and attack points. The chaos theorems hold in $d=1$, so by themselves they are not evidence of many ground states (Chapter 6); what carries information about multiplicity is the edge overlap at small perturbation, which is a statement about the low-degree part of the spectral measure of a nearest-neighbour product. A dimension-sensitive power-law lower bound for periodic-minus-antiperiodic ground-state energies in $d\geq 3$ would give a rigorous trace of positive stiffness; nothing in the present methods distinguishes dimensions.

16.3.4 Random-field-style multiscale methods

For the random-field Ising model, multiscale and fluctuation methods settled order in $d\geq 3$ and its absence in $d=2$, with quantitative rates (Chapter 10). What transfers to the EA model is known precisely: the Aizenman–Wehr comparison transfers to bond observables [68], and lattice-animal and chaining arguments transfer [26]; the joint flip of spins and disorder is exactly a gauge transformation, with zero surface gain, and there is no monotone coupling of boundary conditions (Chapter 10). Dario, Harel and Peled conjecture that in $d=2$ the dependence of bond observables on the boundary condition vanishes almost surely, not only on average [68, Conj. 2.1]; even in full, this would leave a flip pair undistinguished (Chapter 4).

Frontier and attack points. An Aizenman–Wehr argument for a disorder that is conjugate to the overlap, such as a random replica coupling $\sum_x \lambda_x\sigma_x\tau_x$, would give two-dimensional

insensitivity of overlap densities; as for the random-field model at zero field, it would say nothing at $\lambda = 0$ without a further continuity argument. More broadly, Chapter 15 records the heuristic that every known multiscale supplier uses a resource that the symmetric iid model lacks.

16.3.5 Blue percolation, imbalance and forest comparison

The author's programme (Chapters 8 and 11 to 13) attacks the two halves of the problem with two devices. On the ordering side, the CMR blue clusters: at most two infinite ones [215], both overlap sectors percolating in high dimension [214], and imbalance implying spin-flip-related Gibbs states; on hierarchical lattices order and imbalance are proved [219]. On the non-ordering side, comparison with the arboreal gas and with electrical networks: the forest bound [216], RD on series–parallel networks, and two-dimensional non-order conditional on NWB on boxes (Chapter 11); and exact uniqueness for diluted couplings just above the percolation threshold [217]. Both sides stop at the scale-uniform estimates of Fig. 16.1.

Frontier and attack points. A sector-asymmetric opening estimate for blue clusters, from a boundary signal or a reference alignment that survives to a fixed scale, is what persistent imbalance needs (Chapter 8); the influence-flow bound on one non-series–parallel infinite family (Open problem 16.3); a model-specific input for blue percolation in $d = 6$, where for unit couplings the Machta–Newman–Stein auxiliary model is subcritical (Chapter 8); and large-β lower bounds for the two-dimensional arboreal gas, which the random-cluster variants of the comparison would need.

16.4 Newman's research problems

Newman's lecture notes of 1997 [171] list research problems that still frame the subject. Table 16.2 quotes those on low-temperature states, with their status as of this review. Here $\beta_c \le \bar{\beta}_c$ are the two critical values of Chapter 1 (a.s. uniqueness for all smaller β, and a.s. non-uniqueness), N is the number of pure states, and κ_J is the metastate.

16.5 Summary

The open problems nearest to present methods are in two dimensions at zero temperature, on the comparison side in two dimensions, and on hierarchical and long-range proxies. Each would sharpen a picture that numerics already support, and none would yet settle the central question. That question asks for control across scales that uses the exact fair average, on the cubic lattice, where the resources that make the proxies tractable are absent. The barriers of Chapter 15 say what such control cannot be; the map of Fig. 16.1 says which single estimate would suffice on each route.

	Problem [171]	Status
4.2 (p. 48)	"Determine (either for all μ or at least for some μ with $\mu(\mathbb{R} \setminus \{0\}) > p_c(2) = 1/2$) whether $\beta_c(2, \mu) = \infty$, or at least whether $\bar{\beta}_c(2, \mu) = \infty$."	Open for $\pm J$ and Gaussian; the "for some μ" form is answered for diluted laws just above p_c (Chapter 13); constraints on metastates (Chapter 4)
4.3 (p. 48)	"Determine whether for some d and some μ, $\bar{\beta}_c(d, \mu) < \infty$, or at least whether $\beta_c(d, \mu) < \infty$."	Open; for symmetric laws order is proved only on hierarchical lattices (Chapter 12)
4.5 (p. 50)	"Determine whether $N(d, \mu, \beta) = \infty$, or at least whether $N(d, \mu, \beta) > 2$, for some d, μ and β."	Open
4.7 (p. 57)	"Prove that $(J, P_{J,L}) \to K^\dagger$ for some probability measure $K^\dagger$ on $(\Omega^\dagger, \mathcal{F}^\dagger)$."	Open, except where the Gibbs state is a.s. unique (Chapters 1 and 9)
4.10 (p. 58)	Almost-sure convergence of the empirical metastate along the full sequence of volumes, without passing to a subsequence.	Open; the mean-field analogue is negative (Chapter 3)
4.14 (p. 64)	"Determine whether for some d, μ and β, the metastate κ_J is (for $\bar{\nu}$-a.e. J) supported on (a) many Γ's, (b) Γ's which decompose into more than two pure states, (c) Γ's which satisfy the other requirements of the nonstandard SK picture."	Open; all four metastate classes remain possible (Chapter 3)

Table 16.2: Newman's research problems on low-temperature states [171], quoted with printed page numbers (μ is Newman's notation for the coupling law, ν in this review; Appendix C translates his metastate notation); Research Problem 4.10 is paraphrased because its notation needs the empirical construction of Chapter 3.

Appendix A

Proofs moved out of the main text

A.1 Proof of the quenched free energy theorem, parts (b) and (c)

Proof of Theorem 1.12(b)–(c). (b) By translation invariance of the law of J, Ψ takes the same value on translates, and Λ_L is a translate of B_{2L+1}. Jensen's inequality for the uniform measure on $\{-1,+1\}^R$ gives $\log Z^{\mathrm{f}}_{R,J} \geq |R|\log 2$ for every finite R, since uniform averages of $\sigma_x\sigma_y$ vanish, and $\log Z^{\mathrm{f}}_{R,J} \leq |R|\log 2 + \beta\sum_{e\in E(R)}|J_e|$ gives $\Psi_R \leq |R|(\log 2 + d\beta\mathbb{E}|J_e|)$. For $n \geq \ell$, B_n contains $\lfloor n/\ell\rfloor^d$ disjoint translates of B_ℓ, so by (a), with the remainder R contributing $\Psi_R \geq 0$, $\Psi_{B_n} \geq \lfloor n/\ell\rfloor^d \Psi_{B_\ell}$. Dividing by n^d, letting $n\to\infty$ and taking the supremum over ℓ gives $\liminf_n n^{-d}\Psi_{B_n} \geq \sup_\ell \ell^{-d}\Psi_{B_\ell} \geq \limsup_n n^{-d}\Psi_{B_n}$, and the supremum is finite by the upper bound. Periodic boundary conditions add $d(2L+1)^{d-1}$ edges, each changing $\log Z$ by at most $\beta|J_e|$; the expected change is $o(|\Lambda_L|)$.

(c) Let $f(J) = \log Z^{\mathrm{f}}_{\Lambda_L,J}$. Replacing one coupling J_e by an independent copy J'_e changes f by at most $\beta|J_e - J'_e|$, because $|\partial f/\partial J_e| = \beta|\langle\sigma_x\sigma_y\rangle| \leq \beta$. The Efron–Stein inequality [82, 226] gives $\operatorname{Var} f \leq \frac{1}{2}\sum_{e\in E(\Lambda_L)}\mathbb{E}[\beta^2(J_e - J'_e)^2] \leq d\beta^2\mathbb{E}J_e^2\,|\Lambda_L|$, so $\operatorname{Var}(f/|\Lambda_L|) \leq d\beta^2\mathbb{E}J_e^2/|\Lambda_L|$, which gives the L^2 statement. For $d \geq 2$, $\sum_L |\Lambda_L|^{-1} < \infty$, and Chebyshev's inequality with the Borel–Cantelli lemma gives $|\Lambda_L|^{-1}(f - \mathbb{E}f) \to 0$ almost surely; combine with (b). For $d = 1$ the free chain is explicit, $\log Z^{\mathrm{f}}_{\Lambda_L,J} = \log 2 + \sum_e \log(2\cosh\beta J_e)$, and the strong law of large numbers applies under $\mathbb{E}|J_e| < \infty$. □

A.2 The bridge from the finite-volume target to the pressure target

We prove Corollary 1.13: for symmetric ν with $\mathbb{E}|J_e| < \infty$, the two-replica limit $p^{(2)}(\beta,\lambda)$ exists, is even and convex in λ, and its right derivative at 0 is at least $\limsup_\ell \mathbb{E}\langle|q_{B_\ell}|\rangle^{\mathrm{f}}_{B_\ell}$.

Proof of Corollary 1.13. The two-replica weight has the form of Lemma 1.11 with the bounded functions $f_e = \sigma_x\sigma_y + \tau_x\tau_y$ multiplying the centred couplings βJ_e, and with the non-random, site-local term $\bar{U} = -\lambda\sum_x \sigma_x\tau_x$, which does not couple different blocks of a partition of Λ. The proof of Theorem 1.12(a)–(b) therefore applies, with $\log Z \geq 2|R|\log 2$ and $\log Z \leq 2|R|\log 2 + |\lambda||R| + 2\beta\sum_{e\in E(R)}|J_e|$: the limit exists and equals $\sup_\ell \ell^{-d}\Psi^{(2)}_{B_\ell}(\beta,\lambda)$. Evenness follows from $\tau \mapsto -\tau$, convexity from that of each $\log Z$. For the inequality, fix ℓ and $\lambda > 0$. Since the free

measure is flip invariant, $\langle e^{\lambda \ell^d q}\rangle = \langle \cosh(\lambda \ell^d q)\rangle \geq \frac{1}{2}\langle e^{\lambda \ell^d |q|}\rangle \geq \frac{1}{2} e^{\lambda \ell^d \langle |q|\rangle}$ by Jensen's inequality, with $q = q_{B_\ell}$. Hence

$$\Psi^{(2)}_{B_\ell}(\beta,\lambda) - \Psi^{(2)}_{B_\ell}(\beta,0) = \mathbb{E}\log\langle e^{\lambda \ell^d q_{B_\ell}}\rangle^{\mathrm{f}}_{B_\ell} \;\geq\; \lambda \ell^d\, \mathbb{E}\langle |q_{B_\ell}|\rangle^{\mathrm{f}}_{B_\ell} - \log 2.$$

By the supremum formula, $p^{(2)}(\beta,\lambda) \geq \ell^{-d}\Psi^{(2)}_{B_\ell}(\beta,\lambda)$, while $\varepsilon_\ell = p^{(2)}(\beta,0) - \ell^{-d}\Psi^{(2)}_{B_\ell}(\beta,0) \to 0$. Therefore $p^{(2)}(\beta,\lambda) - p^{(2)}(\beta,0) \geq \lambda\, \mathbb{E}\langle |q_{B_\ell}|\rangle - \ell^{-d}\log 2 - \varepsilon_\ell$ for every ℓ. Let $\ell \to \infty$ along a subsequence realizing the lim sup, divide by λ and let $\lambda \downarrow 0$. The second inequality in (1.7) is $|q| \geq q^2$. Since Λ_L is a translate of B_{2L+1}, the right-hand side dominates $\limsup_L \mathbb{E}\langle q^2_{\Lambda_L}\rangle$. □

A.3 Proof of Newman's uniqueness theorem

We prove Theorem 1.16 by exploring the ferromagnetic cluster of the boundary while building both bond measures on one probability space. We follow [171, pp. 43–46], using the monotonicity of Lemma 1.15(ii) in place of Newman's explicit treatment of generalized boundary conditions.

Proof of Theorem 1.16. Exploration coupling. Fix an order of E^*_Λ and independent uniform variables $M_1, M_2, \dots$ on $[0,1]$. We reveal edges one at a time. At step j let C_{j-1} be the set of sites joined to $\partial\Lambda$ by edges already revealed to be open in ω^{w} (with $\partial\Lambda \subset C_{j-1}$), let e_j be the first unrevealed edge with an endpoint in C_{j-1}, and set

$$\omega_{e_j} = \mathbf{1}\{M_j \leq \varphi^\eta(\omega_{e_j} = 1 \mid \text{revealed } \omega)\}, \qquad \omega^{\mathrm{w}}_{e_j} = \mathbf{1}\{M_j \leq \varphi^{\mathrm{w}}(\omega_{e_j} = 1 \mid \text{revealed } \omega^{\mathrm{w}})\}.$$

By Lemma 1.15(ii) and induction, $\omega \leq \omega^{\mathrm{w}}$ on all revealed edges. Stop when no unrevealed edge meets $C = C_{j-1}$; then sample the unrevealed edges of ω from $\varphi^\eta(\cdot \mid \text{revealed } \omega)$ and those of ω^{w} from $\varphi^{\mathrm{w}}(\cdot \mid \text{revealed } \omega^{\mathrm{w}})$, independently. Revealing coordinates in an order that depends only on what has been revealed, and sampling each from its conditional law, produces a sample of the joint law [171, pp. 44–45]; so $\omega \sim \varphi^\eta_{\Lambda,J}$ and $\omega^{\mathrm{w}} \sim \varphi^{\mathrm{w}}_{\Lambda,|J|}$. Finally colour the ω-clusters by Lemma 1.14(iii) with independent fair coins, so that $(\sigma,\omega) \sim P^\eta_{\Lambda,J}$.

What the exploration leaves behind. At the stopping time C is the set of sites joined to $\partial\Lambda$ in ω^{w}, the revealed edges are exactly those with an endpoint in C, and every revealed edge from C to $D = \Lambda \setminus C$ is closed in ω^{w}, hence closed in ω. Given the revealed information, the weight in (ii) of Lemma 1.14 factorizes: $k_\Lambda(\omega)$ is the number of clusters of ω inside D plus a revealed constant, and $\omega \in U^\eta_J$ holds if and only if the revealed part is admissible and the restriction of ω to $E(D)$ is free of frustrated cycles. So, conditionally, $\omega|_{E(D)}$ has the law of the free random-cluster measure of J on D, and, by Lemma 1.14(iii) applied on D, σ_D has the law $\mu^{\beta,\mathrm{f}}_{D,J}$. The law of D is that of Λ minus the wired ferromagnetic cluster of $\partial\Lambda$, and does not depend on η. Hence

$$\langle \sigma_A\rangle^\eta_{\Lambda,J} = \sum_{D_0 \supseteq A} \varphi^{\mathrm{w}}_{\Lambda,|J|}(\Lambda \setminus C = D_0)\, \langle \sigma_A\rangle^{\mathrm{f}}_{D_0,J} \;+\; \mathbb{E}\big[\sigma_A\,;\, A \cap C \neq \emptyset\big],$$

and the first term is the same for every η. Since $|\sigma_A| = 1$ and $\mathbb{P}(A \cap C \neq \emptyset) \leq \sum_{x\in A} \varphi^{\mathrm{w}}_{\Lambda,|J|}(x \leftrightarrow \partial\Lambda)$, this proves (1.9).

Uniqueness. If $\mu, \mu' \in \mathcal{G}_J(\beta)$, the DLR equations give $\mu(\sigma_A) = \int \langle \sigma_A\rangle^\eta_{\Lambda,J}\, \mu(\mathrm{d}\eta)$ and likewise for μ', so $|\mu(\sigma_A) - \mu'(\sigma_A)|$ is bounded by the right side of (1.9) for every $\Lambda \supset A$. The probabilities $\varphi^{\mathrm{w}}_{\Lambda,|J|}(x \leftrightarrow \partial\Lambda)$ decrease as $\Lambda \uparrow \mathbb{Z}^d$, and for the ferromagnet their limit is the plus-state magnetization at x [171, (3.17), Proposition 3.8], which vanishes when the ferromagnet has a unique Gibbs state. Hence $\mu(\sigma_A) = \mu'(\sigma_A)$ for all finite A, that is, $\mu = \mu'$. □

A.4 The rung computation

Proof of Proposition 1.19. Squared correlations are gauge invariant, so we may gauge the edges 34, 45, 56, 61, 25, which form a spanning tree, to positive signs. What remains are two independent fair signs x on 12 and y on 23 (for $\lambda = 0$ give the absent rung a dummy sign; nothing depends on it). The high-temperature expansion of Z runs over the even subgraphs, which are the empty set, the two squares and the outer six-cycle; that of $Z\langle\sigma_1\sigma_3\rangle$ runs over the four paths from 1 to 3. With the common factor removed,

$$Z \propto 1 + xy\,t^6 + (x+y)\,t^3u, \qquad Z\langle\sigma_1\sigma_3\rangle \propto xy\,t^2 + t^4 + (x+y)\,t^3u.$$

The four sign pairs give $(A \pm Bu)/(C \pm Bu)$ for $x = y = \pm 1$ and $(t^4 - t^2)/(1 - t^6)$ for $x = -y$, which is (1.10). At $t^2 = \frac{1}{2}$ the three ratios are $\frac{10}{13}, \frac{2}{5}, -\frac{2}{7}$ with the rung and $\frac{2}{3}, \frac{2}{3}, -\frac{2}{7}$ without, which gives the stated values. Differentiating (1.10),

$$\partial_u Q = \frac{B^2u\,(C-A)\,\{C^2(C-3A) + (3C-A)B^2u^2\}}{(C^2 - B^2u^2)^3},$$

with $C - A = (1-t^2)(1-t^4) > 0$ and $C - 3A = (1+t^2)(t^4 - 4t^2 + 1)$, which is negative when $t^2 > 2 - \sqrt{3}$; then $\partial_u Q < 0$ for small $u > 0$. □

A.5 The Aizenman–Wehr construction

We prove Theorem 3.5 and then supply the measure-theoretic details used in the proof.

Proof of Theorem 3.5. Tightness. The first marginal is the fixed law ν on the Polish space $\mathbb{R}^{E_d}$, and the second coordinate takes values in the compact space $M_1(\Sigma)$. A family of probability measures on a product whose marginals are tight is tight, so by Prokhorov's theorem some deterministic subsequence converges to a limit $K^\dagger$, whose J-marginal is ν. Disintegrating, $K^\dagger(dJ, d\Gamma) = \nu(dJ)\,\kappa_J(d\Gamma)$ with $J \mapsto \kappa_J$ measurable.

The DLR equations survive the limit. Fix finite sets $\Lambda \subset \Lambda'$ with $\Lambda \cup \partial\Lambda \subset \Lambda'$, where $\partial\Lambda$ denotes the outer vertex boundary. Consider the finitely many equations

$$\Gamma(\sigma_{\Lambda'} = s) = \Gamma(\sigma_{\Lambda'\setminus\Lambda} = s|_{\Lambda'\setminus\Lambda})\;\mu_{\Lambda,J}^{s|_{\partial\Lambda}}(\sigma_\Lambda = s|_\Lambda), \qquad s \in \{\pm 1\}^{\Lambda'}, \tag{A.1}$$

where $\mu^\eta_{\Lambda,J}$ is the Gibbs measure on Λ with boundary spins η. Both sides are continuous functions of Γ and of the finitely many couplings on $E(\Lambda) \cup \partial_e\Lambda$, so the set $F_{\Lambda,\Lambda'}$ of pairs (J,Γ) satisfying (A.1) is closed. Once Λ' lies well inside Λ_L, the torus measure $\mu^{\mathrm{per}}_{L,J}$ satisfies (A.1) exactly, so $(J, \mu^{\mathrm{per}}_{L_n,J}) \in F_{\Lambda,\Lambda'}$ for all large n, and the portmanteau theorem gives $K^\dagger(F_{\Lambda,\Lambda'}) = 1$. By the martingale convergence theorem the set of pairs (J,Γ) with $\Gamma \in \mathcal{G}_J$ is the intersection of the countably many $F_{\Lambda,\Lambda'}$ (see below), so $\int \nu(dJ)\,\kappa_J(\mathcal{G}_J) = 1$. □

Spaces. $\Omega = \mathbb{R}^{E_d}$ with the product topology is Polish. Enumerate the cylinder events of Σ as $C_1, C_2, \ldots$; then $d(\Gamma,\Gamma') = \sum_k 2^{-k}|\Gamma(C_k) - \Gamma'(C_k)|$ metrizes the topology of $M_1(\Sigma)$, which is compact. Weak convergence of probability measures on $\Omega \times M_1(\Sigma)$ is equivalent to convergence of the finite-dimensional distributions of the coordinates $(J_e)_e$ and $(\Gamma(C_k))_k$, which is the formulation of [171, App. B], because the map $(J,\Gamma) \mapsto ((J_e)_e, (\Gamma(C_k))_k)$ is a homeomorphism onto its (closed) image in $\mathbb{R}^{E_d} \times [0,1]^{\mathbb{N}}$.

The finite-volume objects. For fixed L the map $J \mapsto \mu^{\mathrm{per}}_{L,J}$ is continuous: it depends on finitely many couplings, through ratios of finite sums of exponentials. Hence the law of $(J, \mu^{\mathrm{per}}_{L,J})$ is a Borel probability measure on $\Omega \times M_1(\Sigma)$.

The set of Gibbs pairs is closed. Let $\mathfrak{G} = \{(J,\Gamma) : \Gamma \in \mathcal{G}_J\}$. Every Gibbs state satisfies the finite equations (A.1): conditioning the DLR identity on $\sigma_{\Lambda'\setminus\Lambda}$ only, the conditional law of σ_Λ is still $\mu^{\sigma|_{\partial\Lambda}}_{\Lambda,J}$, because the latter depends only on $\sigma|_{\partial\Lambda}$ and $\partial\Lambda \subset \Lambda' \setminus \Lambda$. Conversely, if (J,Γ) satisfies (A.1) for all pairs $\Lambda \subset \Lambda'$, then for each Λ and $t \in \{\pm 1\}^\Lambda$ we have $\Gamma(\sigma_\Lambda = t \mid \mathcal{F}_{\Lambda'\setminus\Lambda}) = \mu^{\sigma|_{\partial\Lambda}}_{\Lambda,J}(t)$ Γ-a.s. for every Λ', where $\mathcal{F}_A$ is the σ-field of the spins in A. Letting Λ' increase to $\mathbb{Z}^d$ along a sequence, Lévy's upward theorem gives $\Gamma(\sigma_\Lambda = t \mid \mathcal{F}_{\Lambda^c}) = \mu^{\sigma|_{\partial\Lambda}}_{\Lambda,J}(t)$ a.s., which is the DLR condition. So $\mathfrak{G} = \bigcap_{\Lambda\subset\Lambda'} F_{\Lambda,\Lambda'}$ is a countable intersection of closed sets, hence closed.

Disintegration. Since Ω and $M_1(\Sigma)$ are Polish, the limit $K^\dagger$ has a regular conditional distribution given the first coordinate: a Borel map $J \mapsto \kappa_J \in M_1(M_1(\Sigma))$ with $K^\dagger(A) = \int \nu(dJ)\,\kappa_J(A_J)$ for Borel A, where $A_J = \{\Gamma : (J,\Gamma) \in A\}$. With $A = \mathfrak{G}$, $J \mapsto \kappa_J(\mathcal{G}_J)$ is measurable and $\int \nu(dJ)\,\kappa_J(\mathcal{G}_J) = K^\dagger(\mathfrak{G}) = 1$, so $\kappa_J(\mathcal{G}_J) = 1$ for ν-a.e. J.

A.6 Ergodicity of iid couplings

Proof of Lemma 3.10. It suffices to treat $f = \mathbf{1}_A$. Given $\varepsilon > 0$ choose a cylinder event C with $\nu(A \triangle C) < \varepsilon$. For $|a|$ large, C and $\tau_a C$ depend on disjoint sets of couplings, so $\nu(C \cap \tau_a C) = \nu(C)^2$. Since $\tau_a A = A$ up to a null set, $\nu(A) = \nu(A \cap \tau_a A)$ is within 2ε of $\nu(C)^2$, hence within 4ε of $\nu(A)^2$. So $\nu(A) \in \{0,1\}$. □

A.7 Stochastic stability

Proof of Proposition 3.26. At fixed J, $\frac{d}{d\beta}\langle f\rangle = \sum_{l\le n}\langle f(-H(\sigma^l))\rangle - n\langle f(-H(\sigma^{n+1}))\rangle$. Average over J and apply Lemma 3.25 to each term, with the roles of the replicas permuted and, for the last term, f regarded as a function of $n+1$ replicas. Collecting terms (the diagonal $l = l'$ contributions cancel) gives $\beta N \Delta_n(f)$. Integrating in β and using $|\mathbb{E}\langle f\rangle| \le \|f\|_\infty$ gives the bound. □

A.8 Concentration of the energy

We prove Lemma 3.28. Write $\langle\cdot\rangle_\beta$ for the Gibbs expectation at inverse temperature β, $\varphi(\beta) = N^{-1}\log Z_\beta$ with $Z_\beta = \sum_\sigma e^{-\beta H(\sigma)}$, and $\psi = \mathbb{E}\varphi$. Then $\varphi'(\beta) = N^{-1}\langle -H\rangle_\beta$ and $\varphi''(\beta) = N^{-1}(\langle H^2\rangle_\beta - \langle H\rangle_\beta^2) \ge 0$, so φ and ψ are convex, and $\psi' = \mathbb{E}\varphi'$. By the triangle inequality,

$$\frac{1}{N}\mathbb{E}\langle|H - \mathbb{E}\langle H\rangle|\rangle \le \frac{1}{N}\mathbb{E}\langle|H - \langle H\rangle|\rangle + \mathbb{E}|\varphi'(\beta) - \psi'(\beta)|.$$

Thermal part. By Cauchy–Schwarz in (β, Gibbs, J),

$$\int_{\beta_1}^{\beta_2} \frac{1}{N}\mathbb{E}\langle|H - \langle H\rangle|\rangle\,d\beta \le (\beta_2-\beta_1)^{1/2}\Big(\frac{1}{N^2}\int_{\beta_1}^{\beta_2}\mathbb{E}\big[\langle H^2\rangle - \langle H\rangle^2\big]\,d\beta\Big)^{1/2},$$

and $\int_{\beta_1}^{\beta_2}\mathbb{E}\big[\langle H^2\rangle - \langle H\rangle^2\big]\,d\beta = \mathbb{E}\langle -H\rangle_{\beta_2} - \mathbb{E}\langle -H\rangle_{\beta_1} \le 2\beta_2 N$ by Lemma 3.25, so the bracket is at most $2\beta_2/N$. This gives the first term of the lemma.

Disorder part. The Gaussian Poincaré inequality $\operatorname{Var} F(J) \le \mathbb{E}|\nabla F(J)|^2$ applied to $F = \log Z_\beta$, whose partial derivatives are $\beta\langle\sigma_e\rangle_\beta$, gives $\operatorname{Var}\varphi(\beta) \le \beta^2/N$, hence $\mathbb{E}|\varphi(\beta)-\psi(\beta)| \le |\beta|/\sqrt{N}$. For convex differentiable functions and $y>0$,

$$\frac{\varphi(\beta)-\varphi(\beta-y)}{y} \le \varphi'(\beta) \le \frac{\varphi(\beta+y)-\varphi(\beta)}{y},$$

and the same for ψ, with $\psi'(\beta-y) \le \frac{\psi(\beta)-\psi(\beta-y)}{y}$ and $\frac{\psi(\beta+y)-\psi(\beta)}{y} \le \psi'(\beta+y)$. Comparing the two chains,

$$|\varphi'(\beta)-\psi'(\beta)| \le \psi'(\beta+y)-\psi'(\beta-y) + \frac{\delta(\beta+y)+2\delta(\beta)+\delta(\beta-y)}{y}, \qquad \delta = |\varphi-\psi|.$$

Integrate over $[\beta_1,\beta_2]$ and take expectations. The first term integrates to $[\psi(\beta_2+y)-\psi(\beta_2-y)]-[\psi(\beta_1+y)-\psi(\beta_1-y)] \le 4y\sup_{|\beta|\le\beta_2+y}|\psi'| \le 8y(\beta_2+y)$, using $|\psi'(\beta)| = |\beta|\,(1-\mathbb{E}\langle R_{12}\rangle_\beta) \le 2|\beta|$ from Lemma 3.25. The second contributes at most $4(\beta_2-\beta_1)(\beta_2+y)/(y\sqrt{N})$. Choosing $y = N^{-1/4} \le 1$,

$$\int_{\beta_1}^{\beta_2} \mathbb{E}|\varphi'-\psi'|\,d\beta \le 8(\beta_2+1)(1+\beta_2-\beta_1)N^{-1/4},$$

which is the second term of the lemma with $C(\beta_2) = 8(\beta_2+1)$.

A.9 The Gaussian lower bound in two dimensions

This section outlines the proof of Proposition 4.8, following [14, Sect. 6.2]. The probabilistic input is a lower bound for the moment generating function of a martingale array, a weak form of the martingale central limit theorem.

Proposition A.1 ([8, Prop. A.2.1]; [14, Prop. 6.3])**.** *For each n let $(X_{n,k})_{k\le n}$ be a martingale with respect to a filtration $(\mathcal{F}_{n,k})_{k\le n}$, with increments $\Delta_{n,k} = X_{n,k}-X_{n,k-1}$ and $X_{n,0}=0$. Suppose that for some $\sigma^2>0$,*

(i) $\mathbb{P}\big(\sum_{k\le n}\mathbb{E}[\Delta_{n,k}^2 \mid \mathcal{F}_{n,k-1}] \le \sigma^2-\delta\big) \to 0$ *for every $\delta>0$, and*

(ii) $\sum_{k\le n}\mathbb{E}[\Delta_{n,k}^2 \mathbf{1}\{|\Delta_{n,k}|>\delta\}] \to 0$ *for every $\delta>0$.*

Then $\liminf_n \mathbb{E}\exp(tX_{n,n}) \ge e^{t^2\sigma^2/2}$ *for every $t\in\mathbb{R}$.*

The martingale. Order the edges of $\mathbb{Z}^2$ lexicographically by their lower-left endpoint, enumerate $E(\Lambda) = \{e_1,\dots,e_n\}$ in this order, and let $Y_{\Lambda,k} = M(F_\Lambda \mid J_{e_1},\dots,J_{e_k})$. Then $X_{n,k} = (Y_{\Lambda,k}-Y_{\Lambda,0})/|E(\Lambda)|^{1/2}$ is a martingale with $X_{n,n} = (G_\Lambda - \mathbb{E}G_\Lambda)/|E(\Lambda)|^{1/2}$. By Lemma 4.6 the increment $\Delta Y_{\Lambda,k}$ is an integral over the revealed coupling of the function $D_k(J) = -\beta[\kappa_J(\Gamma(\sigma_{e_k})) - \kappa'_J(\Gamma'(\sigma_{e_k}))]$, averaged over the unrevealed couplings, so $|\Delta Y_{\Lambda,k}| \le 2\beta(|J_{e_k}|+\mathbb{E}|J|)$.

Condition (ii). From the last bound, $\mathbb{E}[\Delta Y_{\Lambda,k}^4]$ is bounded uniformly in k and Λ when $\mathbb{E}J^4<\infty$. Chebyshev's and the Cauchy–Schwarz inequalities then bound the sum in (ii) by $C\max_k \mathbb{E}[\Delta Y_{\Lambda,k}^4]^{1/2}\,(\mathbb{E}J^2)^{1/2}/(\delta|E(\Lambda)|^{1/2}) \to 0$.

Condition (i). Replace $Y_{\Lambda,k}$ by $\tilde Y_k = M(F_\Lambda \mid J_e,\ e\preceq e_k)$, which conditions on all edges of $\mathbb{Z}^2$ preceding e_k rather than only those in Λ; its increments $\Delta\tilde Y_k$ do not depend on Λ, and the

functions $f_k(J) = \mathbb{E}[\Delta\tilde{Y}_k^2 \mid J_e,\ e \prec e_k]$ are translates of one another (of two functions, one for horizontal and one for vertical edges). By von Neumann's L^2 ergodic theorem, $|E(\Lambda)|^{-1}\sum_k f_k$ converges in L^2 to a constant σ^2. The difference between the conditional variances of $\Delta Y_{\Lambda,k}$ and of $\Delta\tilde{Y}_k$ tends to zero in L^1 uniformly over edges at distance at least $|\Lambda|^{1/4}$ from the boundary, by the L^2 martingale convergence theorem applied in a box of area $|\Lambda|^{1/2}$ around e_k; the remaining edges are a vanishing fraction. This gives (i). Finally $\sigma^2 > 0$, because $\sigma^2 = \lim \mathrm{Var}(G_\Lambda)/|E(\Lambda)|$ and the proof of Theorem 4.4 bounds $\mathrm{Var}(G_\Lambda)/|\Lambda|$ from below by a positive constant for all large Λ under asymmetry on average.

Proposition A.1 now gives the claim with $c = \sigma^2|E(\Lambda)|/(2|\Lambda|)$ in the limit, that is, $c = \sigma^2$ for $d = 2$, since $|E(\Lambda)|/|\Lambda| \to 2$.

A.10 Energy bookkeeping for the dynamics of Chapter 5

This section proves the balance identity (5.5) used in the proof of Theorem 5.17. Let P be the joint law of the couplings J (iid with $\mathbb{E}|J_e| < \infty$), the initial configuration σ^0 (iid symmetric, independent of J) and the clocks and coins of the graphical construction. P is invariant under the translations of $\mathbb{Z}^d$, and the dynamics commutes with them.

Let $N_x(t)$ be the number of clock rings at x in $(0,t]$, a Poisson variable of mean t independent of J. The number of flips at x is at most $N_x(t)$, and each flip changes ε_x and the ε_y, $y \sim x$, by at most $\sum_{z\sim x}|J_{xz}|$. Hence every sum below is dominated by $N_x(t)\sum_{z\sim x}|J_{xz}|$ for some x in the unit ball around the origin, whose expectation is at most $2d\,t\,\mathbb{E}|J_e|$. So all expectations are finite and Fubini's theorem applies.

Only finitely many flips occur in $(0,t]$ in the unit ball around the origin, almost surely, so $\varepsilon_0(t) - \varepsilon_0(0)$ is the sum of the jumps of ε_0. A flip at 0 contributes $\sum_{y\sim 0} J_{0y}\sigma_0\sigma_y = \Delta H_0/2$ (values just before the flip), and a flip at a neighbour y contributes $J_{0y}\sigma_0\sigma_y$. Write $\Phi(x,y) := \sum_{\text{flips at } x \text{ in } (0,t]} J_{xy}\sigma_x\sigma_y$ for neighbours x, y. Then

$$\mathbb{E}\varepsilon_0(t) - \mathbb{E}\varepsilon_0(0) = \mathbb{E}\Big[\sum_{\text{flips at } 0} \tfrac{1}{2}\Delta H_0\Big] + \sum_{y\sim 0}\mathbb{E}\,\Phi(y,0).$$

Translating by $-y$ and using translation invariance, $\mathbb{E}\,\Phi(y,0) = \mathbb{E}\,\Phi(0,-y)$. As y runs over the neighbours of 0, so does $-y$, hence

$$\sum_{y\sim 0}\mathbb{E}\,\Phi(y,0) = \mathbb{E}\sum_{\text{flips at } 0}\ \sum_{z\sim 0} J_{0z}\sigma_0\sigma_z = \mathbb{E}\sum_{\text{flips at } 0}\tfrac{1}{2}\Delta H_0.$$

Adding the two terms gives (5.5).

A.11 Hermite facts used in Chapter 6

The normalised Hermite polynomials satisfy $h_k' = \sqrt{k}\,h_{k-1}$ and $xh_k(x) = \sqrt{k+1}\,h_{k+1}(x) + \sqrt{k}\,h_{k-1}(x)$, so $xh_k - h_k' = \sqrt{k+1}\,h_{k+1}$. For an absolutely continuous g on $\mathbb{R}$ with g and g' of polynomial growth, one-dimensional Gaussian integration by parts gives $\mathbb{E}[g'(x)h_k(x)] = \mathbb{E}[g(x)(xh_k(x) - h_k'(x))] = \sqrt{k+1}\,\mathbb{E}[g(x)h_{k+1}(x)]$. Applying this in the coordinate J_e and Fubini's theorem in the others gives, for every Lipschitz f on $\mathbb{R}^E$,

$$\widehat{\partial_e f}(m) = \sqrt{m_e+1}\,\hat{f}(m+\mathbf{1}_e) \qquad (m \in \mathbb{N}^E). \tag{A.2}$$

This is the identity used in the proof of Theorem 6.3. In that proof, F_β is smooth with bounded derivatives for $\beta < \infty$, and $F_\infty = \min_\sigma H_J(\sigma)$ is Lipschitz with constant $\sqrt{|E|}$, with gradient $-(\sigma_e)_e$ wherever the ground-state pair is unique, which is almost everywhere.

The semigroup relations (6.1) are standard. For the OU perturbation, Mehler's formula gives $\mathbb{E}[h_k(e^{-t}x + \sqrt{1-e^{-2t}}\,y)] = e^{-kt}h_k(x)$ for y standard Gaussian [48, Chapters 2 and 6], and the multivariate statement follows by independence of coordinates. For resampling, the coordinate J_e is kept with probability $1-p = e^{-t}$ and otherwise replaced by an independent copy, on which h_{n_e} has mean zero unless $n_e = 0$. Both relations extend from h_n to L^2 by linearity and continuity, which gives (6.2).

The sums in the proof of Theorem 6.3 have nonnegative terms after the substitution $n = m+\mathbf{1}_e$, so the interchange of the sum over n with the integral over t is justified by Tonelli's theorem, and $\int_0^\infty |n|e^{-|n|t}\,dt = \mathbf{1}\{n \neq 0\}$.

A.12 Proof of Theorem 6.7 for a symmetric law with 0 in its support

We follow [12, Sections 1.2 and 2.1]. For $e = xy \in \Lambda^*$ let A_e be the event $J_{xy} > \sum_{z\sim x, z\neq y} |J_{xz}|$. It is open, depends on $2d$ couplings, and has the same positive probability $\mathbb{P}(A)$ for every e, because the law is symmetric with 0 in its support. On A_e the edge is super-satisfied, so $\sigma_e(J) = +1$ (Lemma 5.7). Off A_e, conditional independence of the two perturbations gives

$$\mathbb{E}\big[\sigma_e(J+t\varepsilon)\sigma_e(J+t\varepsilon');\, J \notin A_e\big] = \mathbb{E}\big[\mathbb{E}[\sigma_e(J+t\varepsilon) \mid J]^2;\, J \notin A_e\big] \geq 0.$$

On A_e the product is $+1$ whenever $J + t\varepsilon$ and $J + t\varepsilon'$ also lie in A_e, so the expectation of the summand is at least $2\mathbb{P}(J,\, J+t\varepsilon,\, J+t\varepsilon' \in A_e) - \mathbb{P}(A)$. Since A_e is open and depends on finitely many coordinates, $\mathbb{P}(J+t\varepsilon \in A_e \mid J \in A_e) \to 1$ as $t \to 0$, at a rate that depends only on the laws of $2d$ coordinates. Choose t_0 so that each of the two exceptional probabilities is below $\frac{\delta}{4}\mathbb{P}(A)$; then each summand exceeds $(1-\delta)\mathbb{P}(A)$.

A.13 Proof of Theorem 6.21

We follow [49, Section 2.7]. Work on the discrete torus V' of side $100L$ centred at 0, let $U = \{-L, \dots, L\}^d$, let ∂U be its edge boundary, and let $k := 4Le_1$. Split $\partial U = S_1 \cup S_2$, where S_1 consists of the edges whose endpoint in U has nonnegative second coordinate. There are paths P_1, P_2 from 0 to k, each of length at most $16L$, such that for every vertex u of P_r the box $W_u := u + V$ contains no edge of S_r (Fig. 6.4; the explicit choice is in Section A.14).

Let J' be J with the signs of the couplings on S_1 reversed, J'' the same on all of ∂U, and $\sigma, \sigma', \sigma''$ the corresponding ground states on the torus. For an edge $f = \{u, v\}$ of P_1, let B_u be the vertex boundary of W_u and let Q_f be the event that the ground state τ of W_u, with the couplings of J inside it, has a value of $\tau_u\tau_v$ that depends on the boundary condition on B_u. The restrictions of σ and σ' to W_u are ground states of W_u for the same couplings, because $J = J'$ there, with boundary conditions $\sigma|_{B_u}$ and $\sigma'|_{B_u}$. Hence $\{\sigma_u\sigma_v \neq \sigma'_u\sigma'_v\} \subseteq Q_f$. By translation invariance and the lattice symmetries, $\mathbb{P}(Q_f) = \mathbb{P}(\mathcal{E})$. Since $\sigma_0\sigma_k = \prod_{f\in P_1}\sigma_f$,

$$\mathbb{P}(\sigma_0\sigma_k \neq \sigma'_0\sigma'_k) \leq \sum_{f\in P_1} \mathbb{P}(Q_f) \leq 16L\,\mathbb{P}(\mathcal{E}).$$

The same argument with P_2, S_2 and the pair (J', J''), which differ only on S_2, gives $\mathbb{P}(\sigma'_0\sigma'_k \neq \sigma''_0\sigma''_k) \leq 16L\,\mathbb{P}(\mathcal{E})$; here we use that J' has the law of J. Finally $J'' = a \circ J$ with $a = 1 - 2\mathbf{1}_U$, so σ'' is σ reversed on U and $\sigma''_0\sigma''_k = -\sigma_0\sigma_k$ always. Therefore $1 \leq 32L\,\mathbb{P}(\mathcal{E})$.

A.14 The paths in the proof of Theorem 6.21

Let $U = \{-L, \dots, L\}^d$, let S_1 be the set of edges of ∂U whose endpoint in U has second coordinate at least 0, and $S_2 := \partial U \setminus S_1$. An edge of S_1 in a direction $\pm e_r$ with $r \neq 2$ has both endpoints with second coordinate at least 0. An edge of S_1 in direction $\pm e_2$ leaves U through its top face $\{x_2 = L\}$, since an edge from a point with $x_2 \geq 0$ downwards stays in U. So every endpoint of an edge of S_1 has second coordinate at least 0. Every edge of ∂U has one endpoint outside U, and all its endpoints have ℓ^∞ norm at most $L + 1$.

The theorem's box V has side length L and is centred at the origin, so $W_u = u + V \subseteq u + [-L/2, L/2]^d$. Let P_1 run from 0 straight down to $-2Le_2$, then to $-2Le_2 + 4Le_1$, then straight up to $k = 4Le_1$, for a length of $8L \leq 16L$. Take a vertex u of P_1.

- If $u = -ye_2$ with $0 \leq y \leq L/2$, then $W_u \subseteq [-L, L]^d$, which is U, so W_u contains no edge of ∂U.
- If $u = -ye_2$ with $y > L/2$, or u lies on the bottom segment $x_2 = -2L$, then every point of W_u has $x_2 < 0$, so W_u contains no edge of S_1.
- If u lies on the last segment, then every point of W_u has $x_1 \geq 4L - L/2 > L + 1$, so W_u meets no endpoint of an edge of ∂U.

The path P_2 is the mirror image of P_1 under $x_2 \mapsto -x_2$. The same three cases, with the bottom face replaced by the top face, show that no W_u along P_2 contains an edge of S_2. All boxes W_u lie within ℓ^∞ distance $5L$ of the origin, so they do not wrap around the torus of side $100L$.

A.15 Annealed insertion tolerance in periodic limits (Chapter 7)

This section supplies the finite-volume details behind part (i) of the uniqueness theorem for the infinite FK cluster in Chapter 7. Let $\mathbb{P}_L$ be the joint law (7.5) of fair ± 1 couplings and frustrated FK bonds on the torus $\mathbb{T}_L^d$, and let $\mathbb{P}$ be a weak limit along $L_n \to \infty$.

The conditional bound passes to the limit. Fix an edge e and a nonnegative function f of finitely many bond variables other than e. For L large the support of f and e sit in $\mathbb{T}_L^d$, and the insertion lemma gives, by conditioning on $\omega_{E\setminus e}$,

$$\mathbb{E}_L[\mathbf{1}\{\omega_e = 1\}f] \;\geq\; \frac{p}{2}\,\mathbb{E}_L[f].$$

Both sides are expectations of local functions, so the inequality holds for $\mathbb{P}$. The class of bounded measurable $f \geq 0$ of $\omega_{E_d\setminus e}$ satisfying it is closed under bounded monotone limits and contains the local functions, so it contains all of them; taking $f = \mathbf{1}_B$ for B in the sigma-field of $\omega_{E_d\setminus e}$ gives $\mathbb{P}(\omega_e = 1 \mid \omega_{E_d\setminus e}) \geq p/2$ almost surely.

Domination in finite volume. Enumerate the edges of $\mathbb{T}_L^d$ as $e_1, \dots, e_N$ and let $U_1, \dots, U_N$ be iid uniform. Sample ω_{e_k} given $\omega_{e_1}, \dots, \omega_{e_{k-1}}$ as $\mathbf{1}\{U_k \leq \pi_k\}$, where π_k is the conditional probability of $\{\omega_{e_k} = 1\}$ given the earlier bonds. Averaging the bound of the insertion lemma over the later bonds shows $\pi_k \geq p/2$, so $\omega_{e_k} \geq \mathbf{1}\{U_k \leq p/2\}$ for every k, and the right side is Bernoulli$(p/2)$ percolation. Stochastic domination is preserved by weak limits, because it is equivalent to $\mathbb{E}f(\omega) \geq \mathbb{E}_{p/2}f$ for all increasing local f, which are continuous. This proves part (i).

A.16 The one-attempt exploration (Chapter 8)

This section proves Proposition 8.9, the embedding step in the proof of the high-dimensional blue-percolation theorem of Chapter 8.

Proof of Proposition 8.9. Keep a first-in first-out queue of reached vertices and a set of attempted vertices, both initially $\{o\}$, and reveal edge variables only when tested. When the head u of the queue has a never-attempted neighbour v, declare v attempted and test uv: with π the conditional probability of $\{\omega_{uv} = 1\}$ given everything revealed so far, reveal ω_{uv} and set $X_v = \mathbf{1}\{\omega_{uv} = 1\}\mathbf{1}\{U \le p/\pi\}$ with a fresh uniform U ($X_v = 0$ if $\pi = 0$, which forces $p = 0$); add v to the queue exactly when $X_v = 1$. Both endpoints of every earlier tested edge were attempted (the first was reached, the second was attempted when tested), so no earlier tested edge contains v: the revealed history is an off-star history at v, and the auxiliary uniforms carry no information about unrevealed edges. Hence $\pi \ge p$, and $X_v = 1$ with conditional probability exactly p whatever the transcript. The next target depends only on the labels and the queue rule, so the labels have the law of an iid Bernoulli(p) field explored from o (unattempted vertices receive independent labels). Degrees are finite, so every reached vertex is processed and the reached set is $C_X(o)$; each reached vertex entered through an open edge, so $C_X(o) \subseteq C_\omega(o)$. Restricting all tests to an induced subgraph changes nothing. □

A.17 Almost-sure coexistence of the two blue sectors (Chapter 8)

This section completes part (i) of the high-dimensional blue-percolation theorem of Chapter 8: positive root probability in each overlap sector is upgraded to the almost-sure existence of an infinite blue cluster in each sector. The argument follows [214, Props. E.3–E.4]. Throughout, $\beta = c/(2dm)$ with $c > 1$ and d large, so that $p = b\,\varrho^{2d-1} > p_c^{\mathrm{site}}(\mathbb{Z}^d)$, and $\mathbf{P} \in \mathcal{L}_{d,\beta}$ is the limit of the torus laws $\mathbf{P}_{L_n}$.

Conditional bounds pass to the limit. Let $e = uv$ and let $h \ge 0$ be a bounded function of finitely many spins off v and finitely many blue indicators off $\mathrm{star}(v)$. The opening bound of Chapter 8, written without division, says $\mathbb{E}_L[h\,B_e] \ge p\,\mathbb{E}_L[h]$ and $\mathbb{E}_L[h\mathbf{1}\{q_v = s\}] \ge \frac{1}{2}\varrho^{2d}\mathbb{E}_L[h]$ for every large torus. Both sides are expectations of local functions, so the inequalities hold under $\mathbf{P}$, and a monotone-class argument extends them to every bounded nonnegative h measurable with respect to the exterior sigma-field $\mathcal{F}_v^{\mathrm{ext}}$ of the limit. Consequently the one-attempt exploration on $\mathbb{Z}^d$ can be run under $\mathbf{P}$, and under $\mathbf{P}(\,\cdot \mid A)$ for any event A of positive probability that belongs to $\mathcal{F}_v^{\mathrm{ext}}$ for every target $v \ne o$: its histories, intersected with A, are again exterior to the target.

Absence of a sector is an exterior event. Fix $s = \pm 1$ and let T_s be the event that no infinite blue cluster has overlap s. Changing the spins at one vertex v and the blue indicators of $\mathrm{star}(v)$ cannot change T_s: an infinite blue cluster through v splits, when v is removed, into finitely many pieces of which at least one is infinite and has the same overlap, and adding v back can join finite clusters only into a finite one. So $T_s \in \mathcal{F}_v^{\mathrm{ext}}$ for every v.

Conclusion. Suppose $\mathbf{P}(T_s) > 0$. The overlap bound gives $\mathbf{P}(T_s \cap \{q_o = s\}) \ge \frac{1}{2}\varrho^{2d}\,\mathbf{P}(T_s) > 0$, and $A = T_s \cap \{q_o = s\}$ is exterior to every target $v \ne o$. Under $\mathbf{P}(\,\cdot \mid A)$ the exploration embeds a root-forced iid site cluster of parameter p in the blue cluster of o, which is therefore infinite with probability at least $\theta^{\mathrm{site}}(p)/p > 0$. On A this infinite cluster has overlap s, contradicting T_s. Hence $\mathbf{P}(T_s) = 0$ for both signs.

A.18 Two routine proofs for the cavity bounds (Chapter 9)

This section gives the bookkeeping behind Lemma 9.7 and the passage to the limit in Corollary 9.9; we use the notation of Section 9.4.

Proof of Lemma 9.7. The term $v = w$ is $\Xi_H(T \setminus \{w\})$. For $v \neq w$ the set $T \triangle \{v\}$ is even and contains w; apply (9.4) at w with $B = (T \triangle \{v\}) \setminus \{w\}$, which is odd, and note that $B \triangle \partial A = ((T \setminus \{w\}) \triangle \partial A) \triangle \{v\}$. Summing over $v \in V(H) \setminus \{w\} = V(H - w)$ gives the inequality. The vertices of $T \setminus \{w\}$ stay adjacent to $W \subseteq W \cup \{w\}$, and ∂A consists of neighbours of w, so the new pairs are non-fresh. A non-fresh w has lost at least one edge, so its degree in H is at most $\Delta - 1$. □

Proof of Corollary 9.9. (i) is Theorem 9.8 with the identity $\mathbb{E}\langle q_\Lambda^2 \rangle = \chi_{\mathrm{SG}}(\Lambda)/|\Lambda|$ of Chapter 1. (ii) For fixed x, $\sigma_0\tau_0\sigma_x\tau_x$ is a bounded local function, so $\mathbf{E}[\sigma_0\tau_0\sigma_x\tau_x] = \lim \mathbb{E}\langle \sigma_0\sigma_x \rangle^2_{(\mathbb{Z}/L\mathbb{Z})^d} \geq 0$ along the subsequence; finite partial sums are therefore at most $\bar{\chi}$. The law $\mathbf{P}$ is translation invariant, so $\mathbf{E}[q^2_{\Lambda_n}] \leq \bar{\chi}/|\Lambda_n|$, the ergodic theorem gives the almost sure limit Q, and Fatou's lemma gives $\mathbf{E}[Q^2] = 0$. □

A.19 Exponential L^2 decay in fixed dimension (Chapter 9)

This section proves Theorem 9.10. We use the notation of Section 9.4: $J_e = \varepsilon_e |J_e|$ with fair signs, $K_e = \beta |J_e|$, $t_e = \tanh K_e$, $\Xi_G(S) = \mathbb{E}\langle \sigma_S \rangle^2_G$, and the weights $\omega_w(A)$ of (9.4). For $r \geq 1$ put

$$\rho_D(r) := \sum_{A \subseteq [D],\ |A| \text{ odd}} r^{|A|}\, \omega(A), \qquad \bar{\rho}_D(r) := \max_{D' \leq D} \rho_{D'}(r),$$

so that $\rho_D(1) = \rho_D$. Each $\rho_D(r)$ is a polynomial in r with finite coefficients when $\mathbb{E}\cosh(2\beta|J|) < \infty$.

Boundary conditions as one vertex. Let $\Lambda \subset \mathbb{Z}^d$ be finite and η a configuration outside Λ. Let G_Λ have vertex set $\Lambda \cup \{g\}$, the lattice edges of Λ, and, for every edge yy' with $y \in \Lambda$ and $y' \notin \Lambda$, an edge yg carrying the coupling $J_{yy'}\eta_{y'}$. Given $\sigma_g = +1$ the Gibbs measure of G_Λ (free boundary conditions) is $\mu_{\Lambda,\eta}$; given $\sigma_g = -1$ it is its image under the global flip; and $\sigma_g = \pm 1$ with probability $\frac{1}{2}$ each. Hence, for every $A \subseteq \Lambda$,

$$\mu_{\Lambda,\eta}(\sigma_A) = \langle \sigma_A\, \sigma_g^{|A|} \rangle_{G_\Lambda}. \tag{A.3}$$

If η is independent of the couplings of the edges with an endpoint in Λ, then conditionally on η the couplings of G_Λ are iid with the law of J, because that law is symmetric. Every vertex of Λ has degree exactly $2d$ in G_Λ; the *ghost* g has unrestricted degree. So it suffices to prove the following statement about finite multigraphs G in which every vertex other than an optional ghost g has degree at most Δ.

Histories. A *state* is a pair (H, S) with $H = G - W$ for a set $W \not\ni g$ of deleted vertices and $S \subseteq V(H)$ of even size. A *rule* assigns to every state with $S \neq \emptyset$ a particle $w \in S \setminus \{g\}$, as a function of (W, S); one exists because $|S| \geq 2$. The *children* of (H, S) are the states $(H - w, (S \setminus \{w\}) \triangle \partial A)$, for A an odd subset of the star of w in H. A *history* is a finite sequence of states, each a child of the previous one, with weight $W(h) = \prod_j \omega_{w_j}(A_j)$; it is *complete* if it ends with $S = \emptyset$.

Lemma A.2 (Expansion inequality)**.** *For every rule, $\Xi_G(S_0) \le \sum_h W(h)$, the sum running over the complete histories from (G, S_0), and it is finite.*

Proof. Induction on $|V(H)|$. If $S = \emptyset$ both sides equal 1. Otherwise apply (9.4) at the particle w chosen by the rule, with $B = S \setminus \{w\}$, which is odd; if w has no edges, both sides vanish. Every step deletes a vertex, so histories have length at most $|V|$. □

Three facts about histories are used. (a) Each edge lies in the star of at most one expanded vertex, since it disappears with its first deleted endpoint. (b) Call a vertex $x \neq g$ of H *fresh* if no neighbour of x in G has been deleted. Particles created by an expansion are neighbours of the vertex just deleted, so they are not fresh, and a vertex never becomes fresh again; a vertex that is not fresh has degree at most $\Delta - 1$ in H. (c) A step replaces S by $S \triangle \partial_E(A)$, where $\partial_E(A) = \{w\} \triangle \partial A$ is the set of odd-degree vertices of the edge set A (the vertex w has odd degree $|A|$ in it). So the edge sets A_j of a complete history have, together, odd-degree set S_0, and $\sum_j |A_j| \ge \ell_G(S_0)$, the least number of edges of a subgraph of G whose odd-degree set is S_0.

Proposition A.3 (Counting)**.** *Let $r \ge 1$ with $\bar\rho_{\Delta-1}(r) < 1$, put $\kappa(r) = \max\{1, \bar\rho_\Delta(r)/\bar\rho_{\Delta-1}(r)\}$, and let $U(S)$ be the number of fresh particles in $S \setminus \{g\}$. Then for every state (G, S_0),*

$$\Xi_G(S_0) \le \frac{\kappa(r)^{U(S_0)}}{1 - \bar\rho_{\Delta-1}(r)}\, r^{-\ell_G(S_0)}.$$

Proof. Let $W_r(h) = r^{\sum_j |A_j|} W(h)$. For one step from (H, S) with particle w: if w is fresh, then $\sum_A r^{|A|}\omega_w(A) \le \bar\rho_\Delta(r) \le \kappa\, \bar\rho_{\Delta-1}(r)$ and U drops by at least one; if w is not fresh, $\sum_A r^{|A|}\omega_w(A) \le \bar\rho_{\Delta-1}(r)$ and U does not increase, by (b). In both cases $\sum_A r^{|A|}\omega_w(A)\, \kappa^{U(S_A)} \le \bar\rho_{\Delta-1}(r)\, \kappa^{U(S)}$, and by induction the histories of length n from (G, S_0) satisfy $\sum_h W_r(h)\kappa^{U(S_n)} \le \bar\rho_{\Delta-1}(r)^n \kappa^{U(S_0)}$. Since $\kappa^U \ge 1$, summing over n bounds the total W_r-weight of complete histories by $\kappa^{U(S_0)}/(1 - \bar\rho_{\Delta-1}(r))$. By (c), $W(h) \le r^{-\ell_G(S_0)} W_r(h)$ for complete h, and Lemma A.2 concludes. □

Proof of Theorem 9.10. With $\Delta = 2d$, the hypothesis gives $\bar\rho_{\Delta-1}(1) < 1$: for a general law by (9.5), and for $\pm J$ because $\rho_{2d-1} = (c^{2d-1} - 1)/(c + 1)$ and the exact ρ_D is nondecreasing in D. Indeed, with $c = \cosh 2\beta$ and $2\cosh^2\beta = c + 1$, $\rho_{2k} = \frac{1}{2}(c^{2k} - 1) \le \rho_{2k+1} = (c^{2k+1} - 1)/(c + 1)$ is equivalent to $(c - 1)(c^{2k} + 1) \ge 0$, and $\rho_{2k+1} \le \rho_{2k+2}$ because $c + 1 \ge 2$. By continuity of the polynomials $\rho_D(r)$ there is $r > 1$ with $\bar\rho_{\Delta-1}(r) < 1$. Put $C = \kappa(r)/(1 - \bar\rho_{\Delta-1}(r)) \ge 1$. Initially every particle other than g is fresh, so $U(S_0) = |S_0 \setminus \{g\}|$, and Proposition A.3 gives $\Xi_G(S_0) \le C^{|S_0 \setminus \{g\}|} r^{-\ell_G(S_0)}$ (if $S_0 = \emptyset$ both sides are 1). For (i) take $S_0 = \{x, y\}$ without ghost, so $\ell_G = \mathrm{dist}_G(x, y)$. For (ii) use (A.3) conditionally on η: $S_0 = \{x, g\}$ for the one-point function, with $\ell_{G_\Lambda}(\{x, g\}) = 1 + \mathrm{dist}(x, \partial^{\mathrm{in}}\Lambda)$, and $S_0 = \{x, y\}$ for the two-point function, with $\ell_{G_\Lambda}(\{x, y\}) = \min\{\mathrm{dist}(x, y), \mathrm{dist}(x, \partial^{\mathrm{in}}\Lambda) + \mathrm{dist}(y, \partial^{\mathrm{in}}\Lambda) + 2\}$; then average over η. □

A.20 The reduction behind the certified windows (Chapter 9)

This section proves the structural facts used by Theorem 9.12, for $\pm J$ couplings with $K = \beta$, and explains how a potential turns them into decay. The computer then checks finitely many inequalities, as described after the theorem. In outline: the worst case over a window is attained with every local coupling at $\pm K$ and every boundary field at the sum of its magnitudes, and deleting vertices or star edges never increases the worst-case loss, so a loss computed on $\mathbb{Z}^d$

minus the lineage bounds the loss on any subgraph (Lemma A.4); all weights are monotone in β, so a certificate at $t_0 = \tanh\beta_0$ covers every $\beta \le \beta_0$ (Lemma A.5). The losses $L(\pi)$ are rational functions of t for windows of radius 1; for radius 2 they are computed by exact integer enumeration over gauge classes of the internal signs and the extremal boundary fields. For Theorem 9.12(b) the type also records the distance class to the nearest side of the box in each axis, and a move onto the frozen boundary is an absorption. Lemma A.6 turns the potential into the bound of (a), and Lemma A.7, which encodes the boundary condition by the ghost vertex of Section A.19, turns it into the bounds of (b). An independent re-implementation in floating point reproduces the loss patterns to $3 \cdot 10^{-13}$ with margins at least $2 \cdot 10^{-4}$.

Proof of Lemma 9.11. By the Markov property, $\mu(\cosh h) = \sum_\eta \mu(\sigma_{\partial R} = \eta)\, \mu_R^\eta(\cosh h) \ge D_{\min}$, since $\cosh h$ depends only on σ_R; the minimum is also over the star signs, so the proof of Lemma 9.6 goes through. If the y_i are distinct and R is independent, then given η the y_i are independent with means m_i, $|m_i| \le T_i$, and $\mu_R^\eta(\cosh h) = \prod_i \cosh K_i \cdot \frac{1}{2}[\prod_i(1+x_i) + \prod_i(1-x_i)]$ with $x_i = m_i \varepsilon_i t_i$. This function is multilinear, so its minimum over the box $|x_i| \le t_i T_i$, which contains every admissible (x_i), is attained at a vertex $x_i = s_i \tanh\alpha_i$. □

Lemma A.4 (Extremal couplings and deletions)**.** *In the setting of Lemma 9.11, fix all couplings but one, either a coupling J of an edge meeting R or a field on a vertex of R coming from the boundary. As a function of $\tilde t = \tanh(\beta J)$ (or of the tanh of the field), $\mu_R^\eta(\cosh h)$ has the form $(a + b\tilde t)/(c + e\tilde t)$ with $c + e\tilde t > 0$, so it is monotone on $[-\tanh\beta, \tanh\beta]$. Hence the minimum $D_{\min}$ over couplings of magnitude at most β is attained with every such coupling at $\pm\beta$, and setting a coupling to 0, that is deleting the edge, never lowers $D_{\min}$ and never raises the worst-case loss. Removing a star edge never raises the worst-case loss either and leaves the weights $\prod_{i\in A} t_i^2$ of the remaining subsets A unchanged.*

Proof. With $e^{\beta J \sigma_a \sigma_b} = \cosh(\beta J)(1 + \tilde t \sigma_a \sigma_b)$, both the numerator and the partition function of $\mu_R^\eta(\cosh h)$ are affine in $\tilde t$, and the partition function is positive; a ratio of affine functions with positive denominator is monotone on an interval. The value $\tilde t = 0$ lies in the interval, so the minimum over the interval is at most the value with the edge deleted. For a star edge e_j, the quantity $\mu(\cosh(K_j \varepsilon_j \sigma_{y_j} + r))/\cosh K_j$ is affine in $\varepsilon_j t_j$ because μ does not depend on the star, so its smaller value at $\varepsilon_j = \pm 1$ is at most its value at 0, which is the corresponding quantity without e_j; dividing $\prod_i \cosh K_i$ by it gives the claim. □

Lemma A.5 (Monotonicity in β)**.** *For windows $R = \{y_i\}$ the loss bound of Lemma 9.11 is nondecreasing in β. For windows of radius 2 and $D \le 3$ (all full types in $d = 2$; the root, with $D = 4$, uses $R = \{y_i\}$), if the environment couplings are only required to lie in $[-\beta_0, \beta_0]$, the loss at star coupling $\beta \le \beta_0$ is at most $\max\{1, L(\beta_0)\}$, where $L(\beta_0)$ is the loss at β_0.*

Proof. For the first claim, $\tanh\alpha_i = t \tanh(k_i\beta)$, with k_i the number of other edges at y_i, increases with β, and $\min_s x_s$ is nonincreasing in each α_j. Indeed $x_s = \cosh(\sum_i s_i\alpha_i)/\prod_i \cosh\alpha_i$, and at a minimizer s with $\Sigma = \sum_i s_i\alpha_i \ge 0$, flipping $s_j = +1$ cannot decrease $|\Sigma|$, so $\Sigma \le \alpha_j$ and $\partial_{\alpha_j} \log x_s = s_j \tanh\Sigma - \tanh\alpha_j \le 0$; the lower envelope of finitely many such functions is then nonincreasing. For the second, fix one of the finitely many extremal environment configurations; it fixes the law p of $Z = \sum_i \sigma_{y_i}$, and the loss is $h(\beta) = \cosh^{2D}\beta/(\mathbb{E}_p \cosh\beta Z)^2$. For $D = 3$, using $\sinh 3\beta - \tanh\beta \cosh 3\beta = 2\sinh\beta$, the derivative of $\log h$ has the sign of $1 - 4p(|Z| = 3)$; for $D = 2$ it has the sign of $p(Z = 0) - p(|Z| = 2)$. Neither depends on β, so h is monotone and $h(\beta) \le \max\{h(0), h(\beta_0)\} = \max\{1, h(\beta_0)\}$. Shrinking the environment interval from $[-\beta_0, \beta_0]$ to $[-\beta, \beta]$ only lowers the worst case. □

Lemma A.6 (Potentials give decay)**.** *Let types, available moves and type losses $L(\pi)$ be as in Section 9.5, with $L(\pi)$ an upper bound for the layer loss of every particle of type π, in every subgraph of $\mathbb{Z}^d$ and in every graph obtained by deleting vertices from a torus whose side is large enough that the window of a particle, together with the lineage vertices that lie in it, is a copy of the corresponding configuration in $\mathbb{Z}^d$, and for every choice of the signs of the couplings in its window. (For the windows and memories of Section 9.5, which with their boundary have radius at most 3 and whose lineages lie within distance 5 of the particle, side at least 11 suffices.) If $\varphi \geq 1$ satisfies (9.7) for every type, with $\lambda < 1$, then for a finite subgraph of $\mathbb{Z}^d$ or such a torus and $S_0 = \{x, y\}$, $\mathbb{E}\langle\sigma_x\sigma_y\rangle^2 \leq \varphi(\text{root})\lambda^{\text{dist}(x,y)/(2d)}/(1-\lambda)$.*

Proof. Expand with the rule "never expand y", possible because $|S|$ is even and there is no ghost, and give y potential 1. Put $\Phi(S) = \prod_{p\in S}\varphi(\text{type } p)$. In one step at a particle w, the true loss is at most $L(\text{type } w)$ by Lemma A.4, missing neighbours only remove terms, and a child that annihilates removes a factor $\varphi \geq 1$; so $\sum_A(\text{weight of } A)\,\Phi(S_A) \leq \lambda\Phi(S)$. By induction the histories of length n have total weight at most $\lambda^n\Phi(S_0)$, and $\Phi \geq 1$ bounds the weight of complete histories. A complete history uses edge sets whose odd-degree set is $\{x, y\}$ (fact (c) of Section A.19), and each step uses at most $2d$ edges, so it has at least $\text{dist}(x, y)/(2d)$ steps. □

For Theorem 9.12(b) the boundary condition is encoded by the ghost g of Section A.19. The proof of Lemma 9.11 applies when g lies in the vertex boundary of the window R rather than in R, since $\cosh h$ depends only on σ_R and σ_g. Given σ_g, an edge from the particle to g adds a constant to h, so the closed form of Lemma 9.11 holds with $T_i = 1$ for such edges (which need not be distinct), and an edge from a vertex of R to g acts on it as a field. These frozen neighbours and boundary fields are how the certificate of Theorem 9.12(b) sees the boundary.

Lemma A.7 (Potentials give decay: fixed boundary conditions)**.** *Let $d = 2$, let $\Lambda \subset \mathbb{Z}^2$ be a rectangular box whose sides both have at least 8 sites, so that each site of Λ is within distance 3 of at most one side in each axis, and let G_Λ be its graph with ghost g. Let the types of Section 9.5 also record the distance class ($0, 1, 2, 3$ or far) to the nearest side in each axis, let a move along an edge to g be an absorption, whose child enters (9.7) with potential 1, and, where the class of a child is not determined by the type of its parent and the move (a move from the class far towards a side, or, on a side of exactly 8 sites, a move away from that side from distance 3), read $\varphi(\text{child}(\pi, m))$ in (9.7) as the maximum of φ over the classes the child can have; and let $L(\pi)$ be an upper bound for the layer loss, with g in the vertex boundary of the window, of every particle of type π in every graph $G_\Lambda - W$ with $W \subseteq \Lambda$, for every choice of the signs of the couplings in its window. If $\varphi \geq 1$ satisfies (9.7) for every type, with $\lambda < 1$, and $\varphi \leq \varphi_{\max}$ on the types of initial particles, then for every boundary condition η independent of the couplings, with $\delta'_x = \text{dist}(x, \mathbb{Z}^2 \setminus \Lambda)$,*

$$\mathbb{E}\,\mu_{\Lambda,\eta}(\sigma_x)^2 \leq \frac{\varphi_{\max}}{1-\lambda}\,\lambda^{\delta'_x/4}, \qquad \mathbb{E}\,\mu_{\Lambda,\eta}(\sigma_x\sigma_y)^2 \leq \frac{\varphi_{\max}^2}{1-\lambda}\,\lambda^{\min\{\text{dist}(x,y),\,\delta'_x+\delta'_y\}/4}.$$

Proof. Fix η. By (A.3) the left sides are $\Xi_{G_\Lambda}(\{x, g\})$ and $\Xi_{G_\Lambda}(\{x, y\})$, and the couplings of G_Λ are iid $\pm J$. Expand with the rule "never expand g", possible because a nonempty even set contains a particle other than g, and put $\Phi(S) = \prod_{p\in S\setminus\{g\}}\varphi(\text{type } p)$, so that g has potential 1. In one step at a particle w, the true loss is at most $L(\text{type } w)$, missing neighbours only remove terms, a child on g toggles g and multiplies Φ by 1, as the absorption terms of (9.7) assume, a child on another particle removes a factor $\varphi \geq 1$, and a new particle has potential at most the value that (9.7) uses for its move; so $\sum_A(\text{weight of } A)\,\Phi(S_A) \leq \lambda\Phi(S)$. By induction

the histories of length n have total weight at most $\lambda^n\Phi(S_0)$, and $\Phi(S_0)$ is at most $\varphi_{\max}$ for $S_0=\{x,g\}$ and at most $\varphi_{\max}^2$ for $S_0=\{x,y\}$. By facts (a) and (c) of Section A.19, the edges used by a complete history form a subgraph of G_Λ whose odd-degree set is S_0, so they contain a path from x to the other point of S_0. A path from x to g has at least $\operatorname{dist}(x,\partial^{\mathrm{in}}\Lambda)+1=\delta'_x$ edges. A path from x to y either avoids g, and then has at least $\operatorname{dist}(x,y)$ edges because Λ is a rectangle, or passes through g, and then has at least $\delta'_x+\delta'_y$ edges. Each step uses at most 4 edges, so a complete history has at least $\delta'_x/4$, respectively $\min\{\operatorname{dist}(x,y),\delta'_x+\delta'_y\}/4$, steps, and $\sum_{n\ge n_0}\lambda^n\le\lambda^{n_0}/(1-\lambda)$ gives the two bounds given η. Average over η. □

The certificate of Theorem 9.12(b) supplies such $L(\pi)$ and φ at $t_0=4762/10000$, with $\lambda=0.9995$ and $\varphi_{\max}=4.1487331$, and Lemma A.5 extends it to every $\beta\le\operatorname{artanh}t_0$.

A.21 Sharpness of the cavity loss (Chapter 9)

This section proves the first claim of Route warning 9.13. We use the notation of the proof of Lemma 9.6: $h=\sum_i K_i\varepsilon_i\sigma_{y_i}$, $s_A=\prod_{i\in A}\sinh K_i\prod_{i\notin A}\cosh K_i$, and $\varepsilon^A=\prod_{i\in A}\varepsilon_i$.

Lemma A.8 (The loss L_u is sharp for cavity measures)**.** *Let $D\ge1$ and $K_1,\dots,K_D>0$, let $y_1,\dots,y_D,b$ be distinct vertices of H, and for a probability measure μ on the spins of H put*

$$\mathcal{R}(\mu)=\frac{\mathbb{E}_\varepsilon\Big[\big(\mu(\sigma_b\sinh h)/\mu(\cosh h)\big)^2\Big]}{\sum_{A\subseteq[D],\ |A|\ \mathrm{odd}}\prod_{i\in A}t_i^2\,\mu(\sigma_b\sigma_{\partial A})^2}$$

whenever the denominator is positive. Then $\mathcal{R}(\mu)\le L_u$, and the supremum of $\mathcal{R}(\mu)$ over flip-invariant μ equals L_u.

Proof. The upper bound is the proof of Lemma 9.6 with $B=\{b\}$. Numerator and denominator depend on μ only through the law of $(\sigma_{y_1},\dots,\sigma_{y_D},\sigma_b)$, so it suffices to specify that law (and to let the other spins be independent fair signs). Write $\sigma_y^A=\prod_{i\in A}\sigma_{y_i}$, let s^* minimize $|\sum_i s_iK_i|$, and put $m=\cosh(\sum_i s_i^*K_i)$ and $c=\prod_i\cosh K_i$, so that $L_u=c^2/m^2$. Let μ_0 be the law under which $(\sigma_{y_i})_i$ equals s^* or $-s^*$ with probability $\frac12$ each and σ_b is an independent fair sign. Let μ_1 be the law under which $(\sigma_{y_i})_i$ is uniform on $\{\pm1\}^D$ and, given it, σ_b has conditional mean $F=\kappa\sum_{A\text{ odd}}s_A^{-1}\sigma_y^A$, with $\kappa>0$ small enough that $|F|\le1$. The law μ_0 is flip-invariant, and so is μ_1 because F is odd; hence so is $\mu_\delta=(1-\delta)\mu_0+\delta\mu_1$ for $\delta\in(0,1)$.

Under μ_0, σ_b is a fair sign independent of the σ_{y_i}, so $\mu_0(\sigma_b\sigma_y^A)=0$; under μ_1, by orthonormality of the characters, $\mu_1(\sigma_b\sigma_y^A)=\kappa/s_A$ for odd A. By the expansion of $\sinh h$ in the proof of Lemma 9.6 and $\sum_{A\text{ odd}}\varepsilon^A=\frac12\big[\prod_i(1+\varepsilon_i)-\prod_i(1-\varepsilon_i)\big]$,

$$\mu_\delta(\sigma_b\sinh h)=\delta\kappa\sum_{A\text{ odd}}\varepsilon^A=\begin{cases}\pm\delta\kappa\,2^{D-1} & \text{if }\varepsilon=\pm(1,\dots,1),\\ 0 & \text{otherwise.}\end{cases}$$

At $\varepsilon=\pm(1,\dots,1)$, $\mu_0(\cosh h)=m$ and $\mu_1(\cosh h)=c$, so $\mu_\delta(\cosh h)=(1-\delta)m+\delta c$. The numerator of $\mathcal{R}(\mu_\delta)$ is therefore $2^{1-D}(\delta\kappa2^{D-1})^2/((1-\delta)m+\delta c)^2$. In the denominator each odd A contributes $\prod_{i\in A}t_i^2\,(\delta\kappa/s_A)^2=\delta^2\kappa^2/c^2$, and there are 2^{D-1} of them. Hence $\mathcal{R}(\mu_\delta)=c^2/((1-\delta)m+\delta c)^2$, which tends to $c^2/m^2=L_u$ as $\delta\to0$. □

In the extremal measures the correlation of σ_b with the star is tuned to the star magnitudes and concentrated on the two sign patterns at which $\mu(\cosh h)$ is smallest. The lemma concerns the class of all flip-invariant measures, to every one of which the proof of Lemma 9.6 applies; it does not say how close Gibbs measures of subgraphs of $\mathbb{Z}^d$ come to these extremal measures.

A.22 Tilted sign laws on three parallel edges (Chapter 11)

This section proves Proposition 11.5.

Proof of Proposition 11.5. Here $Z_\varepsilon = 4\cosh(K\Sigma)$ with $\Sigma = \varepsilon_1+\varepsilon_2+\varepsilon_3$, which is ± 3 for two sign patterns and ± 1 for six. With $C_1 = \cosh K$ and $C_3 = \cosh 3K$,

$$1 - \mathbb{E}_{\mathbb{P}_\alpha} c^2 = \frac{6C_1^\alpha \operatorname{sech}^2 K + 2C_3^\alpha \operatorname{sech}^2 3K}{6C_1^\alpha + 2C_3^\alpha}.$$

Since $1+3y = 3\cosh^2 K - 2 \le 3\cosh^2 K$ and the denominator exceeds $2C_3^\alpha$, $(1-\mathbb{E}_{\mathbb{P}_\alpha} c^2)(1+3y) \le 9(C_1/C_3)^\alpha + 3\cosh^2 K/\cosh^2 3K$, which tends to 0 as $K\to\infty$. Hence $1-\mathbb{E}_{\mathbb{P}_\alpha} c^2 < 1/(1+3y)$ for large K. □

A.23 Restoring one edge (Chapter 11)

This section proves Lemma 11.7.

Proof of Lemma 11.7. Let $e = ab$. The weight of G is that of $G-e$ times $\exp(\varepsilon_e K_e \sigma_a\sigma_b)$, a function of $\sigma_a\sigma_b$; so conditionally on $\sigma_a\sigma_b$ the two measures agree, and $\langle \sigma_u\sigma_v \mid \sigma_a\sigma_b\rangle_G = \alpha_e + B_e\sigma_a\sigma_b$ with α_e, B_e not depending on ε_e. Let b_e be the effective coupling of $G-e$ between a and b, $\tanh b_e = \langle\sigma_a\sigma_b\rangle_{G-e}$; then $\langle\sigma_a\sigma_b\rangle_G = \tanh(b_e+\varepsilon_e K_e)$ and $c = \alpha_e + B_e\tanh(b_e+\varepsilon_e K_e)$. By Lemma 11.3,

$$(D_e c)^2 = B_e^2 \operatorname{Var}_{\varepsilon_e} \tanh(b_e+\varepsilon_e K_e) \le y_e\, \mathbb{E}_{\varepsilon_e}\big[B_e^2 \operatorname{sech}^2(b_e+\varepsilon_e K_e)\big].$$

Here $B_e^2\operatorname{sech}^2(b_e+\varepsilon_e K_e) = B_e^2 \operatorname{Var}_G(\sigma_a\sigma_b)$ is the variance under G of the conditional expectation $\langle\sigma_u\sigma_v \mid \sigma_a\sigma_b\rangle_G$, which is at most $\operatorname{Var}_G(\sigma_u\sigma_v) = 1-c^2$. □

A.24 The closure theorem (Chapter 11)

This section proves Theorem 11.19. Part (a) is short:

Proof of (a). Every u–v path of $N_1 \parallel N_2$ lies in one part, so the union of independent copies of the two random sets separates, with marginals $w^1 \oplus w^2$. For $e \in N_1$, K_2 does not depend on ε_e, so $D_e \tanh(b+K_1+K_2) = D_e\tanh((b+K_2)+K_1)$, and averaging first over the signs of N_1,

$$\sum_{e\in N_1} w_e^1\, I_e(b) = \mathbb{E}_{N_2} \sum_{e\in N_1} w_e^1\, I_e^{N_1}(b+K_2) \le \rho_1\, \mathbb{E}_{N_2} S_{N_1}(b+K_2) = \rho_1\, S(b),$$

by tameness of N_1 at the coupling $b+K_2$. Adding the same bound for N_2 gives the claim. □

The series step is the heart of the matter. We prove part (b): if N_1 and N_2 are tame with (ρ_1, w^1) and (ρ_2, w^2), then their series composition $M = N_1 - N_2$ is tame with $(\rho_1\rho_2/(\rho_1+\rho_2),\ \alpha w^1 \oplus (1-\alpha)w^2)$, $\alpha = \rho_2/(\rho_1+\rho_2)$. Throughout, $\tau_i = c_{N_i}$ is the correlation across N_i (a function of the signs of N_i only), $a = \tanh b$, and

$$\varphi_a(x) = \frac{a+x}{1+ax}, \qquad \lambda_a(t) = \frac{t(1-a^2)}{1-a^2t^2} = \tfrac{1}{2}(\varphi_a(t) - \varphi_a(-t)),$$

so that $\tanh(b+K)=\varphi_a(\tanh K)$ and $\lambda_a(\tanh K)$ is the sign-flip derivative of $\tanh(b+\varepsilon K)$. We use repeatedly the identity

$$1-\varphi_c(\tau)^2=\frac{(1-c^2)(1-\tau^2)}{(1+c\tau)^2}\qquad(|c|,|\tau|<1).\tag{A.4}$$

Lemma A.9 (Möbius factorization)**.** *For $|a|,|s|<1$ and $\tau\in(-1,1)$, $\varphi_a(s\tau)=\lambda_a(s)\,\varphi_{as}(\tau)+\mu_a(s)$ with $\mu_a(s)=a(1-s^2)/(1-a^2s^2)$.*

Proof. Over the common denominator $(1-a^2s^2)(1+as\tau)$ the right side has numerator $s(1-a^2)(as+\tau)+a(1-s^2)(1+as\tau)=(1-a^2s^2)(a+s\tau)$. □

Lemma A.10 (Influences in series)**.** *For an edge e of N_1, with $b_2'=\operatorname{artanh}(a\tau_2)$,*

$$I_e^M(b)=\mathbb{E}_{N_2}\big[\lambda_a(\tau_2)^2\,I_e^{N_1}(b_2')\big],\qquad \mathbb{E}_{N_1}\big[1-\varphi_{a\tau_2}(\tau_1)^2\big]=S_{N_1}(b_2'),$$

where $\mathbb{E}_{N_i}$ averages the signs of N_i. The same holds with N_1 and N_2 exchanged.

Proof. $\tanh(b+K_M)=\varphi_a(\tau_1\tau_2)$. By Lemma A.9 with $s=\tau_2$, which does not depend on ε_e, $D_e\varphi_a(\tau_1\tau_2)=\lambda_a(\tau_2)\,D_e\varphi_{a\tau_2}(\tau_1)$, and $\varphi_{a\tau_2}(\tau_1)=\tanh(b_2'+K_{N_1})$. Square and average over the signs of N_1, then of N_2. The second identity is the definition of S_{N_1} at the coupling b_2'. □

Lemma A.11 (Series slack)**.** *For all $a,x,z\in(-1,1)$,*

$$\lambda_a(z)^2\big(1-\varphi_{az}(x)^2\big)+\lambda_a(x)^2\big(1-\varphi_{ax}(z)^2\big)\;\le\;1-\varphi_a(xz)^2.$$

Proof. Write $p=x^2$, $q=z^2$, $A=a^2$. By (A.4) the first term on the left is $q(1-A)^2(1-p)/\big((1-Aq)(1+axz)^2\big)$, the second is the same with p and q exchanged, and the right side is $(1-A)(1-pq)/(1+axz)^2$. Using $(1-A)/(1-Aq)=1-A(1-q)/(1-Aq)$ and the same with p,

$$\text{right}-\text{left}=\frac{1-A}{(1+axz)^2}\Big[(1-p)(1-q)+\frac{Aq(1-p)(1-q)}{1-Aq}+\frac{Ap(1-p)(1-q)}{1-Ap}\Big]\;\ge\;0.\qquad\square$$

At $a=0$ the lemma says $q(1-p)+p(1-q)\le 1-pq$, that is, $(1-p)(1-q)\ge 0$.

Proof of part (b) of the closure theorem. Let C_1,C_2 be the random separating sets of N_1,N_2, independent of each other and of the signs, and choose C_1 with probability α and C_2 otherwise, by an independent coin. Each choice meets every u–v path of M, and the marginals are $\alpha w^1\oplus(1-\alpha)w^2$. By Lemma A.10 and tameness of N_1 at the coupling b_2', which depends only on the signs of N_2,

$$\sum_{e\in N_1}\alpha w_e^1\,I_e^M(b)=\alpha\,\mathbb{E}_{N_2}\Big[\lambda_a(\tau_2)^2\sum_{e\in N_1}w_e^1 I_e^{N_1}(b_2')\Big]\le\alpha\rho_1\,\mathbb{E}\big[\lambda_a(\tau_2)^2\big(1-\varphi_{a\tau_2}(\tau_1)^2\big)\big],$$

and symmetrically the N_2 block is at most $(1-\alpha)\rho_2\,\mathbb{E}[\lambda_a(\tau_1)^2(1-\varphi_{a\tau_1}(\tau_2)^2)]$. With $\alpha=\rho_2/(\rho_1+\rho_2)$ both coefficients equal $\rho=\rho_1\rho_2/(\rho_1+\rho_2)$, and Lemma A.11, applied pointwise in (τ_1,τ_2), bounds the sum by $\rho\,\mathbb{E}[1-\varphi_a(\tau_1\tau_2)^2]=\rho\,S_M(b)$. If $\rho_1=0<\rho_2$ then $\alpha=1$ and $\rho=0$, and the same computation gives 0 for both blocks; if $\rho_1=\rho_2=0$ both blocks vanish for any α. □

A.25 The edge of the random-cluster family (Chapter 11)

This section proves Proposition 11.24.

Proof of Proposition 11.24. (a) Two parallel edges with activities a, b; put $x = a^2b^2 \in (0,1)$. Multiplying the weights $1, y_a, y_b, qy_ay_b$ of $\emptyset, \{a\}, \{b\}, \{a,b\}$ by $(1-a^2)(1-b^2)$ gives total weight proportional to $1+(q-1)x$ and $1-\varphi_q(u \leftrightarrow v) = (1-a^2)(1-b^2)/(1+(q-1)x)$, while $1-\mathcal{Q} = \frac{1}{2}[\operatorname{sech}^2(K_a+K_b) + \operatorname{sech}^2(K_a-K_b)] = (1-a^2)(1-b^2)(1+x)/(1-x)^2$. When $1+(q-1)x > 0$, C_q is therefore equivalent to $(1+x)(1+(q-1)x) \geq (1-x)^2$, that is, to $x\,[(q+2)+(q-2)x] \geq 0$. At $q=0$ this holds, with slack $2x(1-x)$. For $-2 < q < 0$ it fails for $(2+q)/(2-q) < x < 1/(1-q)$, a nonempty window since $(2+q)(1-q) < 2-q$; for $q \leq -2$ it fails for every $0 < x < 1/(1-q)$. (b) φ_q is the random-cluster measure with $p = qy/(1+qy)$ and $q \leq 1$, which stochastically dominates Bernoulli(p) percolation [109, Theorem 3.21]; and $p > \frac{1}{2} = p_c(\mathbb{Z}^2)$ [137]. □

A.26 The random-cluster electrical factor (Chapter 11)

This section records the obstruction to extending the forest bound of Chapter 11 to random-cluster weights with $q > 0$, in the notation of Proposition 11.24.

Observation A.12 (The random-cluster electrical factor is unbounded). For $q > 0$ and constant weight y, φ_q has weight $w^{|\omega|}q^{k(\omega)}$, $w = qy$, $k(\omega)$ the number of clusters, and the argument of Theorem 11.11 gives

$$\frac{\varphi_q(u \leftrightarrow v)}{\varphi_q(u \not\leftrightarrow v)} = \mathbb{E}\Big[\frac{C_w(G[S])}{q\,C_{2,w}(G[S];u|v)} \;\Big|\; u \not\leftrightarrow v\Big], \qquad S = S_u \cup S_v,$$

where $C_w(H)$ sums $w^{|A|}$ over connected spanning subgraphs A of H and $C_{2,w}$ over spanning subgraphs with two components separating u from v (checked in exact arithmetic on random graphs for $q \in (0,1)$; not yet independently reviewed). The electrical factor ρ_w defined by $C_w/C_{2,w} = w\,\rho_w\,C_{\mathrm{eff}}^{\mathrm{unit}}$ tends to 1 as $w \to 0$ (Lemma 11.10), recovering Theorem 11.11, and a bound on it uniform in the subgraph would extend the forest argument to $q > 0$; but it is unbounded on subgraphs of $\mathbb{Z}^2$. For $w > 1$, at $q = 1$ the left side concerns Bernoulli percolation with $p = w/(1+w) > \frac{1}{2}$, which does not decay in a box of side N, while the unit conductance between distant points is at most $8/\ln N$ up to constants, so $\sup_S \rho_w \gtrsim \ln N$; for $w \leq 1$ the same is expected from the predicted self-dual threshold, and transfer-matrix computations on strips of width W show $\rho_w \approx W^{cw}$, $c \approx 0.5$, at small w. Both parts rest on derivations and numerics that have not been independently reviewed. The forest proof works because the factor is exactly 1 at $q = 0$.

A.27 The uniform-flux XY comparison on a ring of squares (Chapter 11)

This section records the planar failure cited in the route warning on exact $\pm J$ cancellations of Chapter 11.

Observation A.13 (The $\pm J$ form also fails on a subgraph of $\mathbb{Z}^2$). At homogeneous activity $t = 0.9$ the comparison with the uniform-flux XY glass fails on a ring of twelve corner-touching squares

in an 8×8 box, closed by one edge, by about 5×10^{-5} (exact on the Ising side; Monte Carlo on the XY side, 138 standard errors; not independently reviewed); in the $\pm J$ glass each square cancels exactly with probability $\frac{1}{2}$. On full boxes no violation has been found (up to 5×6).

A.28 Exotic states of long-range chains (Chapter 12)

We prove part (a) of Theorem 12.15 for one fixed $\hat{\sigma}$ and part (c), and indicate what the statement for all $\hat{\sigma}$ at once and part (b) need, following [102, §§3–4].

The forest. Each site x receives an infinite set U_x of children with $\sum_{y \in U_x} |c_{y-x}| = \infty$; the sets U_x are disjoint, and the directed graph $x \to y$ ($y \in U_x$) is a union of disjoint trees whose roots are exactly the points of S. The forest can be chosen with the additional property $\sum_{w \in S} [\sum_{y \in T_w} c_{y-x}^2]^{1/2} < \infty$ for every x, where T_w is the tree of w; its existence under the stated hypotheses is a deterministic construction by successive halving [102, §3].

(a) for fixed $\hat{\sigma}$. Set $\sigma_w^* = \hat{\sigma}_w$ for $w \in S$, and recursively $\sigma_y^* = \operatorname{sgn}(J_{xy}) \sigma_x^*$ for $y \in U_x$ (with $\operatorname{sgn} 0 := 1$). It suffices to show $\sigma_x^* h_x = +\infty$ for every x, where $h_x = \lim_m \sum_{y \in \Lambda_m} J_{xy} \sigma_y^*$: then $\Delta H_A = 2 \sum_{x \in A} \sigma_x^* h_x - 2 \sum_{x,y \in A} J_{xy} \sigma_x^* \sigma_y^* = +\infty$ for every nonempty finite A. Split the sum defining $\sigma_x^* h_x$ into the tree neighbours $y \in U_x \cup \{\text{parent of } x\}$ and the rest. For tree neighbours $J_{xy} \sigma_x^* \sigma_y^* = |J_{xy}|$, and $\sum_{y \in U_x} |c_{y-x}| |Z_{xy}| = \infty$ almost surely: the truncated sums $\sum |c_{y-x}| \min(|Z_{xy}|, 1)$ have means tending to infinity and variances bounded by $\sum_y c_y^2$. For the rest, the sign $\sigma_x^* \sigma_y^*$ is a product of signs of couplings on tree edges, none of which is $\{x, y\}$, so it is independent of the family $(Z_{xy})_y$ of non-tree couplings at x. Conditionally on the tree couplings, the partial sums are sums of independent centred terms with $\sum_y c_{y-x}^2 < \infty$, which converge almost surely. So $\sigma_x^* h_x = +\infty$ almost surely for each of the countably many x.

All $\hat{\sigma}$ at once. The tree-neighbour part does not depend on $\hat{\sigma}$. The rest splits into one sum per tree, and changing $\hat{\sigma}$ multiplies the sum over T_w by ± 1 as a whole. The summability of the per-tree L^2 norms, $\sum_{w \in S} [\sum_{y \in T_w} c_{y-x}^2]^{1/2} < \infty$, with Doob's maximal inequality gives $\sum_w \sup_m |V_w^{(m)}| < \infty$ almost surely for the per-tree partial sums $V_w^{(m)}$, so by dominated convergence the sum over trees converges for every sign choice simultaneously [102, (3.8)–(3.13)].

(b). Give each site only finitely many children, as many as needed to make the coherent field larger than θ_A while keeping it finite [102, §4].

(c). For finite sets A, A' the cocycle identity $\Delta H_A(\tilde{\sigma}) + \Delta H_{A'}(\tilde{\sigma}^A) = \Delta H_{A \triangle A'}(\tilde{\sigma})$ holds, because each term is a limit along Λ_m of finite-volume energy differences, which are additive. Fix a finite Λ and condition on the spins off Λ; under μ they equal $\tilde{\sigma}^{A_0}$ off Λ for some finite A_0, and the compatible configurations are $\tilde{\sigma}^A$ with $A \setminus \Lambda = A_0 \setminus \Lambda$. For two of them, the cocycle identity gives $\mu(\tilde{\sigma}^A)/\mu(\tilde{\sigma}^{A'}) = e^{-\beta \Delta H_{A \triangle A'}(\tilde{\sigma}^{A'})}$ with $A \triangle A' \subset \Lambda$, which is the ratio of the finite-volume Gibbs weights in Λ with boundary condition $\tilde{\sigma}^{A_0}$ (an allowed boundary condition, since h_x changes by a finite amount under a finite flip). So μ satisfies the DLR equations. For the last assertion, pick uncountably many $\hat{\sigma}$ that pairwise differ at infinitely many sites of S; the resulting $\tilde{\sigma}$ are not finite perturbations of each other, so the measures in (c) are supported on disjoint sets.

A.29 The certified bound for biased couplings (Chapter 14)

We sketch the structure of the proof of Theorem 14.9; the local lemmas are long and are not reproduced. The model is the square lattice with iid couplings ν_p, $p \geq 0.9537$, at fixed $\beta \geq 10^6$.

Gauge. Colour $\mathbb{Z}^2$ black and white. At a black site v put $\eta_v = -1$ if at least three of its four couplings are negative and $\eta_v = +1$ otherwise, and put $\eta = +1$ at white sites. In the gauge-transformed couplings $J' = J^\eta$ the boundary condition η becomes plus, the overlap is unchanged, and the four-bond stars of distinct black sites are independent, with an explicit law in which at most two bonds of a star are negative.

Contours and local flip families. As in the proof of Theorem 14.8, $\sigma_x = -1$ under the plus boundary condition forces a simple dual polygon γ around x whose crossed edges all disagree. The bound $\mu(E_\gamma) \le e^{-2\beta S_\gamma}$ of the interior flip is improved by three further injective flip families: flips of short "arms" next to flippable corners of γ, flips of three- and four-site caps, and strict local deformations of γ. Each family gives a separate upper bound on $\mu(E_\gamma)$ of the form $e^{-2\beta S_\gamma}/(\text{number of images})$, up to small errors controlled at $\beta \ge 10^6$; the bounds are interpolated into one weight W_γ that factorizes over a window of nine consecutive signed edges of γ.

Star law. A polygon visits a black star in a run of one to three consecutive edges or in two opposite edges. Replacing the joint law of each opposite pair by independent signs with the same marginals can only increase the expected weight, by a mixed-difference inequality; what remains is a product of exact run laws.

Finite-memory matrix. Self-avoiding polygons are encoded by histories of fifteen consecutive dual edges with no repeated vertex. The expected weight of a polygon is a cyclic product of entries of a nonnegative rational matrix $\mathcal{M}$ on 453,233,520 states, whose entries combine the sign probabilities of the run law, a tilt, and the window factors. The computer checks $\mathcal{M}v \le \frac{9999}{10000}v$ for an explicit positive integer vector v, which bounds the trace of $\mathcal{M}^n$, and hence the expected weight of all polygons of length n around a site, by a constant times $(9999/10000)^n$. The generating program and a separately written checker, which share the definitions of the row weights, verify the inequality in exact arithmetic.

From summability to order. The favourable-patch argument of Theorem 14.8 turns contour summability into $\mathbb{E}\langle q_\Lambda \rangle \ge c_0$ for the boundary condition $\eta(J)$. Comparison with free boundary conditions changes the two-replica quenched pressure by $O(|\partial\Lambda|/|\Lambda|)$ uniformly in the replica coupling λ, and convexity in λ together with Jensen's inequality gives $P(\beta,\lambda) - P(\beta,0) \ge \beta\lambda c_0$ for the pressure $P = -\beta f^{(2)}$, so the right derivative at $\lambda = 0$ is positive and does not depend on the boundary condition. Compactness of the joint law of couplings and finite-volume Gibbs measures, and ergodicity of the disorder, give almost-sure non-uniqueness, as at the end of the proof of Theorem 14.8.

Appendix B

Computer-assisted certificates

Several results of the review are proved with a computer. This appendix records, for each of them, what the computer checks, in which arithmetic, how the check was verified independently, and where the programs are. The mathematical reduction from the theorem to the finite computation is in the chapter cited (or in the part of Appendix A to which it defers), or, where a row says so, in the source cited there; it is part of the proof and is not repeated here.

Two conventions hold throughout. A *certificate* is a finite list of inequalities between explicitly given numbers, decided either in exact integer or rational arithmetic, or in interval arithmetic with outward (directed) rounding, so that no conclusion rests on an unenclosed floating-point value; floating-point searches are used only to find the data that the certificate then checks. An *independent check* is either a second program written without sharing code with the first, or a re-implementation by an auditor from the written reduction; we say which, and when the check is in floating point we say so.

B.1 Certified theorems

Tables B.1 and B.2 list the theorems of the review whose proofs contain a certificate, and the certified cases of conjectures stated in the text.

Where the programs are. The programs and frozen inputs for the diamond-lattice certificates are in the public repository `https://github.com/PeaBrane/mk-spin-glass-order` at commit `21744230ff7d`, as stated in [219]. The programs for the diluted windows are part of the Zenodo deposit of [217]; they are also in the repository `https://github.com/PeaBrane/cmr-blue-percolation` at revision `f0a590cb168c`, directory `diluted_uniqueness`. The checks of finite χ_{SG} and exact balance at the explicit points are in the same repository at commit `8da06c6`. The programs for the L^2 windows and for the comparison certificates (forest domination on the Wheatstone bridge, on K_4 and on small grids, and the flow-weighted influence bound) are in the companion code record of this review [220], with a table that maps each result to its program and expected output. The biased-order certificate is in [213]. The certificates for explicit blue percolation in $d = 7$–12 belong to a revised version of [214], which is not yet public, and are not part of these records.

Result	What is checked	Arithmetic	Independent check
Both-sector blue percolation, fair ± 1, seventeen explicit (d, t) with $d = 7, \dots, 12$ (Chapter 8)	local opening floors for the overlap-revealed exploration and a one-sided Holley comparison, as finitely many inequalities between sums of rationals; second-moment criteria for weighted oriented paths ($d = 9$–12) or paths with short lateral excursions ($d = 7$–9)	exact rational; rational enclosures, or directed rounding followed by an exact comparison	two independently written programs for each pair; adversarial audits
Finite $\chi_{\mathrm{SG}} < 7/2$ and exact balance at the same seventeen points (Chapter 8)	the high-temperature criterion $F(2\beta)^{2d-1} < 12/5$ at every point, with $F(u) = \mathbb{E}\cosh(uW)$ for the coupling variable W of Chapter 8; the criterion and its reduction to $\chi_{\mathrm{SG}} < 7/2$ and exact balance are in the revised version of [214], not yet public, and are not reproduced in Chapter 8; the row is pending until that version is public	exact rational and integer	two programs in different languages; an auditor's recomputation
L^2 non-order windows for $\pm J$ on $\mathbb{Z}^2$ ($T \geq 1.90926$) and $\mathbb{Z}^3$ ($T \geq 2.73214$), free boundary conditions and tori of side at least 11 (Chapter 9)	after the layer-cavity reduction of Chapter 9: the losses of all window types, and the supersolution inequality for the type potential at $t_0 = 0.4806$ ($d = 2$) and 0.3505 ($d = 3$); largest ratio 0.99930010	exact rational; the potential is rounded up from a floating-point iteration and then checked exactly	auditor's re-implementation in floating point (loss patterns to $3 \cdot 10^{-13}$, margins at least $2 \cdot 10^{-4}$)
L^2 window for $\pm J$ on rectangular boxes of $\mathbb{Z}^2$ with both sides at least 8 and fixed coupling-independent boundary conditions, $T \geq 1.93028$ (Chapter 9)	the same inequalities with types that record the distance to the boundary, and absorption at the frozen boundary	exact rational	exhaustive check of all type transitions; exact re-certification; floating-point re-implementation

Table B.1: Certified theorems: blue percolation and high-temperature windows. Except in the second row, the reduction from each theorem to its certificate is proved in the chapter cited or, for the third and fourth rows, in the section of Appendix A to which that chapter defers it.

B.2 Exact finite computations inside proofs

Some propositions are proved by an exact computation on a small graph that is short enough to be checked by hand in principle, but was checked by computer. Computations reported in the review as observations or numerical evidence (for instance, several witnesses in the kill-test suite of Chapter 15, or the Monte Carlo comparison with the XY model) are not certificates in the sense of this appendix and are labelled where they appear.

- *Reflection positivity on the* 2×4 *box* (Chapter 15). At $\beta = \log 2$ every Boltzmann weight is 2^{-10} times a power of 4, so the averaged overlap law is a sum over the 1024 sign configurations of ratios of integers. The reflection form of the witness is computed as the exact rational

$$\mathbb{E}_P[f\,\theta f] = -\frac{467221206243297124026292225185554816 79}{138715107517119843472935230100788281250 00} \approx -0.0033682;$$

the 1024 configurations fall into six classes with the same partition function. An independent derivation through high-temperature boundary polynomials and the eight gauge sectors of the box, without enumerating spins, gives the same fraction and the same six classes.

- *Positive association of the averaged overlap law on the free square* (Chapter 1). For every pair of increasing events the covariance is a rational function of t^2 whose numerator has nonnegative integer coefficients and whose denominator is positive; all numerator coefficients are checked.

- *The six-cycle* (Chapter 15) and the *small counterexamples* of Chapter 1 have proofs by hand; their closed forms were compared with exhaustive enumeration of all sign and spin configurations.

Result	What is checked	Arithmetic	Independent check
Forest domination (FD) on the Wheatstone bridge and on K_4, every pair and all activities (Chapters 11 and 15)	nonnegativity on the unit cube of the polynomial numerator of $P^{\mathrm{AG}}(u \leftrightarrow v) - \mathbb{E}_\varepsilon\langle\sigma_u\sigma_v\rangle^2$	exact integer Bernstein coefficients, with subdivision and blow-ups at the zeros	a second verifier sharing no code (own numerators and Bernstein transforms); negative controls rejected
The same comparison at every homogeneous activity on the free 2×3, 2×4, 2×5, 3×3 and 3×4 grids and the 3×3 torus (Chapter 15)	positivity of the difference as a rational function of t for every pair	exact Bernstein tests and root counting	separate implementation (Walsh–Hadamard transforms over the cycle space)
The flow-weighted influence bound on the Wheatstone bridge between its terminals, all magnitudes (Chapter 11)	after a two-branch reduction (by the sign of the potential drop across the bridge edge), nonnegativity on the unit cube of a polynomial in the five variables $\tanh^2 K_e$	exact Bernstein coefficients on charts (one chart of the first certificate: exact monomial domination)	a second certificate for its own polynomial, sharing no code; reduction checked by hand
Spin-glass order, a giant blue pole cluster and blue imbalance for Gaussian couplings on the diamond lattice with s branches, at $T \le T^*(s)$ (Chapter 12)	for $(s, T^*) = (3, \frac{4}{15})$, $(4, \frac{6}{7})$, $(5, \frac{6}{5})$, $(7, \frac{7}{4})$, $(8, 2)$: the coupling-growth recursion for the laws of effective pole couplings and the per-level conditions	exact rational, with erf, exponentials and logarithms enclosed by interval arithmetic with directed rounding	a published checker with must-fail inputs; a clean-room implementation for $s = 5, 7$
Uniqueness windows for diluted $\pm J$ couplings: $p < \frac{1}{2} + g^*$, $g^* \approx 1.16927 \cdot 10^{-5}$, in $d = 2$ (Chapter 13); $p < p_c(3) + 7.4055 \cdot 10^{-9}$ in $d = 3$ [217, Thm. 1.4]	$d = 2$: the exact inequalities between counts of box configurations, one for each non-crossing partition of the boundary, described in Chapter 13; $d = 3$: a unit-cube certificate, with its reduction in [217]; code and data in [218]	exact rational	two engines for $d = 2$ sharing no code
Overlap order (positive thermodynamic Edwards–Anderson response) and non-uniqueness for biased $\pm J$ on $\mathbb{Z}^2$ with $p \ge 0.9537$ and $\beta \ge 10^6$ (Chapter 14)	a transfer-matrix inequality $\mathcal{M}v \le \frac{9999}{10000}v$ on 453,233,520 rows for the weighted contour transfer with memory 15; largest ratio ≈ 0.998946	exact rational	the generator's exact verifier and a second verifier (shared row-weight definitions) give the same exact ratio [212, 213]

Table B.2: Certified theorems: comparison inequalities, hierarchical order, diluted uniqueness and biased order.

Appendix C

Notation across programmes

The review states every imported result in one notation. This appendix records how the main programmes write the same objects, so that a statement can be checked against its source without re-deriving the translation. The review's own conventions are those of Chapter 1: $H_J(\sigma) = -\sum_{xy} J_{xy}\sigma_x\sigma_y$ with β multiplying H_J; iid couplings with law ν; $\nu_p = p\,\delta_{+1} + (1-p)\,\delta_{-1}$ for biased $\pm J$ (so p is the fraction of ferromagnetic bonds); $\langle\cdot\rangle_J$ for Gibbs averages and $\mathbb{E}$ for the disorder average; replicas σ, τ; site and edge overlaps q_Λ and q^e_Λ (the latter normalized by the number of edges); FK probabilities $p_e = 1 - e^{-2\beta|J_e|}$ on satisfied edges; metastates κ_J; and the Nishimori line $e^{-2\beta_{\mathrm{N}}} = (1-p)/p$.

C.1 Translation tables

Object	As written in the source	In this review
Newman's lecture notes and the frustrated FK papers [170, 171]		
coupling law	μ (and ν for the probability on the disorder space)	ν (and $\mathbb{P}$)
FK edge probability	$1 - e^{-\beta\lvert J_b\rvert}$	$p_e = 1 - e^{-2\beta\lvert J_e\rvert}$ (their β is twice ours)
critical values	$\beta_c \le \bar{\beta}_c$: a.s. uniqueness below, a.s. non-uniqueness above	the same
metastate	$K^\dagger$ (joint limit), κ_J (its conditional law)	κ_J
Newman–Stein, Arguin–Newman–Stein–Wehr, Newman–Read–Stein		
states	Γ drawn from κ_J; pure states α with weights W_α; ρ_J the metastate average	the same symbols; ρ_J is the barycenter
overlaps	$q_{\alpha\alpha'}$ (site), $q^{(e)}_{\alpha\beta}$ (edge, a spatial average)	site and edge overlaps of two states
ground states	GSP, a ground-state pair	ground-state pair
free-energy difference	$F_L(\Gamma, \Gamma')$, with boundary conditions drawn from Γ and Γ'; measure $M = \nu(\mathrm{d}J)\kappa_J \times \kappa'_J$	the same (Chapter 4)
incongruent	edge overlap strictly below the edge self-overlap	disagreement on a positive density of edges
coupling law	continuous and symmetric unless stated	stated in each theorem

Table C.1: Translations for Newman's notes and the metastate literature.

Object	As written in the source	In this review
Chatterjee and the zero-temperature chaos papers [46, 49, 56]		
perturbation	$p \in (0,1)$; OU form $(1-p)J + \sqrt{2p-p^2}\,J'$, i.e. time $t = -\log(1-p)$; or resampling with probability p	Chapter 6 uses both, with p and t as in the source
overlaps	R site overlap on the interior; $Q_{1,2}$ bond overlap	q_Λ, q^e_Λ
sign convention	Chatterjee: $H = -\sum J\sigma\sigma$; Chen–Kim–Sen and Basu Roy Chowdhury–Ganguly maximize $+\sum J\sigma\sigma$	minimize H_J
other symbols	β is also a growth constant of the graph in [49]	β is always inverse temperature
Contucci and coauthors [61, 63]		
overlap	covariance of the Hamiltonian normalized by the number of sites, so the link overlap is divided by $\lvert\Lambda\rvert$	q^e_Λ is divided by $\lvert E(\Lambda)\rvert$, a factor d
quenched state	$\langle\cdot\rangle = \mathrm{Av}\,\Omega(\cdot)$, a finite-volume, disorder-averaged replica measure	$\mathbb{E}[\langle\cdot\rangle_J]$ in finite volume
pressure	$P = \mathrm{Av}\log Z$ with β absorbed into the interaction strengths λ_X	quenched pressure (Chapter 1)

Table C.2: Translations for the chaos and overlap-identity literature.

C.2 The author's papers

The author's papers [214–217, 219], on which Chapters 8, 11 to 13 and 15 draw, depart from the conventions above in two places.

- Migdal–Kadanoff lattices in [219] have scale factor 2 and n branches; the review writes (b, s) as in Chapter 12.
- β is the uniform edge weight, or a uniform bound on the edge weights, of the arboreal gas in [216]; the review writes the arboreal-gas weights as y_e and keeps β for the inverse temperature.

C.3 Framings

The same fact is often stated in different frames, and a translation of symbols is not enough to compare two statements. Five differences recur.

(i) *States, metastates, or joint limits.* Newman–Stein results concern metastates built from coupling-independent boundary conditions; Gandolfi–Newman–Stein work at fixed J with extremal states; the blue-percolation results concern stationary periodic joint limits of (J, σ, τ). The review says which in each statement.

(ii) *Finite or infinite volume.* Chatterjee's theorems and the overlap identities of Contucci and coauthors are finite-volume statements; a metastate statement needs a limit that these papers do not take.

(iii) *Quenched or annealed.* The averaged two-replica law contains Z_J^{-2}; replica calculations at integer n and annealed laws do not, which changes positivity properties (Chapter 15).

(iv) *Spin or inference language.* On the Nishimori line the same identities are stated for spins, for planted signals and for decoders.

Object	As written in the source	In this review
Graphical representations: Machta–Newman–Stein, Gandolfi–Newman–Stein, Gandolfi		
CMR colours	blue with probability $1-e^{-4\beta\lvert J\rvert}$ on doubly satisfied edges, red with $1-e^{-2\beta\lvert J\rvert}$ on singly satisfied ones	the same (Chapter 7)
signed FK	$\pm J$ at fixed J; $p=1-e^{-2\beta}$; companions C^A_∞, C^B_∞; λ the bias of the coupling law	p_e with $\lvert J_e\rvert=1$; the bias is p of ν_p
states	extremal or product-extremal states at fixed J	stated in each theorem
Gauge theory and the Nishimori line [198]		
coupling and bias	$H=-J\sum\tau_{ij}S_iS_j$, $K=\beta J$; $p=P(\tau_{ij}=+1)$	ν_p with the same p
Nishimori line	$K=K_p$, $e^{2K_p}=p/(1-p)$; Gaussian: $J_0/J^2=\beta$	$\beta_{\mathrm{N}}(p)$; Gaussian $N(J_0,\Delta^2)$: $\beta_{\mathrm{N}}=J_0/\Delta^2$
other papers	some authors, and several network-model papers, use p for the antiferromagnetic fraction	replace p by $1-p$
language	spin language, equivalently Bayes-optimal inference ("planted" couplings)	both, identified in Chapter 14
Hierarchical lattices [59, 201, 202, 219]		
lattice parameters	Okuyama–Ohzeki: s parallel branches of b bonds in series, $d=1+\log s/\log b$; Collet–Eckmann and [219]: scale factor 2 and n branches, $d=1+\log_2 n$	(b,s): scale factor b, s branches, as in Chapter 12
effective couplings	Okuyama–Ohzeki: energy-valued $J_n(\beta)$; Collet–Eckmann and [219]: dimensionless ξ_n, with $\tanh\xi=\tanh\xi_1\tanh\xi_2$ in series	dimensionless; $J_n(\beta)=\xi_n/\beta$
order parameter	$\mathbb{E}\langle\sigma_A\sigma_B\rangle^2$ between the poles, $\mathbb{E}\langle Q_n^2\rangle$, or $\mathbb{E}\langle R_{12}^2\rangle$	stated with the boundary condition at the poles
Long-range chains [102, 203]		
couplings	Gandolfi–Newman–Stein: $J_{xy}=c_{y-x}Z_{xy}$ with $\lvert c_r\rvert\sim r^{-a}$; Okuyama–Ohzeki: variance $\propto r^{-\alpha}$	$a=\alpha/2$
states	DLR states for a fixed order of summation; weak uniqueness for coupling-independent boundary conditions	stated in each theorem
The random-field model [8, 26, 68]		
model	ferromagnetic $J=1$ plus field εh_v; boundary influence $\mu^+(\sigma_o=1)-\mu^-(\sigma_o=1)$	Chapter 10
satisfied edges	Dario–Harel–Peled: edges with $\sigma_x=\sigma_y$ (their base Hamiltonian is zero)	$J_{xy}\sigma_x\sigma_y>0$
spin-glass couplings	Berger–Tessler: weights w_e; unsatisfied iff $w_e\sigma_i\sigma_j<0$, on the dual lattice	the same, with $w=J$
Numerical work		
units and exponents	T in units of the coupling scale; θ, also y, the stiffness exponent; d_s or d_f the interface dimension; D for the dimension	T, θ, d_s, d (Chapter 2)
bias	p the ferromagnetic fraction (for example in studies of the $\pm J$ model at $p=0.7$)	p of ν_p

Table C.3: Translations for graphical representations, gauge theory, hierarchical and long-range proxies, the random-field model and numerics.

(v) *Which order.* “Order” may mean a pressure cusp, a stationary overlap moment or a finite-volume overlap moment (the three targets of Chapter 1); results proved for one target are not silently transferred to another.

Appendix D

Cluster-representation formulas

This appendix collects, in the review's notation, the formulas of the graphical representations used in Chapters 7 and 8, with pointers to the statements where they are proved or cited. Throughout, $G = (V, E)$ is a finite graph with free boundary condition, $J \in \mathbb{R}^E$, $\beta \geq 0$, and

$$p_e = 1 - e^{-2\beta|J_e|}, \qquad r_e = e^{-2\beta|J_e|} = 1 - p_e.$$

An edge $e = xy$ is satisfied by σ if $J_e\sigma_x\sigma_y > 0$. For two replicas σ, τ, $q_x = \sigma_x\tau_x$. $k(\omega)$ is the number of components of (V, ω), isolated vertices included; FF_J is the set of J-frustration-free edge sets and $\mathrm{sgn}_J(x, y; \omega)$ the relative parity (Definition 7.1).

D.1 One replica: the frustrated Edwards–Sokal coupling

Object	Formula
joint weight	$\Phi_{G,J}(\sigma,\omega) \propto \prod_e \big[r_e \mathbf{1}\{\omega_e = 0\} + p_e \mathbf{1}\{\omega_e = 1,\ e \text{ satisfied}\}\big]$, normalization $e^{-\beta\sum_e |J_e|} Z^\beta_{G,J}$
spin marginal	$\mu^\beta_{G,J}$
bond marginal	$\varphi_{G,J}(\omega) \propto \prod_e p_e^{\omega_e} r_e^{1-\omega_e}\, 2^{k(\omega)} \mathbf{1}\{\omega \in \mathrm{FF}_J\}$
bonds given spins	independent; $\omega_e \sim \mathrm{Bernoulli}(p_e)$ on satisfied edges, 0 otherwise
spins given bonds	one fair sign per cluster C, $\sigma_y = \sigma_{x_C}\,\mathrm{sgn}_J(x_C, y; \omega)$
two-point function	$\langle\sigma_x\sigma_y\rangle_J = \mathbb{E}_{\varphi_{G,J}}[\mathrm{sgn}_J(x, y; \omega)]$, $\lvert\langle\sigma_x\sigma_y\rangle_J\rvert \leq \varphi_{G,J}(x \leftrightarrow y)$
conditioned form	$\varphi_{G,J} = \varphi_{G,\lvert J\rvert}(\cdot \mid \mathrm{FF}_J)$, $Z^\beta_{G,J} = Z^\beta_{G,\lvert J\rvert}\,\varphi_{G,\lvert J\rvert}(\mathrm{FF}_J)$
conditionally unfrustrated Bernoulli	$\psi_{J,p}(\omega) \propto p^{\lvert\omega\rvert}(1-p)^{\lvert E\rvert - \lvert\omega\rvert}\mathbf{1}\{\omega \in \mathrm{FF}_J\}$ (no factor $2^{k(\omega)}$)

These are recalled in Section 7.2 and Remark 7.2 (proved in Chapter 1), after [170], [2, Thm. 7.2] and [71, (2.1), (2.6)].

Annealed edge law for fair ± 1 *couplings* [101, Thm. 3]. With $\mathbb{P}_G(J, \omega) = 2^{-|E|}\varphi_{G,J}(\omega)$ and $p = 1 - e^{-2\beta}$, the conditional odds of $\{\omega_e = 1\}$ given $\omega_{E\setminus e}$, for fixed couplings off e, are

$$\frac{p}{2(1-p)} \quad \text{if the endpoints of } e \text{ are not connected off } e, \qquad \frac{p/(1-p)}{1 + Z_+/Z_-} \quad \text{otherwise,}$$

where $Z_\pm$ is the normalization with J_e equal to $\pm$ the relative parity of the endpoints; since $Z_+/Z_- \leq 1/(1-p)$, the odds are at least $p/(2-p)$ and the probability at least $p/2$ (Lemma 7.3).

Frustrated Peierls estimate (compare [71, (2.8)]). For a dual circuit image $\gamma \subset \mathbb{Z}^2$, uniformly in the configuration off γ,

$$\varphi_{\Lambda,J}(\gamma \equiv 0 \mid \omega_{\Lambda\setminus\gamma}) \le \Big(\frac{2(1-p)}{2-p}\Big)^{|\gamma|/2}$$

(Theorem 7.6).

D.2 The Nishimori line

For iid couplings with law $\nu_p = p\delta_{+1} + (1-p)\delta_{-1}$ and $\beta = \beta_{\mathrm{N}}(p)$, $e^{-2\beta_{\mathrm{N}}} = (1-p)/p$:

$$\mathbb{P}(J) = (2\cosh\beta)^{-|E|} e^{\beta\sum_e J_e},$$

$$\mathbb{E}g(J) = 2^{-|V|}(2\cosh\beta)^{-|E|}\sum_J g(J)\, Z^{\beta}_{G,J} \qquad (g \text{ gauge invariant}),$$

$$\mathbb{E}\,\varphi_{G,J}(\omega) = (\tanh\beta)^{|\omega|}(1-\tanh\beta)^{|E|-|\omega|}, \qquad \tanh\beta_{\mathrm{N}} = 2p-1,$$

the last also for every boundary condition that does not depend on the random couplings (Theorem 7.7; [99, Thm. 5], [158, §2]).

D.3 Signed FK

For $J \in \{-1,1\}^{E_d}$ and frustration-free ω with at most one infinite cluster C_∞, the signed companions are $C^A_\infty = \{y \in C_\infty : \operatorname{sgn}_J(x_0,y;\omega) = +1\}$ and $C^B_\infty = C_\infty \setminus C^A_\infty$ for a root $x_0 \in C_\infty$. For two replicas, the paired coherent companions are the classes of $\{s(x,y)=+1\}$ on $C_\infty(1)\cap C_\infty(2)$, where $s(x,y) = \operatorname{sgn}_J(x,y;\omega^1)\operatorname{sgn}_J(x,y;\omega^2)$. For extremal states with the hypotheses of Theorems 7.17 and 7.18,

$$|m_\Gamma| = \mathbb{E}|d_\Gamma(C^A_\infty) - d_\Gamma(C^B_\infty)|, \qquad q_\Gamma = \varepsilon_J\,\mathbb{E}[D_\Gamma],$$

[104, Thms. 5.7 and 8.5]. In the union coupling below, $s(x,y) = q_x q_y$ (Observation 7.19).

D.4 Two replicas: double FK and the CMR colouring

Let σ,τ be independent with law $\mu^{\beta}_{G,J}$. Relative to $e^{2\beta|J_e|}$, the two-replica Boltzmann weight of an edge is 1, r_e or r_e^2 according as it is satisfied in two, one or no replicas. The CMR colouring [55], [151, (7)–(13)] splits these weights as follows [215, Def. 1.3]:

	satisfied by both	by σ only	by τ only	by neither
blue	$1-r_e^2$	0	0	0
red	0	$r_e(1-r_e)$	$r_e(1-r_e)$	0
closed	r_e^2	r_e^2	r_e^2	r_e^2
column sum	1	r_e	r_e	r_e^2

So, given the spins, a doubly satisfied edge is blue with probability $1-e^{-4\beta|J_e|}$, a singly satisfied edge is red with probability p_e, and grey means blue or red.

Object	Formula
local constraints	$q_x = q_y$ on blue edges, $q_x = -q_y$ on red edges
cycle constraints	(C1) every grey cycle has an even number of red edges; (C2) every all-blue cycle has $\prod \operatorname{sgn} J_e = +1$
bond marginal	$\propto \prod_e w(c_e)\,\mathbf{1}_{(\mathrm{C1})}\mathbf{1}_{(\mathrm{C2})}\,2^{K(c)}2^{C_{\mathrm{b}}(c)}$, with $w(\text{blue}) = 1 - r_e^2$, $w(\text{red}) = r_e(1-r_e)$, $w(\text{closed}) = r_e^2$, and K, C_{b} the numbers of grey and blue components [215, Rem. 1.4]
overlap two-point function	$\langle\sigma_x\sigma_y\rangle_J^2 = \mathbf{P}_J(x \overset{\text{grey}}{\leftrightarrow} y,\ \text{even red parity}) - \mathbf{P}_J(x \overset{\text{grey}}{\leftrightarrow} y,\ \text{odd red parity})$ [151, (16)]
grey-cluster flip	reversing τ on a grey cluster preserves the weight; reversing both replicas on a blue cluster preserves the weight off blue (Proposition 7.12, Lemma 8.1)
double FK	$\omega^1 \cap \omega^2$; a doubly satisfied edge is doubly occupied with probability p_e^2
union coupling	with ω^1, ω^2 the ES bonds of σ and τ: blue $= (\omega^1 \cup \omega^2) \cap \{\text{doubly satisfied}\}$, red $= (\omega^1 \cup \omega^2) \cap \{\text{singly satisfied}\}$, grey $= \omega^1 \cup \omega^2$ (Observation 7.13)
I-replica rule	$p_{xy} = 1 - \exp\bigl(-2\beta J_{xy} \sum_{i=1}^{I} \sigma_x^{(i)}\sigma_y^{(i)}\bigr)$ when the exponent is positive, 0 otherwise; $I = 2$ is the blue rule [166, (13)]
Houdayer clusters	connected components of $\{q = +1\}$ and of $\{q = -1\}$; each blue cluster lies in one of them [120]

D.5 Two copies and random currents

Gandolfi's two-copy representation [100]. For two independent copies with a common boundary condition, condition on the site sums $\sigma_x + \tau_x$; for ±1 spins the overlap region is $\{q = +1\}$ and the non-overlap region $\{q = -1\}$ carries a symmetric Gibbs field with a random-cluster representation by active bonds. For local events A, B with supports Λ_A, Λ_B, $|\mu(A \cap B) - \mu(A)\mu(B)| \le \mathbb{P}(\Lambda_A \overset{\text{act}}{\leftrightarrow} \Lambda_B)$ [100, Thm. 3.1]; for the $\pm J$ model the active bonds of the non-overlap region are the blue bonds there (Theorem 7.20).

Frustration-adjusted currents [2, Thm. 7.5]. With sourceless currents n_1, n_2 of the ferromagnet $|J|$ and weights $w(n) = \prod_e (\beta|J_e|)^{n_e}/n_e!$,

$$Z^{\beta}_{G,J} / Z^{\beta}_{G,|J|} = \mathbb{P}^{\varnothing,\varnothing}(n_1 + n_2 \text{ is frustration-free}).$$

D.6 Constants of the blue-percolation proofs

For a symmetric law $J_e = \varepsilon_e W_e$ (Section 8.4; [214, §§2, 5 and App. F]):

leaf identity
$\mathbf{P}(e \text{ blue} \mid \text{exterior}, W_e) = \tanh(\beta W_e)$, since $\left(\frac{1+t}{2}\right)^2 \frac{4t}{(1+t)^2} = t$ and $1 - e^{-4a} = \frac{4t}{(1+t)^2}$ for $t = \tanh a$

sign-averaged insertion density
$K_t(r, s; z) = \frac{1}{2}\frac{(1+tr)(1+ts)}{(1+tz)^2} + \frac{1}{2}\frac{(1-tr)(1-ts)}{(1-tz)^2} = \frac{\cosh^2 h}{\cosh^2 a}\cosh(a(r+s) - 2h) \in [\operatorname{sech}^2 a, \cosh 4a]$, $h = \operatorname{artanh}(tz)$

expansion
$K_t = 1 + t^2(rs - 2z(r+s) + 3z^2) + O(t^4)$

averaged constants
$b = \mathbb{E}\tanh\beta W$, $\ell = \mathbb{E}\operatorname{sech}^2 \beta W$, $M = \mathbb{E}\cosh 4\beta W$, $\varrho = \ell/M$

opening bounds (degree Δ)
$b\,\varrho^{\Delta-1} \le \mathbf{P}(B_e = 1 \mid H) \le b\,\varrho^{-(\Delta-1)}$, $\mathbf{P}(q_v = s \mid H) \ge \frac{1}{2}\varrho^{\Delta}$

site parameters on $\mathbb{Z}^d$
$p = b\,\varrho^{2d-1}$ (embedded iid site cluster), $u = b\,\varrho^{-(2d-1)}$ (path bound $2d(2d-1)^{n-1}u^n$)

centred susceptibility
$F(u) = \mathbb{E}\cosh uW$, $A_\beta = \frac{1}{2}(F(4\beta) + F(2\beta))$, $B_\beta = \frac{1}{2}(F(4\beta) - F(2\beta))$, $\kappa_\Delta = \Delta B_\beta A_\beta^{\Delta-1}$; $\kappa_\Delta < 1 \Rightarrow \sup_x \sum_y \mathbb{E}\langle\sigma_x\sigma_y\rangle_J^2 \le (1-\kappa_\Delta)^{-1}$

mixed site–bond comparison
$L_\beta = \frac{1}{2}\mathbb{E}(1 - e^{-4\beta W})$, $R_\beta = F(4\beta)/F(2\beta)$, $a_\Delta = (1 + R_\beta^{\Delta})^{-1}$, $p_\Delta = a_\Delta L_\beta$; for unit couplings $p_\Delta \le t/(1 + 3\Delta t^2) \le (2\sqrt{3\Delta})^{-1}$

Appendix E

Sources of the numerics

Chapter 2 uses a short table of numerical values to calibrate the rigorous programme. This appendix gives, for each entry, the value with the error bar as quoted by the authors, the model and the method, the system sizes where the source states them, and the place in the source where the value appears. All values are heuristic: they come from finite-size scaling of simulations, from exact ground states at finite sizes, from approximate ground-state algorithms, or from field-theoretic estimates, and none is used as an input to a proof. Temperatures are in units of the coupling scale ($|J_e| = 1$ for $\pm J$, unit variance for Gaussian couplings); θ is the stiffness exponent, also written y.

Two remarks help in reading the table. First, the two-dimensional values distinguish continuous from $\pm J$ couplings: for Gaussian couplings $\theta < 0$, so domain walls become cheaper at large scales and $T_c = 0$; for $\pm J$ couplings the domain-wall energy saturates, the zero-temperature order is entropic, and ground states are degenerate, so zero-temperature theorems for continuous laws do not transfer to $\pm J$. Second, the three-dimensional critical temperatures from different groups agree within their error bars for $\pm J$ couplings (1.1019(29) and 1.109(10)), while the exponents carry larger systematic uncertainties from corrections to scaling, which each source discusses.

Other numerical statements. The review quotes a few further numerical facts, each labelled as heuristic where it appears: the Monte Carlo comparison with the XY model in Chapters 11 and 15 and the floating-point screens of Chapter 8. The evidence on the nature of the low-temperature phase, which is discussed in the text of Chapter 2 rather than tabulated, is summarized with its sources at the end of Chapter 4.

Quantity	Value as quoted	Model, method, sizes	Source, locator
$d = 2$, Gaussian	$\theta = -0.2793(3)$; domain-wall dimension $d_f = 1.27319(9)$; for $\pm J$, $d_f = 1.279(2)$	exact ground states by mappings to matching, square lattices up to $10\,000 \times 10\,000$ spins; a windowing method for fully periodic samples; finite-size scaling	[139], abstract
$d = 2$, Gaussian	$\theta = -0.266(2)$ and $\theta = -0.282(2)$ for two types of boundary-condition change	exact ground states by matching, up to $N = 480^2$ spins	[116], abstract
$d = 2$, $\pm J$	domain-wall energy saturates, $\theta = 0$	as above	[116], abstract
$d = 2$, $\pm J$	entropic stiffness exponent $\theta_S = 0.50 \pm 0.01$; order at $T = 0$ only	ground-state entropy of domain walls, power-law fit for $L = 32$ to 256	[233], p. 2
$d = 3$, $\pm J$	$T_c = 1.1019(29)$, $\nu = 2.562(42)$, $\eta = -0.3900(36)$, correction exponent $\omega = 1.12(10)$	Monte Carlo on the Janus computer, lattices thermalized up to $L = 40$; finite-size scaling with leading corrections	[127], abstract
$d = 3$, $\pm J$	$\beta_c = 0.902(8)$, $T_c = 1.109(10)$; $\nu = 2.45(15)$, $\eta = -0.375(10)$	Monte Carlo up to $L = 28$ (and $L = 20$ at ferromagnetic fraction 0.7); quartic cumulants and ξ/L	[118], Sect. VI A and abstract
$d = 3$, Gaussian	$T_c = 0.951(9)$, $\nu = 2.44(9)$, $\eta = -0.37(5)$	parallel-tempering Monte Carlo, $L = 3$ to 16 for Gaussian couplings	[135], Table I
$d = 4$, $\pm J$	$T_c = 2.03 \pm 0.03$, $\nu = 1.00 \pm 0.10$	parallel-tempering Monte Carlo up to volume 10^4 ($L = 10$)	[155], p. 5
Stiffness, $d = 3, \ldots, 6$	$y_3 = 0.24(1)$, $y_4 = 0.61(2)$, $y_5 = 0.88(5)$, $y_6 = 1.1(1)$	defect-energy widths on bond-diluted lattices near the zero-temperature glass transition, exact graph reduction, $L \leq 30$	[31], abstract
Lower critical dimension	$d_l = 2.5$; "exactly" 5/2; "supports" 2.5	domain-wall free energy in a field-theoretic calculation; stiffness exponents in all dimensions combined with theory; free-energy barriers in $d = 3, 4$	[32, 93, 154], abstracts
Upper critical dimension, $h = 0$	$d_u = 6$	renormalization group	[115] (not opened for this review)
$d = 3$, replicon exponent	$\theta(q = 0) = 0.38(2)$ (droplet picture: 0 asymptotically)	recalled from an earlier non-equilibrium computation of the same collaboration	[125], eq. (19)
Interface dimension, $d = 2, \ldots, 6$	$d_s = 1.2529(14)$, 2.5256(30), 3.7358(36), 4.884(60), 5.9899(60) (strong-disorder RG); 1.2196(11), 2.4962(19), 3.7190(47), 4.9068(32), 6.0023(22) (greedy algorithm)	periodic against antiperiodic domain walls; approximate ground-state algorithms	[243], Table II

Table E.1: The numerical values of Chapter 2 with their sources. Error bars are those quoted by the authors.

Bibliography

[1] D. Ahlberg, M. Deijfen, and M. Sfragara. Chaos, concentration and multiple valleys in first-passage percolation. *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques*, 62:919–940, 2026. doi: 10.1214/24-AIHP1542.

[2] M. Aizenman. Geometric analysis of Ising models, Part III. *Preprint*, 2025. arXiv:2509.02850v3.

[3] M. Aizenman and P. Contucci. On the stability of the quenched state in mean-field spin-glass models. *Journal of Statistical Physics*, 92:765–783, 1998.

[4] M. Aizenman and H. Duminil-Copin. Marginal triviality of the scaling limits of critical 4D Ising and $\lambda\phi_4^4$ models. *Ann. of Math.*, 194:163–235, 2021. doi: 10.4007/annals.2021.194.1.3.

[5] M. Aizenman and G. Grimmett. Strict monotonicity for critical points in percolation and ferromagnetic models. *J. Stat. Phys.*, 63(5–6):817–835, 1991. doi: 10.1007/BF01029985.

[6] M. Aizenman and R. Peled. A power-law upper bound on the correlations in the 2D random field Ising model. *Commun. Math. Phys.*, 372:865–892, 2019. doi: 10.1007/s00220-019-03450-3.

[7] M. Aizenman and S. Warzel. Kac–Ward formula and its extension to order-disorder correlators through a graph zeta function. *J. Stat. Phys.*, 173:1755–1778, 2018. doi: 10.1007/s10955-018-2184-9. Locators refer to arXiv v4.

[8] M. Aizenman and J. Wehr. Rounding effects of quenched randomness on first-order phase transitions. *Communications in Mathematical Physics*, 130:489–528, 1990. doi: 10.1007/BF02096933.

[9] M. Aizenman, J. T. Chayes, L. Chayes, and C. M. Newman. The phase boundary in dilute and random Ising and Potts ferromagnets. *J. Phys. A: Math. Gen.*, 20(5):L313–L318, 1987. doi: 10.1088/0305-4470/20/5/010.

[10] M. Aizenman, M. Harel, and R. Peled. Exponential decay of correlations in the 2D random field Ising model. *J. Stat. Phys.*, 180:304–331, 2020. doi: 10.1007/s10955-019-02401-5.

[11] L.-P. Arguin and M. Damron. On the number of ground states of the Edwards–Anderson spin glass model. *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques*, 50(1):28–62, 2014. doi: 10.1214/12-AIHP499.

[12] L.-P. Arguin and J. Hanson. On absence of disorder chaos for spin glasses on $\mathbb{Z}^d$. *Electronic Communications in Probability*, 25:1–12, 2020. doi: 10.1214/20-ECP311.

[13] 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. *Communications in Mathematical Physics*, 300(3):641–657, 2010. doi: 10.1007/s00220-010-1130-8.

[14] L.-P. Arguin, C. M. Newman, D. L. Stein, and J. Wehr. Fluctuation bounds for interface free energies in spin glasses. *Journal of Statistical Physics*, 156:221–238, 2014.

[15] L.-P. Arguin, C. M. Newman, and D. L. Stein. Thermodynamic identities and symmetry breaking in short-range spin glasses. *Physical Review Letters*, 115:187202, 2015.

[16] L.-P. Arguin, C. M. Newman, D. L. Stein, and J. Wehr. Zero-temperature fluctuations in short-range spin glasses. *Journal of Statistical Physics*, 163:1069–1078, 2016. doi: 10.1007/s10955-016-1516-x.

[17] L.-P. Arguin, C. M. Newman, and D. L. Stein. A relation between disorder chaos and incongruent states in spin glasses on $\mathbb{Z}^d$. *Comm. Math. Phys.*, 367:1019–1043, 2019. doi: 10.1007/s00220-019-03418-3.

[18] A. Auffinger, G. Ben Arous, and J. Černý. Random matrices and complexity of spin glasses. *Comm. Pure Appl. Math.*, 66:165–201, 2013. doi: 10.1002/cpa.21422.

[19] J. E. Avron, G. Roepstorff, and L. S. Schulman. Ground state degeneracy and ferromagnetism in a spin glass. *J. Stat. Phys.*, 26:25–36, 1981. doi: 10.1007/BF01106783.

[20] P. Balister, B. Bollobás, and O. Riordan. Essential enhancements revisited. *Preprint*, 2014. arXiv:1402.0834.

[21] F. Barahona. On the computational complexity of Ising spin glass models. *Journal of Physics A: Mathematical and General*, 15:3241–3253, 1982. doi: 10.1088/0305-4470/15/10/028.

[22] M. Basu Roy Chowdhury and S. Ganguly. Spectral properties of the zero temperature Edwards–Anderson model, 2025. arXiv:2507.10507v1.

[23] R. Bauerschmidt, N. Crawford, T. Helmuth, and A. Swan. Random spanning forests and hyperbolic symmetry. *Comm. Math. Phys.*, 381(3):1223–1261, 2021. doi: 10.1007/s00220-020-03921-y. arXiv:1912.04854; theorem numbers refer to the arXiv version.

[24] V. Beffara and H. Duminil-Copin. The self-dual point of the two-dimensional random-cluster model is critical for $q \geq 1$. *Probab. Theory Related Fields*, 153:511–542, 2012. doi: 10.1007/s00440-011-0353-8.

[25] I. Benjamini, G. Kalai, and O. Schramm. Noise sensitivity of Boolean functions and applications to percolation. *Publications Mathématiques de l'IHÉS*, 90:5–43, 1999. doi: 10.1007/BF02698830.

[26] N. Berger and R. J. Tessler. No percolation in low temperature spin glass. *Electronic Journal of Probability*, 22:1–19, 2017. doi: 10.1214/17-EJP103.

[27] A. Berretti. Some properties of random Ising models. *J. Stat. Phys.*, 38:483–496, 1985. doi: 10.1007/BF01010473.

[28] L. Bieche, J. P. Uhry, R. Maynard, and R. Rammal. On the ground states of the frustration model of a spin glass by a matching method of graph theory. *Journal of Physics A: Mathematical and General*, 13: 2553–2576, 1980. doi: 10.1088/0305-4470/13/8/005.

[29] A. Billoire, L. A. Fernandez, A. Maiorano, E. Marinari, V. Martin-Mayor, J. Moreno-Gordo, G. Parisi, F. Ricci-Tersenghi, and J. J. Ruiz-Lorenzo. Numerical construction of the Aizenman–Wehr metastate. *Phys. Rev. Lett.*, 119:037203, 2017. doi: 10.1103/PhysRevLett.119.037203.

[30] P. M. Bleher, J. Ruiz, and V. A. Zagrebnov. On the purity of the limiting Gibbs state for the Ising model on the Bethe lattice. *J. Stat. Phys.*, 79(1–2):473–482, 1995. doi: 10.1007/BF02179399.

[31] S. Boettcher. Stiffness exponents for lattice spin glasses in dimensions $d = 3, \ldots, 6$. *Eur. Phys. J. B*, 38: 83–91, 2004. doi: 10.1140/epjb/e2004-00102-5.

[32] S. Boettcher. Stiffness of the Edwards–Anderson model in all dimensions. *Phys. Rev. Lett.*, 95:197205, 2005. doi: 10.1103/PhysRevLett.95.197205.

[33] A. Bovier. Statistical mechanics of disordered systems, 2001. Lecture notes, MaPhySto, University of Copenhagen.

[34] A. Bovier. *Statistical Mechanics of Disordered Systems*. Cambridge University Press, 2006. doi: 10.1017/CBO9780511616808.

[35] A. J. Bray and M. A. Moore. Lower critical dimension of Ising spin glasses: a numerical study. *J. Phys. C: Solid State Phys.*, 17(18):L463–L468, 1984. doi: 10.1088/0022-3719/17/18/004.

[36] A. J. Bray and M. A. Moore. Chaotic nature of the spin-glass phase. *Phys. Rev. Lett.*, 58(1):57–60, 1987. doi: 10.1103/PhysRevLett.58.57.

[37] J. Bricmont and A. Kupiainen. Phase transition in the 3d random field Ising model. *Commun. Math. Phys.*, 116:539–572, 1988. doi: 10.1007/BF01224901.

[38] R. M. Burton and M. Keane. Density and uniqueness in percolation. *Communications in Mathematical Physics*, 121:501–505, 1989. doi: 10.1007/BF01217735.

[39] F. Camia and D. L. Stein. Probability theory in statistical physics, percolation, and other random topics: the work of C. Newman. Preprint, 2020. arXiv:2002.00342.

[40] M. Campanino, E. Olivieri, and A. C. D. van Enter. One dimensional spin glasses with potential decay $1/r^{1+\epsilon}$. Absence of phase transitions and cluster properties. *Commun. Math. Phys.*, 108:241–255, 1987. doi: 10.1007/BF01210614.

[41] S. Caracciolo, J. L. Jacobsen, H. Saleur, A. D. Sokal, and A. Sportiello. Fermionic field theory for trees and forests. *Phys. Rev. Lett.*, 93:080601, 2004. doi: 10.1103/PhysRevLett.93.080601.

[42] M. Castellana, A. Barra, and F. Guerra. Free-energy bounds for hierarchical spin models. *J. Stat. Phys.*, 155(2):211, 2014. doi: 10.1007/s10955-014-0951-9.

[43] T. Castellani and A. Cavagna. Spin-glass theory for pedestrians. *J. Stat. Mech.*, 2005:P05012, 2005. doi: 10.1088/1742-5468/2005/05/P05012.

[44] J. T. Chalker. On the lower critical dimensionality of the Ising model in a random field. *J. Phys. C: Solid State Phys.*, 16:6615–6622, 1983. doi: 10.1088/0022-3719/16/34/011.

[45] P. Charbonneau, E. Marinari, M. Mézard, G. Parisi, F. Ricci-Tersenghi, G. Sicuro, and F. Zamponi, editors. *Spin Glass Theory and Far Beyond: Replica Symmetry Breaking after 40 Years*. World Scientific, 2023. doi: 10.1142/13341.

[46] S. Chatterjee. Disorder chaos and multiple valleys in spin glasses, 2009. arXiv:0907.3381v4.

[47] S. Chatterjee. The Ghirlanda–Guerra identities without averaging, 2009. arXiv:0911.4520.

[48] S. Chatterjee. *Superconcentration and Related Topics*. Springer Monographs in Mathematics. Springer International Publishing, 2014. doi: 10.1007/978-3-319-03886-5.

[49] S. Chatterjee. Spin glass phase at zero temperature in the Edwards–Anderson model, 2023. arXiv:2301.04112v6 (2 May 2026); cited in version 6.

[50] S. Chatterjee and S. Ray. On the stability of solutions to random optimization problems under small perturbations, 2024. arXiv:2410.21513v1.

[51] J. T. Chayes, L. Chayes, and J. Fröhlich. The low-temperature behavior of disordered magnets. *Commun. Math. Phys.*, 100:399–437, 1985. doi: 10.1007/BF01206137.

[52] J. T. Chayes, L. Chayes, J. P. Sethna, and D. J. Thouless. A mean field spin glass with short-range interactions. *Commun. Math. Phys.*, 106:41–89, 1986. doi: 10.1007/BF01210926.

[53] L. Chayes and J. Machta. Graphical representations and cluster algorithms I. Discrete spin systems. *Physica A*, 239:542–601, 1997. doi: 10.1016/S0378-4371(96)00438-4.

[54] L. Chayes and J. Machta. Graphical representations and cluster algorithms II. *Physica A*, 254:477–516, 1998. doi: 10.1016/S0378-4371(97)00637-7.

[55] L. Chayes, J. Machta, and O. Redner. Graphical representations for Ising systems in external fields. *J. Stat. Phys.*, 93:17–32, 1998. doi: 10.1023/B:JOSS.0000026726.43558.80. arXiv:cond-mat/9806312.

[56] W.-K. Chen, H. Kim, and A. Sen. Disorder chaos in short-range, diluted, and Lévy spin glasses. *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques*, 62:2199–2222, 2026. doi: 10.1214/25-AIHP1583.

[57] C. Chilin, E. Marinari, V. Martín-Mayor, G. Parisi, J. J. Ruiz-Lorenzo, and D. Yllanes. On the true low-energy excitations of the three-dimensional spin glass, 2026.

[58] P. D. Coddington and L. Han. Generalized cluster algorithms for frustrated spin models. *Phys. Rev. B*, 50: 3058–3067, 1994. doi: 10.1103/PhysRevB.50.3058.

[59] P. Collet and J.-P. Eckmann. A spin-glass model with random couplings. *Commun. Math. Phys.*, 93:379–406, 1984. doi: 10.1007/BF01258536.

[60] P. Contucci and C. Giardinà. Spin-glass stochastic stability: a rigorous proof. *Annales Henri Poincaré*, 6: 915–923, 2005.

[61] P. Contucci and C. Giardinà. The Ghirlanda–Guerra identities. *Journal of Statistical Physics*, 126:917–931, 2007.

[62] P. Contucci and C. Giardinà. *Perspectives on Spin Glasses*. Cambridge University Press, 2013. doi: 10.1017/CBO9781139049306.

[63] P. Contucci and J. Lebowitz. Correlation inequalities for spin glasses. *Ann. Henri Poincaré*, 8:1461–1467, 2007. doi: 10.1007/s00023-007-0342-8.

[64] P. Contucci, F. Unguendoli, and C. Vernia. Lack of monotonicity in spin glass correlation functions. *J. Phys. A: Math. Theor.*, 41(38):385001, 2008. doi: 10.1088/1751-8113/41/38/385001.

[65] P. Contucci, E. Mingione, and S. Starr. Factorization properties in d-dimensional spin glasses. Rigorous results and some perspectives. *Journal of Statistical Physics*, 151:809–829, 2013.

[66] O. Couronné. New upper bound for the connective constant for square-lattice self-avoiding walks. *J. Stat. Phys.*, 192:153, 2025. doi: 10.1007/s10955-025-03542-6.

[67] J. T. Cox, A. Gandolfi, P. S. Griffin, and H. Kesten. Greedy lattice animals I: upper bounds. *Ann. Appl. Probab.*, 3:1151–1169, 1993. doi: 10.1214/aoap/1177005277.

[68] P. Dario, M. Harel, and R. Peled. Quantitative disorder effects in low-dimensional spin systems, 2024. arXiv:2101.01711v3.

[69] J. R. L. de Almeida and D. J. Thouless. Stability of the Sherrington–Kirkpatrick solution of a spin glass model. *J. Phys. A: Math. Gen.*, 11(5):983–990, 1978. doi: 10.1088/0305-4470/11/5/028.

[70] E. De Santis. Strict inequality for phase transition between ferromagnetic and frustrated systems. *Electron. J. Probab.*, 6:1–27, 2001. doi: 10.1214/EJP.v6-79.

[71] E. De Santis and A. Gandolfi. Bond percolation in frustrated systems. *Ann. Probab.*, 27:1781–1808, 1999. doi: 10.1214/aop/1022874815.

[72] B. Dembin and C. Garban. Superconcentration for minimal surfaces in first passage percolation and disordered Ising ferromagnets. *Probability Theory and Related Fields*, 190:675–702, 2024. doi: 10.1007/s00440-023-01252-2.

[73] E. Dennis, A. Kitaev, A. Landahl, and J. Preskill. Topological quantum memory. *J. Math. Phys.*, 43: 4452–4505, 2002. doi: 10.1063/1.1499754.

[74] J. Ding and M. Wirth. Correlation length of the two-dimensional random field Ising model via greedy lattice animal. *Duke Math. J.*, 172:1781–1811, 2023.

[75] J. Ding and J. Xia. Exponential decay of correlations in the two-dimensional random field Ising model. *Inventiones mathematicae*, 224:999–1045, 2021. doi: 10.1007/s00222-020-01024-y.

[76] J. Ding and Z. Zhuang. Long range order for random field Ising and Potts models, 2021. arXiv:2110.04531.

[77] J. Ding, Y. Liu, and A. Xia. Long range order for three-dimensional random field Ising model throughout the entire low temperature regime, 2022. arXiv:2209.13998.

[78] H. Duminil-Copin, A. Raoufi, and V. Tassion. Sharp phase transition for the random-cluster and Potts models via decision trees. *Ann. of Math.*, 189:75–99, 2019. doi: 10.4007/annals.2019.189.1.2.

[79] F. J. Dyson. Existence of a phase-transition in a one-dimensional Ising ferromagnet. *Commun. Math. Phys.*, 12:91–107, 1969. doi: 10.1007/BF01645907.

[80] R. G. Edwards and A. D. Sokal. Generalization of the Fortuin–Kasteleyn–Swendsen–Wang representation and Monte Carlo algorithm. *Phys. Rev. D*, 38(6):2009–2012, 1988. doi: 10.1103/PhysRevD.38.2009.

[81] S. F. Edwards and P. W. Anderson. Theory of spin glasses. *J. Phys. F: Metal Phys.*, 5(5):965–974, 1975. doi: 10.1088/0305-4608/5/5/017.

[82] B. Efron and C. Stein. The jackknife estimate of variance. *Ann. Statist.*, 9(3):586–596, 1981. doi: 10.1214/aos/1176345462.

[83] C. Efthymiou and K. Zampetakis. On sampling diluted spin-glasses with unbounded interactions, 2026.

[84] W. Evans, C. Kenyon, Y. Peres, and L. J. Schulman. Broadcasting on trees and the Ising model. *Ann. Appl. Probab.*, 10(2):410–433, 2000. doi: 10.1214/aoap/1019487349.

[85] L. A. Fernandez, V. Martin-Mayor, G. Parisi, and B. Seoane. Temperature chaos in 3D Ising spin glasses is driven by rare events. *EPL*, 103:67003, 2013. doi: 10.1209/0295-5075/103/67003.

[86] D. S. Fisher and D. A. Huse. Ordered phase of short-range Ising spin-glasses. *Phys. Rev. Lett.*, 56(15): 1601–1604, 1986. doi: 10.1103/PhysRevLett.56.1601.

[87] D. S. Fisher and D. A. Huse. Absence of many states in realistic spin glasses. *J. Phys. A: Math. Gen.*, 20 (15):L1005–L1010, 1987. doi: 10.1088/0305-4470/20/15/013.

[88] D. S. Fisher and D. A. Huse. Equilibrium behavior of the spin-glass ordered phase. *Phys. Rev. B*, 38(1): 386–411, 1988. doi: 10.1103/PhysRevB.38.386.

[89] D. S. Fisher, J. Fröhlich, and T. Spencer. The Ising model in a random magnetic field. *J. Stat. Phys.*, 34: 863–870, 1984. doi: 10.1007/BF01009445.

[90] C. M. Fortuin and P. W. Kasteleyn. On the random-cluster model. I. Introduction and relation to other models. *Physica*, 57(4):536–564, 1972. doi: 10.1016/0031-8914(72)90045-6.

[91] S. Franz and F. L. Toninelli. Finite-range spin glasses in the Kac limit: free energy and local observables. *J. Phys. A: Math. Gen.*, 37:7433–7446, 2004. doi: 10.1088/0305-4470/37/30/003.

[92] S. Franz and F. L. Toninelli. The Kac limit for finite-range spin glasses. *Phys. Rev. Lett.*, 92:030602, 2004. doi: 10.1103/PhysRevLett.92.030602.

[93] S. Franz, G. Parisi, and M. A. Virasoro. Interfaces and lower critical dimension in a spin glass model. *J. Phys. I France*, 4(11):1657–1667, 1994. doi: 10.1051/jp1:1994213.

[94] G. Franzese. Cluster analysis for percolation on two-dimensional fully frustrated system, 1996. arXiv:cond-mat/9606003.

[95] S. Friedli and Y. Velenik. *Statistical Mechanics of Lattice Systems: A Concrete Mathematical Introduction*. Cambridge University Press, 2017. doi: 10.1017/9781316882603.

[96] J. Fröhlich and J. Z. Imbrie. Improved perturbation expansion for disordered systems: beating Griffiths singularities. *Commun. Math. Phys.*, 96:145–180, 1984. doi: 10.1007/BF01240218.

[97] J. Fröhlich and B. Zegarlinski. The high-temperature phase of long-range spin glasses. *Commun. Math. Phys.*, 110:121–155, 1987. doi: 10.1007/BF01209020.

[98] D. Gamarnik. The overlap gap property: A topological barrier to optimizing over random structures. *Proc. Natl. Acad. Sci. USA*, 118:e2108492118, 2021. doi: 10.1073/pnas.2108492118.

[99] A. Gandolfi. A remark on gauge symmetries in Ising spin glasses. *Probab. Theory Related Fields*, 114: 419–430, 1999. doi: 10.1007/s004400050230.

[100] A. Gandolfi. Random-cluster correlation inequalities for Gibbs fields. *J. Stat. Phys.*, 173:249–267, 2018. doi: 10.1007/s10955-018-2130-x. arXiv:1804.04247.

[101] A. Gandolfi, M. S. Keane, and C. M. Newman. Uniqueness of the infinite component in a random graph

with applications to percolation and spin glasses. *Probab. Theory Related Fields*, 92:511–527, 1992. doi: 10.1007/BF01274266.

[102] A. Gandolfi, C. M. Newman, and D. L. Stein. Exotic states in long-range spin glasses. *Commun. Math. Phys.*, 157:371–387, 1993. doi: 10.1007/BF02099766. Locators refer to the Berkeley Statistics Technical Report No. 374 (1993).

[103] A. Gandolfi, C. M. Newman, and D. L. Stein. Zero-temperature dynamics of $\pm J$ spin glasses and related models. *Communications in Mathematical Physics*, 214:373–387, 2000. doi: 10.1007/PL00005535.

[104] A. Gandolfi, C. M. Newman, and D. L. Stein. Signed FK percolation in Ising spin glasses. *Preprint*, 2026. arXiv:2610.24629.

[105] C. Garban, G. Pete, and O. Schramm. The Fourier spectrum of critical percolation. *Acta Mathematica*, 205: 19–104, 2010. doi: 10.1007/s11511-010-0051-x.

[106] H.-O. Georgii. *Gibbs Measures and Phase Transitions*. de Gruyter, 1988. doi: 10.1515/9783110850147.

[107] S. Ghirlanda and F. Guerra. General properties of overlap probability distributions in disordered spin systems. Towards Parisi ultrametricity. *Journal of Physics A*, 31(46):9149–9155, 1998.

[108] R. B. Griffiths. Nonanalytic behavior above the critical point in a random Ising ferromagnet. *Phys. Rev. Lett.*, 23:17–19, 1969. doi: 10.1103/PhysRevLett.23.17.

[109] G. Grimmett. *The Random-Cluster Model*. Grundlehren der mathematischen Wissenschaften. Springer, Berlin, 2006. doi: 10.1007/978-3-540-32891-9.

[110] G. R. Grimmett and A. M. Stacey. Critical probabilities for site and bond percolation models. *Ann. Probab.*, 26:1788–1812, 1998. doi: 10.1214/aop/1022855883.

[111] F. Guerra. Broken replica symmetry bounds in the mean field spin glass model. *Comm. Math. Phys.*, 233 (1):1–12, 2003. doi: 10.1007/s00220-002-0773-5.

[112] F. Guerra and F. L. Toninelli. The high temperature region of the Viana–Bray diluted spin glass model. *J. Stat. Phys.*, 115:531–555, 2004. doi: 10.1023/B:JOSS.0000019815.11115.54.

[113] N. Halberstam and T. Hutchcroft. Uniqueness of the infinite tree in low-dimensional random forests. *Probab. Math. Phys.*, 5:1185–1216, 2024. arXiv:2302.12224.

[114] F. Hamze, D. C. Jacob, A. J. Ochoa, D. Perera, W. Wang, and H. G. Katzgraber. From near to eternity: spin-glass planting, tiling puzzles, and constraint-satisfaction problems. *Phys. Rev. E*, 97:043303, 2018. doi: 10.1103/PhysRevE.97.043303.

[115] A. B. Harris, T. C. Lubensky, and J.-H. Chen. Critical properties of spin-glasses. *Phys. Rev. Lett.*, 36(8): 415–418, 1976. doi: 10.1103/PhysRevLett.36.415.

[116] A. K. Hartmann and A. P. Young. Lower critical dimension of Ising spin glasses. *Phys. Rev. B*, 64:180404, 2001. doi: 10.1103/PhysRevB.64.180404.

[117] M. Hasenbusch, F. Parisen Toldin, A. Pelissetto, and E. Vicari. Multicritical Nishimori point in the phase diagram of the $\pm J$ Ising model on a square lattice. *Phys. Rev. E*, 77:051115, 2008. doi: 10.1103/PhysRevE.77.051115.

[118] M. Hasenbusch, A. Pelissetto, and E. Vicari. The critical behavior of three-dimensional Ising spin glass models. *Phys. Rev. B*, 78:214205, 2008. doi: 10.1103/PhysRevB.78.214205.

[119] T. Horiguchi and T. Morita. Existence of the ferromagnetic phase in competing random-bond Ising models on the square lattice for several types of probability distribution of exchange interaction. *J. Phys. A: Math. Gen.*, 15:3551–3560, 1982. doi: 10.1088/0305-4470/15/11/033.

[120] J. Houdayer. A cluster Monte Carlo algorithm for 2-dimensional spin glasses. *Eur. Phys. J. B*, 22:479, 2001.

[121] D. A. Huse and D. S. Fisher. Pure states in spin glasses. *J. Phys. A: Math. Gen.*, 20(15):L997–L1003, 1987. doi: 10.1088/0305-4470/20/15/012.

[122] J. Z. Imbrie. The ground state of the three-dimensional random-field Ising model. *Commun. Math. Phys.*, 98:145–176, 1985. doi: 10.1007/BF01220505.

[123] Y. Imry and S.-K. Ma. Random-field instability of the ordered state of continuous symmetry. *Physical Review Letters*, 35(21):1399, 1975.

[124] T. S. Jackson and N. Read. Theory of minimum spanning trees I: mean-field theory and strongly disordered spin-glass model. *Phys. Rev. E*, 81:021130, 2010. doi: 10.1103/PhysRevE.81.021130.

[125] Janus Collaboration, R. Alvarez Baños, A. Cruz, L. A. Fernandez, et al. Nature of the spin-glass phase at experimental length scales. *J. Stat. Mech.*, page P06026, 2010. doi: 10.1088/1742-5468/2010/06/P06026.

[126] Janus Collaboration, R. A. Baños, A. Cruz, L. A. Fernandez, et al. Thermodynamic glass transition

in a spin glass without time-reversal symmetry. *Proc. Natl. Acad. Sci. USA*, 109:6452–6456, 2012. doi: 10.1073/pnas.1203295109.

[127] Janus Collaboration, M. Baity-Jesi, R. A. Baños, A. Cruz, L. A. Fernandez, et al. Critical parameters of the three-dimensional Ising spin glass. *Phys. Rev. B*, 88:224416, 2013. doi: 10.1103/PhysRevB.88.224416.

[128] T. Jörg. Cluster Monte Carlo algorithms for diluted spin glasses. *Prog. Theor. Phys. Suppl.*, 157:349, 2005. arXiv:cond-mat/0410328.

[129] T. Jörg and F. Ricci-Tersenghi. Entropic effects in the very low temperature regime of diluted Ising spin glasses with discrete couplings. *Phys. Rev. Lett.*, 100(17):177203, 2008. doi: 10.1103/PhysRevLett.100.177203. arXiv:0707.0480.

[130] D. Kandel and E. Domany. General cluster Monte Carlo dynamics. *Phys. Rev. B*, 43:8539–8548, 1991. doi: 10.1103/PhysRevB.43.8539.

[131] D. Kandel, R. Ben-Av, and E. Domany. Cluster dynamics for fully frustrated systems. *Phys. Rev. Lett.*, 65: 941–944, 1990. doi: 10.1103/PhysRevLett.65.941.

[132] Y. Kasai and A. Okiji. Percolation problem describing $\pm J$ Ising spin glass system. *Progr. Theoret. Phys.*, 79 (5):1080–1094, 1988. doi: 10.1143/PTP.79.1080.

[133] Y. Kasai, S. Miyazima, and I. Syozi. Ising lattices with random arrangements of ferromagnetic and antiferromagnetic bonds. *Prog. Theor. Phys.*, 42(1):1–8, 1969. doi: 10.1143/PTP.42.1.

[134] H. G. Katzgraber and F. Krzakala. Temperature and disorder chaos in three-dimensional Ising spin glasses. *Phys. Rev. Lett.*, 98:017201, 2007. doi: 10.1103/PhysRevLett.98.017201.

[135] H. G. Katzgraber, M. Körner, and A. P. Young. Universality in three-dimensional Ising spin glasses: a Monte Carlo study. *Phys. Rev. B*, 73:224432, 2006. doi: 10.1103/PhysRevB.73.224432.

[136] S. Keiser and J. Machta. High resolution study of the 2D ANNNI model using a two-replica cluster algorithm and population annealing. *Preprint*, 2026. arXiv:2606.00882.

[137] H. Kesten. The critical probability of bond percolation on the square lattice equals $\frac{1}{2}$. *Comm. Math. Phys.*, 74(1):41–59, 1980. doi: 10.1007/BF01197577.

[138] H. Kesten. Asymptotics in high dimensions for percolation. In G. Grimmett and D. Welsh, editors, *Disorder in Physical Systems*, pages 219–240. Oxford University Press, 1990.

[139] H. Khoshbakht and M. Weigel. Domain-wall excitations in the two-dimensional Ising spin glass. *Phys. Rev. B*, 97:064410, 2018. doi: 10.1103/PhysRevB.97.064410.

[140] H. Kitatani. Griffiths inequalities for Ising spin glasses on the Nishimori line. *J. Phys. Soc. Jpn.*, 78:044714, 2009. doi: 10.1143/JPSJ.78.044714.

[141] H. Kitatani, H. Nishimori, and A. Aoki. Inequalities for the local energy of random Ising models. *J. Phys. Soc. Japan*, 76:074711, 2007. doi: 10.1143/JPSJ.76.074711.

[142] F. R. Klausen and N. Kravitz. Half-plane non-coexistence without FKG, 2026. arXiv:2602.12261.

[143] G. Kotliar, P. W. Anderson, and D. L. Stein. One-dimensional spin-glass model with long-range random interactions. *Phys. Rev. B*, 27:602–605, 1983. doi: 10.1103/PhysRevB.27.602.

[144] K. Krishnan and G. Ray. Uniqueness and CLT for the ground state of the disordered monomer-dimer model on $\mathbb{Z}^d$. *International Mathematics Research Notices*, 2025:rnaf169, 2025. doi: 10.1093/imrn/rnaf169.

[145] F. Krzakala and O. C. Martin. Spin and link overlaps in 3-dimensional spin glasses. *Phys. Rev. Lett.*, 85: 3013, 2000. doi: 10.1103/PhysRevLett.85.3013.

[146] D. Larson, H. G. Katzgraber, M. A. Moore, and A. P. Young. Spin glasses in a field: three and four dimensions as seen from one space dimension. *Phys. Rev. B*, 87:024414, 2013. doi: 10.1103/PhysRevB.87.024414.

[147] G. F. Lawler, O. Schramm, and W. Werner. Conformal invariance of planar loop-erased random walks and uniform spanning trees. *Ann. Probab.*, 32:939–995, 2004. doi: 10.1214/aop/1079021469.

[148] P. Le Doussal and A. B. Harris. Location of the Ising spin-glass multicritical point on Nishimori's line. *Phys. Rev. Lett.*, 61:625–628, 1988. doi: 10.1103/PhysRevLett.61.625.

[149] M. Loebl. Ground state incongruence in 2D spin glasses revisited. *The Electronic Journal of Combinatorics*, 11:R40, 2004. doi: 10.37236/1793.

[150] R. Lyons and Y. Peres. *Probability on Trees and Networks*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, Cambridge, 2016. doi: 10.1017/9781316672815.

[151] J. Machta, C. M. Newman, and D. L. Stein. The percolation signature of the spin glass transition. *J. Stat. Phys.*, 130:113–128, 2008. doi: 10.1007/s10955-007-9446-2. arXiv:0707.0073; page numbers refer to the arXiv version.

[152] J. Machta, C. M. Newman, and D. L. Stein. Percolation in the Sherrington–Kirkpatrick spin glass. In V. Sidoravicius and M. E. Vares, editors, *In and Out of Equilibrium 2*, volume 60 of *Progress in Probability*, pages 527–542. Birkhäuser, 2008. doi: 10.1007/978-3-7643-8786-0_24. arXiv:0710.1399; theorem numbers refer to the arXiv version.

[153] J. Machta, C. M. Newman, and D. L. Stein. A percolation-theoretic approach to spin glass phase transitions. Preprint, 2008. arXiv:0805.0794.

[154] A. Maiorano and G. Parisi. Support for the value 5/2 for the spin glass lower critical dimension at zero magnetic field. *Proc. Natl. Acad. Sci. USA*, 115:5129–5134, 2018. doi: 10.1073/pnas.1720832115.

[155] E. Marinari and F. Zuliani. Numerical simulations of the 4D Edwards–Anderson spin glass with binary couplings. *J. Phys. A: Math. Gen.*, 32(43):7447, 1999. doi: 10.1088/0305-4470/32/43/301.

[156] V. Martin-Mayor, J. J. Ruiz-Lorenzo, B. Seoane, and A. P. Young. Numerical simulations and replica symmetry breaking, 2022. Chapter in *Spin Glass Theory and Far Beyond* (World Scientific, 2023).

[157] D. C. Mattis. Solvable spin systems with random interactions. *Phys. Lett. A*, 56(5):421–422, 1976. doi: 10.1016/0375-9601(76)90396-0.

[158] C. Mazza. On the mean pair correlation function of $\pm J$ Ising spin glasses. *J. Stat. Phys.*, 92:559–569, 1998.

[159] O. A. McBryan and T. Spencer. On the decay of correlations in SO(n)-symmetric ferromagnets. *Comm. Math. Phys.*, 53(3):299–302, 1977. doi: 10.1007/BF01609854.

[160] W. L. McMillan. Scaling theory of Ising spin glasses. *J. Phys. C: Solid State Phys.*, 17(18):3179–3187, 1984. doi: 10.1088/0022-3719/17/18/010.

[161] M. Mézard, G. Parisi, and M. A. Virasoro. *Spin Glass Theory and Beyond*. World Scientific, Singapore, 1987.

[162] L. H. Miranda-Filho and Y. Jin. Evidence of the de Almeida–Thouless transition in three-dimensional spin glasses, 2026.

[163] A. Montanari. Optimization of the Sherrington–Kirkpatrick Hamiltonian. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1417–1433, 2019. doi: 10.1109/FOCS.2019.00087.

[164] M. A. Moore. The droplet-scaling versus replica symmetry breaking debate in spin glasses revisited. *Phys. Rev. E*, 103:062111, 2021. doi: 10.1103/PhysRevE.103.062111.

[165] S. Morita, H. Nishimori, and P. Contucci. Griffiths inequalities for the Gaussian spin glass. *J. Phys. A*, 37: L203, 2004. doi: 10.1088/0305-4470/37/18/L03.

[166] L. Münster and M. Weigel. Cluster percolation in the two-dimensional Ising spin glass. *Phys. Rev. E*, 107: 054103, 2023. doi: 10.1103/PhysRevE.107.054103. arXiv:2212.13087.

[167] L. Münster and M. Weigel. Cluster percolation in the three-dimensional $\pm J$ random-bond Ising model. *Phys. Rev. E*, 113:024139, 2026. doi: 10.1103/y464-wxbk. arXiv:2511.05748.

[168] S. Nanda, C. M. Newman, and D. L. Stein. Dynamics of Ising spin systems at zero temperature. In R. Minlos, S. Shlosman, and Y. Suhov, editors, *On Dobrushin's Way (from Probability Theory to Statistical Physics)*, volume 198 of *American Mathematical Society Translations: Series 2*, pages 183–194. American Mathematical Society, 2000. doi: 10.1090/trans2/198/13.

[169] C. S. J. A. Nash-Williams. Random walk and electric currents in networks. *Math. Proc. Cambridge Philos. Soc.*, 55(2):181–194, 1959. doi: 10.1017/S0305004100033879.

[170] C. M. Newman. Disordered Ising systems and random cluster representations. In G. Grimmett, editor, *Probability and Phase Transition*, volume 420 of *NATO ASI Series C*, pages 247–260. Kluwer, Dordrecht, 1994. doi: 10.1007/978-94-015-8326-8_15.

[171] C. M. Newman. *Topics in Disordered Systems*. Lectures in Mathematics ETH Zürich. Birkhäuser, Basel, 1997. doi: 10.1007/978-3-0348-8912-4.

[172] C. M. Newman and D. L. Stein. Multiple states and thermodynamic limits in short-ranged Ising spin-glass models. *Phys. Rev. B*, 46(2):973–982, 1992. doi: 10.1103/PhysRevB.46.973.

[173] C. M. Newman and D. L. Stein. Non-mean-field behavior of realistic spin glasses. *Phys. Rev. Lett.*, 76(3): 515–518, 1996. doi: 10.1103/PhysRevLett.76.515.

[174] C. M. Newman and D. L. Stein. Spatial inhomogeneity and thermodynamic chaos. *Physical Review Letters*, 76:4821–4824, 1996.

[175] C. M. Newman and D. L. Stein. Metastate approach to thermodynamic chaos. *Physical Review E*, 55: 5194–5211, 1997.

[176] C. M. Newman and D. L. Stein. Short-range spin glasses. Contribution of 18 September 1998 to the

open-problems page edited by M. Aizenman, 1998. URL https://web.math.princeton.edu/~aizenman/OpenProblems_MathPhys/9803.SpinGlass.html.

[177] C. M. Newman and D. L. Stein. Simplicity of state and overlap structure in finite-volume realistic spin glasses. *Physical Review E*, 57:1356, 1998.

[178] C. M. Newman and D. L. Stein. Thermodynamic chaos and the structure of short-range spin glasses. In A. Bovier and P. Picco, editors, *Mathematics of Spin Glasses and Neural Networks*, pages 243–287. Birkhäuser, Boston, 1998.

[179] C. M. Newman and D. L. Stein. Nature of ground state incongruence in two-dimensional spin glasses. *Physical Review Letters*, 84:3966–3969, 2000.

[180] C. M. Newman and D. L. Stein. Zero-temperature dynamics of Ising spin systems following a deep quench: results and open problems. *Physica A: Statistical Mechanics and its Applications*, 279:159–168, 2000. doi: 10.1016/S0378-4371(99)00511-7.

[181] C. M. Newman and D. L. Stein. Are there incongruent ground states in 2D Edwards–Anderson spin glasses? *Communications in Mathematical Physics*, 224:205–218, 2001. doi: 10.1007/PL00005586.

[182] C. M. Newman and D. L. Stein. Interfaces and the question of regional congruence in spin glasses. *Physical Review Letters*, 87:077201, 2001. doi: 10.1103/PhysRevLett.87.077201.

[183] C. M. Newman and D. L. Stein. Ordering and broken symmetry in short-ranged spin glasses. *J. Phys.: Condens. Matter*, 15(32):R1319–R1364, 2003. doi: 10.1088/0953-8984/15/32/202.

[184] C. M. Newman and D. L. Stein. Short-range spin glasses: the metastate approach. Preprint, 2004. arXiv:cond-mat/0411426.

[185] C. M. Newman and D. L. Stein. Short-range spin glasses: results and speculations. In E. Bolthausen and A. Bovier, editors, *Spin Glasses*, volume 1900 of *Lecture Notes in Mathematics*, pages 159–175. Springer, Berlin, 2007. doi: 10.1007/978-3-540-40908-3_7.

[186] C. M. Newman and D. L. Stein. *Spin Glasses and Complexity*. Princeton University Press, 2013. doi: 10.23943/princeton/9780691147338.001.0001.

[187] C. M. Newman and D. L. Stein. Free energy difference fluctuations in short-range spin glasses. *Journal of Statistical Physics*, 191:128, 2024. doi: 10.1007/s10955-024-03334-4.

[188] C. M. Newman and D. L. Stein. Overlap structure and free energy fluctuations in short-range spin glasses. *Journal of Physics A*, 57:11LT01, 2024.

[189] C. M. Newman and D. L. Stein. Ground state excitations and energy fluctuations in short-range spin glasses. *Journal of Statistical Physics*, 193:70, 2026. doi: 10.1007/s10955-026-03631-0.

[190] C. M. Newman, N. Read, and D. L. Stein. Metastates and replica symmetry breaking. In *Spin Glass Theory and Far Beyond*, pages 697–718. World Scientific, 2023. doi: 10.1142/9789811273926_0033.

[191] C. M. Newman, N. Read, and D. L. Stein. Proof of single-replica equivalence in short-range spin glasses. *Physical Review Letters*, 130:077102, 2023.

[192] V. Q. Nguyen. Superconcentration and chaos in Bernoulli percolation, 2026. arXiv:2602.01036v1.

[193] H. Nishimori. *Statistical Physics of Spin Glasses and Information Processing: An Introduction*. Oxford University Press, 2001. Page numbers refer to the printed book.

[194] H. Nishimori. Duality in finite-dimensional spin glasses. *J. Stat. Phys.*, 126:977, 2007. doi: 10.1007/s10955-006-9156-1. Locators refer to arXiv v3.

[195] H. Nishimori and K. Nemoto. Duality and multicritical point of two-dimensional spin glasses. *J. Phys. Soc. Jpn.*, 71:1198, 2002. doi: 10.1143/JPSJ.71.1198.

[196] H. Nishimori and D. Sherrington. Absence of replica symmetry breaking in a region of the phase diagram of the Ising spin glass. In P. Sollich et al., editors, *Disordered and Complex Systems*, volume 553 of *AIP Conf. Proc.*, page 67, 2001. doi: 10.1063/1.1358165.

[197] H. Nishimori and K. Y. M. Wong. Statistical mechanics of image restoration and error-correcting codes. *Phys. Rev. E*, 60:132, 1999. doi: 10.1103/PhysRevE.60.132.

[198] H. Nishimori, M. Ohzeki, and M. Okuyama. Temperature chaos as a logical consequence of the reentrant transition in spin glasses. *Phys. Rev. E*, 112:044140, 2025. doi: 10.1103/qp1w-qcbs.

[199] M. Ohzeki. Locations of multicritical points for spin glasses on regular lattices. *Phys. Rev. E*, 79:021129, 2009. doi: 10.1103/PhysRevE.79.021129.

[200] M. Okuyama and M. Ohzeki. Inequality for local energy of Ising models with quenched randomness and its application. *J. Phys. Soc. Japan*, 89:064704, 2020. doi: 10.7566/JPSJ.89.064704.

[201] M. Okuyama and M. Ohzeki. Absence of spin-glass order on Migdal–Kadanoff hierarchical lattices near three dimensions. *Preprint*, 2026. arXiv:2607.15676.

[202] M. Okuyama and M. Ohzeki. Existence and absence of spin-glass phases on Migdal–Kadanoff hierarchical lattices: A complete classification, 2026. arXiv:2610.32291.

[203] M. Okuyama and M. Ohzeki. Existence of a phase transition in the one-dimensional Ising spin glass model with long-range interactions on the Nishimori line. *J. Phys. A: Math. Theor.*, 59:305204, 2026. doi: 10.1088/1751-8121/ae8c95. arXiv:2604.07130; locators refer to v2.

[204] M. Okuyama and M. Ohzeki. Griffiths inequalities and Gibbs–Bogoliubov inequality for general gauge glasses with Gaussian disorder on Nishimori line, 2026. arXiv:2604.25697.

[205] M. Palassini and A. P. Young. Nature of the spin glass state. *Phys. Rev. Lett.*, 85:3017, 2000. doi: 10.1103/PhysRevLett.85.3017.

[206] D. Panchenko. A note on Talagrand's positivity principle. *Electron. Commun. Probab.*, 12:401–410, 2007.

[207] D. Panchenko. The Parisi ultrametricity conjecture. *Ann. of Math.*, 177(1):383–393, 2013.

[208] D. Panchenko. *The Sherrington–Kirkpatrick Model*. Springer Monographs in Mathematics. Springer, 2013. doi: 10.1007/978-1-4614-6289-7.

[209] D. Panchenko. The Parisi formula for mixed p-spin models. *Ann. Probab.*, 42(3):946–958, 2014. doi: 10.1214/12-AOP800.

[210] G. Parisi. Infinite number of order parameters for spin-glasses. *Phys. Rev. Lett.*, 43(23):1754–1756, 1979. doi: 10.1103/PhysRevLett.43.1754.

[211] G. Parisi. Order parameter for spin-glasses. *Phys. Rev. Lett.*, 50(24):1946–1948, 1983. doi: 10.1103/PhysRevLett.50.1946.

[212] Y. R. Pei. A computer-assisted sufficient condition for Edwards–Anderson order in the square-lattice $\pm J$ model. Zenodo preprint, 2026. doi: 10.5281/zenodo.24059073.

[213] Y. R. Pei. Certificate data for Edwards–Anderson order in the square-lattice $\pm J$ model. Zenodo, 2026. doi: 10.5281/zenodo.24059144.

[214] Y. R. Pei. Towards the spin-glass transition in finite dimensions via blue percolation. *Preprint*, 2026. arXiv:2609.22301; numbering refers to version 1.

[215] Y. R. Pei. At most two infinite blue clusters in the CMR representation of the Edwards–Anderson spin glass. *Preprint*, 2026. arXiv:2605.17338.

[216] Y. R. Pei. An effective-conductance bound for the arboreal gas. Preprint, Zenodo, 2026. doi: 10.5281/zenodo.24047742.

[217] Y. R. Pei. Exact frustration cancellation and uniqueness at all positive temperatures for diluted $\pm J$ spin glasses slightly above the percolation threshold. Preprint, Zenodo, 2026. doi: 10.5281/zenodo.24049737.

[218] Y. R. Pei. Certificates for explicit uniqueness windows of diluted $\pm J$ spin glasses. Zenodo, 2026. doi: 10.5281/zenodo.24057994.

[219] Y. R. Pei. Positive-temperature spin-glass order on the three-dimensional Migdal–Kadanoff lattice, 2026. arXiv:2609.38220.

[220] Y. R. Pei. Certificate programs for the review "Dancing with frustration: short-range Ising spin glasses". Zenodo, 2026. doi: 10.5281/zenodo.24061127.

[221] N. Read. Short-range Ising spin glasses: the metastate interpretation of replica symmetry breaking. *Phys. Rev. E*, 90:032142, 2014. doi: 10.1103/PhysRevE.90.032142.

[222] N. Read and A. W. W. Ludwig. Absence of a metallic phase in random-bond Ising models in two dimensions: applications to disordered superconductors and paired quantum Hall states. *Phys. Rev. B*, 63:024404, 2000. doi: 10.1103/PhysRevB.63.024404.

[223] O. Redner, J. Machta, and L. F. Chayes. Graphical representations and cluster algorithms for critical points with fields. *Phys. Rev. E*, 58:2749, 1998. arXiv:cond-mat/9802063.

[224] D. Sherrington and S. Kirkpatrick. Solvable model of a spin-glass. *Phys. Rev. Lett.*, 35(26):1792–1796, 1975. doi: 10.1103/PhysRevLett.35.1792.

[225] S. Smirnov. Critical percolation in the plane: conformal invariance, Cardy's formula, scaling limits. *C. R. Acad. Sci. Paris Sér. I Math.*, 333:239–244, 2001. doi: 10.1016/S0764-4442(01)01991-7.

[226] J. M. Steele. An Efron–Stein inequality for nonsymmetric statistics. *Ann. Statist.*, 14(2):753–758, 1986. doi: 10.1214/aos/1176349952.

[227] D. L. Stein. Frustration and fluctuations in systems with quenched disorder, 2014. arXiv:1406.0816; published in *PWA90: A Lifetime of Emergence* (World Scientific, 2016), pp. 169–186.

[228] E. Subag. Following the ground states of full-RSB spherical spin glasses. *Comm. Pure Appl. Math.*, 74: 1021–1044, 2021. doi: 10.1002/cpa.21922.

[229] R. H. Swendsen and J.-S. Wang. Replica Monte Carlo simulation of spin-glasses. *Phys. Rev. Lett.*, 57:2607, 1986.

[230] R. H. Swendsen and J.-S. Wang. Nonuniversal critical dynamics in Monte Carlo simulations. *Phys. Rev. Lett.*, 58(2):86–88, 1987. doi: 10.1103/PhysRevLett.58.86.

[231] M. Talagrand. The Parisi formula. *Ann. of Math.*, 163(1):221–263, 2006. doi: 10.4007/annals.2006.163.221.

[232] M. Talagrand. *Mean Field Models for Spin Glasses*. Springer, 2011. doi: 10.1007/978-3-642-15202-3.

[233] C. K. Thomas, D. A. Huse, and A. A. Middleton. Zero and low temperature behavior of the two-dimensional $\pm J$ Ising spin glass. *Phys. Rev. Lett.*, 107:047203, 2011. doi: 10.1103/PhysRevLett.107.047203.

[234] T. S. Tomar. Certified lower bounds on the optimal-decoding threshold of the surface code. aiXiv:2611.00584v2, 2026.

[235] J. van den Berg. A uniqueness condition for Gibbs measures, with application to the 2-dimensional Ising antiferromagnet. *Comm. Math. Phys.*, 152:161–166, 1993. doi: 10.1007/BF02097061.

[236] J. van den Berg and C. Maes. Disagreement percolation in the study of Markov fields. *Ann. Probab.*, 22: 749–763, 1994. doi: 10.1214/aop/1176988728.

[237] A. C. D. van Enter and R. B. Griffiths. The order parameter in a spin glass. *Comm. Math. Phys.*, 90: 319–327, 1983. doi: 10.1007/BF01206885.

[238] A. C. D. van Enter and J. L. van Hemmen. The thermodynamic limit for long-range random systems. *J. Stat. Phys.*, 32:141–152, 1983. doi: 10.1007/BF01009426.

[239] B. Vedula, M. A. Moore, and A. Sharma. Evidence that the de Almeida–Thouless transition disappears below six dimensions. *Phys. Rev. E*, 110(5):054131, 2024. doi: 10.1103/PhysRevE.110.054131.

[240] P. A. Vuillermot. Thermodynamics of quenched random spin systems, and application to the problem of phase transitions in magnetic (spin) glasses. *J. Phys. A: Math. Gen.*, 10:1319–1333, 1977. doi: 10.1088/0305-4470/10/8/011.

[241] C. Wang, J. Harrington, and J. Preskill. Confinement-Higgs transition in a disordered gauge theory and the accuracy threshold for quantum memory. *Ann. Physics*, 303:31–58, 2003.

[242] W. Wang, J. Machta, and H. G. Katzgraber. Evidence against a mean field description of short-range spin glasses revealed through thermal boundary conditions. *Phys. Rev. B*, 90:184412, 2014. doi: 10.1103/PhysRevB.90.184412.

[243] W. Wang, M. A. Moore, and H. G. Katzgraber. Fractal dimension of interfaces in Edwards–Anderson spin glasses for up to six space dimensions. *Phys. Rev. E*, 97:032104, 2018. doi: 10.1103/PhysRevE.97.032104.

[244] J. Wehr. On the number of infinite geodesics and ground states in disordered systems. *Journal of Statistical Physics*, 87:439–447, 1997. doi: 10.1007/BF02181495.

[245] J. Wehr and M. Aizenman. Fluctuations of extensive functions of quenched random couplings. *Journal of Statistical Physics*, 60:287–306, 1990. doi: 10.1007/BF01314921.

[246] B. Yucesoy, H. G. Katzgraber, and J. Machta. Evidence of non-mean-field-like low-temperature behavior in the Edwards–Anderson spin-glass model. *Phys. Rev. Lett.*, 109:177204, 2012. doi: 10.1103/PhysRevLett.109.177204.

[247] Z. Zhu, A. J. Ochoa, and H. G. Katzgraber. Efficient cluster algorithm for spin glasses in any space dimension. *Phys. Rev. Lett.*, 115:077201, 2015. arXiv:1501.05630.