

Belief propagation on the random k-SAT model

Citation for published version (APA):

Coja-Oghlan, A., Müller, N., & Ravelomanana, J. B. (2022). Belief propagation on the random k-SAT model. *Annals of Applied Probability*, 32(5), 3718-3796. <https://doi.org/10.1214/21-AAP1772>

Document license:

TAVERNE

DOI:

[10.1214/21-AAP1772](https://doi.org/10.1214/21-AAP1772)

Document status and date:

Published: 01/10/2022

Document Version:

Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

BELIEF PROPAGATION ON THE RANDOM k -SAT MODEL

BY AMIN COJA-OGHLAN^{1,a}, NOELA MÜLLER^{2,c} AND JEAN B. RAVELOMANANA^{1,b}

¹Fakultät für Informatik, TU Dortmund, ^aamin.coja-oghlan@tu-dortmund.de, ^bjean.ravelomanana@tu-dortmund.de

²Department of Mathematics and Computer Science, Eindhoven University of Technology, ^cn.s.muller@tue.nl

Corroborating a prediction from statistical physics, we prove that the belief propagation message passing algorithm approximates the partition function of the random k -SAT model well for all clause/variable densities and all inverse temperatures for which a modest absence of long-range correlations condition is satisfied. This condition is known as “replica symmetry” in physics language. From this result we deduce that a replica symmetry breaking phase transition occurs in the random k -SAT model at low temperature for clause/variable densities below but close to the satisfiability threshold.

CONTENTS

1. Introduction	3719
1.1. Background and motivation	3719
1.2. Belief propagation	3720
1.3. Replica symmetry breaking	3722
2. Overview	3723
2.1. The second moment bound	3723
2.2. Belief propagation on trees	3725
2.3. The Bethe free energy	3726
2.4. Replica symmetry breaking	3727
3. Discussion	3728
4. Preliminaries and notation	3730
5. Moment calculations	3732
5.1. Overview	3732
5.2. Proof of Proposition 5.1	3734
5.3. Proof of Proposition 5.2	3736
5.3.1. Proof of Proposition 5.8	3737
5.3.2. Proof of Proposition 5.9	3749
5.3.3. Proof of Proposition 5.10	3751
5.4. Proof of Corollary 5.3	3753
6. Proof of Proposition 2.3	3754
6.1. Outline	3754
6.2. Proof of Proposition 6.2	3757
6.2.1. Overview	3757
6.2.2. Proof of Lemma 6.5	3759
6.2.3. Proof of Lemma 6.8	3760
6.2.4. Proof of Lemma 6.9	3761
6.3. Proof of Proposition 6.3	3763
6.4. Proof of Proposition 6.4	3767
6.4.1. Overview	3767
6.4.2. Proof of Lemma 6.16	3768
6.4.3. Proof of Lemma 6.17	3770
7. Proof of Proposition 2.5	3775
8. Replica symmetry breaking	3778
8.1. Proof of Proposition 2.7	3778

Received November 2020; revised October 2021.

MSC2020 subject classifications. 68Q87, 60C05.

Key words and phrases. k -SAT, belief propagation, replica symmetry breaking, phase transition.

8.1.1. Proof of Lemma 8.2 3780
 8.1.2. Proof of Lemma 8.4 3780
 8.1.3. Proof of Lemma 8.7 3781
 8.1.4. Proof of Lemma 8.8 3785
 8.1.5. Proof of Lemma 8.5 3786
 8.1.6. Proof of Lemma 8.6 3787
 8.2. Proof of Proposition 2.6 3789
 Appendix: Proof of Lemma 8.1 3791
 Acknowledgments 3794
 Funding 3794
 References 3794

1. Introduction.

1.1. *Background and motivation.* According to a prominent physics prediction the belief propagation message passing algorithm renders a good approximation to the partition function of locally tree-like graphical models that do not exhibit long-range correlations [33]. Turning this somewhat vague intuition into a mathematical theorem has been a major open problem at the junction of computer science and probability theory, specifically spin glass theory, for quite some time [34]. The random k -SAT model is one of the specific examples to which both communities have directed a large amount of effort [3, 4, 19, 24, 30, 35, 38, 41, 43, 46].

Corroborating the physics conjecture, we prove that an extremely modest absence of long-range correlations condition known as “replica symmetry” precipitates the success of belief propagation for the random k -SAT model. The replica symmetry condition is generally deemed to be necessary, too [33]. Apart from significantly advancing the mathematical understanding of belief propagation, this result allows for an intriguing application. Namely, by way of characterising the fixed point of the belief propagation message passing process precisely, we deduce that replica symmetry fails to hold for clause-to-variable densities near the satisfiability threshold. Thus, we prove that the random k -SAT model undergoes a replica symmetry breaking phase transition for clause-to-variable ratios close to but below the satisfiability threshold.

To appraise ourselves of the random k -SAT model, let $V_n = \{x_1, \dots, x_n\}$ be a set of n Boolean variables. We represent their possible values “true” and “false” by ± 1 . Also let $k \geq 3$ be an integer, let $d > 0$ be a real and let m be a Poisson variable with mean dn/k . The random k -SAT formula comprises m clauses a_1, \dots, a_m . For each clause a_i we independently choose a family $(x_{ij})_{1 \leq j \leq k} \in V_n^k$ of k variables uniformly without replacement. Additionally, let $(J_{ij})_{i,j \geq 1}$ be a family of independent ± 1 -variables with mean zero. Combinatorially a_i represents the Boolean clause comprising the k variables x_{i1}, \dots, x_{ik} with signs J_{i1}, \dots, J_{ik} . Thus, x_{ij} appears as a positive literal in a_i if $J_{ij} = 1$, and a_i features the negative literal $\neg x_{ij}$ otherwise. Hence, a Boolean assignment $\sigma \in \{\pm 1\}^{V_n}$ satisfies clause a_i (“ $\sigma \models a_i$ ”) if $\max_{j=1, \dots, k} J_{ij} \sigma_{x_{ij}} = 1$. Finally, $\Phi = \Phi_k(n, m)$ is the conjunction of all the m clauses, that is,

$$\Phi = \bigwedge_{i=1}^m a_i = \bigwedge_{i=1}^m (J_{i1}x_{i1} \vee \dots \vee J_{ik}x_{ik}).$$

Further, given an inverse temperature parameter $\beta > 0$, the Boltzmann distribution of the model reads

$$(1.1) \quad \mu_{\Phi, \beta}(\sigma) = \frac{1}{Z(\Phi, \beta)} \prod_{i=1}^m \exp(-\beta \mathbb{1}\{\sigma \not\models a_i\}) \quad (\sigma \in \{\pm 1\}^{V_n}),$$

where

$$Z(\Phi, \beta) = \sum_{\tau \in \{\pm 1\}^{V_n}} \exp\left(-\beta \sum_{i=1}^m \mathbb{1}\{\tau \not\models a_i\}\right).$$

Thus, the Boltzmann weight of an assignment σ contains an $\exp(-\beta)$ penalty factor for every violated clause. In effect, as β increases, the distribution assigns greater weight to “more satisfying” assignments. As always, the partition function $Z(\Phi, \beta)$ accounts for the total weight.

The random k -SAT model undergoes a satisfiability phase transition at a certain critical value of d called the *satisfiability threshold*. To elaborate, observe that d gauges the average number of clauses in which a given Boolean variable appears. Clearly, as variables appear in more and more clauses it becomes harder to satisfy all these clauses simultaneously. Indeed, for large enough $k \geq 3$ there exists a threshold $d_{\text{SAT}}(k)$ such that Φ admits an assignment that satisfies all clauses asymptotically almost surely if $d < d_{\text{SAT}}(k)$, while for $d > d_{\text{SAT}}(k)$ no satisfying assignment exists a.a.s. The value of $d_{\text{SAT}}(k)$ is known precisely but the formula is quite complicated [24]; asymptotically in the limit of large k we have

$$(1.2) \quad d_{\text{SAT}}(k) = 2^k k \log 2 - \frac{1 + \log 2}{2} k + o(1).$$

The regime $d < d_{\text{SAT}}(k)$ is of fundamental interest in computer science to assess the power and the limitations of algorithms for finding, counting and sampling solutions to the k -SAT problem, the cornerstone of computational complexity theory [4]. Therefore, we will investigate the Boltzmann distribution for $d < d_{\text{SAT}}(k)$ for varying values of β . As we increase β we effectively scan the energy landscape that an algorithm has to traverse on its quest for satisfying assignments. In particular, we investigate the performance of the belief propagation message passing algorithm. With what regimes of d, β can the algorithm cope? Is belief propagation fit to approximate the partition function $Z(\Phi, \beta)$? Does there exist a critical value of β where long-range correlations emerge?

1.2. Belief propagation. The belief propagation algorithm associates two “messages” $\mu_{\Phi, \beta, a_i \rightarrow x_{ij}, t}(\pm 1), \mu_{\Phi, \beta, x_{ij} \rightarrow a_i, t}(\pm 1) \in (0, 1)$ with each interacting clause/variable pair (a_i, x_{ij}) . The messages are indexed by time $t \geq 0$ and always normalised such that

$$(1.3) \quad \mu_{\Phi, \beta, a_i \rightarrow x_{ij}, t}(1) + \mu_{\Phi, \beta, a_i \rightarrow x_{ij}, t}(-1) = \mu_{\Phi, \beta, x_{ij} \rightarrow a_i, t}(1) + \mu_{\Phi, \beta, x_{ij} \rightarrow a_i, t}(-1) = 1.$$

The first message $\mu_{\Phi, \beta, a_i \rightarrow x_{ij}, t}(\pm 1)$ is directed from the clause to the variable. The other one travels in the reverse direction. The messages are updated iteratively. Initially, all messages are set to $1/2$, that is,

$$(1.4) \quad \mu_{\Phi, \beta, a_i \rightarrow x_{ij}, 0}(\pm 1) = \mu_{\Phi, \beta, x_{ij} \rightarrow a_i, 0}(\pm 1) = 1/2 \quad \text{for all } 1 \leq i \leq m, 1 \leq j \leq k.$$

Furthermore, for integers $t \geq 0$ and $s = \pm 1$ we inductively define

$$(1.5) \quad \mu_{\Phi, \beta, a_i \rightarrow x_{ij}, t+1}(s) \propto \sum_{\sigma \in \{\pm 1\}^k} \mathbb{1}\{\sigma_j = s\} \exp(-\beta \mathbb{1}\{\sigma \not\models a_i\}) \prod_{\substack{1 \leq h \leq k \\ h \neq j}} \mu_{\Phi, \beta, x_{ih} \rightarrow a_i, t}(\sigma_h),$$

$$(1.6) \quad \mu_{\Phi, \beta, x_{ij} \rightarrow a_i, t+1}(s) \propto \prod_{\substack{1 \leq h \leq m \\ h \neq i}} \prod_{\substack{1 \leq \ell \leq k \\ x_{h\ell} = x_{ij}}} \mu_{\Phi, \beta, a_h \rightarrow x_{h\ell}, t+1}(s).$$

Here the α -symbol hides the normalisation required to bring about (1.3). Finally, the estimate of the partition function after t iterations reads

$$\begin{aligned}
 \mathcal{B}_t = & \sum_{i=1}^n \log \left[\sum_{s=\pm 1} \prod_{\substack{1 \leq h \leq m, 1 \leq j \leq k \\ x_{hj}=x_i}} \mu_{\Phi, \beta, a_h \rightarrow x_i, t}(s) \right] \\
 (1.7) \quad & + \sum_{i=1}^m \log \left[\sum_{\sigma \in \{\pm 1\}^k} e^{-\beta \mathbb{1}_{\{\sigma \neq a_i\}}} \prod_{j=1}^k \mu_{\Phi, \beta, x_{ij} \rightarrow a_i, t}(\sigma_j) \right] \\
 & - \sum_{i=1}^m \sum_{j=1}^k \log \left[\sum_{s=\pm 1} \mu_{\Phi, \beta, a_i \rightarrow x_{ij}, t}(s) \mu_{\Phi, \beta, x_{ij} \rightarrow a_i, t}(s) \right].
 \end{aligned}$$

This expression is called the *Bethe free energy* in physics jargon. An excellent in-depth discussion of belief propagation, including a derivation of (1.5)–(1.7), can be found in [34], Chapter 14.

The key feature of all the above formulas is that they are governed by the *local* structure of the k -SAT formula. For instance, (1.5) involves only the messages sent out by the variables which appear in clause a_i . Similarly, (1.6) comes in terms of the messages sent out by the clauses in which variable x_{ij} appears. Therefore, we can reasonably hope that belief propagation represents local dependencies accurately, but hardly that the messages can faithfully capture long-range correlations. In fact, one of the most important predictions about the random k -SAT model holds that a very weak “absence of long-range correlations” condition suffices for the success of belief propagation [33]. Specifically, let $\sigma = \sigma_{\Phi, \beta}$ denote a sample from the Boltzmann distribution $\mu_{\Phi, \beta}$. Then following [33] we say that the random k -SAT model with parameters d, β is *replica symmetric* if

$$(1.8) \quad \lim_{n \rightarrow \infty} \mathbb{E} |\mu_{\Phi, \beta}(\{\sigma_{x_1} = \sigma_{x_2} = 1\}) - \mu_{\Phi, \beta}(\{\sigma_{x_1} = 1\})\mu_{\Phi, \beta}(\{\sigma_{x_2} = 1\})| = 0.$$

In words, in the experiment of first drawing a formula Φ on the variable set $V_n = \{x_1, \dots, x_n\}$ and then an assignment $\sigma \in \{\pm 1\}^{V_n}$ from the distribution $\mu_{\Phi, \beta}$, the event $\{\sigma_{x_1} = 1\}$ that the first variable $x_1 \in V_n$ is set to “true” is asymptotically (in the limit of large n) independent of the event $\{\sigma_{x_2} = 1\}$ that the second variable $x_2 \in V_n$ is set to “true”. In order to justify the term “absence of long-range correlations”, we represent the formula Φ by a graph $G(\Phi)$ that induces a metric on the variables and clauses. Specifically, the *factor graph* $G(\Phi)$ of Φ is a bipartite graph with one class of vertices comprising the variables x_1, \dots, x_n and a second class of vertices comprising the clauses a_1, \dots, a_m . A clause a_i and a variable x_j are connected by an edge if x_j appears in a_i . Routine arguments then show that with respect to the shortest path metric induced by this graph, the first two variables x_1 and x_2 typically are distance $\Omega(\log n)$ apart. Thus, x_1, x_2 are typically far apart, and if (1.8) holds then their truth values are typically asymptotically independent as well. By contrast, we cannot expect that variables that are near to others in the factor graph metric have asymptotically independent truth values. For instance, in the unlikely event that the first two variables x_1, x_2 are both adjacent to the same clause, their truth values will clearly be correlated as they will strive to satisfy their common adjacent clause. Finally, because the random formula Φ is invariant under permutations of the variables, there is nothing special about x_1, x_2 . In fact, (1.8) is equivalent to the statement

$$(1.9) \quad \lim_{n \rightarrow \infty} \frac{1}{n^2} \mathbb{E} \sum_{i, j=1}^n |\mu_{\Phi, \beta}(\{\sigma_{x_i} = \sigma_{x_j} = 1\}) - \mu_{\Phi, \beta}(\{\sigma_{x_i} = 1\})\mu_{\Phi, \beta}(\{\sigma_{x_j} = 1\})| = 0$$

that the truth values of all but $o(n^2)$ variable pairs x_i, x_j decorrelate. In summary, (1.8) rules out extensive long-range correlations, albeit in a very weak sense. In particular, (1.8) is far more modest a condition than classical spatial mixing properties such as Gibbs uniqueness or nonreconstruction [26, 33].

The following theorem vindicates the prediction that (1.8) is a sufficient condition for the success of belief propagation for all $\beta \geq 1$ and for all d up to within a whisker of the satisfiability threshold $d_{\text{SAT}}(k)$.

THEOREM 1.1. *There exists a constant $k_0 \geq 3$ such that for any $k \geq k_0, \beta \geq 1$ and any*

$$(1.10) \quad d \leq d^* = d^*(k) = k2^k \log 2 - 10k^2$$

the following is true: if (1.8) is satisfied then

$$\lim_{t \rightarrow \infty} \limsup_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} |\mathcal{B}_t - \log Z(\Phi, \beta)| = 0.$$

Theorem 1.1 is a conditional result: for all $d \leq d^*$ and all $\beta \geq 1$ for large enough t the Bethe free energy formula (1.7) estimates the logarithm of the partition function up to an additive error of $o(n)$ provided that (1.8) holds. At this point (1.8) is known to be satisfied only for values of d much smaller than d^* ; the best current bound yields $d \leq \log k$ [38]. However, (1.8) is expected to hold for all $\beta > 0$ and all $d \leq d^*$ (and in fact for slightly larger d) [33]. Yet conceptually the point that Theorem 1.1 makes is that the modest condition (1.8) is the *only* requirement for the success of belief propagation. In other words, belief propagation launched from the trivial initial condition (1.4) does faithfully capture the short-range effects of the random k -SAT model. A further strength of Theorem 1.1 is that the result covers all reasonable values of β . Indeed, the assumption $\beta \geq 1$ is harmless as the most interesting regime should be that of large β , where the satisfiability condition really bites, known as the “low temperature” regime in physics terminology.

1.3. Replica symmetry breaking. The proof of Theorem 1.1 has an unconditional consequence. Namely, we can turn the tables and prove that (1.8) fails to be satisfied for d close to the satisfiability threshold $d_{\text{SAT}}(k)$.

THEOREM 1.2. *There exist sequences $\varepsilon_k \rightarrow 0$ and $\beta_0(k) > 0$ such that the following is true. Assume that $\beta > \beta_0(k)$ and*

$$(1.11) \quad 2^k k \log 2 - k(3 + \varepsilon_k) \log 2/2 \leq d \leq d_{\text{SAT}}.$$

Then

$$(1.12) \quad \limsup_{n \rightarrow \infty} \mathbb{E} |\mu_{\Phi, \beta}(\{\sigma_{x_1} = \sigma_{x_2} = 1\}) - \mu_{\Phi, \beta}(\{\sigma_{x_1} = 1\})\mu_{\Phi, \beta}(\{\sigma_{x_2} = 1\})| > 0 \quad \text{and}$$

$$(1.13) \quad \liminf_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} [\mathcal{B}_t - \log Z(\Phi, \beta)] > 0 \quad \text{uniformly for all } t > 0.$$

The asymptotic value $2^k k \log 2 - 3k \log 2/2$ from (1.11) was predicted via physics methods as the threshold for replica symmetry to break [33]. Thus, Theorem 1.2 confirms this conjecture. Indeed, (1.12) shows that very strong long-range correlations start to emerge in the “low temperature” (viz. large β) regime. Together with known results on the structure of asymptotic Gibbs measures, (1.12) shows that the Boltzmann distribution decomposes into several “pure states” in a certain precise sense; see [9, 21] for a detailed discussion. Additionally, (1.13) implies that beyond (1.11) belief propagation ceases to yield a good approximation to the partition function.

Theorem 1.2 also sheds new light on the satisfiability threshold. Namely, Theorem 1.2 establishes for the first time that replica symmetry breaking occurs in the random k -SAT problem strictly prior to the k -SAT threshold from (1.2), which exceeds the bound from (1.11) by an additive $\log(2) - 1/2 \approx 0.19$. Hence, Theorem 1.2 demonstrates that the random k -SAT model really is conceptually richer than models like random k -XORSAT or random 2-SAT, whose satisfiability thresholds were found much earlier [14, 23, 25, 31, 44].

We proceed to outline the proof strategy behind Theorems 1.1 and 1.2. Subsequently we discuss how the contributions of this paper compare to prior work.

2. Overview. The proof of Theorem 1.1 has three basic ingredients. First we need a rough estimate of $Z(\Phi, \beta)$, which we derive via a subtle second moment calculation. From this estimate we will deduce that “most” variable marginals under the Boltzmann distribution $\mu_{\Phi, \beta}$ are close to $1/2$ a.a.s. Second, we investigate the belief propagation message passing scheme on a random Galton–Watson tree that mimics the local geometry of the random k -SAT formula Φ . Specifically, we will use contraction arguments to show that if belief propagation launches from messages that are mostly close to $1/2$, the message passing scheme will rapidly approach a fixed point. Third, we will combine these two facts with probabilistic invariance properties of the random formula Φ to complete the proof of Theorem 1.1.

The proof of Theorem 1.2 is an afterthought to the proof of the first theorem. Indeed, the proof of Theorem 1.1 renders an implicit formula for the value that $Z(\Phi, \beta)$ must take if (1.8) is satisfied. To obtain Theorem 1.2 we calculate this value explicitly for large β . To refute (1.8) we then compare this result with the upper bound that the interpolation method from mathematical physics yields.

2.1. *The second moment bound.* A natural first stab at estimating $Z(\Phi, \beta)$ is to calculate its first two moments. The first moment is easy. Indeed, because any specific assignment $\sigma \in \{\pm 1\}^{V_n}$ satisfies a random clause with probability $1 - 2^{-k}$ and because the clauses are independent, the linearity of expectation gives

$$(2.1) \quad \log \mathbb{E}[Z(\Phi, \beta) \mid \mathbf{m}] = n \log 2 + \mathbf{m} \log(1 - 2^{-k}(1 - e^{-\beta})).$$

Hence, Markov’s inequality immediately implies that $\log Z(\Phi, \beta) \leq n \log 2 + \frac{dn}{k} \log(1 - (1 - e^{-\beta})2^{-k}) + o(n)$.

Moving on to the second moment and using the linearity of expectation and independence once more, we find

$$(2.2) \quad \begin{aligned} \mathbb{E}[Z(\Phi, \beta)^2 \mid \mathbf{m}] &= \sum_{\sigma, \tau \in \{\pm 1\}^{V_n}} \mathbb{E}[e^{-\beta \sum_{i=1}^m \mathbb{1}\{\sigma \neq a_i\} + \mathbb{1}\{\tau \neq a_i\}} \mid \mathbf{m}] \\ &= \sum_{\sigma, \tau \in \{\pm 1\}^{V_n}} \mathbb{E}[e^{-\beta(\mathbb{1}\{\sigma \neq a_1\} + \mathbb{1}\{\tau \neq a_1\})}]^m. \end{aligned}$$

To evaluate the right-hand side we define the *overlap* of two assignments $\sigma, \tau \in \{\pm 1\}^{V_n}$ as

$$(2.3) \quad \alpha(\sigma, \tau) = \frac{1}{n} \sum_{i=1}^n \frac{1 + \sigma_{x_i} \tau_{x_i}}{2} = \frac{1}{n} \sum_{i=1}^n \mathbb{1}\{\sigma_{x_i} = \tau_{x_i}\}.$$

A straightforward application of inclusion/exclusion then reveals that

$$(2.4) \quad \mathbb{E}[e^{-\beta(\mathbb{1}\{\sigma \neq a_1\} + \mathbb{1}\{\tau \neq a_1\})}] = 1 - 2^{1-k}(1 - e^{-\beta}) + 2^{-k} \alpha(\sigma, \tau)^k (1 - e^{-\beta})^2 + O(1/n).$$

Hence, it seems like a good idea to reorder the sum (2.2) according to the overlap. We thus sum on $\alpha \in [0, 1]$ such that αn is an integer. Since there are $2^n \binom{n}{\alpha n}$ pairs σ, τ with overlap α ,

(2.4) yields

$$(2.5) \quad \mathbb{E}[Z(\Phi, \beta)^2 \mid \mathbf{m}] = \exp(O(m/n)) \cdot 2^n \sum_{\alpha} \binom{n}{\alpha n} [1 - 2^{1-k}(1 - e^{-\beta}) + 2^{-k}\alpha^k(1 - e^{-\beta})^2]^m.$$

Taking logarithms in (2.5), assuming $\mathbf{m} = dn/k + o(n)$ and replacing the sum by a max, we obtain

$$(2.6) \quad \frac{1}{n} \log \mathbb{E}[Z(\Phi, \beta)^2 \mid \mathbf{m}] = \max_{\alpha \in (0,1)} f(\alpha) + o(1),$$

a.a.s., where

$$f(\alpha) = \log 2 - \alpha \log \alpha - (1 - \alpha) \log(1 - \alpha) + \frac{d}{k} \log(1 - 2^{1-k}(1 - e^{-\beta}) + 2^{-k}\alpha^k(1 - e^{-\beta})^2).$$

Thus, the maximiser α in (2.6) represents the overlap value that renders the dominant contribution to the second moment. Since the entropy function $-\alpha \log \alpha - (1 - \alpha) \log(1 - \alpha)$ attains its maximum at $\alpha = 1/2$, while the term $\log(1 - 2^{1-k}(1 - e^{-\beta}) + 2^{-k}\alpha^k(1 - e^{-\beta})^2)$ is strictly increasing in α , the dominant overlap value inevitably exceeds $1/2$. In effect, since $f(1/2)$ equals twice the right-hand side of (2.1) and $\max_{\alpha} f(\alpha) > f(1/2)$, the second moment $\mathbb{E}[Z(\Phi, \beta)^2]$ exceeds the square $\mathbb{E}[Z(\Phi, \beta)]^2$ of the first moment by an exponential factor for all $d, \beta > 0$. While it is still possible to salvage *some* estimate of $\log Z(\Phi, \beta)$ from (2.1)–(2.6), its quality deteriorates rapidly as β increases. In the extreme case $\beta = \infty$ of “hard” constraints this issue was already highlighted in the seminal work [3] where Achlioptas and Moore pioneered the second moment method for random k -SAT.

To remedy this problem we take a leaf out of earlier work on “hard” random k -SAT [5, 19]. Instead of applying the second moment method directly to $Z(\Phi, \beta)$, we consider a suitably truncated random variable. Its second moment is asymptotically bounded by the square of the first moment and we obtain the following explicit lower bound.

PROPOSITION 2.1. *Let $\beta \geq 1, k \geq k_0$ and $d < d^*$ and let $p \in (0, 1)$ be the unique root of*

$$(2.7) \quad 1 - 2p - (1 - e^{-\beta})(1 - p)^k = 0;$$

then

$$(2.8) \quad \liminf_{n \rightarrow \infty} \frac{1}{n} \mathbb{E}[\log Z(\Phi, \beta)] \geq \left(1 - \frac{(k-1)d}{k}\right) \log 2 - \frac{d}{2} \log p - \frac{d}{2} \log(1 - p) + \frac{d}{k} \log p.$$

We apply Proposition 2.1 to estimate the Boltzmann marginals. More precisely, recalling that σ signifies a sample from $\mu_{\Phi, \beta}$, we care to learn the marginal probabilities $\mu_{\Phi, \beta}(\{\sigma_{x_i} = 1\})$ that specific variables x_i take the value “true”. Hence, with δ_z denoting the probability measure on \mathbb{R} that places mass one on the number z , let

$$\pi_{\Phi, \beta} = \frac{1}{n} \sum_{i=1}^n \delta_{\mu_{\Phi, \beta}(\{\sigma_{x_i}=1\})} \in \mathcal{P}(0, 1)$$

be the empirical distribution of these marginals. We say that a probability measure π on $[0, 1]$ has *slim tails* if

$$(2.9) \quad \pi \left(\left[0, \frac{1}{2} - 2^{-k/10}\right] \cup \left[\frac{1}{2} + 2^{-k/10}, 1\right] \right) \leq 2^{-k/10}.$$

Additionally, π has *very slim tails* if (2.9) holds with the right-hand side replaced by $2^{-k/9}$.

COROLLARY 2.2. *Suppose that $\beta \geq 1$, $k \geq k_0$ and $d < d^*$ and that (1.8) is satisfied. Then $\pi_{\Phi, \beta}$ has very slim tails a.a.s.*

The proofs of Proposition 2.1 and Corollary 2.2 can be found in Section 5.

2.2. *Belief propagation on trees.* As a next step we analyse belief propagation on a Galton–Watson tree that mimics the local structure of the random formula Φ . To elaborate, recall the factor graph representation $G(\Phi)$ of Φ from Section 1.2. For a variable x_j we denote by ∂x_j the set of adjacent clauses. Moreover, to keep track of the order as variables appear in clauses we write $\partial_h a_i$ for the h th variable in clause a_i and ∂a_i for the set of all variables that occur in a_i . Finally, for an adjacent clause/variable pair (a, x) we let $J_{ax} = \text{sign}(a, x) \in \{\pm 1\}$ signify the sign with which x appears in a .

The graph $G(\Phi)$ induces a metric on the set of variables and clauses. Moreover, it is well known that $G(\Phi)$ contains only a small number of, say, $o(\log n)$ cycles of bounded length. Hence, for any specific variable x_i and for any fixed radius $t > 0$ the depth- t neighbourhood of x_i in $G(\Phi)$ is a tree a.a.s. The distribution of this tree can be characterised precisely by a two-type Galton–Watson process. The two types are variables and clauses, of course. The process starts from a single root variable x_0 . Moreover, the offspring of a variable is a $\text{Po}(d)$ number of clauses. Furthermore, a clause begets $k - 1$ variables. Let T signify the resulting (quite possibly infinite) tree. Also let $V(T)$, $C(T)$ be the sets of variables and clauses of T , respectively. As in the case of the random formula Φ we use the ∂ -symbol to denote adjacencies. Finally, to turn T into a k -SAT formula, we choose for each adjacent clause/variable pair $(a, x) \in C(T) \times V(T)$ a sign $J_{ax} \in \{\pm 1\}$ uniformly and independently.

It is well known that the graph $G(\Phi)$ converges locally to the random tree T in the sense that for any specific variable node x_i , $1 \leq i \leq n$, and for any fixed radius t the depth- t neighbourhood of x_i and the depth- t neighbourhood of the root x_0 of T can be coupled such that both coincide a.a.s. Therefore, in order to investigate the first t rounds of belief propagation Φ as per (1.4)–(1.6), we just need to study belief propagation on T .

Hence, we proceed to define belief propagation messages on T . Generalising (1.4), we allow for an arbitrary probability distribution π on $[0, 1]$ from which we draw the initial messages. Thus, for any adjacent a, x we draw $\mu_{T, \beta, \pi, x \rightarrow a, 0}(1)$, $\mu_{T, \beta, \pi, a \rightarrow x, 0}(1)$ independently from π and set

$$\mu_{T, \beta, \pi, x \rightarrow a, 0}(-1) = 1 - \mu_{T, \beta, \pi, x \rightarrow a, 0}(1), \quad \mu_{T, \beta, \pi, a \rightarrow x, 0}(-1) = 1 - \mu_{T, \beta, \pi, a \rightarrow x, 0}(1).$$

Further, for $t \geq 0$, $s = \pm 1$ and adjacent a, x we inductively define

$$(2.10) \quad \mu_{T, \beta, \pi, a \rightarrow x, t+1}(s) \propto \sum_{\sigma \in \{\pm 1\}^{\partial a}} \mathbb{1}\{\sigma_x = s\} e^{-\beta \mathbb{1}\{\sigma \neq a\}} \prod_{y \in \partial a \setminus \{x\}} \mu_{T, \beta, \pi, y \rightarrow a, t}(\sigma_y),$$

$$(2.11) \quad \mu_{T, \beta, \pi, x \rightarrow a, t+1}(s) \propto \prod_{b \in \partial x \setminus \{a\}} \mu_{T, \beta, \pi, b \rightarrow x, t+1}(s).$$

Finally, the belief propagation estimate of the marginal of x_0 after $t + 1$ rounds reads

$$(2.12) \quad \mu_{T, \beta, \pi, x_0, t+1}(s) \propto \prod_{b \in \partial x_0} \mu_{T, \beta, \pi, b \rightarrow x_0, t+1}(s) \quad (s = \pm 1).$$

Let $\pi_0 = \delta_{1/2}$ be the probability distribution on $(0, 1)$ that places all mass on $1/2$. The following proposition shows that in the limit of large t , any distribution with slim tails yields the same belief propagation marginal as π_0 .

PROPOSITION 2.3. Assume that $d \leq d_{\text{SAT}}(k)$ and $\beta \geq 1$. Then uniformly for all π with slim tails and $k \geq k_0$ we have

$$\lim_{t \rightarrow \infty} \mathbb{E} |\mu_{\mathcal{T}, \beta, \pi, x_0, t}(1) - \mu_{\mathcal{T}, \beta, \pi_0, x_0, t}(1)| = 0.$$

Furthermore, the sequence $(\mu_{\mathcal{T}, \beta, \pi_0, x_0, t}(1))_{t \geq 1}$ converges weakly to a probability measure $\pi_{d, \beta}^*$ with slim tails.

The proof of Proposition 2.3 can be found in Section 6.

2.3. *The Bethe free energy.* It is known that under the replica symmetry assumption (1.8) the partition function $Z(\Phi, \beta)$ can be approximated well in terms of certain “pseudo-messages”. To be precise, consider a clause a_i and a variable x_{ij} that appears in it. Then we define the pseudo-message $\mu_{\Phi, \beta, x_{ij} \rightarrow a_i}$ as the Boltzmann marginal of x_{ij} in the formula $\Phi - a_i$ obtained by deleting clause a_i . Thus, $\mu_{\Phi, \beta, x_{ij} \rightarrow a_i}(\pm 1) = \mu_{\Phi - a_i, \beta}(\{\sigma_{x_{ij}} = \pm 1\})$. Similarly, we define the reverse message $\mu_{\Phi, \beta, a_i \rightarrow x_{ij}}$ as the marginal of x_{ij} in the formula obtained from Φ by deleting all clauses in which the variable x_{ij} appears apart from a_i . In symbols,

$$\mu_{\Phi, \beta, a_i \rightarrow x_{ij}}(s) = \mu_{\Phi - (\partial x_{ij} \setminus \{a_i\}), \beta}(\{\sigma_{x_{ij}} = s\}) \quad (s = \pm 1).$$

A result about general random factor graph models from [20] implies that the pseudo-messages yield the following approximation to the partition function if (1.8) is satisfied.

LEMMA 2.4 ([20], Corollary 1.2). Let

$$\begin{aligned} \mathcal{B}(\Phi, \beta) &= \sum_{i=1}^n \log \left[\sum_{\sigma = \pm 1} \prod_{a \in \partial x_i} \mu_{\Phi, \beta, a \rightarrow x_i}(\sigma) \right] \\ (2.13) \quad &+ \sum_{i=1}^m \log \left[\sum_{\sigma \in \{\pm 1\}^{\partial a_i}} \exp(-\beta \mathbb{1}\{\sigma \neq a_i\}) \prod_{x \in \partial a_i} \mu_{\Phi, \beta, x \rightarrow a_i}(\sigma_x) \right] \\ &- \sum_{i=1}^n \sum_{a \in \partial x_i} \log \left[\sum_{\sigma = \pm 1} \mu_{\Phi, \beta, x_i \rightarrow a}(\sigma) \mu_{\Phi, \beta, a \rightarrow x_i}(\sigma) \right]. \end{aligned}$$

If (1.8) holds and $\lim_{n \rightarrow \infty} \mathcal{B}(\Phi, \beta)/n = b \in \mathbb{R}$ in probability, then $\lim_{n \rightarrow \infty} \frac{1}{n} \log Z(\Phi, \beta) = b$ in probability.

Apart from the formula (2.13) it is known that the pseudo-messages form an approximate fixed point of the belief propagation recurrence (1.5)–(1.6) if (1.8) is satisfied [20], Theorem 1.1. In light of the contraction property of the belief propagation iteration that Proposition 2.3 provides it is therefore tempting to think that the messages obtained after a large enough number t of iterations of (1.4)–(1.6) should be about the same as the pseudo-messages. Yet this conclusion is anything but immediate. First, Proposition 2.3 establishes contraction only under the assumption that belief propagation launches from a set of *independent* initial messages. Second, the proposition requires that the distribution of these initial messages has slim tails. Thus, for all we know the belief propagation equations on Φ could have many fixed points that do not fall into the basin of attraction of the all- $\frac{1}{2}$ initialisation (1.4). Nonetheless, combining a subtle coupling argument with the tail bound for $\pi_{\Phi, \beta}$ from Corollary 2.2 we can link the pseudo-message with the belief propagation fixed point that we approach from the naive initialisation (1.4). We can thus relate the pseudo-messages and the real ones as follows.

PROPOSITION 2.5. *If (1.8) is satisfied, $d < d^*$, $\beta \geq 1$ and $k \geq k_0$, then*

$$\lim_{t \rightarrow \infty} \limsup_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} \left[\sum_{i=1}^n \sum_{a \in \partial x_i} |\mu_{\Phi, \beta, x_i \rightarrow a}(1) - \mu_{\Phi, \beta, x_i \rightarrow a, t}(1)| + |\mu_{\Phi, \beta, a \rightarrow x_i}(1) - \mu_{\Phi, \beta, a \rightarrow x_i, t}(1)| \right] = 0.$$

Equipped with Lemma 2.4 and Proposition 2.5 we have only very little work left to complete the proof of Theorem 1.1. Indeed, Lemma 2.4 shows how $\mathcal{B}(\Phi, \beta)$ approximates $\log Z(\Phi, \beta)$. Moreover, the only difference between $\mathcal{B}(\Phi, \beta)$ and the expression $\mathcal{B}_t(\Phi, \beta)$ that appears in Theorem 1.1 is that the latter comes in terms of the belief propagation messages $\mu_{\Phi, \beta, x \rightarrow a, t}, \mu_{\Phi, \beta, a \rightarrow x, t}$ rather than the actual standard messages $\mu_{\Phi, \beta, x \rightarrow a}, \mu_{\Phi, \beta, a \rightarrow x}$. But Proposition 2.5 shows that the belief propagation messages approximate the standard messages well. Hence, we are left to show that the approximation is good enough and that the Bethe free energy is sufficiently continuous. We will carry these steps out in Section 7, thereby completing the proof of Theorem 1.1.

2.4. *Replica symmetry breaking.* The proof of Theorem 1.2 consists of two parts: an unconditional upper bound on $\mathbb{E}[\log Z(\Phi, \beta)]$ and a lower bound that is conditional on the assumption (1.8) of replica symmetry. To state these bounds we define a functional $\mathfrak{B}_{d, \beta}$ on the space of probability measures on the unit interval. This functional can be viewed as the $n \rightarrow \infty$ limit of the Bethe free energy functional from (1.7).

Hence, let π be a probability measure on $[0, 1]$. Let $(\rho_{\pi, i, j})_{i, j \geq 1}$ be an array of independent random variables with distribution π . Furthermore, let $(J_{i, j})_{i, j \geq 1}$ be an array of Rademacher variables with mean zero, mutually independent and independent of the $\rho_{\pi, i, j}$. Additionally, let

$$(2.14) \quad \mu_{\pi, i, j} = \frac{1 + J_{i, j}(2\rho_{\pi, i, j} - 1)}{2} = \begin{cases} \rho_{\pi, i, j} & \text{if } J_{i, j} = 1, \\ 1 - \rho_{\pi, i, j} & \text{if } J_{i, j} = -1. \end{cases}$$

Further, let γ^+, γ^- be two $\text{Po}(d/2)$ variables, mutually independent and independent of everything else. We define

$$(2.15) \quad \mathfrak{B}_{d, \beta}(\pi) = \mathbb{E} \left[\log \left(\prod_{i=1}^{\gamma^+} 1 - (1 - e^{-\beta}) \prod_{j=1}^{k-1} \mu_{\pi, i, j} + \prod_{i=1}^{\gamma^-} 1 - (1 - e^{-\beta}) \prod_{j=1}^{k-1} \mu_{\pi, i + \gamma^-, j} \right) - \frac{d(k-1)}{k} \log \left(1 - (1 - e^{-\beta}) \prod_{j=1}^k \mu_{\pi, 1, j} \right) \right].$$

Using the so-called 1-step replica symmetry breaking interpolation method from mathematical physics, we obtain the following upper bound. Recall $\pi_{d, \beta}^*$ from Proposition 2.3.

PROPOSITION 2.6. *Assume that d satisfies (1.11) and that $\beta > \beta_0(k)$ for a large enough $\beta_0(k)$. Then*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E}[\log Z(\Phi, \beta)] < \mathfrak{B}_{d, \beta}(\pi_{d, \beta}^*).$$

We remark that the limit on the left-hand side is known to exist [40].

On the other hand, a careful study facilitated by the ideas from the proof of Proposition 2.1 and by Proposition 2.3 shows that under the assumption (1.8) even in the regime $d^*(k) < d <$

$d_{\text{SAT}}(k)$ the only conceivable scenario is that the actual pseudo-messages nearly coincide with the messages that result from the fixed point iteration (1.4)–(1.6). Combining this fact with Lemma 2.4, we obtain the following lower bound.

PROPOSITION 2.7. *Assume that d satisfies (1.11) and that $\beta \geq \beta_0(k)$ for a large enough $\beta_0(k)$. If (1.8) holds, then*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E}[\log Z(\Phi, \beta)] \geq \mathfrak{B}_{d,\beta}(\pi_{d,\beta}^*).$$

Naturally, Propositions 2.6 and 2.7 imply that (1.8) cannot be satisfied under the assumptions of Theorem 1.2. The detailed proofs of the propositions as well as of Theorem 1.2 can be found in Section 8.

3. Discussion. Early experimental work [13] inspired the hunt for the satisfiability threshold, a pursuit conducted by many authors over many years (e.g., [4–6, 11, 12, 19, 27–29]) and which culminated in the aforementioned work of Ding, Sly and Sun [24]. Prior to the seminal work of Achlioptas and Moore [3], lower bounds on the k -SAT threshold were based on the analysis of simple satisfiability algorithms [6, 11, 12, 28]. However, all known algorithms fail to find satisfying assignments efficiently for densities well below the satisfiability threshold. The best algorithmic result today reaches up to $d \sim 2^k \log k$ [15], undershooting $d_{\text{SAT}}(k)$ by almost a factor of k . Even satisfiability algorithms that employ message passing techniques such as belief propagation guided decimation fail to beat this bound [16, 32]. Furthermore, the current algorithmic threshold of $2^k \log k$ marks the onset of combinatorial effects that conceivably stymie various algorithmic techniques [2].

Matters get worse when it comes to algorithms for counting or sampling satisfying assignments. From a worst-case viewpoint this task is the epitome of the complexity class $\#\text{P}$, the counterpart of the notorious complexity class NP in the realm of counting and sampling [47]. Hence, we expect that counting or sampling satisfying assignments is conceptually far more challenging than “merely” finding one. In the context of random k -SAT a first important contribution is due to Montanari and Shah [38], who showed that belief propagation does the trick up to $d \sim \log k$. Their analysis is based on Gibbs uniqueness, an extremely strong spatial mixing property that fails to hold for (much) larger d . In particular, Gibbs uniqueness implies condition (1.8). Moreover, in the case $k = 2$ Gibbs uniqueness holds up to the satisfiability threshold [1].

Recently Galanis, Goldberg, Guo and Yang [30] proposed a fully polynomial-time approximation scheme for computing $Z(\Phi, \infty)$ for large enough k and $d \leq 2^{k/301}$. While avoiding an explicit connection to the replica symmetry assumption (1.8), they seize upon the technique of Moitra [36] for approximately counting satisfying assignments of formulas with bounded variable degrees. It would be interesting to see if this approach extends to finite β and if a proof of (1.8) can be salvaged from the techniques from [30, 36].

A key feature of the k -SAT problem that goes a long way to explaining the technical difficulty of the model is that the marginals of the Boltzmann distribution vary from one variable to another. This manifests itself in the fact that the limiting distribution $\pi_{d,\beta}^*$ from Proposition 2.3 is nontrivial [37]. The distribution is generally a mixture of a discrete and a continuous probability measure. As a result, we do not expect that there is a simple analytic expression for the limiting value of the Bethe free energy in Theorem 1.1. On a technical level one of the main contributions of this article is that we manage to deal with the inherent asymmetry of the problem. By comparison, the replica symmetric regime of symmetric problems where the relevant belief propagation fixed point is trivial (e.g., the uniform distribution) is well understood. In this case the counterpart of Theorem 1.1 is essentially trivial

and the existence and location of the replica symmetry breaking phase have been established precisely [9, 17, 22]. But of course quite a few prominent problems do not possess symmetry. For instance, apart from the k -SAT model the hard-core model on random graphs springs to mind.

Beyond probabilistic combinatorics and computer science, the random k -SAT model has been studied as a diluted spin glass model. Using a combination of the interpolation method and spatial mixing arguments, Panchenko and Talagrand [41, 43, 46] studied the model in the case of very low d and/or β , known as the “high-temperature” version of the model. Additionally, Panchenko [40] obtained a variational formula for $\lim_{n \rightarrow \infty} n^{-1} \mathbb{E}[\log Z(\Phi, \beta)]$. However, this formula does not easily reveal the connection with belief propagation, nor the existence or location of the replica symmetry breaking phase transition. Generally speaking, the analysis of diluted models with sparse interactions appears to lead to less explicit solutions than in the case of models with full interactions such as the Sherrington–Kirkpatrick model [39].

In a few models that bear similarity with random k -SAT it has been possible to move beyond the replica symmetric phase. For example, the 1-step replica symmetry breaking phase has been investigated in detail in the random regular k -NAESAT model, where even the existence of a Gardner (or full replica symmetry breaking) transition has been established rigorously [10, 45]. The work of Sly, Sun and Zhang [45] on the 1-step RSB formula for the free energy combines the interpolation method with the second moment method. Moreover, the proof design from [10] is somewhat reminiscent of the strategy that we pursue here to establish Theorem 1.2, but the details are very different. More specifically, conceptually [10] deals with a more complex question, namely 1-step versus 2-step replica symmetry breaking, whereas here we are concerned with replica symmetry versus 1-step replica symmetry breaking. That said, the model that we study here is more intricate than the regular k -NAESAT model, which enjoys relatively strong symmetry properties.

What the present strategy and that pursued in [10] have in common is that we use contraction techniques to pin down the conceivable solutions to the simpler recurrence (replica symmetry in our case and 1-step RSB in [10]). In addition, both the present work and [10] use the interpolation method to obtain a contradiction. Yet a difference is that here the reference point of the analysis is the replica symmetry condition (1.8), whereas the starting point in [10] is the 1RSB formula for the ground state energy. In particular, the proof of Theorem 1.2 requires a lower bound as well as an upper bound on the partition function. Moreover, in order to refute the replica symmetric scenario we need to investigate two different conceivable solutions to replica symmetric ansatz, respectively two fixed points of belief propagation. The first of these corresponds to the scenario that typical pairs of satisfying assignments are essentially orthogonal; this case we tackle via the contraction method. The second scenario is that typical Boltzmann samples have a very high overlap. This case requires delicate combinatorial expansion arguments.

Finally, Panchenko [42] studied the random k -SAT model in the limit of large d . The main result is that $n^{-1} \mathbb{E}[\log Z(\Phi, \beta)]$ approaches the solution to a certain fully connected k -spin model in the limit of large d ; that is, a model of Sherrington–Kirkpatrick type that lives on a complete hypergraph. These models are currently better understood than models on sparse random graphs and, in particular, there are formulas for the approximate value of the partition function that match the so-called full replica symmetry breaking predictions from physics [39]. However, the regime of d where Panchenko’s results bite is well beyond the k -SAT threshold and, indeed, a full replica symmetry breaking regime is not expected to occur in the satisfiable phase of the random k -SAT model [35].

4. Preliminaries and notation. A k -SAT formula Φ with a set $V(\Phi)$ of variables and a set $C(\Phi)$ of clauses can be represented by a bipartite graph $G(\Phi)$ known as the *factor graph*. In this graph there is an edge between $x \in V(\Phi)$ and $a \in C(\Phi)$ iff x occurs in a . For a vertex v of the graph we let ∂v denote the set of neighbours. Furthermore, for an integer $\ell \geq 1$ we let $\partial^\ell v$ be the set of vertices at distance precisely ℓ from v . Where necessary we annotate Φ to clarify the reference to the formula. In addition, for a formula Φ and an assignment $\sigma \in \{\pm 1\}^{V(\Phi)}$ we let

$$\mathcal{H}_\Phi(\sigma) = \sum_{a \in C(\Phi)} \mathbb{1}\{\sigma \not\models a\}$$

be the number of clauses that σ violates; the function \mathcal{H}_Φ is known as the *Hamiltonian* in physics jargon.

We always write $V_n = \{x_1, \dots, x_n\}$ for the variable set of the random formula Φ . For each of the corresponding $2n$ literals $x_i, \neg x_i$ we denote by d_i^\pm the degree of that literal, that is, the number of clauses where the literal appears. For the entire literal degree sequence we introduce the symbol $\mathbf{d} = (d_i^\pm)_{i \in [n]}$. Additionally, let \mathfrak{D} be the σ -algebra generated by \mathbf{d} ; observe that the total number m of clauses is \mathfrak{D} -measurable.

We will use the following important theorem about the random k -SAT model.

THEOREM 4.1 ([24]). *There exist a number $k_0 > 3$ and a sequence $d_{\text{SAT}}(k) = k2^k \log 2 - k(1 + \log 2)/2 + o(1)$ such that for all $k \geq k_0$ a.a.s. Φ has a satisfying assignment if $d < d_{\text{SAT}}(k)$ and fails to possess one if $d > d_{\text{SAT}}(k)$.*

Theorem 4.1 implies that $Z(\Phi, \beta) \geq Z(\Phi, \infty) \geq 1$ for all $\beta > 0$ and all $d < d_{\text{SAT}}(k)$ a.a.s.

Throughout the paper we will be dealing a fair bit with probability distributions on discrete cubes. For finite sets $\Omega, V \neq \emptyset$ we let $\mathcal{P}(\Omega^V)$ be the set of all probability measures on Ω^V . For a set $U \subset V$ and $\mu \in \mathcal{P}(\Omega^V)$ we let $\mu_U \in \mathcal{P}(\Omega^U)$ be the joint distribution of the coordinates $u \in U$ under μ . If $U = \{u_1, \dots, u_\ell\}$ is given explicitly, we use the shorthand $\mu_U = \mu_{u_1, \dots, u_\ell}$. Moreover, a distribution $\mu \in \mathcal{P}(\Omega^V)$ is (ε, ℓ) -extremal if

$$\sum_{U \subset V: |U|=\ell} d_{\text{TV}}\left(\mu_U, \bigotimes_{u \in U} \mu_u\right) < \varepsilon \binom{|V|}{\ell}.$$

Thus, for most sets $U \subset V$ with ℓ elements, the joint distribution μ_U is close in total variation to the product distribution with the same marginals $\mu_u, u \in U$. If $\ell = 2$ we just call μ ε -extremal. Hence, because Φ is invariant under permutations of the variables, the condition (1.8) posits that the Boltzmann distribution $\mu_{\Phi, \beta}$ is $o(1)$ -extremal a.a.s. We will apply the following result repeatedly.

LEMMA 4.2 ([8]). *For any $\Omega \neq \emptyset, \varepsilon > 0, \ell \geq 3$ there exists $\delta > 0$ such that for all sets V with $|V| > 1/\delta$ any δ -extremal $\mu \in \mathcal{P}(\Omega^V)$ is (ε, ℓ) -extremal.*

Let $\mu, \nu \in \mathcal{P}(\Omega^V)$ and $c > 0$. We say that μ is c -contiguous w.r.t. ν if $\mu(\mathcal{E}) \leq c\nu(\mathcal{E})$ for any $\mathcal{E} \subset \Omega^V$.

LEMMA 4.3 ([21]). *For any $\Omega \neq \emptyset, c > 0, \varepsilon > 0$ there exists $\delta > 0$ such that for all sets V with $|V| > 1/\delta$, any δ -extremal $\mu \in \mathcal{P}(\Omega^V)$ and any $\nu \in \mathcal{P}(\Omega^V)$ that is c -contiguous w.r.t. μ the following statements are true.*

- (i) ν is ε -extremal.
- (ii) $\sum_{v \in V} d_{\text{TV}}(\mu_v, \nu_v) < \varepsilon|V|$.

For a probability measure μ on a discrete space Ω and function $X : \Omega \rightarrow \mathbb{R}$ we introduce the bracket notation

$$\langle X, \mu \rangle = \sum_{\omega \in \Omega} X(\omega)\mu(\omega).$$

Thus, $\langle X, \mu \rangle$ is the mean of X w.r.t. μ . More generally, if $Y : \Omega^\ell \rightarrow \mathbb{R}$ for some $\ell \geq 1$, then

$$\langle Y, \mu \rangle = \sum_{\omega_1, \dots, \omega_\ell \in \Omega} Y(\omega_1, \dots, \omega_\ell) \prod_{i=1}^\ell \mu(\omega_i)$$

is the expectation of Y w.r.t. $\mu^{\otimes \ell}$.

Apart from discrete distributions we will also be working with spaces of continuous probability measures. For a Polish space \mathfrak{E} let $\mathcal{P}(\mathfrak{E})$ be the space of all probability measures on \mathfrak{E} . In addition, for a subspace $\mathfrak{E} \subset \mathbb{R}$ we introduce the L^r -Wasserstein space $\mathcal{W}_r(\mathfrak{E})$ as the space of all probability distributions $\mu \in \mathcal{P}(\mathfrak{E})$ with $\int_{\mathfrak{E}} |x|^r d\mu(x) < \infty$. We endow this space with the Wasserstein metric W_r , thereby turning $\mathcal{W}_r(\mathfrak{E})$ into a complete metric space. We recall that the Wasserstein metric is defined as

$$W_r(\mu, \nu) = \inf \left\{ \left(\int_{\mathfrak{E} \times \mathfrak{E}} |x - y|^r d\gamma(x, y) \right)^{1/r} : \gamma \in \mathcal{P}(\mathfrak{E} \times \mathfrak{E}) \text{ is a coupling of } \mu, \nu \right\}.$$

Throughout the paper we use the O -notation to refer to asymptotics as either n, k or β get large. By default O -symbols refer to the limit as $n \rightarrow \infty$. However, where the expression inside the $O(\cdot)$ -symbol depends on k or β but not on n , it is understood that we mean to take the respective parameter to infinity instead of n . Where there is a risk of ambiguity we make the reference explicit by adding a subscript. Thus, $O_k(1)$ stands for an expression that remains bounded in the limit of large k . Naturally, the same conventions also apply to $o(\cdot), \Omega(\cdot), \Theta(\cdot)$. In addition, we use symbols such as $\tilde{O}(\cdot)$ to suppress logarithmic factors. For example, $\tilde{O}(n)$ can be written out as $O(n \log^{O(1)} n)$ while $\tilde{O}(2^k)$ stands for a term of order $k^{O(1)} 2^k$. Moreover, in the entire paper we tacitly assume that k, n exceed large enough absolute constants k_0, n_0 , respectively.

We also need a few basic large deviations inequalities. Recall that the Kullback–Leibler divergence of two probability measures $\mu, \nu \in \mathcal{P}(\Omega)$ is defined as

$$D_{\text{KL}}(\mu \parallel \nu) = \sum_{\omega \in \Omega} \mu(\omega) \log \frac{\mu(\omega)}{\nu(\omega)} \in [0, \infty],$$

with the conventions $0 \log 0 = 0 \log \frac{0}{0} = 0$.

LEMMA 4.4 (“Chernoff bound”). *Suppose that X has a binomial distribution $\text{Bin}(n, p)$. Then*

$$\begin{aligned} \mathbb{P}[X \geq qn] &\leq \exp(-nD_{\text{KL}}(\text{Be}(q) \parallel \text{Be}(p))) \quad \text{if } q > p, \\ \mathbb{P}[X \leq qn] &\leq \exp(-nD_{\text{KL}}(\text{Be}(q) \parallel \text{Be}(p))) \quad \text{if } q < p. \end{aligned}$$

LEMMA 4.5 (“Bennett’s inequality”). *Suppose that X is a $\text{Po}(\lambda)$ variable. Then*

$$\begin{aligned} \mathbb{P}(X \geq \lambda + x) &\leq \exp\left(x - (\lambda + x) \log\left(1 + \frac{x}{\lambda}\right)\right) \\ (4.1) \qquad \qquad &\leq \exp\left(-\frac{x^2}{2\lambda + 2x/3}\right) \quad \text{for any } x \geq 0, \end{aligned}$$

$$\begin{aligned}
 \mathbb{P}(X \leq \lambda - x) &\leq \exp\left(-x - (\lambda - x) \log\left(1 - \frac{x}{\lambda}\right)\right) \\
 &\leq \exp\left(-\frac{x^2}{2\lambda}\right) \text{ for any } 0 \leq x < \lambda.
 \end{aligned}
 \tag{4.2}$$

Finally, we remind ourselves of the well-known fact that belief propagation “is exact on trees”. To be precise, let T be a k -SAT formula whose factor graph $G(T)$ is a tree. Then we can introduce belief propagation on T via (2.10)–(2.12). The following statement provides that for large enough t these recurrences render the Boltzmann marginals of T .

THEOREM 4.6 ([34], Theorem 14.4). *Assume T is a k -SAT instance whose factor graph $G(T)$ is a tree and that t exceeds the diameter of $G(T)$. Then for all variables x of T and any $\beta > 0$ we have $\mu_{T,\beta}(\{\sigma_x = 1\}) = \mu_{T,\beta,x,t}(1)$.*

5. Moment calculations. In this section we prove Proposition 2.1 and Corollary 2.2. Unless specified otherwise we tacitly assume that $d < d^*$. Moreover, recalling the definition (2.7) of $p = p(k, \beta)$, we introduce

$$u = u(k, \beta) = \frac{1 - 2p}{2p(e^\beta - 1)} \in (0, 1).
 \tag{5.1}$$

5.1. Overview. As we saw in Section 2.1, we cannot hope to prove Proposition 2.1 by simply calculating the second moment of the partition function $Z(\Phi, \beta)$. This is because the expression (2.6) for the second moment attains its maximum at a value of α strictly greater than $1/2$. To solve this problem we will replace $Z(\Phi, \beta)$ by a modified random variable for which the overlap value $\alpha = 1/2$ dominates by design. The precise construction of this random variable borrows an idea from the work of Achlioptas and Peres [5] on k -SAT with hard constraints (i.e., $\beta = \infty$). Namely, we call an assignment $\sigma \in \{\pm 1\}^{V_n}$ *balanced* if

$$\sum_{x \in V_n} \sigma_x (d_x^+ - d_x^-) = \begin{cases} 0 & \text{if } km \text{ is even,} \\ 1 & \text{otherwise.} \end{cases}
 \tag{5.2}$$

Hence, if we inspect the truth values of the km literals as they appear in the m clauses, we observe as many “true” as “false” literals, up to an additive error of one. Further, we call a balanced assignment σ *strongly balanced* if

$$\left| \sum_{x \in V_n} \sigma_x \mathbb{1}\{d_x^+ = d^+, d_x^- = d^-\} \right| \leq \sqrt{n} \text{ for all integers } d^+, d^- \geq 0.
 \tag{5.3}$$

Thus, under a strongly balanced assignment about half the variables with each possible degree constellation (d^+, d^-) are set to “true”, up to an error of $O(\sqrt{n})$. Let BAL denote the set of all strongly balanced assignments. Now our modified version of the partition function reads

$$Z_{\text{bal}}(\Phi, \beta) = \exp(-\beta um) \sum_{\sigma \in \text{BAL}} \mathbb{1}\left\{\sum_{i=1}^m \mathbb{1}\{\sigma \not\equiv a_i\} = \lceil um \rceil\right\}.$$

Thus, we confine ourselves to strongly balanced assignments that leave precisely $\lceil um \rceil$ clauses unsatisfied. Naturally, it will emerge in due course that the choice (5.1) of u maximises the mean of $Z_{\text{bal}}(\Phi, \beta)$. The following two propositions, which we prove in Sections 5.2 and 5.3, render the first and the second moment of $Z_{\text{bal}}(\Phi, \beta)$.

PROPOSITION 5.1. *A.a.s. we have*

$$\frac{1}{n} \log \mathbb{E}[Z_{\text{bal}}(\Phi, \beta) \mid \mathfrak{D}] = \left(1 - \frac{(k-1)d}{k}\right) \log 2 - \frac{d}{2} \log(p(1-p)) + \frac{d}{k} \log p + o(1).$$

PROPOSITION 5.2. *A.a.s. we have*

$$\frac{1}{2n} \log \mathbb{E}[Z_{\text{bal}}(\Phi, \beta)^2 \mid \mathcal{D}] = \left(1 - \frac{(k-1)d}{k}\right) \log 2 - \frac{d}{2} \log(p(1-p)) + \frac{d}{k} \log p + o(1).$$

The proofs of Propositions 5.1–5.2 are generalisations of the moment calculations from [22], where assignments that satisfy *all* clauses were counted. A significant complication here is that a certain number of clauses are left unsatisfied. This introduces a further dimension to the second moment analysis, namely the number of clauses that are left unsatisfied under both assignments, leaving us with a technically far more challenging task. Proposition 2.1 is an easy consequence of Propositions 5.1 and 5.2 and the Paley–Zygmund and Azuma–Hoeffding inequalities.

PROOF OF PROPOSITION 2.1. The Paley–Zygmund inequality implies that

$$\mathbb{P}[Z_{\text{bal}}(\Phi, \beta) \geq \mathbb{E}[Z_{\text{bal}}(\Phi, \beta) \mid \mathcal{D}]/4 \mid \mathcal{D}] \geq \frac{\mathbb{E}[Z_{\text{bal}}(\Phi, \beta) \mid \mathcal{D}]^2}{4\mathbb{E}[Z_{\text{bal}}(\Phi, \beta)^2 \mid \mathcal{D}]}.$$

Hence, Proposition 5.2 shows that a.a.s.

$$(5.4) \quad \mathbb{P}[Z_{\text{bal}}(\Phi, \beta) \geq \mathbb{E}[Z_{\text{bal}}(\Phi, \beta) \mid \mathcal{D}]/4 \mid \mathcal{D}] \geq \exp(o(n)).$$

Further, combining (5.4) with Proposition 5.1 and using the trivial inequality $Z(\Phi, \beta) \geq Z_{\text{bal}}(\Phi, \beta)$, we obtain

$$(5.5) \quad \mathbb{P}\left[n^{-1} \log Z(\Phi, \beta) \geq \left(1 - \frac{(k-1)d}{k}\right) \log 2 - \frac{d}{2} \log(p(1-p)) + \frac{d}{k} \log p + o(1)\right] \geq \exp(o(n)).$$

Moreover, because adding or removing a single clause can alter the value of the partition function by no more than a factor of $\exp(\pm\beta)$, the Azuma–Hoeffding inequality shows that for any $t > 0$,

$$(5.6) \quad \mathbb{P}[|\log Z(\Phi, \beta) - \mathbb{E}[\log Z(\Phi, \beta) \mid \mathbf{m}]| \geq t \mid \mathbf{m}] \leq 2 \exp(-t^2/(2\beta^2 \mathbf{m})).$$

Thus, combining (5.5) and (5.6), we conclude that

$$\liminf_{n \rightarrow \infty} n^{-1} \mathbb{E}[\log Z(\Phi, \beta)] \geq \left(1 - \frac{(k-1)d}{k}\right) \log 2 - \frac{d}{2} \log(p(1-p)) + \frac{d}{k} \log p + o(1),$$

as desired. \square

Let us move on to the proof of Corollary 2.2 concerning the tails of the distribution of Boltzmann marginals. Combining Proposition 2.1 with the Azuma–Hoeffding inequality as in (5.6), we conclude that

$$(5.7) \quad \begin{aligned} & Z(\Phi, \beta) \\ & \geq \exp\left(n \left[\left(1 - \frac{(k-1)d}{k}\right) \log 2 - \frac{d}{2} \log(p(1-p)) + \frac{d}{k} \log p + o(1) \right]\right) \quad \text{a.a.s.} \end{aligned}$$

Of course, this bound directly yields a lower bound on the corresponding sum over pairs of assignments, namely

$$(5.8) \quad \begin{aligned} & Z(\Phi, \beta)^2 \\ & = \sum_{\sigma, \tau} e^{-\beta \sum_{i=1}^m \mathbb{1}\{\sigma \neq a_i\} + \mathbb{1}\{\tau \neq a_i\}} \\ & \geq \exp\left(2n \left[\left(1 - \frac{(k-1)d}{k}\right) \log 2 - \frac{d}{2} \log(p(1-p)) + \frac{d}{k} \log p + o(1) \right]\right) \quad \text{a.a.s.} \end{aligned}$$

Let us compare this bound with the expansion (2.5) of the second moment. The contribution to (2.5) of a specific overlap value α is bounded by $\exp(n(f(\alpha) + o(1)))$. Comparing these estimates carefully, we will discover that the total contribution of all overlap values α that differ significantly from $1/2$ is tiny by comparison to (5.8). As a consequence, a.a.s. the overlap of two independent random samples σ, σ' drawn from the Boltzmann distribution must be close to $1/2$. The following corollary provides a precise statement of this observation.

COROLLARY 5.3. *We have $\mathbb{E}[\mu_{\Phi, \beta}(\{|\alpha(\sigma, \sigma') - 1/2| > k^9 2^{-k/2}\})] = o(1)$.*

Finally, Corollary 2.2 is an easy consequence of Corollary 5.3 and general facts about Boltzmann distributions. The details can be found in Section 5.4.

5.2. *Proof of Proposition 5.1.* As a first step we estimate the number of balanced assignments.

LEMMA 5.4. *A.a.s. we have $|\text{BAL}| = 2^{n+o(n)}$.*

PROOF. The Chernoff bound shows that $d_x^+, d_x^- \leq \log n$ a.a.s. for all $x \in V_n$. Moreover, Chebyshev’s inequality easily shows that for a uniformly random $\zeta \in \{\pm 1\}^{V_n}$, for any $d^+, d^- \leq \log n$ we have

$$(5.9) \quad \mathbb{P}\left[\left|\sum_{x \in V_n} \zeta_x \mathbb{1}\{d_x^+ = d^+, d_x^- = d^-\}\right| \leq \sqrt{n} \mid \mathfrak{D}\right] = \Omega(1).$$

Consequently, ζ satisfies (5.3) with probability $\exp(O(\log^2 n))$. Further, the central limit theorem shows that a.a.s.

$$(5.10) \quad \mathbb{P}\left[\left|\sum_{x \in V_n} (1 - \mathbb{1}\{d_x^+ = 1, d_x^- = 0\}) \zeta_x (d_x^+ - d_x^-)\right| \leq \sqrt{n}/2 \mid \zeta \text{ satisfies (5.3), } \mathfrak{D}\right] = \Omega(1).$$

Finally, there are $\Theta(n)$ variables $x \in V_n$ such that $d_x^+ = 1, d_x^- = 0$ a.a.s. and hence Stirling’s formula shows that for any integer h with $|h| \leq \sqrt{n}/2$ we have

$$(5.11) \quad \mathbb{P}\left[\sum_{x \in V_n} \mathbb{1}\{d_x^+ = 1, d_x^- = 0\} \zeta_x = h \mid \zeta \text{ satisfies (5.3), } \mathfrak{D}\right] = \Omega(n^{-1/2}).$$

Since $\sum_{x \in V_n} (1 - \mathbb{1}\{d_x^+ = 1, d_x^- = 0\}) \zeta_x (d_x^+ - d_x^-)$ and $\sum_{x \in V_n} \mathbb{1}\{d_x^+ = 1, d_x^- = 0\} \zeta_x$ are conditionally independent given \mathfrak{D} , (5.9)–(5.11) imply that a.a.s. $\mathbb{P}[\zeta \in \text{BAL} \mid \mathfrak{D}] = \exp(o(n))$, whence the assertion is immediate. \square

Let us now fix any strongly balanced assignment σ . Given \mathfrak{D} and σ , the only randomness left is the way in which the positive and negative occurrences of the individual variables are matched to the clauses. To be precise, since we only need to know the number of clauses that will be left unsatisfied, we do not care about the identity of the underlying variable of a literal in a given clause, but only about its truth value. Therefore, we can think of the positive and negative variable occurrences as tokens that are labelled either “true” or “false”. Hence, a variable x gives rise to d_x^+ “true” and d_x^- “false” tokens if $\sigma_x = 1$, and to d_x^+ “false” and d_x^- “true” tokens if $\sigma_x = -1$. In effect, we just need to study the number of clauses that receive k “false” tokens if we put the km tokens down randomly upon the m clauses. In fact, since σ is strongly balanced, we know that the precise number of “true” tokens equals

$$(5.12) \quad \sum_{x \in V_n} d_x^+ \mathbb{1}\{\sigma_x = 1\} + d_x^- \mathbb{1}\{\sigma_x = -1\} = \lceil km/2 \rceil.$$

Of course, the remaining $\lfloor km/2 \rfloor$ tokens must be “false”.

To study this token shuffling experiment we introduce an auxiliary probability space. Let $(\chi_{i,j})_{i,j \geq 1}$ be a family of Rademacher variables such that $\mathbb{P}[\chi_{i,j} = 1] = p$ for all i, j . The idea is that $\chi_{i,1}, \dots, \chi_{i,k}$ represent the k tokens that clause i receives. Of course, in order to faithfully represent the token experiment we need to ensure that (5.12) is satisfied, that is, that the total number of $+1$ -tokens comes to $\lceil km/2 \rceil$. Thus, we need to condition on the event

$$\mathcal{B} = \left\{ \sum_{i=1}^m \sum_{j=1}^k \mathbb{1}\{\chi_{i,j} = 1\} = \lceil km/2 \rceil \right\}.$$

We need to compute the conditional probability that given \mathcal{B} the total number of clauses that only receive -1 -tokens equals $\lceil um \rceil$. Hence, introducing

$$(5.13) \quad \mathcal{S} = \left\{ \sum_{i=1}^m \mathbb{1}\left\{ \max_{j \in [k]} \chi_{i,j} = -1 \right\} = \lceil um \rceil \right\},$$

we obtain

$$\mathbb{E}[Z_{\text{bal}}(\Phi, \beta) \mid \mathcal{D}] = \exp(-\beta um) |\text{BAL}| \cdot \frac{\mathbb{P}[\mathcal{S} \cap \mathcal{B} \mid \mathbf{m}]}{\mathbb{P}[\mathcal{B} \mid \mathbf{m}]}.$$

Since Lemma 5.4 already shows that $|\text{BAL}| = 2^{n+o(n)}$, the remaining challenge is to calculate $\mathbb{P}[\mathcal{S} \mid \mathcal{B}, \mathcal{D}]$. To this end we calculate $\mathbb{P}[\mathcal{B} \mid \mathcal{D}]$, $\mathbb{P}[\mathcal{S} \mid \mathcal{D}]$ and $\mathbb{P}[\mathcal{B} \mid \mathcal{S}, \mathcal{D}]$ and use Bayes' formula.

LEMMA 5.5. We have $\mathbb{P}[\mathcal{B} \mid \mathcal{D}] = \binom{km}{\lceil km/2 \rceil} p^{\lceil km/2 \rceil} (1-p)^{km - \lceil km/2 \rceil}$.

PROOF. This is because the random variables $\chi_{i,j}$ are mutually independent. \square

LEMMA 5.6. We have $\mathbb{P}[\mathcal{S} \mid \mathcal{D}] = \binom{m}{\lceil um \rceil} (1-p)^{k \lceil um \rceil} (1 - (1-p)^k)^{m - \lceil um \rceil}$.

PROOF. Because the $\chi_{i,j}$ are independent, for any given index $i \in [m]$ we have $\mathbb{P}[\max_{j \in [k]} \chi_{i,j} = -1] = (1-p)^k$, independently of all others. Thus, the number of $i \in [m]$ with $\max_{j \in [k]} \chi_{i,j} = -1$ has distribution $\text{Bin}(m, (1-p)^k)$. \square

LEMMA 5.7. We have $\mathbb{P}[\mathcal{B} \mid \mathcal{S}, \mathcal{D}] = \Theta(m^{-1/2})$.

PROOF. Due to (2.7) and the choice of u a.a.s. we have

$$\begin{aligned} \mathbb{E} \left[\sum_{i=1}^m \sum_{j=1}^k \mathbb{1}\{\chi_{i,j} = 1\} \mid \mathcal{S}, \mathcal{D} \right] &= \frac{(m - \lceil um \rceil)kp}{1 - (1-p)^k} \\ &= km \left(1 - \frac{(1-p)^k}{2p \exp(\beta)} \right) \cdot \frac{p}{1 - (1-p)^k} + O(1) \\ &= \frac{km}{2} + O(1). \end{aligned}$$

Therefore, the assertion follows from the local limit theorem for sums of independent random variables. \square

PROOF OF PROPOSITION 5.1. Combining Lemmas 5.4, 5.5, 5.6 and 5.7, we conclude that a.a.s.

$$\begin{aligned} \log \mathbb{E}[Z_{\text{bal}}(\Phi, \beta) \mid \mathcal{D}] &= (n - km) \log 2 - \frac{km}{2} \log(p(1 - p)) \\ &\quad + kum \log(1 - p) + (1 - u)m \log(1 - (1 - p)^k) \\ &\quad - \beta um + \log \binom{m}{\lceil um \rceil} + o(n) \\ &= n \left[(1 - d) \log 2 - \frac{d}{2} \log p + d \left(\frac{(1 - p)^k}{2p \exp(\beta)} - \frac{1}{2} \right) \log(1 - p) \right. \\ &\quad \left. + \frac{d}{k} \left(1 - \frac{(1 - p)^k}{2p \exp(\beta)} \right) \log(1 - (1 - p)^k) \right. \\ &\quad \left. - \frac{\beta du}{k} - \frac{d}{k} (u \log(u) + (1 - u) \log(1 - u)) + o(1) \right]. \end{aligned}$$

Simplifying the above using the definition of u and (2.7) yields the desired expression. \square

5.3. Proof of Proposition 5.2. The weighted overlap of two truth assignments $\sigma, \tau \in \{\pm 1\}^{V_n}$ is defined as

$$\omega(\sigma, \tau) = \frac{1}{km} \sum_{x \in V_n} \mathbb{1}\{\sigma_x = \tau_x = 1\} d_x^+ + \mathbb{1}\{\sigma_x = \tau_x = -1\} d_x^-.$$

Thus, the weighted overlap equals the fraction of literal occurrences that evaluate to “true” under both σ, τ . Let $\mathcal{O} = \mathcal{O}(d) = \{\omega(\sigma, \tau) : \sigma, \tau \in \{\pm 1\}^{V_n}\}$ be the set of all conceivable weighted overlaps. Introducing

$$\begin{aligned} \mathbf{E}(\omega) &= \sum_{\sigma, \tau \in \text{BAL}} \mathbb{1}\{\omega(\sigma, \tau) = \omega\} \exp(-2\beta um) \\ (5.14) \quad &\cdot \mathbb{P} \left[\sum_{i=1}^m \mathbb{1}\{\sigma \neq a_i\} = \sum_{i=1}^m \mathbb{1}\{\tau \neq a_i\} = \lceil um \rceil \mid \mathcal{D} \right], \end{aligned}$$

we can then write the second moment as $\mathbb{E}[Z_{\text{bal}}(\Phi, \beta)^2 \mid \mathcal{D}] = \sum_{\omega \in \mathcal{O}} \mathbf{E}(\omega)$.

We will use two separate arguments to estimate $\mathbf{E}(\omega)$ for different regimes of ω . The first regime that we consider is ω close to $1/4$. This will turn out to be the dominant case.

PROPOSITION 5.8. A.a.s.,

$$\max\{\mathbf{E}(\omega) : \omega \in \mathcal{O}, |\omega - 1/4| \leq k^{100} 2^{-k/2}\} \leq \exp(o(n)) \mathbb{E}[Z_{\text{bal}}(\Phi, \beta) \mid \mathcal{D}]^2.$$

The proof of Proposition 5.8 can be found in Section 5.3.1. Moving on to weighted overlaps far from $1/4$, we will derive the following bound on $\mathbf{E}(\omega)$ in terms of the function $f(\alpha)$ from (2.6).

PROPOSITION 5.9. A.a.s.

$$\max\{\mathbf{E}(\omega) : \omega \in \mathcal{O}, |\omega - 1/4| > k^{100} 2^{-k/2}\} \leq \exp(n \max\{f(\alpha) : \alpha \in [1/2 + k^{90} 2^{-k/2}, 1]\}).$$

We prove Proposition 5.9 in Section 5.3.2. Finally, in Section 5.3.3 we will bound $f(\alpha)$ as follows.

PROPOSITION 5.10. *We have*

$$\begin{aligned} & \max\{f(\alpha) : \alpha \in [1/2 + k^{90}2^{-k/2}, 1]\} \\ & < 2(1 - (k - 1)d/k) \log 2 - d \log(p(1 - p)) + 2d \log(p)/k. \end{aligned}$$

Proposition 5.2 is an easy consequence of Propositions 5.9–5.10.

PROOF OF PROPOSITION 5.2. Combining Propositions 5.1, 5.9 and 5.10, we conclude that a.a.s.

$$(5.15) \quad \max\{\mathbf{E}(\omega) : \omega \in \mathcal{O}, |\omega - 1/4| > k^{100}2^{-k/2}\} \leq \exp(-\Omega(n))\mathbb{E}[Z_{\text{bal}}(\Phi, \beta) \mid \mathfrak{D}]^2.$$

Since $|\mathcal{O}| = O(n)$ a.a.s., Proposition 5.8, (5.14) and (5.15) imply that

$$\mathbb{E}[Z_{\text{bal}}(\Phi, \beta)^2 \mid \mathfrak{D}] \leq \exp(o(n))\mathbb{E}[Z_{\text{bal}}(\Phi, \beta) \mid \mathfrak{D}]^2$$

a.a.s., as desired. \square

5.3.1. *Proof of Proposition 5.8.* As in the proof of Proposition 5.1 we begin by estimating the number of pairs of balanced assignments with a given weighted overlap. Subsequently we will switch to an auxiliary probability space to calculate the probability that both such assignments happen to violate precisely $\lceil um \rceil$ clauses. Hence, draw $\tau, \tau' \in \text{BAL}$ uniformly and independently.

LEMMA 5.11. *A.a.s. we have $\mathbb{P}[|\omega(\tau, \tau') - 1/4| > \varepsilon \mid \mathfrak{D}] \leq 2 \exp(-\frac{\varepsilon^2 m^2}{4d^2 n} + o(n))$ for all $\varepsilon > 0$.*

PROOF. Let $\zeta, \zeta' \in \{\pm 1\}^{V_n}$ be drawn uniformly and independently. Then Lemma 5.4 implies that a.a.s.

$$(5.16) \quad \mathbb{P}[|\omega(\tau, \tau') - 1/4| > \varepsilon \mid \mathfrak{D}] \leq \exp(o(n))\mathbb{P}[|\omega(\zeta, \zeta') - 1/4| > \varepsilon \mid \mathfrak{D}].$$

Furthermore, since the pairs $(\zeta_x, \zeta'_x) \in \{\pm 1\}^2$ are mutually independent and changing (ζ_x, ζ'_x) can alter $\omega(\zeta, \zeta')$ by at most $d_x/(km)$, the Azuma–Hoeffding inequality yields

$$(5.17) \quad \mathbb{P}[|\omega(\zeta, \zeta') - 1/4| > \varepsilon \mid \mathfrak{D}] \leq 2 \exp\left(-\frac{\varepsilon^2 (km)^2}{2 \sum_{x \in V_n} d_x^2}\right).$$

Finally, since $\sum_{x \in V_n} d_x^2 \leq 2d^2 n$ a.a.s., (5.17) and (5.16) imply the assertion. \square

Like in Section 5.2 we now fix any two assignments τ, τ' with a given weighted overlap ω such that

$$(5.18) \quad |\omega - 1/4| \leq k^{100}2^{-k/2}.$$

Given $\tau, \tau', \mathfrak{D}$, the experiment of actually constructing the random formula Φ boils down to matching the km literal slots in the m clauses with the positive/negative occurrences of the variables x_1, \dots, x_n . But once again we do not actually care to know the identities of the literals in the various clauses, but only their truth value combinations under τ, τ' . Hence, instead of actually matching literals to clauses, we might as well think of merely tossing tokens that indicate the truth value combinations $(1, 1), (1, -1), (-1, 1), (-1, -1)$ of the literals onto the clauses. To be precise, because τ, τ' are balanced, the fractions of tokens of each of the four types work out to be

$$(5.19) \quad \begin{aligned} \omega_{11} &= \omega, & \omega_{1-1} &= \omega_{-11} = \frac{1}{2} - \omega_{11} + \frac{\mathbb{1}\{km \text{ is odd}\}}{km}, \\ \omega_{-1-1} &= 1 - \omega_{11} - 2\omega_{1-1}. \end{aligned}$$

Hence, we just need to work out the probability that if we randomly put down $km\omega_{11}$, $km\omega_{1-1}$, $km\omega_{-11}$, $km\omega_{-1-1}$ tokens of these four types onto the m clauses, precisely $\lceil um \rceil$ clauses will receive k tokens of type either $(-1, 1)$ or $(-1, -1)$ and, symmetrically, precisely $\lceil um \rceil$ clauses will receive tokens of type $(1, -1)$ or $(-1, -1)$ only.

As in the first moment calculation, in order to calculate this probability it is convenient to move to an auxiliary probability space. Specifically, let $(p_{11}, p_{1-1}, p_{-11}, p_{-1-1}) \in (0, 1)^4$ be a probability distribution on $\{\pm 1\}^2$, that is, $p_{11} + p_{1-1} + p_{-11} + p_{-1-1} = 1$, such that $p_{1-1} = p_{-11}$; we will choose expedient values of p_{11}, \dots, p_{-1-1} in due course. Moreover, let $(\chi_{ij}, \chi'_{ij})_{i,j \geq 1}$ be a sequence of i.i.d. random pairs $(\chi_{ij}, \chi'_{ij}) \in \{\pm 1\}^2$ such that

$$(5.20) \quad \mathbb{P}[\chi_{ij} = s, \chi'_{ij} = t] = p_{st} \quad (s, t = \pm 1).$$

Further, let

$$\mathcal{B}^\otimes = \left\{ \sum_{i=1}^m \sum_{j=1}^k \chi_{ij} = \sum_{i=1}^m \sum_{j=1}^k \chi'_{ij} = \mathbb{1}\{km \text{ is odd} \} \right\}$$

and

$$\mathcal{R}^\otimes(\omega) = \left\{ \sum_{i=1}^m \sum_{j=1}^k \mathbb{1}\{\chi_{ij} = \chi'_{ij} = 1\} = \omega km \right\} \cap \mathcal{B}^\otimes.$$

Then the sequence $(\chi_{ij}, \chi'_{ij})_{i \in [m], j \in [k]}$ given $\mathcal{R}^\otimes(\omega)$ is distributed precisely as a random sequence of $\{\pm 1\}^2$ tokens comprising precisely $km\omega_{11}, \dots, km\omega_{-1-1}$ tokens of each type. Hence, letting

$$\mathcal{S}^\otimes = \left\{ \sum_{i=1}^m \mathbb{1}\left\{ \max_{j \in [k]} \chi_{ij} = -1 \right\} = \sum_{i=1}^m \mathbb{1}\left\{ \max_{j \in [k]} \chi'_{ij} = -1 \right\} = \lceil um \rceil \right\},$$

we obtain

$$(5.21) \quad \mathbb{E}[E(\omega) \mid \mathcal{D}] = |\text{BAL}|^2 \mathbb{P}[\omega(\boldsymbol{\tau}, \boldsymbol{\tau}') = \omega \mid \mathcal{D}] \cdot \mathbb{P}[\mathcal{S}^\otimes \mid \mathcal{R}^\otimes(\omega)] \exp(-2\beta um).$$

Since Lemmas 5.4 and 5.11 already yield the first two factors on the right-hand side, we are left to compute $\mathbb{P}[\mathcal{S}^\otimes \mid \mathcal{R}^\otimes(\omega)]$.

As in the first moment calculations we will calculate this conditional probability via Bayes' formula. Hence, we begin by computing the unconditional probabilities of $\mathcal{R}^\otimes(\omega)$ and \mathcal{S}^\otimes .

LEMMA 5.12. *We have*

$$\begin{aligned} \mathbb{P}[\mathcal{R}^\otimes(\omega) \mid \mathcal{D}] &= \exp(o(n)) \binom{km}{\omega km, (1/2 - \omega)km, (1/2 - \omega)km, \omega km} \\ &\cdot (p_{11}p_{-1-1})^{\omega km} (p_{1-1}p_{-11})^{(1/2-\omega)km}. \end{aligned}$$

PROOF. This is an immediate consequence of the fact that the random pairs $(\chi_{ij}, \chi'_{ij})_{i,j}$ are mutually independent and that the individual pairs (χ_{ij}, χ'_{ij}) are drawn from the distribution (5.20). \square

To calculate the probability of \mathcal{S}^\otimes we need to introduce one extra parameter. Namely, for $s \in [0, u]$ we let

$$\mathcal{S}^\otimes(s) = \mathcal{S}^\otimes \cap \left\{ \sum_{i=1}^m \mathbb{1}\left\{ \max_{j \in [k]} \chi_{ij} = \max_{j \in [k]} \chi'_{ij} = -1 \right\} = \lceil sm \rceil \right\}.$$

Thus, s specifies the fraction of clauses that receive $(-1, -1)$ tokens only. Of course, we have the bound

$$(5.22) \quad \mathbb{P}[\mathcal{S}^\otimes] \leq m \max_{s \in [0, u]} \mathbb{P}[\mathcal{S}^\otimes(s)].$$

LEMMA 5.13. *For any $s \in [0, u]$ we have*

$$\begin{aligned} \mathbb{P}[\mathcal{S}^\otimes(s) \mid \mathfrak{D}] &= \exp(o(n)) \binom{m}{sm, (u-s)m, (u-s)m, (1-2u+s)m} \\ &\cdot p_{-1-1}^{ksm} ((p_{-1-1} + p_{1-1})^k - p_{-1-1}^k)^{2(u-s)m} \\ &\cdot (1 - (p_{1-1} + p_{-1-1})^k - (p_{-11} + p_{-1-1})^k + p_{-1-1}^k)^{(1-2u+s)m}. \end{aligned}$$

PROOF. The vectors $(\chi_{i1}, \chi'_{i1}, \dots, \chi_{ik}, \chi'_{ik})_{i \in [m]}$ are mutually independent. Moreover, (5.20) provides that

$$(5.23) \quad \mathbb{P}[\chi_{i1} = \chi'_{i1} = \dots = \chi_{ik} = \chi'_{ik} = -1] = p_{-1-1}^k,$$

$$(5.24) \quad \mathbb{P}[\chi_{i1} = \dots = \chi_{ik} = -1 \wedge \exists j \in [k] : \chi'_{ij} = 1] = (p_{-11} + p_{-1-1})^k - p_{-1-1}^k,$$

$$(5.25) \quad \mathbb{P}[\chi'_{i1} = \dots = \chi'_{ik} = -1 \wedge \exists j \in [k] : \chi_{ij} = 1] = (p_{-11} + p_{-1-1})^k - p_{-1-1}^k,$$

$$(5.26) \quad \mathbb{P}[\exists j, l \in [k] : \chi_{ij} = \chi'_{il} = 1] = 1 - (p_{1-1} + p_{-1-1})^k - (p_{-11} + p_{-1-1})^k + p_{-1-1}^k.$$

Since $\mathcal{S}^\otimes(s)$ asks that (5.23) occur for $\lceil sm \rceil$ indices $i \in [m]$, that (5.24) and (5.25) occur for $\lceil um \rceil - \lceil sm \rceil$ indices and that, naturally, (5.26) occur for the remaining $m - 2\lceil um \rceil + \lceil sm \rceil$ indices, we obtain the assertion. \square

As in the first moment calculation we are going to apply Bayes' rule

$$(5.27) \quad \mathbb{P}[\mathcal{S}^\otimes \mid \mathcal{R}^\otimes(\omega)] = \frac{\mathbb{P}[\mathcal{S}^\otimes(\omega)]}{\mathbb{P}[\mathcal{R}^\otimes(\omega)]} \cdot \mathbb{P}[\mathcal{R}^\otimes(\omega) \mid \mathcal{S}^\otimes]$$

to calculate the probability on the left-hand side. But while we easily obtained succinct expressions for the unconditional probabilities $\mathbb{P}[\mathcal{S}^\otimes(\omega)]$ and $\mathbb{P}[\mathcal{R}^\otimes(\omega)]$ for any choice of p_{11}, \dots, p_{-1-1} , calculating $\mathbb{P}[\mathcal{R}^\otimes(\omega) \mid \mathcal{S}^\otimes]$ for general choices of these parameters appears to be tricky. Yet it turns out that for a diligent choice of the $p_{\pm 1 \pm 1}$ we will obtain $\mathbb{P}[\mathcal{R}^\otimes(\omega) \mid \mathcal{S}^\otimes] = \exp(o(n))$. To work out these $p_{\pm 1 \pm 1}$, we need to calculate the conditional expectation of the number of pairs (χ_{ij}, χ'_{ij}) for which specific $(\pm 1, \pm 1)$ -values materialise given \mathcal{S}^\otimes . Thus, for $v, w = \pm 1$ let

$$(5.28) \quad X_{vw} = \frac{1}{km} \sum_{i=1}^m \sum_{j=1}^k \mathbb{1}\{\chi_{ij} = v, \chi'_{ij} = w\}.$$

For brevity we introduce $p_1 = p_{11} + p_{1-1} = p_{11} + p_{-11}$ and $p_{-1} = p_{-11} + p_{-1-1} = p_{1-1} + p_{-1-1}$.

LEMMA 5.14. *A.a.s. for any $s \in [0, u]$ we have*

$$(5.29) \quad \mathbb{E}[X_{11} \mid \mathfrak{D}, \mathcal{S}^\otimes(s)] = \frac{(1 - 2u + s)p_{11}}{1 - 2p_{-1}^k + p_{-1-1}^k} + O(1/n),$$

$$(5.30) \quad \begin{aligned} \mathbb{E}[X_{1-1} \mid \mathfrak{D}, \mathcal{S}^\otimes(s)] &= \frac{(u-s)p_{1-1}p_{-1}^{k-1}}{p_{-1}^k - p_{-1-1}^k} + \frac{(1-2u+s)p_{1-1}(1-p_{-1}^{k-1})}{1 - 2p_{-1}^k + p_{-1-1}^k} \\ &+ O(1/n). \end{aligned}$$

PROOF. By linearity of expectation we just need to contemplate the $\{\pm 1\}^2$ -sequence $(\chi_{11}, \chi'_{11}), \dots, (\chi_{1k}, \chi'_{1k})$ that represents the first clause. Given $\mathcal{S}^\otimes(s)$ the event $\max_{j \in [k]} \chi_{1j} = \max_{j \in [k]} \chi'_{1j} = 1$ has probability $1 - 2u + s + O(1/m) = 1 - 2u + s + O(1/n)$. Furthermore, the conditional probability that $\chi_{11} = \chi'_{11} = 1$ given $\max_{j \in [k]} \chi_{1j} = \max_{j \in [k]} \chi'_{1j} = 1$ equals $p_{11}/(1 - 2p_{-1}^k + p_{-1-1}^k)$ because the pairs $(\chi_{1j}, \chi'_{1j})_j$ are mutually independent, whence we obtain (5.29).

Similar steps yield (5.30). For with probability $u - s + O(1/m)$ we have $\max_{j \in [k]} \chi_{1j} = -\max_{j \in [k]} \chi'_{1j} = 1$ and given this event the probability that $(\chi_{11}, \chi'_{11}) = (1, -1)$ equals $p_{1-1}p_{-1}^{k-1}/(p_{-1}^k - p_{-1-1}^k)$; hence the first summand in (5.30). Further, as in the previous paragraph, the event $\max_{j \in [k]} \chi_{1j} = \max_{j \in [k]} \chi'_{1j} = 1$ has probability $1 - 2u + s + O(1/m)$ and then the conditional probability of $(\chi_{11}, \chi'_{11}) = (1, -1)$ works out to be $p_{1-1}(1 - p_{-1}^{k-1})/(1 - 2p_{-1}^k + p_{-1-1}^k)$. \square

LEMMA 5.15. For any $r \in [1/4 - 2^{-k/3}, 1/4 + 2^{-k/3}]$, $s \in [0, u]$ the system of equations

$$(5.31) \quad \frac{(1 - 2u + s)p_{11}}{1 - 2p_{-1}^k + p_{-1-1}^k} = r,$$

$$\frac{(u - s)p_{1-1}p_{-1}^{k-1}}{p_{-1}^k - p_{-1-1}^k} + \frac{(1 - 2u + s)p_{1-1}(1 - p_{-1}^{k-1})}{1 - 2p_{-1}^k + p_{-1-1}^k} = 1/2,$$

$$(5.32) \quad p_{1-1} = p_{-11},$$

$$p_{11} + p_{1-1} + p_{-11} + p_{-1-1} = 1$$

possesses a unique solution $p_{\pm 1 \pm 1} \in [\frac{1}{4} - 2^{-k/3-1}, \frac{1}{4} + 2^{-k/3-1}]$.

PROOF. The proof is based on the inverse function theorem. Implementing the last two constraints (5.32), we substitute $p_{-11} = p_{1-1}$ and $p_{-1-1} = 1 - 2p_{1-1} - p_{11}$. Hence, the remaining free variables are s, p_{11}, p_{1-1} and we need to work out the Jacobi matrix of the map

$$g : \mathbb{R}^3 \rightarrow \mathbb{R}^3, \quad (p_{11}, p_{1-1}, s) \mapsto (g_1, g_2, g_3),$$

where

$$g_1 = \frac{(1 - 2u + s)p_{11}}{1 - 2p_{-1}^k + p_{-1-1}^k}, \quad g_2 = \frac{(u - s)p_{1-1}p_{-1}^{k-1}}{p_{-1}^k - p_{-1-1}^k} + \frac{(1 - 2u + s)p_{1-1}(1 - p_{-1}^{k-1})}{1 - 2p_{-1}^k + p_{-1-1}^k},$$

$$g_3 = s.$$

The partial derivatives of g_1 come to

$$(5.33) \quad \frac{\partial}{\partial p_{11}} \frac{(1 - 2u + s)p_{11}}{1 - 2p_{-1}^k + p_{-1-1}^k} = \frac{1 - 2u + s}{1 - 2p_{-1}^k + p_{-1-1}^k}$$

$$- \frac{k(1 - 2u + s)p_{11}(2p_{-1}^{k-1} - p_{-1-1}^{k-1})}{(1 - 2p_{-1}^k + p_{-1-1}^k)^2},$$

$$(5.34) \quad \frac{\partial}{\partial p_{1-1}} \frac{(1 - 2u + s)p_{11}}{1 - 2p_{-1}^k + p_{-1-1}^k} = - \frac{2(1 - 2u + s)k(p_{-1}^{k-1} - p_{-1-1}^{k-1})}{(1 - 2p_{-1}^k + p_{-1-1}^k)^2},$$

$$(5.35) \quad \frac{\partial}{\partial s} \frac{(1 - 2u + s)p_{11}}{1 - 2p_{-1}^k + p_{-1-1}^k} = \frac{p_{11}}{1 - 2p_{-1}^k + p_{-1-1}^k}.$$

Moreover, the first summand of g_2 has derivatives

$$(5.36) \quad \frac{\partial}{\partial p_{11}} \frac{(u-s)p_{1-1}p_{-1}^{k-1}}{p_{-1}^k - p_{-1-1}^k} = -\frac{(k-1)(u-s)p_{1-1}p_{-1}^{k-2}}{p_{-1}^k - p_{-1-1}^k} + \frac{k(u-s)p_{1-1}p_{-1}^{k-1}(p_{-1}^{k-1} - p_{-1-1}^{k-1})}{(p_{-1}^k - p_{-1-1}^k)^2},$$

$$(5.37) \quad \frac{\partial}{\partial p_{1-1}} \frac{(u-s)p_{1-1}p_{-1}^{k-1}}{p_{-1}^k - p_{-1-1}^k} = \frac{(u-s)(p_{-1}^{k-1} - (k-1)p_{1-1}p_{-1}^{k-2})}{p_{-1}^k - p_{-1-1}^k} + \frac{k(u-s)p_{1-1}p_{-1}^{k-1}(p_{-1}^{k-1} - 2p_{-1-1}^{k-1})}{(p_{-1}^k - p_{-1-1}^k)^2},$$

$$(5.38) \quad \frac{\partial}{\partial s} \frac{(u-s)p_{1-1}p_{-1}^{k-1}}{p_{-1}^k - p_{-1-1}^k} = -\frac{p_{1-1}p_{-1}^{k-1}}{p_{-1}^k - p_{-1-1}^k}.$$

Finally, for the second summand we obtain

$$(5.39) \quad \begin{aligned} \frac{\partial}{\partial p_{11}} \frac{(1-2u+s)p_{1-1}(1-p_{-1}^{k-1})}{1-2p_{-1}^k + p_{-1-1}^k} &= \frac{(1-2u+s)(k-1)p_{1-1}p_{-1}^{k-2}}{1-2p_{-1}^k + p_{-1-1}^k} \\ &\quad - \frac{(1-2u+s)kp_{1-1}(1-p_{-1}^{k-1})(2p_{-1}^{k-1} - p_{-1-1}^{k-1})}{(1-2p_{-1}^k + p_{-1-1}^k)^2}, \end{aligned}$$

$$(5.40) \quad \begin{aligned} \frac{\partial}{\partial p_{1-1}} \frac{(1-2u+s)p_{1-1}(1-p_{-1}^{k-1})}{1-2p_{-1}^k + p_{-1-1}^k} &= \frac{(1-2u+s)(1-p_{-1}^{k-1} + 2(k-1)p_{1-1}p_{-1}^{k-2})}{1-2p_{-1}^k + p_{-1-1}^k} \\ &\quad - \frac{2(1-2u+s)kp_{1-1}(1-p_{-1}^{k-1})(p_{-1}^{k-1} - p_{-1-1}^{k-1})}{(1-2p_{-1}^k + p_{-1-1}^k)^2}, \end{aligned}$$

$$(5.41) \quad \frac{\partial}{\partial s} \frac{(1-2u+s)p_{1-1}(1-p_{-1}^{k-1})}{1-2p_{-1}^k + p_{-1-1}^k} = \frac{p_{1-1}(1-p_{-1}^{k-1})}{1-2p_{-1}^k + p_{-1-1}^k}.$$

Hence, for $p_{11} = \frac{1}{4} + O(k^{-4})$, $p_{1-1} = \frac{1}{4} + O(k^{-4})$ we obtain

$$(5.42) \quad Dg = \begin{pmatrix} 1 + \tilde{O}(2^{-k}) & \tilde{O}(2^{-k}) & p_{11} + \tilde{O}(2^{-k}) \\ \tilde{O}(2^{-k}) & 1 + \tilde{O}(2^{-k}) & p_{1-1} + \tilde{O}(2^{-k}) \\ 0 & 0 & 1 \end{pmatrix}.$$

Consequently, (5.51) and the inverse function theorem yield

$$(5.43) \quad Dg^{-1} = \begin{pmatrix} 1 + \tilde{O}(2^{-k}) & \tilde{O}(2^{-k}) & -p_{11} + \tilde{O}(2^{-k}) \\ \tilde{O}(2^{-k}) & 1 + \tilde{O}(2^{-k}) & -p_{1-1} + \tilde{O}(2^{-k}) \\ 0 & 0 & 1 \end{pmatrix},$$

whence the assertion follows. \square

Let $\mathbf{p} = \mathbf{p}(r, s) = (p_{11}, p_{1-1}, p_{-11}, p_{-1-1})$ denote the solution to (5.31)–(5.32) provided by Lemma 5.15. By construction, $p_{-11} = p_{1-1}$ and $p_{-1-1} = 1 - 2p_{1-1} - p_{11}$. Let us make a note of the first and second derivatives of p_{11}, p_{1-1} .

COROLLARY 5.16. For all $0 \leq s \leq u$ and $\omega = \frac{1}{2} + O(2^{-k/3})$ we have

$$(5.44) \quad \frac{\partial p_{11}}{\partial \omega} = 1 + \tilde{O}(2^{-k}), \quad \frac{\partial p_{1-1}}{\partial \omega} = -1 + \tilde{O}(2^{-k}),$$

$$\frac{\partial p_{11}}{\partial s} = -p_{11} + \tilde{O}(2^{-k}), \quad \frac{\partial p_{1-1}}{\partial s} = -p_{1-1} + \tilde{O}(2^{-k}),$$

$$(5.45) \quad \frac{\partial^2 p_{11}}{\partial \omega^2}, \frac{\partial^2 p_{11}}{\partial \omega \partial s}, \frac{\partial^2 p_{1-1}}{\partial \omega^2}, \frac{\partial^2 p_{1-1}}{\partial \omega \partial s} = \tilde{O}(2^{-k}), \quad \frac{\partial^2 p_{11}}{\partial s^2}, \frac{\partial^2 p_{1-1}}{\partial s^2} = \tilde{O}(1).$$

PROOF. The assertions regarding the first derivative are immediate from the Jacobi matrix (5.43) of g . Furthermore, to obtain the bounds on the second derivatives we need to estimate the derivatives of the entries of Dg^{-1} with respect to ω, s . Let

$$\begin{aligned} \mathbf{a} &= \frac{\partial g_1}{\partial p_{11}} = 1 + \tilde{O}(2^{-k}), & \mathbf{b} &= \frac{\partial g_1}{\partial p_{1-1}} = \tilde{O}(2^{-k}), & \mathbf{c} &= \frac{\partial g_1}{\partial s} = -p_{11} + \tilde{O}(2^{-k}), \\ \mathbf{d} &= \frac{\partial g_2}{\partial p_{11}} = \tilde{O}(2^{-k}), & \mathbf{e} &= \frac{\partial g_2}{\partial p_{1-1}} = 1 + \tilde{O}(2^{-k}), & \mathbf{f} &= \frac{\partial g_2}{\partial s} = -p_{1-1} + \tilde{O}(2^{-k}) \end{aligned}$$

denote the entries in the first two rows of Dg . Revisiting the explicit expressions (5.33)–(5.41) for these partial derivatives, we verify that the second partial derivatives satisfy

$$\begin{aligned} &\frac{\partial \mathbf{a}}{\partial p_{11}}, \frac{\partial \mathbf{a}}{\partial p_{1-1}}, \frac{\partial \mathbf{b}}{\partial p_{11}}, \frac{\partial \mathbf{b}}{\partial p_{1-1}}, \frac{\partial \mathbf{c}}{\partial p_{11}}, \frac{\partial \mathbf{c}}{\partial p_{1-1}}, \frac{\partial \mathbf{d}}{\partial p_{11}}, \frac{\partial \mathbf{d}}{\partial p_{1-1}}, \frac{\partial \mathbf{e}}{\partial p_{11}}, \frac{\partial \mathbf{e}}{\partial p_{1-1}}, \frac{\partial \mathbf{f}}{\partial p_{11}}, \frac{\partial \mathbf{f}}{\partial p_{1-1}} \\ &= \tilde{O}(2^{-k}), \\ &\frac{\partial \mathbf{a}}{\partial s}, \frac{\partial \mathbf{d}}{\partial s}, \frac{\partial \mathbf{e}}{\partial s} = \tilde{O}(1), \quad \frac{\partial \mathbf{b}}{\partial s} = \tilde{O}(2^{-k}), \quad \frac{\partial \mathbf{c}}{\partial s}, \frac{\partial \mathbf{f}}{\partial s} = 0. \end{aligned}$$

Consequently, using (5.44) and the chain rule, we obtain

$$\begin{aligned} &\frac{\partial}{\partial \omega} \left[\left(\frac{\mathbf{e}}{\mathbf{ae} - \mathbf{bd}} \right) \Big|_{p_{11}=p_{11}, p_{1-1}=p_{1-1}} \right] \\ (5.46) \quad &= \frac{\frac{\partial \mathbf{e}}{\partial p_{11}} \Big|_{p_{11}=p_{11}, p_{1-1}=p_{1-1}} \frac{\partial p_{11}}{\partial \omega} + \frac{\partial \mathbf{e}}{\partial p_{1-1}} \Big|_{p_{11}=p_{11}, p_{1-1}=p_{1-1}} \frac{\partial p_{1-1}}{\partial \omega}}{\mathbf{ae} - \mathbf{bd}} \\ &+ \frac{\mathbf{e} \left(\frac{\partial(\mathbf{ae} - \mathbf{bd})}{\partial p_{11}} \Big|_{p_{11}=p_{11}, p_{1-1}=p_{1-1}} \frac{\partial p_{11}}{\partial \omega} + \frac{\partial(\mathbf{ae} - \mathbf{bd})}{\partial p_{1-1}} \Big|_{p_{11}=p_{11}, p_{1-1}=p_{1-1}} \frac{\partial p_{1-1}}{\partial \omega} \right)}{(\mathbf{ae} - \mathbf{bd})^2} \\ &= \tilde{O}(2^{-k}). \end{aligned}$$

Similarly,

$$(5.47) \quad \frac{\partial}{\partial \omega} \left[\left(\frac{\mathbf{b}}{\mathbf{ae} - \mathbf{bd}} \right) \Big|_{p_{11}=p_{11}, p_{1-1}=p_{1-1}} \right], \frac{\partial}{\partial \omega} \left[\left(\frac{\mathbf{d}}{\mathbf{ae} - \mathbf{bd}} \right) \Big|_{p_{11}=p_{11}, p_{1-1}=p_{1-1}} \right] = \tilde{O}(2^{-k}),$$

$$(5.48) \quad \frac{\partial}{\partial \omega} \left[\left(\frac{\mathbf{bf} - \mathbf{ce}}{\mathbf{ae} - \mathbf{bd}} \right) \Big|_{p_{11}=p_{11}, p_{1-1}=p_{1-1}} \right], \frac{\partial}{\partial \omega} \left[\left(\frac{\mathbf{a}}{\mathbf{ae} - \mathbf{bd}} \right) \Big|_{p_{11}=p_{11}, p_{1-1}=p_{1-1}} \right] = \tilde{O}(2^{-k}),$$

$$(5.49) \quad \frac{\partial}{\partial \omega} \left[\left(\frac{\mathbf{af} - \mathbf{cd}}{\mathbf{ae} - \mathbf{bd}} \right) \Big|_{p_{11}=p_{11}, p_{1-1}=p_{1-1}} \right] = \tilde{O}(2^{-k}),$$

$$(5.50) \quad \frac{\partial}{\partial s} \left[\left(\frac{bf - ce}{ae - bd} \right) \Big|_{p_{11}=p_{11}, p_{1-1}=p_{1-1}} \right], \frac{\partial}{\partial s} \left[\left(\frac{af - cd}{ae - bd} \right) \Big|_{p_{11}=p_{11}, p_{1-1}=p_{1-1}} \right] = \tilde{O}(1).$$

Finally, we remind ourselves of the following elementary formula: for any s, t, u, v, w, x such that $sw - tv \neq 0$,

$$(5.51) \quad \begin{pmatrix} s & t & u \\ v & w & x \\ 0 & 0 & 1 \end{pmatrix}^{-1} = \begin{pmatrix} \frac{w}{sw - tv} & -\frac{t}{sw - tv} & \frac{tx - uw}{sw - tv} \\ -\frac{v}{sw - tv} & \frac{s}{sw - tv} & -\frac{sx - uv}{sw - tv} \\ 0 & 0 & 1 \end{pmatrix}.$$

Combining (5.46)–(5.50) with the formula (5.51) for the Jacobian of g^{-1} , we obtain (5.45). □

We next verify that with the choice $p_{\pm 1 \pm 1} = p_{\pm 1 \pm 1}$ we may neglect the conditional probability $\mathbb{P}[\mathcal{R}(\omega) \mid \mathcal{D}, \mathcal{S}(s)]$.

LEMMA 5.17. *A.a.s. at the point $p_{\pm 1 \pm 1} = p_{\pm 1 \pm 1}$ we have*

$$\mathbb{P}[\mathcal{R}^\otimes(\omega) \mid \mathcal{D}, \mathcal{S}^\otimes(s)] = \exp(o(n)).$$

PROOF. Given $\mathcal{S}^\otimes(s)$ the random variables $X_{\pm 1 \pm 1}$ from (5.28) can be written as sums of m independent random variables. Indeed, because the pairs (χ_{ij}, χ'_{ij}) are identically distributed, instead of conditioning on $\mathcal{S}^\otimes(s)$ we may condition on the event $\mathcal{S}_0^\otimes(s)$ that:

- $\chi_{ij} = \chi'_{ij} = -1$ for $i = 1, \dots, \lceil sm \rceil, j \in [k]$,
- $\max_{j \in [k]} \chi_{ij} = 1$ and $\max_{j \in [k]} \chi'_{ij} = -1$ for $i = \lceil sm \rceil + 1, \dots, \lceil um \rceil$,
- $\max_{j \in [k]} \chi_{ij} = -1$ and $\max_{j \in [k]} \chi'_{ij} = 1$ for $i = \lceil um \rceil + 1, \dots, 2\lceil um \rceil - \lceil sm \rceil$,
- $\max_{j \in [k]} \chi_{ij} = \max_{j \in [k]} \chi'_{ij} = 1$ for $i = 2\lceil um \rceil - \lceil sm \rceil, \dots, m$.

Evidently, given $\mathcal{S}_0^\otimes(s)$ the random variables $X_{vw}(i) = \frac{1}{km} \sum_{j=1}^k \mathbb{1}\{\chi_{ij} = v, \chi'_{ij} = w\}$ with $v, w = \pm 1$ are independent for all $i \in [m]$. Moreover, $X_{vw} = \sum_{i=1}^m X_{vw}(i)$ and due to (5.29)–(5.30) the choice $p_{\pm 1 \pm 1} = p_{\pm 1 \pm 1}$ ensures that

$$(5.52) \quad \mathbb{E}[X_{vw} \mid \mathcal{S}_0^\otimes(s)] = \omega_{vw} + O(1/n).$$

Hence, assuming that $m \sim dn/k$ is about as large as its expectation, we can apply the local limit theorem for sums of independent random variables to the conditional random variables X_{vw} given $\mathcal{S}_{0,m}$ to conclude that

$$\begin{aligned} \mathbb{P}[\mathcal{R}^\otimes(\omega) \mid \mathcal{D}, \mathcal{S}^\otimes(s)] &= \mathbb{P}[\forall v, w \in \{\pm 1\} : X_{vw} = \omega_{vw} \mid \mathcal{D}, \mathcal{S}^\otimes(s)] \\ &= \Omega(n^{-3/2}) = \exp(o(n)), \end{aligned}$$

thereby completing the proof. □

Combining Lemmas 5.12, 5.13 and 5.17, we finally obtain a handy bound on $\mathbb{P}[\mathbf{E}(\omega) \mid \mathcal{D}]$. Indeed, keeping in mind our convention that $p_{-11} = p_{1-1}$ and $p_{-1-1} = 1 - p_{11} - 2p_{1-1}$, we introduce $\mathfrak{F} := \mathfrak{F}(\omega, s, p_{11}, p_{1-1})$ defined by

$$(5.53) \quad \begin{aligned} \mathfrak{F} &= -D_{\text{KL}}(s, u - s, u - s, 1 - 2u + s \parallel p_{-1-1}^k, p_{-1}^k - p_{-1-1}^k, p_{-1}^k - p_{-1-1}^k, \\ &1 - 2p_{-1}^k + p_{-1-1}^k) \\ &+ kD_{\text{KL}}(\omega, 1/2 - \omega, 1/2 - \omega, \omega \parallel p_{11}, p_{1-1}, p_{1-1}, p_{-1-1}). \end{aligned}$$

Moreover, with $\mathfrak{p} = \mathfrak{p}(\omega, s)$ the solution to (5.31)–(5.32) from Lemma 5.15, we let

$$(5.54) \quad F(\omega, s) = \mathfrak{F}(\omega, s, \mathfrak{p}_{11}, \mathfrak{p}_{1-1}).$$

COROLLARY 5.18. *A.a.s. we have $m^{-1} \log \mathbb{P}[\mathcal{S}^\otimes \mid \mathcal{R}^\otimes(\omega), \mathfrak{D}] \leq \max_{s \in [0, u]} F(\omega, s) + o(1)$.*

Thus, the next item on the agenda is to find the stationary points of $F(\omega, s)$ for ω close to $1/4$. We begin by exhibiting an explicit stationary point.

LEMMA 5.19. *We have $DF(1/4, u^2) = 0$ and*

$$\frac{d}{k} F(1/4, u^2) = -\frac{2(k-1)d}{k} \log 2 - d \log(p(1-p)) + \frac{2d}{k} \log p + \frac{2d}{k} \beta u.$$

PROOF. With p the solution to (2.7), we verify directly that at the point $\omega = 1/4, s = u^2$ the solution to (5.31) reads

$$(5.55) \quad p_{11} = p^2, \quad p_{1-1} = p_{-11} = p(1-p), \quad p_{-1-1} = (1-p)^2.$$

Hence, we obtain the formula for $F(1/4, u^2)$ by simply plugging (5.55) into (5.54). Moreover, using the formulas $\frac{\partial}{\partial y} y \log \frac{y}{z} = 1 + \log \frac{y}{z}, \frac{\partial}{\partial z} y \log \frac{y}{z} = -\frac{y}{z}$ we compute

$$(5.56) \quad \frac{\partial \mathfrak{F}}{\partial \omega} = k \left[2 \log 2 + \log \frac{\omega}{p_{11}} + \log \frac{\omega}{1 - p_{11} - 2p_{1-1}} - 2 \log \frac{1 - 2\omega}{p_{1-1}} \right].$$

Substituting (5.55) into (5.56), we find

$$(5.57) \quad \left. \frac{\partial \mathfrak{F}}{\partial \omega} \right|_{\omega=1/4, p_{11}=p^2, p_{1-1}=p(1-p)} = 0,$$

and similarly

$$(5.58) \quad \begin{aligned} \frac{\partial \mathfrak{F}}{\partial s} = & -\log \frac{s}{(1 - 2p_{1-1} - p_{11})^k} + 2 \log \frac{u - s}{(1 - p_{1-1} - p_{11})^k - (1 - 2p_{1-1} - p_{11})^k} \\ & - \log \frac{1 - 2u + s}{1 - 2(1 - p_{11} - p_{1-1})^k + (1 - 2p_{1-1} - p_{11})^k}. \end{aligned}$$

As substituting $s = u^2, p_{11} = p^2$ and $p_{1-1} = p(1-p)$ into the last expression yields zero, we conclude that

$$(5.59) \quad \left. \frac{\partial \mathfrak{F}}{\partial s} \right|_{s=u^2, p_{11}=p^2, p_{1-1}=p(1-p)} = 0.$$

Moreover,

$$(5.60) \quad \begin{aligned} \frac{\partial \mathfrak{F}}{\partial p_{11}} = & -\frac{ks}{1 - p_{11} - 2p_{1-1}} \\ & - \frac{2(u-s)k((1 - 2p_{1-1} - p_{11})^{k-1} - (1 - p_{11} - p_{1-1})^{k-1})}{(1 - 2p_{1-1} - p_{11})^k - (1 - p_{11} - p_{1-1})^k} \\ & - \frac{(1 - 2u + s)k((1 - 2p_{1-1} - p_{11})^{k-1} - 2(1 - p_{11} - p_{1-1})^{k-1})}{1 - 2(1 - p_{11} - p_{1-1})^k + (1 - 2p_{1-1} - p_{11})^k} \\ & - \frac{(1 - 2p_{11} - 2p_{1-1})k\omega}{p_{11}(1 - 2p_{1-1} - p_{11})}. \end{aligned}$$

Hence,

$$\begin{aligned} \frac{\partial \mathfrak{F}}{\partial p_{11}} \Big|_{\substack{\omega=1/4, s=u^2 \\ p_{11}=p^2 \\ p_{1-1}=p(1-p)}} &= -\frac{u^2 k}{(1-p)^2} - \frac{2u(1-u)k(1-(1-p)^{k-1})}{(1-p)(1-(1-p)^k)} \\ &\quad + \frac{k(1-u)^2(1-p)^{k-1}(2-(1-p)^{k-1})}{(1-(1-p)^k)^2} - \frac{k(1-2p)}{4p^2(1-p)^2}. \end{aligned}$$

Further, recalling that $u = (1 - 2p)/(2p(\exp(\beta) - 1))$ and that p is the solution to (2.7) and hence $e^\beta = \frac{(1-p)^k}{2p-1+(1-p)^k}$, we obtain

$$(5.61) \quad \frac{\partial \mathfrak{F}}{\partial p_{11}} \Big|_{\substack{\omega=1/4, s=u^2 \\ p_{11}=p^2 \\ p_{1-1}=p(1-p)}} = 0.$$

In addition,

$$\begin{aligned} (5.62) \quad \frac{\partial \mathfrak{F}}{\partial p_{1-1}} &= -\frac{2ks}{1-2p_{1-1}-p_{11}} \\ &\quad - \frac{2k(u-s)(2(1-2p_{1-1}-p_{11})^{k-1} - (1-p_{11}-p_{1-1})^{k-1})}{(1-2p_{1-1}-p_{11})^k - (1-p_{11}-p_{1-1})^k} \\ &\quad - \frac{2k(1-2u+s)((1-2p_{1-1}-p_{11})^{k-1} - (1-p_{11}-p_{1-1})^{k-1})}{1-2(1-p_{11}-p_{1-1})^k + (1-2p_{1-1}-p_{11})^k} \\ &\quad + k \left(\frac{2\omega}{1-2p_{1-1}-p_{11}} - \frac{1-2\omega}{p_{1-1}} \right). \end{aligned}$$

Hence, using $u = (1 - 2p)/(2p(\exp(\beta) - 1))$ and (2.7), we obtain

$$(5.63) \quad \frac{\partial \mathfrak{F}}{\partial p_{1-1}} \Big|_{\substack{\omega=1/4, s=u^2 \\ p_{11}=p^2 \\ p_{1-1}=p(1-p)}} = 0.$$

Combining (5.57), (5.59), (5.61) and (5.63) with the chain rule, we conclude that $DF(1/4, u^2) = 0$. \square

We are now going to compare $\max_{s \in [0, u]} F(\omega, s)$ for ω close to $1/4$ with $F(1/4, u^2)$. To this end we need to get a handle on the value of the maximiser s of $F(\omega, s)$ for a given ω . Hence, we investigate the second partial derivatives of the function \mathfrak{F} from (5.53). Let $g := g(\omega, p_{11}, p_{1-1})$ and $h := h(s, p_{11}, p_{1-1})$ be defined as

$$\begin{aligned} g &= D_{\text{KL}}(\omega, 1/2 - \omega, 1/2 - \omega, \omega \parallel p_{11}, p_{1-1}, p_{1-1}, p_{-1-1}), \\ h &= -D_{\text{KL}}(s, u - s, u - s, 1 - 2u + s \parallel p_{-1-1}^k, p_{-1}^k - p_{-1-1}^k, p_{-1}^k - p_{-1-1}^k, \\ &\quad 1 - 2p_{-1}^k + p_{-1-1}^k) \end{aligned}$$

which respectively denote the two constituent terms of (5.53).

LEMMA 5.20. For $\omega = \frac{1}{4} + \tilde{O}(2^{-k/2})$ and $0 \leq s \leq u$ we have

$$\begin{aligned} \frac{\partial^2 g}{\partial \omega^2} \Big|_{\substack{p_{11}=p_{11} \\ p_{1-1}=p_{1-1}}} &+ \frac{\partial^2 g}{\partial p_{11}^2} \Big|_{\substack{p_{11}=p_{11} \\ p_{1-1}=p_{1-1}}} \left(\frac{\partial p_{11}}{\partial \omega} \right)^2 + \frac{\partial^2 g}{\partial p_{1-1}^2} \Big|_{\substack{p_{11}=p_{11} \\ p_{1-1}=p_{1-1}}} \left(\frac{\partial p_{1-1}}{\partial \omega} \right)^2 \\ &+ 2 \frac{\partial^2 g}{\partial \omega \partial p_{11}} \Big|_{\substack{p_{11}=p_{11} \\ p_{1-1}=p_{1-1}}} \frac{\partial p_{11}}{\partial \omega} + 2 \frac{\partial^2 g}{\partial \omega \partial p_{1-1}} \Big|_{\substack{p_{11}=p_{11} \\ p_{1-1}=p_{1-1}}} \frac{\partial p_{1-1}}{\partial \omega} \end{aligned}$$

$$\begin{aligned}
 &+ 2 \frac{\partial^2 g}{\partial p_{11} \partial p_{1-1}} \Big|_{\substack{p_{11}=p_{11} \\ p_{1-1}=p_{1-1}}} \frac{\partial p_{11}}{\partial \omega} \frac{\partial p_{1-1}}{\partial \omega} \\
 &= \tilde{O}(4^{-k}).
 \end{aligned}$$

PROOF. Direct calculations reveal

$$(5.64) \quad \frac{\partial^2 g}{\partial \omega^2} = \frac{2}{\omega(1-2\omega)}, \quad \frac{\partial^2 g}{\partial \omega \partial p_{11}} = -\frac{1-2p_{1-1}-2p_{11}}{p_{11}(1-2p_{1-1}-p_{11})},$$

$$\frac{\partial^2 g}{\partial \omega \partial p_{1-1}} = \frac{2(1-p_{1-1}-p_{11})}{p_{1-1}(1-2p_{1-1}-p_{11})},$$

$$(5.65) \quad \frac{\partial^2 g}{\partial p_{11}^2} = \frac{\omega(1-2p_{11}-2p_{1-1}+2p_{11}^2+4p_{11}p_{1-1})}{p_{11}^2(1-2p_{1-1}-p_{11})^2},$$

$$\frac{\partial^2 g}{\partial p_{11} \partial p_{1-1}} = \frac{2\omega}{(1-2p_{1-1}-p_{11})^2},$$

$$(5.66) \quad \frac{\partial^2 g}{\partial p_{1-1}^2} = \frac{4(1-\omega)p_{1-1}^2 - 4(1-2\omega)(1-p_{11})p_{1-1} + (1-2\omega)(p_{11}-1)^2}{p_{1-1}^2(1-2p_{1-1}-p_{11})^2}.$$

Moreover, since Corollary 5.16 shows that $\frac{\partial p_{11}}{\partial \omega} = 1 + \tilde{O}(2^{-k})$, $\frac{\partial p_{1-1}}{\partial \omega} = -1 + \tilde{O}(2^{-k})$, we obtain

$$(5.67) \quad p_{11} = \omega + \tilde{O}(2^{-k}), \quad p_{1-1} = \frac{1}{2} - \omega + \tilde{O}(2^{-k}).$$

Substituting (5.67) into (5.64)–(5.66) yields the assertion. \square

LEMMA 5.21. For $0 \leq s \leq u$ and $\omega = \frac{1}{4} + \tilde{O}(2^{-k/2})$ we have $\frac{\partial^2}{\partial s^2} F(\omega, s) = -\Omega(1/s)$.

PROOF. Combining (5.65) and (5.66) with Corollary 5.16, we find

$$\begin{aligned}
 (5.68) \quad &\frac{\partial^2 g}{\partial p_{11}^2} \Big|_{\substack{p_{11}=p_{11} \\ p_{1-1}=p_{1-1}}} \left(\frac{\partial p_{11}}{\partial s} \right)^2 + \frac{\partial^2 g}{\partial p_{1-1}^2} \Big|_{\substack{p_{11}=p_{11} \\ p_{1-1}=p_{1-1}}} \left(\frac{\partial p_{1-1}}{\partial s} \right)^2 \\
 &+ 2 \frac{\partial^2 g}{\partial p_{11} \partial p_{1-1}} \Big|_{\substack{p_{11}=p_{11} \\ p_{1-1}=p_{1-1}}} \frac{\partial p_{11}}{\partial s} \frac{\partial p_{1-1}}{\partial s} = \tilde{O}(1).
 \end{aligned}$$

Moreover, we compute

$$\begin{aligned}
 (5.69) \quad &\frac{\partial^2 h}{\partial s^2} = \frac{1}{s} - \frac{2}{u-s} - \frac{1}{1-2u+s}, \\
 &\frac{\partial^2 h}{\partial p_{11} \partial s}, \frac{\partial^2 h}{\partial p_{1-1} \partial s}, \frac{\partial^2 h}{\partial p_{11}^2}, \frac{\partial^2 h}{\partial p_{1-1}^2}, \frac{\partial^2 h}{\partial p_{11} \partial p_{1-1}} = \tilde{O}(1).
 \end{aligned}$$

Since $s \leq u = \tilde{O}(2^{-k})$, the first term in (5.69) is of order $-\Omega(1/s) = -\tilde{\Omega}(2^k)$. Therefore, (5.69) and Corollary 5.16 shows

$$\begin{aligned}
 & \frac{\partial^2 h}{\partial s^2} + \frac{\partial^2 h}{\partial p_{11}^2} \Big|_{p_{1-1} = p_{1-1}}^{p_{11} = p_{11}} \left(\frac{\partial p_{11}}{\partial s} \right)^2 + \frac{\partial^2 h}{\partial p_{1-1}^2} \Big|_{p_{1-1} = p_{1-1}}^{p_{11} = p_{11}} \left(\frac{\partial p_{1-1}}{\partial s} \right)^2 \\
 (5.70) \quad & + 2 \frac{\partial^2 h}{\partial s \partial p_{11}} \Big|_{p_{1-1} = p_{1-1}}^{p_{11} = p_{11}} \frac{\partial p_{11}}{\partial s} + 2 \frac{\partial^2 h}{\partial s \partial p_{1-1}} \Big|_{p_{1-1} = p_{1-1}}^{p_{11} = p_{11}} \frac{\partial p_{1-1}}{\partial s} \\
 & + 2 \frac{\partial^2 h}{\partial p_{11} \partial p_{1-1}} \Big|_{p_{1-1} = p_{1-1}}^{p_{11} = p_{11}} \frac{\partial p_{11}}{\partial s} \frac{\partial p_{1-1}}{\partial s} = -\tilde{\Omega}(1/s).
 \end{aligned}$$

Combining (5.68) and (5.70), we see that

$$\begin{aligned}
 & \frac{\partial^2 \mathfrak{F}}{\partial s^2} + \frac{\partial^2 \mathfrak{F}}{\partial p_{11}^2} \Big|_{p_{1-1} = p_{1-1}}^{p_{11} = p_{11}} \left(\frac{\partial p_{11}}{\partial s} \right)^2 + \frac{\partial^2 \mathfrak{F}}{\partial p_{1-1}^2} \Big|_{p_{1-1} = p_{1-1}}^{p_{11} = p_{11}} \left(\frac{\partial p_{1-1}}{\partial s} \right)^2 \\
 (5.71) \quad & + 2 \frac{\partial^2 \mathfrak{F}}{\partial \omega \partial p_{11}} \Big|_{p_{1-1} = p_{1-1}}^{p_{11} = p_{11}} \frac{\partial p_{11}}{\partial s} + 2 \frac{\partial^2 \mathfrak{F}}{\partial s \partial p_{1-1}} \Big|_{p_{1-1} = p_{1-1}}^{p_{11} = p_{11}} \frac{\partial p_{1-1}}{\partial s} \\
 & + 2 \frac{\partial^2 \mathfrak{F}}{\partial p_{11} \partial p_{1-1}} \Big|_{p_{1-1} = p_{1-1}}^{p_{11} = p_{11}} \frac{\partial p_{11}}{\partial s} \frac{\partial p_{1-1}}{\partial s} = -\tilde{\Omega}(1/s).
 \end{aligned}$$

Furthermore, the expressions (5.60) and (5.62) for the partial derivatives of \mathfrak{F} and Corollary 5.16 yield

$$(5.72) \quad \frac{\partial \mathfrak{F}}{\partial p_{11}} \Big|_{p_{1-1} = p_{1-1}}^{p_{11} = p_{11}} \frac{\partial^2 p_{11}}{\partial s^2}, \frac{\partial \mathfrak{F}}{\partial p_{1-1}} \Big|_{p_{1-1} = p_{1-1}}^{p_{11} = p_{11}} \frac{\partial^2 p_{1-1}}{\partial s^2} = \tilde{O}(2^{-k}).$$

Hence, combining (5.71) and (5.72) with Faà di Bruno’s rule, we conclude that $\frac{\partial^2}{\partial s^2} F(\omega, s) = -\Omega(1/s)$. \square

COROLLARY 5.22. *For any $\omega = \frac{1}{4} + \tilde{O}(2^{-k/2})$ the function $s \in [0, u] \mapsto F(\omega, s)$ attains its unique maximum at $s = \tilde{O}(4^{-k})$.*

PROOF. We consider the first derivative $\frac{\partial}{\partial s} F(\omega, s)$ for $\omega = \frac{1}{4} + \tilde{O}(2^{-k/2})$. The partial derivatives of \mathfrak{F} , which we computed in (5.56), (5.58), (5.60) and (5.62), satisfy

$$\begin{aligned}
 & \frac{\partial \mathfrak{F}}{\partial s} \Big|_{p_{1-1} = p_{1-1}}^{p_{11} = p_{11}} = -\log \frac{s}{4^{-k} + \tilde{O}(8^{-k})} + 2 \log \frac{u - s}{2^{-k} + \tilde{O}(4^{-k})} - \log \frac{1 - 2u + s}{1 + \tilde{O}(2^{-k})}, \\
 (5.73) \quad & \frac{\partial \mathfrak{F}}{\partial p_{11}} \Big|_{p_{1-1} = p_{1-1}}^{p_{11} = p_{11}}, \frac{\partial \mathfrak{F}}{\partial p_{1-1}} \Big|_{p_{1-1} = p_{1-1}}^{p_{11} = p_{11}} = \tilde{O}(2^{-k}).
 \end{aligned}$$

Hence, Corollary 5.16 and the chain rule yield

$$(5.74) \quad \frac{\partial F(\omega, s)}{\partial s} = \log \frac{(u - s)^2}{s(1 - 2u + s)} + \tilde{O}(2^{-k}).$$

Since $u = \tilde{O}(2^{-k})$, (5.74) shows that the equation $\frac{\partial F(\omega, s)}{\partial s} = 0$ is satisfied only for $s = \tilde{O}(4^{-k})$. Thus, the assertion follows from Lemma 5.21. \square

Having estimated the maximiser s of $F(\omega, s)$, we now bound the Hessian $D^2 F(\omega, s)$ of $F(\omega, s)$.

LEMMA 5.23. We have $D^2F(\omega, s) \leq \tilde{O}(4^{-k})\text{id}$ for all $\omega = \frac{1}{4} + \tilde{O}(2^{-k/2})$ and $s = \tilde{O}(4^{-k})$.

The proof requires two intermediate steps.

CLAIM 5.24. For $\omega = \frac{1}{4} + \tilde{O}(2^{-k/2})$ and $s \leq u$ we have

$$\frac{\partial \mathfrak{F}}{\partial p_{11}} \Big|_{p_{1-1} = p_{1-1}}^{p_{11} = p_{11}} \frac{\partial^2 p_{11}}{\partial \omega^2}, \frac{\partial \mathfrak{F}}{\partial p_{11}} \Big|_{p_{1-1} = p_{1-1}}^{p_{11} = p_{11}} \frac{\partial^2 p_{11}}{\partial \omega \partial s} = \tilde{O}(4^{-k})$$

and

$$\frac{\partial \mathfrak{F}}{\partial p_{1-1}} \Big|_{p_{1-1} = p_{1-1}}^{p_{11} = p_{11}} \frac{\partial^2 p_{1-1}}{\partial \omega^2}, \frac{\partial \mathfrak{F}}{\partial p_{1-1}} \Big|_{p_{1-1} = p_{1-1}}^{p_{11} = p_{11}} \frac{\partial^2 p_{1-1}}{\partial \omega \partial s} = \tilde{O}(4^{-k}).$$

PROOF. The claim follows from (5.73) and (5.45). \square

CLAIM 5.25. For $s = \tilde{O}(4^{-k})$ and $\omega = \frac{1}{4} + \tilde{O}(2^{-k/2})$ we have $\frac{\partial^2}{\partial \omega \partial s} F(\omega, s) = \tilde{O}(1)$.

PROOF. The second derivative of g with respect to s and ω is bounded because of Corollary 5.16 and as

$$(5.75) \quad \begin{aligned} \frac{\partial^2 g}{\partial s \partial \omega} &= -\frac{1}{p_{11}} \frac{\partial p_{11}}{\partial s} + \frac{2}{p_{1-1}} \frac{\partial p_{1-1}}{\partial s} \\ &+ \frac{1}{1 - 2p_{-1-1} - p_{11}} \frac{\partial p_{11}}{\partial s} + \frac{2}{1 - 2p_{-1-1} - p_{11}} \frac{\partial p_{1-1}}{\partial s}. \end{aligned}$$

So is the contribution of the first derivative. Thus, we just need to investigate the second derivative of h . The contribution of $\frac{\partial^2 h}{\partial p_{11}^2}, \frac{\partial^2 h}{\partial p_{1-1}^2}, \frac{\partial^2 h}{\partial p_{11} \partial p_{1-1}}$ is bounded by (5.69) and $\frac{\partial^2 h}{\partial \omega \partial s} = 0$. Therefore, the assertion follows from Faà di Bruno’s rule. \square

PROOF OF LEMMA 5.23. Because the Kullback–Leibler divergence is convex and because $\frac{\partial^2 F}{\partial s^2} < 0$ by Lemma 5.21, Lemma 5.20 and Claims 5.24–5.25 imply that

$$D^2F(\omega, s) \leq \mathfrak{H} = \begin{pmatrix} \mathfrak{x} & \eta \\ \eta & \mathfrak{z} \end{pmatrix} \quad \text{where } \mathfrak{x} = \tilde{O}(4^{-k}), \eta = \tilde{O}(1), \mathfrak{z} = \tilde{O}(4^k).$$

The eigenvalues of \mathfrak{H} work out to be $\frac{1}{2}(\mathfrak{x} + \mathfrak{z} \pm \sqrt{(\mathfrak{x} - \mathfrak{z})^2 + 4\eta^2})$. Therefore, the smaller eigenvalue of $D^2F(\omega, s)$ has size $-\tilde{\Omega}(4^k)$, while the large one is upper bounded by $\tilde{O}(4^{-k})$. Consequently, $D^2F \leq \tilde{O}(4^{-k})\text{id}$. \square

PROOF OF PROPOSITION 5.8. Corollary 5.18 shows that

$$m^{-1} \log \mathbb{P}[\mathbf{E}(\omega) \mid \mathfrak{D}] \leq \max_{s \in [0, u]} F(\omega, s) + o(1) \quad \text{a.a.s.}$$

Hence, with $s^*(\omega) = \tilde{O}(4^{-k})$ the unique maximiser of $s \mapsto F(\omega, s)$, we obtain a.a.s.

$$(5.76) \quad \begin{aligned} m^{-1} \log \mathbb{P}[\mathcal{S}^{\otimes} \mid \mathcal{R}^{\otimes}(\omega), \mathfrak{D}] &\leq F(\omega, s^*(\omega)) + o(1) \\ &\leq F(1/4, u^2) + (F(\omega, s^*(\omega)) - F(1/4, u^2)). \end{aligned}$$

Furthermore, Lemma 5.23 and Taylor’s formula imply that

$$(5.77) \quad F(\omega, s^*(\omega)) - F(1/4, u^2) \leq \tilde{O}(4^{-k}) \cdot (\omega - 1/4)^2.$$

Therefore, combining (5.21) and (5.76)–(5.77) with Lemma 5.11, we conclude that a.a.s.

$$(5.78) \quad \frac{\mathbb{E}[\mathbf{E}(\omega) \mid \mathfrak{D}]}{\mathbb{E}[\mathbf{E}(1/4) \mid \mathfrak{D}]} \leq \exp \left[-(\omega - 1/4)^2 \left(\frac{m^2}{d^2 n} + m \tilde{O}(4^{-k}) \right) + o(n) \right] \\ = \exp \left[-(\omega - 1/4)^2 (\Omega(k^{-2}n) + \tilde{O}(2^{-k})) + o(n) \right] \leq \exp(o(n)).$$

Finally, the assertion follows from Lemma 5.4, Lemma 5.19 and (5.78). \square

5.3.2. *Proof of Proposition 5.9.* We need to connect the weighted with the unweighted overlap. To this end we approximate for a given weighted overlap the dominant contributing unweighted overlap. Let $P(\ell) = \mathbb{P}[\text{Po}(d/2) = \ell]$.

LEMMA 5.26. *A.a.s. we have*

$$(5.79) \quad \sum_{\ell, \ell' \geq 0} (\ell + \ell' + 1) \left| P(\ell)P(\ell') - \frac{1}{n} \sum_{x \in V_n} \mathbb{1}\{\mathbf{d}_x^+ = \ell, \mathbf{d}_x^- = \ell'\} \right| \leq \sqrt{n} \log^4 n,$$

and $\max\{\mathbf{d}_x^+, \mathbf{d}_x^- : x \in V_n\} \leq \log n$.

PROOF. This follows from routine concentration bounds. Indeed, for $\ell, \ell' \leq \log n$ a straightforward application of Chebyshev’s inequality shows that

$$\left| \sum_{x \in V_n} \mathbb{1}\{\mathbf{d}_x^+ = \ell, \mathbf{d}_x^- = \ell'\} - n P(\ell)P(\ell') \right| \leq \sqrt{n} \log n$$

a.a.s. Moreover, Bennett’s inequality shows that $\mathbf{d}_x^+, \mathbf{d}_x^- \leq \log n$ for all $x \in V_n$ a.a.s. \square

If the condition (5.79) is satisfied, then we can express the most likely overlap α that gives rise to a given weighted overlap ω in terms of a neat optimisation problem:

$$\mathfrak{M}(\omega) = \max \left\{ -2 \sum_{d^+, d^- \geq 0} P(d^+)P(d^-) (\alpha_{11}(d^+, d^-) \log \alpha_{11}(d^+, d^-) \right. \\ \left. + (1/2 - \alpha_{11}(d^+, d^-)) \log(1/2 - \alpha_{11}(d^+, d^-)) \right\} \\ \text{s.t. } \sum_{d^+, d^- \geq 0} P(d^+)P(d^-) (d^+ + d^-) \alpha_{11}(d^+, d^-) = d\omega, \\ \forall d^+, d^- \geq 0 : 0 \leq \alpha_{11}(d^+, d^-) \leq 1/2.$$

Here the variable $\alpha_{11}(d^+, d^-)$ represents the fraction of variables with degrees d^+, d^- that get both set to “true”. Let $N(d^+, d^-)$ the number of variables x with $\mathbf{d}_x^+ = d^+, \mathbf{d}_x^- = d^-$. Because we only count assignments that satisfy the strongly balanced condition (5.3), there remain $N(d^+, d^-)(1/2 - \alpha_{11}(d^+, d^-))$ variables of degree (d^+, d^-) set to $(1, -1)$, another $N(d^+, d^-)(1/2 - \alpha_{11}(d^+, d^-))$ set to $(-1, 1)$, while the remaining $N(d^+, d^-)(\alpha_{11}(d^+, d^-))$ ones are set to $(-1, -1)$. Hence, assuming (5.79) the total number of such assignments comes to

$$(5.80) \quad \left(\alpha_{11} N(d^+, d^-), (1/2 - \alpha_{11}) N(d^+, d^-), (1/2 - \alpha_{-11})(d^+, d^-) N(d^+, d^-), \alpha_{11}(d^+, d^-) N(d^+, d^-) \right) \\ = \exp(-2n P(d^+)P(d^-) (\alpha_{11} \log \alpha_{11} + (1/2 - \alpha_{11}) \log(1/2 - \alpha_{11})) + o(n)).$$

In other words, $\mathfrak{M}(\omega)$ asks to choose $\alpha_{11}(d^+, d^-)$ so as to maximise the total number of possible assignments with weighted overlap ω . The following lemma shows that for ω far from $1/4$, the optimal solution to $\mathfrak{M}(\omega)$ renders an unweighted overlap $2 \sum_{d^+, d^-} \alpha_{11}(d^+, d^-)$ far from $1/2$.

LEMMA 5.27. For $|\omega - 1/4| > k^{100}2^{-k/2}$ the optimal solution to $\mathfrak{M}(\omega)$ satisfies

$$\left| \frac{1}{4} - \sum_{d^+, d^-} P(d^+)P(d^-)\alpha_{11}(d^+, d^-) \right| \geq k^{95}2^{-k/2}.$$

PROOF. We set up the Lagrangian

$$\begin{aligned} \mathcal{L} = & -2 \sum_{d^+, d^- \geq 0} P(d^+)P(d^-)[\alpha_{11}(d^+, d^-) \log \alpha_{11}(d^+, d^-) \\ & + (1/2 - \alpha_{11}(d^+, d^-)) \log(1/2 - \alpha_{11}(d^+, d^-))] \\ & - \lambda \left[\sum_{d^+, d^- \geq 0} P(d^+)P(d^-)(d^+ + d^-)\alpha_{11}(d^+, d^-) - d\omega \right]. \end{aligned}$$

The derivatives

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \alpha_{11}(d^+, d^-)} &= -2P(d^+)P(d^-) \left[\log \frac{2\alpha_{11}(d^+, d^-)}{1 - 2\alpha_{11}(d^+, d^-)} + \lambda(d^+ + d^-) \right], \\ \frac{\partial \mathcal{L}}{\partial \lambda} &= d\omega - \sum_{d^+, d^- \geq 0} P(d^+)P(d^-)(d^+ + d^-)\alpha_{11}(d^+, d^-) \end{aligned}$$

vanish iff

$$(5.81) \quad \sum_{d^+, d^- \geq 0} P(d^+)P(d^-)(d^+ + d^-)\alpha_{11}(d^+, d^-) = d\omega,$$

$$\alpha_{11}(d^+, d^-) = \frac{1 - \tanh(\lambda(d^+ + d^-)/2)}{4}.$$

Substituting the expression for $\alpha_{11}(d^+, d^-)$ into the left equation, we obtain

$$(5.82) \quad \omega = \frac{1}{4} - \frac{1}{4de^d} \sum_{d^+, d^- \geq 0} \frac{d^+ + d^-}{d^+!d^-!} \left(\frac{d}{2}\right)^{d^++d^-} \tanh \frac{\lambda(d^+ + d^-)}{2}.$$

Since the right-hand side is strictly increasing in λ , this equation has a unique solution for any $\alpha \in [0, 1/2]$. In fact, for the derivative of the right-hand side, setting

$$\mathfrak{J} := \sum_{d^+, d^- \geq 0} \binom{d}{2}^{d^++d^-} \frac{d^+ + d^-}{d^+!d^-!} \tanh \frac{\lambda(d^+ + d^-)}{2},$$

we obtain

$$\frac{\partial}{\partial \lambda} \mathfrak{J} = \sum_{d^+, d^- \geq 0} \binom{d}{2}^{d^++d^-} \frac{(d^+ + d^-)^2}{2(d^+!d^-!)} \left(1 - \tanh^2 \frac{\lambda(d^+ + d^-)}{2} \right).$$

In particular, for $\omega = 0$ the choice $\lambda = 0$ solves (5.82). Moreover, the derivatives of ω obtained through (5.82) can be estimated as

$$\begin{aligned} \frac{\partial \omega}{\partial \lambda} &= -\frac{1}{8d} \sum_{d^+, d^- \geq 0} P(d^+)P(d^-)(d^+ + d^-)^2 \left(1 - \tanh^2 \frac{\lambda(d^+ + d^-)}{2} \right) \\ &= -\frac{d+1}{8} + \frac{1}{32d} \lambda^2 \mathbb{E}[(d^+ + d^-)^4] + O(d^6 \lambda^3) \end{aligned}$$

$$= -\frac{d+1}{8} + \frac{\lambda^2}{32}(d^3 + 6d^2 + 7d + 1) + O(d^5\lambda^3),$$

$$\frac{\partial^2\omega}{\partial\lambda^2} = \frac{\lambda}{16}(d^3 + 6d^2 + 7d + 1) + O(d^5\lambda^2) = \Theta(\lambda d^3).$$

Because $\omega = \tilde{O}(2^{-k/2})$ the inverse function theorem yields $\frac{\partial\lambda}{\partial w} = \frac{8}{d+1} = \Theta(\frac{1}{d})$. Thus, $\frac{\partial^2\omega}{\partial\lambda^2} = \Theta(d2^{-k})$. Substituting the approximation for λ that we get from $\frac{\partial\lambda}{\partial w} = \frac{8}{d+1}$ back into (5.81) and summing on d^+, d^- completes the proof. \square

PROOF OF PROPOSITION 5.9. This is now an immediate consequence of Lemma 5.26, Lemma 5.27 and the elementary bound (2.5). \square

5.3.3. *Proof of Proposition 5.10.* Toward the proof of the proposition we first derive an explicit approximation of the term on the right-hand side. We begin by estimating p from (2.7).

LEMMA 5.28. We have $p = \frac{1}{2} - (1 - e^{-\beta})2^{-k-1} + (1 - e^{-\beta})^2k2^{-2k-2} + O(k^22^{-3k})$.

PROOF. The choice (2.7) ensures that

$$e^\beta = \frac{(1-p)^k}{2p-1+(1-p)^k} = 1 - \frac{2p-1}{2p-1+(1-p)^k} = 1 - \frac{2(p-\frac{1}{2})}{2(p-\frac{1}{2})+(\frac{1}{2}-(p-\frac{1}{2}))^k}.$$

Hence, $p \in (\frac{1}{2} - 2^{-k}(1 - e^{-\beta}), \frac{1}{2})$ and thus $q = p - 1/2 \in (-2^{-k}(1 - e^{-\beta}), 0)$ is the solution to $(2q + (\frac{1}{2} - q)^k)(e^\beta - 1) + 2q = 0$. Using the binomial expansion $(\frac{1}{2} - q)^k = 2^{-k} - k2^{-k-1}q + O(k^22^{-k}q^2)$ we obtain

$$e^\beta - 1 = -\frac{2q}{2q + 2^{-k} - k2^{-k-1}q + O(k^22^{-k}q^2)}.$$

Hence,

$$q = \frac{(e^\beta - 1)2^{-k}}{k2^{-k-1}(e^\beta - 1) - 2e^\beta + O(k^22^{-2k})}$$

$$= -\frac{(e^\beta - 1)2^{-k}}{2e^\beta} + k2^{-2k-2}\left(\frac{e^\beta - 1}{e^\beta}\right)^2 + O(k^22^{-3k}),$$

which implies the assertion. \square

COROLLARY 5.29. We have

$$\left(1 - \frac{(k-1)d}{k}\right) \log 2 - \frac{d}{2} \log(p(1-p)) + \frac{d}{k} \log p$$

$$= \log 2 - \frac{d}{k}(1 - e^{-\beta})2^{-k} + \frac{d(2k-1)}{2k}2^{-2k}(1 - e^{-\beta})^2 + O(dk^22^{-3k}).$$

PROOF. Using the approximation for p from Lemma 5.28, we obtain

$$-\frac{d}{2} \log(p(1-p)) = d \log 2 - \frac{d}{2} \log(1 - (1 - e^{-\beta})2^{-k} + (1 - e^{-\beta})^2k2^{-2k-1} + O(k^22^{-3k}))$$

$$- \frac{d}{2} \log(1 + (1 - e^{-\beta})2^{-k} - (1 - e^{-\beta})^2k2^{-2k-1} + O(k^22^{-3k})).$$

Moreover,

$$\begin{aligned} & \log(1 \mp (1 - e^{-\beta})2^{-k} \pm (1 - e^{-\beta})^2 k 2^{-2k-1} + O(k^2 2^{-3k})) \\ &= \mp (1 - e^{-\beta})2^{-k} \pm (1 - e^{-\beta})^2 k 2^{-2k-1} - 2^{-2k-1} (1 - e^{-\beta})^2 + O(k^2 2^{-3k}). \end{aligned}$$

Hence,

$$(5.83) \quad -\frac{d}{2} \log(p(1-p)) = d \log 2 + d 2^{-2k-1} (1 - e^{-\beta})^2 + O(dk^2 2^{-3k}).$$

Further, using (2.7), we obtain $2p = 1 - (1 - e^{-\beta})(1-p)^k$ and thus

$$(5.84) \quad \log(2p) = -(1 - e^{-\beta})2^{-k} + (k-1)2^{-2k-1}(1 - e^{-\beta})^2 + O(k^2 2^{-3k}).$$

Combining (5.83) and (5.84) completes the proof. \square

Having estimated the expression on the right-hand side of Proposition 5.10, we proceed to investigate the function $f(\alpha)$. Its derivatives read

$$(5.85) \quad f'(\alpha) = \log \frac{1-\alpha}{\alpha} + \frac{d\alpha^{k-1}(1 - e^{-\beta})^2}{2^k(1 - 2^{1-k}(1 - e^{-\beta}) + \alpha^k 2^{-k}(1 - e^{-\beta})^2)},$$

$$(5.86) \quad \begin{aligned} f''(\alpha) = & -\frac{1}{\alpha} - \frac{1}{1-\alpha} + \frac{(k-1)d\alpha^{k-2}(1 - e^{-\beta})^2}{2^k(1 - 2^{1-k}(1 - e^{-\beta}) + 2^{-k}\alpha^k(1 - e^{-\beta})^2)} \\ & - \frac{k d \alpha^{2(k-1)}(1 - e^{-\beta})^4}{2^{2k}(1 - 2^{1-k}(1 - e^{-\beta}) + 2^{-k}\alpha^k(1 - e^{-\beta})^2)^2}. \end{aligned}$$

CLAIM 5.30. *We have $f(\alpha) \leq f(1-\alpha)$ for all $\alpha < 1/2$. Moreover, f is concave on the interval $[1/2, 1/2 + o(k^{-3})]$, where it attains a local maximum at $\alpha^* = \frac{1}{2} + O(d4^{-k})$ with*

$$(5.87) \quad f(\alpha^*) = 2 \log 2 - \frac{2d(1 - e^{-\beta})}{2^k} (1 + 2^{-k-1}(1 - e^{-\beta})) + O(k^8 4^{-k}).$$

PROOF. The first assertion follows immediately from the symmetry of the entropy function and the fact that $\alpha \mapsto 1 - 2^{1-k}(1 - \exp(-\beta)) + 2^{-k}\alpha^k(1 - \exp(-\beta))^2$ is increasing. We also read off of (5.85) that $f'(1/2) = O(d4^{-k})$, while (5.86) shows that $f''(\alpha) = -4 + o(1)$ if $\alpha = 1/2 + o(k^{-3})$. Hence, f attains a local maximum at $\alpha^* = 1/2 + O(d4^{-k})$. Finally, Taylor’s formula shows that $f(\alpha^*) = f(1/2) + O(k^4 d^2 16^{-k})$, whence we obtain (5.87). \square

CLAIM 5.31. *The function $f(\alpha)$ is monotonically decreasing on $(1/2 + k^{-4}, 1 - 2 \log(k)/k)$.*

PROOF. For $\alpha \in [1/2 + k^{-4}, 0.99]$ we see that $\log((1-\alpha)/\alpha) = -\Omega(k^8)$ while

$$(5.88) \quad \frac{d\alpha^{k-1}(1 - e^{-\beta})^2}{2^k(1 - 2^{1-k}(1 - e^{-\beta}) + \alpha^k 2^{-k}(1 - e^{-\beta})^2)} = \exp(-\Omega(k)).$$

Hence, (5.85) shows that f is decreasing on this interval. Similarly, for $\alpha \in (0.99, 1 - 2 \log(k)/k)$ we obtain $\log((1-\alpha)/\alpha) = -\Omega(1)$, while the left-hand side of (5.88) is of order $o(1)$. Hence, $f'(\alpha) < 0$. \square

We are going to prove that for $d \leq d^*$ the maximum of $f(\alpha)$ is approximately equal to $f(1/2)$ by comparing the function value $f(1/2)$ with the function values $f(\alpha)$ for $\alpha > 1/2$. Since the function $\alpha \mapsto 1 - 2^{1-k}(1 - \exp(-\beta)) + 2^{-k}\alpha^k(1 - \exp(-\beta))^2$ is monotonically increasing, we may assume that $d = d^*$. Actually, in order to facilitate the proof of Theorem 1.2, in some of the estimates below we will allow for d to take values up to $d_{\text{SAT}}(k)$.

CLAIM 5.32. Assume that $d^* \leq d \leq d_{\text{SAT}}(k)$. Then the function $f(\alpha)$ has only one stationary point in the Interval $[1 - 2 \log(k)/k, 1 - k^{-3/2}]$, which is a local minimum.

PROOF. Substituting $\alpha = 1 - \varepsilon$ into (5.85) we can write

$$(5.89) \quad f'(\alpha) = \log \varepsilon - \log(1 - \varepsilon) + \frac{d(1 - e^{-\beta})^2 \exp(\varepsilon - k\varepsilon + O(k\varepsilon^2))}{2^k(1 + O(2^{-k}))}.$$

Therefore, for $d^* \leq d \leq d_{\text{SAT}}(k)$ the only solution to $f'(\alpha) = 0$ is such that $\varepsilon = (\log(k) + O(\log \log k))/k$. Furthermore, for the root of $f'(\alpha)$ we read off (5.86) that $f''(\alpha) = \Omega(k \log k) > 0$. \square

CLAIM 5.33. Assume that $d^* \leq d \leq d_{\text{SAT}}(k)$. Then the function $f(\alpha)$ has only one stationary point in the interval $[1 - k^{-3/2}, 1]$, namely a local maximum at $\alpha_* = 1 - 2^{-k(1-e^{-\beta})^2+o(1)}$.

PROOF. Contemplating (5.89), we see that for $\alpha = 1 - \varepsilon$ with $\varepsilon < k^{-3/2}$ the last summand has the form $(1 + o(1))(1 - e^{-\beta})^2 k \log 2$. Hence, the only root of $f'(\alpha)$ in this interval occurs at $\varepsilon = \varepsilon_* = 2^{-k(1-e^{-\beta})^2+o(1)}$. A glimpse at (5.86) reveals that $f''(1 - \varepsilon_*) < 0$. Hence, $\alpha_* = 1 - \varepsilon_*$ is a local maximum. \square

CLAIM 5.34. Assume that $d \leq d^*$. Then the global maximum of $f(\alpha)$ is attained at α^* from Claim 5.30.

PROOF. In light of Claims 5.31–5.33 we just need to compare $f(\alpha_*)$ and $f(\alpha^*)$ for $d = d^*$. Let us estimate them:

$$(5.90) \quad \begin{aligned} f(\alpha_*) &\leq \log 2 + \varepsilon_*(1 - \log \varepsilon_*) \\ &\quad + \frac{d^*}{k} \log(1 - 2^{1-k}(1 - e^{-\beta}) + 2^{-k}(1 - \varepsilon_*)^k(1 - e^{-\beta})^2) \\ &\leq e^{-2\beta} \log 2 + 10k2^{-k}(1 - e^{-2\beta}) + o(k2^{-k}), \end{aligned}$$

$$(5.91) \quad \begin{aligned} f(\alpha^*) &= 2 \log 2 + \frac{2d}{k} \log(1 - (1 - e^{-\beta})2^{-k}) \\ &= 2e^{-\beta} \log 2 + 20k2^{-k}(1 - e^{-\beta}) + o(k2^{-k}). \end{aligned}$$

Combining (5.90) and (5.91), we see that $f(\alpha^*) > f(\alpha_*)$ for $\beta \geq 1$. \square

PROOF OF PROPOSITION 5.10. Combining Claims 5.31–5.34, we are left to merely compare $f(\alpha^*)$ and the right-hand side expression from Proposition 5.10. Comparing the approximation of the latter supplied by Corollary 5.29 with (5.87) completes the proof. \square

5.4. Proof of Corollary 5.3. As an immediate consequence of Proposition 2.1, (5.15) and Markov’s inequality we obtain that a.a.s. the random formula Φ satisfies

$$\sum_{\sigma, \tau \in \{\pm 1\}^n} \mathbb{1}\{|\sigma \cdot \tau| > k^{100}2^{-k/2}n\} \prod_{i=1}^m \exp(-\beta(\mathbb{1}\{\sigma \not\equiv a_i\} + \mathbb{1}\{\tau \not\equiv a_i\})) = o(Z(\Phi, \beta)^2).$$

Dividing by $Z(\Phi, \beta)^2$, we obtain

$$(5.92) \quad \mu_{\Phi, \beta}(\{|\sigma \cdot \sigma'| > nk^{100}2^{-k/2}\}) = o(1),$$

whence Corollary 5.3 is immediate.

To prove Corollary 2.2 we combine (5.92) with the following general lemma.

CLAIM 5.35. *Suppose that $\nu \in \mathcal{P}(\{\pm 1\}^n)$ satisfies $\nu(\{|\boldsymbol{\sigma} \cdot \boldsymbol{\sigma}'| \geq \varepsilon n\}) < 1/2$ and*

$$(5.93) \quad \sum_{i,j=1}^n |\nu(\{\boldsymbol{\sigma}_i = \boldsymbol{\sigma}_j = 1\}) - \nu(\{\boldsymbol{\sigma}_i = 1\})\nu(\{\boldsymbol{\sigma}_j = 1\})| = o(n^2).$$

Then $\sum_{i=1}^n (\nu(\{\boldsymbol{\sigma}_i = 1\}) - 1/2)^2 < \varepsilon n$.

PROOF. The assumption (5.93) implies together with [21], Lemma 2.11, that the product measure $\nu \otimes \nu$ has the same property, that is, that for two independent samples $\boldsymbol{\sigma}, \boldsymbol{\sigma}'$ from ν and for any $s, s', t, t' \in \{\pm 1\}$ we have

$$(5.94) \quad \sum_{i,j=1}^n |\nu \otimes \nu(\{\boldsymbol{\sigma}_i = s, \boldsymbol{\sigma}_j = s', \boldsymbol{\sigma}'_i = t, \boldsymbol{\sigma}'_j = t'\}) - \nu \otimes \nu(\{\boldsymbol{\sigma}_i = s, \boldsymbol{\sigma}'_i = t\}) \cdot \nu \otimes \nu(\{\boldsymbol{\sigma}_j = s', \boldsymbol{\sigma}'_j = t'\})| = o(n^2).$$

Now, for $i \in [n]$ let $p_i = \nu(\{\boldsymbol{\sigma}_i = 1\})$. Then

$$\langle \boldsymbol{\sigma}_i \cdot \boldsymbol{\sigma}'_i, \nu \rangle = p_i^2 + (1 - p_i)^2 - 2p_i(1 - p_i) = (1 - 2p_i)^2 = 4\left(p_i - \frac{1}{2}\right)^2.$$

Hence, $\langle \sum_{i=1}^n \boldsymbol{\sigma}_i \cdot \boldsymbol{\sigma}'_i, \nu \rangle = 4 \sum_{i=1}^n (p_i - \frac{1}{2})^2$. Therefore, (5.94) and Chebychev's inequality show

$$\nu\left(\left\{\sum_{i=1}^n \boldsymbol{\sigma}_i \cdot \boldsymbol{\sigma}'_i \geq \sum_{i=1}^n \left(p_i - \frac{1}{2}\right)^2\right\}\right) \geq 1/2.$$

Consequently, the assumption $\nu(\{|\boldsymbol{\sigma} \cdot \boldsymbol{\sigma}'| \geq \varepsilon n\}) < 1/2$ implies that $\sum_{i=1}^n (p_i - \frac{1}{2})^2 \leq \varepsilon n$, as desired. \square

PROOF OF COROLLARY 2.2. The corollary is an immediate consequence of (5.92) and Claim 5.35. \square

6. Proof of Proposition 2.3. Throughout this section we assume that $d \leq d_{\text{SAT}}(k)$.

6.1. *Outline.* The proof of Proposition 2.3 relies on a contraction argument. Specifically, we will be able to describe the distribution of the BP marginal $\mu_{T, \beta, \pi, x_0, t}(1)$ at the root of the random tree T , which we aim to compute, in terms of an operator \mathcal{R} on the space $\mathcal{P}([0, 1])$ of probability measures on the unit interval. To define this operator let $\boldsymbol{\gamma}^+, \boldsymbol{\gamma}^-$ be $\text{Po}(d/2)$ variables and given $\mu \in \mathcal{P}([0, 1])$ let $\boldsymbol{\eta} = (\boldsymbol{\eta}_{ij}^+, \boldsymbol{\eta}_{ij}^-)_{i,j \geq 1}$ be random variables with distribution μ . All these random variables are mutually independent. Then $\mathcal{R}(\mu) \in \mathcal{P}([0, 1])$ is the law of the random variable

$$(6.1) \quad R(\boldsymbol{\gamma}^+, \boldsymbol{\gamma}^-, \boldsymbol{\eta}) = \frac{\prod_{i=1}^{\boldsymbol{\gamma}^+} (1 - (1 - e^{-\beta}) \prod_{j=1}^{k-1} \boldsymbol{\eta}_{ij}^+)}{\prod_{i=1}^{\boldsymbol{\gamma}^+} (1 - (1 - e^{-\beta}) \prod_{j=1}^{k-1} \boldsymbol{\eta}_{ij}^+) + \prod_{i=1}^{\boldsymbol{\gamma}^-} (1 - (1 - e^{-\beta}) \prod_{j=1}^{k-1} \boldsymbol{\eta}_{ij}^-)} \in (0, 1).$$

We write $\mathcal{R}^t(\cdot)$ for the t -fold iteration of \mathcal{R} .

We are going to investigate the operator \mathcal{R} on a subspace of $\mathcal{P}([0, 1])$. Specifically, let \mathcal{P}^* be the space of all probability measures $\mu \in \mathcal{P}([0, 1])$ such that $\mu([0, x]) = \mu([1 - x, 1])$ for all $x \in [0, 1]$. Because $\boldsymbol{\gamma}^+$ and $\boldsymbol{\gamma}^-$ are identically distributed, (6.1) ensures that \mathcal{R} maps the subspace \mathcal{P}^* into itself. Further, for any probability measure $\pi \in \mathcal{P}([0, 1])$ we obtain a probability measure $\pi^* \in \mathcal{P}^*$ as follows. Draw \mathbf{X} from π and independently draw a Rademacher variable \mathbf{J} with $\mathbb{E}[\mathbf{J}] = 0$. Then π^* is the distribution of $(1 + \mathbf{J}(2\mathbf{X} - 1))/2$. The following observation links \mathcal{R} to the belief propagation message passing scheme on T .

LEMMA 6.1. For any $\pi \in \mathcal{P}([0, 1])$ and any $t \geq 1$ the random variable $\mu_{T, \beta, \pi, x_0, t}(1)$ has distribution $\mathcal{R}^t(\pi^*)$.

PROOF. We first consider the case $t = 1$. The construction of T ensures that the root x_0 has $\text{Po}(d/2)$ children a with $J_{ax_0} = 1$ and an independent number of $\text{Po}(d/2)$ children a with $J_{ax_0} = -1$. Hence, these numbers can be coupled with γ^+, γ^- . Now consider a child a of x_0 with $J_{ax_0} = 1$. Then a has children y_1, \dots, y_{k-1} and the signs J_{ay_i} are independent Rademacher variables. Hence, if X_1, \dots, X_{k-1} are independent samples from π , then the message that a passes to x_0 reads

$$\begin{aligned} &\mu_{T, \beta, \pi, a \rightarrow x_0, 1}(s) \\ &\propto \mathbb{1}\{J_{ax_0} = s\} \\ &\quad + \mathbb{1}\{J_{ax_0} = -s\} \left(1 - (1 - e^{-\beta}) \prod_{i=1}^{k-1} (\mathbb{1}\{J_{ay_i} = 1\}(1 - X_i) + \mathbb{1}\{J_{ay_i} = -1\}X_i) \right), \end{aligned}$$

where $s = \pm 1$. Because the J_{ay_i} and the X_i are independent, the distribution of $\mu_{T, \beta, \pi, a \rightarrow x_0, 1}(s)$ can alternatively be described as follows: let X_1^*, \dots, X_{k-1}^* be independent samples from π^* ; then $\mu_{T, \beta, \pi, a \rightarrow x_0, 1}(s)$ is distributed as

$$\frac{\mathbb{1}\{J_{ax_0} = s\} + \mathbb{1}\{J_{ax_0} = -s\}(1 - (1 - e^{-\beta}) \prod_{i=1}^{k-1} X_i^*)}{2 - (1 - e^{-\beta}) \prod_{i=1}^{k-1} X_i^*}.$$

Consequently, $\mu_{T, \beta, \pi, x_0, t}(1)$ has distribution $\mathcal{R}(\pi^*) = \mathcal{R}^t(\pi^*)$ for $t = 1$. Finally, a straightforward induction extends this statement to all $t \geq 1$. \square

To prove Proposition 2.3 we first study the operator \mathcal{R} on a subspace $\mathcal{P}^\dagger \subset \mathcal{P}^*$ of probability measures that satisfy a certain tail bound. Specifically, we define \mathcal{P}^\dagger as the space of all measures $\mu \in \mathcal{P}^*$ such that

$$(6.2) \quad \mu\left(\left[\frac{e^s}{1 + e^s}, 1\right]\right) \leq \exp(-s2^{k/4}) \quad \text{for all } s \geq 2^{-k/4}.$$

For future reference we observe that the function

$$(6.3) \quad \begin{aligned} \varphi : \mathbb{R} &\rightarrow (0, 1), & z &\mapsto \frac{e^z}{1 + e^z} \quad \text{is a bijection with inverse} \\ \varphi^{-1} : (0, 1) &\rightarrow \mathbb{R}, & y &\mapsto \log \frac{y}{1 - y}. \end{aligned}$$

The operator \mathcal{R} maps the subspace \mathcal{P}^\dagger into itself.

PROPOSITION 6.2. For every $\mu \in \mathcal{P}^\dagger$ we have $\mathcal{R}(\mu) \in \mathcal{P}^\dagger$.

The proof of Proposition 6.2 can be found in Section 6.2. Next we show that \mathcal{R} is a contraction on \mathcal{P}^\dagger . Specifically, in Section 6.3 we will prove the following.

PROPOSITION 6.3. For $d \leq d_{\text{SAT}}(k)$ the operator \mathcal{R} is a contraction on \mathcal{P}^\dagger with respect to the W_r -metric, where r is the smallest even integer greater than $2^{k/10}$.

Finally, in Section 6.4 we will prove that a sufficient number of iterations of \mathcal{R} will map any distribution π with slim tails to a distribution that “almost” belongs to the subspace \mathcal{P}^\dagger .

To formalise this, for a random variable $\eta \in [0, 1]$ and a number $\varepsilon > 0$ let $\tilde{\eta}_\varepsilon$ be a random variable whose distribution is characterised by

$$(6.4) \quad \begin{aligned} \mathbb{P}[\tilde{\eta}_\varepsilon \in A] &= \mathbb{P}[\eta \in A \cap [\varepsilon, 1 - \varepsilon]] \\ &+ \mathbb{1}\left\{\frac{1}{2} \in A\right\} \mathbb{P}[\eta \notin [\varepsilon, 1 - \varepsilon]] \quad \text{for any measurable } A \subset [0, 1]. \end{aligned}$$

In words, $\tilde{\eta}_\varepsilon$ is obtained by truncating η at ε and $1 - \varepsilon$ and shifting the lost probability mass to $\frac{1}{2}$.

PROPOSITION 6.4. *For any $\varepsilon > 0$ there exists $\ell_0 = \ell_0(\varepsilon) > 0$ such that for all $\ell > \ell_0$ and all probability measures π with slim tails the following two statements hold.*

- (i) *Let ξ be a random variable with distribution $\mathcal{R}^\ell(\pi)$. Then $\mathbb{P}[\xi \notin [\varepsilon, 1 - \varepsilon]] \leq \varepsilon$.*
- (ii) *The distribution of $\tilde{\xi}_\varepsilon$ belongs to \mathcal{P}^\dagger .*

We prove Proposition 6.4 in Section 6.4. These statements now easily imply Proposition 2.3.

PROOF OF PROPOSITION 2.3. Let $\varepsilon > 0$ and let π, π' be two distributions with slim tails. Further, choose a large enough integer $\ell = \ell(\varepsilon) > 0$, a small enough $\delta = \delta(\varepsilon, \ell) > 0$ and finally a large enough integer $L = L(\varepsilon, \ell, \delta) > 0$. Suppose we initialise the variable-to-parent messages for the variables at distance $2(\ell + L)$ from the root of T with independent messages drawn from π . Consider the variables at intermediate distance precisely 2ℓ from the root of T and denote the set of these variables by U . For each $y \in U$, let b_y be its parent clause. Let $\xi_y^{(L)}$ be the message that y sends to its parent b_y with the π -initialisation at level $2(\ell + L)$. Moreover, let $\tilde{\xi}_{y,\delta}^{(L)}$ be the corresponding δ -truncated message as per (6.4). Define $\xi'_y{}^{(L)}, \tilde{\xi}'_{y,\delta}{}^{(L)}$ analogously with the initial messages at distance $2(\ell + L)$ from the root drawn independently from π' . By the construction of the Galton–Watson tree T the size of U has mean $(kd)^\ell$. Thus, providing $\ell = \ell(\varepsilon)$ is chosen large enough, Markov’s inequality implies that

$$(6.5) \quad \mathbb{P}[|U| \leq 2\varepsilon^{-1}(dk)^\ell] > 1 - \varepsilon/2.$$

Introducing the event

$$\mathcal{U} = \bigcap_{y \in U} (\{\xi_y^{(L)} = \tilde{\xi}_{y,\delta}^{(L)}\} \cap \{\xi'_y{}^{(L)} = \tilde{\xi}'_{y,\delta}{}^{(L)}\})$$

we thus obtain

$$(6.6) \quad \mathbb{P}[\mathcal{U} \mid U] \geq 1 - \sum_{y \in U} (\mathbb{P}[\xi_y^{(L)} \neq \tilde{\xi}_{y,\delta}^{(L)} \mid U] + \mathbb{P}[\xi'_y{}^{(L)} \neq \tilde{\xi}'_{y,\delta}{}^{(L)} \mid U]).$$

Now, because both π and π' have slim tails, Proposition 6.4 implies that for large enough $L = L(\varepsilon, \ell, \delta)$ and any $y \in U$ we have

$$(6.7) \quad \mathbb{P}[\xi_y^{(L)} \notin [\delta, 1 - \delta] \mid U] \leq \delta \quad \text{and} \quad \mathbb{P}[\xi'_y{}^{(L)} \notin [\delta, 1 - \delta] \mid U] \leq \delta.$$

Indeed, Proposition 6.4 is applicable here because the tree pending on any variable node $y \in U$ at distance precisely 2ℓ from the root is an independent copy of the Galton–Watson tree T . Combining (6.5)–(6.7) and choosing $\delta = \delta(\varepsilon, \ell) \ll \varepsilon(dk)^{-\ell}$ and $L = L(\delta)$ appropriately, we obtain

$$(6.8) \quad \mathbb{P}[\mathcal{U}] \geq 1 - \frac{\varepsilon}{2} - 4\delta\varepsilon^{-1}(dk)^\ell > 1 - \varepsilon.$$

(A careful analysis of the proof of Proposition 6.4 actually shows that

$$L \geq C + (\log 2)^{-1} \log \log \frac{1 - \delta}{\delta}$$

is sufficient for given δ and a constant C that is independent of δ .) Finally, given \mathcal{U} all messages at level 2ℓ both with initial distribution π as well as with initial distribution π' belong to \mathcal{P}^\dagger . Proposition 6.3 therefore shows that given \mathcal{U} , the conditional L^r -distance at the root of T is bounded by $(1 - \alpha)^\ell$ for some fixed $\alpha = \alpha(k, d, \beta) > 0$. Therefore, the L^r -distance between the two messages at the root is bounded by $\varepsilon + (1 - \alpha)^\ell$. Thus, to complete the proof we just need to choose $\ell = \ell(\varepsilon) > 0$ large enough. \square

The proofs of Propositions 6.2 and 6.3 are adaptations of the proofs of [24], Lemmas 4.2 and 4.4, where a related but slightly more intricate distributional recursion is analysed. Moreover, the proof of Proposition 6.4 combines ideas and observations from the proof of Proposition 6.2 such as the random walk analysis on typical events and the stochastic domination argument with a novel bootstrapping idea to track the tightening and proliferation of certain tail bounds as \mathcal{R} is iteratively applied. This allows to extend the conclusion of Proposition 6.3 to the relevant initial distributions. Let us delve into the details.

6.2. *Proof of Proposition 6.2.* Throughout this section we assume that $\mu \in \mathcal{P}^\dagger$. The proof of Proposition 6.2 follows along the steps of [24], Lemma 4.2, where Ding, Sly and Sun analyse a more complex distributional recursion related to the survey propagation algorithm. More precisely, they show that the application of their operator preserves a tail bound similar to (6.2). Here we follow the steps of their proof closely to derive the corresponding statement for the conceptually simpler belief propagation operator. That said, here and there some extra care is required because we work with an arbitrary temperature parameter $0 < \beta < \infty$ while in [24] it suffices to study zero temperature.

6.2.1. *Overview.* We aim to show that $\hat{\mu} = \mathcal{R}(\mu)$ belongs to \mathcal{P}^\dagger as well. Letting

$$\Pi^+ = \prod_{i=1}^{\gamma^+} \left(1 - (1 - e^{-\beta}) \prod_{j=1}^{k-1} \eta_{ij}^+ \right), \quad \Pi^- = \prod_{i=1}^{\gamma^-} \left(1 - (1 - e^{-\beta}) \prod_{j=1}^{k-1} \eta_{ij}^- \right),$$

we see from the definition of \mathcal{R} that

$$(6.9) \quad \hat{\mu} \left(\left[\frac{e^s}{1 + e^s}, 1 \right] \right) = \mathbb{P}(\log \Pi^+ - \log \Pi^- \geq s).$$

Both $-\log \Pi^+$ and $-\log \Pi^-$ are sums of a random number of nonnegative i.i.d. random variables, and $-\log \Pi^+$ and $-\log \Pi^-$ are identically distributed. Hence, estimating the probability (6.9) is a bit like estimating the probability that a weighted symmetric random walk strays far from the origin. To bound this probability we first bound the large deviations of the individual summands. More specifically, for $i \geq 1$, set

$$X_i^\pm = -\log \left(1 - (1 - e^{-\beta}) \prod_{j=1}^{k-1} \eta_{ij}^\pm \right) \geq 0 \quad \text{so that}$$

$$-\log \Pi^+ = \sum_{i=1}^{\gamma^+} X_i^+, \quad -\log \Pi^- = \sum_{i=1}^{\gamma^-} X_i^-.$$

Let $X \stackrel{d}{=} X_1^+$ denote a generic summand. In Section 6.2.2 we are going to prove the following.

LEMMA 6.5.

- (i) For all $t \geq 1$ we have $\mathbb{P}(X \geq t) \leq \exp(-\frac{t}{2}(k-1)2^{k/4})$.
- (ii) For all $\varepsilon \in [2^{-k/9}, 1]$ we have $\mathbb{P}(X \geq (1 - e^{-\beta})2^{-(1-\varepsilon)(k-1)}) \leq 2^{k-1} \exp(-(k-1)(\frac{\varepsilon}{3})^2 2^{k/4})$.
- (iii) For all $\varepsilon \in [2^{-k/9}, 1]$ we have $\mathbb{P}[X \leq (1 - e^{-\beta})2^{-(1+\varepsilon)(k-1)}] \leq k \exp(-2^{k/4}\varepsilon)$.

Lemma 6.5 implies the following estimate of $\mathbb{E}[X]$.

COROLLARY 6.6. We have $\mathbb{E}[X] = (1 - e^{-\beta})2^{-(k-1)} + O(k2^{-10k/9})$.

PROOF. Applying Lemma 6.5 with $\varepsilon = 2^{-k/9}$, we obtain

$$\begin{aligned} \mathbb{E}[X] &= \mathbb{E}[X\mathbb{1}\{X \in [0, (1 - e^{-\beta})2^{-(1+\varepsilon)(k-1)}]\}] \\ &\quad + \mathbb{E}[X\mathbb{1}\{X \in ((1 - e^{-\beta})2^{-(1+\varepsilon)(k-1)}, (1 - e^{-\beta})2^{-(1-\varepsilon)(k-1)})\}] \\ &\quad + \mathbb{E}[X\mathbb{1}\{X \in [(1 - e^{-\beta})2^{-(1-\varepsilon)(k-1)}, 1)\}] + \mathbb{E}[X\mathbb{1}\{X \geq 1\}] \\ &= (1 - e^{-\beta})2^{-(k-1)} + O(k2^{-(k-1)-k/9}), \end{aligned}$$

as claimed. \square

In order to estimate the difference between $\log \Pi^+$ and $\log \Pi^-$ we are going to replace X by the truncated random variable

$$(6.10) \quad \bar{X} = X\mathbb{1}\{X \leq (1 - e^{-\beta})2^{-(9/10)(k-1)}\}.$$

Lemma 6.5 implies the following bound on the difference between X and \bar{X} .

COROLLARY 6.7. We have $\mathbb{E}[X - \bar{X}] \leq \exp(-\Omega(k2^{k/4}))$.

PROOF. Applying Lemma 6.5 (i) and (ii) with $\varepsilon = 1/10$, we obtain

$$\begin{aligned} \mathbb{E}[X - \bar{X}] &\leq \mathbb{P}(X \geq (1 - e^{-\beta})2^{-(9/10)(k-1)}) + \mathbb{E}[X\mathbb{1}\{X \geq 1\}] \\ &\leq 2^k \exp\left(-\frac{k-1}{900}2^{k/4}\right) + 2 \exp(-(k-1)2^{k/4-1}), \end{aligned}$$

as claimed. \square

With these preparations in place we prove the desired tail bound in two steps. First we consider the case $s \in [2^{-k/4}, 1]$. To be precise, in Section 6.2.3 we are going to prove the following.

LEMMA 6.8. For $s \in [2^{-k/4}, 1]$ we have $\hat{\mu}([\frac{e^s}{1+e^s}, 1]) \leq \exp(-s2^{k/4})$.

Finally, in Section 6.2.4 we are going to deal with $s > 1$.

LEMMA 6.9. For $s > 1$ we have $\hat{\mu}([\frac{e^s}{1+e^s}, 1]) \leq \exp(-s2^{k/4})$.

PROOF OF PROPOSITION 6.2. The assertion follows immediately from Lemmas 6.8 and 6.9. \square

6.2.2. *Proof of Lemma 6.5.* We prove the three statements separately. With respect to (i) we notice that $X \leq \beta$ deterministically. Hence, the assertion trivially holds for $t \geq \beta$. Now consider $t \in [1, \beta)$. Because the random variables η_{1j}^+ are bounded by one and independent, we obtain

$$\begin{aligned} \mathbb{P}[X \geq t] &= \mathbb{P}\left[\prod_{j=1}^{k-1} \eta_{1j}^+ \geq \frac{1 - e^{-t}}{1 - e^{-\beta}}\right] \leq \mathbb{P}\left(\forall j \in [k - 1] : \eta_{1j}^+ \geq \frac{1 - e^{-t}}{1 - e^{-\beta}}\right) \\ (6.11) \quad &= \mu\left(\left[\frac{1 - e^{-t}}{1 - e^{-\beta}}, 1\right]\right)^{k-1}. \end{aligned}$$

Further, since $\frac{1 - e^{-t}}{1 - e^{-\beta}} < 1$ for $t < \beta$, $\varphi^{-1}\left(\frac{1 - e^{-t}}{1 - e^{-\beta}}\right) \geq t/2$ and because $\mu \in \mathcal{P}^\dagger$, (6.11) yields $\mathbb{P}[X \geq t] \leq \exp(-\frac{t}{2}(k - 1)2^{k/4})$, as claimed.

We proceed to (ii). Since $1 - \exp(-yz) \geq y(1 - \exp(-z))$ for any $y \in (0, 1]$, $z \geq 0$, we obtain

$$\begin{aligned} \mathbb{P}(X \geq (1 - e^{-\beta})2^{(\varepsilon-1)(k-1)}) &= \mathbb{P}\left(\prod_{j=1}^{k-1} \eta_{1j}^+ \geq \frac{1 - \exp(-(1 - e^{-\beta})2^{(\varepsilon-1)(k-1)})}{1 - e^{-\beta}}\right) \\ (6.12) \quad &\leq \mathbb{P}\left(\prod_{j=1}^{k-1} \eta_{1j}^+ \geq 1 - \exp(-2^{(\varepsilon-1)(k-1)})\right). \end{aligned}$$

Further, since $\eta_{1j}^+ \leq 1$ for all j , for any $a \in [0, 1]$ and $b \in (0, k - 2]$ we have

$$(6.13) \quad \mathbb{P}\left(\prod_{j=1}^{k-1} \eta_{1j}^+ \geq a\right) \leq \mathbb{P}(|\{j \in [k - 1] : \eta_{1j}^+ \geq a^{\frac{1}{k-1-b}}\}| \geq b).$$

Moreover, for large k we have

$$(6.14) \quad 1 - \exp(-2^{(\varepsilon-1)(k-1)}) \geq 2^{-(k-1)(1-\varepsilon/3)^2}.$$

Combining (6.12), (6.13) and (6.14), we obtain

$$\begin{aligned} \mathbb{P}(X \geq (1 - e^{-\beta})2^{(\varepsilon-1)(k-1)}) &\leq \mathbb{P}\left(\prod_{j=1}^{k-1} \eta_{1j}^+ \geq 2^{-(k-1)(1-\varepsilon/3)^2}\right) \\ (6.15) \quad &\leq \mathbb{P}(|\{j \in [k - 1] : \eta_{1j}^+ \geq 2^{\varepsilon/3-1}\}| \geq (k - 1)\varepsilon/3) \\ &\leq \sum_{(k-1)\varepsilon/3 \leq j \leq k-1} \binom{k-1}{j} \mu([2^{\varepsilon/3-1}, 1])^j \\ &\leq 2^{k-1} \mu([2^{\varepsilon/3-1}, 1])^{(k-1)\varepsilon/3}. \end{aligned}$$

Since $\varphi^{-1}(2^{\varepsilon/3-1}) \geq \frac{2\varepsilon}{3} \log 2 \geq 2^{-k/4}$ for all $\varepsilon \in [2^{-k/9}, 1]$, the assumption $\mu \in \mathcal{P}^\dagger$ and (6.15) imply

$$(6.16) \quad \mathbb{P}(X \geq (1 - e^{-\beta})2^{(\varepsilon-1)(k-1)}) \leq 2^{k-1} \exp\left(- (k - 1) \frac{\varepsilon}{3} 2^{k/4} \varphi^{-1}(2^{\varepsilon/3-1})\right).$$

Finally, since $\varphi^{-1}(2^{\varepsilon/3-1}) \geq \frac{2\varepsilon}{3} \log 2$ for $\varepsilon \in [0, 1]$, the second assertion follows from (6.16).

Coming to the statement (iii), we use the inequality $1 - e^{-x} \leq x$ to obtain

$$\begin{aligned} \mathbb{P}(X \leq (1 - e^{-\beta})2^{-(\varepsilon+1)(k-1)}) &\leq \mathbb{P}(1 - e^{-X} \leq (1 - e^{-\beta})2^{-(\varepsilon+1)(k-1)}) \\ (6.17) \quad &\leq \mathbb{P}(\exists j \in [k - 1] : \eta_{1j}^+ \leq 2^{-(\varepsilon+1)}) \\ &\leq (k - 1) \mu([0, 2^{-(\varepsilon+1)}]). \end{aligned}$$

Since $\varphi^{-1}(1 - 2^{-(\varepsilon+1)}) = \log(2^{\varepsilon+1} - 1) \geq \varepsilon \geq 2^{-k/4}$, the fact that $\mu \in \mathcal{P}^\dagger$ implies that

$$(6.18) \quad (k - 1)\mu([0, 2^{-(\varepsilon+1)}]) \leq k \exp(-2^{k/4}\varepsilon).$$

The assertion follows from (6.17) and (6.18).

6.2.3. *Proof of Lemma 6.8.* We need to bound the tails of $-\log \Pi^+$, which is the sum of a Poisson number of i.i.d. copies of \mathbf{X} . Having derived tail bounds for the individual summands \mathbf{X} and an approximation of $\mathbb{E}[\mathbf{X}]$ already, we are going to deal with large deviations of the number of summands next. Bennett’s inequality Lemma 4.5 directly implies that if $d = O(k2^k)$, then

$$(6.19) \quad \mathbb{P}(|\boldsymbol{\gamma}^+ - d/2| \geq k^2 2^{5k/8}) \leq \exp(-\Omega(k^3 2^{k/4})).$$

In analogy to (6.10), for $i \geq 1$, set

$$(6.20) \quad \bar{\mathbf{X}}_i^+ = \mathbf{X}_i^+ \mathbb{1}\{\mathbf{X}_i^+ \leq (1 - e^{-\beta})2^{-(9/10)(k-1)}\}.$$

On the event

$$\mathcal{E} = \{|\boldsymbol{\gamma}^+ - d/2| \leq k^2 2^{5k/8}\} \cap \{\max\{\mathbf{X}_1^+, \dots, \mathbf{X}_{\boldsymbol{\gamma}^+}^+\} \leq (1 - e^{-\beta})2^{-(9/10)(k-1)}\}$$

Corollaries 6.6 and 6.7 yield

$$(6.21) \quad \begin{aligned} & \left| \sum_{i=1}^{\boldsymbol{\gamma}^+} (\mathbf{X}_i^+ - \mathbb{E}[\mathbf{X}_i^+]) - \sum_{i=1}^{d/2} (\bar{\mathbf{X}}_i^+ - \mathbb{E}[\bar{\mathbf{X}}_i^+]) \right| \\ & \leq \frac{d}{2} \mathbb{E}[\mathbf{X} - \bar{\mathbf{X}}] + \left| \boldsymbol{\gamma}^+ - \frac{d}{2} \right| \cdot |\mathbb{E}[\mathbf{X}] + \max\{\mathbf{X}_1^+, \dots, \mathbf{X}_{\boldsymbol{\gamma}^+}^+\}| \\ & \leq \frac{d}{2} \exp(-\Omega(k2^{k/4})) + k^2 2^{5k/8} (2^{-(k-1)} + O(k2^{-10k/9})) + 2^{-(9/10)(k-1)} \\ & \leq k^{-5} 2^{-1-k/4}. \end{aligned}$$

Therefore, using Lemma 6.5, Corollary 6.6, (6.19) and (6.21), we obtain for any $t > 0$,

$$(6.22) \quad \begin{aligned} & \mathbb{P}\left(\left| -\log \Pi^+ - \frac{d}{2} \mathbb{E}[\mathbf{X}] \right| \geq t\right) \\ & = \mathbb{P}\left(\left| \sum_{i=1}^{\boldsymbol{\gamma}^+} (\mathbf{X}_i^+ - \mathbb{E}[\mathbf{X}_i^+]) + \mathbb{E}[\mathbf{X}] \left(\boldsymbol{\gamma}^+ - \frac{d}{2}\right) \right| \geq t\right) \\ & \leq 1 - \mathbb{P}(\mathcal{E}) + \mathbb{P}\left(\mathcal{E} \cap \left\{ \left| \sum_{i=1}^{\boldsymbol{\gamma}^+} (\mathbf{X}_i^+ - \mathbb{E}[\mathbf{X}_i^+]) - \sum_{i=1}^{d/2} (\bar{\mathbf{X}}_i^+ - \mathbb{E}[\bar{\mathbf{X}}_i^+]) \right. \right. \right. \\ & \quad \left. \left. \left. + \sum_{i=1}^{d/2} (\bar{\mathbf{X}}_i^+ - \mathbb{E}[\bar{\mathbf{X}}_i^+]) + \mathbb{E}[\mathbf{X}] \left(\boldsymbol{\gamma}^+ - \frac{d}{2}\right) \right| \geq t \right\}\right) \\ & \leq \exp(-\Omega(k2^{k/4})) + \mathbb{P}\left(\left| \sum_{i=1}^{d/2} (\bar{\mathbf{X}}_i - \mathbb{E}[\bar{\mathbf{X}}_i]) \right| \geq t - k^{-5} 2^{-k/4}\right). \end{aligned}$$

The last probability involving a sum with a deterministic number of bounded summands can be bounded by the Azuma–Hoeffding inequality, which yields for $t = \Omega(2^{-k/4})$

$$\begin{aligned} \mathbb{P}\left(\left| \sum_{i=1}^{d/2} (\bar{\mathbf{X}}_i - \mathbb{E}[\bar{\mathbf{X}}_i]) \right| \geq t - k^{-5} 2^{-k/4}\right) & \leq 2 \exp\left(-\frac{(t - k^{-5} 2^{-k/4})^2}{d(1 - e^{-\beta})2^{2-(18/10)(k-1)}}\right) \\ & = \exp(-\Omega(t^2 k^{-1} 2^{4k/5})). \end{aligned}$$

Combining this estimate with (6.22), for $t = \Omega(2^{-k/4})$ we obtain

$$(6.23) \quad \mathbb{P}\left(\left|-\log \Pi^+ - \frac{d}{2}\mathbb{E}[\mathbf{X}]\right| \geq t\right) \leq \exp(-\Omega(k2^{k/4})) + \exp(-\Omega(t^2k^{-1}2^{4k/5})) \\ \leq \exp(-\Omega(k2^{k/4})) + \exp(-\Omega(tk2^{k/4})).$$

Because Π^-, Π^+ are identically distributed, (6.23) implies that

$$(6.24) \quad \mathbb{P}\left(\left|\log \Pi^- + \frac{d}{2}\mathbb{E}[\mathbf{X}]\right| \geq t\right) \leq \exp(-\Omega(k2^{k/4})) + \exp(-\Omega(tk2^{k/4}))$$

as well. Finally, for $2^{-k/4} \leq t \leq 1$ the second terms in (6.23)–(6.24) dominate. Therefore, for $s \in [2^{-k/4}, 1]$, we obtain

$$\hat{\mu}\left(\left[\frac{e^s}{1+e^s}, 1\right]\right) = \mathbb{P}\left(\left(-\log \Pi^- - \frac{d}{2}\mathbb{E}[\mathbf{X}]\right) - \left(-\log \Pi^+ - \frac{d}{2}\mathbb{E}[\mathbf{X}]\right) \geq s\right) \\ \leq \exp(-\Omega(sk2^{k/4})) \leq \exp(-s2^{k/4}),$$

as claimed.

6.2.4. *Proof of Lemma 6.9.* For $s \geq 1$ let $\mathcal{E}_s = \{|\boldsymbol{\gamma}^+ - d/2| \leq s2^{9k/10}k^{-5/4}\}$. Using Bennett’s inequality Lemma 4.5, we obtain for $s \geq 1$

$$(6.25) \quad 1 - \mathbb{P}(\mathcal{E}_s) \leq \exp(-\Omega(s2^{4k/5}k^{-7/2})).$$

Combining (6.25) and (6.20) with the Azuma–Hoeffding inequality, we obtain for $s \geq 1$

$$(6.26) \quad \mathbb{P}\left(\left|\sum_{i=1}^{\boldsymbol{\gamma}^+} \bar{X}_i^+ - \frac{d}{2}\mathbb{E}[\bar{X}]\right| \geq \frac{s}{k}\right) \leq \mathbb{P}\left(\left|\sum_{i=1}^{d/2} (\bar{X}_i^+ - \mathbb{E}[\bar{X}])\right| \geq \frac{s}{2k}\right) \\ + \mathbb{P}\left(|\boldsymbol{\gamma}^+ - d/2|(1 - e^{-\beta})2^{-(9/10)(k-1)} \geq \frac{s}{2k}\right) \\ \leq 1 - \mathbb{P}(\mathcal{E}_s) + 2 \exp\left(-\frac{s^2}{4(1 - e^{-\beta})^2k^2d2^{-(9/5)(k-1)}}\right) \\ \leq \exp(-s2^{2k/3}).$$

We are left to bound the difference between $-\log \Pi^+$ and the sum of truncated random variables. We write

$$(6.27) \quad \sum_{i=1}^{\boldsymbol{\gamma}^+} \mathbf{X}_i^+ - \sum_{i=1}^{\boldsymbol{\gamma}^+} \bar{X}_i^+ \leq \sum_{i=1}^{\boldsymbol{\gamma}^+} \mathbb{1}\{\mathbf{X}_i^+ \geq (1 - e^{-\beta})2^{-(9/10)(k-1)}\} + \sum_{i=1}^{\boldsymbol{\gamma}^+} \mathbb{1}\{\mathbf{X}_i^+ > 1\}\mathbf{X}_i^+.$$

We now compare the right-hand side with a more accessible distribution. Lemma 6.5 shows that for $c := 1/1000$, for all $a \geq 1$,

$$(6.28) \quad \mathbb{P}(\mathbf{X}_i^+ \geq (1 - e^{-\beta})2^{-(9/10)(k-1)}) \leq \exp(-ck2^{k/4}), \\ \mathbb{P}(\mathbf{X}_i^+ \geq a) \leq \exp(-2ack2^{k/4}).$$

Set $\vartheta = ck2^{k/4}$. Then (6.28) shows that for $z \geq 0$,

$$\mathbb{P}(\mathbb{1}\{\mathbf{X}_i^+ > (1 - e^{-\beta})2^{-(9/10)(k-1)}\}(1 + \mathbb{1}\{\mathbf{X}_i^+ > 1\}\mathbf{X}_i^+) > z) \\ \leq \begin{cases} \mathbb{P}(\mathbf{X}_i^+ > (1 - e^{-\beta})2^{-(9/10)(k-1)}) \leq e^{-\vartheta} & \text{if } z < 1, \\ \mathbb{P}(\mathbf{X}_i^+ \geq 1) \leq e^{-2\vartheta} & \text{if } 1 \leq z < 2, \\ \mathbb{P}(\mathbf{X}_i^+ \geq z - 1) \leq e^{-2(z-1)\vartheta} \leq e^{-z\vartheta} & \text{if } z \geq 2. \end{cases}$$

Hence, we can estimate X_i^+ as follows. Let $(I_i^+)_{i \geq 1}$ be a sequence of $\text{Be}(e^{-\vartheta})$ random variables, let $(Z_i^+)_{i \geq 1}$ be a sequence of exponential random variables with mean $1/\vartheta$ and let \bar{y}^+ be a $\text{Po}(d/2)$ random variable, all mutually independent. Then for all $i \geq 1, z \geq 0$ we have

$$\begin{aligned} \mathbb{P}(\mathbb{1}\{X_i^+ > (1 - e^{-\beta})2^{-(9/10)(k-1)}\}(1 + \mathbb{1}\{X_i^+ > 1\}X_i^+) > z) &\leq \mathbb{P}(I_i^+(1 + Z_i^+) > z) \\ &= \begin{cases} e^{-\vartheta}, & z < 1, \\ e^{-z\vartheta}, & z \geq 1. \end{cases} \end{aligned}$$

Thus, $\mathbb{1}\{X_i^+ > (1 - e^{-\beta})2^{-(9/10)(k-1)}\}(1 + \mathbb{1}\{X_i^+ > 1\}X_i^+)$ is stochastically dominated by $I_i^+(1 + Z_i^+)$. Therefore, we also obtain stochastic dominance for the sums of these random variables, that is,

$$(6.29) \quad \sum_{i=1}^{\gamma^+} \mathbb{1}\{X_i^+ > (1 - e^{-\beta})2^{-(9/10)(k-1)}\}(1 + \mathbb{1}\{X_i^+ > 1\}X_i^+) \leq \Sigma_1^+ + \Sigma_2^+$$

where

$$\Sigma_1^+ = \sum_{i=1}^{\bar{y}^+} I_i^+, \quad \Sigma_2^+ = \sum_{i=1}^{\bar{y}^+} I_i^+ Z_i^+.$$

Clearly, Σ_1^+ has distribution $\text{Po}(de^{-\vartheta}/2)$. Bennett’s inequality (4.1) therefore yields for $s \geq 1$,

$$(6.30) \quad \begin{aligned} \mathbb{P}\left(\Sigma_1^+ - \frac{d}{2}e^{-\vartheta} \geq s/\sqrt{k}\right) &\leq \exp\left(-\frac{s}{\sqrt{k}}\left(\log\left(\frac{e^{\vartheta}}{d\sqrt{k}}\right) - 1\right)\right) \\ &\leq \exp\left(-\frac{s}{2\sqrt{k}}(\vartheta - \log(d\sqrt{k}))\right) \leq \exp(-sk^{1/3}2^{k/4}). \end{aligned}$$

Let us now turn our attention to Σ_2^+ .

CLAIM 6.10. *For all $s > d/(2\vartheta e^{\vartheta})$ we have $\mathbb{P}(\Sigma_2^+ \geq s) \leq \exp(-\vartheta s(1 - \sqrt{\frac{d}{2s\vartheta e^{\vartheta}}})^2)$.*

PROOF. Given Σ_1^+, Σ_2^+ is a sum of Σ_1^+ independent exponential random variables with parameter ϑ and therefore $\Gamma(\Sigma_1^+, \vartheta)$ -distributed (where Σ_1^+ and ϑ denote the form and scale parameters, respectively). Therefore, for $0 \leq t < \vartheta$,

$$\mathbb{E}[\exp(t\Sigma_2^+)] = \mathbb{E}[\mathbb{E}[\exp(t\Gamma(\Sigma_1^+, \vartheta)) \mid \Sigma_1^+]] = \mathbb{E}[(1 - t/\vartheta)^{-\Sigma_1^+}] = \exp\left(\frac{dt}{2e^{\vartheta}(\vartheta - t)}\right).$$

Consequently, for $s > 0$ and $0 < t < \vartheta$,

$$\mathbb{P}(\Sigma_2^+ \geq s) \leq \frac{\mathbb{E}[\exp(t\Sigma_2^+)]}{\exp(st)} = \exp\left(\frac{dt}{2e^{\vartheta}(\vartheta - t)} - st\right).$$

With the choice $t^* = \vartheta - \sqrt{(d\vartheta e^{-\vartheta})/(2s)}$, which lies between 0 and ϑ for $s > \frac{d}{2}\vartheta^{-1}e^{-\vartheta}$, we find

$$\mathbb{P}(\Sigma_2^+ \geq s) \leq \mathbb{E}[\exp(t^*\Sigma_2^+ - t^*s)] = \exp\left(-\vartheta s\left(1 - \sqrt{\frac{d}{2s\vartheta e^{\vartheta}}}\right)^2\right),$$

as desired. \square

PROOF OF LEMMA 6.9. Claim 6.10 implies that for all $s \geq 1$,

$$(6.31) \quad \mathbb{P}(\Sigma_2^+ \geq s/\sqrt{k}) \leq \exp(-sk^{1/3}2^{k/4}).$$

Since the same estimates hold for $-\log \Pi^-$, (6.26), (6.30) and (6.31) show that for $s \geq 1$,

$$\begin{aligned} \hat{\mu}\left(\left[\frac{e^s}{1+e^s}, 1\right]\right) &\leq 2\mathbb{P}\left(\left|\log \Pi^+ + \frac{d}{2}\mathbb{E}[\bar{X}]\right| \geq \frac{s}{2}\right) \\ &\leq 2\left(\mathbb{P}\left(\left|\Sigma_1^+ + \Sigma_2^+\right| \geq \frac{s}{4}\right) + \mathbb{P}\left(\left|\sum_{i=1}^{\gamma^+} \bar{X}_i^+ - \frac{d}{2}\mathbb{E}[\bar{X}]\right| \geq \frac{s}{4}\right)\right) \\ &\leq \exp(-sk^{1/4}2^{k/4}), \end{aligned}$$

as claimed. \square

6.3. *Proof of Proposition 6.3.* The proof is an adaptation of the proof of [24], Lemma 4.4, where the authors showed a corresponding statement for the survey propagation distributional recursion, which is more complicated than the present belief propagation recurrence. Thus, we follow the path beaten in [24], simplifying the argument where possible. Throughout this section r denotes the smallest even integer greater than $2^{k/10}$. We set out to prove that on \mathcal{P}^\dagger the operator \mathcal{R} is a W_r -contraction. Hence, we consider two probability distributions $\rho, \rho' \in \mathcal{P}^\dagger$. Let $(\eta_{ij}^+, \chi_{ij}^+)_{i,j \geq 1}$, $(\eta_{ij}^-, \chi_{ij}^-)_{i,j}$ be independent identically distributed pairs of numbers in $[0, 1]$ such that the first components η_{ij}^\pm all have distribution ρ and the second components χ_{ij}^\pm have distribution ρ' and $\mathbb{E}[|\eta_{ij}^\pm - \chi_{ij}^\pm|^r]^{1/r} = W_r(\rho, \rho')$. Let $\eta = (\eta_{ij}^+, \eta_{ij}^-)_{i,j}$, $\chi = (\chi_{ij}^+, \chi_{ij}^-)_{i,j}$ and, recalling (6.1), define

$$\hat{\eta} = R(\boldsymbol{\gamma}^+, \boldsymbol{\gamma}^-, \eta), \quad \hat{\chi} = R(\boldsymbol{\gamma}^+, \boldsymbol{\gamma}^-, \chi).$$

To prove Proposition 6.3 we are going to couple $\hat{\eta}, \hat{\chi}$ such that

$$(6.32) \quad \mathbb{E}[|\hat{\eta} - \hat{\chi}|^r] \leq 2^{-kr/11} \mathbb{E}[|\eta_{11}^+ - \chi_{11}^+|^r] = W_r(\rho, \rho')^r.$$

To construct the coupling, let $\boldsymbol{\gamma} \stackrel{d}{=} \text{Po}(d)$. Moreover, let $(J_i)_{i \geq 1}$ be the uniform i.i.d. signs; $\boldsymbol{\gamma}$ and the J_i are mutually independent, and independent of everything else. Then, setting $\boldsymbol{\gamma}^+ = \sum_{i=1}^{\boldsymbol{\gamma}} \mathbb{1}\{J_i = 1\}$, $\boldsymbol{\gamma}^- = \sum_{i=1}^{\boldsymbol{\gamma}} \mathbb{1}\{J_i = -1\}$ ensures that $\boldsymbol{\gamma}^+, \boldsymbol{\gamma}^-$ are independent and $\text{Po}(d/2)$ variables. Further, let

$$\begin{aligned} P^\pm &= \prod_{i=1}^{\boldsymbol{\gamma}^\pm} \left(1 - (1 - e^{-\beta}) \prod_{j=1}^{k-1} \eta_{ij}^\pm\right), & \hat{\eta} &= \frac{P^+}{P^+ + P^-}, \\ Q^\pm &= \prod_{i=1}^{\boldsymbol{\gamma}^\pm} \left(1 - (1 - e^{-\beta}) \prod_{j=1}^{k-1} \chi_{ij}^\pm\right), & \hat{\chi} &= \frac{Q^+}{Q^+ + Q^-}. \end{aligned}$$

Then

$$(6.33) \quad \hat{\eta} - \hat{\chi} = \frac{P^+}{P^+ + P^-} - \frac{Q^+}{Q^+ + Q^-} = \varphi\left(\log \frac{P^+}{P^-}\right) - \varphi\left(\log \frac{Q^+}{Q^-}\right).$$

Finally, to estimate the last expression we introduce

$$\Delta_i = \log\left(1 - (1 - e^{-\beta}) \prod_{j=1}^{k-1} \chi_{ij}^+\right) - \log\left(1 - (1 - e^{-\beta}) \prod_{j=1}^{k-1} \eta_{ij}^+\right).$$

Since $0 \leq \varphi'(x) \leq 1$ for all x , (6.33) and the mean value theorem imply

$$(6.34) \quad \mathbb{E}[|\hat{\eta} - \hat{\chi}|^r] \leq \mathbb{E}\left[\left|\log \frac{P^+}{P^-} - \log \frac{Q^+}{Q^-}\right|^r\right] = \mathbb{E}\left[\left|\sum_{i=1}^{\gamma} J_i \Delta_i\right|^r\right].$$

Several steps are required to bound the right-hand side of (6.34).

LEMMA 6.11. *We have $\mathbb{E}[|\sum_{i=1}^{\gamma} J_i \Delta_i|^r] \leq (r - 1)!! \mathbb{E}[\gamma^{r/2}] \mathbb{E}[\Delta_1^r]$.*

PROOF. Since r is even we have

$$(6.35) \quad \begin{aligned} \mathbb{E}\left[\left|\sum_{i=1}^{\gamma} J_i \Delta_i\right|^r\right] &= \mathbb{E}\left[\sum_{i_1, \dots, i_r=1}^{\gamma} \prod_{\ell=1}^r J_{i_\ell} \Delta_{i_\ell}\right] \\ &= e^{-d} \sum_{\gamma=0}^{\infty} \frac{d^\gamma}{\gamma!} \sum_{i_1, \dots, i_r=1}^{\gamma} \mathbb{E}\left[\prod_{\ell=1}^r J_{i_\ell}\right] \mathbb{E}\left[\prod_{\ell=1}^r \Delta_{i_\ell}\right]. \end{aligned}$$

Let us now fix $\gamma \geq 0$ and $i_1, \dots, i_r \in [\gamma]$. Moreover, for $h \in [\gamma]$ let $N_h = N_h(i_1, \dots, i_r)$ be the number of indices ℓ such that $i_\ell = h$. Then

$$(6.36) \quad \begin{aligned} E(i_1, \dots, i_r) &= \mathbb{E}\left[\prod_{\ell=1}^r J_{i_\ell}\right] \mathbb{E}\left[\prod_{\ell=1}^r \Delta_{i_\ell}\right] = \mathbb{E}\left[\prod_{h=1}^{\gamma} J_h^{N_h}\right] \mathbb{E}\left[\prod_{h=1}^{\gamma} \Delta_h^{N_h}\right] \\ &= \prod_{h=1}^{\gamma} \mathbb{E}[J_h^{N_h}] \mathbb{E}[\Delta_h^{N_h}]. \end{aligned}$$

In particular, $E(i_1, \dots, i_r) = 0$ unless all N_h are even. Furthermore, if all N_h are even, then Jensen's inequality yields

$$(6.37) \quad E(i_1, \dots, i_r) = \prod_{h=1}^{\gamma} \mathbb{E}[\Delta_h^{N_h}] \leq \mathbb{E}[\Delta_1^r].$$

Finally, let $f(\gamma, r) = |\{(i_1, \dots, i_r) \in [\gamma]^r : N_h \bmod 2 = 0 \text{ for all } h \in [\gamma]\}|$ be the number of index sequences in which each index appears an even number of times. Combining (6.35)–(6.37), we conclude that

$$(6.38) \quad \mathbb{E}\left[\left|\sum_{i=1}^{\gamma} J_i \Delta_i\right|^r\right] \leq e^{-d} \mathbb{E}[\Delta_1^r] \sum_{\gamma=0}^{\infty} \frac{d^\gamma f(\gamma, r)}{\gamma!}.$$

We now claim that for all γ ,

$$(6.39) \quad f(\gamma, r) \leq (r - 1)!! \gamma^{r/2}.$$

To see this consider a family (i_1, \dots, i_r) such that all N_h are even and think of $1, \dots, r$ as the vertices of a complete graph. Since all N_h are even we can find a perfect matching of $1, \dots, r$ such that any two indices u, v that are matched satisfy $i_u = i_v$. Moreover, if we label the matching edge uv with the value $i_u = i_v \in [\gamma]$, then (i_1, \dots, i_r) can be recovered from the labelled matching. Because the total number of perfect matchings equals $(r - 1)!!$ and there are $\gamma^{r/2}$ possible labellings, we obtain (6.39). Finally, combining (6.38) and (6.39), we obtain

$$\mathbb{E}\left[\left|\sum_{i=1}^{\gamma} J_i \Delta_i\right|^r\right] \leq \sum_{\gamma=0}^{\infty} \frac{d^\gamma \gamma^{r/2} (r - 1)!!}{\gamma! \exp(d)} \mathbb{E}[\Delta_1^r] = (r - 1)!! \mathbb{E}[\gamma^{r/2}] \mathbb{E}[\Delta_1^r],$$

as claimed. \square

Hence, we are left to bound $\mathbb{E}[\gamma^{r/2}]$ and $\mathbb{E}[\Delta_1^r]$.

LEMMA 6.12. We have $\mathbb{E}[\boldsymbol{\gamma}^{r/2}] \leq \left(\frac{r(d+1)}{2}\right)^{r/2}$.

PROOF. Let $\left\{ \begin{smallmatrix} N \\ K \end{smallmatrix} \right\}$ denote the Stirling number of the second kind. Because the factorial moments of the Poisson variable $\boldsymbol{\gamma}$ satisfy $\mathbb{E}[\prod_{j=1}^h \boldsymbol{\gamma} - j + 1] = d^h$, we obtain

$$\begin{aligned} \mathbb{E}[\boldsymbol{\gamma}^{r/2}] &= \sum_{h=0}^{r/2} \left\{ \begin{smallmatrix} r/2 \\ h \end{smallmatrix} \right\} \mathbb{E} \left[\prod_{j=1}^h \boldsymbol{\gamma} - j + 1 \right] = \sum_{h=0}^{r/2} \left\{ \begin{smallmatrix} r/2 \\ h \end{smallmatrix} \right\} d^h \\ &\leq \sum_{h=0}^{r/2} \binom{r/2}{h} h^{r/2-h} d^h \leq \left(\frac{r(d+1)}{2} \right)^{r/2}, \end{aligned}$$

as desired. \square

To bound $\mathbb{E}[\Delta_1^r]$ set $\boldsymbol{U}_j = \prod_{h=1}^{j-1} \boldsymbol{\chi}_{1h}^+ \prod_{h=j+1}^{k-1} \boldsymbol{\eta}_{1h}^+$ for $j \in [k-1]$. Writing a telescoping sum and applying the mean value theorem, we obtain

$$\begin{aligned} |\Delta_1| &\stackrel{d}{=} \left| \log \left(1 - (1 - e^{-\beta}) \prod_{j=1}^{k-1} \boldsymbol{\eta}_{1j}^+ \right) - \log \left(1 - (1 - e^{-\beta}) \prod_{j=1}^{k-1} \boldsymbol{\chi}_{1j}^+ \right) \right| \\ (6.40) \quad &= \left| \sum_{j=1}^{k-1} \log \frac{1 - (1 - e^{-\beta}) \boldsymbol{U}_j \boldsymbol{\eta}_{1j}^+}{1 - (1 - e^{-\beta}) \boldsymbol{U}_j \boldsymbol{\chi}_{1j}^+} \right| \leq \sum_{j=1}^{k-1} \left| \log \left(\frac{1 - (1 - e^{-\beta}) \boldsymbol{U}_j \boldsymbol{\eta}_{1j}^+}{1 - (1 - e^{-\beta}) \boldsymbol{U}_j \boldsymbol{\chi}_{1j}^+} \right) \right| \\ &\leq (1 - e^{-\beta}) \sum_{j=1}^{k-1} \frac{\boldsymbol{U}_j}{1 - \boldsymbol{U}_j} |\boldsymbol{\eta}_{1j}^+ - \boldsymbol{\chi}_{1j}^+|. \end{aligned}$$

Since $|\sum_{j=1}^{k-1} z_j|^r \leq k^{r-1} \sum_{j=1}^{k-1} z_j^r$ for all $z_1, \dots, z_{k-1} \geq 0$ by Hölder’s inequality, (6.40) implies that

$$\begin{aligned} (6.41) \quad \mathbb{E}[|\Delta_1|^r] &\leq (1 - e^{-\beta})^r \mathbb{E} \left[\left| \sum_{j=1}^{k-1} \frac{\boldsymbol{U}_j}{1 - \boldsymbol{U}_j} |\boldsymbol{\eta}_{1j}^+ - \boldsymbol{\chi}_{1j}^+| \right|^r \right] \\ &\leq k^{r-1} (1 - e^{-\beta})^r \sum_{j=1}^{k-1} \mathbb{E} \left[\left| \frac{\boldsymbol{U}_j}{1 - \boldsymbol{U}_j} \right|^r \right] \mathbb{E}[|\boldsymbol{\eta}_{1j}^+ - \boldsymbol{\chi}_{1j}^+|^r] \\ (6.42) \quad &\leq k^r \mathbb{E}[|\boldsymbol{\eta}_{1j}^+ - \boldsymbol{\chi}_{1j}^+|^r] \max_{j=1, \dots, k-1} \mathbb{E} \left[\left| \frac{\boldsymbol{U}_j}{1 - \boldsymbol{U}_j} \right|^r \right]. \end{aligned}$$

So, by Cauchy–Schwarz

$$(6.43) \quad \mathbb{E}[|\Delta_1|^r] \leq k^r \mathbb{E}[|\boldsymbol{\eta}_{1j}^+ - \boldsymbol{\chi}_{1j}^+|^r] \max_{j=1, \dots, k-1} \mathbb{E}[\boldsymbol{U}_j^{2r}]^{1/2} \mathbb{E}[(1 - \boldsymbol{U}_j)^{-2r}]^{1/2}.$$

Hence, we need to bound $\mathbb{E}[(1 - \boldsymbol{U}_j)^{-2r}]$ and $\mathbb{E}[\boldsymbol{U}_j^{2r}]$. To this end we need the following estimate.

LEMMA 6.13. For any $\mu \in \mathcal{P}^\dagger$ and $s \in [2^{-k/4}, 1]$, we have $\mu([1/2 + s, 1]) \leq \exp(-s2^{k/4})$.

PROOF. We notice that $e^s/(1 + e^s) \leq 1/2 + s$. Therefore, (6.2) implies that

$$\mu([1/2 + s, 1]) \leq \mu \left(\left[\frac{e^s}{1 + e^s}, 1 \right] \right) \leq \exp(-s2^{k/4}),$$

as claimed. \square

COROLLARY 6.14. *We have $\mathbb{E}[(1 - U_j)^{-2r}] \leq 5^r$ and $\mathbb{E}[U_j^{2r}] \leq 2^{-9(k-2)r/5}$ for all $j \in [k - 1]$.*

PROOF. Suppose that $z \geq 2$ satisfies $\log(z - 1) \geq 2^{-k/4}$. Then (6.2) ensures that

$$\begin{aligned}
 \mathbb{P}[(1 - U_j)^{-1} \geq z] &= \mathbb{P}\left[\prod_{h=1}^{j-1} \chi_{1h}^+ \prod_{h=j+1}^{k-1} \eta_{1h}^+ \geq 1 - \frac{1}{z}\right] \\
 &\leq \mathbb{P}[\varphi(\log(z - 1)) \leq \chi_{11}^+]^{j-1} \mathbb{P}[\varphi(\log(z - 1)) \leq \eta_{11}^+]^{k-1-j} \\
 &\leq \exp(-(k - 2) \log(z - 1) 2^{k/4}) \leq (z - 1)^{-2^{k/4}}.
 \end{aligned}
 \tag{6.44}$$

Hence, letting $\xi = 2 + 2^{-k/5}$, we obtain

$$\begin{aligned}
 \mathbb{E}[(1 - U_j)^{-2r}] &= \mathbb{E}[(1 - U_j)^{-2r} \mathbb{1}\{(1 - U_j)^{-1} < \xi\}] \\
 &\quad + \mathbb{E}[(1 - U_j)^{-2r} \mathbb{1}\{(1 - U_j)^{-1} \geq \xi\}] \\
 &\leq 2^{2r+1} + 2r \int_{\xi}^{\infty} z^{2r-1} \mathbb{P}[(1 - U_j)^{-1} \geq z] dz \\
 &\stackrel{(6.44)}{\leq} 2^{2r+1} + 2r \int_{\xi}^{\infty} \frac{z^{2r-1} dz}{(z - 1)^{2^{k/4}}} \\
 &\leq 2^{2r+1} + 2r \int_1^{\infty} \frac{(z + 1)^{2r-1}}{z^{2^{k/4}}} dz \\
 &\leq 2^{2r+1} + 2r \int_1^{\infty} \frac{(2z)^{2r-1}}{z^{2^{k/4}}} dz \\
 &= 2^{2r+1} + r4^r \int_1^{\infty} \frac{dz}{z^{2^{k/4} - (2r-1)}} \\
 &= 2^{2r+1} + \frac{r4^r}{2^{k/4} - 2r} \leq 5^r.
 \end{aligned}
 \tag{6.45}$$

Next, in order to bound $\mathbb{E}[U_j^{2r}]$, we use the fact that all η_{ij}^+ are bounded above by 1 and Lemma 6.13 to compute

$$\begin{aligned}
 \mathbb{E}[(\eta_{11}^+)^{2r}] &= \mathbb{E}\left[(\eta_{11}^+)^{2r} \mathbb{1}\left\{(\eta_{11}^+)^{2r} \leq \frac{(1 + 2^{-k/30})^{2r}}{2^{2r}}\right\}\right] \\
 &\quad + \mathbb{E}\left[(\eta_{11}^+)^{2r} \mathbb{1}\left\{(\eta_{11}^+)^{2r} > \frac{(1 + 2^{-k/30})^{2r}}{2^{2r}}\right\}\right] \\
 &\leq \frac{(1 + 2^{-k/30})^{2r}}{2^{2r}} + \mathbb{P}\left[\eta_{11}^+ > \frac{1 + 2^{-k/30}}{2}\right] \\
 &\leq \frac{\exp(2r2^{-k/30})}{2^{2r}} + \exp(-2^{-k/30-1}2^{k/4}) \leq 2^{-9r/5}.
 \end{aligned}
 \tag{6.46}$$

Since due to Proposition 6.2 the same bound holds for $\mathbb{E}[(\chi_{11}^+)^{2r}]$, we obtain $\mathbb{E}[U_j^{2r}] \leq 2^{-9(k-2)r/5}$. \square

COROLLARY 6.15. *We have $\mathbb{E}[\Delta_1^r] \leq (5k^2)^{r/2} 2^{-9(k-2)r/10} \mathbb{E}[|\eta_{11}^+ - \chi_{11}^+|^r]$.*

PROOF. This is an immediate consequence of (6.43) and Corollary 6.14. \square

PROOF OF PROPOSITION 6.3. Combining (6.34), Lemmas 6.11, 6.12 and Corollary 6.15 and recalling that $d \leq k^2 2^k$, we obtain

$$\begin{aligned} \mathbb{E}[|\hat{\eta} - \hat{\chi}|^r] &\leq (r-1)! \left(\frac{r(d+1)}{2}\right)^{r/2} k^r 5^{r/2} 2^{-9r(k-2)/10} \mathbb{E}[|\eta_{1j}^+ - \chi_{1j}^+|^r] \\ &\leq \frac{r!}{(r/2)!} \left(\frac{r(d+1)}{2}\right)^{r/2} 2^{-0.89kr} \mathbb{E}[|\eta_{1j}^+ - \chi_{1j}^+|^r] \\ &\leq r^{r/2} 2^{-0.4kr} \mathbb{E}[|\eta_{1j}^+ - \chi_{1j}^+|^r] \leq 2^{-0.2kr} \mathbb{E}[|\eta_{1j}^+ - \chi_{1j}^+|^r], \end{aligned}$$

which implies (6.32). Hence, \mathcal{R} is a W_r -contraction. \square

6.4. *Proof of Proposition 6.4.* Throughout this section, additionally to $d \leq d_{k\text{-SAT}}$, we assume that π is a probability distribution with slim tails. Let η be a random variable with distribution π ; then η satisfies

$$(6.47) \quad \mathbb{P}[|\eta - 1/2| \geq 2^{-k/10}] \leq 2^{-k/10}.$$

6.4.1. *Overview.* Our aim is to study the distribution $\mathcal{R}^\ell(\pi)$. Because of the inherent symmetry of the operator \mathcal{R} , we may assume that $\pi \in \mathcal{P}^\star$. In the following, let $\xi^{(\ell)}$ be a random variable with distribution $\mathcal{R}^\ell(\pi)$. To prove the proposition, we need to bound the tails of $\xi^{(\ell)}$. We proceed in two steps. First, in Section 6.4.2 we derive the following estimate.

LEMMA 6.16. *There exists an integer $L = L(k)$ such that for all $\ell \geq L$ and all $s \in [2^{-k/4}, 1]$ the random variable $\xi^{(\ell)}$ satisfies*

$$(6.48) \quad \mathbb{P}\left[\left|\log \frac{\xi^{(\ell)}}{1 - \xi^{(\ell)}}\right| \geq s\right] \leq \exp(-s2^{k/4}).$$

Moreover, in Section 6.4.3 we will prove the following. Let us introduce the shorthand $\xi = \xi^{(1)}$ for a random variable with distribution $\mathcal{R}(\pi)$.

LEMMA 6.17. *Assume that for a number $s \geq 1/2$ we have*

$$\mathbb{P}\left[\left|\log \frac{\eta}{1 - \eta}\right| \geq t\right] \leq \exp(-t2^{k/4}) \quad \text{for all } t \in [2^{-k/4}, s].$$

Then

$$\mathbb{P}\left[\left|\log \frac{\xi}{1 - \xi}\right| \geq 2s\right] \leq \exp(-2s \cdot 2^{k/4}).$$

PROOF OF PROPOSITION 6.4. Given $\varepsilon > 0$ choose $\ell_0 = \ell_0(\varepsilon) > 0$ large enough and assume that $\ell \geq \ell_0(\varepsilon)$. As in (6.4), set

$$\tilde{\xi}_\varepsilon^{(\ell)} = \xi^{(\ell)} \mathbb{1}\{\xi^{(\ell)} \in [\varepsilon, 1 - \varepsilon]\} + \frac{1}{2} \mathbb{1}\{\xi^{(\ell)} \notin [\varepsilon, 1 - \varepsilon]\}.$$

Then Lemmas 6.16 and 6.17 imply that $\tilde{\xi}_\varepsilon^{(\ell)}$ satisfies (6.2). Hence, $\mathcal{L}(\tilde{\xi}_\varepsilon^{(\ell)}) \in \mathcal{P}^\dagger$; this establishes part (ii) of the proposition. Furthermore, Lemma 6.17 also implies that

$$\mathbb{P}[\xi^{(\ell)} \notin [\varepsilon, 1 - \varepsilon]] < \varepsilon$$

for large enough ℓ , which is part (i). \square

6.4.2. *Proof of Lemma 6.16.* The following lemma summarises the key step towards the proof of Lemma 6.16.

LEMMA 6.18. *Let $\ell \geq 1$ be an integer such that $2^{-k/10-2(\ell-1)} \geq 2^{-k/4}$ and assume that*

$$(6.49) \quad \mathbb{P}[|\boldsymbol{\eta} - 1/2| \geq 2^{-k/10-2(\ell-1)}] \leq \max\{2^{-\frac{k}{10}(\frac{k}{50})^{\ell-1}}, \exp(-2^{k/4})\}.$$

Then $\mathbb{P}[|\boldsymbol{\xi} - 1/2| \geq 2^{-k/10-2\ell}] \leq \max\{2^{-\frac{k}{10}(\frac{k}{50})^\ell}, \exp(-2^{k/4})\}.$

The proof of Lemma 6.18 proceeds in three steps. We continue to let $(\boldsymbol{\eta}_{i,j}^\pm)_{i,j \geq 1}$ be an array of independent random variables distributed as $\boldsymbol{\eta}$. Also let $\boldsymbol{\gamma}^+, \boldsymbol{\gamma}^-$ be independent Poisson variables with mean $d/2$, independent of $(\boldsymbol{\eta}_{i,j}^\pm)_{i,j \geq 1}$, so that

$$(6.50) \quad \log \frac{\boldsymbol{\xi}}{1 - \boldsymbol{\xi}} \stackrel{d}{=} \sum_{i=1}^{\boldsymbol{\gamma}^+} \log \left(1 - (1 - e^{-\beta}) \prod_{j=1}^{k-1} \boldsymbol{\eta}_{i,j}^+ \right) - \sum_{i=1}^{\boldsymbol{\gamma}^-} \log \left(1 - (1 - e^{-\beta}) \prod_{j=1}^{k-1} \boldsymbol{\eta}_{i,j}^- \right).$$

For $i \geq 1$ let

$$\mathcal{A}_{i,\ell}^\pm = \left\{ \sum_{j=1}^{k-1} \mathbb{1} \left[\left| \boldsymbol{\eta}_{i,j}^\pm - \frac{1}{2} \right| \geq 2^{-k/10-2\ell+2} \right] \leq \frac{k-1}{10} \right\}$$

be the event that a ‘‘clause’’ i receives no more than $(k-1)/10$ ‘‘atypical messages’’. Moreover, let

$$\mathcal{A}_\ell = \bigcup_{m_1=0}^\infty \bigcup_{m_2=0}^\infty \left(\{\boldsymbol{\gamma}^+ = m_1, \boldsymbol{\gamma}^- = m_2\} \cap \bigcap_{i=1}^{m_1} \mathcal{A}_{i,\ell}^+ \cap \bigcap_{i=1}^{m_2} \mathcal{A}_{i,\ell}^- \right)$$

be the event that there is no clause with too many atypical messages. Our first goal is to bound the probability that \mathcal{A}_ℓ fails to occur given that for each message the probability of being atypical is small.

LEMMA 6.19. *Assume that for some $\frac{k}{10} \log 2 \leq z \leq 2^{k/4}$ we have*

$$(6.51) \quad \mathbb{P}[|\boldsymbol{\eta} - 1/2| \geq 2^{-k/10-2\ell+2}] \leq \exp(-z).$$

Then $\mathbb{P}[\mathcal{A}_\ell] \geq 1 - \exp(-zk/50)/3.$

PROOF. We first estimate $\mathbb{P}[\mathcal{A}_{i,\ell}^\pm]$. Since the $\boldsymbol{\eta}_{i,1}^\pm, \dots, \boldsymbol{\eta}_{i,(k-1)}^\pm$ are independent, the assumption (6.51) implies together with the Chernoff bound that

$$\begin{aligned} 1 - \mathbb{P}[\mathcal{A}_{i,\ell}^\pm] &\leq \mathbb{P}\left(\text{Bin}(k-1, \exp(-z)) > \frac{k-1}{10}\right) \\ &\leq \exp(-(k-1)D_{\text{KL}}(1/10 \parallel \exp(-z))) \leq \exp\left(-\frac{k-1}{20}z\right). \end{aligned}$$

Hence, by (4.1) and subadditivity,

$$(6.52) \quad \begin{aligned} 1 - \mathbb{P}[\mathcal{A}_\ell] &\leq \mathbb{P}[\boldsymbol{\gamma}^+ + \boldsymbol{\gamma}^- > 2k2^k] + \mathbb{P}[\{\boldsymbol{\gamma}^+ + \boldsymbol{\gamma}^- \leq 2k2^k\} \setminus \mathcal{A}_\ell] \\ &\leq \exp\left(-\frac{3}{8}k2^k\right) + 2k2^k \exp\left(-\frac{k-1}{20}z\right) \\ &\leq \exp\left(-\frac{3}{8}k2^k\right) + \exp(-zk/40). \end{aligned}$$

As $z \leq 2^{k/4} < 3k2^k/8$, (6.52) implies the claim. \square

We now analyse the sum (6.50) on the event \mathcal{A}_ℓ . Let $c_k = (\frac{1}{2} + 2^{-k/10-2\ell+2})^{9(k-1)/10}.$

CLAIM 6.20. For $s \geq 2^{-k/4}$ we have $\mathbb{P}[\mathbb{1}_{\mathcal{A}_\ell} \cdot |\log \frac{\xi}{1-\xi}| \geq s] \leq \exp(-s2^{k/2} - 1)$.

PROOF. On \mathcal{A}_ℓ we have $\prod_{j=1}^{k-1} \eta_{i,j}^+ \leq c_k$ for all $i \leq \gamma^+$ and $\prod_{j=1}^{k-1} \eta_{i,j}^- \leq c_k$ for all $i \leq \gamma^-$. Letting

$$Y_i^\pm = \log \left(1 - (1 - e^{-\beta}) \prod_{j=1}^{k-1} \eta_{i,j}^\pm \right) \mathbb{1} \left\{ \prod_{j=1}^{k-1} \eta_{i,j}^\pm \leq c_k \right\}$$

we thus obtain

$$\begin{aligned} & \mathbb{P} \left[\mathbb{1}_{\mathcal{A}_\ell} \cdot \left| \log \frac{\xi}{1-\xi} \right| \geq s \right] \\ &= \mathbb{P} \left[\mathcal{A}_\ell \cap \left\{ \left| \sum_{i=1}^{\gamma^+} Y_i^+ - \sum_{i=1}^{\gamma^-} Y_i^- \right| \geq s \right\} \right] \\ (6.53) \quad & \leq \mathbb{P} \left[\left| \sum_{i=1}^{\gamma^+} (Y_i^+ - \mathbb{E}[Y_i^+]) - \sum_{i=1}^{\gamma^-} (Y_i^- - \mathbb{E}[Y_i^-]) \right. \right. \\ & \quad \left. \left. + \mathbb{E}[Y_1^+] \left(\gamma^+ - \frac{d}{2} \right) - \mathbb{E}[Y_1^-] \left(\gamma^- - \frac{d}{2} \right) \right| \geq s \right] \\ & \leq 2\mathbb{P} \left[\left| \sum_{i=1}^{\gamma^+} Y_i^+ - \mathbb{E}[Y_i^+] \right| \geq \frac{s}{4} \right] + 2\mathbb{P} \left[\left| \mathbb{E}[Y_1^+] \left(\gamma^+ - \frac{d}{2} \right) \right| \geq \frac{s}{4} \right]. \end{aligned}$$

To estimate the last summand we first bound $|\mathbb{E}[Y_1^+]|$: using $-\log(1-x) \leq x + O(x^2)$, we obtain

$$|\mathbb{E}[Y_1^+]| \leq -\log(1 - (1 - e^{-\beta})c_k) \leq \frac{5}{4}c_k$$

for k sufficiently large. Therefore, Bennett’s inequality Lemma 4.5 shows that for $s \geq 2^{-k/4}$,

$$\begin{aligned} (6.54) \quad & \mathbb{P} \left[\left| \mathbb{E}[Y_1^+] \left(\gamma^+ - \frac{d}{2} \right) \right| \geq \frac{s}{4} \right] \\ & \leq \mathbb{P} \left[\left| \gamma^+ - \frac{d}{2} \right| \geq \frac{s}{5c_k} \right] \leq \mathbb{P} \left[\left| \gamma^+ - \frac{d}{2} \right| \geq \frac{s}{10} 2^{9(k-1)/10} \right] \leq \exp(-s2^{k/2} - 2). \end{aligned}$$

For the second last summand in (6.53), we condition on γ^+ and apply the Azuma-Hoeffding inequality. The definition of the random variable Y_i^+ ensures that

$$|Y_i^+ - \mathbb{E}[Y_i^+]| \leq -\log(1 - (1 - e^{-\beta})c_k) \leq c_k + O(c_k^2).$$

Hence, as in the computation towards (6.52) for $s \geq 2^{-k/4}$ we obtain

$$\begin{aligned} (6.55) \quad & \mathbb{P} \left[\left| \sum_{i=1}^{\gamma^+} Y_i^+ - \mathbb{E}[Y_i^+] \right| \geq \frac{s}{4} \right] \\ & \leq \mathbb{P}[\gamma^+ > k2^k] + \mathbb{P} \left[\{ \gamma^+ \leq k2^k \} \cap \left\{ \left| \sum_{i=1}^{\gamma^+} (Y_i^+ - \mathbb{E}[Y_i^+]) \right| \geq \frac{s}{4} \right\} \right] \\ & \leq \exp\left(-\frac{3}{8}k2^k\right) + 2\exp\left(-\frac{s^2}{100k2^kc_k^2}\right) \leq \exp(-s2^{k/2} - 3). \end{aligned}$$

Combining (6.53), (6.54) and (6.55) completes the proof. \square

The last ingredient that we need for the proof of Lemma 6.18 is the following.

CLAIM 6.21. *For each $s \in [0, 4]$ we have $\mathbb{P}(|\xi - \frac{1}{2}| \geq \frac{s}{4}) \leq \mathbb{P}(|\log \frac{\xi}{1-\xi}| \geq s)$.*

PROOF. For all $s \geq 0$ we have $\frac{e^s}{1+e^s} \leq \frac{1}{2} + \frac{s}{4}$. Therefore,

$$\mathbb{P}\left(\xi - \frac{1}{2} \geq \frac{s}{4}\right) \leq \mathbb{P}\left(\xi - \frac{1}{2} \geq \frac{e^s}{1+e^s} - \frac{1}{2}\right) = \mathbb{P}\left(\log\left(\frac{\xi}{1-\xi}\right) \geq s\right).$$

The symmetry of ξ , that is, that due to the definition (6.1) of \mathcal{R} the random variables ξ and $1 - \xi$ are identically distributed, therefore implies the claim. \square

PROOF OF LEMMA 6.18. Assume that (6.49) is satisfied for some $\ell \geq 1$. Then also the assumption of Lemma 6.19 with

$$z = \min\left\{\frac{k}{10}\left(\frac{k}{50}\right)^{\ell-1} \log 2, 2^{k/4}\right\}$$

is satisfied, and we will use this estimate to bootstrap (6.49). Indeed, Lemma 6.19 and Claims 6.20 and 6.21 yield

$$\mathbb{P}\left(|\xi - \frac{1}{2}| \geq 2^{-k/10-2\ell}\right) \leq \mathbb{P}\left(\left|\log \frac{\xi}{1-\xi}\right| \geq 2^{-k/10-2\ell+2}\right) \leq \max\{2^{-\frac{k}{10}\left(\frac{k}{50}\right)^\ell}, \exp(-2^{k/4})\},$$

as claimed. \square

PROOF OF LEMMA 6.16. The proof is by induction on ℓ . Since we assume that η satisfies (6.47), (6.49) holds for $\ell = 1$. Therefore, we may repeatedly apply Lemma 6.18 until for the first time $2^{-\frac{k}{10}-2\ell} < 2^{-k/4}$, which happens after $\lceil 3k/40 \rceil - 1$ steps. At this point, also

$$(6.56) \quad \max\{2^{-\frac{k}{10}\left(\frac{k}{50}\right)^\ell}, \exp(-2^{k/4})\} = \exp(-2^{k/4}).$$

Therefore, (6.56) implies that $\mathbb{P}(|\xi^{(\ell)} - \frac{1}{2}| \geq s) \leq \exp(-s2^{k/4})$ for all $s \in [2^{-k/4}, 1]$, whence (6.48) holds for all $\ell \geq \lceil 3k/40 \rceil + 1$. \square

6.4.3. *Proof of Lemma 6.17.* We combine some of the elements of the proofs of Lemmas 6.2 and 6.16. We continue to denote by $(\eta_{i,j}^\pm)_{i,j \geq 1}$ an array of independent copies of η . Moreover, γ^+, γ^- are Poisson variables with mean $d/2$, independent of $(\eta_{i,j}^\pm)_{i,j \geq 1}$. Further, for $s \geq 2^{-k/4}$ and $i = 1, \dots, \gamma^\pm$ let

$$\mathcal{A}_{i,s}^\pm = \left\{ \sum_{j=1}^{k-1} \mathbb{1}\left\{ \left| \log \frac{\eta_{i,j}^\pm}{1-\eta_{i,j}^\pm} \right| \geq s \right\} \leq \frac{k-1}{10} \right\} \quad \text{and}$$

$$\mathcal{A}_s = \bigcup_{m_1=0}^\infty \bigcup_{m_2=0}^\infty \left(\{\gamma^+ = m_1, \gamma^- = m_2\} \cap \bigcap_{i=1}^{m_1} \mathcal{A}_{i,s}^+ \cap \bigcap_{i=1}^{m_2} \mathcal{A}_{i,s}^- \right).$$

For the event \mathcal{A}_s we prove an analogue of Lemma 6.19.

LEMMA 6.22. *Assume that for some $s \geq \frac{1}{2}$ we have*

$$\mathbb{P}\left[\left|\log \frac{\eta}{1-\eta}\right| \geq s\right] \leq \exp(-s2^{k/4}).$$

Then $\mathbb{P}(\mathcal{A}_s) \geq 1 - \frac{1}{3} \exp(-s2^{1+k/4})$.

PROOF. Because the $\eta_{i,j}^\pm$ are mutually independent, the Chernoff bound yields

$$(6.57) \quad \mathbb{P}[\mathcal{A}_{i,s}^\pm] \geq 1 - \exp\left(-\frac{s(k-1)}{20}2^{k/4}\right).$$

Further, for $s > 1/2$ Lemma 4.5 and (6.57) yield

$$\begin{aligned} 1 - \mathbb{P}(\mathcal{A}_s) &\leq \mathbb{P}[\boldsymbol{y}^+ + \boldsymbol{y}^- > 4sk2^k] + \mathbb{P}[\{\boldsymbol{y}^+ + \boldsymbol{y}^- \leq 4sk2^k\} \setminus \mathcal{A}_s] \\ &\leq \exp\left(-\left(\log(4s) - \frac{4s-1}{4s}\right)4sk2^k\right) + 4sk2^k \exp\left(-\frac{k-1}{20}s2^{k/4}\right) \\ &\leq \frac{1}{3} \exp(-2s2^{k/4}), \end{aligned}$$

as desired. \square

Consider the nonnegative random variables

$$X_i^\pm = -\log\left(1 - (1 - e^{-\beta}) \prod_{j=1}^{k-1} \eta_{ij}^\pm\right) \geq 0, \quad Y_i^\pm = X_i^\pm \mathbb{1}\{X_i^\pm \leq (1 - e^{-\beta})2^{-9(k-1)/10}\}.$$

Then

$$\begin{aligned} \mathbb{P}\left(\mathbb{1}_{\mathcal{A}_s} \log \frac{\xi}{1-\xi} \geq 2s\right) &\leq \mathbb{P}\left(\mathcal{A}_s \cap \left\{\sum_{i=1}^{y^-} X_i^- - \sum_{i=1}^{y^+} X_i^+ \geq 2s\right\}\right) \\ &\leq \mathbb{P}\left(\mathcal{A}_s \cap \left\{\sum_{i=1}^{y^-} (X_i^- - Y_i^-) - \sum_{i=1}^{y^+} (X_i^+ - Y_i^+) + \sum_{i=1}^{y^-} (Y_i^- - \mathbb{E}[Y_i^-]) \right. \right. \\ &\quad \left. \left. - \sum_{i=1}^{y^+} (Y_i^+ \mathbb{E}[Y_i^+]) + \mathbb{E}[Y_1^-](y^- - y^+) \geq 2s\right\}\right). \end{aligned}$$

So,

$$(6.58) \quad \begin{aligned} \mathbb{P}\left(\mathbb{1}_{\mathcal{A}_s} \log \frac{\xi}{1-\xi} \geq 2s\right) &\leq \mathbb{P}\left(\mathcal{A}_s \cap \left\{\sum_{i=1}^{y^-} (X_i^- - Y_i^-) \geq \frac{s}{3}\right\}\right) \\ &\quad + 2 \cdot \mathbb{P}\left(\left|\sum_{i=1}^{y^-} (Y_i^- - \mathbb{E}[Y_i^-])\right| \geq \frac{s}{3}\right) \\ &\quad + 2 \cdot \mathbb{P}\left(\mathbb{E}[Y_1^-] \left|y^- - \frac{d}{2}\right| \geq \frac{s}{3}\right). \end{aligned}$$

We proceed to bound the three terms on the right-hand side of (6.58) separately, starting with the one in the middle.

LEMMA 6.23. We have $\mathbb{P}[|\sum_{i=1}^{y^-} Y_i^- - \mathbb{E}[Y_i^-]| \geq \frac{s}{3}] \leq \frac{1}{9} \exp(-2s2^{k/4})$.

PROOF. Lemma 4.5 and Azuma’s inequality yield

$$\begin{aligned} \mathbb{P}\left(\left|\sum_{i=1}^{y^-} Y_i^- - \mathbb{E}[Y_i^-]\right| \geq \frac{s}{3}\right) &\leq \mathbb{P}(y^- > 2sk2^k) \\ &\quad + \mathbb{P}\left(\{y^- \leq 2sk2^k\} \cap \left\{\left|\sum_{i=1}^{y^-} Y_i^- - \mathbb{E}[Y_i^-]\right| \geq \frac{s}{3}\right\}\right) \end{aligned}$$

$$\begin{aligned} &\leq \exp\left(-\left(\log(4s) - \frac{4s-1}{4s}\right)2sk2^k\right) \\ &\quad + 2 \exp\left(-\frac{s^2}{36sk2^k2^{-9(k-1)/5}}\right) \leq \frac{1}{9} \exp(-2s2^{k/4}), \end{aligned}$$

as claimed. \square

The rightmost term from (6.58) is next.

LEMMA 6.24. *We have $\mathbb{P}[\mathbb{E}[\mathbf{Y}_1^-]|\boldsymbol{\gamma}^- - \frac{d}{2}| \geq \frac{s}{3}] \leq \frac{1}{9} \exp(-2s2^{k/4})$.*

PROOF. The definition of \mathbf{Y}_1^\pm ensures that $\mathbb{E}[\mathbf{Y}_1^-] \leq 2^{-9(k-1)/10}$. Therefore, Lemma 4.5 yields

$$\mathbb{P}\left(\mathbb{E}[\mathbf{Y}_1^-] \mid \boldsymbol{\gamma}^- - \frac{d}{2} \geq \frac{s}{3}\right) \leq \mathbb{P}\left(\boldsymbol{\gamma}^- - \frac{d}{2} \geq \frac{s}{3}2^{9(k-1)/10}\right) \leq \frac{1}{9} \exp(-2s2^{k/4}),$$

as claimed. \square

Finally, the following lemma deals with the first term from (6.58).

LEMMA 6.25. *We have $\mathbb{P}[\mathcal{A}_s \cap \{\sum_{i=1}^{\boldsymbol{\gamma}^-} (\mathbf{X}_i^- - \mathbf{Y}_i^-) \geq \frac{s}{3}\}] \leq \frac{1}{9} \exp(-2s2^{k/4})$.*

The proof of Lemma 6.25 requires several steps.

CLAIM 6.26. *Assume that $s \geq 1/2$ and that for all $t \in [2^{-k/4}, s]$ we have*

$$\mathbb{P}\left(\log \frac{\eta}{1-\eta} \geq t\right) \leq \exp(-t2^{k/4}).$$

Then $\mathbb{P}(\mathbf{X}_1^- \geq t) \leq \exp(-\frac{t}{2}(k-1)2^{k/4})$ for all $t \in [1, 2s]$.

PROOF. For $t \in [1, 2s]$ we obtain

$$\begin{aligned} \mathbb{P}[\mathbf{X}_1^- \geq t] &\leq \mathbb{P}\left[\forall j \in [k-1]: \eta_{1j}^- \geq \frac{1-e^{-t}}{1-e^{-\beta}}\right] \\ &= \mathbb{P}\left[\log \frac{\eta}{1-\eta} \geq \log \frac{1-e^{-t}}{e^{-t}-e^{-\beta}}\right]^{k-1} \leq \mathbb{P}\left[\log \frac{\eta}{1-\eta} \geq \frac{t}{2}\right]^{k-1} \\ &\leq \exp\left(-\frac{t}{2}(k-1)2^{k/4}\right), \end{aligned}$$

as desired. \square

Thus, we have a tail bound for \mathbf{X}_1^- up to $2s$. To bound the probability that \mathbf{X}_1^- grows even larger, we are going to condition on the event \mathcal{A}_s . Indeed, on \mathcal{A}_s for all $i = 1, \dots, \boldsymbol{\gamma}^-$ we have

$$(6.59) \quad \mathbf{X}_i^- = -\log\left(1 - (1 - e^{-\beta}) \prod_{j=1}^{k-1} \eta_{ij}^-\right) \leq -\log\left(1 - (1 - e^{-\beta}) \left(\frac{e^s}{1+e^s}\right)^{9(k-1)/10}\right).$$

The following two claims show that the \mathbf{X}_i^- are bounded by $2s$ deterministically on \mathcal{A}_s .

CLAIM 6.27. *For all $s \leq \log k$ we have $-\log(1 - (1 - e^{-\beta}) (\frac{e^s}{1+e^s})^{9(k-1)/10}) \leq 1$.*

PROOF. This is equivalent to showing that

$$1 - (1 - e^{-\beta}) \left(\frac{e^s}{1 + e^s} \right)^{9(k-1)/10} \geq \frac{1}{e}.$$

The left-hand side is strictly decreasing in s and thus it is sufficient to show the claim for $s = \log k$, in which case

$$\begin{aligned} 1 - (1 - e^{-\beta}) \left(\frac{e^{\log k}}{1 + e^{\log k}} \right)^{9(k-1)/10} &= 1 - (1 - e^{-\beta}) \left(1 - \frac{1}{k+1} \right)^{9(k-1)/10} \\ &\geq 1 - (1 - e^{-\beta}) \exp\left(-\frac{9(k-1)}{10(k+1)} \right) \geq \frac{1}{2}, \end{aligned}$$

as desired. \square

CLAIM 6.28. For $s > \log k$ we have $-\log(1 - (1 - e^{-\beta}) (\frac{e^s}{1+e^s})^{9(k-1)/10}) \leq s$.

PROOF. We have

$$\left(1 - \frac{1}{1 + e^s} \right)^{9(k-1)/10} \leq 1 - \left(\frac{9}{10} \right)^2 (k-1)e^{-s} + \frac{1}{2} \left(\frac{9(k-1)}{10} \right)^2 e^{-2s}.$$

Therefore,

$$\begin{aligned} &-\log\left(1 - (1 - e^{-\beta}) \left(\frac{e^s}{1 + e^s} \right)^{9(k-1)/10} \right) \\ &\leq -\log\left(1 - \left(\frac{e^s}{1 + e^s} \right)^{9(k-1)/10} \right) \\ &\leq -\log\left(\left(\frac{9}{10} \right)^2 (k-1)e^{-s} - \frac{1}{2} \left(\frac{9(k-1)}{10} \right)^2 e^{-2s} \right) \\ &= -2\log\left(\frac{9}{10} \right) - \log(k-1) + s - \log\left(1 - \frac{k-1}{2} e^{-s} \right) \leq s, \end{aligned}$$

as claimed. \square

PROOF OF LEMMA 6.25. On the event \mathcal{A}_s , by Claim 6.27 and 6.28 we have

$$\begin{aligned} \sum_{i=1}^{\gamma^-} X_i^- - Y_i^- &\leq \sum_{i=1}^{\gamma^-} \mathbb{1}\{X_i^- \geq (1 - e^{-\beta})2^{-9(k-1)/10}\} \\ &\quad + \sum_{i=1}^{\gamma^-} X_i^- \mathbb{1}\{X_i^- \in [1, s]\} \mathbb{1}\{X_i \geq (1 - e^{-\beta})2^{-9(k-1)/10}\}. \end{aligned}$$

We now reprove Lemma 6.5, part (ii), with $\varepsilon = 1/10$ to get an upper bound on

$$\mathbb{P}[X_i^- \geq (1 - e^{-\beta})2^{-9(k-1)/10}].$$

First of all, as in (6.12), rearranging and elimination of β gives

$$(6.60) \quad \mathbb{P}[X_1^- \geq (1 - e^{-\beta})2^{-(9/10)(k-1)}] \leq \mathbb{P}\left[\prod_{j=1}^{k-1} \eta_{1,j}^+ \geq 1 - \exp(-2^{-(9/10)(k-1)}) \right].$$

Further, since $\eta_{1j} \leq 1$ for all j , for any $a \in [0, 1]$ and $b \in (0, k - 2]$ we have

$$(6.61) \quad \mathbb{P} \left[\prod_{j=1}^{k-1} \eta_{1j}^+ \geq a \right] \leq \mathbb{P} [|\{j \in [k - 1] : \eta_{1j}^+ \geq a^{\frac{1}{k-1-b}}\}| \geq b].$$

Moreover, for large k we have

$$(6.62) \quad 1 - \exp(-2^{-(9/10)(k-1)}) \geq 2^{-(k-1)(29/30)^2},$$

which was proved in Section 6.2.2. Combining (6.60), (6.61) and (6.62), we obtain

$$\begin{aligned} \mathbb{P}(X_1^- \geq (1 - e^{-\beta})2^{-(9/10)(k-1)}) &\leq \mathbb{P} \left(\prod_{j=1}^{k-1} \eta_{1j}^+ \geq 2^{-(k-1)(29/30)^2} \right) \\ &\leq \mathbb{P} (|\{j \in [k - 1] : \eta_{1j}^+ \geq 2^{-29/30}\}| \geq (k - 1)/30) \\ &\leq \sum_{(k-1)/30d \leq j \leq k-1} \binom{k-1}{j} \mathbb{P}[\eta_{11}^+ \geq 2^{-29/30}]^j \\ &\leq 2^{k-1} \mathbb{P}[\eta_{11}^+ \geq 2^{-29/30}]^{(k-1)/30}. \end{aligned}$$

Since $1/2 \geq \varphi^{-1}(2^{-29/30}) \geq 2^{-k/4}$, the assumption of Lemma 6.17 implies that

$$\mathbb{P}(X \geq (1 - e^{-\beta})2^{-(9/10)(k-1)}) \leq 2^{k-1} \exp\left(-\frac{(k-1)}{30} 2^{k/4} \varphi^{-1}(2^{-29/30})\right).$$

Finally, since $\varphi^{-1}(2^{-29/30}) \geq \frac{2}{30} \log 2$ and for $k \geq 100$ and $c := 1/1000$

$$(6.63) \quad \mathbb{P}[X_1^- \geq (1 - e^{-\beta})2^{-9(k-1)/10}] \leq \exp(-ck2^{k/4}).$$

Moreover, for all $a \in [1, s]$, Lemma 6.26 implies that also

$$(6.64) \quad \mathbb{P}(X_1^- \geq a) \leq \exp(-2cak2^{k/4}).$$

Set $\vartheta = ck2^{k/4}$. Then (6.63), the definition of the random variable and (6.64) show that for $z \geq 0$,

$$\begin{aligned} &\mathbb{P}(\mathbb{1}\{X_i^- > (1 - e^{-\beta})2^{-(9/10)(k-1)}\} (1 + \mathbb{1}\{X_i^- \in [1, s]\} X_i^-) > z) \\ &\leq \begin{cases} \mathbb{P}(X_i^- > (1 - e^{-\beta})2^{-(9/10)(k-1)}) \leq e^{-\vartheta} & \text{if } z < 1, \\ \mathbb{P}(X_i^- \geq 1) \leq e^{-2\vartheta} & \text{if } 1 \leq z < 2, \\ \mathbb{P}(\mathbb{1}\{X_i^- \in [1, s]\} X_i^- \geq z - 1) \leq e^{-z\vartheta} & \text{if } z \geq 2. \end{cases} \end{aligned}$$

Hence, we can estimate these random variables as follows. Let $(I_i^-)_{i \geq 1}$ be a sequence of $\text{Be}(e^{-\vartheta})$ random variables, let $(Z_i^-)_{i \geq 1}$ be a sequence of exponential random variables with mean $1/\vartheta$ and let $\bar{\gamma}^-$ be a $\text{Po}(d/2)$ random variable, all mutually independent. Then for all $i \geq 1, z \geq 0$ we have

$$\begin{aligned} \mathbb{P}(\mathbb{1}\{X_i^- > (1 - e^{-\beta})2^{-(9/10)(k-1)}\} (1 + \mathbb{1}\{X_i^- \in [1, s]\} X_i^-) > z) &\leq \mathbb{P}(I_i^- (1 + Z_i^-) > z) \\ &= \begin{cases} e^{-\vartheta}, & z < 1, \\ e^{-z\vartheta}, & z \geq 1. \end{cases} \end{aligned}$$

Thus, $\mathbb{1}\{X_i^- > (1 - e^{-\beta})2^{-(9/10)(k-1)}\} (1 + \mathbb{1}\{X_i^- > 1\} X_i^-)$ is stochastically dominated by $I_i^- (1 + Z_i^-)$. Therefore, we also obtain stochastic dominance for the sums of these random

variables, that is,

$$(6.65) \quad \sum_{i=1}^{\gamma^-} \mathbb{1}(X_i^- > (1 - e^{-\beta})2^{-(9/10)(k-1)})(1 + \mathbb{1}(X_i^- > 1)X_i^-) \leq \Sigma_1^- + \Sigma_2^-,$$

where

$$\Sigma_1^- = \sum_{i=1}^{\bar{\gamma}^-} I_i^- \quad \text{and} \quad \Sigma_2^- = \sum_{i=1}^{\bar{\gamma}^-} I_i^- Z_i^-.$$

Hence, as in the proof of Lemma 6.2, we can stochastically dominate the difference $\sum_{i=1}^{\gamma^-} (X_i^- - Y_i^-)$ by a sum of a Poisson random variable and a random variable that is Gamma distributed, conditionally on the Poisson variable. Thus, we obtain

$$\mathbb{P}\left[\mathcal{A}_s \cap \left\{ \sum_{i=1}^{\gamma^-} (X_i^- - Y_i^-) \geq s/3 \right\}\right] \leq \mathbb{P}[\Sigma_1^- + \Sigma_2^- \geq s/3] \leq \mathbb{P}[\Sigma_1^- \geq s/6] + \mathbb{P}[\Sigma_2^- \geq s/6],$$

where Σ_1^- has distribution $\text{Po}(\frac{d}{2} \exp(-ck2^{k/4}))$ and Σ_2^- has distribution $\Gamma(\Sigma_1, 2ck2^{k/4})$. Bennett's inequality yields

$$\begin{aligned} \mathbb{P}[\Sigma_1^- \geq s/6] &\leq \exp\left(\frac{s}{6} - \frac{d}{2} \exp(-ck2^{k/4}) + \frac{s}{6} \log\left(\frac{d}{2} \exp(-ck2^{k/4})\right) - \frac{s}{6} \log\left(\frac{s}{6}\right)\right) \\ &\leq \frac{1}{18} \exp(-2s2^{k/4}). \end{aligned}$$

Moreover, we again set $\vartheta := ck2^{k/4}$. Then $d/(\vartheta \exp(\vartheta)) \rightarrow 0$ as $k \rightarrow \infty$, and Claim 6.10 yields that for k sufficiently large,

$$\begin{aligned} \mathbb{P}(\Sigma_2^- \geq s/6) &\leq \exp\left(-ck2^{k/4} \frac{s}{6} \left(1 - 2\sqrt{\frac{d}{12s\vartheta \exp(\vartheta)}} + \frac{d}{12s\vartheta \exp(\vartheta)}\right)\right) \\ &\leq \frac{1}{18} \exp(-2s2^{k/4}), \end{aligned}$$

which completes the proof. \square

PROOF OF LEMMA 6.17. Combining (6.58) with the estimates from Lemma 6.22, 6.23, 6.24 and 6.25 yields

$$\mathbb{P}\left[\left|\log \frac{\xi}{1-\xi}\right| \geq 2s\right] \leq 1 - \mathbb{P}[\mathcal{A}_s] + \mathbb{P}\left[\mathbb{1}_{\mathcal{A}_s} \log \frac{\xi}{1-\xi} \geq 2s\right] \leq \exp(-2s \cdot 2^{k/4}),$$

as claimed. \square

7. Proof of Proposition 2.5. Throughout this section we assume that (1.8) is satisfied. We start by estimating the difference of the actual variable-to-clause messages and the pseudo-messages.

LEMMA 7.1. For any $\varepsilon > 0$ there is t_0 such that for $t > t_0$ and for large enough n we have

$$\mathbb{E} \sum_{i=1}^n \sum_{a \in \partial x_i} |\mu_{\Phi, \beta, x_i \rightarrow a}(1) - \mu_{\Phi, \beta, x_i \rightarrow a, t}(1)| < \varepsilon n.$$

PROOF. Observe that the double sum amounts to choosing a random clause \mathbf{a} of Φ and then a random variable \mathbf{x} that appears in \mathbf{a} . Therefore, it suffices to prove that

$$(7.1) \quad \mathbb{E}|\mu_{\Phi,\beta,\mathbf{x}\rightarrow\mathbf{a}}(1) - \mu_{\Phi,\beta,\mathbf{x}\rightarrow\mathbf{a},t}(1)| = o_t(1).$$

To prove (7.1), we recall that by construction the pseudo messages satisfy

$$(7.2) \quad \mu_{\Phi,\beta,\mathbf{x}\rightarrow\mathbf{a}}(1) = \mu_{\Phi-\mathbf{a},\beta}(\{\sigma_{\mathbf{x}} = 1\});$$

in other words, the pseudo messages from \mathbf{x} to \mathbf{a} is just the Boltzmann marginal of \mathbf{x} in the formula $\Phi - \mathbf{a}$ obtained by removing \mathbf{a} . Moreover, to get a handle on the message $\mu_{\Phi,\beta,\mathbf{x}\rightarrow\mathbf{a},t}(1)$ we define

$$\mu_{\Phi,\beta,\mathbf{x},t}(s) = \frac{\prod_{b \in \partial x_i} \mu_{\Phi,\beta,b \rightarrow \mathbf{x},t}(s)}{\prod_{b \in \partial \mathbf{x}} \mu_{\Phi,\beta,b \rightarrow \mathbf{x},t}(1) + \prod_{b \in \partial \mathbf{x}} \mu_{\Phi,\beta,b \rightarrow \mathbf{x},t}(-1)} \quad \text{for } s \in \{\pm 1\}.$$

Thus, unlike in the original definition (1.6) of the belief propagation message, we do not drop any neighbour of \mathbf{x} in the products on the right-hand side. Furthermore, for any t , for large enough n the depth- $2t + 2$ neighbourhood of \mathbf{x} in $G(\Phi)$ is acyclic a.a.s. In this event the construction (1.5)–(1.6) of the belief propagation messages ensures that for any $b \in \partial \mathbf{x} \setminus \{\mathbf{a}\}$ we have $\mu_{\Phi,\beta,b \rightarrow \mathbf{x},t}(s) = \mu_{\Phi-\mathbf{a},\beta,b \rightarrow \mathbf{x},t}(s)$. Therefore, we obtain

$$(7.3) \quad \begin{aligned} \mu_{\Phi,\beta,\mathbf{x}\rightarrow\mathbf{a},t}(s) &= \frac{\prod_{b \in \partial \mathbf{x} \setminus \{\mathbf{a}\}} \mu_{\Phi,\beta,b \rightarrow \mathbf{x},t}(s)}{\prod_{b \in \partial \mathbf{x} \setminus \{\mathbf{a}\}} \mu_{\Phi,\beta,b \rightarrow \mathbf{x},t}(1) + \prod_{b \in \partial \mathbf{x} \setminus \{\mathbf{a}\}} \mu_{\Phi,\beta,b \rightarrow \mathbf{x},t}(-1)} \\ &= \mu_{\Phi-\mathbf{a},\beta,\mathbf{x},t}(s). \end{aligned}$$

Furthermore, because the total number m of clauses of the random formula Φ is a Poisson variable with standard deviation $\Theta(\sqrt{n})$, the random formulas Φ and $\Phi - \mathbf{a}$ have total variation distance $o(1)$. Therefore, there exists a coupling of Φ and $\Phi - \mathbf{a}$ such that these two random formulas coincide a.a.s. Consequently, (7.2) and (7.3) show that

$$(7.4) \quad \begin{aligned} \mathbb{E}|\mu_{\Phi,\beta,\mathbf{x}\rightarrow\mathbf{a}}(1) - \mu_{\Phi,\beta,\mathbf{x}\rightarrow\mathbf{a},t}(1)| &= \mathbb{E}|\mu_{\Phi-\mathbf{a},\beta}(\{\sigma_{\mathbf{x}} = 1\}) - \mu_{\Phi-\mathbf{a},\beta,\mathbf{x},t}(1)| \\ &= \mathbb{E}|\mu_{\Phi,\beta}(\{\sigma_{\mathbf{x}} = 1\}) - \mu_{\Phi,\beta,\mathbf{x},t}(1)| + o(1). \end{aligned}$$

We thus proceed to show that

$$(7.5) \quad \mathbb{E}|\mu_{\Phi,\beta}(\{\sigma_{\mathbf{x}} = 1\}) - \mu_{\Phi,\beta,\mathbf{x},t}(1)| = o_t(1).$$

Picking a large enough $t > 0$ and assuming that n is sufficiently large, we may condition on the event \mathcal{S} that the depth- $2t$ neighbourhood of x_1 in the factor graph $G(\Phi)$ is acyclic and that the total number of variables and clauses in this neighbourhood is bounded by $(kd)^{2t}$. To prove (7.5) let Φ^- denote the random formula obtained by deleting all clauses b_1, \dots, b_ℓ at distance exactly $2t - 1$ from x_1 in Φ . Further, obtain Φ^+ from Φ^- by inserting new clauses b'_1, \dots, b'_ℓ instead such that:

- each b'_i is connected with the same variable y_i at distance $2t - 2$ from x_1 as b_i with the same sign, that is, $\text{sign}(y_i, b_i) = \text{sign}(y_i, b'_i)$.
- the other variables $y_{ij}, j \in [k - 1]$, that occur in the clauses b'_i and their signs are chosen uniformly and independently of the b_i .

Then Φ^+ and Φ are identically distributed. Therefore, to prove (7.5) we just need to show that

$$(7.6) \quad \mathbb{E}[|\mu_{\Phi^+,\beta}(\{\sigma_{x_1} = 1\}) - \mu_{\Phi^+,\beta,x_1,t}(1)| \mid \mathcal{S}] = o_t(1).$$

Because the formula Φ^- is obtained from Φ by merely deleting a bounded number of at most $(kd)^{2t}$ clauses, by their definition (1.1) the Boltzmann distributions $\mu_{\Phi,\beta}$ and $\mu_{\Phi^+,\beta}$

are mutually $1/\delta$ -contiguous for some $\delta = \delta(t) > 0$. The assumption (1.8) and Lemma 4.3 therefore imply that a.a.s.

$$(7.7) \quad \sum_{i=1}^n |\mu_{\Phi, \beta}(\{\sigma_{x_i} = 1\}) - \mu_{\Phi^-, \beta}(\{\sigma_{x_i} = 1\})| = o(n),$$

$$(7.8) \quad \sum_{i,j=1}^n |\mu_{\Phi^-, \beta}(\{\sigma_{x_i} = \sigma_{x_j} = 1\}) - \mu_{\Phi^-, \beta}(\{\sigma_{x_i} = 1\})\mu_{\Phi^-, \beta}(\{\sigma_{x_j} = 1\})| = o(n^2).$$

In particular, (7.7) ensures together with Corollary 2.2 that the empirical distribution $\pi_{\Phi^-, \beta}$ of the Boltzmann marginals of Φ^- has slim tails a.a.s. Furthermore, Lemma 4.2 implies that a.a.s. over the choice of Φ^- and of the random attachment points y_{ij} of the clauses b'_i in Φ^+ for a sample σ^- from $\mu_{\Phi^-, \beta}$ we have

$$(7.9) \quad \left| \mu_{\Phi^-, \beta}(\{\forall i \in [\ell], j \in [k-1] : \sigma_{y_{ij}}^- = \sigma_{ij}\}) - \prod_{i=1}^{\ell} \prod_{j=1}^{k-1} \mu_{\Phi^-, \beta}(\{\sigma_{y_{ij}}^- = \sigma_{ij}\}) \right| = o(1)$$

for all $\sigma = (\sigma_{ij}) \in \{\pm 1\}^{(k-1)\ell}$. In other words, the joint distribution of the $\sigma_{y_{ij}}^-$ factorises.

Finally, now that we have pinned down the distribution of the boundary condition $(\sigma_{y_{ij}}^-)_{i,j}$, we can easily get a handle on the marginal of x_1 . Namely, let Φ' denote the formula comprising all clauses and variables at distance at most $2t$ from x_1 in Φ^+ . Then (7.9) shows that for a sample σ^+ from $\mu_{\Phi^+, \beta}$ a.a.s.

$$(7.10) \quad \begin{aligned} &\mu_{\Phi^+, \beta}(\{\sigma_{x_1}^+ = s\}) \\ &\propto o(1) + \sum_{\sigma \in \{\pm 1\}^{V(\Phi')}} \mathbb{1}\{\sigma_{x_1} = s\} \mu_{\Phi', \beta}(\sigma) \prod_{i=1}^{\ell} \prod_{j=1}^{k-1} \mu_{\Phi^-, \beta}(\{\sigma_{y_{ij}}^- = \sigma_{y_{ij}}\}) \end{aligned}$$

for all $s = \pm 1$. Since Φ' is acyclic a.a.s. and belief propagation is exact on acyclic factor graphs by Theorem 4.6, (7.10) shows that the Boltzmann marginal $\mu_{\Phi^+, \beta, x_1}(1)$ can be computed by running t iterations of belief propagation, with the boundary messages initialised by the marginals $\mu_{\Phi^-, \beta, y_{ij}}$. Moreover, because the marginals $\mu_{\Phi^-, \beta, y_{ij}}$ are actually independent samples from the slim-tailed distribution $\pi_{\Phi^-, \beta}$ and since the tree Φ' is asymptotically distributed as the Galton-Watson tree T , Proposition 2.3 implies the desired bound (7.5). \square

LEMMA 7.2. *For any $\varepsilon > 0$ there is t_0 such that for $t > t_0$ and for large enough n we have*

$$\mathbb{E} \sum_{i=1}^n \sum_{a \in \partial x_i} |\mu_{\Phi, \beta, a \rightarrow x_i}(1) - \mu_{\Phi, \beta, a \rightarrow x_i, t}(1)| < \varepsilon n.$$

PROOF. While we could repeat a similar argument as in the proof of Lemma 7.1, there is a shorter route that uses the belief propagation recurrence. Specifically, [20], Theorem 1.1, implies together with the assumption (1.8) that a.a.s. for all but $o(n)$ adjacent clause/variable pairs a, x we have

$$(7.11) \quad \mu_{\Phi, \beta, a \rightarrow x}(s) \propto o(1) + \sum_{\sigma \in \{\pm 1\}^{\partial a}} \mathbb{1}\{\sigma_x = s\} \exp(-\beta \mathbb{1}\{\sigma \not\models a\}) \prod_{y \in \partial a \setminus \{x\}} \mu_{\Phi, \beta, y \rightarrow a}(\sigma_y).$$

Since Lemma 7.1 shows that for large enough t a.a.s. we have

$$\mu_{\Phi, \beta, y \rightarrow a}(1) = \mu_{\Phi, \beta, y \rightarrow a, t-1}(1) + o_t(1)$$

for all $y \in \partial a$ and since (7.11) matches (1.5), we conclude that

$$\mu_{\Phi, \beta, a \rightarrow x}(1) = \mu_{\Phi, \beta, a \rightarrow x, t}(1) + o_t(1)$$

a.a.s. \square

PROOF OF PROPOSITION 2.5. The proposition is an immediate consequence of Lemmas 7.1 and 7.2. \square

Finally, for later use we make a note of the following consequence of the arguments presented in this section. We recall the distribution $\pi_{d, \beta}^*$ from Proposition 2.3.

COROLLARY 7.3. *Assume that (1.8) holds, that $d \leq d_{\text{SAT}}(k)$ and that the event \mathfrak{E} that $\pi_{\Phi, \beta}$ has very slim tails satisfies $\limsup_{n \rightarrow \infty} \mathbb{P}[\mathfrak{E}] > 0$. Then along any subsequence where $\mathbb{P}[\mathfrak{E}] > 0$ we have $\mathbb{E}[W_1(\pi_{\Phi, \beta}, \pi_{d, \beta}^*) \mid \mathfrak{E}] = o(1)$.*

PROOF. The same argument as in the proof of Lemma 7.1 shows that $\pi_{\Phi, \beta}$ can be coupled within total variation distance $o(1)$ to coincide with the distribution of a random variable-to-clause message $\mu_{\Phi, \beta, x \rightarrow a}(1)$. Furthermore, Lemma 7.1 implies that this message is well approximated by $\mu_{\Phi, \beta, x \rightarrow a, t}(1)$ for a sufficiently large t a.a.s. Moreover, due to local weak convergence of the factor graph, the depth- $2t$ neighbourhood of x converges weakly in probability to the top $2t$ layers of T . Therefore, the reattachment argument from the proof of Lemma 7.1 implies together with the contraction result from Proposition 2.3 that $\mu_{\Phi, \beta, x \rightarrow a, t}(1)$ converges weakly to $\pi_{d, \beta}^*$ on \mathfrak{E} . \square

8. Replica symmetry breaking. In this section we prove Theorem 1.2 by way of establishing Propositions 2.7 and 2.6; the former is required to derive the latter. Throughout this section we assume that d, β satisfy the assumptions of Theorem 1.2 for suitable sequences $\varepsilon_k, \beta_0(k)$. Let $c > 0$ be such that $d/k = 2^k \log 2 - c$.

8.1. Proof of Proposition 2.7. The proposition asserts a lower bound on $\mathbb{E}[\log Z(\Phi, \beta)]$ under the assumption that the replica symmetry condition (1.8) is satisfied. The starting point for the proof is the following statement, which is implicit in [40]. For the convenience of the reader a self-contained proof is contained in the Appendix.

LEMMA 8.1. *Assume that (1.8) is satisfied. Then*

$$\liminf_{n \rightarrow \infty} \frac{1}{n} \mathbb{E}[\log Z(\Phi, \beta)] \geq \liminf_{n \rightarrow \infty} \mathbb{E}[\mathfrak{B}_{d, \beta}(\pi_{\Phi, \beta})].$$

As Lemma 8.1 provides a lower bound on $\mathbb{E}[\log Z(\Phi, \beta)]$ in terms of the Bethe free energy of the empirical distribution $\pi_{\Phi, \beta}$, the logical next step for us is to get a handle on $\pi_{\Phi, \beta}$.

To this end we are going to harness some of the intermediate results from the second moment calculation from Section 5, particularly Proposition 5.10. Of course the techniques deployed in that section are relatively crude. Instead of dealing with as fine-grained an object as the empirical marginal distribution, the moment calculation has the overlap as its protagonist. More precisely, let

$$\alpha = \alpha(\Phi, \beta) = \langle \alpha(\sigma, \sigma'), \mu_{\Phi, \beta} \rangle$$

denote the average overlap of two independent random samples σ, σ' from the Boltzmann distribution $\mu_{\Phi, \beta}$. In Section 8.1.1 we will derive the following consequence of Proposition 5.10.

LEMMA 8.2. *We have $\alpha \in (1/2 - k^{100}2^{-k/2}, 1/2 + k^{100}2^{-k/2}) \cup (1 - k^22^{-k}, 1)$ a.a.s.*

Furthermore, a simple consequence of (1.8) is that the overlap concentrates about its expectation.

LEMMA 8.3 ([18], Corollary 1.14). *If (1.8) holds, then*

$$\lim_{n \rightarrow \infty} \mathbb{E} \langle |\alpha(\sigma, \sigma') - \alpha|, \mu_{\Phi, \beta} \rangle = 0.$$

Combining Lemmas 8.2 and 8.3, we see that there are two possibilities: either the typical inner product of two Boltzmann samples is close to zero, that is, typical Boltzmann samples are essentially orthogonal. Or the inner product is close to one, in which case σ, σ' largely agree. It is easy to see that in the latter case many of the Boltzmann marginals $\mu_{\Phi, \beta}(\{\sigma_{x_i} = 1\})$ are strongly polarised, that is, $\mu_{\Phi, \beta}(\{\sigma_{x_i} = 1\})$ is fairly close to either zero or one. To be more precise, it is very easy to derive from Lemma 8.3 that if $\alpha \geq 1 - 2^{4-k}$, say, then all but $2^{-0.99k}n$ marginals $\mu_{\Phi, \beta}(\{\sigma_{x_i} = 1\})$ either belong to the interval $(0, 2^{-0.99k})$ or to the interval $(1 - 2^{-0.99k}, 1)$.

But unfortunately this estimate is far too rough to be useful. Indeed, recalling Lemma 8.1, we need to estimate the empirical marginal distribution $\pi_{\Phi, \beta}$ precisely enough to actually estimate $\mathfrak{B}_{d, \beta}(\pi_{\Phi, \beta})$. Due to the $e^{-\beta}$ terms that occur in $\mathfrak{B}_{d, \beta}$, the mere knowledge that most marginals belong to $(0, 2^{-0.99k}) \cup (1 - 2^{-0.99k}, 1)$ does not suffice to calculate the Bethe free energy as even a single unlucky $e^{-\beta}$ term might have a huge impact on the expression (2.15). Yet remarkably, thanks to a delicate expansion argument in Section 8.1.2 we will be able to bootstrap on the rough estimate and derive a much tighter estimate of the Boltzmann marginals. Let \mathfrak{A} be the event that

$$(8.1) \quad \frac{1}{n} \sum_{i=1}^n \mathbb{1} \{ \mu_{\Phi, \beta}(\{\sigma_{x_i} = 1\}) \in (0, \exp(-\beta)) \cup (1 - \exp(-\beta), 1) \} \geq 1 - 2^{-0.98k}.$$

LEMMA 8.4. *If (1.8) holds then $\mathbb{P}[\{\alpha \geq 1 - k^22^{-k}\} \setminus \mathfrak{A}] = o(1)$.*

Thus, combining Lemmas 8.2 and 8.4, we learn that unless α is close to $1/2$, most Boltzmann marginals are actually extremely polarised. This polarisation is strong enough for us to derive the following explicit lower bound on the Bethe free energy, whose proof can be found in Section 8.1.5.

LEMMA 8.5. *On the event \mathfrak{A} we have $\mathfrak{B}_{d, \beta}(\pi_{\Phi, \beta}) \geq 2^{-k}(c - \log 2/2 + o(1))$.*

Hence, we are left to deal with the scenario that α is close to $1/2$, that is, $\alpha \in (1/2 - k^{100}2^{-k/2}, 1/2 + k^{100}2^{-k/2})$ as in Lemma 8.2. In this case it is not difficult to verify that the empirical marginal distribution $\pi_{\Phi, \beta}$ has very slim tails. Consequently, Corollary 7.3 shows that $\pi_{\Phi, \beta}$ is close to the distribution $\pi_{d, \beta}^*$ from Proposition 2.3. Further, in Section 8.1.6 we will be able to derive the following estimate of the latter distribution’s Bethe free energy.

LEMMA 8.6. *We have $\mathfrak{B}(\pi_{d, \beta}^*) = 2^{-k}(c - \log 2/2) + o(2^{-k})$.*

With these ingredients we can now deduce Proposition 2.7.

PROOF OF PROPOSITION 2.7. Assume that (1.8) holds a.a.s. Let \mathfrak{E} be the event that $\pi_{\Phi, \beta}$ has very slim tails. Then Claim 5.35 and Lemma 8.2 imply that $\mathbb{P}[\mathfrak{E} \cup \mathfrak{A}] = 1 - o(1)$. There are three cases to consider.

Case 1: $\mathbb{P}[\mathfrak{E}] = o(1)$ then the assertion follows from Lemma 8.5 and Lemma 8.6 in combination with Lemma 8.1.

Case 2: $\mathbb{P}[\mathfrak{A}] = o(1)$ the assertion follows from Corollary 7.3, Lemma 8.6 and Lemma 8.1.

Case 3: neither $\mathbb{P}[\mathfrak{A}] = o(1)$ nor $\mathbb{P}[\mathfrak{E}] = o(1)$ in this case we combine Corollary 7.3, Lemma 8.6, Lemma 8.5 and Lemma 8.1.

Thus, in any case we obtain the desired lower bound on $\mathbb{E}[\log Z(\Phi, \beta)]$. \square

8.1.1. *Proof of Lemma 8.2.* We use the techniques and results from Section 5.3.3 to estimate $\mathbb{E}[\mu_{\Phi, \beta}(\{\alpha(\sigma, \sigma') \notin \mathcal{A}\})]$ with $\mathcal{A} = (1/2 - k^{100}2^{-k/2}, 1/2 + k^{100}2^{-k/2}) \cup (1 - k^{200}2^{-k}, 1)$. Recalling the function $f(\alpha)$ from (2.6), we see that

$$(8.2) \quad \frac{1}{n} \log \mathbb{E}[Z(\Phi, \beta)^2 \mu_{\Phi, \beta}(\{\alpha(\sigma, \sigma') \notin \mathcal{A}\})] \leq \max_{\alpha \in \mathcal{A}} f(\alpha) + o(1).$$

Hence, because Theorem 4.1 implies that $Z(\Phi, \beta) \geq 1$ a.a.s., it suffices to show that $f(\alpha) < 0$ for all $\alpha \notin \mathcal{A}$. Indeed, Claims 5.30–5.33 reduce our task to proving that $f(1/2 + k^{100}2^{-k/2}) < 0$ and $f(1 - 2^{4-k}) < 0$. Applying Taylor’s formula, we obtain $f(1/2 + k^{100}2^{-k/2}) \leq f(1/2) - k^{200}2^{-k} < 0$ and

$$\begin{aligned} f(1 - 2^{4-k}) &\leq \log 2 + k^3 2^{-k} \log(2) - k^2 2^{-k} \log(k) \\ &\quad + \frac{d}{k} \log(1 - (1 - e^{-\beta})2^{1-k} + (1 - e^{-\beta})^2 2^{-k} (1 - k^2 2^{-k})^k) \\ &\leq -k^2 2^{-k} \log(k) + \left(\frac{d}{k} - 2^k \log 2\right) (1 - e^{-\beta})2^{-k} + O(2^{-k}) < 0. \end{aligned}$$

Thus, the assertion follows from (8.2).

8.1.2. *Proof of Lemma 8.4.* We seize upon the expansion properties of the hypergraph underlying the random formula Φ . To set up the necessary terminology let Φ be any k -CNF formula on the variable set $V_n = \{x_1, \dots, x_n\}$ and let $\sigma \in \{\pm 1\}^{V_n}$ be a truth assignment. We say that a variable x_i supports a clause a of Φ under σ if $x_i \in \partial a$, $\text{sign}(x_i, a) = \sigma_{x_i}$ and $\text{sign}(x, a) \neq \sigma_x$ for all $x \in \partial a \setminus \{x_i\}$. Hence, x_i contributes the single true literal of a . Let $\text{supp}_{\Phi, \sigma}(x_i)$ be the set of all clauses that x_i supports. Further, call a set $S \subset V_n$ of variables stable in (Φ, σ) if:

ST1 every $x \in S$ supports at least $10^{-5}k$ clauses that contain variables from S only, and

ST2 no $x \in S$ appears in more than $10^{-6}k$ clauses a that fail to contain a variable $y \in S$ with $\text{sign}(y, a) = \sigma_y$.

Since the union of two stable sets is stable, we denote by $S(\Phi, \sigma)$ the largest stable set of (Φ, σ) . The following lemma, whose proof we defer to Section 8.1.3, asserts that an assignment σ drawn from the Boltzmann distribution of the random formula Φ induces a very large stable set a.a.s.

LEMMA 8.7. We have $\mathbb{E}[\mu_{\Phi, \beta}(\{|S(\Phi, \sigma)| \geq (1 - 2^{-0.99k})n\})] \sim 1$.

In light of Lemma 8.7 we call a formula Φ normal if

$$\langle \mathbb{1}\{|S(\Phi, \sigma)| \geq (1 - 2^{-0.99k})n\}, \mu_{\Phi, \beta} \rangle \sim 1,$$

that is, if its typical Boltzmann samples induce stable sets as large as promised by Lemma 8.7. Furthermore, we call Φ separable if

$$\mu_{\Phi, \beta} \left(\left\{ 1 - k^3 2^{-k} \leq \alpha(\sigma, \sigma') \wedge \sum_{y \in S(\Phi, \sigma)} \mathbb{1}\{\sigma_y \neq \sigma'_y\} > ne^{-10\beta} \right\} \right) = o(1).$$

Hence, it is unlikely that σ, σ' have a high overlap but differ on a lot of variables from $S(\Phi, \sigma)$. In Section 8.1.4 we are going to prove that Φ is separable a.a.s.

LEMMA 8.8. *The random formula Φ is separable a.a.s.*

Lemma 8.4 is now an easy consequence of Lemmas 8.7 and 8.8.

PROOF OF LEMMA 8.4. Thanks to Lemmas 8.7 and 8.8 we may assume that Φ is normal and separable. We also assume that the event $\{\alpha \geq 1 - k^2 2^{-k}\}$ occurs and that the replica symmetry condition (1.8) holds, that is,

$$(8.3) \quad \frac{1}{n^2} \sum_{i,j=1}^n d_{\text{TV}}(\mu_{\Phi,\beta,x_i,x_j}, \mu_{\Phi,\beta,x_i} \otimes \mu_{\Phi,\beta,x_j}) = o(1).$$

Draw a random σ from $\mu_{\Phi,\beta}$ and let $\mathcal{V} = \mathcal{V}(\sigma)$ be the set of all variables $x \in S(\Phi, \sigma)$ such that $\mu_{\Phi,\beta,x}(1) \in (0, e^{-\beta}) \cup (1 - e^{-\beta}, 1)$. We claim that $|\mathcal{V}| \geq (1 - e^{-\beta})|S(\Phi, \sigma)|$ a.a.s. over the choice of σ , which would clearly imply the lemma. Indeed, let

$$\mathfrak{W} = \sum_{i=1}^n \mathbb{1}\{\mu_{\Phi,\beta}(\{\sigma_{x_i} = 1\}) \in (0, \exp(-\beta)) \cup (1 - \exp(-\beta), 1)\}.$$

Because of our assumption that the event $\{\alpha \geq 1 - k^2 2^{-k}\}$ occurs it remains to prove that \mathfrak{W} occurs a.a.s. Lemma 8.7 shows that $|S(\Phi, \sigma)| \geq (1 - 2^{-0.99k})n$ a.a.s. Therefore, we obtain

$$\mathfrak{W} \geq |\mathcal{V}| \geq (1 - e^{-\beta})|S(\Phi, \sigma)| \geq (1 - e^{-\beta})(1 - 2^{-0.99k})n \geq (1 - 2^{-0.98k})n,$$

for β large enough a.a.s. To verify this claim assume that $|\mathcal{V}| < (1 - e^{-\beta})|S(\Phi, \sigma)|$ and draw a second, independent sample σ' from $\mu_{\Phi,\beta}$. Let X be the number of variables $x \in \mathcal{V}$ such that $\sigma_x \neq \sigma'_x$. Then the asymptotic pairwise independence property (8.3) implies together with Chebyshev's inequality that a.a.s. over the choice of σ, σ' we have

$$\begin{aligned} X &\geq |\mathcal{V}| - \sum_{x \in \mathcal{V}} (\mu_{\Phi,\beta,x}(1)^2 + \mu_{\Phi,\beta,x}(-1)^2) + o(n) \\ &= 2 \sum_{x \in \mathcal{V}} \mu_{\Phi,\beta,x}(1)\mu_{\Phi,\beta,x}(-1) + o(n) \geq \exp(-2\beta)n + o(n), \end{aligned}$$

in contradiction to separability. \square

8.1.3. *Proof of Lemma 8.7.* We prove the lemma by way of a distribution on random k -CNF formulas known as the *planted model*. Recall that $m = \text{Po}(dn/k)$ is a Poisson variable and consider the following experiment.

PL1 draw a truth assignment $\sigma^* \in \{\pm 1\}^V$ uniformly at random

PL2 then draw a k -CNF $\Phi^* = \Phi^*(\sigma^*)$ with $m \sim \text{Po}(dn/k)$ clauses from the distribution

$$\mathbb{P}[\Phi^* = \Phi \mid m, \sigma^*] = \frac{\mathbb{P}[\Phi = \Phi \mid m] \exp(-\beta \mathcal{H}_\Phi(\sigma^*))}{(1 - (1 - e^{-\beta})2^{-k})^m}.$$

The planted model (Φ^*, σ^*) is a tried and tested device for studying the Boltzmann distribution of random formulas [2]. Indeed, while it is difficult to tackle the Boltzmann distribution directly, the planted model is amenable to the toolbox of probabilistic combinatorics thanks to its constructive definition PL1–PL2. The following statement ties the two models together.

LEMMA 8.9. *Let \mathcal{E} be a set of formula/assignment pairs. Then*

$$\mathbb{E}[\mathbb{1}\{(\Phi, \sigma) \in \mathcal{E}\}, \mu_{\Phi, \beta} \mid \mathbf{m}] \leq \mathbb{E}[Z(\Phi, \beta) \mid \mathbf{m}] \mathbb{P}[(\Phi^*, \sigma^*) \in \mathcal{E} \mid \mathbf{m}] + o(1).$$

PROOF. Let $\mathcal{E}' = \mathcal{E} \cap \{Z(\Phi, \beta) \geq 1\}$. Then Theorem 4.1 implies that

$$(8.4) \quad \mathbb{E}[\mathbb{1}\{(\Phi, \sigma) \in \mathcal{E}\}, \mu_{\Phi, \beta} \mid \mathbf{m}] = \mathbb{E}[\mathbb{1}\{(\Phi, \sigma) \in \mathcal{E}'\}, \mu_{\Phi, \beta} \mid \mathbf{m}] + o(1).$$

Furthermore, the definition PL1–PL2 of the planted model ensures that

$$(8.5) \quad \begin{aligned} & \mathbb{E}[\mathbb{1}\{(\Phi, \sigma) \in \mathcal{E}'\}, \mu_{\Phi, \beta} \mid \mathbf{m}] \\ &= \sum_{(\Phi, \sigma) \in \mathcal{E}'} \mathbb{P}[\Phi = \Phi \mid \mathbf{m}] \mu_{\Phi, \beta}(\sigma) \\ &= \sum_{(\Phi, \sigma) \in \mathcal{E}'} \mathbb{P}[\Phi = \Phi \mid \mathbf{m}] e^{-\beta \mathcal{H}_{\Phi}(\sigma)} / Z(\Phi, \beta) \\ &\leq \sum_{(\Phi, \sigma) \in \mathcal{E}} \mathbb{P}[\Phi = \Phi \mid \mathbf{m}] e^{-\beta \mathcal{H}_{\Phi}(\sigma)} \\ &\leq 2^n (1 - (1 - e^{-\beta}) 2^{-k})^m \sum_{(\Phi, \sigma) \in \mathcal{E}} \mathbb{P}[(\Phi^*, \sigma^*) = (\Phi, \sigma) \mid \mathbf{m}] \\ &= \mathbb{E}[Z(\Phi, \beta) \mid \mathbf{m}] \mathbb{P}[(\Phi^*, \sigma^*) \in \mathcal{E} \mid \mathbf{m}]. \end{aligned}$$

The assertion follows from (8.4) and (8.5). \square

To facilitate the use of the planted model we make a note of the following easy upper bound.

LEMMA 8.10. *We have $2(1 - (1 - e^{-\beta}) 2^{-k})^{d/k} \leq \exp(2^{-k-1})$.*

PROOF. Using the bound $d/k \geq 2^k \log 2 - 3 \log(2)/2$ we obtain in the limit of large β ,

$$\limsup_{\beta \rightarrow \infty} \log 2 + \frac{d}{k} \log(1 - (1 - e^{-\beta}) 2^{-k}) \leq \log 2 - \frac{d}{k} [2^{-k} + 2^{-2k-1}] \leq \frac{\log 2}{2^{k+1}},$$

as desired. \square

As a final preparation we reformulate the second part PL2 of the experiment above as follows.

PL2a for each of the \mathbf{m} clauses a_1, \dots, a_m of Φ^* draw the k -tuple of variables that occur in the clause uniformly and independently.

PL2b subsequently, once more independently for each $i \in [\mathbf{m}]$, draw the signs with which the variables appear in clause a_i such that $\mathbb{P}[\sigma^* \not\models a_i \mid \mathbf{m}, \sigma^*] = e^{-\beta} / (2^k - 1 + e^{-\beta})$.

The distributions produced by PL2 and PL2a–PL2b coincide because the clauses of Φ are mutually independent.

We proceed to exhibit a large stable set. The following lemma shows that in (Φ^*, σ^*) most variables support a good number of clauses. To be precise, for a variable x let s_x be the number of clauses a of Φ^* to which x contributes the only literal that is satisfied under σ^* .

LEMMA 8.11. *We have $\mathbb{P}[\sum_{x \in V_n} \mathbb{1}\{s_x < 10^{-4}k\} > 2^{-0.997k} n] < \exp(-n/2^k)$.*

PROOF. Because the total number of clauses of Φ^* is Poisson, the random variables $(s_x)_{x \in V_n}$ are mutually independent Poissons. Moreover, PL2b shows that $\mathbb{E}[s_x] \sim d/(2^k - 1 + e^{-\beta}) = k \log(2) + O(2^{-k})$ for every x . Therefore, Bennett’s inequality from Lemma 4.5 yields

$$(8.6) \quad \mathbb{P}[s_x < 10^{-4}k] \leq 2^{-0.998k}.$$

Furthermore, due to the independence of the s_x the sum $\sum_{x \in V_n} \mathbb{1}\{s_x < 10^{-4}k\}$ is a binomial variable. Since (8.6) shows that its mean is bounded by $n2^{-0.998k}$, the assertion follows from the Chernoff bound. \square

For a variable x let u_x be the number of clauses of Φ^* in which x occurs and that σ^* fails to satisfy.

LEMMA 8.12. *We have $\mathbb{P}[\sum_{x \in V_n} \mathbb{1}\{u_x > 0\} > 2^{-0.997k}n] < \exp(-n/2^k)$.*

PROOF. Let m_0 be the total number of clauses of Φ^* that σ^* fails to satisfy. If $\sum_{x \in V_n} \mathbb{1}\{u_x > 0\} > 2^{-0.997k}n$, then $m_0 \geq 2^{-0.997k}n/k$. But PL2b ensures that $m_0 \sim \text{Po}(dne^{-\beta}/(k(2^k - 1 + e^{-\beta})))$. Therefore, by Bennett’s inequality,

$$\begin{aligned} \mathbb{P}\left[\sum_{x \in V_n} \mathbb{1}\{u_x > 0\} > 2^{-0.997k}n\right] &\leq \mathbb{P}[\text{Po}(dne^{-\beta}/(k(2^k - 1 + e^{-\beta}))) > 2^{-0.997k}n/k] \\ &\leq \exp(-n/2^k), \end{aligned}$$

providing that β is sufficiently large, as claimed. \square

The following lemma shows that the planted model possesses a large stable set with very high probability.

LEMMA 8.13. *On the event $m \sim dn/k$ we have*

$$\mathbb{P}[|S(\Phi^*, \sigma^*)| < (1 - 2^{-0.99k})n \mid m] \leq 4 \exp(-n/2^k).$$

PROOF. Due to symmetry we may condition on the event $\sigma_x^* = 1$ for all variables x . Starting from the set S_0 of all variables x such that $s_x \geq 10^{-4}k$ and $u_x = 0$, we attempt to construct a large stable set. To this end, we iteratively obtain S_{i+1} from S_i by removing an arbitrary variable $y \in S_i$ that violates one of the conditions ST1–ST2. Hence, either y supports fewer than $10^{-5}k$ clauses comprising variables from S_i only, or y appears negatively in more than $10^{-6}k$ clauses that contain at least one positive literal but whose positive literals stem from $V_n \setminus S_i$ only. Of course, once no such variable y is left the process stops. Let T be the stopping time of the process. By Lemmas 8.11 and 8.12 we may assume that $|S_0| \leq (1 - 2^{-0.997k+1})n$. Moreover, by construction the final set S_T is stable and has size at least $|S_0| - T$. Therefore, we just need to bound the probability of the event $\{T > 2^{-0.997k+1}n\}$.

Hence, let $t = \lfloor 2^{-0.997k+1}n \rfloor$, set $\theta = t/n$ and let $R = V_n \setminus S_t$. Then R contains the set $S_0 \setminus S_t$ of variables that our process removes by time t as well as the variables $V_n \setminus S_0$ that were excluded from the beginning. Since $|S_0| \leq (1 - 2^{-0.997k+1})n$ and $t \leq 2^{-0.997k+1}n$ we have

$$(8.7) \quad |R| \geq 2t.$$

Further, let X be the number of clauses that are supported by a variable from R and contain a second variable from R . Also let \mathcal{C} be the set of clauses that contain at least one variable

from R positively and at least one variable from R negatively but none from $S_t = V_n \setminus R$ positively. Moreover, let Y be the number of R - \mathcal{C} -edges in $G(\Phi)$. By construction, if $T > t$ then either $X > 10^{-7}kt$ or $Y > 10^{-7}kt$. Thus, letting

$$\mathcal{E} = \{\mathbf{m} \sim dn/k, |S_0| \leq (1 - 2^{-0.997k+1})n\},$$

we have

$$(8.8) \quad \mathbb{P}[T > 2^{-0.997k+1}n \mid \mathcal{E}] \leq \mathbb{P}[X > 10^{-7}ktn \mid \mathcal{E}] + \mathbb{P}[Y > 10^{-7}kt \mid \mathcal{E}].$$

In light of (8.7), to bound the first probability $\mathbb{P}[X > 10^{-7}ktn \mid \mathcal{E}]$ we estimate the probability that there exists a set $\mathcal{R} \subset V_n$ of size $|\mathcal{R}| = 2t$ such that the number $\mathcal{X}_{\mathcal{R}}$ of clauses supported by a variable from \mathcal{R} that contain a second variable from \mathcal{R} exceeds $\ell = 10^{-7}kt$. Also let $\mathcal{X} = \mathcal{X}_{\{x_1, \dots, x_{2t}\}}$. Since $\mathbb{P}[\mathcal{E}] \sim 1$, we obtain the upper bound

$$(8.9) \quad \begin{aligned} \mathbb{P}[X > 10^{-7}ktn \mid \mathcal{E}] &\leq 2\mathbb{P}[\exists \mathcal{R} : \mathcal{X}_{\mathcal{R}} > \ell \mid \mathbf{m} \sim dn/k] \\ &\leq 2 \binom{n}{2t} \mathbb{P}[\mathcal{X} > \ell \mid \mathbf{m} \sim dn/k]. \end{aligned}$$

Furthermore, the last probability is easy to estimate. Indeed, due to PL2b the probability that a single clause is supported by a variable from \mathcal{R} and features a second variable from \mathcal{R} negatively is bounded by

$$p = \frac{4k(k-1)\theta^2}{2^k - 1 + e^{-\beta}}.$$

Consequently, since the \mathbf{m} clauses are drawn independently, \mathcal{X} is stochastically dominated by a binomial variable $\text{Bin}(\mathbf{m}, p)$. Combining (8.9) with the Chernoff bound, we therefore obtain

$$(8.10) \quad \mathbb{P}[X > 10^{-7}ktn \mid \mathcal{E}] \leq 2 \left(\frac{en}{2t}\right)^{2t} \mathbb{P}[\text{Bin}(2dn/k, p) > \ell] \leq \exp(-n/2^k).$$

Moving on to Y , we consider an arbitrary set \mathcal{R} as above, define $\mathcal{C}_{\mathcal{R}}$ as above as the set of clauses that contain at least two variables from \mathcal{R} but in which no variable from $V_n \setminus \mathcal{R}$ occurs positively and let $\mathcal{Y}_{\mathcal{R}}$ be the number of \mathcal{R} - $\mathcal{C}_{\mathcal{R}}$ -edges in $G(\Phi)$. Observe that $\mathcal{Y}_{\mathcal{R}}/k \leq \mathcal{C}_{\mathcal{R}} \leq \mathcal{Y}_{\mathcal{R}}/2$. Thanks to symmetry it suffices to consider $\mathcal{Y} = \mathcal{Y}_{\{x_1, \dots, x_{2t}\}}$ and we obtain

$$(8.11) \quad \begin{aligned} \mathbb{P}[Y > 10^{-7}tn \mid \mathcal{E}] &\leq 2\mathbb{P}[\exists \mathcal{R} : \mathcal{Y}_{\mathcal{R}} > \ell \mid \mathbf{m} \sim dn/k] \\ &\leq 2 \binom{n}{2t} \mathbb{P}[\mathcal{Y} > \ell \mid \mathbf{m} \sim dn/k]. \end{aligned}$$

We bound the last probability by

$$(8.12) \quad \begin{aligned} \mathbb{P}[\mathcal{Y} > \ell \mid \mathbf{m} \sim dn/k] &\leq \sum_{\ell/k \leq M \leq \ell/2} \binom{2dn/k}{M} \binom{kM}{\ell} 2^{-kM} (2\theta)^\ell \\ &\leq 2 \binom{2dn/k}{\ell/2} \binom{k\ell/2}{\ell} 2^{-k\ell/2} (2\theta)^\ell \leq (10^{10}k\theta)^{\ell/2}. \end{aligned}$$

Finally, the assertion follows from (8.8), (8.10) and (8.12). \square

PROOF OF LEMMA 8.7. The assertion is an immediate consequence of Lemma 8.9, Lemma 8.10 and Lemma 8.13. \square

8.1.4. *Proof of Lemma 8.8.* We treat two regimes of distances $\sum_{x \in S(\Phi, \sigma)} \mathbb{1}\{\sigma_x \neq \sigma'_x\}$ separately. Let us begin with very small distances.

LEMMA 8.14. *A.a.s. the random formula Φ has the following property. For any set $\mathcal{V} \subset V_n$ of size $|\mathcal{V}| \leq 10^{-9}k^{-1}2^{-k}n$ the number of clauses in which at least two variables from \mathcal{V} occur is bounded above by $10^{-7}k|\mathcal{V}|$.*

PROOF. Fix any such set \mathcal{V} and let $v = |\mathcal{V}|/n$ and $\lambda = 10^{-7}k$. Clearly, we may condition on the event that $m \sim dn/k \leq 2^k n$. Because the clauses are drawn independently, given m the number $X_{\mathcal{V}}$ of clauses that contain two variables from \mathcal{V} is stochastically dominated by a $\text{Bin}(m, k^2v^2)$ variable. Therefore,

$$(8.13) \quad \mathbb{P}[X_{\mathcal{V}} > \lambda vn \mid m] \leq \binom{m}{\lambda vn} (k^2v^2)^{\lambda vn} \leq \left(\frac{e2^k k^2 v}{\lambda}\right)^{\lambda vn}.$$

Thanks to the assumption on v , combining (8.13) with a union bound on sets \mathcal{V} completes the proof. \square

COROLLARY 8.15. *A.a.s. we have*

$$\mu_{\Phi, \beta} \left(\left\{ 10^{-9}k^{-1}2^{-k}n > \sum_{x \in S(\Phi, \sigma)} \mathbb{1}\{\sigma_x \neq \sigma'_x\} > ne^{-10\beta} \right\} \right) = o(1).$$

PROOF. We may condition on Φ possessing the property quoted in Lemma 8.14. For a pair of assignments σ, σ' let σ'' be the assignment $\sigma''_x = \sigma_x$ for all $x \in S(\Phi, \sigma)$ and $\sigma''_x = \sigma'_x$ for all $x \notin S(\Phi, \sigma)$. Also let \mathcal{V} be the set of all variables $x \in S(\Phi, \sigma)$ such that $\sigma_x \neq \sigma'_x$. We claim that

$$(8.14) \quad \mathcal{H}_{\Phi}(\sigma'') \leq \mathcal{H}_{\Phi}(\sigma') - 10^{-7}k|\mathcal{V}|.$$

Indeed, by ST1 every $x \in S(\Phi, \sigma)$ supports at least $10^{-5}k$ clauses. If $\sigma'_x \neq \sigma_x$, then σ' can only satisfy those clauses supported by x that contain a second variable from \mathcal{V} . But Lemma 8.14 shows that there are no more than $10^{-7}k|\mathcal{V}|$ such clauses. Furthermore, ST2 ensures that there are no more than $10^{-6}k|\mathcal{V}|$ clauses that σ' satisfies and that σ'' fails to satisfy. Thus, we obtain (8.14). Finally, (8.14) implies

$$\begin{aligned} & \sum_{e^{-10\beta}n < t < 10^{-9}k^{-1}2^{-k}n} \mu_{\Phi, \beta} \left(\left\{ \sum_{x \in S(\Phi, \sigma)} \mathbb{1}\{\sigma_x \neq \sigma'_x\} = t \right\} \right) \\ & \leq \sum_{e^{-10\beta}n < t} \binom{n}{t} \exp(-10^{-7}\beta kt) = o(1), \end{aligned}$$

as desired. \square

We proceed to assignment pairs that differ on an intermediate number of variables from the stable set.

LEMMA 8.16. *The pair (Φ^*, σ^*) has the following property with probability at least $1 - \exp(-n/2^k)$. Let σ be any assignment such that*

$$t = \sum_{x \in S(\Phi^*, \sigma^*)} \mathbb{1}\{\sigma_x \neq \sigma^*_x\} \in (10^{-9}k^{-1}2^{-k}n, k^{-4}n).$$

Furthermore, let $\sigma'_x = \sigma^*_x$ for all $x \in S(\Phi^*, \sigma^*)$ and set $\sigma'_x = \sigma_x$ for all $x \notin S(\Phi^*, \sigma^*)$. Then $\mathcal{H}_{\Phi^*}(\sigma) \geq \mathcal{H}_{\Phi^*}(\sigma') + 10^{-6}kt$.

PROOF. Let $\lambda = 10^{-7}k$. We pursue a similar strategy as in the previous lemma, but this time we confine ourselves to the clauses supported by variables from $S(\Phi^*, \sigma^*)$. Indeed, by ST1 every variable $x \in S(\Phi^*, \sigma^*)$ supports at least $10^{-5}k$ clauses. Hence, if $\sigma_x \neq \sigma_x^*$, then for σ to satisfy such a clause, it must contain another variable y such that $\sigma_y \neq \sigma_y^*$ negatively.

Consequently, to prove the assertion it suffices to show that Φ^* has the following property a.a.s. Let Φ^* be the sub-formula obtained by retaining only those clauses that contain a single true literal under σ^* . Then the probability that there exists a set \mathcal{V} of variables of size $10^{-9}k^{-1}2^{-k} \leq |\mathcal{V}|/n \leq k^{-4}$ such that the number $X_{\mathcal{V}}$ of clauses of Φ^* that contain one variables from \mathcal{V} negatively and one positively exceeds $\lambda|\mathcal{V}|$ is upper bounded by $\exp(-n/2^k)$.

Thus, fix a set \mathcal{V} as above and let $v = |\mathcal{V}|/n$. The number of clauses of Φ^* that contain one variable from \mathcal{V} positively and one negatively is stochastically dominated by $\text{Po}(pdn/k)$ with $p = k^2v^2/(2^k - 1 + e^{-\beta})$. Therefore, Bennett’s inequality shows that

$$(8.15) \quad \mathbb{P}[X_{\mathcal{V}} > \lambda vn] \leq \exp(10^{-8}kv \log(dp/(k\lambda v))).$$

Combining (8.15) with a union bound on sets \mathcal{V} completes the proof. \square

COROLLARY 8.17. *A.a.s. we have $\mu_{\Phi, \beta}(\{k^{-4}n > \sum_{x \in S(\Phi, \sigma)} \mathbb{1}\{\sigma_x \neq \sigma'_x\}\} \geq nk^{-1/2}2^{-kn}) = o(1)$.*

PROOF. Invoking Lemmas 8.9 and 8.10, we extend the statement of Lemma 8.16 from the planted model (Φ^*, σ^*) to the random pair (Φ, σ) . Then we follow the steps of the proof of Corollary 8.15. \square

PROOF OF LEMMA 8.8. The assertion follows from Corollaries 8.15 and 8.17. \square

8.1.5. *Proof of Lemma 8.5.* Let μ be a random variable with distribution $\pi_{\Phi, \beta}$ and let $J = \pm 1$ be an independent random variable with $\mathbb{E}[J] = 0$. Moreover, let $\pi'_{\Phi, \beta}$ be the distribution of $(1 + J(2\mu - 1))/2$. Further, let $(\mu_{i,j})_{i,j}$ be a family of independent samples from $\pi'_{\Phi, \beta}$ and let γ^{\pm} be two independent $\text{Po}(d/2)$ variables. We are going to estimate the two contributions to the Bethe free energy separately. The following claim deals with the second part.

CLAIM 8.18. *We have $-\frac{d(k-1)}{k}\mathbb{E}(\log 1 - (1 - e^{-\beta}) \prod_{j=1}^k \mu_{1,j}) \geq \frac{d(k-1+o_k(1))}{k^2}\beta + o_{\beta}(1)$.*

PROOF. Let \mathcal{A} be the event that $\mu_{1,j} \geq 1 - e^{-\beta}$ for all $j \in [k]$ and let $\bar{\mathcal{A}}$ be the complement of \mathcal{A} . Then (8.1) implies together with the fact that $\mu_{i,j}$ and $1 - \mu_{i,j}$ are identically distributed that

$$(8.16) \quad \mathbb{P}[\mathcal{A}] \geq 2^{-k} + O(2^{-1.9k}).$$

Further, we have

$$(8.17) \quad \begin{aligned} 1 - (1 - e^{-\beta}) \prod_{j=1}^k \mu_{1,j} &\leq 1 - (1 - e^{-\beta})^{k+1} \leq (k + 1)e^{-\beta} \quad \text{on } \mathcal{A}, \\ 1 - (1 - e^{-\beta}) \prod_{j=1}^k \mu_{1,j} &\leq 1 \quad \text{on } \bar{\mathcal{A}}. \end{aligned}$$

Combining (8.16) and (8.17), we obtain the assertion. \square

Let

$$\Pi^+ = \prod_{i=1}^{\gamma^+} \left(1 - (1 - e^{-\beta}) \prod_{j=1}^{k-1} \mu_{i,j} \right), \quad \Pi^- = \prod_{i=1}^{\gamma^-} \left(1 - (1 - e^{-\beta}) \prod_{j=1}^{k-1} \mu_{i+\gamma^+,j} \right).$$

CLAIM 8.19. We have $\mathbb{E} \log [\Pi^+ + \Pi^-] \geq \beta(-d2^{-k} + \Omega(\sqrt{k})) + o_\beta(\beta)$.

PROOF. As in the previous proof we are going to separate the clause terms $\prod_{j=1}^{k-1} \mu_{i,j}$ with all $\mu_{i,j}$ close to one from the rest. Specifically, let \mathbf{g}_1^+ be the number of indices $i \leq \gamma^+$ such that $\mu_{i,j} \geq 1 - e^{-\beta}$ for all $j \in [k - 1]$. Moreover, let \mathbf{g}_0^+ be the number $i \leq \gamma^+$ such that $\mu_{i,j} \leq e^{-\beta}$ for some $j \in [k - 1]$ and let $\mathbf{g}_*^+ = \gamma^+ - \mathbf{g}_1^+ - \mathbf{g}_0^+$. Also define $\mathbf{g}_0^-, \mathbf{g}_1^-, \mathbf{g}_*^-$ analogously for the second summand and let $\mathbf{g}_* = \mathbf{g}_*^+ + \mathbf{g}_*^-, \mathbf{g}_0 = \mathbf{g}_0^+ + \mathbf{g}_0^-$. Then we obtain the lower bounds

$$\Pi^\pm \geq \exp(-\beta(\mathbf{g}_1^\pm + \mathbf{g}_*^\pm) - \mathbf{g}_0^\pm).$$

Hence,

$$(8.18) \quad \log(\Pi^+ + \Pi^-) \geq \log(\Pi^+ \vee \Pi^-) \geq -\beta(\mathbf{g}_1^+ \wedge \mathbf{g}_1^- + \mathbf{g}_*) - \mathbf{g}_0.$$

Recalling (8.1) and using Poisson thinning, we can view $\mathbf{g}_1^{\pm 1}, \mathbf{g}_*$ and \mathbf{g}_0 as independent Poissons with means

$$(8.19) \quad \lambda_1^+ = \lambda_1^- \leq \frac{d}{2^k + 2^{-1.9k}}, \quad \lambda_* \leq d2^{-1.9k}, \quad \lambda_0 \leq d.$$

Additionally, invoking the normal approximation to the Poisson distribution, we obtain

$$(8.20) \quad \mathbb{E}[\mathbf{g}_1^+ \wedge \mathbf{g}_1^-] \leq \lambda_1^+ - \Omega(\sqrt{\lambda_1^+}) = 2^{-k}d - \Omega(\sqrt{k}).$$

Finally, combining (8.18)–(8.20) we obtain the assertion. \square

PROOF OF LEMMA 8.5. The lemma is an immediate consequence of Claims 8.18 and 8.19. \square

8.1.6. *Proof of Lemma 8.6.* Let $(\mu_{i,j})_{i,j \geq 1}$ be a sequence of independent samples from $\pi_{d,\beta}^*$ and let γ^\pm be two independent Poisson variables with mean $d/2$. Then

$$\begin{aligned} \mathcal{B}(\pi_{d,\beta}^*) &= \mathbb{E} \left[\log \left(\prod_{i=1}^{\gamma^+} \left(1 - (1 - e^{-\beta}) \prod_{j=1}^{k-1} \mu_{i,j} \right) + \prod_{i=1}^{\gamma^-} \left(1 - (1 - e^{-\beta}) \prod_{j=1}^{k-1} \mu_{\gamma^++i,j} \right) \right) \right] \\ &\quad - \frac{d(k-1)}{k} \mathbb{E} \left[\log \left(1 - (1 - e^{-\beta}) \prod_{j=1}^k \mu_{1,j} \right) \right]. \end{aligned}$$

Hence, for large enough β we obtain

$$(8.21) \quad \begin{aligned} \mathcal{B}(\pi_{d,\beta}^*) &= \mathbb{E} \left[\log \left(\prod_{i=1}^{\gamma^+} \left(1 - \prod_{j=1}^{k-1} \mu_{i,j} \right) + \prod_{i=1}^{\gamma^-} \left(1 - \prod_{j=1}^{k-1} \mu_{\gamma^++i,j} \right) \right) \right] \\ &\quad - \frac{d(k-1)}{k} \mathbb{E} \left[\log \left(1 - \prod_{j=1}^k \mu_{1,j} \right) \right] + o(2^{-k}). \end{aligned}$$

We expand the two terms on the right-hand side separately. Due to the independence of the $\mu_{1,j}$ and (6.2) we obtain

$$\begin{aligned} & \frac{d(k-1)}{k} \mathbb{E} \left[\log \left(1 - \prod_{j=1}^k \mu_{1,j} \right) \right] \\ &= -\frac{d(k-1)}{k} \left[\mathbb{E} \left(\prod_{j=1}^k \mu_{1,j} \right) + \frac{1}{2} \mathbb{E} \left(\prod_{j=1}^k \mu_{1,j}^2 \right) + o \left(\mathbb{E} \left(\prod_{j=1}^k \mu_{1,j}^3 \right) \right) \right] \\ &= -\frac{d(k-1)}{k} \left[\mathbb{E}[\mu_{1,1}]^k + \frac{1}{2} \mathbb{E}(\mu_{1,1}^2)^k + o(\mathbb{E}(\mu_{1,1}^3)^k) \right]. \end{aligned}$$

Hence, because the construction of $\pi_{d,\beta}^*$ ensures that $\mu_{1,1}$ and $1 - \mu_{1,1}$ are identically distributed and thus $\mathbb{E}[\mu_{1,1}] = 1/2$ and because $\pi_{d,\beta}^*$ satisfies (6.2), we obtain

$$(8.22) \quad \frac{d(k-1)}{k} \mathbb{E} \left[\log \left(1 - \prod_{j=1}^k \mu_{1,j} \right) \right] = -\frac{d(k-1)}{k} [2^{-k} + 2^{-2k-1} + o(4^{-k})].$$

Moving on to the other term, we set $\Pi^+ = \prod_{i=1}^{\gamma^+} (1 - \prod_{j=1}^{k-1} \mu_{i,j})$ and $\Pi^- = \prod_{i=1}^{\gamma^-} (1 - \prod_{j=1}^{k-1} \mu_{\gamma^++i,j})$. Then

$$\begin{aligned} & \mathbb{E} \left[\log \left(\prod_{i=1}^{\gamma^+} \left(1 - \prod_{j=1}^{k-1} \mu_{i,j} \right) + \prod_{i=1}^{\gamma^-} \left(1 - \prod_{j=1}^{k-1} \mu_{\gamma^++i,j} \right) \right) \right] = \mathbb{E}[\log(\Pi^+ + \Pi^-)] \\ (8.23) \quad &= \mathbb{E}[\log \Pi^+] + \mathbb{E} \left[\log \left(2 + \frac{\Pi^-}{\Pi^+} - 1 \right) \right] \\ &= \log(2) + \frac{d}{2} \mathbb{E} \log \left(1 - \prod_{j=1}^{k-1} \mu_{1,i} \right) + \mathbb{E} \left[\log \left(1 + \frac{1}{2} \left(\frac{\Pi^-}{\Pi^+} - 1 \right) \right) \right]. \end{aligned}$$

Further,

$$(8.24) \quad \begin{aligned} \frac{d}{2} \mathbb{E} \log \left(1 - \prod_{j=1}^{k-1} \mu_{1,i} \right) &= -\frac{d}{2} [-2^{1-k} - 2^{1-2k} + o(4^{-k})] \\ &= -d2^{-k} - d2^{-2k} + o(2^{-k}). \end{aligned}$$

Moreover, using the inequality $\log(1+x) - x + x^2/2 \leq |x|^3$, we obtain

$$\mathbb{E} \left[\log \left(1 + \frac{1}{2} \left(\frac{\Pi^-}{\Pi^+} - 1 \right) \right) \right] \leq \frac{1}{2} \left(\frac{\Pi^-}{\Pi^+} - 1 \right) - \frac{1}{8} \left(\frac{\Pi^-}{\Pi^+} - 1 \right)^2 + o \left(\left(\frac{\Pi^-}{\Pi^+} - 1 \right)^3 \right).$$

Now,

$$\begin{aligned} \frac{\Pi^-}{\Pi^+} &= \exp \left(\log \frac{\Pi^-}{\Pi^+} \right) = \exp \left(\sum_{i=1}^{\gamma^+} \log \left(1 - \prod_{j=1}^{k-1} \mu_{i,j} \right) - \sum_{i=1}^{\gamma^-} \log \left(1 - \prod_{j=1}^{k-1} \mu_{i+\gamma^+,j} \right) \right) \\ &= 1 + \sum_{i=1}^{\gamma^+} \log \left(1 - \prod_{j=1}^{k-1} \mu_{i,j} \right) - \sum_{i=1}^{\gamma^-} \log \left(1 - \prod_{j=1}^{k-1} \mu_{i+\gamma^+,j} \right) \end{aligned}$$

$$\begin{aligned}
 &+ \frac{1}{2} \left(\sum_{i=1}^{\gamma^+} \log \left(1 - \prod_{j=1}^{k-1} \mu_{i,j} \right) - \sum_{i=1}^{\gamma^-} \log \left(1 - \prod_{j=1}^{k-1} \mu_{i+\gamma^+,j} \right) \right)^2 \\
 &+ O \left(\sum_{h \geq 3} \frac{1}{h!} \left(\sum_{i=1}^{\gamma^+ + \gamma^-} \log \left(1 - \prod_{j=1}^{k-1} \mu_{i,j} \right) \right)^h \right).
 \end{aligned}$$

Hence, using that the γ^\pm and the $\mu_{1,i,j}, \mu_{2,i,j}$ are identically distributed, we obtain

$$\begin{aligned}
 &\mathbb{E} \left[\log \left(1 + \frac{1}{2} \left(\frac{\Pi^-}{\Pi^+} - 1 \right) \right) \right] \\
 (8.25) \quad &= \frac{1}{8} \mathbb{E} \left[\left(\sum_{i=1}^{\gamma^+} \log \left(1 - \prod_{j=1}^{k-1} \mu_{i,j} \right) - \sum_{i=1}^{\gamma^-} \log \left(1 - \prod_{j=1}^{k-1} \mu_{i+\gamma^+,j} \right) \right)^2 \right] \\
 &+ O \left(\sum_{h \geq 3} \frac{1}{h!} \mathbb{E} \left[\left(\sum_{i=1}^{\gamma^+ + \gamma^-} \log \left(1 - \prod_{j=1}^{k-1} \mu_{i,j} \right) \right)^h \right] \right).
 \end{aligned}$$

Further, using the tail bound (6.2) for $\pi_{d,\beta}^*$ we obtain

$$\begin{aligned}
 &\mathbb{E} \left[\left(\sum_{i=1}^{\gamma^+} \log \left(1 - \prod_{j=1}^{k-1} \mu_{i,j} \right) - \sum_{i=1}^{\gamma^-} \log \left(1 - \prod_{j=1}^{k-1} \mu_{i+\gamma^+,j} \right) \right)^2 \right] \\
 (8.26) \quad &= \mathbb{E}[(\gamma^+ - \gamma^-)^2] 2^{2-2k} + o(2^{-k}) \\
 &= d 2^{2-2k} + o(2^{-k}).
 \end{aligned}$$

Similarly, the tail bound (6.2) and Bennett’s inequality imply that

$$(8.27) \quad \sum_{h \geq 3} \frac{1}{h!} \mathbb{E} \left[\left(\sum_{i=1}^{\gamma^+ + \gamma^-} \log \left(1 - \prod_{j=1}^{k-1} \mu_{i,j} \right) \right)^h \right] = o(2^{-k}).$$

Combining (8.25)–(8.27) we get

$$(8.28) \quad \mathbb{E} \left[\log \left(1 + \frac{1}{2} \left(\frac{\Pi^-}{\Pi^+} - 1 \right) \right) \right] = 2^{-1-2k} d + o(2^{-k}).$$

Finally, combining (8.21), (8.22), (8.23), (8.24) and (8.28), we obtain the assertion.

8.2. Proof of Proposition 2.6. The proof hinges on the so-called “1-step replica symmetry breaking interpolation method” from mathematical physics. Specifically, we seize upon the following result. Recall that γ^\pm signify independent $\text{Po}(d/2)$ variables and that $(\mu_{\pi,i,j})_{i,j}$ is a sequence of independent samples from a distribution π .

THEOREM 8.20 ([43]). *For any $y > 0, \beta > 0$, any probability distribution π on $[0, 1]$ and any $n \geq 1$ we have*

$$\begin{aligned}
 &\frac{y}{n} \mathbb{E}[\log Z(\Phi, \beta)] \\
 &\leq \mathbb{E} \left[\log \mathbb{E} \left[\left(\prod_{i=1}^{\gamma^+} 1 - (1 - e^{-\beta}) \prod_{j=1}^{k-1} \mu_{\pi,i,j} + \prod_{i=1}^{\gamma^-} 1 - (1 - e^{-\beta}) \prod_{j=1}^{k-1} \mu_{\pi,i+\gamma^+,j} \right)^y \middle| \gamma^+, \gamma^- \right] \right] \\
 &\quad - \frac{d(k-1)}{k} \log \mathbb{E} \left[\left(1 - (1 - e^{-\beta}) \prod_{j=1}^k \mu_{\pi,1,j} \right)^y \right].
 \end{aligned}$$

We apply Theorem 8.20 with the specific choice $\pi = \frac{1}{2}(\delta_1 + \delta_0)$. For the last expression we obtain

$$\begin{aligned}
 & -\frac{d(k-1)}{k} \log \mathbb{E} \left[\left(1 - (1 - e^{-\beta}) \prod_{j=1}^k \mu_{\pi,1,j} \right)^y \right] \\
 (8.29) \quad & = -\frac{d(k-1)}{k} \log(1 - 2^{-k}) + o(2^{-k}) \\
 & = -\frac{d(k-1)}{k} (-2^{-k} - 2^{-2k-1}) + o(2^{-k}) \\
 & = 2^{-k}d + 2^{-2k-1}d - \log(2) + c2^{-k} - 2^{-k-1} \log(2) + o(2^{-k}).
 \end{aligned}$$

Further, to estimate the first term let $(\mu_{\pi,i,j,h})_{i,j,h}$ be additional independent samples from π and

$$\Pi_+ = \prod_{i=1}^{\mathbf{y}^+} \left(1 - \prod_{j=1}^{k-1} \mu_{\pi,i,j,1} \right), \quad \Pi_- = \prod_{i=1}^{\mathbf{y}^-} \left(1 - \prod_{j=1}^{k-1} \mu_{\pi,i,j,2} \right).$$

Then for large β we have

$$\begin{aligned}
 (8.30) \quad & \mathbb{E} \left[\log \mathbb{E} \left[\left(\prod_{i=1}^{\mathbf{y}^+} 1 - (1 - e^{-\beta}) \prod_{j=1}^{k-1} \mu_{\pi,i,j,1} + \prod_{i=1}^{\mathbf{y}^-} 1 - (1 - e^{-\beta}) \prod_{j=1}^{k-1} \mu_{\pi,i,j,2} \right)^y \mid \mathbf{y}^\pm \right] \right] \\
 & = \mathbb{E}[\log \mathbb{E}[(\Pi^+ + \Pi^-)^y] \mid \mathbf{y}^\pm] + o(2^{-k}).
 \end{aligned}$$

Furthermore, Π^\pm are $\{0, 1\}$ -valued random variables and $\mathbb{E}[\Pi_\pm \mid \mathbf{y}^\pm] = (1 - 2^{1-k})^{\mathbf{y}^\pm}$. Therefore,

$$\begin{aligned}
 & \mathbb{E}[\log \mathbb{E}[(\Pi^+ + \Pi^-)^y] \mid \mathbf{y}^+, \mathbf{y}^-] \\
 & = \mathbb{E} \left[\log \left(\sum_{s=\pm 1} (1 - 2^{1-k})^{\mathbf{y}^s} (1 - (1 - 2^{1-k})^{\mathbf{y}^{-s}}) + 2^y (1 - 2^{1-k})^{\mathbf{y}^+ + \mathbf{y}^-} \right) \right] \\
 (8.31) \quad & = \mathbb{E}[\mathbf{y}^+ \log(1 - 2^{1-k})] \\
 & \quad + \mathbb{E}[\log(1 - (1 - 2^{1-k})^{\mathbf{y}^-} + (1 - (1 - 2^{1-k})^{\mathbf{y}^+})(1 - 2^{1-k})^{\mathbf{y}^- - \mathbf{y}^+} \\
 & \quad + 2^y (1 - 2^{1-k})^{\mathbf{y}^-})] \\
 & = \frac{d}{2} \log(1 - 2^{1-k}) \\
 & \quad + \mathbb{E}[\log(1 - (1 - 2^{1-k})^{\mathbf{y}^-} + (1 - (1 - 2^{1-k})^{\mathbf{y}^+})(1 - 2^{1-k})^{\mathbf{y}^- - \mathbf{y}^+} \\
 & \quad + 2^y (1 - 2^{1-k})^{\mathbf{y}^-})].
 \end{aligned}$$

Applying Bennett’s inequality, we obtain

$$\begin{aligned}
 & \mathbb{E}[\log(1 - (1 - 2^{1-k})^{\mathbf{y}^-} + (1 - (1 - 2^{1-k})^{\mathbf{y}^+})(1 - 2^{1-k})^{\mathbf{y}^- - \mathbf{y}^+} + 2^y (1 - 2^{1-k})^{\mathbf{y}^-})] \\
 (8.32) \quad & = \mathbb{E}[\log(1 - 2^{-k} + (1 - 2^{-k})(1 - 2^{1-k})^{\mathbf{y}^- - \mathbf{y}^+} + 2^{y-k})] + o(2^{-k}) \\
 & = \log(2) + \mathbb{E}[\log(1 - 2^{-k} + 2^{y-1-k} + (1 - 2^{1-k})^{\mathbf{y}^- - \mathbf{y}^+})] + o(2^{-k}) \\
 & = \log(2) - 2^{-k} + 2^{y-1-k} + \mathbb{E} \left[\log \left(1 + \frac{1}{2} ((1 - 2^{1-k})^{\mathbf{y}^- - \mathbf{y}^+} - 1) \right) \right] + o(2^{-k}).
 \end{aligned}$$

Further, once more by Bennett’s inequality,

$$\begin{aligned}
 (8.33) \quad & \mathbb{E} \left[\log \left(1 + \frac{1}{2} \left((1 - 2^{1-k})^{\mathbf{y}^- - \mathbf{y}^+} - 1 \right) \right) \right] \\
 &= \frac{1}{2} \mathbb{E} \left[(1 - 2^{1-k})^{\mathbf{y}^- - \mathbf{y}^+} - 1 \right] - \frac{1}{4} \mathbb{E} \left[\left((1 - 2^{1-k})^{\mathbf{y}^- - \mathbf{y}^+} - 1 \right)^2 \right] + o(2^{-k}).
 \end{aligned}$$

Since $\mathbf{y}^+, \mathbf{y}^-$ are $\text{Po}(d/2)$ variables, we obtain

$$\begin{aligned}
 (8.34) \quad & \mathbb{E} \left[(1 - 2^{1-k})^{\mathbf{y}^- - \mathbf{y}^+} - 1 \right] \\
 &= \mathbb{E} \left[\exp \left((\mathbf{y}^- - \mathbf{y}^+) \log(1 - 2^{1-k}) \right) - 1 \right] \\
 &= \mathbb{E} \left[(\mathbf{y}^+ - \mathbf{y}^-) \log(1 - 2^{1-k}) + \frac{1}{2} (\mathbf{y}^+ - \mathbf{y}^-)^2 \log^2(1 - 2^{1-k}) \right] \\
 &= \mathbb{E} \left[(\mathbf{y}^+ - \mathbf{y}^-)^2 \right] 2^{1-2k} = d 2^{1-2k}.
 \end{aligned}$$

Similarly,

$$(8.35) \quad \mathbb{E} \left[\left((1 - 2^{1-k})^{\mathbf{y}^- - \mathbf{y}^+} - 1 \right)^2 \right] = \mathbb{E} \left[(\mathbf{y}^+ - \mathbf{y}^-)^2 \right] 2^{1-2k} = d 2^{1-2k}.$$

Plugging (8.34) and (8.35) into (8.33), we obtain

$$(8.36) \quad \mathbb{E} \left[\log \left(1 + \frac{1}{2} \left((1 - 2^{1-k})^{\mathbf{y}^- - \mathbf{y}^+} - 1 \right) \right) \right] = d 2^{-1-2k}.$$

Finally, combining (8.29), (8.30), (8.31), (8.32) and (8.36), we get

$$\begin{aligned}
 (8.37) \quad & \mathbb{E} \left[\log \mathbb{E} \left[\left(\prod_{i=1}^{\mathbf{y}^+} 1 - (1 - e^{-\beta}) \prod_{j=1}^{k-1} \mu_{i,j,1} + \prod_{i=1}^{\mathbf{y}^-} 1 - (1 - e^{-\beta}) \prod_{j=1}^{k-1} \mu_{i,j,2} \right)^y \middle| \mathbf{y}^+, \mathbf{y}^- \right] \right. \\
 & \quad \left. - \frac{d(k-1)}{ky} \log \mathbb{E} \left[\left(1 - (1 - e^{-\beta}) \prod_{j=1}^k \mu_j \right)^y \right] \right] \\
 & \leq \frac{c - 1 + 2^{y-1} - \log(2)/2}{2^k y} + o(2^{-k}).
 \end{aligned}$$

To complete the proof, we observe that the function $y \mapsto (c - 1 + 2^{y-1} - \log(2)/2)/y$ attains its minimum at $y < 1$ if $c < 3 \log 2/2$. Since the function value for $y = 1$ comes to $c - \log(2)/2$, (8.37) shows together with Theorem 8.20 that for any $c < 3 \log 2/2$ we have $n^{-1} \mathbb{E}[\log Z(\Phi, \beta)] \leq 2^{-k} (c - \log(2)/2 - \Omega(1))$. Hence, the assertion follows from Proposition 2.7.

APPENDIX: PROOF OF LEMMA 8.1

We include a full proof of Lemma 8.1 for the sake of completeness. The argument is an adaptation of the proofs from [20]. Recall that $G(\Phi)$ is the factor graph obtained from a CNF-formula Φ and set $\Omega = \{\pm 1\}$.

LEMMA A.1 ([20], Lemma 3.1). *For any integer $L > 0$ and any $\alpha > 0$ there exist $\varepsilon = \varepsilon(\alpha, L)$, $n_0 = n_0(\varepsilon, L)$ such that the following is true. Suppose $G(\Phi)$ is the factor graph corresponding to any formula Φ with $n > n_0$ variables. Moreover, assume that $\mu_{\Phi, \beta}$ is ε -extremal. Let $G^*(\Phi)$ be obtained from $G(\Phi)$ by adding L constraints nodes b_1, \dots, b_L arbitrarily and denote by Φ^* the formula corresponding to $G^*(\Phi)$. Then, $\mu_{\Phi^*, \beta}$ is α -extremal and*

$$(A.1) \quad \sum_{x \in V(G(\Phi))} d_{\text{TV}}(\mu_{\Phi, \beta, x}, \mu_{\Phi^*, \beta, x}) < \alpha n.$$

In the following, let Φ_n denote a random formula with n variables x_1, \dots, x_n . Now, we will proceed to the proof of Lemma 8.1. Following Aizenman–Sims–Starr [7], we are going to show that

$$(A.2) \quad \liminf_{n \rightarrow \infty} \mathbb{E} \left[\log \frac{Z(\Phi_n, \beta)}{Z(\Phi_{n-1}, \beta)} \right] \geq \liminf_{n \rightarrow \infty} \mathbb{E} [\mathfrak{B}_{d, \beta}(\pi_{\Phi, \beta})].$$

The assertion then follows by summing on n . To prove (A.2), we will couple the random variables $Z(\Phi_n, \beta)$ and $Z(\Phi_{n-1}, \beta)$ by way of a third formula $\hat{\Phi}$. Specifically, let $\hat{\Phi}$ be the random formula with variables x_1, \dots, x_n obtained by including $\hat{m} = \text{Po}(\hat{d}n/k)$ independent random clauses, where

$$\hat{d} = d \frac{n-1}{n-k}.$$

Further, set $q = (n-k)/(n-1)$ and let Φ' be a random formula obtained from $\hat{\Phi}$ by deleting each clause with probability $1-q$ independently. Let A be the set of clauses removed from $\hat{\Phi}$ to obtain Φ' . In addition, obtain Φ'' from $\hat{\Phi}$ by selecting a variable \mathbf{x} uniformly at random and removing all constraints $a \in \partial_{\hat{\Phi}} \mathbf{x}$ along with \mathbf{x} itself. Then Φ' is distributed as Φ_n and Φ'' is distributed as Φ_{n-1} . Thus, $Z(\hat{\Phi}_n, \beta)$ is distributed as $Z(\Phi', \beta)$ and $Z(\Phi_{n-1}, \beta)$ is distributed as $Z(\Phi'', \beta)$.

FACT A.2. The two random formulas $\hat{\Phi}, \Phi_n$ have total variance distance $o(1)$.

PROOF. Given that $\hat{m} = m$ both formulas are identically distributed. Moreover, the random variable m is Poisson distributed with mean dn/k , which has total variation distance $o(1)$ from the distribution of \hat{m} . \square

For a clause b of $\hat{\Phi}$, we define

$$S(b) = \log \left[\sum_{\sigma \in \Omega^{\partial b}} e^{-\beta \mathbb{1}\{\sigma \neq b\}} \prod_{y \in \partial b} \mu_{\hat{\Phi}, \beta, y \rightarrow b}(\sigma_y) \right].$$

LEMMA A.3. A.a.s. we have $\log(Z(\hat{\Phi}, \beta)/Z(\Phi', \beta)) = o(1) + \sum_{a \in A} S(a)$.

PROOF. Given $\varepsilon > 0$ let $L = L(\varepsilon) > 0$ be a large enough number, let $\gamma = \gamma(\varepsilon, L) > \delta = \delta(\gamma) > 0$ be small enough and assume that n is sufficiently large. Let $X = |A|$, X is distributed as $\text{Po}((1-q)\hat{d}n/k) = \text{Po}(d \frac{(k-1)n}{k(n-k)})$. Then, the construction of Φ' ensures that

$$(A.3) \quad \mathbb{P}[X > L] < \varepsilon.$$

Instead of thinking of Φ' as being obtained from $\hat{\Phi}$ by removing X random clauses, we can think of $\hat{\Phi}$ as being obtained from Φ' by adding X random clauses a_1, \dots, a_X . More precisely, let $\Phi'_0 = \Phi'$ and $\Phi'_i = \Phi'_{i-1} \wedge a_i$ for $i \in [X]$. Then given X the triple $(\Phi', \hat{\Phi}, A)$ has the same distribution as $(\Phi', \Phi'_X, \{a_1, \dots, a_X\})$. Moreover, because $q\hat{d}n/k = dn/k$, Φ' has the same distribution as Φ_n . Therefore, our assumption (1.8) implies that Φ' is $o(1)$ -extremal a.a.s. Hence, Lemma A.1 implies that Φ'_{i-1} remains $o(1)$ -extremal a.a.s for any $1 \leq i \leq \min\{X, L\}$. Consequently, Lemma 4.2 implies that Φ'_{i-1} is $(o(1), k)$ -extremal a.a.s. Since ∂a_i is chosen uniformly and independently of a_1, \dots, a_{i-1} , Markov’s inequality shows that for every

$$\mathbb{P} \left[\sum_{\tau \in \Omega^k} \left| \sum_{\sigma \in \Omega^n} \mathbb{1}\{\forall y \in \partial a_i : \sigma_y = \tau_y\} \mu_{\Phi'_{i-1}, \beta}(\sigma) - \prod_{y \in \partial a_i} \mu_{\Phi'_{i-1}, \beta, y}(\tau_y) \right| \geq \delta \right] < \varepsilon,$$

for n large enough. Further, since the clauses $(a_i)_{i \in [X]}$ are chosen independently and because $\mu_{\hat{\Phi}, y \rightarrow a_i}(\tau_y)$ is the marginal in the formula $\hat{\Phi} - a_i$, (A.1) and (A.3) imply that

$$\mathbb{P}\left[\forall i \in [X] : \sum_{\tau \in \Omega^k} \left| \prod_{y \in \partial a_i} \mu_{\hat{\Phi}, \beta, y \rightarrow a_i}(\tau_y) - \prod_{y \in \partial a_i} \mu_{\Phi'_{i-1}, \beta, y}(\tau_y) \right| \geq \delta \right] < 2\varepsilon.$$

Hence, with probability at least $1 - 3\varepsilon$ the bound

$$(A.4) \quad \sum_{\tau \in \Omega^k} \left| \sum_{\sigma \in \Omega^n} \mathbb{1}\{\forall y \in \partial a_i : \sigma_y = \tau_y\} \mu_{\Phi'_{i-1}, \beta}(\sigma) - \prod_{y \in \partial a_i} \mu_{\hat{\Phi}, \beta, y \rightarrow a_i}(\tau_y) \right| < 2\delta$$

holds for all $i \in [X]$ simultaneously. Further, the definition (1.1) of the partition function shows that for any $i \in [X]$,

$$\frac{Z(\Phi'_i, \beta)}{Z(\Phi'_{i-1}, \beta)} = \sum_{\sigma \in \Omega^{\partial a_i}} e^{-\beta \mathbb{1}\{\sigma \neq a_i\}} \sum_{\tau \in \Omega^n} \mathbb{1}\{\forall y \in \partial a_i : \tau_y = \sigma_y\} \mu_{\Phi'_{i-1}, \beta}(\tau).$$

Thus, if (A.4) holds and if δ is chosen sufficiently small then

$$\left| \frac{Z(\Phi'_i, \beta)}{Z(\Phi'_{i-1}, \beta)} - \sum_{\sigma \in \Omega^{\partial a_i}} e^{-\beta \mathbb{1}\{\sigma \neq a_i\}} \prod_{y \in \partial a_i} \mu_{\hat{\Phi}, \beta, y \rightarrow a_i}(\sigma_y) \right| < \gamma.$$

Finally, the assertion follows by taking logarithms and summing over $i = 1, \dots, X$. \square

LEMMA A.4. *Let $U = \bigcup_{a \in \partial_{\hat{\Phi}} \mathbf{x}} \partial a$. Then a.a.s we have*

$$\log \frac{Z(\hat{\Phi}, \beta)}{Z(\Phi'', \beta)} = o(1) + \log \sum_{\tau \in \Omega^U} \prod_{a \in \partial_{\hat{\Phi}} \mathbf{x}} e^{-\beta \mathbb{1}\{\tau \neq a\}} \prod_{y \in \partial a \setminus \mathbf{x}} \mu_{\hat{\Phi}, \beta, y \rightarrow a}(\tau_y).$$

PROOF. Given $\varepsilon > 0$ let $L = L(\varepsilon) > 0$ be large enough, let $\gamma = \gamma(\varepsilon, L) > \delta = \delta(\gamma) > 0$ be small enough and assume that n is sufficiently large. Letting $Y = |\partial_{\hat{\Phi}} \mathbf{x}|$, we can pick L large enough so that

$$(A.5) \quad \mathbb{P}[Y > L] < \varepsilon.$$

As in the previous proof, we think of $\hat{\Phi}$ as being obtained from Φ'' by adding a new variable \mathbf{x} and Y independent clauses a_1, \dots, a_Y such that $\mathbf{x} \in \partial a_i$ for all i . Then assumption (1.8), Lemma A.1 and Lemma 4.2 imply that

$$(A.6) \quad \mathbb{P}\left[\sum_{\tau \in \Omega^{U \setminus \{\mathbf{x}\}}} \left| \sum_{\sigma \in \Omega^n} \mathbb{1}\{\forall y \in U \setminus \{\mathbf{x}\} : \sigma_y = \tau_y\} \mu_{\Phi'', \beta}(\sigma) - \prod_{i=1}^Y \prod_{y \in \partial a_i \setminus \{\mathbf{x}\}} \mu_{\hat{\Phi}, \beta, y \rightarrow a_i}(\tau_y) \right| \geq \delta \mid Y \leq L \right] = o(1).$$

In addition, (1.1) yields

$$\frac{Z(\hat{\Phi}, \beta)}{Z(\Phi'', \beta)} = \sum_{\tau \in \Omega^U} \prod_{i=1}^Y e^{-\beta \mathbb{1}\{\tau \neq a_i\}} \sum_{\sigma \in \Omega^n} \mathbb{1}\{\forall y \in U \setminus \{\mathbf{x}\} : \sigma_y = \tau_y\} \mu_{\Phi'', \beta}(\sigma).$$

Hence, (A.5) and (A.6) show that with probability at least $1 - 2\varepsilon$,

$$\left| \frac{Z(\hat{\Phi}, \beta)}{Z(\Phi'', \beta)} - \sum_{\tau \in \Omega^U} \prod_{i=1}^Y e^{-\beta \mathbb{1}\{\tau \neq a_i\}} \prod_{y \in \partial a_i \setminus \{\mathbf{x}\}} \mu_{\hat{\Phi}, \beta, y \rightarrow a_i}(\tau_y) \right| < \gamma.$$

The assertion follows by taking logarithm. \square

CLAIM A.5. *If a_1, \dots, a_Y are the clauses containing \mathbf{x} and $U = \bigcup_{i=1}^Y \partial_{\hat{\Phi}} a_i$ then a.a.s.*

$$\begin{aligned} & \sum_{\tau \in \Omega^U} \prod_{i=1}^Y e^{-\beta \mathbb{1}\{\tau \neq a_i\}} \prod_{y \in \partial a_i \setminus \{\mathbf{x}\}} \mu_{\hat{\Phi}, \beta, y \rightarrow a_i}(\tau_y) \\ &= \sum_{\tau_{\mathbf{x}} = \pm 1} \prod_{i=1}^Y \sum_{\tau \in \Omega^{\partial a_i \setminus \{\mathbf{x}\}}} e^{-\beta \mathbb{1}\{\tau \neq a_i\}} \prod_{y \in \partial a_i \setminus \{\mathbf{x}\}} \mu_{\hat{\Phi}, \beta, y \rightarrow a_i}(\tau_y). \end{aligned}$$

PROOF. With probability $1 - o(1)$ for all $1 \leq i < j \leq Y$ we have $\partial a_i \cap \partial a_j \setminus \{\mathbf{x}\} = \emptyset$. □

PROOF OF LEMMA 8.1. Recall that $\pi_{\Phi, \beta} = \frac{1}{n} \sum_{i=1}^n \delta_{\mu_{\Phi, \beta}(\{\sigma_{x_i}=1\})}$. Moreover, let $(\rho_{\pi, i, j})_{i, j \geq 1}$ be an array of independent random variables with distribution $\pi_{\Phi, \beta}$ and define $(\mu_{i, j})_{i, j \geq 1}$ as in (2.14). Additionally, let $(\hat{\mu}_{i, j})_{i, j}$ be a family of independent random variables defined accordingly for $\pi_{\hat{\Phi}, \beta}$. Then Lemma A.1 shows that $W_2(\pi_{\hat{\Phi}, \beta}, \pi_{\Phi, \beta}) = o(1)$. Therefore, using Lemmas A.1 and A.3 and Wald’s identity, we can write

$$\begin{aligned} \mathbb{E} \log \frac{Z(\hat{\Phi}, \beta)}{Z(\Phi', \beta)} &= \frac{d(k-1)}{k} \mathbb{E} \log \left[1 - (1 - e^{-\beta}) \prod_{i=1}^k \hat{\mu}_{1, i} \right] + o(1) \\ &= \frac{d(k-1)}{k} \mathbb{E} \log \left[1 - (1 - e^{-\beta}) \prod_{i=1}^k \mu_{1, i} \right] + o(1). \end{aligned} \tag{A.7}$$

Similarly, Lemmas A.1 and A.4 and Claim A.5 yield

$$\begin{aligned} & \mathbb{E} \log \frac{Z(\hat{\Phi}, \beta)}{Z(\Phi'', \beta)} \\ &= \mathbb{E} \left[\log \left(\prod_{i=1}^{\gamma^+} 1 - (1 - e^{-\beta}) \prod_{j=1}^{k-1} \hat{\mu}_{i, j} + \prod_{i=1}^{\gamma^-} 1 - (1 - e^{-\beta}) \prod_{j=1}^{k-1} \hat{\mu}_{i+\gamma^+, j} \right) \right] \\ &+ o(1) \\ &= \mathbb{E} \left[\log \left(\prod_{i=1}^{\gamma^+} 1 - (1 - e^{-\beta}) \prod_{j=1}^{k-1} \mu_{i, j} + \prod_{i=1}^{\gamma^-} 1 - (1 - e^{-\beta}) \prod_{j=1}^{k-1} \mu_{i+\gamma^+, j} \right) \right] \\ &+ o(1). \end{aligned} \tag{A.8}$$

Finally, combining (A.7) and (A.8) completes the proof. □

Acknowledgments. The authors would like to thank the anonymous referees and the Editor for their constructive comments that improved the quality of this paper.

Funding. The first author was supported by DFG CO 646/3 and DFG CO 646/4. The second author was supported in part by ERC-Grant 772606-PTRCSP. The third author was supported by DFG CO 646/4.

REFERENCES

[1] ACHLIOPTAS, D., COJA-OGHLAN, A., HAHN-KLIMROTH, M., LEE, J., MÜLLER, N., PENSCHUCK, M. and ZHOU, G. The number of satisfying assignments of random 2-SAT formulas. *Random Structures Algorithms* **58** 609–647. <https://doi.org/10.1002/rsa.20993>

- [2] ACHLIOPTAS, D. and COJA-OGHLAN, A. (2008). Algorithmic barriers from phase transitions. In *Proc. 49th FOCS* 793–802.
- [3] ACHLIOPTAS, D. and MOORE, C. (2006). Random k -SAT: Two moments suffice to cross a sharp threshold. *SIAM J. Comput.* **36** 740–762. MR2263010 <https://doi.org/10.1137/S0097539703434231>
- [4] ACHLIOPTAS, D., NAOR, A. and PERES, Y. (2005). Rigorous location of phase transitions in hard optimization problems. *Nature* **435** 759–764.
- [5] ACHLIOPTAS, D. and PERES, Y. (2004). The threshold for random k -SAT is $2^k \log 2 - O(k)$. *J. Amer. Math. Soc.* **17** 947–973. MR2083472 <https://doi.org/10.1090/S0894-0347-04-00464-3>
- [6] ACHLIOPTAS, D. and SORKIN, G. B. (2000). Optimal myopic algorithms for random 3-SAT. In *41st Annual Symposium on Foundations of Computer Science (Redondo Beach, CA, 2000)* 590–600. IEEE Comput. Soc. Press, Los Alamitos, CA. MR1931856 <https://doi.org/10.1109/SFCS.2000.892327>
- [7] AIZENMAN, M., SIMS, R. and STARR, S. (2003). An extended variational principle for the SK spin-glass model. *Phys. Rev. B* **68** 214403.
- [8] BAPST, V. and COJA-OGHLAN, A. (2016). Harnessing the Bethe free energy. *Random Structures Algorithms* **49** 694–741. MR3570985 <https://doi.org/10.1002/rsa.20692>
- [9] BAPST, V. and COJA-OGHLAN, A. (2016). The condensation phase transition in the regular k -SAT model. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques. LIPIcs. Leibniz Int. Proc. Inform.* **60** Art. No. 22, 18 pp. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern. MR3566764
- [10] BARTHA, Z., SUN, N. and ZHANG, Y. (2019). Breaking of 1RSB in random regular MAX-NAE-SAT. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science* 1405–1416. IEEE Comput. Soc. Press, Los Alamitos, CA. MR4228233
- [11] BRODER, A. Z., FRIEZE, A. M. and UPFAL, E. (1993). On the satisfiability and maximum satisfiability of random 3-CNF formulas. In *Proceedings of the Fourth Annual ACM-SIAM Symposium on Discrete Algorithms (Austin, TX, 1993)* 322–330. ACM, New York. MR1213244
- [12] CHAO, M.-T. and FRANCO, J. (1986). Probabilistic analysis of two heuristics for the 3-satisfiability problem. *SIAM J. Comput.* **15** 1106–1118. MR0861375 <https://doi.org/10.1137/0215080>
- [13] CHEESEMAN, P., KANEFSKY, B. and TAYLOR, W. (1991). Where the really hard problems are. In *Proc. IJCAI* 331–337.
- [14] CHVATAL, V. and REED, B. (1992). Mick gets some (the odds are on his side). In *Proc. 33th FOCS* 620–627.
- [15] COJA-OGHLAN, A. (2010). A better algorithm for random k -SAT. *SIAM J. Comput.* **39** 2823–2864. MR2645890 <https://doi.org/10.1137/09076516X>
- [16] COJA-OGHLAN, A. (2017). Belief propagation guided decimation fails on random formulas. *J. ACM* **63** Art. 49, 55 pp. MR3614857 <https://doi.org/10.1145/3005398>
- [17] COJA-OGHLAN, A., EFTHYMIU, C., JAAFARI, N., KANG, M. and KAPETANOPOULOS, T. (2018). Charting the replica symmetric phase. *Comm. Math. Phys.* **359** 603–698. MR3783558 <https://doi.org/10.1007/s00220-018-3096-x>
- [18] COJA-OGHLAN, A. and HAHN-KLIMROTH, M. (2021). The cut metric for probability distributions. *SIAM J. Discrete Math.* **35** 1096–1135. MR4266221 <https://doi.org/10.1137/19M126548X>
- [19] COJA-OGHLAN, A. and PANAGIOTOU, K. (2016). The asymptotic k -SAT threshold. *Adv. Math.* **288** 985–1068. MR3436404 <https://doi.org/10.1016/j.aim.2015.11.007>
- [20] COJA-OGHLAN, A. and PERKINS, W. (2018). Belief propagation on replica symmetric random factor graph models. *Ann. Inst. Henri Poincaré D* **5** 211–249. MR3813215 <https://doi.org/10.4171/AIHPD/53>
- [21] COJA-OGHLAN, A. and PERKINS, W. (2019). Bethe states of random factor graphs. *Comm. Math. Phys.* **366** 173–201. MR3919446 <https://doi.org/10.1007/s00220-019-03387-7>
- [22] COJA-OGHLAN, A. and WORMALD, N. (2018). The number of satisfying assignments of random regular k -SAT formulas. *Combin. Probab. Comput.* **27** 496–530. MR3816318 <https://doi.org/10.1017/S0963548318000263>
- [23] DIETZFELBINGER, M., GOERDT, A., MITZENMACHER, M., MONTANARI, A., PAGH, R. and RINK, M. (2010). Tight thresholds for cuckoo hashing via XORSAT. In *Proc. 35th ICALP* 213–225.
- [24] DING, J., SLY, A. and SUN, N. (2015). Proof of the satisfiability conjecture for large k . In *STOC'15—Proceedings of the 2015 ACM Symposium on Theory of Computing* 59–68. ACM, New York. MR3388183
- [25] DUBOIS, O. and MANDLER, J. (2002). The 3-XORSAT threshold. In *Proc. 43rd FOCS* 769–778.
- [26] EFTHYMIU, C., HAYES, T. P., ŠTEFANKOVIČ, D., VIGODA, E. and YIN, Y. (2019). Convergence of MCMC and loopy BP in the tree uniqueness region for the hard-core model. *SIAM J. Comput.* **48** 581–643. MR3945256 <https://doi.org/10.1137/17M1127144>
- [27] FRIEDGUT, E. (1999). Sharp thresholds of graph properties, and the k -sat problem. *J. Amer. Math. Soc.* **12** 1017–1054. MR1678031 <https://doi.org/10.1090/S0894-0347-99-00305-7>

- [28] FRIEZE, A. and SUEN, S. (1996). Analysis of two simple heuristics on a random instance of k -SAT. *J. Algorithms* **20** 312–355. MR1379227 <https://doi.org/10.1006/jagm.1996.0016>
- [29] FRIEZE, A. and WORMALD, N. C. (2005). Random k -SAT: A tight threshold for moderately growing k . *Combinatorica* **25** 297–305. MR2141660 <https://doi.org/10.1007/s00493-005-0017-3>
- [30] GALANIS, A., GOLDBERG, L. A., GUO, H. and YANG, K. (2021). Counting solutions to random CNF formulas. *SIAM J. Comput.* **50** 1701–1738. MR4344439 <https://doi.org/10.1137/20M1351527>
- [31] GOERDT, A. (1996). A threshold for unsatisfiability. *J. Comput. System Sci.* **53** 469–486. MR1423858 <https://doi.org/10.1006/jcss.1996.0081>
- [32] HETTERICH, S. (2016). Analysing survey propagation guided decimation on random formulas. In *43rd International Colloquium on Automata, Languages, and Programming. LIPIcs. Leibniz Int. Proc. Inform.* **55** Art. No. 65, 12 pp. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern. MR3577126
- [33] KRZAKALA, F., MONTANARI, A., RICCI-TERSENGHI, F., SEMERJIAN, G. and ZDEBOROVÁ, L. (2007). Gibbs states and the set of solutions of random constraint satisfaction problems. *Proc. Natl. Acad. Sci. USA* **104** 10318–10323. MR2317690 <https://doi.org/10.1073/pnas.0703685104>
- [34] MÉZARD, M. and MONTANARI, A. (2009). *Information, Physics, and Computation. Oxford Graduate Texts.* Oxford Univ. Press, Oxford. MR2518205 <https://doi.org/10.1093/acprof:oso/9780198570837.001.0001>
- [35] MÉZARD, M., PARISI, G. and ZECCHINA, R. (2002). Analytic and algorithmic solution of random satisfiability problems. *Science* **297** 812–815.
- [36] MOITRA, A. (2019). Approximate counting, the Lovász local lemma, and inference in graphical models. *J. ACM* **66** Art. 10, 25 pp. MR3936375 <https://doi.org/10.1145/3268930>
- [37] MONASSON, R. and ZECCHINA, R. (1996). Entropy of the K -satisfiability problem. *Phys. Rev. Lett.* **76** 3881–3885. MR1388240 <https://doi.org/10.1103/PhysRevLett.76.3881>
- [38] MONTANARI, A. and SHAH, D. (2007). Counting good truth assignments of random k -SAT formulae. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms* 1255–1264. ACM, New York. MR2485278
- [39] PANCHENKO, D. (2013). *The Sherrington–Kirkpatrick Model.* Springer, New York.
- [40] PANCHENKO, D. (2013). Spin glass models from the point of view of spin distributions. *Ann. Probab.* **41** 1315–1361. MR3098679 <https://doi.org/10.1214/11-AOP696>
- [41] PANCHENKO, D. (2014). On the replica symmetric solution of the K -sat model. *Electron. J. Probab.* **19** no. 67, 17 pp. MR3248196 <https://doi.org/10.1214/EJP.v19-2963>
- [42] PANCHENKO, D. (2018). On the K -sat model with large number of clauses. *Random Structures Algorithms* **52** 536–542. MR3783209 <https://doi.org/10.1002/rsa.20748>
- [43] PANCHENKO, D. and TALAGRAND, M. (2004). Bounds for diluted mean-fields spin glass models. *Probab. Theory Related Fields* **130** 319–336. MR2095932 <https://doi.org/10.1007/s00440-004-0342-2>
- [44] PITTEL, B. and SORKIN, G. B. (2016). The satisfiability threshold for k -XORSAT. *Combin. Probab. Comput.* **25** 236–268. MR3455676 <https://doi.org/10.1017/S0963548315000097>
- [45] SLY, A., SUN, N. and ZHANG, Y. (2016). The number of solutions for random regular NAE-SAT. In *57th Annual IEEE Symposium on Foundations of Computer Science—FOCS 2016* 724–731. IEEE Computer Soc., Los Alamitos, CA. MR3631035 <https://doi.org/10.1109/FOCS.2016.82>
- [46] TALAGRAND, M. (2001). The high temperature case for the random K -sat problem. *Probab. Theory Related Fields* **119** 187–212. MR1818246 <https://doi.org/10.1007/PL00008758>
- [47] VALIANT, L. G. (1979). The complexity of enumeration and reliability problems. *SIAM J. Comput.* **8** 410–421. MR0539258 <https://doi.org/10.1137/0208032>