Combinatorics, Probability and Computing (2026), **35**, pp. 433–485
doi:10.1017/S0963548326100406

CAMBRIDGE UNIVERSITY PRESS

ARTICLE

The random *k*-SAT Gibbs uniqueness threshold revisited

Arnab Chatterjee[1], Amin Coja-Oghlan[1], Catherine Greenhill[2], Vincent Pfenninger[3], Maurice Rolvien[4], Pavel Zakharov[1] and Kostas Zampetakis[1]

[1]TU Dortmund University, Germany, [2]UNSW Sydney, Australia, [3]TU Graz, Austria, and [4]University of Hamburg, Germany
Corresponding author: Amin Coja-Oghlan; Email: amin.coja-oghlan@tu-dortmund.de

(Received 3 June 2025; revised 5 November 2025; accepted 12 November 2025; first published online 27 March 2026)

Abstract
We prove that for any $k \geq 3$ for clause/variable ratios up to the Gibbs uniqueness threshold of the corresponding Galton-Watson tree, the number of satisfying assignments of random k-SAT formulas is given by the 'replica symmetric solution' predicted by physics methods [Monasson, Zecchina: Phys. Rev. Lett. **76** (1996)]. Furthermore, while the Gibbs uniqueness threshold is still not known precisely for any $k \geq 3$, we derive new lower bounds on this threshold that improve over prior work [Montanari and Shah: SODA (2007)]. The improvement is significant particularly for small k.

Keywords: random satisfiability; random graphs; Gibbs measures
2020 MSC Codes: 68Q87, 60C05, 68R05

1. Introduction

1.1 Background and motivation

Going back to experimental work from the 1990s, the most prominent question concerning random k-SAT has been to pinpoint the satisfiability threshold, defined as the largest density m/n of clauses m to variables n up to which satisfying assignments likely exist [6, 18]. Currently, the satisfiability threshold is known precisely in the case of $k = 2$ [21, 41] and for $k \geq k_0$ with k_0 an undetermined (large) constant [34]. The latter result confirms 'predictions' based on an analytic but non-rigorous physics technique called the 'cavity method'. Indeed, the cavity method predicts the satisfiability threshold for every $k \geq 3$ [52], but random k-SAT for 'small' $k \geq 3$ appears to be a particularly hard nut to crack. Additionally, according to the cavity method several phase transitions precede the satisfiability threshold and are expected to impact, among other things, the performance of algorithms [48]. One of these phase transitions, the *Gibbs uniqueness transition*, pertains to a spatial mixing property that also plays a pivotal role in the computational complexity of counting and sampling [62].

From a statistical physics viewpoint, the satisfiability threshold is only the second most important quantity associated with random k-SAT. The first place firmly belongs to the typical number of satisfying assignments, known as the *partition function* in physics parlance [51]. All the other predictions, including the location of the satisfiability threshold, ultimately derive from the formula for the number of satisfying assignments or closely related variables [50]. Yet there has been little progress on confirming the physics formula for the number of satisfying assignments rigorously.

© The Author(s), 2026. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

Three prior contributions stand out. First, a proof technique called the 'interpolation method' turns the physics prediction into a rigorous upper bound [36, 42, 61].[1] Second, in the case $k = 2$, conceptually much simpler than $k \geq 3$, the physics formula has been proved correct [3]. Third, Montanari and Shah [57] proved that also for $k \geq 3$ for certain clause/variable densities the 'replica symmetric solution' from physics correctly approximates the number of 'good' assignments that satisfy all but $o(n)$ clauses. However, it seems difficult to estimate the gap between the number of such 'good' assignments and the number of actual satisfying assignments. A rigorous method to this effect would likely imply the existence of uniform satisfiability thresholds for all $k \geq 3$, thereby resolving a long-standing conundrum [12, 37]. The proof of Montanari and Shah is based on the aforementioned Gibbs uniqueness property.

The aim of the present paper is to determine the number of *actual* satisfying assignments of random k-SAT formulas for clause/variable densities up to the Gibbs uniqueness threshold. Specifically, we verify that the 'replica symmetric solution' from [55, 56] yields the correct answer for any $k \geq 3$ right up to the Gibbs uniqueness threshold, even though the precise value of this threshold is not currently known. Additionally, we derive a new lower bound on the Gibbs uniqueness threshold. The improvement is particularly significant for 'small' $k \geq 3$. Combining these two results, we obtain the first rigorous formula for the number of satisfying assignments of random k-SAT formula for a non-trivial regime of clause/variable densities. Crucially, the result covers meaningful clause/variable densities even for small $k \geq 3$.

1.2 Results

Let $\mathbf{\Phi} = \mathbf{\Phi}_{d,k}(n)$ be the random k-CNF on n Boolean variables $x_1, \ldots, x_n$ with $m = m_n \sim \mathrm{Po}(dn/k)$ clauses $a_1, \ldots, a_m$. The clauses a_i are drawn independently and uniformly from the set of all $2^k\binom{n}{k}$ possible clauses with k distinct variables. Hence, the parameter d prescribes the expected number of clauses in which a given variable appears. Let $S(\mathbf{\Phi})$ be the set of satisfying assignments of $\mathbf{\Phi}$ and let $Z(\mathbf{\Phi}) = |S(\mathbf{\Phi})|$. We encode the Boolean values 'true' and 'false' by $+1$ and -1, respectively. Since right up to the satisfiability threshold $Z(\mathbf{\Phi})$ is of order $\exp(\Theta(n))$ w.h.p. for trivial reasons,[2] our objective is to study the random variable $n^{-1}\log Z(\mathbf{\Phi})$ as $n \to \infty$.

1.2.1 The number of satisfying assignments up to the Gibbs uniqueness threshold

The first main result vindicates the 'replica symmetric solution' for values of d up to the Gibbs uniqueness threshold of the Galton-Watson tree that mimics the local topology of $\mathbf{\Phi}$. Let us define these concepts precisely.

We begin with the Galton-Watson tree $\mathbb{T} = \mathbb{T}_{d,k}$, which is generated by a two-type branching process. The two types are *variable nodes* and *clause nodes*. The process starts with a single root variable node $\mathfrak{x}$. The offspring of any variable node is a $\mathrm{Po}(d)$ number of clause nodes, while every clause node begets precisely $k-1$ variable nodes. Additionally, independently for each clause node a and every variable node x that is either a child or the parent of a a *sign*, denoted $\mathrm{sign}(x, a) \in \{\pm 1\}$, is chosen uniformly at random. The resulting random tree $\mathbb{T}$ models the local structure of the random formula $\mathbf{\Phi}$ in the sense of local weak convergence [9, 49].[3]

Next, we define the Gibbs uniqueness property on the tree $\mathbb{T}$. For an integer $\ell \geq 0$ let $\mathbb{T}^{(\ell)}$ be the finite tree obtained by removing all variable and clause nodes at a distance greater than 2ℓ from the root $\mathfrak{x}$. We identify the finite tree $\mathbb{T}^{(\ell)}$ with a Boolean formula whose variables/clauses

[1] Strictly speaking, the contributions [36, 42, 61] deal with the 'random k-SAT model at positive tempertature', see Section 2.6. In Corollary 2.2 below we combine the interpolation method with a concentration argument to bound the number of actual satisfying assignments.

[2] For example, w.h.p. there are $\Omega(n)$ variables that do not appear in any clause.

[3] Corollary 3.7 below provides a precise statement to this effect.

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

are precisely the variable/clause nodes of $\mathbb{T}^{(\ell)}$. Let $S(\mathbb{T}^{(\ell)}) \neq \emptyset$ be the set of satisfying assignments of this formula and let $\boldsymbol{\tau}^{(\ell)} \in S(\mathbb{T}^{(\ell)})$ be a uniformly random satisfying assignment. Moreover, let $\partial^{2\ell}\mathfrak{x}$ be the set of variable nodes of $\mathbb{T}^{(\ell)}$ at distance precisely 2ℓ from the root $\mathfrak{x}$. Then for given d, k the tree $\mathbb{T} = \mathbb{T}_{d,k}$ has the *Gibbs uniqueness property* if

$$\lim_{\ell\to\infty} \mathbb{E}\left[\max_{\tau\in S(\mathbb{T}^{(\ell)})} \left|\mathbb{P}\left[\boldsymbol{\tau}^{(\ell)}(\mathfrak{x}) = 1 \mid \mathbb{T}\right] - \mathbb{P}\left[\boldsymbol{\tau}^{(\ell)}(\mathfrak{x}) = 1 \mid \mathbb{T},\ \forall x \in \partial^{2\ell}\mathfrak{x} : \boldsymbol{\tau}^{(\ell)}(x) = \tau(x)\right]\right|\right] = 0 \quad \text{(see [47])}. \tag{1.1}$$

In words, in the limit of large ℓ the truth value $\boldsymbol{\tau}^{(\ell)}(\mathfrak{x})$ of the root $\mathfrak{x}$ is asymptotically independent of the truth values $\{\boldsymbol{\tau}^{(\ell)}(x)\}_{x\in\partial^{2\ell}\mathfrak{x}}$ of the variables at distance 2ℓ from $\mathfrak{x}$. In light of the above, for any $k \geq 2$ we further define $d_{\mathrm{uniq}}(k)$ as

$$d_{\mathrm{uniq}}(k) = \inf\{d > 0 : \text{condition (1.1) fails to hold for } d, k\}. \tag{1.2}$$

It is easy to see that $d_{\mathrm{uniq}}(k)$ is strictly positive and finite for any $k \geq 2$. Indeed, in Theorem 1.2 we will derive explicit lower bounds on $d_{\mathrm{uniq}}(k)$. However, the exact value of $d_{\mathrm{uniq}}(k)$ is not currently known for any $k \geq 3$.

As a final preparation we need to spell out the 'replica symmetric solution' from [55]. This prediction comes in terms of a distributional fixed point problem, i.e., a fixed point problem on the space $\mathscr{P}(0,1)$ of probability measures on the open unit interval. Specifically, consider the *Belief Propagation operator*

$$\mathrm{BP}_{d,k} : \mathscr{P}(0,1) \to \mathscr{P}(0,1), \quad \pi \mapsto \hat{\pi} = \mathrm{BP}_{d,k}(\pi) \tag{1.3}$$

defined as follows. Let $d^+, d^- \sim \mathrm{Po}(d/2)$ be Poisson variables with expectation $d/2$. Moreover, let $(\boldsymbol{\mu}_{\pi,i,j})_{i,j\geq 1}$ be a sequence of i.i.d. random variables, each following distribution π. All these random variables are mutually independent. Further, let

$$\boldsymbol{\mu}_{\pi,i} = 1 - \prod_{j=1}^{k-1} \boldsymbol{\mu}_{\pi,i,j} \quad \text{for } i \geq 1, \quad \text{and} \quad \hat{\boldsymbol{\mu}}_{\pi} = \frac{\prod_{i=1}^{d^-} \boldsymbol{\mu}_{\pi,2i-1}}{\prod_{i=1}^{d^-} \boldsymbol{\mu}_{\pi,2i-1} + \prod_{i=1}^{d^+} \boldsymbol{\mu}_{\pi,2i}}. \tag{1.4}$$

Then $\hat{\pi}$ is the distribution of $\hat{\boldsymbol{\mu}}_{\pi}$. Furthermore, for a probability measure $\pi \in \mathscr{P}(0,1)$ define the *Bethe free entropy*[4]

$$\mathfrak{B}_{d,k}(\pi) = \mathbb{E}\left[\log\left(\prod_{i=1}^{d^-} \boldsymbol{\mu}_{\pi,2i-1} + \prod_{i=1}^{d^+} \boldsymbol{\mu}_{\pi,2i}\right) - \frac{d(k-1)}{k}\log\left(1 - \prod_{j=1}^{k} \boldsymbol{\mu}_{\pi,1,j}\right)\right], \tag{1.5}$$

provided that the expectation on the r.h.s. exists. Finally, let $\delta_{1/2} \in \mathscr{P}(0,1)$ be the atom at $1/2$ and let us write $\mathrm{BP}_{d,k}^{\ell}$ for the ℓ-fold application of the operator $\mathrm{BP}_{d,k}$.

Theorem 1.1. *Let $k \geq 3$ and assume that $0 < d < d_{\mathrm{uniq}}(k)$. Then the weak limit*

$$\pi_{d,k} = \lim_{\ell\to\infty} \mathrm{BP}_{d,k}^{\ell}(\delta_{1/2}) \in \mathscr{P}(0,1) \tag{1.6}$$

exists and

$$\lim_{n\to\infty} \frac{1}{n} \log Z(\boldsymbol{\Phi}) = \mathfrak{B}_{d,k}(\pi_{d,k}) \text{ in probability.} \tag{1.7}$$

The formula (1.7) matches the prediction from [55] precisely. Of course, part of the assertion of Theorem 1.1 is that the Bethe free entropy $\mathfrak{B}_{d,k}(\pi_{d,k})$ is well defined. Admittedly, the formula (1.7) is not 'explicit'. But the proof of Theorem 1.1 evinces that the convergence (1.6) occurs rapidly.

[4]Throughout the paper log refers to the natural logarithm.

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

Therefore, a randomised algorithm called 'population dynamics' [51] can be used to approximate (1.7) within any desired numerical accuracy.

1.2.2 An improved lower bound on Gibbs uniqueness

The obvious next task is to determine the Gibbs uniqueness threshold $d_{\mathrm{uniq}}(k)$. Currently, its value is known precisely only in the case $k=2$, where $d_{\mathrm{uniq}}(2)=2$ coincides with the random 2-SAT satisfiability threshold [3, 21, 41]. Furthermore, Montanari and Shah [57] proved that the *pure literal threshold*[5] $d_{\mathrm{pure}}(k)$ upper bounds $d_{\mathrm{uniq}}(k)$ for all $k \geq 2$.[6] The value of $d_{\mathrm{pure}}(k)$ admits a neat formula [16, 54]:

$$d_{\mathrm{uniq}}(k) \leq d_{\mathrm{pure}}(k) = \min_{z>0} \frac{z}{(1-\exp(-z/2))^{k-1}}. \tag{1.8}$$

Complementing the upper bound (1.8), Montanari and Shah derived a lower bound $d_{\mathrm{MS}}(k)$:

$$d_{\mathrm{MS}}(k) = \sup\left\{d>0 : d(k-1)(1-\exp(-d/2)/4)(1-\exp(-d/2)/2)^{k-2} < 1\right\} \leq d_{\mathrm{uniq}}(k). \tag{1.9}$$

Unfortunately, the bound (1.9) is tight not even in the case $k=2$, where $d_{\mathrm{uniq}}(2)=2$ while $d_{\mathrm{MS}}(2) \approx 1.16$. That said, the lower and upper bounds $d_{\mathrm{MS}}(k)$ and $d_{\mathrm{pure}}(k)$ match asymptotically in the limit of large k, as

$$d_{\mathrm{MS}}(k), d_{\mathrm{pure}}(k) = (2+o_k(1))\log k, \tag{1.10}$$

with $o_k(1)$ hiding a term that vanishes as $k \to \infty$. The following theorem yields an improved lower bound $d_{\mathrm{con}}(k)$ on $d_{\mathrm{uniq}}(k)$.

Theorem 1.2. *For all $k \geq 3$ we have*

$$d_{\mathrm{uniq}}(k) \geq d_{\mathrm{con}}(k) := \sup\left\{d>0 : \frac{d(k-1)}{2}(1-\exp(-d/2)/2)^{k-2} < 1\right\}. \tag{1.11}$$

An easy calculation reveals that

$$d_{\mathrm{MS}}(k) < d_{\mathrm{con}}(k) \quad \text{for every } k \geq 2. \tag{1.12}$$

Moreover, it is satisfactory that the formula (1.11) reproduces the correct (previously known) threshold $d_{\mathrm{uniq}}(2) = d_{\mathrm{con}}(2) = d_{\mathrm{pure}}(2) = 2$. That said, we have no reason to believe that (1.11) is tight for any $k \geq 3$.

Combining Theorems 1.1 and 1.2, we obtain the following.

Corollary 1.3. *Let $k \geq 3$. If $d < d_{\mathrm{con}}(k)$ then (1.7) holds.*

Corollary 1.3 constitutes the first rigorous result to determine the precise asymptotic value of $\log Z(\boldsymbol{\Phi})$ for a non-trivial regime of d for any $k \geq 3$. To elaborate, the formula (1.7) is trivially true for $d < 1/(k-1)$ because for such d the k-uniform hypergraph induced by the clauses of $\boldsymbol{\Phi}$ has no giant component and Belief Propagation is exact on acyclic graphical models [51]. But Corollary 1.3 applies to d well beyond this threshold, as displayed in Table 1. In particular, in contrast to much of the prior work on random k-SAT, Corollary 1.3 applies to a non-trivial regime of d even for 'small' $k \geq 3$.

Although Table 1 contains the values $d_{\mathrm{MS}}(k)$ from [57] for comparison, we emphasise that Montanari and Shah's result only yields the number of 'good' assignments satisfying all but $o(n)$

[5] This marks the threshold up to which the pure literal algorithm–which repeatedly assigns the preferred value to all variables appearing with a single sign–produces a satisfying assignment w.h.p.

[6] To be precise, Montanari and Shah established an upper bound on the Gibbs uniqueness threshold that turns out to coincide with the pure literal threshold, albeit without pointing out this identity.

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

Table 1. The values of $d_{\mathrm{MS}}(k)$, $d_{\mathrm{con}}(k)$, and $d_{\mathrm{pure}}(k)$ for $2 \le k \le 5$. Additionally, $d_{\mathrm{giant}}(k) = 1/(k-1)$ marks the giant component threshold of the hypergraph induced by the random k-CNF formula. Moreover, $d_{\mathrm{sat}}(k)$ is the satisfiability threshold according to physics predictions [50]. It is not hard to show that $d_{\mathrm{giant}}(k) \le d_{\mathrm{MS}}(k) \le d_{\mathrm{con}}(k) \le d_{\mathrm{uniq}}(k) \le d_{\mathrm{pure}}(k) \le d_{\mathrm{sat}}(k)$, for all $k \ge 2$

k	2	3	4	5
d_{giant}	1.0000	0.5000	0.3333	0.2500
d_{MS}	1.1625	0.8792	0.8695	0.9236
d_{con}	2.0000	1.3431	1.2451	1.2635
d_{pure}	2.0000	4.9108	6.1782	7.0178
d_{sat}	2.0000	12.801	39.724	105.585

clauses, rather than of actual satisfying assignments. In fact, the best prior rigorous bounds on the number of satisfying assignments for $d > 1/(k-1)$ derive from the first and the second moment methods. Specifically, the folklore first moment bound reads

$$\frac{1}{n}\log Z(\mathbf{\Phi}) \le \log 2 + \frac{d}{k}\log\left(1-2^{-k}\right) + o(1) \quad \text{w.h.p.} \tag{1.13}$$

Furthermore, Achlioptas and Peres [7] perform a second moment argument on the number of *balanced* satisfying assignments, i.e., satisfying assignments that enjoy a peculiar additional condition required to keep the second moment under control. They show that w.h.p.

$$\frac{1}{n}\log Z(\mathbf{\Phi}) \ge (1-d)\log 2 + \frac{d}{k}\log\left[\left(\lambda^{1/2}+\lambda^{-1/2}\right)^k - \lambda^{-k/2}\right] + o(1), \quad \text{where} \tag{1.14}$$
$$(1-\lambda)(1+\lambda)^{k-1} = 1,\ \lambda > 0.$$

Figure 1 illustrates the bounds (1.13)–(1.14) along with (1.7) for $k = 3$. As the figure shows, the correct value (1.7) is quite close to the first moment bound. That said, the first moment bound strictly exceeds $\mathfrak{B}_{d,k}(\pi_{d,k})$ for all $d > 0$, $k \ge 3$ [24]. On the other hand, Figure 1 demonstrates that the 'balanced second moment bound' (1.14) significantly undershoots $\mathfrak{B}_{d,3}(\pi_{d,3})$. Recall that Figure 1 is on a logarithmic scale; thus, even small differences translate into exponentially large errors.

1.3 Preliminaries and notation

Let Φ be a Boolean expression in conjunctive normal such that no clause contains the same variable twice. We write $V(\Phi)$ for the set of Boolean variables of Φ and $F(\Phi)$ for the set of clauses. The formula Φ gives rise to a bipartite graph $G(\Phi)$ on the vertex set $V(\Phi) \cup F(\Phi)$ in which a variable x and a clause a are adjacent iff variable x appears in clause a (either positively or negatively). Let $E(\Phi)$ denote the edge set of the graph $G(\Phi)$. Furthermore, for a vertex $v \in V(\Phi) \cup F(\Phi)$ let $\partial_\Phi v$ be the set of neighbours of v; where the reference to Φ is self-evident, we just write ∂v.

The graph $G(\Phi)$ induces a metric on $V(\Phi) \cup F(\Phi)$ by letting $\mathrm{dist}_\Phi(v, w)$ equal the length of the shortest path from v to w. For a vertex v and an integer $\ell \ge 0$ let $\partial_\Phi^\ell v = \partial^\ell v$ be the set of all vertices w at distance precisely ℓ from v.

For a clause a and a variable $x \in \partial a$ we define $\mathrm{sign}_\Phi(x, a) = 1$ if a contains x as a positive literal, and $\mathrm{sign}_\Phi(x, a) = -1$ if a contains the negation $\neg x$. (This is unambiguous because clause a is not allowed to contain both x and $\neg x$.) For a variable $x \in V(\Phi)$ and $s \in \{\pm 1\}$ we let $\partial_\Phi^s x = \partial^s x$ be the set of clauses $a \in \partial_\Phi x$ such that $\mathrm{sign}_\Phi(x, a) = s$. Where convenient we use the shorthand $\partial^\pm x = \partial^{\pm 1} x$. We say that a variable x is *pure* in Φ if $\mathrm{sign}_\Phi(x, a) = \mathrm{sign}_\Phi(x, b)$ for all $a, b \in \partial x$. More specifically,

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

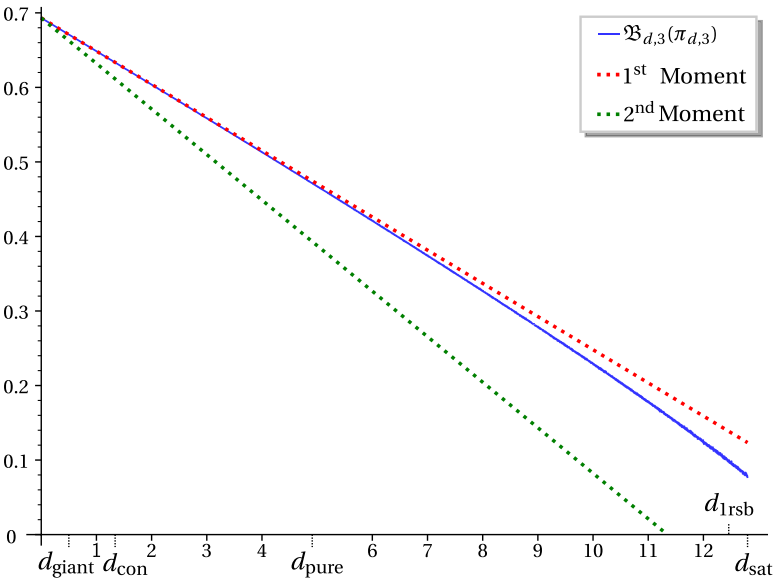


Figure 1. Comparison of $\mathfrak{B}_{d,k}(\pi_{d,k})$ with known bounds for $\lim_{n\to\infty}\frac{1}{n}\log Z(\boldsymbol{\Phi})$ for $k=3$. The red dotted line depicts the first moment upper bound (1.13), while the green dotted line represents the lower bound provided by (1.14). The blue line displays a numerical approximation of $\mathfrak{B}_{d,3}(\pi_{d,3})$. To obtain our values, we generated 10^6 samples from $\pi \approx \mathrm{BP}^{25}_{d,3}(\delta_{1/2})$ and then evaluated the corresponding empirical average of the expression in (1.5).

say that x is a *pure literal* of Φ if $\partial^- x = \emptyset$. Similarly, $\neg x$ is called a pure literal if $\partial^+ x = \emptyset$. A variable or literal that fails to be pure is called *mixed*.

For a literal $l \in \{x, \neg x : x \in V(\Phi)\}$ we let $|l|$ denote the underlying variable; thus, $|x| = |\neg x| = x$ for $x \in V(\Phi)$. Moreover, we define $\mathrm{sign}(x) = 1$ and $\mathrm{sign}(\neg x) = -1$. Further, for a literal l we define $1 \cdot l = l$ and $(-1)\cdot l = \neg l$.

If Φ is satisfiable, then $\boldsymbol{\sigma}_\Phi = (\boldsymbol{\sigma}_\Phi(x))_{x\in V(\Phi)}$ denotes a uniformly random satisfying assignment of Φ. Where the reference to Φ is obvious we just write $\boldsymbol{\sigma}$.

Let μ, ν be two probability measures on $\mathbb{R}^h$, let $q \geq 1$ and assume that $\int_{\mathbb{R}^h} \|x\|_q^q \mathrm{d}\mu(x), \int_{\mathbb{R}^h} \|x\|_q^q \mathrm{d}\nu(x) < \infty$. We recall that the *$L_q$-Wasserstein distance* of μ, ν is defined as

$$W_q(\mu,\nu) = \inf_{(\boldsymbol{\xi},\boldsymbol{\zeta})} \mathbb{E}\left[\|\boldsymbol{\xi}-\boldsymbol{\zeta}\|_q^q\right]^{1/q},$$

where the infimum is taken over all pairs $(\boldsymbol{\xi}, \boldsymbol{\zeta})$ of random variables defined on the same probability space Ω such that $\boldsymbol{\xi}$ has distribution μ and $\boldsymbol{\zeta}$ has distribution ν. If X, Y are random variables with distributions μ, ν, it is convenient to use the shorthand $W_q(X,Y) = W_q(\mu,\nu)$, provided that $\mathbb{E}[\|X\|_q^q], \mathbb{E}[\|Y\|_q^q] < \infty$.

For two random variables X, Y we write $X \sim Y$ if X, Y are identically distributed. Moreover, for a probability distribution μ and a random variable X we write $X \sim \mu$ if X has distribution μ.

We will make repeated use of the following tail bound for Poisson variables.

Lemma 1.4 (Bennett's inequality [14, Theorem 2.9]). *Suppose that $X \sim Po(\lambda)$ with $\lambda > 0$ and let $\varphi(x) = (1+x)\log(1+x) - x$ for $x > -1$. Then*

$$\mathbb{P}[X \geq \lambda + t] \leq \exp(-\lambda\varphi(t/\lambda)) \qquad \text{for any } t > 0,$$
$$\mathbb{P}[X \leq \lambda - t] \leq \exp(-\lambda\varphi(-t/\lambda)) \qquad \text{for any } 0 < t < \lambda.$$

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

For reals a, b we write

$$a \vee b = \max\{a, b\}, \quad a \wedge b = \min\{a, b\}.$$

Unless specified otherwise asymptotic notation $o(\cdot)$, $O(\cdot)$, etc. is understood to refer to the limit $n \to \infty$. The symbol $\tilde{O}(\cdot)$ is understood to swallow polylog(n) terms. Throughout we tacitly assume that n is sufficiently large so that the various estimates are valid. We use the conventions $\log 0 = -\infty$ and $\log \infty = \infty$. Finally, throughout the paper we assume that $k \geq 3$ is a fixed integer.

2. Overview

In this section we survey the proofs of the main results. Subsequently, we discuss further related work. The proof details are deferred to the remaining sections; see Section 2.7 for pointers. We assume throughout that $k \geq 3$.

2.1 Existence of the fixed point and upper bound

As a first step towards the proof of Theorem 1.1 we prove that the limit (1.6) exists for $d < d_{\mathrm{uniq}}(k)$. More precisely, we will establish the following statement.

Proposition 2.1. *For every $k \geq 3$ and every $d < d_{\mathrm{uniq}}(k)$, the W_1-limit $\pi_{d,k} = \lim_{\ell\to\infty} \mathrm{BP}_{d,k}^{\ell}(\delta_{1/2})$ exists and*

$$\mathbb{E}\left[\log^2 \boldsymbol{\mu}_{\pi_{d,k},1,1}\right] + \mathbb{E}\left|\log\left(\prod_{i=1}^{d^-} \boldsymbol{\mu}_{\pi_{d,k},2i} + \prod_{i=1}^{d^+} \boldsymbol{\mu}_{\pi_{d,k},2i-1}\right)\right| + \mathbb{E}\left|\log\left(1 - \prod_{j=1}^{k} \boldsymbol{\mu}_{\pi_{d,k},1,j}\right)\right| < \infty. \tag{2.1}$$

In addition, $\boldsymbol{\mu}_{\pi_{d,k},1,1}$ and $1 - \boldsymbol{\mu}_{\pi_{d,k},1,1}$ are identically distributed.

The existence of the limit $\pi_{d,k}$ is an easy consequence of the Gibbs uniqueness property. As an aside, the limit $\pi_{d,k} = \lim_{\ell\to\infty} \mathrm{BP}_{d,k}^{\ell}(\delta_{1/2})$ is a fixed point of the Belief Propagation operator, i.e.,

$$\pi_{d,k} = \mathrm{BP}_{d,k}(\pi_{d,k}). \tag{2.2}$$

The proof of the bound (2.1) is a bit more subtle and requires a few preparations, but we will come to that. The upshot of (2.1) is that the Bethe free entropy $\mathfrak{B}_{d,k}(\pi_{d,k})$ is well defined.

With the fixed point $\pi_{d,k}$ in hand we can bring to bear the 'interpolation method' to the upper bound the likely value of $\log Z(\boldsymbol{\Phi})$.

Corollary 2.2. *If $d < d_{\mathrm{uniq}}(k)$ then w.h.p. we have $\frac{1}{n}\log Z(\boldsymbol{\Phi}) \leq \mathfrak{B}_{d,k}(\pi_{d,k}) + o(1)$.*

The interpolation method is a mainstay of the study of disordered systems in mathematical physics and has also been used to investigate random constraint satisfaction problems. In particular, the variant of the interpolation method from [61] (in combination with Proposition 2.1) easily implies that

$$\limsup_{n\to\infty} \frac{1}{n}\mathbb{E}\left[\log\left(Z(\boldsymbol{\Phi}) \vee 1\right)\right] \leq \mathfrak{B}_{d,k}(\pi_{d,k});$$

taking the logarithm of $Z(\boldsymbol{\Phi}) \vee 1$ ensures that the expectation is well defined, as it is possible (albeit unlikely for $d < d_{\mathrm{uniq}}(k)$) that $\boldsymbol{\Phi}$ is unsatisfiable. The added value of Corollary 2.2 is that we obtain a bound that holds *with high probability*, rather than just a bound on the expectation. The interpolation method was used in [3] in a similar fashion to prove a 'with high probability' bound on the number of satisfying assignments of random 2-CNFs. The proof of Corollary 2.2 is an adaptation of that argument to $k \geq 3$.

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

2.2 A matching lower bound

The key step towards Theorem 1.1 is to establish a lower bound on $\log Z(\mathbf{\Phi})$ that matches the upper bound from Corollary 2.2. To accomplish this task we employ a coupling argument known as the 'Aizenman-Sims-Starr scheme' in mathematical physics. Its original version was intended to estimate the partition function of the Sherrington-Kirkpatrick model, a spin glass model [8]. But the technique has since been employed in probabilistic combinatorics (e.g., [24, 25, 29]). By comparison to the mathematical physics context, the crucial difference is that here our objective is to count actual satisfying assignments where every single clause imposes a hard constraint, whereas in spin glass theory constraints are soft. The same issue occurred in previous work on the random 2-SAT problem [3]. However, in that case a relatively simple percolation argument was sufficient to deal with the ensuing complications. As we will see, for $k \geq 3$ considerably more care is needed.

But first things first. The basic idea behind the Aizenman-Sims-Starr argument is to perform a kind of induction. Translated to random k-SAT this means that we couple the random k-CNF $\mathbf{\Phi}_{d,k}(n)$ with n variables with the random k-CNF $\mathbf{\Phi}_{d,k}(n+1)$ with $n+1$ variables. Recall that $\mathbf{\Phi}_{d,k}(n)$ comprises $m_n \sim \mathrm{Po}(dn/k)$ independent random clauses. Ultimately Theorem 1.1 is going to be a consequence of Corollary 2.2 and the following statement.

Proposition 2.3. *If* $d < d_{\mathrm{uniq}}(k)$ *then* $\mathbb{E}\left[\log\left(Z(\mathbf{\Phi}_{d,k}(n+1)) \vee 1\right)\right] - \mathbb{E}\left[\log\left(Z(\mathbf{\Phi}_{d,k}(n)) \vee 1\right)\right] = \mathfrak{B}_{d,k}(\pi_{d,k}) + o(1)$.

Once again we work with $Z(\mathbf{\Phi}_{d,k}(n)) \vee 1$ and $Z(\mathbf{\Phi}_{d,k}(n+1)) \vee 1$ to ensure that the expectations are well defined.

To prove Proposition 2.3 we couple the random formulas $\mathbf{\Phi}_{d,k}(n+1)$ and $\mathbf{\Phi}_{d,k}(n)$ as follows.

- **CPL1:** Let $\mathbf{\Phi}'$ be a random k-CNF with variables $x_1, \ldots, x_n$ and $m' \sim \mathrm{Po}(d(n-k+1)/k)$ clauses.
- **CPL2:** Obtain $\mathbf{\Phi}''$ from $\mathbf{\Phi}'$ by adding another $\mathbf{\Delta}'' \sim \mathrm{Po}(d(k-1)/k)$ independent random clauses.
- **CPL3:** Obtain $\mathbf{\Phi}'''$ from $\mathbf{\Phi}'$ by adding one new variable x_{n+1} and $\mathbf{\Delta}''' \sim \mathrm{Po}(d)$ independent random clauses that each contain x_{n+1} and $k-1$ other variables from $\{x_1, \ldots, x_n\}$.

Observe that $\mathbf{\Phi}''$ ultimately has variables $x_1, \ldots, x_n$ and a total of $m_n \sim \mathrm{Po}(dn/k)$ random clauses. Thus, $\mathbf{\Phi}''$ is identical to the random formula $\mathbf{\Phi}_{d,k}(n)$. Similarly, $\mathbf{\Phi}'''$ has the same distribution as $\mathbf{\Phi}_{d,k}(n+1)$. Consequently, we obtain the following.

Fact 2.4. *For any* $d > 0$ *we have* $Z(\mathbf{\Phi}_{d,k}(n)) \sim Z(\mathbf{\Phi}'')$ *and* $Z(\mathbf{\Phi}_{d,k}(n+1)) \sim Z(\mathbf{\Phi}''')$.

The coupling **CPL1–CPL3** reduces the proof of Proposition 2.3 to getting a handle on the differences $\log\left(Z(\mathbf{\Phi}'') \vee 1\right) - \log\left(Z(\mathbf{\Phi}') \vee 1\right)$ and $\log\left(Z(\mathbf{\Phi}''') \vee 1\right) - \log\left(Z(\mathbf{\Phi}') \vee 1\right)$. More precisely, recalling (1.4)–(1.5), we see that Proposition 2.3 is a consequence of the following two statements.

Proposition 2.5. *If* $d < d_{\mathrm{uniq}}(k)$ *then*

$$\mathbb{E}\left[\log\frac{Z(\mathbf{\Phi}'') \vee 1}{Z(\mathbf{\Phi}') \vee 1}\right] = \frac{d(k-1)}{k}\mathbb{E}\left[\log\left(1 - \prod_{j=1}^{k} \boldsymbol{\mu}_{\pi_{d,k},1,j}\right)\right] + o(1). \tag{2.3}$$

Proposition 2.6. *If* $d < d_{\mathrm{uniq}}(k)$ *then*

$$\mathbb{E}\left[\log\frac{Z(\mathbf{\Phi}''') \vee 1}{Z(\mathbf{\Phi}') \vee 1}\right] = \mathbb{E}\left[\log\left(\prod_{i=1}^{d^-} \boldsymbol{\mu}_{\pi_{d,k},2i} + \prod_{i=1}^{d^+} \boldsymbol{\mu}_{\pi_{d,k},2i-1}\right)\right] + o(1). \tag{2.4}$$

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

To prove Propositions 2.5– 2.6 we effectively need to trace the impact that *local* changes have on the number of satisfying assignments. Indeed, under the coupling **CPL1–CPL3**, the formula $\mathbf{\Phi}''$ is obtained from the 'base formula' $\mathbf{\Phi}'$ by adding just a bounded expected number of random clauses. Thus, if we imagine that, as both the first moment upper bound (1.13) and the balanced second moment lower bound (1.14) suggest, each additional random clause typically reduces the number of satisfying assignments by a constant factor, then the quantity $|\log(Z(\mathbf{\Phi}'')/Z(\mathbf{\Phi}'))|$ should be bounded with probability close to one. Similar reasoning applies to $\mathbf{\Phi}'''$.

Yet while with high probability the local changes that turn $\mathbf{\Phi}'$ into $\mathbf{\Phi}''$ or $\mathbf{\Phi}'''$ are indeed benign, because we are dealing with hard constraints there is a non-negligible probability that $\log(Z(\mathbf{\Phi}'')/Z(\mathbf{\Phi}'))$ and $\log(Z(\mathbf{\Phi}''')/Z(\mathbf{\Phi}'))$ could be large. Indeed, a single extra clause might wipe out all satisfying assignments of $\mathbf{\Phi}'$, in which case

$$\log\frac{Z(\mathbf{\Phi}'')\vee 1}{Z(\mathbf{\Phi}')\vee 1} = -\log Z(\mathbf{\Phi}') = -\Omega(n).$$

Hence, we need to argue that such drastic changes are sufficiently rare. The following statement furnishes the necessary tail bound.

Proposition 2.7. *For $d < d_{\mathrm{uniq}}(k)$ we have*

$$\mathbb{E}\left[\left|\log\frac{Z(\mathbf{\Phi}'')\vee 1}{Z(\mathbf{\Phi}')\vee 1}\right|^{3/2} + \left|\log\frac{Z(\mathbf{\Phi}''')\vee 1}{Z(\mathbf{\Phi}')\vee 1}\right|^{3/2}\right] = O(1). \tag{2.5}$$

2.3 Pure literal pursuit

The proof of Proposition 2.7 constitutes the main technical challenge towards the proof of Theorem 1.1. The linchpin of the proof is an algorithm that we call *Pure Literal Pursuit* ('PULP'). Its purpose is to trace the repercussions of setting a relatively small number of variables to specific truth values. More precisely, PULP will allow us to compare the number of satisfying assignments that set a few chosen variables to specific values to the total number of satisfying assignments.

To this end PULP attempts to solve the following optimisation task. Suppose we are given a k-CNF Φ and a set $\mathscr{L}$ of literals of Φ that we deem to be set to 'true'. We would like to identify a superset $\bar{\mathscr{L}} \supseteq \mathscr{L}$ of literals with the following properties; think of $\bar{\mathscr{L}}$ as a 'closure' of $\mathscr{L}$.

PULP1: every clause a of Φ that contains a literal from $\neg\bar{\mathscr{L}} = \{\neg l : l \in \bar{\mathscr{L}}\}$ also contains a literal from $\bar{\mathscr{L}}$.

PULP2: there is no literal l such that $l, \neg l \in \bar{\mathscr{L}}$.

Of course, it may be impossible to satisfy **PULP1** and **PULP2** simultaneously. In this case we ask PULP to report a 'contradiction'. But if **PULP1–PULP2** can be satisfied, we aim to find a closure $\bar{\mathscr{L}}$ of as small size $|\bar{\mathscr{L}}|$ as possible.

The combinatorial idea behind **PULP1–PULP2** is as follows. Deeming the literals from the initial set $\mathscr{L}$ 'true', our goal is to reconcile this assumption with the formula Φ. To this end we enhance the set $\mathscr{L}$. Clearly, any clause that contains the negation $\neg l$ of a literal l that we deem true also needs to contain another literal l' that is set to true. This is what **PULP1** asks. Furthermore, it would be contradictory to deem both l and its negation $\neg l$ true; this is **PULP2**.

The size of the closure $\bar{\mathscr{L}}$ yields a bound on the reduction in the number of satisfying assignments if we indeed insist on all literals $l \in \mathscr{L}$ being set to true. Formally, let $S(\Phi, \mathscr{L})$ be the set of all satisfying assignments $\sigma \in S(\Phi)$ under which all literals $l \in \mathscr{L}$ evaluate to 'true'. Also set $Z(\Phi, \mathscr{L}) = |S(\Phi, \mathscr{L})|$.

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

Algorithm 1. The PULP algorithm

Input: A k-CNF Φ and a set $\mathscr{L}$ of literals.

1: Let $\bar{\mathscr{L}} = \mathscr{L}$;
2: **while** there is a clause a that contains a literal from $\neg\bar{\mathscr{L}}$ but no literal from $\bar{\mathscr{L}}$
3: Pick such a clause a that minimises the distance from the initial set $\mathscr{L} = \{|l| : l \in \mathscr{L}\}$
4: **if** a consists of literals $l \in \neg\bar{\mathscr{L}}$ only
5: return 'contradiction' and halt;
6: **else**
7: choose $x \in \partial a$ with $x, \neg x \notin \bar{\mathscr{L}}$ that minimises $\mathfrak{h}_x(\mathrm{sign}(x, a), \Phi)$ and add $\mathrm{sign}(x, a) \cdot x$ to $\bar{\mathscr{L}}$;
8: **return** $\bar{\mathscr{L}}$

Lemma 2.8. *For any* $\Phi, \mathscr{L}$ *and any* $\bar{\mathscr{L}} \supseteq \mathscr{L}$ *that satisfies* ***PULP1–PULP2*** *we have* $Z(\Phi) \leq 2^{|\bar{\mathscr{L}}|} Z(\Phi, \mathscr{L})$.

In order to identify a 'small' closure $\bar{\mathscr{L}}$ the PULP algorithm resorts to *pure literal elimination*, a simple trick commonplace to satisfiability algorithms. A variable x is *pure* in a CNF formula Φ if $\mathrm{sign}(x, a) = \mathrm{sign}(x, b)$ for any two clauses $a, b \in \partial x$. Clearly, if our objective is to construct a satisfying assignment, we might as well set all pure variables x to the value that satisfies all clauses $a \in \partial x$ and disregard these clauses henceforth. In light of this observation, pure literal elimination repeatedly removes all clauses that contain a pure variable. Naturally, every round of clause removals may create new pure variables, and thus more clauses may be ripe for removal in the next round. For a clause a of the original formula Φ let $\mathfrak{h}_a(\Phi) \geq 1$ be the number of the round at which pure literal elimination removes a. If a is never removed then we set $\mathfrak{h}_a(\Phi) = \infty$.

The PULP algorithm invokes a slightly modified version of pure literal elimination to accommodate the initial set $\mathscr{L}$ of literals. Specifically, for a variable x of a CNF Φ and $s \in \{\pm 1\}$ let $\Phi[x \mapsto s]$ be the CNF obtained by removing all clauses $a \in \partial x$ with $\mathrm{sign}(x, a) = s$ and removing the literal $-s \cdot x$ from all $a \in \partial x$ with $\mathrm{sign}(x, a) = -s$. The definition reflects that if we set x to value s, all $a \in \partial^s x$ will be satisfied, while all $a \in \partial^{-s} x$ will have to be satisfied by one of their other constituent literals. Further, let

$$\mathfrak{h}_x(s, \Phi) = \begin{cases} 0 & \text{if } \partial_\Phi^{-s} x = \emptyset, \\ \max\left\{\mathfrak{h}_a(\Phi[x \mapsto s]) : a \in \partial_\Phi^{-s} x\right\} & \text{otherwise.} \end{cases} \quad \in [0, \infty]. \tag{2.6}$$

We refer to $\mathfrak{h}_x(s, \Phi)$ as the *height of literal* $s \cdot x$ *in* Φ.

The PULP algorithm, displayed as Algorithm 1, harnesses the heights as follows. In its attempt to precipitate **PULP1** and **PULP2** the algorithm iteratively enhances the set $\mathscr{L}$ of literals deemed to be 'true'. For any clause a that violates **PULP1** and that contains a literal $l \notin \neg\mathscr{L}$ the algorithm adds one such literal l *of minimum height* to $\mathscr{L}$. This choice is intended to keep the ultimate size of the closure small; one could say that PULP uses height as a proxy of 'size'. If at any point the algorithm encounters a clause a that consists of literals from $\neg\mathscr{L}$ only, the algorithm reports a contradiction and aborts.

Remark 2.9. To break ties that may occur in the execution of Steps 3 and 7 of PULP we assume that the variables and clauses of Φ are numbered so that Steps 3 and 7 can choose the clause/variable with the smallest number that satisfies the respective requirements. In due course we will run PULP on (finite subtrees of) the Galton-Watson tree $\mathbb{T}$. To number the variables and clauses of $\mathbb{T}$ we equip each of them with an independent Gaussian label. Since $\mathbb{T}$ comprises a countable number of clauses/variables, these labels will almost surely be pairwise distinct.

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

From here on we write $\bar{\mathscr{L}}$ for the set of literals returned by PULP if the algorithm does not encounter a contradiction; in the event of a a contradiction we let $\bar{\mathscr{L}} = \{x, \neg x : x \in V(\Phi)\}$ be the set of all literals. Where the reference to the formula Φ is not entirely obvious, we write $\bar{\mathscr{L}}_\Phi$. The analysis of PULP on the random formula $\boldsymbol{\Phi}'$ furnishes the following bound on $|\bar{\mathscr{L}}|$ in terms of the size of the initial set $\mathscr{L}$. This bound is the key ingredient towards the proof of Proposition 2.7.

Lemma 2.10. *There exists $C = C(d,k) > 0$ such that the following is true. Let $\mathscr{L}$ be a set of literals of $\boldsymbol{\Phi}'$ such that $1 \le |\mathscr{L}| \le \log^2 n$ and such that $\{x_i, \neg x_i\} \not\subseteq \mathscr{L}$ for all $1 \le i \le n$. Then $\mathbb{E}[|\bar{\mathscr{L}}|^{3/2}] \le C|\mathscr{L}|^{3/2}$.*

The proof of Lemma 2.10 is one of the main technical challenges of the present work. The difficulty stems from the stochastic dependencies that are inherent to the PULP algorithm. Specifically, in order to decide which literals to add to the set $\mathscr{L}$, PULP requires knowledge of the heights $\mathfrak{h}_x(\pm 1, \boldsymbol{\Phi}')$. But these heights depend on the other variables $y \in \partial a \setminus \{x\}$, the clauses that these variables y appear in, etc. Furthermore, in its subsequent iterations the algorithm is apt to revisit some of these variables and clauses at a point when their heights have already been revealed. These repetitions rule out an analysis of PULP by way of routine techniques such as the principle of deferred decisions or the differential equations method. The reason why we manage to cope with these complicated dependencies at all is that, remarkably, the heights $\mathfrak{h}_x(\pm 1, \boldsymbol{\Phi}')$ have only a tiny upper tail. More precisely, as we will see the tails of these random variables decay at a *doubly* exponential rate.

Proposition 2.7 follows from the analysis of PULP. The basic idea is to apply the algorithm to an initial set $\mathscr{L}$ of literals that contain one literal from each of the extra clauses that are present in $\boldsymbol{\Phi}''$ or $\boldsymbol{\Phi}'''$ but not in $\boldsymbol{\Phi}'$. With a bit of care the bounds from Lemmas 2.8 and 2.10 then imply (2.5). Finally, the analysis of PULP that leads up to the proof of Lemma 2.10 also implies the necessary tail bounds to verify the bounds from (2.1). Specifically, the proof of Lemma 2.10 proceeds by way of analysing PULP on the Galton-Watson tree $\mathbb{T}_{d,k}$, and the bounds (2.1) come out as a byproduct of that analysis.

2.4 Completing the Aizenman-Sims-Starr scheme

To obtain Propositions 2.5– 2.6 we combine Proposition 2.7 with an analysis of the quotients $Z(\boldsymbol{\Phi}'')/Z(\boldsymbol{\Phi}')$ and $Z(\boldsymbol{\Phi}''')/Z(\boldsymbol{\Phi}')$ on a likely 'good' event. On this good event the empirical distribution of the marginal probabilities $(\mathbb{P}[\boldsymbol{\sigma}_{\boldsymbol{\Phi}'}(x_i) = 1 \mid \boldsymbol{\Phi}'])_{1 \le i \le n}$ of the different variables x_i receiving the value 'true' under a random satisfying assignment is 'close' to the limiting distribution $\pi_{d,k}$ from Proposition 2.1. Additionally, on the good event the joint distribution of the truth values assigned to a moderate number of variables is well approximated by a product measure. Of course, to make this precise we need to investigate the empirical distribution

$$\boldsymbol{\pi}'_n = \frac{1}{n}\sum_{i=1}^{n} \delta_{\mathbb{P}[\boldsymbol{\sigma}_{\boldsymbol{\Phi}'}(x_i)=1 \mid \boldsymbol{\Phi}']} \in \mathscr{P}(0,1) \tag{2.7}$$

of the marginals $(\mathbb{P}[\boldsymbol{\sigma}_{\boldsymbol{\Phi}'}(x_i) = 1 \mid \boldsymbol{\Phi}'])_{1 \le i \le n}$.

Proposition 2.11. *Assume that $d < d_{\mathrm{uniq}}(k)$. Then $\mathbb{E}\left[W_1(\boldsymbol{\pi}'_n, \pi_{d,k})\right] = o(1)$ and for any $\ell = O(1)$ we have*

$$\sum_{\sigma \in \{\pm 1\}^\ell} \mathbb{E}\left|\mathbb{P}\left[\forall 1 \le i \le \ell : \boldsymbol{\sigma}_{\boldsymbol{\Phi}'}(x_i) = \sigma_i \mid \boldsymbol{\Phi}'\right] - \prod_{i=1}^{\ell} \mathbb{P}\left[\boldsymbol{\sigma}_{\boldsymbol{\Phi}'}(x_i) = \sigma_i \mid \boldsymbol{\Phi}'\right]\right| = o(1).$$

The proof of Proposition 2.11 hinges on the Gibbs uniqueness property and the convergence of the local topology of the random formula $\boldsymbol{\Phi}'$ to the Galton-Watson tree $\mathbb{T}_{d,k}$. Together with careful coupling arguments Propositions 2.7– 2.11 imply Propositions 2.5– 2.6. Moreover, in combination

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

with Fact 2.4 these two propositions yield Proposition 2.3. We complete this paragraph by showing how Theorem 1.1 follows from Corollary 2.2 and Proposition 2.3.

Proof of Theorem 1.1. The existence of the limit (1.6) follows from Proposition 2.1. With respect to (1.7), we apply Proposition 2.3 to obtain

$$\frac{1}{n}\mathbb{E}\left[\log\left(1 \vee Z(\mathbf{\Phi}_{d,k}(n))\right)\right] = \frac{1}{n}\sum_{N=0}^{n-1}\left(\mathbb{E}\left[\log\left(1 \vee Z(\mathbf{\Phi}_{d,k}(N+1))\right)\right] - \mathbb{E}\left[\log\left(1 \vee Z(\mathbf{\Phi}_{d,k}(N))\right)\right]\right)$$
$$= \mathfrak{B}_{d,k}(\pi_{d,k}) + o(1). \tag{2.8}$$

Since, conversely, Corollary 2.2 shows that $\frac{1}{n}\log Z(\mathbf{\Phi}) \leq \mathfrak{B}_{d,k}(\pi_{d,k}) + o(1)$ w.h.p. and since $\log Z(\mathbf{\Phi}) \leq n\log 2$ deterministically, the assertion follows from (2.8). □

2.5 Lower-bounding the uniqueness threshold

The proof of Theorem 1.2 combines three ingredients. From the work [3] on random 2-SAT we borrow the idea of constructing an explicit extremal boundary configuration. In effect, in order to prove Gibbs uniqueness we just have to consider one single boundary configuration, instead of an enormous number of possible configurations τ that grows quickly with the height ℓ as in the original definition (1.1). Second, from the work [57] of Montanari and Shah we borrow the idea of expressly considering the effect of pure literals. As it turns out, without explicit consideration of pure literals it seems difficult to even recover the correct asymptotic order (1.8) of the Gibbs uniqueness threshold. Third, and most importantly, the improvement over the bound from [57] stems from a new subtle coupling argument that we will explain in due course.

2.5.1 The extremal boundary condition

An obvious challenge associated with establishing the Gibbs uniqueness property (1.1) seems to be that we need to estimate the marginal of the root variable given *any* possible boundary condition, i.e., given any assignment of the variables at distance 2ℓ from $\mathfrak{x}$. As we expect to see $(d(k-1))^{\ell}$ variables at distance 2ℓ from $\mathfrak{x}$, we thus face a doubly exponential number $2^{(d(k-1))^{\ell}}$ of possible boundary conditions. But fortunately, following [3] we may confine ourselves to just a single, explicit boundary configuration $\boldsymbol{\tau}^{+}$ that satisfies

$$\mathbb{P}\left[\boldsymbol{\tau}^{(\ell)}(r) = 1 \mid \mathbb{T},\ \forall x \in \partial^{2\ell} r : \boldsymbol{\tau}^{(\ell)}(x) = \boldsymbol{\tau}^{+}(x)\right]$$
$$= \max_{\tau \in S(\mathbb{T}^{(\ell)})} \mathbb{P}\left[\boldsymbol{\tau}^{(\ell)}(r) = 1 \mid \mathbb{T},\ \forall x \in \partial^{2\ell} r : \boldsymbol{\tau}^{(\ell)}(x) = \tau(x)\right]. \tag{2.9}$$

Due to the inherent symmetry of the distribution of $\mathbb{T}$ with respect to the signs of the clauses, towards the proof of (1.1) it is sufficient to show that the difference (2.9) vanishes as $\ell \to \infty$.

The extremal boundary condition can be constructed explicitly. Specifically, given $\mathbb{T}^{(\ell)}$ we construct a satisfying assignment $\boldsymbol{\tau}^{+} \in S(\mathbb{T}^{(\ell)})$ by working our way down the tree $\mathbb{T}^{(\ell)}$. We begin by setting $\boldsymbol{\tau}^{+}(\mathfrak{x}) = 1$. Now suppose that for $q \geq 1$, the values of the variables at distance $2(q-1)$ from $\mathfrak{x}$ have been already determined. Let w be a variable at distance $2q$ from $\mathfrak{x}$ with parent clause a and grandparent variable u. Then we define

$$\boldsymbol{\tau}^{+}(w) = \mathrm{sign}(w,a) \cdot \mathbb{1}\{\mathrm{sign}(u,a) \neq \boldsymbol{\tau}^{+}(u)\} - \mathrm{sign}(w,a) \cdot \mathbb{1}\{\mathrm{sign}(u,a) = \boldsymbol{\tau}^{+}(u)\}. \tag{2.10}$$

The idea behind (2.10) is for $\boldsymbol{\tau}^{+}(w)$ to 'nudge' u towards $\boldsymbol{\tau}^{+}(u)$ by making sure that w satisfies clause a if setting u to $\boldsymbol{\tau}^{+}(u)$ does not, and conversely making sure that w fails to satisfy clause a if setting u to $\boldsymbol{\tau}^{+}(u)$ does. A simple induction on ℓ shows that $\boldsymbol{\tau}^{+}$ is a satisfying assignment for which (2.9) holds.

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

Lemma 2.12. *For any integer $\ell \geq 0$ the assignment $\boldsymbol{\tau}^+$ defined via (2.10) satisfies (2.9).*

Hence, proving Theorem 1.2 reduces to establishing the following.

Proposition 2.13. *For $d < d_{\mathrm{con}}(k)$ we have that*

$$\lim_{\ell\to\infty} \mathbb{E}\big[\mathbb{P}\big[\boldsymbol{\tau}^{(\ell)}(\mathfrak{x}) = 1 \mid \mathbb{T},\ \forall x \in \partial^{2\ell}\mathfrak{x} : \boldsymbol{\tau}^{(\ell)}(x) = \boldsymbol{\tau}^+(x)\big] - \mathbb{P}\big[\boldsymbol{\tau}^{(\ell)}(\mathfrak{x}) = 1 \mid \mathbb{T}\big]\big] = 0. \tag{2.11}$$

The proof of Proposition 2.13 may seem delicate because the boundary condition $\boldsymbol{\tau}^+$ depends on the tree $\mathbb{T}^{(\ell)}$. To sidestep this problem, we generalise another idea from the work [3] on random 2-SAT to $k \geq 3$ by introducing a quantity that allows us to prove (2.13) but that behaves 'Markovian' as we pass up and down the tree. Specifically, for a variable x of $\mathbb{T}^{(\ell)}$ let $\mathbb{T}_x^{(\ell)}$ be the subformula of $\mathbb{T}^{(\ell)}$ comprising x and its progeny. Moreover, for a satisfying assignment $\tau \in S(\mathbb{T}^{(\ell)})$ let

$$S\big(\mathbb{T}_x^{(\ell)}, \tau\big) = \big\{\chi \in S\big(\mathbb{T}_x^{(\ell)}\big) : \forall y \in V\big(\mathbb{T}_x^{(\ell)}\big) \cap \partial_{\mathbb{T}}^{2\ell}\mathfrak{x} : \chi_y = \tau_y\big\}, \qquad Z\big(\mathbb{T}_x^{(\ell)}, \tau\big) = \big|S\big(\mathbb{T}_x^{(\ell)}, \tau\big)\big|.$$

In words, $S(\mathbb{T}_x^{(\ell)}, \tau)$ contains all satisfying assignments of $\mathbb{T}_x^{(\ell)}$ that comply with the boundary condition induced by τ. Additionally, for $t = \pm 1$ let

$$S\big(\mathbb{T}_x^{(\ell)}, \tau, t\big) = \big\{\chi \in S\big(\mathbb{T}_x^{(\ell)}, \tau\big) : \chi_x = t\big\}, \qquad Z\big(\mathbb{T}_x^{(\ell)}, \tau, t\big) = \big|S\big(\mathbb{T}_x^{(\ell)}, \tau, t\big)\big|$$

be the set and number of satisfying assignments of $\mathbb{T}_x^{(\ell)}$ that agree with τ on the boundary and assign value t to x. Finally, let

$$\boldsymbol{\eta}_x^{(\ell)} = \log \frac{Z\big(\mathbb{T}_x^{(\ell)}, \boldsymbol{\tau}^+, \boldsymbol{\tau}^+(x)\big)}{Z\big(\mathbb{T}_x^{(\ell)}, \boldsymbol{\tau}^+, -\boldsymbol{\tau}^+(x)\big)} \in \mathbb{R} \cup \{\pm\infty\} \tag{2.12}$$

be the log-likelihood ratio that gauges how likely a random satisfying assignment $\boldsymbol{\tau}$ of $\mathbb{T}_x^{(\ell)}$ subject to the $\boldsymbol{\tau}^+$-boundary condition is to set x to its designated value $\boldsymbol{\tau}^+(x)$ from (2.10). In terms of (2.12), the proof of Proposition 2.13 comes down to showing that for $d < d_{\mathrm{con}}(k)$,

$$\lim_{\ell\to\infty} \boldsymbol{\eta}_{\mathfrak{x}}^{(\ell)} = \log\left(\frac{\boldsymbol{\mu}_{\pi_{d,k},1,1}}{1 - \boldsymbol{\mu}_{\pi_{d,k},1,1}}\right) \quad \text{in distribution.} \tag{2.13}$$

For a start, the following lemma bounds the tails of $\boldsymbol{\eta}_x^{(\ell)}$ for large enough ℓ and x reasonably close to the root variable $\mathfrak{x}$.

Lemma 2.14. *For every $0 < d < d_{\mathrm{con}}(k)$ there exist $c = c(d,k)$ and a sequence $(\varepsilon_t)_t$ with $\lim_{t\to\infty} \varepsilon_t = 0$ such that for any $t > 0$, $\ell > ct^c$ we have*

$$\mathbb{P}\left[\max_{x \in \partial^{2t}\mathfrak{x}} \left|\boldsymbol{\eta}_x^{(\ell)}\right| \leq 2t^c\right] > 1 - \varepsilon_t. \tag{2.14}$$

The proof of Lemma 2.14 rests on combinatorial arguments reminiscent of the analysis of PULP.

A key feature of the definition (2.12) is that the random variables $\boldsymbol{\eta}_x^{(\ell)}$ exhibit a 'reverse Markovian' behaviour. This is because $\boldsymbol{\eta}_x^{(\ell)}$ depends only on $\boldsymbol{\tau}^+(x)$ and the part $\mathbb{T}_x^{(\ell)}$ of the tree pending on x. Furthermore, because the distribution of the random tree $\mathbb{T}_x^{(\ell)}$ is symmetric with respect to sign flips, even the dependence on the value $\boldsymbol{\tau}^+(x)$ can be eliminated. All we need to keep in mind is that the values $\boldsymbol{\tau}^+(y)$ for $y \in V(\mathbb{T}_x^{(\ell)})$ are constructed from the value $\boldsymbol{\tau}^+(x)$ in accordance with the recurrence (2.10). Thus, by flipping all signs in the tree $\mathbb{T}_x^{(\ell)}$ if necessary, we could assume without loss that $\boldsymbol{\tau}^+(x) = 1$ without changing the distribution of $\boldsymbol{\eta}_x^{(\ell)}$ with respect to the randomness of $\mathbb{T}_x^{(\ell)}$. As a consequence, it is possible to set up a recurrence that expresses

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

the log-likelihood ratios $\boldsymbol{\eta}_x^{(\ell)}$ of variables x at distance q from $\mathfrak{x}$ in terms of the $\boldsymbol{\eta}_y^{(\ell)}$ for y at distance $q+2$ from $\mathfrak{x}$.

Due to the recursive nature of the random tree $\mathbb{T}$, it suffices to set up this recurrence for the root $\mathfrak{x}$ of the tree. In other words, to prove (2.13) we just need a recurrence that expresses the distribution of the random variable $\boldsymbol{\eta}_{\mathfrak{x}}^{(\ell+1)}$ in terms of the law of $\boldsymbol{\eta}_{\mathfrak{x}}^{(\ell)}$ for $\ell \geq 0$. A bit of reflection (see Claim 7.1), reveals that the corresponding distributional operator

$$\mathrm{LL}_{d,k}^{+} : \mathscr{P}((-\infty,\infty]) \to \mathscr{P}((-\infty,\infty]), \quad \rho \mapsto \hat{\rho} = \mathrm{LL}_{d,k}^{+}(\rho)$$

has the following shape. For a distribution $\rho \in \mathscr{P}((-\infty,\infty])$ let $(\boldsymbol{\eta}_{\rho,i,j})_{i,j\geq 1}$ be a family of random variables with distribution ρ. Moreover, let $(s_i)_{i\geq 1}$ be a sequence of uniformly random ± 1-valued random variables and let $\boldsymbol{d} \sim \mathrm{Po}(d)$. All of these random variables are mutually independent. Additionally, for $q \geq 0$ and $z_1, \ldots, z_q \in \mathbb{R} \cup \{\pm\infty\}$ define

$$\Gamma(z_1,\ldots,z_q) = \prod_{i=1}^{q} \frac{1+\tan\ h(z_i/2)}{2}. \tag{2.15}$$

Then $\hat{\rho} = \mathrm{LL}_{d,k}^{+}(\rho)$ is the distribution of the random variable

$$-\sum_{i=1}^{d} s_i \cdot \log\left(1-\Gamma\left(s_i \cdot \left(\boldsymbol{\eta}_{\rho,i,1},\ldots,\boldsymbol{\eta}_{\rho,i,k-1}\right)\right)\right). \tag{2.16}$$

Ultimately we will derive (2.13), and thereby Proposition 2.13, from Lemma 2.14 and a contraction argument. However, this is not quite as straightforward as one might be inclined to expect. Indeed, at first glance, a natural approach to proving (2.13) from Lemma 2.14 seems to be to show that $\mathrm{LL}_{d,k}^{+}$ is a contraction, say, with respect to the W_1-metric. This is indeed carried out in [3] for $k=2$, where it is shown that $\mathrm{LL}_{d,2}^{+}$ contracts for all $0 < d < 2$, i.e., right up to the random 2-SAT satisfiability threshold. However, for $k \geq 3$ we can only show that $\mathrm{LL}_{d,k}^{+}$ contracts for $d < 2/(k-1)$, a value well below $d_{\mathrm{con}}(k)$ and short of the correct asymptotic order (1.10).

2.5.2 Pure and mixed literals

To cover a larger range of d we borrow from [57] the idea of expressly taking into account pure literals. To elaborate, while $\mathrm{LL}_{d,k}^{+}$ describes how the law of $\boldsymbol{\eta}_{\mathfrak{x}}^{(\ell+1)}$ results from that of $\boldsymbol{\eta}_{\mathfrak{x}}^{(\ell)}$, the operator fails to take into account that $\mathfrak{x}$ itself as well as some of the grandchildren of $\mathfrak{x}$ in $\mathbb{T}$ may be pure literals. However, the pure literal property has a marked effect on the log-likelihood ratios. For if, say, $\mathfrak{x}$ only appears positively, then a simple double counting argument shows that $\boldsymbol{\eta}_{\mathfrak{x}}^{(\ell)} \geq 0$ for all ℓ. By extension, pure literals among the grandchildren of $\mathfrak{x}$ have a 'dampening' effect and may thus improve the range of d for which we can establish contraction.

For a variable node x of $\mathbb{T}$, let us denote by $\mathbb{T}_x$ the subtree of $\mathbb{T}$ rooted at x and containing its progeny. Leveraging the above observation, we classify a variable x of $\mathbb{T}$ as $\bullet$, $\oplus$, $\ominus$, or $\circ$, depending on whether x appears both positively and negatively in $\mathbb{T}_x$, only positively, only negatively, or whether x has no children at all, respectively. Furthermore, instead of just tracing the law of $\boldsymbol{\eta}_{\mathfrak{x}}^{(\ell)}$ for $\ell \geq 0$, we study the four separate conditional distributions given the type $\bullet$, $\oplus$, $\ominus$ or $\circ$ of $\mathfrak{x}$. Of course, the distribution of $\boldsymbol{\eta}_{\mathfrak{x}}^{(\ell)}$ given type $\circ$ (i.e., $\mathfrak{x}$ has no children) is just the atom at zero for all ℓ.

To describe the evolution of the other distributions we introduce the operator

$$\mathrm{LL}_{d,k}^{\star} : \mathscr{P}(-\infty,\infty] \times \mathscr{P}(0,+\infty] \times \mathscr{P}(-\infty,0] \to \mathscr{P}(-\infty,\infty] \times \mathscr{P}(0,+\infty] \times \mathscr{P}(-\infty,0],$$
$$\text{with } (\rho_\bullet,\rho_\oplus,\rho_\ominus) \mapsto (\hat{\rho}_\bullet,\hat{\rho}_\oplus,\hat{\rho}_\ominus) = \mathrm{LL}_{d,k}^{\star}(\rho_\bullet,\rho_\oplus,\rho_\ominus) \tag{2.17}$$

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

defined as follows. Let $d_+^\star$, $d_+^{\star\prime}$, $d_-^\star$, $d_-^{\star\prime}$ be Poisson variables with parameter $d/2$, conditioned on being positive. Moreover, let $r_1 = (r_{\bullet,1}, r_{\oplus,1}, r_{\ominus,1}, r_{\circ,1})$, $r_2 = (r_{\circ,2}, r_{\oplus,2}, r_{\ominus,2}, r_{\circ,2}), \ldots$ be multinomial variables with $k-1$ trials and probabilities

$$p_\bullet = (1-e^{-d/2})^2, \quad p_\oplus = p_\ominus = e^{-d/2}(1-e^{-d/2}), \quad p_\circ = e^{-d}. \tag{2.18}$$

For $i, j \geq 1$ let $\boldsymbol{\eta}_{\bullet,i,j}$, $\boldsymbol{\eta}_{\oplus,i,j}$, $\boldsymbol{\eta}_{\ominus,i,j}$ be random variables with law $\rho_\bullet$, $\rho_\oplus$, $\rho_\ominus$, respectively. All of the aforementioned random variables are mutually independent. Further, for a sign $\varepsilon \in \{\pm 1\}$ and a vector $r = (r_\bullet, r_\oplus, r_\ominus, r_\circ)$ of non-negative integers with $r_\bullet + r_\oplus + r_\ominus + r_\circ = k-1$ and $i \geq 0$, $1 \leq j \leq 4$ we let

$$\begin{aligned}\boldsymbol{\Xi}_{i,j}(\varepsilon, r) = 1 - \frac{1}{2^{r_\circ}} \cdot \Gamma(\varepsilon(\boldsymbol{\eta}_{\bullet,4i+j,1}, \ldots, \boldsymbol{\eta}_{\bullet,4i+j,r_\bullet}))\Gamma(\varepsilon(\boldsymbol{\eta}_{\oplus,4i+j,1}, \ldots, \boldsymbol{\eta}_{\oplus,4i+j,r_\oplus})) \\ \Gamma(\varepsilon(\boldsymbol{\eta}_{\ominus,4i+j,1}, \ldots, \boldsymbol{\eta}_{\ominus,4i+j,r_\ominus})).\end{aligned} \tag{2.19}$$

The r.h.s. of (2.19) amounts to rewriting the argument of the logarithm in (2.16) when the number of variables of each type is distributed according to r. Finally, let

$$\boldsymbol{\Xi}_{i,j} = \boldsymbol{\Xi}_{i,j}\left((-1)^{j+1}, r_{4i+j}\right). \tag{2.20}$$

Then the operator (2.17) maps $\rho_\bullet$, $\rho_\oplus$, $\rho_\ominus$ to the distributions $\hat{\rho}_\bullet$, $\hat{\rho}_\oplus$, $\hat{\rho}_\ominus$ of the random variables

$$\hat{\rho}_\bullet \sim -\sum_{i=1}^{d_+^\star} \log\, \boldsymbol{\Xi}_{i,1} + \sum_{i=1}^{d_-^\star} \log\, \boldsymbol{\Xi}_{i,2}, \quad \hat{\rho}_\oplus \sim -\sum_{i=1}^{d_+^{\star\prime}} \log\, \boldsymbol{\Xi}_{i,3}, \quad \hat{\rho}_\ominus \sim +\sum_{i=1}^{d_-^{\star\prime}} \log\, \boldsymbol{\Xi}_{i,4}. \tag{2.21}$$

2.5.3 Coupling and contraction

While Montanari and Shah [57] do not write their proof of the lower bound $d_{\mathrm{MS}}(k) \leq d_{\mathrm{uniq}}(k)$ in the language of distributional recurrences, translating their argument to the current formalism evinces two key differences by comparison to the approach that we are going to take. First, Montanari and Shah establish contraction with respect to messages *from clauses to variables*, instead of messages from variables to clauses as considered here.

While this change of perspective may seem innocuous at first, working with respect to variables provides us with greater control over how the change in log-likelihood ratios propagates. In particular, working with variable-to-clause messages and taking into account the four variable types $\bullet, \oplus, \ominus, \circ$ allows us to optimise the metric with respect to which we establish contraction. Hence, for $t > 0$ we endow the space $\mathscr{P}(-\infty, \infty] \times \mathscr{P}(0, +\infty] \times \mathscr{P}(-\infty, 0]$ with the metric

$$\begin{aligned}\mathrm{dist}_t\left((\rho_\bullet, \rho_\oplus, \rho_\ominus), (\rho'_\bullet, \rho'_\oplus, \rho'_\ominus)\right) = \left(1-e^{-t/2}\right) \cdot W_1\left(\rho_\bullet, \rho'_\bullet\right) + e^{-t/2} \cdot W_1\left(\rho_\oplus, \rho'_\oplus\right) \\ + e^{-t/2} \cdot W_1\left(\rho_\ominus, \rho'_\ominus\right).\end{aligned} \tag{2.22}$$

The following proposition summarises the main step towards the proof of Theorem 1.2.

Proposition 2.15. *For every $d < d_{\mathrm{con}}(k)$, the operator $\mathrm{LL}_{d,k}^\star$ is a contraction with respect to the metric dist_d.*

The second key difference between [57] and the present approach will emerge in the proof of Proposition 2.15 itself. As we are about to see, leveraging the four variable types enables us to carry out a sharper bound on the derivative of our operator $\mathrm{LL}_{d,k}^\star$. This comes in the form of a subtle combinatorial coupling between variable types among clauses with opposite signs. To explain this, we recall that $\mathrm{LL}_{d,k}^\star$ describes how the laws of the log-likelihood ratios $\rho_\bullet$, $\rho_\oplus$, and $\rho_\ominus$, evolve given the corresponding laws of the variables in one generation below. Recall also that we are always considering the positive boundary condition, i.e., the one maximising the value of each log-likelihood ratio.

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

Let us write $\rho = (\rho_\bullet, \rho_\oplus, \rho_\ominus)$, $\rho' = (\rho'_\bullet, \rho'_\oplus, \rho'_\ominus)$, and $\hat{\rho}, \hat{\rho}'$ for their corresponding images under the operator $\mathrm{LL}^\star_{d,k}$. We wish to establish that $\mathrm{dist}_d(\hat{\rho}, \hat{\rho}') < c \cdot \mathrm{dist}_d(\rho, \rho')$, for some constant $c = c(d,k) < 1$. We call a clause a positive if it contains its parent variable as a direct literal; otherwise, we call a negative. The change between the output distributions $\hat{\rho}, \hat{\rho}'$ describing the log-likelihood law of, say, variable $\mathfrak{x}$, comes from two sources: the positive and the negative children of $\mathfrak{x}$. Observe that there is no obvious symmetry between the two, as we have imposed the positive boundary condition, and therefore, the influence of positive clauses is typically more pronounced. In turn, the change caused by each clause can be further attributed to that of the $k-1$ grandchildren variables it features. To be more precise, let us consider the contribution of a single positive clause a. Let us write $r = (r_\bullet, r_\oplus, r_\ominus, r_\circ)$ for the type-distribution of the children variables of a, where r follows the law described in (2.18). Consider also an arbitrary enumeration of the variables of each type $t \in \{\bullet, \oplus, \ominus\}$, and write $\mathscr{D}^t_i(z, r; +1)$ for the magnitude of the partial derivative of the message clause a sends to $\mathfrak{x}$, with respect to the message clause a receives from its i-th variable of type t. Then, the expected contribution of clause a to the distance $\mathrm{dist}(\hat{\rho}, \hat{\rho}')$ is bounded in terms of $\mathscr{D}^t_i(z, r; +)$'s as follows

$$\mathbb{E}\left[\sum_{i=1}^{r_\bullet}\left|\int_{\boldsymbol{\eta}_{\bullet,i}}^{\boldsymbol{\eta}'_{\bullet,i}} \mathscr{D}^\bullet_i(w_i, r; +1)\mathrm{d}w_i\right| + \sum_{j=1}^{r_\oplus}\left|\int_{\boldsymbol{\eta}_{\oplus,j}}^{\boldsymbol{\eta}'_{\oplus,j}} \mathscr{D}^\oplus_j(y_j, r; +1)\mathrm{d}y_j\right| + \sum_{\ell=1}^{r_\ominus}\left|\int_{\boldsymbol{\eta}_{\ominus,\ell}}^{\boldsymbol{\eta}'_{\ominus,\ell}} \mathscr{D}^\ominus_\ell(z_\ell, r; +1)\mathrm{d}z_\ell\right|\right], \tag{2.23}$$

where $\boldsymbol{\eta}_{t,i}, \boldsymbol{\eta}'_{t,i}$ follow the law of ρ_t, ρ'_t, respectively. Expanding the expectation with respect to the type-distribution r, and writing $P(r) = \mathbb{P}[r = r]$, for the probability of a vector $r = (r_\bullet, r_\oplus, r_\ominus, r_\circ)$, we rewrite (2.23) as

$$\sum_r P(r)\left(r_\bullet \cdot \mathbb{E}\left|\int_{\boldsymbol{\eta}_{\bullet,1}}^{\boldsymbol{\eta}'_{\bullet,1}} \mathscr{D}^\bullet_1(z, r; +1)\mathrm{d}z\right| + r_\oplus \cdot \mathbb{E}\left|\int_{\boldsymbol{\eta}_{\oplus,1}}^{\boldsymbol{\eta}'_{\oplus,1}} \mathscr{D}^\oplus_1(z, r; +1)\mathrm{d}z\right| + r_\ominus \cdot \mathbb{E}\left|\int_{\boldsymbol{\eta}_{\ominus,1}}^{\boldsymbol{\eta}'_{\ominus,1}} \mathscr{D}^\ominus_1(z, r; +1)\mathrm{d}z\right|\right). \tag{2.24}$$

The expected contribution of a negative clause is given by an expression similar to (2.24), albeit in terms of $\mathscr{D}^t_1(z, r; -)$, i.e., the partial derivative of the message $a \to \mathfrak{x}$, with respect to the message from a variable of type t to clause a. Specifically, the expected contribution of a negative clause reads:

$$\sum_{r'} P(r')\left(r'_\bullet \cdot \mathbb{E}\left|\int_{\boldsymbol{\eta}_{\bullet,1}}^{\boldsymbol{\eta}'_{\bullet,1}} \mathscr{D}^\bullet_1(z, r'; -1)\mathrm{d}z\right| + r'_\oplus \cdot \mathbb{E}\left|\int_{\boldsymbol{\eta}_{\oplus,1}}^{\boldsymbol{\eta}'_{\oplus,1}} \mathscr{D}^\oplus_1(z, r'; -1)\mathrm{d}z\right| + r'_\ominus \cdot \mathbb{E}\left|\int_{\boldsymbol{\eta}_{\ominus,1}}^{\boldsymbol{\eta}'_{\ominus,1}} \mathscr{D}^\ominus_1(z, r'; -1)\mathrm{d}z\right|\right). \tag{2.25}$$

It is not hard to see that pure literals have a 'dampening' effect on each partial derivative $\mathscr{D}^t_1(z, r; \pm)$. Consider a clause a whose children variables are distributed among the different types according to r. Then each derivative in (2.24)–(2.25), can be bounded in terms of the number of pure literals featured in a, excluding the variable of type t with respect to which the derivative is taken. Notice that the operator $\mathrm{LL}^\star_{d,k}$ effectively incorporates the positive boundary condition by imposing the sign of each variable with respect to its parent clause, a, to be $+$ if a is positive, and $-$ if a is negative. With that in mind, we see that if a is a positive clause, the total number of pure literals it contains is just $r_\ominus + r_\circ$. On the other hand, if a is negative, then the total number of pure literals it contains is $r_\oplus + r_\circ$. Bounding *separately* each derivative $\mathscr{D}^t_i(z, r; \pm 1)$ in (2.24)–(2.25), and invoking the mean value theorem, yields an upper bound on for the contraction constant c.

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

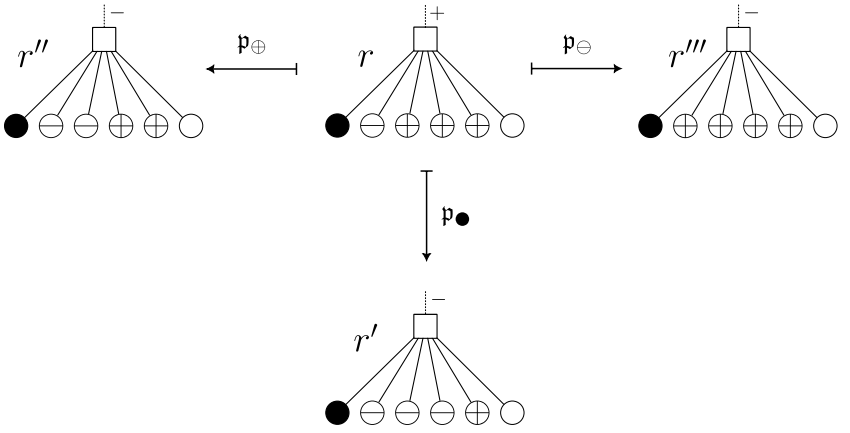


Figure 2. Example of a coupling between derivative terms in (2.24)–(2.25). For vector r and type $t \in \{\bullet, \oplus, \ominus\}$, we pair the term $\mathscr{D}^t(z, r; +1)$ in (2.24) with the term $\mathscr{D}^t(z, \mathfrak{p}_t(r); -1)$ in (2.25).

However, we can do better by partitioning the derivatives in (2.24)–(2.25) into groups, and optimising them *jointly*. Indeed, a careful examination of the expression (2.19), reveals that, for example, any sum of the form $\mathscr{D}^{\oplus}(z, (*, *, r_{\ominus}, r_{\circ}); +1) + \mathscr{D}^{\oplus}(z, (*, r_{\ominus}+1, *, r_{\circ}); -1)$ can be explicitly maximised, and the resulting maximum is smaller than the sum of maxima of the parts. At first sight, this seems to be of little use, if any, as in order to implement such a coupling between terms (2.24)–(2.25), we should also match their coefficients, that is, the quantity $P(r) \cdot r_{\oplus}$, must remain invariant under the coupling. Somewhat unexpectedly, it turns out that the coupling $r \mapsto r'$ with $r' = (r_{\bullet}, r_{\ominus}+1, r_{\oplus}-1, r_{\circ})$, enjoys both features. Similar couplings strategies (depicted in Figure 2) facilitate the maximisation of $\oplus, \ominus$-terms. The full proof of Proposition 2.15 can be found in Section 7.

We conclude the section explaining how Theorem 1.2 follows from the above.

Proof of Theorem 1.2. From the triangle inequality, and Lemma 2.12, it is immediate to obtain Theorem 1.2 from Proposition 2.13. □

2.6 Discussion

The location of the random 2-SAT satisfiability threshold was pinpointed already in the 1990s [21, 41] essentially because the threshold coincides with the giant component phase transition of a directed random graph whose edges correspond to the clauses. This argument also implies that both the pure literal algorithm and another efficient algorithm called unit clause propagation find satisfying assignments up to the satisfiability threshold w.h.p. By contrast, in the case of random k-SAT with $k \geq 3$ the satisfiability threshold is known only for k exceeding an undetermined (but large) constant k_0 [34]. The proof is based on a sophisticated, physics-inspired second moment argument that significantly extends ideas from earlier work [5, 7, 27]. Asymptotically in the limit of large k the satisfiability threshold reads

$$d_{\mathrm{sat}}(k) = k\left[2^k \log 2 - \frac{1+\log 2}{2}\right] + \varepsilon_k, \quad \text{where } \lim_{k\to\infty} \varepsilon_k = 0. \tag{2.26}$$

Even though [5, 7, 27, 34] rely on the second moment method, they do not yield asymptotically tight estimates of the number of satisfying assignments for any regime of d. This is because the second moment method is applied not to the number of satisfying assignments, but to another,

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

exponentially smaller random variable. The assumption that k exceeds a large constant is used critically in [27, 34] to ensure certain concentration and expansion properties.

For $3 \leq k < k_0$ even the *existence* of a uniform satisfiability threshold remains an open problem, although a sharp threshold sequence that may vary with n is known to exist [37]. That said, an upper bound on the satisfiability threshold (sequence) that matches the so-called '1-step replica symmetry breaking' prediction from statistical physics can be verified using the interpolation method from mathematical physics [36, 50, 52, 61]. However, the currently known lower bounds for small k (say, $k = 3, 4, 5$) fall short of this upper bound [7, 43, 47]. For example, in the case $k = 3$ the best current lower bound is $d_{\mathrm{sat}}(3) \geq 10.56$, while $d_{\mathrm{sat}}(3) \approx 12.801$ according to physics predictions [50, 52].

Thus, the satisfiability of random formulas continues to pose a substantial challenge for 'small' $3 \leq k < k_0$. In light of this, a particularly satisfactory aspect of the present results is that they apply and are meaningful for all $k \geq 3$. In fact, comparing the asymptotic bounds (1.10) and (2.26), we see that the Gibbs uniqueness threshold $d_{\mathrm{uniq}}(k)$ is much smaller than $d_{\mathrm{sat}}(k)$ for large k. Thus, Theorems 1.1 and 1.2 cover larger shares of the satisfiable regime of d for smaller values of k; cf. Table 1.

The best current lower bounds on the satisfiability thresholds for $k \geq 4$ are non-constructive. With respect to the algorithmic problem of finding a satisfying assignment of a random k-CNF the best current results for 'small' k are based on simple combinatorial algorithms, analysed via the method of differential equations [38, 43, 47]. Asymptotically for large k the best known efficient algorithm [22] succeeds up to

$$d_{\mathrm{alg}}(k) = (1 - \varepsilon_k)2^k \log k \quad \text{where} \lim_{k\to\infty} \varepsilon_k = 0, \tag{2.27}$$

about a factor of $\log(k)/k$ below (2.26). There is evidence that certain types of algorithms do not succeed for much larger values of d, at least for enough large k [2, 15, 23, 45]. Apart from the task of finding a satisfying assignment, an important line of work deals with the problem of counting and sampling satisfying assignments of random k-CNFs for large k [19, 20, 44]. The best current result [20] covers the regime $d \leq 2^k/k^c$ for an undetermined (large enough) constant $c > 0$. Since for large k the bound $2^k/k^c$ significantly exceeds the pure literal threshold (1.10), it might be an interesting question whether ideas from [19, 20, 44] can be used to verify the replica symmetric solution (1.7) for d beyond the Gibbs uniqueness threshold for large k.

Most of the prior work on the rigorous verification of the replica symmetric solution focuses on a soft version of random k-SAT, the so-called random k-SAT model at *inverse temperature* $\beta > 0$ [58]. The partition function $Z_\beta(\mathbf{\Phi})$ of this model, its the key quantity of interest, is defined as follows. For a clause a of the random formula $\mathbf{\Phi}$ and a truth assignment σ write $\sigma \models a$ if σ satisfies clause a. Then

$$Z_\beta(\mathbf{\Phi}) = \sum_{\sigma\in\{\pm1\}^n} \exp\left(-\beta \sum_{i=1}^{m} \mathbb{1}\{\sigma \not\models a_i\}\right). \tag{2.28}$$

Thus, each assignment σ contributes a summand equal to $\exp(-\beta)$ raised to the power of the number of clauses that σ fails to satisfy. In effect,

$$Z(\mathbf{\Phi}) = \lim_{\beta\to\infty} Z_\beta(\mathbf{\Phi}). \tag{2.29}$$

A line of prior work [13, 60, 63] deals with the derivation of the 'thermodynamic limit'

$$\lim_{n\to\infty} \frac{1}{n}\mathbb{E}\left[\log Z_\beta(\mathbf{\Phi})\right] \tag{2.30}$$

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

for small d and/or small β. Specifically, these works verify that (2.30) is given by the replica symmetric solution at inverse temperature β from [55, 56] under the assumption

$$d(k-1)\min\left\{1, 6\beta\exp(4\beta)\right\} < 1. \tag{2.31}$$

We observe that for large β the bound (2.31) holds only up to the giant component threshold $d = 1/(k-1)$, where the replica symmetric solution trivially follows from the fact that Belief Propagation is exact on acyclic graphical models [51, Theorem 4.1]. That said, a technique called the interpolation method shows that the replica symmetric solution yields an upper bound on (2.30) for all $d, \beta > 0$ [36, 61]. In particular, we will combine the interpolation method with a concentration argument in order to prove Corollary 2.2.

According to physics predictions the 'replica symmetric solution' from [55, 56] yields the correct value of both $\lim_{n\to\infty} n^{-1}\log Z_\beta(\mathbf{\Phi})$ for all $\beta > 0$ and of $\lim_{n\to\infty} n^{-1}\log Z(\mathbf{\Phi})$ for all d up to a threshold $d_{\mathrm{rsb}}(k)$ close to but strictly below the satisfiability threshold $d_{\mathrm{sat}}(k)$ for all $k \geq 3$ [48]. The threshold $d_{\mathrm{rsb}}(k)$ is known as the '1-step replica symmetry breaking phase transition' in physics jargon; its asymptotic value is predicted as

$$d_{\mathrm{rsb}}(k) = k\left[2^k\log 2 - 2\log 2\right] + \varepsilon_k, \quad \text{where } \lim_{k\to\infty}\varepsilon_k = 0. \tag{2.32}$$

Indeed, the interpolation method can be used to verify that the replica symmetric solution ceases to be correct for $d_{\mathrm{rsb}}(k) + \varepsilon_k < d < d_{\mathrm{sat}}(k)$ with $\varepsilon_k \to 0$. Conversely, the replica symmetric solution is known to be correct for all d and $\beta > 0$ where a certain correlation decay condition is satisfied [26], provided that k is large enough. Physics methods predict that this condition holds for all $\beta > 0$ and all $d < d_{\mathrm{rsb}}(k)$ [48].

The aforementioned work of Montanari and Shah [57] also deals with the soft variant of random k-SAT (2.28), but allows for an inverse temperature $\beta = \beta(n)$ that tends to infinity slowly as $n \to \infty$. Specifically, considering a small power $\beta = n^\delta$ enables Montanari and Shah to estimate the number of assignments that satisfy all but $o(n)$ clauses. The proof combines an interpolation on $0 \leq \beta \leq n^\delta$ with a contraction argument that improves over the previous contraction estimates from [13, 60, 63]. Instead of the interpolation on β, in order to prove Theorem 1.1 we use the Aizenman-Sims-Starr scheme. Because we count actual satisfying assignments, this requires the careful combinatorial analysis of tail events, which is where the PULP algorithm and its analysis come in. Additionally, towards the proof of Theorem 1.2 we devise an improved version of the contraction argument from [57]. Following Montanari and Shah, we also take advantage of the impact of pure literals on the Belief Propagation operator. But we develop an improved coupling scheme that yields a better range of d for which contraction occurs. Additionally, once again because we deal with actual satisfying assignments, the proof of the Gibbs uniqueness property involves the analysis of the PULP algorithm on a Galton-Watson tree in order to cope with unlikely events.

By comparison to the 'soft' random k-SAT model (2.28), few prior contributions deal with the actual number $Z(\mathbf{\Phi})$ of satisfying assignments. A result of Abbe and Montanari [1] implies that a deterministic limit (in probability)

$$\lim_{n\to\infty}\frac{1}{n}\log Z(\mathbf{\Phi}) \tag{2.33}$$

exists for Lesbegue-almost all $0 < d < d_{\mathrm{pure}}(k)$ for all $k \geq 2$. However, the proof, which is based on the interpolation method, does not reveal the value of (2.33). In fact, prior to the present work the limit (2.33) was known only in two cases. First, in the trivial regime $d < 1/(k-1)$ below the giant component threshold. Second, in the case $k = 2$ for $0 < d < d_{\mathrm{sat}}(2) = 2$ [3]. In both cases the limit (2.33) coincides with the replica symmetric solution from [55]. Beyond the convergence in probability, $\log Z(\mathbf{\Phi})$ is known to satisfy a central limit theorem in the case $k = 2$ [17].

To compute the limit (2.33) in the case $k = 2$ the contribution [3] employs the Aizenman-Sims-Starr scheme. The couplings that we use towards the proofs of Propositions 2.5– 2.6 generalise the

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

argument from [3] to $k \geq 3$. The main technical novelty lies in the way that moderately unlikely events are treated. Specifically, in the case $k = 2$ the simple Unit Clause propagation algorithm, which essentially boils down to directed reachability, was sufficient to derive a tail bound similar to (and actually stronger than) (2.5). By contrast, since in the case $k \geq 3$ the clauses 'branch out', the analysis of tail events and, accordingly, the derivation of (2.5) is far more delicate. The core of this derivation is the detailed analysis of the PULP algorithm right up to $d_{\mathrm{pure}}(k)$.

Finally, by contrast to random k-SAT the validity of the replica symmetric solution is known for the optimal parameter range for several other random constraint satisfaction problems that enjoy certain symmetry properties. Examples include random graph colouring or random k-NAESAT [11, 25]. Due to the symmetry property[7] the replica symmetric solution simply coincides with the first moment of the number of solutions. In effect, in many symmetric problems it is even possible to precisely determine the limiting distribution of the number of solutions, which superconcentrates on the first moment [24]. By contrast, in random k-SAT the first moment overshoots the typical number of satisfying assignments by an exponential factor [5], which is why random k-SAT is so much more delicate than symmetric problems. That said, there is a regular variant of random k-SAT (where every variable appears an equal number of times positively and negatively) where symmetry and superconcentration are recovered [30].

2.7 Organisation

In the remaining sections we work our way through the proofs of Theorems 1.1 and 1.2. Specifically, in Section 3 we analyse the PULP algorithm introduced in Section 2.3, proving Lemmas 2.8-2.10, which facilitates many of the subsequent results.

In Section 4 we establish Proposition 2.1, verifying that the quantities appearing in Theorem 1.1 are well-defined. The proof of Corollary 2.2 follows in Section 5. Section 6 is devoted to the Aizenman-Sims-Starr scheme, and in particular the proof of Proposition 2.3. There we also complete the proof of Theorem 1.1.

Our final Section 7, deals with the remaining proofs toward establishing Theorem 1.2. We begin by proving Lemma 2.14, showing that the log-likelihood ratios of the random Galton-Watson formula close to the root are bounded w.h.p. This enables us to compare the output distribution of the non-random operator introduced in Section 2.5 with that of actual ratios on the random tree. We then proceed with the proof of Proposition 2.15, and conclude with the proof of (2.13), completing the proof of Theorem 1.2.

3. Analysis of PULP

This section is concerned with the analysis of PULP from Section 2.3. In particular, we prove Lemma 2.10. But let us get the proof of Lemma 2.8 out of the way first.

3.1 Proof of Lemma 2.8

Suppose that $\bar{\mathscr{L}} \supseteq \mathscr{L}$ satisfies **PULP1**–**PULP2**. Let $U = \{|l| : l \in \bar{\mathscr{L}}\}$ be the set of variables underlying the literals $\bar{\mathscr{L}}$. Moreover, let $\chi : U \to \{\pm 1\}$ be the truth assignments under which all literals of $l \in \bar{\mathscr{L}}$ evaluate to 'true'. Due to **PULP2**, the assignment χ is well defined. Moreover, since $\mathscr{L} \subseteq \bar{\mathscr{L}}$, under χ all literals $l \in \mathscr{L}$ evaluate to 'true'. Hence, for a satisfying assignment $\sigma \in S(\Phi)$ define an assignment σ' by letting

$$\sigma'(x) = \mathbb{1}\{x \in U\}\chi(x) + \mathbb{1}\{x \notin U\}\sigma(x).$$

[7] A formal definition of 'symmetry' could be that the uniform distribution on spins is a fixed point of the Belief Propagation operator on the respective Galton-Watson tree.

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

Because $\bar{\mathscr{L}}$ satisfies condition **PULP1**, we have $\sigma' \in S(\Phi, \mathscr{L})$. Finally, because for a satisfying assignment $\tau' \in S(\Phi, \mathscr{L})$ there are no more than $2^{|U|} = 2^{|\bar{\mathscr{L}}|}$ satisfying assignments $\tau \in S(\Phi)$ such that $\tau(x) = \tau'(x)$ for all $x \notin U$, we obtain the desired bound $Z(\Phi) \leq 2^{|\bar{\mathscr{L}}|} Z(\Phi, \mathscr{L})$.

3.2 *Turning a tree to* PULP

While the ultimate goal of this section is to study the PULP algorithm on the random formula $\boldsymbol{\Phi}'$ to prove Lemma 2.10, a necessary preparation is to investigate the algorithm on the random Galton-Watson tree $\mathbb{T} = \mathbb{T}_{d,k}$. Of course, since $\mathbb{T}$ may be infinite we should formally confine ourselves to the finite trees $\mathbb{T}^{(\ell)}$ truncated at the 2ℓ-th level from the root $\mathfrak{x}$. Hence, recalling (2.6), we aim to estimate the height $\mathfrak{h}_{\mathfrak{x}}(s, \mathbb{T}^{(\ell)})$ for finite ℓ. That said, since these random variables are monotonically increasing in ℓ, it makes sense to define

$$\mathfrak{h}_{\mathfrak{x}}(s, \mathbb{T}) = \lim_{\ell \to \infty} \mathfrak{h}_{\mathfrak{x}}(s, \mathbb{T}^{(\ell)}) \in [0, \infty]. \tag{3.1}$$

We point out that for $d < d_{\mathrm{pure}}(k)$ the tails of $\mathfrak{h}_{\mathfrak{x}}(s, \mathbb{T})$ decay at a doubly exponential rate.

Lemma 3.1. *For any $d < d_{\mathrm{pure}}(k)$ there exist $c_1 = c_1(d,k), c_2 = c_2(d,k) > 0$ such that*

$$\mathbb{P}\left[\mathfrak{h}_{\mathfrak{x}}(\pm 1, \mathbb{T}) \geq h\right] \leq c_1 \cdot \exp\left(-\exp\left(c_2 \cdot h\right)\right) \quad \textit{for every } h \geq 1.$$

Proof. By symmetry it suffices to consider $\mathfrak{h}_{\mathfrak{x}}(1, \mathbb{T})$. Thus, let $p_{h,\ell} = \mathbb{P}\left[\mathfrak{h}_{\mathfrak{x}}(1, \mathbb{T}^{(\ell)}) \geq h\right]$. All variables at distance 2ℓ from $\mathfrak{x}$ are leaves and therefore pure in the tree $\mathbb{T}^{(\ell)}$. Consequently, pure literal elimination removes all clauses of $\mathbb{T}^{(\ell)}$ within at most ℓ rounds. Hence, $p_{h,\ell} = 0$ for $h > \ell$. Furthermore, we claim that

$$p_{h,\ell} = \varphi_{d,k}(p_{h-1,\ell-1}), \quad \text{where} \quad \varphi_{d,k}(z) = 1 - \exp\left(-\frac{d}{2} z^{k-1}\right) \qquad (1 \leq h \leq \ell). \tag{3.2}$$

Indeed, if $\mathfrak{h}_{\mathfrak{x}}(1, \mathbb{T}^{(\ell)}) \geq h \geq 1$ then by (2.6) there exists a clause $a \in \partial_{\mathbb{T}} r$ with $\mathrm{sign}(\mathfrak{x}, a) = -1$ such that $\mathfrak{h}_a(\mathbb{T}^{(\ell)}[\mathfrak{x} \mapsto 1]) \geq h - 1$. In other words, pure literal elimination on the sub-tree $\mathbb{T}^{(\ell)}[a]$ of $\mathbb{T}^{(\ell)}$ rooted at clause a and with variable $\mathfrak{x}$ removed takes at least $h-1$ rounds to remove clause a. Consequently, pure literal elimination on $\mathbb{T}^{(\ell)}[a]$ takes at least $h-1$ rounds to remove one of the variables $x \in \partial_{\mathbb{T}} a \setminus \{\mathfrak{x}\}$. In other words, the sub-tree $\mathbb{T}^{(\ell)}[x]$ comprising x and its successors satisfies

$$\mathfrak{h}_x(\mathrm{sign}(x,a), \mathbb{T}^{(\ell)}[x]) \geq h - 1 \quad \text{for every } x \in \partial_{\mathbb{T}} a \setminus \{\mathfrak{x}\}. \tag{3.3}$$

But since $\mathbb{T}$ is a Galton-Watson tree, the sub-tree $\mathbb{T}^{(\ell)}[x]$ has the same distribution as the random tree $\mathbb{T}^{(\ell-1)}$. Hence, (3.3) implies that for every $a \in \partial_{\mathbb{T}} \mathfrak{x}$ with $\mathrm{sign}(x,a) = -1$,

$$\mathbb{P}\left[\mathfrak{h}_a(\mathbb{T}^{(\ell)}[\mathfrak{x} \mapsto 1]) \geq h - 1 \mid \mathbb{T}^{(1)}\right] = p_{h-1,\ell-1}^{k-1}. \tag{3.4}$$

Finally, the construction of $\mathbb{T}$ ensures that the number of $a \in \partial_{\mathbb{T}} \mathfrak{x}$ with $\mathrm{sign}(x,a) = -1$ has distribution $\mathrm{Po}(d/2)$. Therefore, (3.4) shows that

$$p_{h,\ell} = \sum_{i=0}^{\infty} \mathbb{P}\left[\mathrm{Po}(d/2) = i\right] \left(1 - \left(1 - p_{h-1,\ell-1}^{k-1}\right)^i\right) = 1 - \exp\left(-\frac{d}{2} p_{h-1,\ell-1}^{k-1}\right),$$

which completes the proof of (3.2).

Since the sequences $(p_{h,\ell})_\ell$ are non-decreasing, the limits $p_h = \lim_{\ell \to \infty} p_{h,\ell}$ exist. Moreover, (3.2) shows that

$$p_h = \varphi_{d,k}(p_{h-1}) \quad (h \geq 1). \tag{3.5}$$

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

Hence, recalling the definition (1.8) of $d_{\mathrm{pure}}(k)$, we find

$$\left(\frac{p_{h+1}}{p_h}\right)^{k-1} = \left(\frac{\varphi_{d,k}(p_h)}{p_h}\right)^{k-1} = d \cdot \frac{\left(1-\exp\left(-dp_h^{k-1}/2\right)\right)^{k-1}}{dp_h^{k-1}} \leq \frac{d}{d_{\mathrm{pure}}} < 1.$$

Consequently,

$$\lim_{h\to\infty} p_h = 0. \tag{3.6}$$

To complete the proof we expand $\varphi_{d,k}(z)$ around $z = 0$:

$$\varphi_{d,k}(z) = \frac{d}{2}z^{k-1} + O(z^{2k-2}) \quad \text{as } z \to 0. \tag{3.7}$$

Thus, the function $\varphi_{d,k}(z)$ is well approximated by a $(k-1)$-th power. Since $k \geq 3$, combining (3.5)–(3.7), we conclude that for sufficiently large h we have $p_h \leq (d/2+1)p_{h-1}^{k-1}$. Consequently, $p_h \leq c_1 \cdot \exp\left(-\exp\left(c_2 \cdot h\right)\right)$ for suitable $c_1 = c_1(d,k)$ and $c_2 = c_2(d,k)$. □

We remind ourselves that $\overline{\{\pm 1 \cdot \mathfrak{x}\}}_{\mathbb{T}^{(\ell)}}$ signifies the output of PULP run on the formula $\mathbb{T}^{(\ell)}$ with initial literal set $\{\pm 1 \cdot \mathfrak{x}\}$. We extend the definition of the closure to the (possibly infinite) tree $\mathbb{T}$ by letting

$$\overline{\{\pm 1 \cdot \mathfrak{x}\}}_{\mathbb{T}} = \bigcap_{\ell_0 \geq 1} \bigcup_{\ell \geq \ell_0} \overline{\{\pm 1 \cdot \mathfrak{x}\}}_{\mathbb{T}^{(\ell)}}.$$

This definition ensures that if the height $\mathfrak{h}_{\mathfrak{x}}(\pm 1, \mathbb{T})$ from (3.1) is finite, then

$$\overline{\{\pm 1 \cdot \mathfrak{x}\}}_{\mathbb{T}} = \overline{\{\pm 1 \cdot \mathfrak{x}\}}_{\mathbb{T}^{(\ell)}} \quad \text{for all } \ell \geq \mathfrak{h}_{\mathfrak{x}}(\pm 1, \mathbb{T}).$$

In order to estimate the size of this set, we combine Lemma 3.1 with a crude bound on the total number of variable nodes of the Galton-Watson tree $\mathbb{T}^{(\ell)}$. Recall that $V(\mathbb{T}^{(\ell)})$ signifies the set of variable nodes of $\mathbb{T}^{(\ell)}$.

Lemma 3.2. *Let $d > 0$. For any $\ell \geq 1$ and any $t > 100(1+d(k-1))^2$ we have*

$$\mathbb{P}\left[|V(\mathbb{T}^{(\ell)})| > t^{\ell}\right] \leq \ell \exp\left(-t^{\ell/2}/4\right).$$

Proof. Let $N_\ell = |V(\mathbb{T}^{(\ell)})|$ for brevity, set $g = 10(1+d(k-1))$ and notice that $t > g^2$. The construction of the Galton-Watson tree $\mathbb{T}$ ensures that $N_0 = 1$ and that for $\ell \geq 1$ given $N_{\ell-1}$ we have $N_\ell \sim (k-1) \cdot \mathrm{Po}(dN_{\ell-1})$. Therefore, Bennett's inequality shows that

$$\mathbb{P}\left[N_h > g^{h-\ell}t^{\ell} \mid N_{h-1} \leq g^{h-1-\ell}t^{\ell}\right] \leq \exp\left(-\frac{t^{\ell}}{4g^{\ell-h}}\right) \leq \exp\left(-t^{\ell/2}/4\right) \quad (1 \leq h \leq \ell). \tag{3.8}$$

Furthermore, if $N_\ell > t^\ell$ then there exists $1 \leq h \leq \ell$ such that $N_h > g^{h-\ell}t^\ell$ while $N_{h-1} \leq g^{h-1-\ell}t^\ell$. Hence, combining (3.8) with the union bound completes the proof. □

Corollary 3.3. *For any $d < d_{\mathrm{pure}}(k)$ there exists $c_3 = c_3(d,k) > 0$ such that*

$$\mathbb{P}[|\overline{\{\pm 1 \cdot \mathfrak{x}\}}_{\mathbb{T}}| > t] \leq c_3 \exp\left(-t^{1/c_3}\right) \quad \textit{for all } t > 0.$$

Proof. By symmetry it suffices to consider $N = |\overline{\{\mathfrak{x}\}}_{\mathbb{T}}|$. Since Lemma 3.1 shows that $\mathbb{P}[\mathfrak{h}_{\mathfrak{x}}(1,\mathbb{T}) < \infty] = 1$, we may assume from now on that indeed $\mathfrak{h}_{\mathfrak{x}}(1,\mathbb{T}) < \infty$. Moreover, picking $c_3 = c_3(d,k) > 0$ large enough, we may assume that $t > t_0$ for a large $t_0 = t_0(d,k)$. Let $N_\ell = |V(\mathbb{T}^{(\ell)})|$,

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

$p_h = \mathbb{P}\left[\mathfrak{h}_{\mathfrak{x}}(1,\mathbb{T}) = h\right]$ and $g = 10(1+d(k-1))$. It is an immediate consequence of the way that PULP proceeds that for all $l \in \overline{\{\mathfrak{x}\}}_{\mathbb{T}}$ we have $|l| \in V(\mathbb{T}^{(\mathfrak{h}_{\mathfrak{x}}(1,\mathbb{T}))})$. Hence, $N \le N_{\mathfrak{h}_{\mathfrak{x}}(1,\mathbb{T})}$. Therefore, by the law of total probability,

$$\mathbb{P}[N > t] \le \sum_{h \ge 0} S_h, \quad \text{where } S_h = \mathbb{P}\left[\mathfrak{h}_{\mathfrak{x}}(1,\mathbb{T}) = h\right]\mathbb{P}\left[N_h > t \mid \mathfrak{h}_{\mathfrak{x}}(1,\mathbb{T}) = h\right]. \tag{3.9}$$

Depending on the value of t in relation to h, we use either Lemma 3.1 or Lemma 3.2 to bound S_h.

Case 1: $t_0 < t \le g^{2h}$: Lemma 3.1 shows that for certain $c_1, c_2 > 0$ we have

$$S_h \le \mathbb{P}\left[\mathfrak{h}_{\mathfrak{x}}(1,\mathbb{T}) = h\right] \le c_1 \exp\left(-\exp\left(c_2 h\right)\right) \le c_1 2^{-h} \exp\left(-t^{1/c_3}\right), \tag{3.10}$$

provided c_3 is chosen large enough.

Case 2: $t > g^{2h}$: we apply Lemma 3.2 to obtain

$$S_h \le \mathbb{P}[N_h > t] \le h \exp\left(-\sqrt{t}/4\right) \le h 2^{-h} \exp\left(-t^{1/3}\right), \tag{3.11}$$

provided that $t > t_0$ is sufficiently large.

Combining the bounds (3.9)–(3.11) completes the proof. □

Corollary 3.4. *For any $d < d_{\mathrm{pure}}(k)$ we have $\mathbb{E}[|\overline{\{\pm 1 \cdot \mathfrak{x}\}}_{\mathbb{T}}|^2] < \infty$.*

Proof. This is an immediate consequence of Corollary 3.3. □

3.3 Proof of Lemma 2.10

Because the distribution of $\boldsymbol{\Phi}'$ is invariant under variable permutations and inversions, we may assume the initial set $\mathscr{L}$ of literals passed to PULP is just $\mathscr{L} = \{x_1, \ldots, x_L\}$ for an integer $L = \tilde{O}(1)$. For an integer $\ell \ge 1$ let $\boldsymbol{\phi}'_{\ell,L}$ be the sub-formula of $\boldsymbol{\Phi}'$ comprising all clauses and variables at distance at most 2ℓ from $\mathscr{L}$. We recall that this formula has a bipartite graph representation $G(\boldsymbol{\phi}'_{\ell,L})$ with variable nodes $V(\boldsymbol{\phi}'_{\ell,L})$, clause nodes $F(\boldsymbol{\phi}'_{\ell,L})$ and edges $E(\boldsymbol{\phi}'_{\ell,L})$. The *excess* of $\boldsymbol{\phi}'_{\ell,L}$ is defined as

$$X_{\ell,L} = |E(\boldsymbol{\phi}'_{\ell,L})| - |V(\boldsymbol{\phi}'_{\ell,L})| - |F(\boldsymbol{\phi}'_{\ell,L})|.$$

Thus, $X_{\ell,L} = -L$ iff $G(\boldsymbol{\phi}'_{\ell,L})$ consists of L acyclic components.

Lemma 3.5. *Let $d > 0$, $c > 0$ and assume that $L \le \log^c n$ and $\ell \le c \log\log n$. Then*

$$\mathbb{P}\left[X_{\ell,L} > -L\right] = \tilde{O}(n^{-1}), \qquad \mathbb{P}\left[X_{\ell,L} > 1 - L\right] = \tilde{O}(n^{-2}). \tag{3.12}$$

Furthermore, there exists $c_4 = c_4(c,d,k) > 0$ such that

$$\mathbb{P}\left[|V(\boldsymbol{\phi}'_{\ell,L})| + |F(\boldsymbol{\phi}'_{\ell,L})| > \log^{c_4} n\right] = O(n^{-2}). \tag{3.13}$$

Proof. We study breadth first search ('BFS') on the graph $G(\boldsymbol{\Phi}')$ from the start vertices $\mathscr{L}$ by means of a routine deferred decisions argument. Throughout the execution of BFS each variable node is in one of three possible states: unexplored, active, or finished.

Towards the proof of (3.13) we study a 'parallel' version of BFS. More precisely, let $\mathscr{A}_0 = \mathscr{L}$ be the set of initially active variables, let $\mathscr{U}_0 = \{x_1, \ldots, x_n\} \setminus \mathscr{L}$ comprise the initially unexplored variables and let $\mathscr{F}_0 = \emptyset$. Further, for $t \ge 0$ define $\mathscr{A}_{t+1}, \mathscr{U}_{t+1}, \mathscr{F}_{t+1}$ as follows. If $\mathscr{A}_t = \emptyset$ then the process has stopped and we let $\mathscr{A}_{t+1} = \mathscr{A}_t = \emptyset$, $\mathscr{U}_{t+1} = \mathscr{U}_t$, $\mathscr{F}_{t+1} = \mathscr{F}_t$. Otherwise let $\mathscr{A}_{t+1}$ be the set of all variable nodes $y \in \mathscr{U}_t$ such that there exist an active variable node $x \in \mathscr{A}_t$ and a clause a that contains x and y; in symbols, $x, y \in \partial_{\boldsymbol{\Phi}'} a$. Further, let $\mathscr{F}_{t+1} = \mathscr{F}_t \cup \mathscr{A}_t$ and $\mathscr{U}_{t+1} = \mathscr{U}_t \setminus \mathscr{A}_{t+1}$. The BFS exploration occurs 'in parallel' in the sense that all active vertices activate their previously unexplored second neighbours simultaneously.

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

Let $\mathfrak{F}_t$ be the σ-algebra generated by the first t rounds of parallel exploring. Then the distribution of $|\mathscr{A}_{t+1}|$ given $\mathfrak{F}_t$ is stochastically dominated by a random variable with distribution $(k-1)\mathrm{Po}(d|\mathscr{A}_t|)$. This is because by the construction of the formula $\boldsymbol{\Phi}'$ the total number of clauses containing a given variable node has distribution $\mathrm{Po}(d(1-(k-1)/n))$. Hence, for any $u>0$ we have

$$\mathbb{P}\left[|\mathscr{A}_{t+1}| > u \mid \mathfrak{F}_t\right] \le \mathbb{P}\left[(k-1)\mathrm{Po}(d|\mathscr{A}_t|) > u\right]. \tag{3.14}$$

To complete the proof of (3.13) we mimic the argument from the proof of Lemma 3.2. Thus, let $u = \log^{c_4-3} n$ for a large enough $c_4 = c_4(c,d,k)$ and set $g = 10(1+d(k-1))$. Since $\ell \le c\log\log n$, the bound (3.14) and Bennett's inequality show that

$$\mathbb{P}\left[|\mathscr{A}_{t+1}| > g^{t+1-\ell}u \mid |\mathscr{A}_t| \le g^{t-\ell}u\right] \le \exp\left(-\frac{u}{4g^{\ell-t+1}}\right) \le \exp\left(-\sqrt{u}/4\right) = O(n^{-3}) \quad (0 \le t < \ell).$$

Hence, taking a union bound on $0 \le t < \ell$ and observing that $|V(\boldsymbol{\phi}'_{\ell,L})| \subseteq \mathscr{A}_0 \cup \cdots \cup \mathscr{A}_\ell$, we obtain

$$\mathbb{P}\left[|V(\boldsymbol{\phi}'_{\ell,L})| > u\log n\right] = \tilde{O}(n^{-3}). \tag{3.15}$$

Finally, another application of Bennett's inequality demonstrates that with probability $1-O(n^{-2})$ no variable of $\boldsymbol{\Phi}'$ appears in more than $\log n$ clauses. Thus, $|F(\boldsymbol{\phi}'_{\ell,L})| \le |V(\boldsymbol{\phi}'_{\ell,L})|\log n$. Hence, (3.15) implies (3.13).

We are left to establish (3.12). The way we set up the BFS process implies that there are only two ways in which excess edges can come about. First, there may be clauses a with $\partial_{\boldsymbol{\Phi}'}a \subseteq \mathscr{A}_t \cup \mathscr{A}_{t+1}$ such that $|\partial_{\boldsymbol{\Phi}'}a \cap \mathscr{A}_t| \ge 2$. Given that $|\mathscr{A}_t \cup \mathscr{A}_{t+1}| \le \log^{c_4} n$, the number of such a with $|\partial_{\boldsymbol{\Phi}'}a \cap \mathscr{A}_t| = 2$ has distribution $\mathrm{Po}(\tilde{O}(1/n))$, and the number of a with $|\partial_{\boldsymbol{\Phi}'}a \cap \mathscr{A}_t| > 2$ has distribution $\mathrm{Po}(\tilde{O}(1/n^2))$. The second possibility is that for a variable $x \in \mathscr{A}_{t+1}$ there exist clauses $a, b \in \partial_{\boldsymbol{\Phi}'}x$ with $\partial_{\boldsymbol{\Phi}'}a, \partial_{\boldsymbol{\Phi}'}b \subseteq \mathscr{A}_t \cup \mathscr{A}_{t+1}$. Once again the number of such clauses has distribution $\mathrm{Po}(\tilde{O}(1/n))$ given $|\mathscr{A}_t \cup \mathscr{A}_{t+1}| \le \log^{c_4} n$. Furthermore, excess inducing clauses occur independently at different rounds t of the BFS process. Thus, (3.12) follows from (3.13). □

We proceed to derive bounds on $|\bar{\mathscr{Z}}| = |\mathscr{Z}_{\boldsymbol{\Phi}'}|$ depending on the value of the excess. To deal with the case of excess $-L$, let $\Lambda = \Theta(\log\log n)$ and let $(\mathbb{T}[i])_{i\ge1}$ be a sequence of independent copies of the random tree $\mathbb{T}$. In the case that the excess $X_{\Lambda,L}$ equals $-L$, the bound on $|\bar{\mathscr{Z}}|$ follows from the fact that the Galton-Watson tree $\mathbb{T}$ captures the local structure of the graph $G(\boldsymbol{\Phi}')$ in combination with the bound from Corollary 3.3. More precisely, the following is true.

Lemma 3.6. *For any $0 < d < d_{\mathrm{uniq}}(k)$ and $c > 0$ there exists $\zeta = \zeta(c,d,k) > 0$ such that with $\Lambda = \lceil \zeta \log\log n \rceil$ uniformly for all $1 \le L \le \log^c n$ and all $u > 0$ we have*

$$\mathbb{P}\left[\mathbb{1}\{X_{\Lambda,L} = -L\}|\bar{\mathscr{Z}}| > u\right] \le \mathbb{P}\left[\sum_{i=1}^{L} |\overline{\{\mathfrak{x}\}}_{\mathbb{T}[i]}| > u\right] + O(n^{-2}).$$

Proof. We begin by coupling the random formula $\boldsymbol{\phi}'_{\ell,1}$ with the Galton-Watson tree $\mathbb{T}^{(\ell)}[1]$ for $0 \le \ell \le \Lambda$. The coupling operates in accordance with the iterations of the BFS process from the proof of Lemma 3.5. Under the coupling some of the variable and clause nodes of $\boldsymbol{\phi}'_{\ell,1}$ and of the tree $\mathbb{T}^{(\ell)}[1]$ are identical, but both $\mathbb{T}^{(\ell)}[1]$ and $\boldsymbol{\phi}'_{\ell,1}$ may contain additional clauses or variables. These additional clauses/variables result from excess edges of $G(\boldsymbol{\phi}'_{\ell,1})$, i.e., edges that close cycles or merge different components in the course of the BFS process.

For $\ell = 0$ we just identify the start variable x_1 with the root $\mathfrak{x}$ of the Galton-Watson tree $\mathbb{T}[1]$. Going from ℓ to $\ell+1$, we remember the sets $\mathscr{A}_\ell, \mathscr{A}_{\ell+1}$ from the proof of Lemma 3.5. For each variable $x \in \mathscr{A}_\ell$ let $\mathscr{C}_x$ be the set of clauses $a \in \partial_{\boldsymbol{\Phi}'}x$ such that $|\partial_{\boldsymbol{\Phi}'}a \cap \mathscr{A}_{\ell+1}| = k-1$ and also such that none of the variables $y \in \mathscr{A}_{\ell+1} \cap \partial_{\boldsymbol{\Phi}'}a$ appear in another clause $b \ne a$ with $\partial_{\boldsymbol{\Phi}'}b \subseteq \mathscr{A}_\ell \cup \mathscr{A}_{\ell+1}$.

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

In other words, $\mathscr{C}_x$ contains all clauses $a \in \partial_{\mathbf{\Phi}'} x$ that do not induce excess edges. Let $d_x = |\mathscr{C}_x|$ be the number of such clauses. As we pointed out in the proof of Lemma 3.5, d_x is stochastically dominated by a $\mathrm{Po}(d)$ variable. Hence, there is a random variable d'_x such that $d_x + d'_x \sim \mathrm{Po}(d)$.

For any variable $x \in \mathscr{A}_\ell$ that is also a variable node of $\mathbb{T}^{(\ell)}[1]$ we add all clauses $a \in \mathscr{C}_x$ and the $k-1$ variables $y \in \partial_{\mathbf{\Phi}'} a \cap \mathscr{A}_{\ell+1}$ to $\mathbb{T}^{(\ell+1)}[1]$. Additionally, $\mathbb{T}^{(\ell+1)}[1]$ contains d'_x independent random clauses that contain x and $k-1$ new variable nodes without a counterpart in $\boldsymbol{\phi}'_{\ell+1,1}$. Finally, to complete $\mathbb{T}^{(\ell+1)}[1]$ every variable y of $\mathbb{T}^{(\ell)}[1]$ at distance precisely 2ℓ from r such that $y \notin V(\boldsymbol{\phi}'_{\ell,1})$ independently begets $\mathrm{Po}(d)$ offspring clause nodes, each containing $k-1$ new variable nodes that do not belong to $V(\boldsymbol{\phi}'_{\ell+1,1})$.

The coupling ensures that $\boldsymbol{\phi}'_{\Lambda,1}$ is a sub-formula of $\mathbb{T}^{(\Lambda)}[1]$ unless $X_{\Lambda,1} > -L$. The extension of this coupling to $\mathscr{L} = \{x_1, \ldots, x_L\}$ is straightforward. We simply perform BFS exploration from the start variables $x_1, \ldots, x_L$ one after the other. Given that $X_{\Lambda,L} = -L$, we thus couple the sub-formula of $\mathbf{\Phi}'$ explored from each x_i with $\mathbb{T}^{(\Lambda)}[i]$ for $1 \le i \le L$ such that $\boldsymbol{\phi}'_{\Lambda,L}$ is contained in the union of $\mathbb{T}^{(\Lambda)}[1], \ldots, \mathbb{T}^{(\Lambda)}[L]$. Finally, we obtain independent copies $\mathbb{T}[1], \ldots, \mathbb{T}[L]$ of the (possibly infinite) tree $\mathbb{T}$ by continuing the Galton-Watson processes $\mathbb{T}^{(\Lambda)}[i]$ independently for depths $\ell > \Lambda$.

The remaining task is to compare $|\bar{\mathscr{L}}|$ with $\sum_{i=1}^{L} |\overline{\{\mathfrak{x}\}}_{\mathbb{T}[i]}|$. If $X_{\Lambda,L} = -L$ and if $\mathfrak{h}_{\mathfrak{x}}(1, \mathbb{T}[i]) < \Lambda$ for all $1 \le i \le L$, then the coupling ensures that all clauses and variables of $\boldsymbol{\phi}'_{\Lambda,L}$ are contained in the disjoint union of the trees $\mathbb{T}[1], \ldots, \mathbb{T}[L]$, and thus $|\bar{\mathscr{L}}| \le \sum_{i=1}^{L} |\overline{\{\mathfrak{x}\}}_{\mathbb{T}[i]}|$. Therefore, for any $u > 0$ we have

$$\mathbb{P}\left[\mathbb{1}\left\{X_{\Lambda,L} = -L, \max_{1\le i\le L} \mathfrak{h}_{\mathfrak{x}}(1, \mathbb{T}[i]) < \Lambda\right\} |\bar{\mathscr{L}}| > u\right] \le \mathbb{P}\left[\sum_{i=1}^{L} |\overline{\{\mathfrak{x}\}}_{\mathbb{T}[i]}| > u\right]. \tag{3.16}$$

Furthermore, since $\Lambda \ge \zeta \log\log n$ for a large $\zeta > 0$, Lemma 3.1 ensures that

$$\mathbb{P}\left[\max_{1\le i\le L} \mathfrak{h}_{\mathfrak{x}}(1, \mathbb{T}[i]) \ge \Lambda\right] \le O(n^{-2}). \tag{3.17}$$

Combining (3.16) and (3.17) completes the proof. □

For later reference we make a note of the following immediate consequence of the coupling from the proof of Lemma 3.6. For two rooted Boolean formulas ϕ, ϕ' we write $\phi \cong \phi'$ if there is an isomorphism of ϕ and ϕ' that preserves the root variable. We consider the random formula $\boldsymbol{\phi}'_{\ell,1}$ rooted at x_1.

Corollary 3.7. *For every $\ell \ge 0$ and any fixed tree T we have* $\left|\mathbb{P}\left[\mathbb{T}^{(\ell)} \cong T\right] - \mathbb{P}\left[\boldsymbol{\phi}'_{\ell,1} \cong T\right]\right| = o(1)$.

From here on, we set $\Lambda = \lceil c_5 \log\log n \rceil$ for a large enough $c_5 = c_5(d,k) > 0$. We obtain the following bound on the second moment of $|\bar{\mathscr{L}}|$ on the event that the excess equals $-L$.

Corollary 3.8. *For any $0 < d < d_{\mathrm{uniq}}(k)$ and any $1 \le L \le \log^2 n$ we have* $\mathbb{E}[\mathbb{1}\{X_{\Lambda,L} = -L\} \cdot |\bar{\mathscr{L}}|^2] = O(1)$.

Proof. Since $|\bar{\mathscr{L}}| \le 2n$ deterministically, this is an immediate consequence of Corollary 3.3 and Lemma 3.6. □

As a next step we deal with the case that the excess equals $1 - L$. More precisely, with $c_6 = c_6(d,k) \gg c_5$ a large enough constant let $\Lambda^+ = \lceil c_6 \log\log n \rceil$. We are going to bound $|\bar{\mathscr{L}}|$ on the event that $X_{\Lambda,L} = X_{\Lambda^+,L} = 1 - L$. The proof combines the bound on the probability of this event from Lemma 3.5 with a crude bound on $|\bar{\mathscr{L}}|$. To elaborate, since Lemma 3.5 shows that the event $X_{\Lambda,L} = X_{\Lambda^+,L} = 1 - L$ has probability $\tilde{O}(n^{-1})$, we can essentially get away with simply bounding $|\bar{\mathscr{L}}|$ by the total number of variables within a $2\Lambda^+$ radius around the start variables $\mathscr{L}$.

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

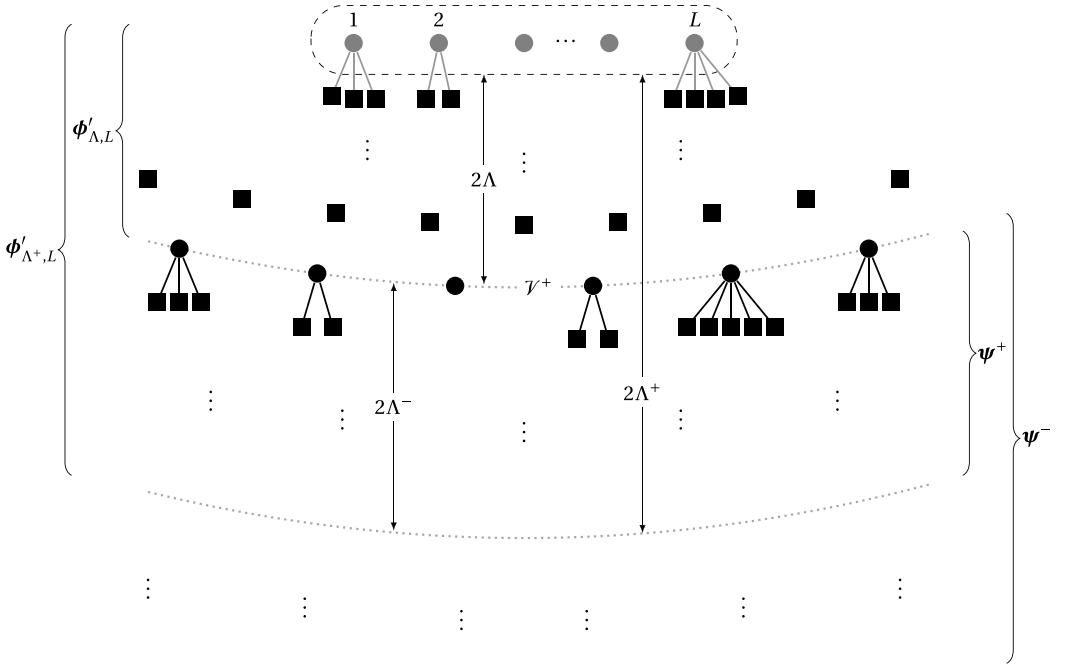


Figure 3. A sketch depicting the subformulas $\boldsymbol{\psi}^+$, $\boldsymbol{\psi}^-$, $\boldsymbol{\phi}'_{\Lambda,L}$, and $\boldsymbol{\phi}'_{\Lambda^+,L}$ of $\boldsymbol{\Phi}'$ constructed above.

Indeed, as Lemma 3.5 shows, this number of variables is very likely polylogarithmic in n. Working out the details, we obtain the following.

Lemma 3.9. *Let* $0 < d < d_{\mathrm{uniq}}(k)$ *and let* $1 \leq L \leq \log^2 n$. *Then* $\mathbb{E}\left[\mathbb{1}\{X_{\Lambda,L} = X_{\Lambda^+,L} = 1 - L\}|\bar{\mathscr{L}}|^{3/2}\right] = o(1)$.

Proof. Let $\mathscr{V}^+ = V(\boldsymbol{\phi}'_{\Lambda,L}) \setminus V(\boldsymbol{\phi}'_{\Lambda-1,L})$ and obtain $\boldsymbol{\psi}^-$ from $\boldsymbol{\Phi}'$ by deleting all variables from $V(\boldsymbol{\phi}'_{\Lambda-1,L})$ and all clauses from $F(\boldsymbol{\phi}'_{\Lambda,L})$. Further, let $\Lambda^- = \Lambda^+ - \Lambda$ and let $\boldsymbol{\psi}^+$ be the sub-formula of $\boldsymbol{\psi}^-$ comprising all clauses and variables of $\boldsymbol{\psi}^-$ with distance at most $2\Lambda^-$ from $\mathscr{V}^+$ (see Figure 3 below). If $X_{\Lambda,L} = X_{\Lambda^+,L} = 1 - L$ then

$$|V(\boldsymbol{\phi}'_{\Lambda,L})| + |F(\boldsymbol{\phi}'_{\Lambda,L})| - |E(\boldsymbol{\phi}'_{\Lambda,L})| = L - 1, \tag{3.18}$$

$$|V(\boldsymbol{\psi}^+)| + |F(\boldsymbol{\psi}^+)| - |E(\boldsymbol{\psi}^+)| = |\mathscr{V}^+|. \tag{3.19}$$

Moreover, Lemma 3.5 shows that for suitable $c'_5 = c'_5(d, k, c_5)$, $c'_6 = c'_6(d, k, c_6) > 0$ we have

$$\mathbb{P}\left[|V(\boldsymbol{\phi}'_{\Lambda,L})| + |F(\boldsymbol{\phi}'_{\Lambda,L})| > \log^{c'_5} n\right] = O(n^{-2}), \tag{3.20}$$

$$\mathbb{P}\left[|V(\boldsymbol{\phi}'_{\Lambda^+,L})| + |F(\boldsymbol{\phi}'_{\Lambda^+,L})| > \log^{c'_6} n\right] = O(n^{-2}). \tag{3.21}$$

Let $\hat{\mathscr{L}} \subseteq \bar{\mathscr{L}}$ be the set of literals l that were added to the output set $\bar{\mathscr{L}}$ by Step 7 of PULP by way of clauses $a \in F(\boldsymbol{\phi}'_{\Lambda,L})$, i.e., at distance less than 2Λ from the initial set $\mathscr{L}$. Let $\mathscr{V}^- = \{|l| : l \in \hat{\mathscr{L}}\} \cap \mathscr{V}^+$ be the set of all variables at distance 2Λ from $\mathscr{L}$ in $\boldsymbol{\Phi}'$ that underlie a literal from $\hat{\mathscr{L}}$. If $X_{\Lambda,L} = 1 - L$, the variables and clauses at distance at most 2Λ from $\mathscr{L}$ do not cause PULP to run into a contradiction, because each clause contains $k \geq 3$ literals. Therefore, there does not exist a variable x such that both x and $\neg x$ belong to $\hat{\mathscr{L}}$. Hence, because the signs of $\boldsymbol{\Phi}'$ are uniformly random and PULP proceeds in a BFS order, we may assume without loss of generality that $\hat{\mathscr{L}}$ contains positive

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

literals only. Thus,

$$\mathscr{V}^- = \hat{\mathscr{L}} \cap \mathscr{V}^+ \subseteq \mathscr{V}^+. \tag{3.22}$$

We now apply Lemma 3.6 to the random formula $\boldsymbol{\psi}^-$. Specifically, let $\bar{\mathscr{L}}^+$ be the output of PULP on the formula $\boldsymbol{\psi}^-$ with the start set $\mathscr{L}^+$ comprised by the positive literals of $\mathscr{V}^+$. Further, let $\mathfrak{E}$ be the event that $|V(\boldsymbol{\phi}'_{\Lambda,L})| + |F(\boldsymbol{\phi}'_{\Lambda,L})| \le \log^{c'_5} n$. Let

$$X^+ = |E(\boldsymbol{\psi}^+)| - |V(\boldsymbol{\psi}^+)| - |F(\boldsymbol{\psi}^+)|$$

be the excess of $\boldsymbol{\psi}^+$. Since $0 < d < d_{\mathrm{uniq}}(k)$, given $\mathfrak{E}$ the formula $\boldsymbol{\psi}^-$ has the same distribution as a random k-CNF with $n^- = n - O(\log^c n)$ variables and $m^- \sim \mathrm{Po}(d^- n^-/k)$ random clauses, with $0 < d^- = d + o(1) < d_{\mathrm{uniq}}(k)$. Hence, assuming that $c_6 = c_6(d,k) > c'_5$ is sufficiently large, Lemma 3.6 shows that

$$\mathbb{P}\left[\mathbb{1}\mathfrak{E} \cap \{X^+ = -|\mathscr{V}^+|\} \cdot |\bar{\mathscr{L}}^+| > u\right] \le \mathbb{P}\left[\sum_{1 \le i \le \log^{c'_5} n} |\overline{\{\mathfrak{x}\}}_{\mathbb{T}[i]}| > u\right] + O(n^{-2}) \quad (u > 0). \tag{3.23}$$

Combining (3.23) with Corollary 3.3, we conclude that for a large enough $c_7 = c_7(d,k) > c'_6$,

$$\mathbb{P}\left[\mathbb{1}\mathfrak{E} \cap \{X^+ = -|\mathscr{V}^+|\} \cdot |\bar{\mathscr{L}}^+| > \log^{c_7} n\right] = O(n^{-2}). \tag{3.24}$$

In light of above, we see that

$$\begin{aligned}
&\mathbb{E}\left[\mathbb{1}\{X_{\Lambda,L} = X_{\Lambda^+,L} = 1 - L\}|\bar{\mathscr{L}}|^{3/2}\right] \\
&\quad \le \mathbb{E}\left[\mathbb{1}\mathfrak{E} \cap \{X_{\Lambda,L} = X_{\Lambda^+,L} = 1 - L\}|\bar{\mathscr{L}}|^{3/2}\right] + (2n)^{3/2}(1 - \mathbb{P}[\mathfrak{E}]), && [\text{since } |\bar{\mathscr{L}}| \le 2n] \\
&\quad \le \mathbb{E}\left[\mathbb{1}\mathfrak{E} \cap \{X_{\Lambda,L} = X_{\Lambda^+,L} = 1 - L\}(|\bar{\mathscr{L}}^+| + \log^{c'_5} n)^{3/2}\right] + o(1), && [\text{from (3.20)}] \\
&\quad \le \mathbb{E}\left[\mathbb{1}\mathfrak{E} \cap \{X^+ = -|\mathscr{V}^+|\}(|\bar{\mathscr{L}}^+| + \log^{c'_5} n)^{3/2}\right] + o(1), && [\text{from (3.18), (3.19)}] \\
&\quad \le \mathbb{P}\left[\mathbb{1}\mathfrak{E} \cap \{X^+ = -|\mathscr{V}^+|\} \cdot |\bar{\mathscr{L}}^+| > \log^{c_7} n\right](2n)^{3/2} \\
&\qquad + \mathbb{P}\left[\{X_{\Lambda,L} = 1 - L\}\right](2\log^{c_7} n)^{3/2} + o(1) && [\text{total probability}] \\
&\quad = o(1) && [\text{from (3.24), Lemma 3.5}]
\end{aligned}$$

completing the proof. □

Proof of Lemma 2.10. The assertion is immediate from Corollary 3.8, Lemma 3.9, Lemma 3.5 and the deterministic bound $|\bar{\mathscr{L}}| \le 2n$. □

4. Proof of Proposition 2.1

Let $\pi^{(\ell)}_{d,k} = \mathrm{BP}^{\ell}_{d,k}(\delta_{1/2})$ be the distribution obtained after ℓ iterations of $\mathrm{BP}_{d,k}(\cdot)$, with the convention $\pi^{(0)}_{d,k} = \delta_{1/2}$. We recall $(\boldsymbol{\mu}_{\pi,i,j})_{i,j\ge 1}$ signify independent random variables with distribution π.

Fact 4.1. *For all $\ell \ge 0$ the random variables $\boldsymbol{\mu}_{\pi^{(\ell)}_{d,k},1,1}$ and $1 - \boldsymbol{\mu}_{\pi^{(\ell)}_{d,k},1,1}$ are identically distributed.*

Proof. This is an immediate consequence of the fact that the random variables d^+, d^- from the definition (1.3)–(1.4) are identically distributed. □

While the following is a direct consequence of the fact that Belief Propagation is 'exact on trees' (see [51, Chapter 14] for precise statements), we carry out a detailed proof for the sake of completeness. Following the conventions from Section 1.2.1, we continue to denote by $\boldsymbol{\tau}^{(\ell)}$ a random satisfying assignment of the k-CNF $\mathbb{T}^{(\ell)} = \mathbb{T}^{(\ell)}_{d,k}$.

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

Fact 4.2. *For all $\ell \geq 0$, $d > 0$ we have $\mathbb{P}\left[\boldsymbol{\tau}^{(\ell)}(\mathfrak{x}) = 1 \mid \mathbb{T}\right] \sim \pi_{d,k}^{(\ell)}$.*

Proof. We proceed by induction on ℓ. As $\pi_{d,k}^{(0)} = \delta_{1/2}$, for $\ell = 0$ there is nothing to show. To go from $\ell - 1$ to $\ell \geq 1$, for a clause $a \in \partial_{\mathbb{T}}\mathfrak{x}$ and a variable $y \in \partial_{\mathbb{T}}a \setminus \{\mathfrak{x}\}$ let $\mathbb{T}_{y\to a}$ be the component of the forest $\mathbb{T} - a$ obtained by removing clause a that contains variable y. We consider y the root of $\mathbb{T}_{y\to a}$. Further, obtain $\mathbb{T}_{y\to a}^{(\ell-1)}$ from $\mathbb{T}_{y\to a}$ by deleting all clauses and variables at a distance greater than $2(\ell-1)$ from y. Additionally, for $s \in \{\pm 1\}$ let

$$Z^{(\ell)}(s) = \left|\left\{\sigma \in S\left(\mathbb{T}^{(\ell)}\right) : \sigma(\mathfrak{x}) = s\right\}\right|, \quad Z_{y\to a}^{(\ell-1)}(s) = \left|\left\{\sigma \in S\left(\mathbb{T}_{y\to a}^{(\ell-1)}\right) : \sigma(y) = s\right\}\right|. \tag{4.1}$$

In words, $Z^{(\ell)}(s)$ is the number of satisfying assignments of $\mathbb{T}^{(\ell)}$ that set the root $\mathfrak{x}$ to s, and $Z_{y\to a}^{(\ell-1)}(s)$ is the corresponding quantity for the sub-tree $\mathbb{T}_{y\to a}^{(\ell-1)}$.

Clearly, setting $\mathfrak{x}$ to $s \in \{\pm 1\}$ immediately satisfies all clauses $a \in \partial_{\mathbb{T}}^{s}\mathfrak{x}$. By contrast, once $\mathfrak{x}$ is assigned the value $+1$ each clause $a \in \partial_{\mathbb{T}}^{-s}\mathfrak{x}$ needs to be satisfied by setting some other variable $y \in \partial_{\mathbb{T}}a \setminus \{\mathfrak{x}\}$ to the value $\mathrm{sign}(y, a)$. Hence,

$$\begin{aligned} Z^{(\ell)}(s) = \Bigg[\prod_{a\in\partial_{\mathbb{T}}^{+}\mathfrak{x}} \prod_{y\in\partial_{\mathbb{T}}a\setminus\{\mathfrak{x}\}} \sum_{t\in\{\pm1\}} Z_{y\to a}^{(\ell-1)}(t)\Bigg] \cdot \Bigg[\prod_{a\in\partial_{\mathbb{T}}^{-}\mathfrak{x}} \Bigg(\prod_{y\in\partial_{\mathbb{T}}a\setminus\{\mathfrak{x}\}} \sum_{t\in\{\pm1\}} Z_{y\to a}^{(\ell-1)}(t) \\ - \prod_{y\in\partial_{\mathbb{T}}a\setminus\{\mathfrak{x}\}} Z_{y\to a}^{(\ell-1)}\left(-\mathrm{sign}(y,a)\right)\Bigg)\Bigg]. \end{aligned} \tag{4.2}$$

Furthermore, the definition of the Galton-Watson tree $\mathbb{T}$ ensures that the sub-trees $\mathbb{T}_{y\to a}^{(\ell-1)}$ are independent copies of $\mathbb{T}^{(\ell-1)}$. Hence, by induction we have

$$\frac{Z_{y\to a}^{(\ell-1)}(1)}{\sum_{s\in\{\pm1\}} Z_{y\to a}^{(\ell-1)}(s)} \sim \pi_{d,k}^{(\ell-1)} \quad \text{for all } a \in \partial_{\mathbb{T}}\mathfrak{x},\ y \in \partial_{\mathbb{T}}a \setminus \{\mathfrak{x}\}, \tag{4.3}$$

and the random variables $Z_{y\to a}^{(\ell-1)}(1)/\sum_{s\in\{\pm1\}} Z_{y\to a}^{(\ell-1)}(s)$ are mutually independent. Combining (4.2)–(4.3) with Fact 4.1, we finally obtain

$$\begin{aligned} \mathbb{P}\left[\boldsymbol{\tau}^{(\ell)}(\mathfrak{x}) = 1 \mid \mathbb{T}\right] &= \frac{Z^{(\ell)}(1)}{\sum_{s\in\{\pm1\}} Z^{(\ell)}(s)} \\ &\sim \frac{\prod_{i=1}^{d^-}\left[1 - \prod_{j=1}^{k-1} \boldsymbol{\mu}_{\pi_{d,k}^{(\ell-1)},2i-1,j}\right]}{\prod_{i=1}^{d^-}\left[1 - \prod_{j=1}^{k-1} \boldsymbol{\mu}_{\pi_{d,k}^{(\ell-1)},2i-1,j}\right] + \prod_{i=1}^{d^+}\left[1 - \prod_{j=1}^{k-1} \boldsymbol{\mu}_{\pi_{d,k}^{(\ell-1)},2i,j}\right]} \sim \pi_{d,k}^{(\ell)}, \end{aligned}$$

thereby completing the induction. □

Combining the combinatorial interpretation of the distributions $\pi_{d,k}^{(\ell)}$ with the Gibbs uniqueness property, we proceed to show that the sequence $(\pi_{d,k}^{(\ell)})_{\ell}$ converges in the weak topology. To this end, it suffices to show that the sequence is Cauchy with respect to the Wasserstein W_1 metric.

Lemma 4.3. *If $d < d_{\mathrm{uniq}}(k)$ then $(\pi_{d,k}^{(\ell)})_{\ell\geq 0}$ is a W_1-Cauchy sequence.*

Proof. If $d < d_{\mathrm{uniq}}(k)$ then the random tree $\mathbb{T} = \mathbb{T}_{d,k}$ enjoys the Gibbs uniqueness property; hence, (1.1) is satisfied. Consequently, given $0 < \varepsilon < 1$ we can choose $\ell_0 = \ell_0(d, k, \varepsilon) > 0$ large

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

enough so that the event

$$\mathfrak{U}_{\varepsilon,\ell} = \left\{ \max_{\tau \in S(\mathbb{T}^{(\ell)})} \left| \mathbb{P}\left[\boldsymbol{\tau}^{(\ell)}(\mathfrak{x}) = 1 \mid \mathbb{T}\right] - \mathbb{P}\left[\boldsymbol{\tau}^{(\ell)}(\mathfrak{x}) = 1 \mid \mathbb{T},\ \forall x \in \partial^{2\ell}\mathfrak{x} : \boldsymbol{\tau}^{(\ell)}(x) = \tau(x)\right] \right| > \varepsilon \right\}$$

has probability

$$\mathbb{P}\left[\mathfrak{U}_{\varepsilon,\ell}\right] < \varepsilon, \quad \text{for all } \ell \geq \ell_0. \tag{4.4}$$

Now suppose that $\ell_0 \leq \ell < \ell'$. Let $\boldsymbol{\tau}^{(\ell)}, \boldsymbol{\tau}^{(\ell')}$ be independent uniformly random satisfying assignments of $\mathbb{T}^{(\ell)}$ and $\mathbb{T}^{(\ell')}$, respectively. We claim that

$$\mathbb{P}\left[\left|\mathbb{P}\left[\boldsymbol{\tau}^{(\ell)}(\mathfrak{x}) = 1 \mid \mathbb{T}\right] - \mathbb{P}\left[\boldsymbol{\tau}^{(\ell')}(\mathfrak{x}) = 1 \mid \mathbb{T}\right]\right| > \varepsilon\right] < \varepsilon. \tag{4.5}$$

To see this, let $\boldsymbol{\tau}^{(\ell,\ell')} = (\boldsymbol{\tau}^{(\ell')}(x))_{x \in \partial^{2\ell}_{\mathbb{T}}\mathfrak{x}}$ comprise the truth values that $\boldsymbol{\tau}^{(\ell')}$ assigns to the variables at distance exactly 2ℓ from $\mathfrak{x}$. Then

$$\begin{aligned}\mathbb{P}\left[\boldsymbol{\tau}^{(\ell')}(\mathfrak{x}) = 1 \mid \mathbb{T}\right] &= \mathbb{E}\left[\mathbb{P}\left[\boldsymbol{\tau}^{(\ell')}(\mathfrak{x}) = 1 \mid \mathbb{T}, \boldsymbol{\tau}^{(\ell,\ell')}\right] \mid \mathbb{T}\right] \\ &= \mathbb{E}\left[\mathbb{P}\left[\boldsymbol{\tau}^{(\ell)}(\mathfrak{x}) = 1 \mid \mathbb{T}, \boldsymbol{\tau}^{(\ell,\ell')},\ \forall x \in \partial^{2\ell}\mathfrak{x} : \boldsymbol{\tau}^{(\ell)}(x) = \boldsymbol{\tau}^{(\ell,\ell')}_x\right] \mid \mathbb{T}\right].\end{aligned}$$

Hence, for every $T \in \mathfrak{U}_{\varepsilon,\ell}$ we have

$$\left|\mathbb{P}\left[\boldsymbol{\tau}^{(\ell')}(\mathfrak{x}) = 1 \mid \mathbb{T} = T\right] - \mathbb{P}\left[\boldsymbol{\tau}^{(\ell)}(\mathfrak{x}) = 1 \mid \mathbb{T} = T\right]\right| \leq \varepsilon. \tag{4.6}$$

Thus, (4.5) follows from (4.4) and (4.6).

Finally, since Fact 4.2 demonstrates that $\mathbb{P}\left[\boldsymbol{\tau}^{(\ell)}(\mathfrak{x}) = 1 \mid \mathbb{T} = T\right] \sim \pi^{(\ell)}_{d,k}$ and $\mathbb{P}\left[\boldsymbol{\tau}^{(\ell')}(\mathfrak{x}) = 1 \mid \mathbb{T} = T\right] \sim \pi^{(\ell')}_{d,k}$, (4.5) shows that

$$W_1\left(\pi^{(\ell)}_{d,k}, \pi^{(\ell')}_{d,k}\right) < 2\varepsilon \quad \text{for all } \ell_0 \leq \ell < \ell'.$$

Hence, the sequence $(\pi^{(\ell)}_{d,k})_\ell$ is Cauchy. □

We are left to bound the lower tail of the limiting distribution $\pi_{d,k} = \lim_{\ell\to\infty} \pi^{(\ell)}_{d,k}$.

Lemma 4.4. *If $d < d_{\mathrm{uniq}}(k)$ then $\mathbb{E}\log^2 \boldsymbol{\mu}_{\pi_{d,k},1,1} < \infty$.*

Proof. We are going to bound $\mathbb{E}\log^2 \boldsymbol{\mu}_{\pi^{(\ell)}_{d,k},1,1}$ and subsequently invoke the monotone convergence theorem to complete the proof. First, we note that for all $\ell \geq 0$ we have

$$\begin{aligned}\mathbb{E}\log^2 \boldsymbol{\mu}_{\pi^{(\ell)}_{d,k},1,1} &= \mathbb{E}\log^2 \frac{Z(\mathbb{T}^{(\ell)}, \{\mathfrak{x}\})}{Z(\mathbb{T}^{(\ell)})} && \text{[by Fact 4.2]} \\ &\leq \mathbb{E}\left[|\overline{\{\mathfrak{x}\}}_{\mathbb{T}}|^2\right] && \text{[by Lemma 2.8]} \\ &< \infty && \text{[by Corollary 3.4].}\end{aligned}$$

Since $\pi_{d,k}$ is the weak limit of $(\pi^{(\ell)}_{d,k})_\ell$, we conclude that for any $N \in \mathbb{N}$,

$$\mathbb{E}\left[N \wedge \log^2 \boldsymbol{\mu}_{\pi_{d,k},1,1}\right] = \lim_{\ell\to\infty} \mathbb{E}\left[N \wedge \log^2 \boldsymbol{\mu}_{\pi^{(\ell)}_{d,k},1,1}\right] \leq \mathbb{E}\left[|\overline{\{\mathfrak{x}\}}_{\mathbb{T}}|^2\right] < \infty. \tag{4.7}$$

Finally, applying the monotone convergence theorem to the limit $N \to \infty$, we see that the uniform bound (4.7) implies the assertion. □

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

Proof of Proposition 2.1. In light of Fact 4.1 and Lemmas 4.3 and 4.4, it only remains to show that

$$\mathbb{E}\left|\log\left(\prod_{i=1}^{d^-}\boldsymbol{\mu}_{\pi_{d,k},2i}+\prod_{i=1}^{d^+}\boldsymbol{\mu}_{\pi_{d,k},2i-1}\right)\right|<\infty \qquad \text{and} \qquad \mathbb{E}\left|\log\left(1-\prod_{j=1}^{k}\boldsymbol{\mu}_{\pi_{d,k},1,j}\right)\right|<\infty.$$

Recall the definition of $\boldsymbol{\mu}_{\pi_{d,k},i}$ from (1.4). Using Fact 4.1 and Lemma 4.4, we obtain

$$\begin{aligned}\mathbb{E}\left|\log\left(\prod_{i=1}^{d^-}\boldsymbol{\mu}_{\pi_{d,k},2i}+\prod_{i=1}^{d^+}\boldsymbol{\mu}_{\pi_{d,k},2i-1}\right)\right| &\le \log(2)+\mathbb{E}\left|\log\prod_{i=1}^{d^-}\boldsymbol{\mu}_{\pi_{d,k},2i}\right|\\ &\le \log(2)+\frac{d}{2}\mathbb{E}\left|\log\boldsymbol{\mu}_{\pi_{d,k},1}\right|\\ &\le \log(2)+\frac{d}{2}\sqrt{\mathbb{E}\left|\log^2\boldsymbol{\mu}_{\pi_{d,k},1,1}\right|}<\infty,\end{aligned}$$

yielding the first inequality. Similarly, invoking Fact 4.1 and Lemma 4.4 for the second l.h.s. above gives

$$\mathbb{E}\left|\log\left(1-\prod_{j=1}^{k}\boldsymbol{\mu}_{\pi_{d,k},1,j}\right)\right|\le\mathbb{E}\left|\log\left(1-\boldsymbol{\mu}_{\pi_{d,k},1,1}\right)\right|=\mathbb{E}\left|\log\left(\boldsymbol{\mu}_{\pi_{d,k},1,1}\right)\right|\le\sqrt{\mathbb{E}\left|\log^2\boldsymbol{\mu}_{\pi_{d,k},1,1}\right|}<\infty,$$

thereby completing the proof. □

5. Proof of Corollary 2.2

In order to turn the estimate of the expectation of $\log 1\vee Z(\boldsymbol{\Phi})$ provided by Proposition 2.3 into a 'with high probability' statement, we harness a 'soft' version of the k-SAT problem where violated clauses are discouraged but not strictly forbidden. To be precise, for a k-CNF Φ and a real $\beta>0$ define

$$Z_\beta(\Phi)=\sum_{\sigma\in\{\pm1\}^{V(\Phi)}}\prod_{a\in F(\Phi)}\exp\left(-\beta\mathbb{1}\{\sigma\not\models a\}\right). \tag{5.1}$$

Thus, each satisfying assignment contributes one to the sum on the r.h.s. of (5.1), while the contribution of assignments that violate a number M of clauses equals $\exp(-\beta M)$. The value $Z_\beta(\boldsymbol{\Phi})$, called the *partition function of the random k-SAT model at inverse temperature* β, has received a considerable amount of attention in the mathematical physics literature (see, e.g., [58]). Crucially, by means of an interpolation argument [36, 42] it is possible to prove the following.

Theorem 5.1 ([60, Theorem 1]). *For any $k\ge3$, any $\beta>0$ and any probability measure π on $[0,1]$ we have*

$$\begin{aligned}\frac{1}{n}\mathbb{E}\left[\log Z_\beta(\boldsymbol{\Phi})\right]\le\mathbb{E}\Bigg[&\log\left(\prod_{i=1}^{d^-}\boldsymbol{\mu}_{\beta,\pi,2i}+\prod_{i=1}^{d^+}\boldsymbol{\mu}_{\beta,\pi,2i-1}\right)\\ &-\frac{d(k-1)}{k}\log\left(1-\left(1-e^{-\beta}\right)\prod_{j=1}^{k}\boldsymbol{\mu}_{\pi,1,j}\right)\Bigg],\quad \textit{where}\end{aligned} \tag{5.2}$$

$$\boldsymbol{\mu}_{\beta,\pi,i}=1-(1-\exp(-\beta))\prod_{j=1}^{k-1}\boldsymbol{\mu}_{\pi,i,j} \qquad \textit{for } i\ge1.$$

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

We emphasise that the bound (5.2) holds for any $n \geq k$ *without* an error term. We also notice that by the monotone convergence theorem for the measure $\pi = \pi_{d,k}$ from Theorem 1.1 we have

$$\lim_{\beta\to\infty} \mathbb{E}\left[\log\left(\prod_{i=1}^{d^-} \boldsymbol{\mu}_{\beta,\pi_{d,k},2i} + \prod_{i=1}^{d^+} \boldsymbol{\mu}_{\beta,\pi_{d,k},2i-1}\right) - \frac{d(k-1)}{k}\log\left(1-\left(1-e^{-\beta}\right)\prod_{j=1}^{k}\boldsymbol{\mu}_{\pi_{d,k},1,j}\right)\right]$$
$$= \mathbb{E}\left[\log\left(\prod_{i=1}^{d^-} \boldsymbol{\mu}_{\pi_{d,k},2i} + \prod_{i=1}^{d^+} \boldsymbol{\mu}_{\pi_{d,k},2i-1}\right) - \frac{d(k-1)}{k}\log\left(1-\prod_{j=1}^{k}\boldsymbol{\mu}_{\pi_{d,k},1,j}\right)\right] = \mathfrak{B}_{d,k}(\pi_{d,k}). \tag{5.3}$$

The reason why we proceed by way of the 'soft' model with $\beta < \infty$ is that for this model a routine application of Azuma-Hoeffding implies the following concentration bound.

Lemma 5.2. *For any fixed $\beta > 0$ we have $\mathbb{P}\left[\left|\log Z_\beta(\boldsymbol{\Phi}) - \mathbb{E}\log Z_\beta(\boldsymbol{\Phi})\right| > \sqrt{n}\log n\right] = o(1/n)$.*

Proof. The clauses of the random formula $\boldsymbol{\Phi}$ are drawn independently, and adding or removing a single clause can alter the value of $\log Z_\beta(\cdot)$ by no more than $\pm\beta$. □

Proof of Corollary 2.2. We proceed with a proof by contradiction. In particular, towards a contradiction, assume there exists an $\varepsilon > 0$ such that for infinitely many $n \geq 1$ we have

$$\mathbb{P}\left[\frac{1}{n}\log Z(\boldsymbol{\Phi}) > \mathfrak{B}_{d,k}(\pi_{d,k}) + \varepsilon\right] > \varepsilon. \tag{5.4}$$

Moreover, by (5.3) we can find a $\beta_0 > 0$ such that for every $\beta \geq \beta_0$ we have

$$\left|\mathbb{E}\left[\log\left(\prod_{i=1}^{d^-} \boldsymbol{\mu}_{\beta,\pi_{d,k},2i} + \prod_{i=1}^{d^+} \boldsymbol{\mu}_{\beta,\pi_{d,k},2i-1}\right) - \frac{d(k-1)}{k}\log\left(1-\left(1-e^{-\beta}\right)\prod_{j=1}^{k}\boldsymbol{\mu}_{\pi_{d,k},1,j}\right)\right]\right.$$
$$\left. - \mathfrak{B}_{d,k}(\pi_{d,k})\right| < \varepsilon/3. \tag{5.5}$$

Invoking Lemma 5.2 for $\beta = \beta_0$ and sufficiently large n gives

$$\mathbb{P}\left[\frac{1}{n}\log Z_{\beta_0}(\boldsymbol{\Phi}) > \frac{1}{n}\mathbb{E}\log Z_{\beta_0}(\boldsymbol{\Phi}) + \varepsilon/3\right] \leq \varepsilon/3. \tag{5.6}$$

The definition (5.1) of the partition function ensures that $Z_\beta(\boldsymbol{\Phi}) \geq Z(\boldsymbol{\Phi})$ for all $\beta > 0$. Therefore, combining (5.4)–(5.6), and Theorem 5.1 we see that for large enough n the following holds with probability at least $1 - \frac{2}{3}\varepsilon$:

$$\frac{1}{n}\log Z(\boldsymbol{\Phi}) \leq \frac{1}{n}\log Z_{\beta_0}(\boldsymbol{\Phi}) \leq \frac{1}{n}\mathbb{E}\log Z_{\beta_0}(\boldsymbol{\Phi}) + \frac{\varepsilon}{3} \leq \mathfrak{B}_{d,k}(\pi_{d,k}) + \frac{2}{3}\varepsilon,$$

contradicting our assumption, and thus completing the proof. □

6. Proof of Proposition 2.3

In this section we prove Propositions 2.5 and 2.6, which in light of Fact 2.4, imply Proposition 2.3. Both proofs follow a similar structure and make use of Propositions 2.7 and 2.11, which we therefore prove first.

6.1 Proof of Proposition 2.7

We show that both terms of (2.5) have finite expectation. Let us begin with the first one.

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

Lemma 6.1. *If* $d < d_{\mathrm{uniq}}(k)$ *then* $\mathbb{E}\left[\left|\log \frac{Z(\mathbf{\Phi}'')\vee 1}{Z(\mathbf{\Phi}')\vee 1}\right|^{3/2}\right] = O(1)$.

Proof. Since $\mathbf{\Phi}''$ is obtained from $\mathbf{\Phi}'$ by adding clauses, we have

$$0 \le Z(\mathbf{\Phi}'') \le Z(\mathbf{\Phi}') \le 2^n. \tag{6.1}$$

Hence,

$$\log \frac{Z(\mathbf{\Phi}'')\vee 1}{Z(\mathbf{\Phi}')\vee 1} = 0 \quad \text{if } Z(\mathbf{\Phi}') = 0. \tag{6.2}$$

Therefore, we may assume from now on that $\mathbf{\Phi}'$ is satisfiable.

The number $\mathbf{\Delta}'' \sim \mathrm{Po}(d(k-1)/k)$ of new clauses is a Poisson variable with bounded mean. Therefore, Bennett's inequality shows that $\mathbb{P}\left[\mathbf{\Delta}'' > \log n\right] = O(n^{-2})$. Since (6.1) shows that $|\log((Z(\mathbf{\Phi}'')\vee 1)/(Z(\mathbf{\Phi}')\vee 1))|^{3/2} \le n^{3/2}$, we conclude that

$$\mathbb{E}\left[\mathbb{1}\left\{\mathbf{\Delta}'' > \log n\right\} \cdot \left|\log \frac{Z(\mathbf{\Phi}'')\vee 1}{Z(\mathbf{\Phi}')\vee 1}\right|^{3/2}\right] = o(1). \tag{6.3}$$

Further, let $c_1, \ldots, c_{\mathbf{\Delta}''}$ be the new clauses added by **CPL2**. Let $x_{1,1}, \ldots, x_{1,k}, \ldots, x_{\mathbf{\Delta}'',1}, \ldots, x_{\mathbf{\Delta}'',k}$ be their constituent variables and let $\mathscr{X} = \{x_{1,1}, \ldots, x_{1,k}, \ldots, x_{\mathbf{\Delta}'',1}, \ldots, x_{\mathbf{\Delta}'',k}\}$. Since the clauses $c_1, \ldots, c_{\mathbf{\Delta}''}$ are chosen uniformly and independently, a routine balls-into-bins consideration shows that

$$\mathbb{P}\left[|\mathscr{X}| \le k(\mathbf{\Delta}'' - 1) \mid \mathbf{\Delta}'' \le \log n\right] = \tilde{O}(n^{-2}). \tag{6.4}$$

Now, consider the 'good' event

$$\mathfrak{G} = \left\{Z(\mathbf{\Phi}') > 0,\ \mathbf{\Delta}'' \le \log n,\ |\mathscr{X}| > k(\mathbf{\Delta}'' - 1)\right\}.$$

Combining (6.1)–(6.4), we see that

$$\mathbb{E}\left[(1 - \mathbb{1}_{\mathfrak{G}}) \cdot \left|\log \frac{Z(\mathbf{\Phi}'')\vee 1}{Z(\mathbf{\Phi}')\vee 1}\right|^{3/2}\right] = o(1). \tag{6.5}$$

Hence, we are left to bound $\mathbb{E}[\mathbb{1}_{\mathfrak{G}} \cdot |\log((Z(\mathbf{\Phi}'')\vee 1)/(Z(\mathbf{\Phi}')\vee 1))|^{3/2}]$. If $\mathfrak{G}$ occurs and thus $|\mathscr{X}| > k(\mathbf{\Delta}'' - 1)$, then there exists a set of literals $\mathscr{L} \subseteq \{x_{1,1}, \neg x_{1,1}, \ldots, x_{1,k}, \neg x_{1,k}, \ldots, x_{\mathbf{\Delta}'',1}, \neg x_{\mathbf{\Delta}'',1}, \ldots, x_{\mathbf{\Delta}'',k}, \neg x_{\mathbf{\Delta}'',k}\}$ such that

- every clause c_i contains a literal from $\mathscr{L}$ $(1 \le i \le \mathbf{\Delta}'')$, and
- there does not exist $x \in \mathscr{X}$ such that $x \in \mathscr{L}$ and $\neg x \in \mathscr{L}$.

Moreover, on $\mathfrak{G}$ we have $|\mathscr{L}| \le |\mathscr{X}| \le k \log n$. Let $\bar{\mathscr{L}} = \bar{\mathscr{L}}_{\mathbf{\Phi}'}$ be the output of PULP on $(\mathbf{\Phi}', \mathscr{L})$. Then Lemma 2.8 shows that

$$\mathbb{E}\left[\mathbb{1}_{\mathfrak{G}} \cdot \left|\log \frac{Z(\mathbf{\Phi}'')\vee 1}{Z(\mathbf{\Phi}')\vee 1}\right|^{3/2}\right] \le \mathbb{E}\left[\mathbb{1}_{\mathfrak{G}} \cdot |\bar{\mathscr{L}}|^{3/2}\right]. \tag{6.6}$$

Furthermore, since by **CPL2** the new clauses $c_1, \ldots, c_{\mathbf{\Delta}''}$ are chosen independently of the formula $\mathbf{\Phi}'$, Lemma 2.10 implies that there exists $C = C(d,k) > 0$ such that

$$\mathbb{E}\left[\mathbb{1}_{\mathfrak{G}} \cdot |\bar{\mathscr{L}}|^{3/2} \mid \mathbf{\Delta}''\right] \le C \cdot (\mathbf{\Delta}'')^{3/2}. \tag{6.7}$$

Combining (6.6)–(6.7) and recalling that $\mathbf{\Delta}'' \sim \mathrm{Po}(d(k-1)/k)$, we obtain

$$\mathbb{E}\left[\mathbb{1}_{\mathfrak{G}} \cdot \left|\log \frac{Z(\mathbf{\Phi}'')\vee 1}{Z(\mathbf{\Phi}')\vee 1}\right|^{3/2}\right] = O(1). \tag{6.8}$$

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

Finally, the assertion follows from (6.5) and (6.8). □

We move on to the second term of (2.5).

Lemma 6.2. *If* $d < d_{\mathrm{uniq}}(k)$ *then* $\mathbb{E}\left[\left|\log \frac{Z(\mathbf{\Phi}''')\vee 1}{Z(\mathbf{\Phi}')\vee 1}\right|^{3/2}\right] = O(1)$.

Proof. We proceed similarly as in the proof of Lemma 6.1. The construction in **CPL3** ensures that $\mathbf{\Phi}'''$ contains one additional variable x_{n+1} and $\mathbf{\Delta}''' \sim \mathrm{Po}(d)$ new clauses $b_1, \ldots, b_{\mathbf{\Delta}'''}$ that each contain x_{n+1} and $k-1$ other variables. Let $x_{1,1}, \ldots, x_{1,k-1}, \ldots, x_{\mathbf{\Delta}''',1}, \ldots, x_{\mathbf{\Delta}''',k-1} \in \{x_1, \ldots, x_n\}$ be the variables among $x_1, \ldots, x_n$ that appear in $b_1, \ldots, b_{\mathbf{\Delta}'''}$ and let $\mathscr{X} = \{x_{1,1}, \ldots, x_{\mathbf{\Delta}''',k-1}\}$. Then

$$0 \le Z(\mathbf{\Phi}'') \le 2Z(\mathbf{\Phi}') \le 2^{n+1}. \tag{6.9}$$

Hence, if $\mathbf{\Phi}'$ is unsatisfiable, then so is $\mathbf{\Phi}'''$ and thus

$$\log \frac{Z(\mathbf{\Phi}''')\vee 1}{Z(\mathbf{\Phi}')\vee 1} = 0 \qquad \text{if } Z(\mathbf{\Phi}') = 0. \tag{6.10}$$

Furthermore, since $\mathbf{\Delta}''' \sim \mathrm{Po}(d)$, Bennett's inequality shows that $\mathbb{P}\left[\mathbf{\Delta}''' > \log n\right] = O(n^{-2})$. Therefore, (6.9) shows that

$$\mathbb{E}\left[\mathbb{1}\left\{\mathbf{\Delta}''' > \log n\right\} \cdot \left|\log \frac{Z(\mathbf{\Phi}''')\vee 1}{Z(\mathbf{\Phi}')\vee 1}\right|^{3/2}\right] = o(1). \tag{6.11}$$

Moreover, since the $k-1$ variables among $x_1, \ldots, x_n$ that appear in the clauses $b_1, \ldots, b_{\Delta'''}$ are chosen uniformly and independently, a simple balls-into-bins argument shows that

$$\mathbb{P}\left[|\mathscr{X}| \le (k-1)(\mathbf{\Delta}''-1) \mid \mathbf{\Delta}'' \le \log n\right] = \tilde{O}(n^{-2}). \tag{6.12}$$

Hence, consider the event

$$\mathfrak{G} = \left\{Z(\mathbf{\Phi}') > 0,\ \mathbf{\Delta}''' \le \log n,\ |\mathscr{X}| > (k-1)(\mathbf{\Delta}''-1)\right\}.$$

Combining (6.9)–(6.12), we obtain

$$\mathbb{E}\left[(1-\mathbb{1}\mathfrak{G}) \cdot \left|\log \frac{Z(\mathbf{\Phi}''')\vee 1}{Z(\mathbf{\Phi}')\vee 1}\right|^{3/2}\right] = o(1). \tag{6.13}$$

Furthermore, if the event $\mathfrak{G}$ occurs, then there exists a set $\mathscr{L} \subseteq \{x, \neg x : x \in \mathscr{X}\}$ of literals such that each clause b_i, $1 \le i \le \mathbf{\Delta}'''$, contains a literal $l \in \mathscr{L}$ and such that $\{x, \neg x\} \not\subseteq \mathscr{L}$ for all $x \in \mathscr{X}$. Hence, with $\bar{\mathscr{L}} = \bar{\mathscr{L}}_{\mathbf{\Phi}'}$ the output of PULP on $(\mathbf{\Phi}', \mathscr{L})$, Lemma 2.8 shows that

$$\mathbb{E}\left[\mathbb{1}\mathfrak{G} \cdot \left|\log \frac{Z(\mathbf{\Phi}''')\vee 1}{Z(\mathbf{\Phi}')\vee 1}\right|^{3/2}\right] \le \mathbb{E}\left[\mathbb{1}\mathfrak{G} \cdot |\bar{\mathscr{L}}|^{3/2}\right]. \tag{6.14}$$

Furthermore, since the clauses $b_1, \ldots, b_{\mathbf{\Delta}'''}$ are drawn independently of $\mathbf{\Phi}'''$, Lemma 2.10 shows that there exists $C = C(d,k) > 0$ such that

$$\mathbb{E}\left[\mathbb{1}\mathfrak{G} \cdot |\bar{\mathscr{L}}|^{3/2} \mid \mathbf{\Delta}'''\right] \le C \cdot (\mathbf{\Delta}''')^{3/2}. \tag{6.15}$$

Finally, since $\mathbf{\Delta}''' \sim \mathrm{Po}(d)$, the assertion follows from (6.13), (6.14) and (6.15). □

Proof of Proposition 2.7. The proposition follows immediately from Lemmas 6.1–6.2. □

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

6.2 Proof of Proposition 2.11

Let $\pi_{d,k}^{(\ell)} = \mathrm{BP}_{d,k}^{\ell}(\delta_{1/2})$ be the result of an ℓ-fold application of the operator $\mathrm{BP}_{d,k}$ from (1.3) to the point mass at $1/2$. Also recall from (2.7) that $\boldsymbol{\pi}'_n$ denotes the empirical distribution of the marginals $(\mathbb{P}[\boldsymbol{\sigma}_{\boldsymbol{\Phi}'}(x_i) = 1 \mid \boldsymbol{\Phi}'])_{1 \le i \le n}$.

Lemma 6.3. *Suppose that $d < d_{\mathrm{uniq}}(k)$. For any $\varepsilon > 0$ there exists $\ell_0 = \ell_0(d, k, \varepsilon) > 0$ such that for all $\ell \ge \ell_0$ we have*

$$\mathbb{E}\left[W_1\left(\boldsymbol{\pi}'_n, \pi_{d,k}^{(\ell)}\right) \mid Z(\boldsymbol{\Phi}') > 0\right] < \varepsilon + o(1).$$

Proof. Assume that $\ell \ge \ell_0$ for a large enough $\ell_0 = \ell_0(d, k, \varepsilon) > 0$. Since $d < d_{\mathrm{uniq}}(k)$ and since $\mathbb{T} = \mathbb{T}_{d,k}$ is a Galton-Watson tree in which every variable node has $\mathrm{Po}(d)$ clause nodes as offspring and the offspring of every clause node consists of $k-1$ variable nodes, there exists a set $\mathscr{T}_\ell$ of trees, with $|\mathscr{T}_\ell| = O(1)$, such that the following hold:

T0: for every $T \in \mathscr{T}_\ell$ we have $\mathbb{P}\left[\mathbb{T}^{(\ell)} = T\right] > 0$.

T1: $\mathbb{P}\left[\mathbb{T}^{(\ell)} \in \mathscr{T}_\ell\right] > 1 - \varepsilon$.

T2: given $\mathbb{T}^{(\ell)} \in \mathscr{T}_\ell$ we have

$$\max_{\tau \in S(\mathbb{T}^{(\ell)})} \left|\mathbb{P}\left[\boldsymbol{\tau}^{(\ell)}(\mathfrak{x}) = 1 \mid \mathbb{T}^{(\ell)}\right] - \mathbb{P}\left[\boldsymbol{\tau}^{(\ell)}(\mathfrak{x}) = 1 \mid \mathbb{T}^{(\ell)},\ \forall x \in \partial^{2\ell}\mathfrak{x} : \boldsymbol{\tau}^{(\ell)}(x) = \tau(x)\right]\right| < \varepsilon.$$

For a variable node x_i of $\boldsymbol{\Phi}'$ obtain $\boldsymbol{\phi}'_\ell(x_i)$ from $\boldsymbol{\Phi}'$ by deleting all variables and clauses at distance greater than 2ℓ from x_i. We consider x_i being the root of $\boldsymbol{\phi}'_\ell(x_i)$. Moreover, for a tree $T \in \mathscr{T}_\ell$ let $\mathscr{V}_T$ be the set of variable nodes x_i, $1 \le i \le n$, such that $\boldsymbol{\phi}'_\ell(x_i) \cong T$; thus, there is an isomorphism of the CNFs T and $\boldsymbol{\phi}'_\ell(x_i)$ that maps the root $\mathfrak{x}$ of T to x_i. Consider the event

$$\mathfrak{T}_\ell = \left\{\sum_{T \in \mathscr{T}_\ell} \left|\mathbb{P}\left[\mathbb{T}^{(\ell)} \cong T\right] - |\mathscr{V}_T|/n\right| < \varepsilon\right\}. \tag{6.16}$$

Then Corollary 3.7 implies that

$$\mathbb{P}\left[\mathfrak{T}_\ell\right] = 1 - o(1) \quad \text{for every } \ell \ge 0. \tag{6.17}$$

We now claim that

$$\left|\mathbb{P}\left[\boldsymbol{\tau}^{(\ell)}(\mathfrak{x}) = 1 \mid \mathbb{T}^{(\ell)} = T\right] - \mathbb{P}\left[\boldsymbol{\sigma}_{\boldsymbol{\Phi}'}(x_i) = 1 \mid \boldsymbol{\Phi}'\right]\right| < \varepsilon \quad \text{for all } T \in \mathscr{T}_\ell,\ x_i \in \mathscr{V}_T. \tag{6.18}$$

To see this, let $S_\ell(\boldsymbol{\Phi}', x_i)$ be the set of all assignments $\sigma \in \{\pm 1\}^{\partial^{2\ell} x_i}$ of the variables at distance 2ℓ from x_i in $\boldsymbol{\Phi}'$ such that there exists a satisfying assignment $\sigma' \in S(\boldsymbol{\Phi}')$ with $\sigma'(y) = \sigma(y)$ for all $y \in \partial^{2\ell} x_i$. Then the law of total probability shows that

$$\mathbb{P}\left[\boldsymbol{\sigma}_{\boldsymbol{\Phi}'}(x_i) = 1 \mid \boldsymbol{\Phi}'\right] = \sum_{\sigma \in S_\ell(\boldsymbol{\Phi}', x_i)} \mathbb{P}\left[\boldsymbol{\sigma}_{\boldsymbol{\Phi}'}(x_i) = 1 \mid \boldsymbol{\Phi}',\ \forall y \in \partial^{2\ell} x_i : \boldsymbol{\sigma}_{\boldsymbol{\Phi}'}(y) = \sigma(y)\right] \mathbb{P}\left[\forall y \in \partial^{2\ell} x_i : \boldsymbol{\sigma}_{\boldsymbol{\Phi}'}(y) = \sigma(y) \mid \boldsymbol{\Phi}'\right]. \tag{6.19}$$

Further, since for $T \in \mathscr{T}_\ell$ and $x_i \in \mathscr{V}_T$ we have $\boldsymbol{\phi}'_\ell(x_i) \cong T$, condition **T2** implies that

$$\left|\mathbb{P}\left[\boldsymbol{\sigma}_{\boldsymbol{\Phi}'}(x_i) = 1 \mid \boldsymbol{\Phi}',\ \forall y \in \partial^{2\ell} x_i : \boldsymbol{\sigma}_{\boldsymbol{\Phi}'}(y) = \sigma(y)\right] - \mathbb{P}\left[\boldsymbol{\tau}^{(\ell)}(\mathfrak{x}) = 1 \mid \mathbb{T}^{(\ell)} = T\right]\right| < \varepsilon. \tag{6.20}$$

Combining (6.19) and (6.20), we obtain (6.18).

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

To complete the proof, we recall from Fact 4.2 that $\pi^{(\ell)}_{d,k}$ is precisely the distribution of $\mathbb{P}\left[\boldsymbol{\tau}^{(\ell)}(\mathfrak{x}) = 1 \mid \mathbb{T}^{(\ell)}\right]$. Therefore, coupling the formulas $\mathbb{T}^{(\ell)}$, $\boldsymbol{\Phi}'$ on the event $\mathfrak{T}_\ell$ we have

$$\begin{aligned} W_1(\boldsymbol{\pi}'_n, \pi^{(\ell)}_{d,k}) &\le \mathbb{P}\left[\mathbb{T}^{(\ell)} \notin \mathscr{T}_\ell\right] + \frac{1}{n}\sum_{T\in\mathscr{T}_\ell}\sum_{x\in\mathscr{V}_T}\left|\mathbb{P}\left[\boldsymbol{\tau}^{(\ell)}(\mathfrak{x}) = 1 \mid \mathbb{T}^{(\ell)} = T\right]\right. \\ &\quad \left. -\mathbb{P}\left[\boldsymbol{\sigma}_{\boldsymbol{\Phi}'}(x) = 1 \mid \boldsymbol{\Phi}'\right]\right| + \varepsilon && \text{[by (6.16)]} \\ &\le 3\varepsilon && \text{[by } \mathbf{T1} \text{ and (6.18)].} \end{aligned}$$

Combining this bound with (6.17) completes the proof. □

Proof of Proposition 2.11. The first assertion follows from Proposition 2.1, Lemma 6.3 and the fact that, since $0 < d < d_{\mathrm{uniq}}(k) < d_{\mathrm{sat}}(k)$, we have that $\mathbb{P}\left[Z(\boldsymbol{\Phi}') > 0\right] = 1 - o(1)$.

The second follows a routine argument, which we present below for the case $\ell = 2$ and it is standard to extend to any finite ℓ (see [29, Proposition 2.5]). Let $t = \Theta(\log\log n)$ and recall the definitions of $\boldsymbol{\phi}'_t(x_i)$, $\mathscr{T}_t$ and $S_t(\boldsymbol{\Phi}', x_i)$ from the proof of Lemma 6.3. Consider the event $\mathfrak{D} = \{\boldsymbol{\phi}'_t(x_1), \boldsymbol{\phi}'_t(x_2)$ are disjoint tree formulas$\}$.

From Lemma 3.5, we have that $\mathbb{P}[\mathfrak{D}] = 1 - o(1)$. On the event $\mathfrak{D}$, Lemma 6.3 implies that for every $\sigma_1, \sigma_2 \in \{\pm 1\}$, and $\tau_1 \in S_t(\boldsymbol{\Phi}', x_1)$, $\tau_2 \in S_t(\boldsymbol{\Phi}', x_2)$ we have

$$\left|\mathbb{P}\left[\boldsymbol{\sigma}_{\boldsymbol{\Phi}'}(x_i) = \sigma_i \mid \boldsymbol{\Phi}',\ \forall y \in \partial^{2t} x_i : \boldsymbol{\sigma}_{\boldsymbol{\Phi}'}(y) = \tau_i(y)\right] - \mathbb{P}\left[\boldsymbol{\sigma}_{\boldsymbol{\Phi}'}(x_i) = \sigma_i \mid \boldsymbol{\Phi}'\right]\right| = o(1), \quad \text{for } i = 1, 2. \tag{6.21}$$

Therefore, from the law of total probability and the triangle inequality we see that for every $\sigma_1, \sigma_2 \in \{\pm 1\}$

$$\begin{aligned} &\left|\mathbb{P}\left[\boldsymbol{\sigma}_{\boldsymbol{\Phi}'}(x_1) = \sigma_1, \boldsymbol{\sigma}_{\boldsymbol{\Phi}'}(x_2) = \sigma_2 \mid \boldsymbol{\Phi}'\right] - \mathbb{P}\left[\boldsymbol{\sigma}_{\boldsymbol{\Phi}'}(x_1) = \sigma_1 \mid \boldsymbol{\Phi}'\right]\mathbb{P}\left[\boldsymbol{\sigma}_{\boldsymbol{\Phi}'}(x_2) = \sigma_2 \mid \boldsymbol{\Phi}'\right]\right| \\ &\le \Big|\left|\mathbb{P}\left[\boldsymbol{\sigma}_{\boldsymbol{\Phi}'}(x_1) = \sigma_1 \mid \boldsymbol{\Phi}', \boldsymbol{\sigma}_{\boldsymbol{\Phi}'}(x_2) = \sigma_2\right] - \mathbb{E}_{\tau_1,\tau_2}\left[\mathbb{P}\left[\boldsymbol{\sigma}_{\boldsymbol{\Phi}'}(x_1) = \sigma_1, \boldsymbol{\sigma}_{\boldsymbol{\Phi}'}(x_2) = \sigma_2 \mid \boldsymbol{\Phi}', \tau_1, \tau_2\right]\right]\right| \\ &\quad - \left|\mathbb{E}_{\tau_1,\tau_2}\left[\mathbb{P}\left[\boldsymbol{\sigma}_{\boldsymbol{\Phi}'}(x_1) = \sigma_1 \mid \boldsymbol{\Phi}', \tau_1\right]\mathbb{P}\left[\boldsymbol{\sigma}_{\boldsymbol{\Phi}'}(x_2) = \sigma_2 \mid \boldsymbol{\Phi}', \tau_2\right]\right] - \mathbb{P}\left[\boldsymbol{\sigma}_{\boldsymbol{\Phi}'}(x_2) = \sigma_2 \mid \boldsymbol{\Phi}'\right]\right. \\ &\qquad \left.\mathbb{P}\left[\boldsymbol{\sigma}_{\boldsymbol{\Phi}'}(x_2) = \sigma_2 \mid \boldsymbol{\Phi}'\right]\right|\Big| \\ &= \Big|\left|\mathbb{E}_{\tau_1,\tau_2}\left[\mathbb{P}\left[\boldsymbol{\sigma}_{\boldsymbol{\Phi}'}(x_1) = \sigma_1 \mid \boldsymbol{\Phi}', \tau_1\right]\mathbb{P}\left[\boldsymbol{\sigma}_{\boldsymbol{\Phi}'}(x_2) = \sigma_2 \mid \boldsymbol{\Phi}', \tau_2\right]\right] - \mathbb{P}\left[\boldsymbol{\sigma}_{\boldsymbol{\Phi}'}(x_2) = \sigma_2 \mid \boldsymbol{\Phi}'\right]\right. \\ &\qquad \left.\mathbb{P}\left[\boldsymbol{\sigma}_{\boldsymbol{\Phi}'}(x_2) = \sigma_2 \mid \boldsymbol{\Phi}'\right]\right|\Big| \\ &\le \mathbb{E}_{\tau_1}\left|\mathbb{P}\left[\boldsymbol{\sigma}_{\boldsymbol{\Phi}'}(x_1) = \sigma_1 \mid \boldsymbol{\Phi}', \tau_1\right] - \mathbb{P}\left[\boldsymbol{\sigma}_{\boldsymbol{\Phi}'}(x_1) = \sigma_1 \mid \boldsymbol{\Phi}'\right]\right| + \mathbb{E}_{\tau_2}\left|\mathbb{P}\left[\boldsymbol{\sigma}_{\boldsymbol{\Phi}'}(x_2) = \sigma_2 \mid \boldsymbol{\Phi}', \tau_2\right]\right. \\ &\quad \left.-\mathbb{P}\left[\boldsymbol{\sigma}_{\boldsymbol{\Phi}'}(x_2) = \sigma_2 \mid \boldsymbol{\Phi}'\right]\right| \\ &= o(1). && \text{[by (6.21)]} \end{aligned}$$

Summing over the four sign combinations of σ_1, σ_2 gives the desired result. □

6.3 Proof of Proposition 2.5

As in the proof of Lemma 6.1 let $c_1, \ldots, c_{\boldsymbol{\Delta}''}$ be the new clauses added by **CPL2** and let $x_{1,1}, \ldots, x_{1,k}, \ldots, x_{\boldsymbol{\Delta}'',1}, \ldots, x_{\boldsymbol{\Delta}'',k}$ be their constituent variables. Let $\mathscr{X} = \{x_{1,1}, \ldots, x_{1,k}, \ldots, x_{\boldsymbol{\Delta}'',1}, \ldots, x_{\boldsymbol{\Delta}'',k}\}$. For $\varepsilon > 0$ and $z \in \mathbb{R}$ define $\lambda_\varepsilon(z) = \log(z \vee \varepsilon)$. Finally, let $(s_i)_{i\ge 0}$ be a sequence of uniformly random ± 1-valued random variables, mutually independent and independent of all other randomness.

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

Lemma 6.4. *Assume that $d < d_{\mathrm{uniq}}(k)$. There exists $B = B(d,k) > 0$ such that for all $0 < \varepsilon < 1$ we have*

$$\limsup_{n\to\infty} \mathbb{E}\left[\left(\sum_{i=1}^{\boldsymbol{\Delta}''} \lambda_\varepsilon\left(1 - \prod_{j=1}^{k} \mathbb{P}\left[\boldsymbol{\sigma}(x_{i,j}) \neq \mathrm{sign}(x_{i,j}, c_i) \mid \boldsymbol{\Phi}'\right]\right)\right)^2 \mid Z(\boldsymbol{\Phi}') > 0\right] \leq B.$$

Proof. Given $Z(\boldsymbol{\Phi}') > 0$ we have

$$0 \geq \lambda_\varepsilon\left(1 - \prod_{j=1}^{k} \mathbb{P}\left[\boldsymbol{\sigma}(x_{1,j}) \neq \mathrm{sign}(x_{1,j}, c_1) \mid \boldsymbol{\Phi}'\right]\right) \geq \lambda_\varepsilon\left(1 - \mathbb{P}\left[\boldsymbol{\sigma}(x_{1,1}) \neq \mathrm{sign}(x_{1,1}, c_1) \mid \boldsymbol{\Phi}'\right]\right). \tag{6.22}$$

Recalling that $\boldsymbol{\Delta}'' \sim \mathrm{Po}(d(k-1)/k)$, we combine (6.22) with Cauchy-Schwarz to obtain $B' = B'(d,k) > 0$ such that

$$\begin{aligned}\mathbb{E}\left[\left(\sum_{i=1}^{\boldsymbol{\Delta}''} \lambda_\varepsilon\left(1 - \prod_{j=1}^{k} \mathbb{P}\left[\boldsymbol{\sigma}(x_{i,j}) \neq \mathrm{sign}(x_{i,j}, c_i) \mid \boldsymbol{\Phi}'\right]\right)\right)^2 \mid Z(\boldsymbol{\Phi}') > 0\right] \\ \leq B' \cdot \mathbb{E}\left[\lambda_\varepsilon\left(1 - \mathbb{P}\left[\boldsymbol{\sigma}(x_{1,1}) \neq \mathrm{sign}(x_{1,1}, c_1) \mid \boldsymbol{\Phi}'\right]\right)^2 \mid Z(\boldsymbol{\Phi}') > 0\right].\end{aligned} \tag{6.23}$$

Further, since the function λ_ε is bounded and continuous for every $\varepsilon > 0$ and since $\mathrm{sign}(x_{1,1}, c_1)$ is chosen independently of $\boldsymbol{\Phi}'$, Proposition 2.11 shows that for any $\varepsilon > 0$,

$$\begin{aligned}\mathbb{E}\left[\lambda_\varepsilon\left(1 - \mathbb{P}\left[\boldsymbol{\sigma}(x_{1,1}) \neq \mathrm{sign}(x_{1,1}, c_1) \mid \boldsymbol{\Phi}'\right]\right)^2 \mid Z(\boldsymbol{\Phi}') > 0\right] &= \mathbb{E}\left[\lambda_\varepsilon\left(\boldsymbol{\mu}_{\pi_{d,k},1,1}\right)^2\right] + o(1) \\ &\leq \mathbb{E}\left[\log\left(\boldsymbol{\mu}_{\pi_{d,k},1,1}\right)^2\right] + o(1).\end{aligned} \tag{6.24}$$

Since Proposition 2.1 shows that $\mathbb{E}\left[\log\left(\boldsymbol{\mu}_{\pi_{d,k},1,1}\right)^2\right] = O(1)$, the assertion follows from (6.23) and (6.24). □

Lemma 6.5. *Assume that $d < d_{\mathrm{uniq}}(k)$. For any $\delta > 0$ there exists $\varepsilon_0 > 0$ such that for all $\varepsilon_0 > \varepsilon > 0$ we have*

$$\limsup_{n\to\infty} \left|\mathbb{E}\left[\log \frac{Z(\boldsymbol{\Phi}'') \vee 1}{Z(\boldsymbol{\Phi}') \vee 1}\right] - \frac{d(k-1)}{k}\mathbb{E}\left[\lambda_\varepsilon\left(1 - \prod_{j=1}^{k} \mathbb{P}\left[\boldsymbol{\sigma}(x_j) = s_j \mid \boldsymbol{\Phi}'\right]\right) \mid Z(\boldsymbol{\Phi}') > 0\right]\right| < \delta.$$

Proof. We choose small enough $\xi = \xi(d,k,\delta) > \zeta(\xi) > \eta = \eta(\zeta) > \varepsilon_0 = \varepsilon_0(\eta) > 0$, let $0 < \varepsilon < \varepsilon_0$ and assume that $n \geq n_0(\varepsilon)$ is large enough. Also let $\gamma = \gamma(n) = o(1)$ be a sequence that tends to zero sufficiently slowly. Additionally, let $\mathfrak{E}$ be the event that all of the following conditions occur.

- **E1:** $Z(\boldsymbol{\Phi}') > 0$.
- **E2:** $\boldsymbol{\Delta}'' \leq \zeta^{-1}$.
- **E3:** $|\mathscr{X}| = k\boldsymbol{\Delta}''$.
- **E4:** $\max_{x\in\mathscr{X}, s\in\{\pm1\}} \mathbb{P}[\boldsymbol{\sigma}(x) = s \mid \boldsymbol{\Phi}'] \leq 1 - \eta$.
- **E5:** $\sum_{\tau\in\{\pm1\}^{\mathscr{X}}} \left|\mathbb{P}[\forall x \in \mathscr{X}: \boldsymbol{\sigma}(x) = \tau(x) \mid \boldsymbol{\Phi}'] - \prod_{x\in\mathscr{X}} \mathbb{P}[\boldsymbol{\sigma}(x) = \tau(x) \mid \boldsymbol{\Phi}']\right| < \gamma$.

We claim that

$$\mathbb{P}[\mathfrak{E}] \geq 1 - 2\xi + o(1). \tag{6.25}$$

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

Indeed, since $0 < d < d_{\mathrm{uniq}}(k) < d_{\mathrm{sat}}(k)$, we have that $\mathbb{P}\left[Z(\mathbf{\Phi}') > 0\right] = 1 - o(1)$. Moreover, since $\mathbf{\Delta}'' \sim \mathrm{Po}(d(k-1)/k)$, Markov's inequality shows that $\mathbb{P}\left[\mathbf{\Delta}'' > \zeta^{-1}\right] \leq \zeta d < \xi$. Further, since the new clauses $c_1, \ldots, c_{\mathbf{\Delta}''}$ are chosen independently, we have $\mathbb{P}\left[|\mathscr{X}| = k\mathbf{\Delta}'' \mid \mathbf{\Delta}'' \leq \zeta^{-1}\right] = 1 - O(1/n)$.

Moreover, per Proposition 2.11 we see that the joint distribution on the assignments over $\mathscr{X}$ must be approximately the product measure. The tails of the limiting distribution of the latter are controlled by (2.1). Therefore, for small enough η we should have

$$\mathbb{P}\left[\max_{x \in \mathscr{X}, s \in \{\pm 1\}} \mathbb{P}[\boldsymbol{\sigma}(x) = s \mid \mathbf{\Phi}'] \leq 1 - \eta \mid \mathbf{\Delta}'' \leq \zeta^{-1},\ Z(\mathbf{\Phi}') > 0\right] \geq 1 - \xi.$$

Similarly, Proposition 2.11 shows together with Markov's inequality that

$$\mathbb{P}\left[\mathbf{E5} \text{ occurs} \mid \mathbf{\Delta}'' \leq \zeta^{-1},\ Z(\mathbf{\Phi}') > 0\right] = 1 - o(1),$$

provided that $\gamma \to \infty$ sufficiently slowly. Thus, we obtain (6.25).

Furthermore, (6.25) implies together with Proposition 2.7 and Hölder's inequality that

$$\mathbb{E}\left|(1 - \mathbb{1}\mathfrak{E}) \cdot \log \frac{Z(\mathbf{\Phi}'')}{Z(\mathbf{\Phi}')}\right| \leq \delta/3 + o(1), \tag{6.26}$$

provided that $\xi = \xi(d, k, \delta) > 0$ is small enough. Analogously, (6.25), Lemma 6.4 and Cauchy-Schwarz yield

$$\mathbb{E}\left|(1 - \mathbb{1}\mathfrak{E})\lambda_\varepsilon\left(1 - \prod_{j=1}^{k} \mathbb{P}\left[\boldsymbol{\sigma}(x_j) = s_j \mid \mathbf{\Phi}'\right]\right)\right| \leq \delta/3 + o(1). \tag{6.27}$$

Thus, we confine ourselves to the event $\mathfrak{E}$, on which we have $Z(\mathbf{\Phi}'), Z(\mathbf{\Phi}'') > 0$ due to **E1**, **E3**, **E4** and **E5**. Hence,

$$\begin{aligned}
\log \frac{Z(\mathbf{\Phi}'') \vee 1}{Z(\mathbf{\Phi}') \vee 1} = \log \frac{Z(\mathbf{\Phi}'')}{Z(\mathbf{\Phi}')} &= \log \sum_{\tau \in \{\pm 1\}^{\mathscr{X}}} \mathbb{1}\left\{\tau \models c_1, \ldots, c_{\mathbf{\Delta}''}\right\} \mathbb{P}[\forall x \in \mathscr{X} : \boldsymbol{\sigma}(x) = \tau(x) \mid \mathbf{\Phi}'] \\
&= \log \sum_{\tau \in \{\pm 1\}^{\mathscr{X}}} \mathbb{1}\left\{\tau \models c_1, \ldots, c_{\mathbf{\Delta}''}\right\} \prod_{x \in \mathscr{X}} \mathbb{P}[\boldsymbol{\sigma}(x) = \tau(x) \mid \mathbf{\Phi}'] + o(1) \qquad [\text{by } \mathbf{E4}, \mathbf{E5}] \\
&= \sum_{i=1}^{\mathbf{\Delta}''} \log\left[1 - \prod_{j=1}^{k} \mathbb{P}\left[\boldsymbol{\sigma}(x_{i,j}) \neq \mathrm{sign}(x_{i,j}, c_i) \mid \mathbf{\Phi}'\right]\right] + o(1) \qquad [\text{by } \mathbf{E3}].
\end{aligned} \tag{6.28}$$

Further, **E4** ensures that for any $1 \leq i \leq \Delta''$,

$$\left|\log\left[1 - \prod_{j=1}^{k} \mathbb{P}\left[\boldsymbol{\sigma}(x_{i,j}) \neq \mathrm{sign}(x_{i,j}, c_i) \mid \mathbf{\Phi}'\right]\right] - \lambda_\varepsilon\left[1 - \prod_{j=1}^{k} \mathbb{P}\left[\boldsymbol{\sigma}(x_{i,j}) \neq \mathrm{sign}(x_{i,j}, c_i) \mid \mathbf{\Phi}'\right]\right]\right| < \xi. \tag{6.29}$$

Thus, combining (6.28) and (6.29), we obtain

$$\mathbb{E}\left|\mathbb{1}\mathfrak{E}\left(\log \frac{Z(\mathbf{\Phi}'') \vee 1}{Z(\mathbf{\Phi}') \vee 1} - \sum_{i=1}^{\mathbf{\Delta}''} \lambda_\varepsilon\left(1 - \prod_{j=1}^{k} \mathbb{P}\left[\boldsymbol{\sigma}(x_{i,j}) \neq \mathrm{sign}(x_{i,j}, c_i) \mid \mathbf{\Phi}'\right]\right)\right)\right| < \delta/3 + o(1). \tag{6.30}$$

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

Further, combining (6.26) and (6.30) with Lemma 6.4, we obtain

$$\left|\mathbb{E}\left[\log\frac{Z(\boldsymbol{\Phi}'')\vee 1}{Z(\boldsymbol{\Phi}')\vee 1}\right]-\mathbb{E}\left[\sum_{i=1}^{\boldsymbol{\Delta}''}\lambda_\varepsilon\left(1-\prod_{j=1}^{k}\mathbb{P}\left[\boldsymbol{\sigma}(x_{i,j})\neq\operatorname{sign}(x_{i,j},c_i)\mid\boldsymbol{\Phi}'\right]\right)\mid Z(\boldsymbol{\Phi}')>0\right]\right|<\delta+o(1). \tag{6.31}$$

Finally, since the clauses $c_1,\dots,c_{\boldsymbol{\Delta}''}$ are drawn uniformly and independently and since the distribution of $\boldsymbol{\Phi}'$ is invariant under permutation of the variable nodes, we find

$$\begin{aligned}&\mathbb{E}\left[\sum_{i=1}^{\boldsymbol{\Delta}''}\lambda_\varepsilon\left(1-\prod_{j=1}^{k}\mathbb{P}\left[\boldsymbol{\sigma}(x_{i,j})\neq\operatorname{sign}(x_{i,j},c_i)\mid\boldsymbol{\Phi}'\right]\mid Z(\boldsymbol{\Phi}')>0\right)\right]\\&\quad=\frac{d(k-1)}{k}\mathbb{E}\left[\lambda_\varepsilon\left(1-\prod_{j=1}^{k}\mathbb{P}\left[\boldsymbol{\sigma}(x_j)=s_j\right]\mid\boldsymbol{\Phi}'\right)\mid Z(\boldsymbol{\Phi}')>0\right].\end{aligned} \tag{6.32}$$

Combining (6.31) and (6.32) completes the proof. □

Proof of Proposition 2.5. Proposition 2.11 shows together with Lemma 6.5 that

$$\mathbb{E}\left[\log\frac{Z(\boldsymbol{\Phi}'')\vee 1}{Z(\boldsymbol{\Phi}')\vee 1}\right]=\frac{d(k-1)}{k}\mathbb{E}\left[\lambda_\varepsilon\left(1-\prod_{j=1}^{k}\boldsymbol{\mu}_{\pi_{d,k},1,j}\right)\right]+o_\varepsilon(1), \tag{6.33}$$

with $o_\varepsilon(1)$ hiding a term that vanishes in the limit $\varepsilon\to 0$. Furthermore, in light of (2.1) the monotone convergence theorem yields

$$\mathbb{E}\left[\log\left(1-\prod_{j=1}^{k}\boldsymbol{\mu}_{\pi_{d,k},1,j}\right)\right]=\lim_{\varepsilon\to 0}\mathbb{E}\left[\lambda_\varepsilon\left(1-\prod_{j=1}^{k}\boldsymbol{\mu}_{\pi_{d,k},1,j}\right)\right]. \tag{6.34}$$

The assertion follows from (6.33) and (6.34). □

6.4 *Proof of Proposition 2.6*

We adapt the steps from Section 6.3 to the coupling of $\boldsymbol{\Phi}'$, $\boldsymbol{\Phi}'''$. Recall that the latter is obtained by adding to $\boldsymbol{\Phi}'$ a single variable x_{n+1} along with $\boldsymbol{\Delta}'''$ clauses $b_1,\dots,b_{\boldsymbol{\Delta}'''}$ that each contain x_{n+1} and $k-1$ other variables. Thus, let $x_{1,1},\dots,x_{1,k-1},\dots,x_{\boldsymbol{\Delta}''',1},\dots,x_{\boldsymbol{\Delta}''',k-1}\in\{x_1,\dots,x_n\}$ be the variables other than x_{n+1} that appear in $b_1,\dots,b_{\boldsymbol{\Delta}'''}$ and let $\mathscr{X}=\{x_{1,1},\dots,x_{\boldsymbol{\Delta}''',k-1}\}$ be the set comprising all these variables.

Lemma 6.6. *Assume that* $0<d<d_{\mathrm{uniq}}(k)$. *There exists* $B=B(d,k)>0$ *such that for all* $0<\varepsilon<1$ *we have*

$$\limsup_{n\to\infty}\mathbb{E}\left[\lambda_\varepsilon\left(\sum_{s\in\{\pm1\}}\prod_{i=1}^{\boldsymbol{\Delta}'''}\left(1-\mathbb{1}\{\operatorname{sign}(x_{n+1},b_i)\neq s\}\prod_{j=1}^{k-1}\mathbb{P}[\boldsymbol{\sigma}(x_{i,j})\neq\operatorname{sign}(x_{i,j},b_i)\mid\boldsymbol{\Phi}']\right)\right)^2\right.$$
$$\left.\mid Z(\boldsymbol{\Phi}')>0\right]\leq B.$$

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

Proof. Given that $\boldsymbol{\Phi}'$ is satisfiable, and noticing that λ_ε is increasing, and $\varepsilon \in (0,1)$, we see that

$$\begin{aligned}
0 \wedge \lambda_\varepsilon &\left(\sum_{s\in\{\pm 1\}} \prod_{i=1}^{\boldsymbol{\Delta}'''} \left(1 - \mathbb{1}\{\operatorname{sign}(x_{n+1}, b_i) \neq s\} \prod_{j=1}^{k-1} \mathbb{P}[\boldsymbol{\sigma}(x_{i,j}) \neq \operatorname{sign}(x_{i,j}, b_i) \mid \boldsymbol{\Phi}'] \right) \right) \\
&\geq \lambda_\varepsilon \left(\sum_{s\in\{\pm 1\}} \prod_{i=1}^{\boldsymbol{\Delta}'''} 1 - \prod_{j=1}^{k-1} \mathbb{P}[\boldsymbol{\sigma}(x_{i,j}) \neq \operatorname{sign}(x_{i,j}, b_i) \mid \boldsymbol{\Phi}'] \right) \\
&\geq \lambda_\varepsilon \left(\prod_{i=1}^{\boldsymbol{\Delta}'''} 1 - \mathbb{P}[\boldsymbol{\sigma}(x_{i,1}) \neq \operatorname{sign}(x_{i,1}, b_i) \mid \boldsymbol{\Phi}'] \right) \\
&= \lambda_\varepsilon \left(\prod_{i=1}^{\boldsymbol{\Delta}'''} \mathbb{P}[\boldsymbol{\sigma}(x_{i,1}) = \operatorname{sign}(x_{i,1}, b_i) \mid \boldsymbol{\Phi}'] \right) \geq \sum_{i=1}^{\boldsymbol{\Delta}'''} \lambda_\varepsilon(\mathbb{P}[\boldsymbol{\sigma}(x_{i,1}) = \operatorname{sign}(x_{i,1}, b_i) \mid \boldsymbol{\Phi}']).
\end{aligned} \tag{6.35}$$

We also notice that

$$0 \vee \lambda_\varepsilon \left(\sum_{s\in\{\pm 1\}} \prod_{i=1}^{\boldsymbol{\Delta}'''} \left(1 - \mathbb{1}\{\operatorname{sign}(x_{n+1}, b_i) \neq s\} \prod_{j=1}^{k-1} \mathbb{P}[\boldsymbol{\sigma}(x_{i,j}) \neq \operatorname{sign}(x_{i,j}, b_i) \mid \boldsymbol{\Phi}'] \right) \right) < 1. \tag{6.36}$$

In light of the above, we now bound

$$\begin{aligned}
&\mathbb{E}\left[\lambda_\varepsilon \left(\sum_{s\in\{\pm 1\}} \prod_{i=1}^{\boldsymbol{\Delta}'''} \left(1 - \mathbb{1}\{\operatorname{sign}(x_{n+1}, b_i) \neq s\} \prod_{j=1}^{k-1} \mathbb{P}[\boldsymbol{\sigma}(x_{i,j}) \neq \operatorname{sign}(x_{i,j}, b_i) \mid \boldsymbol{\Phi}'] \right) \right)^2 \mid Z(\boldsymbol{\Phi}') > 0 \right] \\
&\leq \mathbb{E}\left[1 + \left(\sum_{i=1}^{\boldsymbol{\Delta}'''} \lambda_\varepsilon(\mathbb{P}[\boldsymbol{\sigma}(x_{i,1}) = \operatorname{sign}(x_{i,1}, b_i) \mid \boldsymbol{\Phi}'] \right)^2 \mid Z(\boldsymbol{\Phi}') > 0 \right] \qquad [\text{from (6.35), (6.36)}] \\
&\leq d(d+1)\mathbb{E}\big[1 + \big(\lambda_\varepsilon(\mathbb{P}[\boldsymbol{\sigma}(x_{1,1}) = \operatorname{sign}(x_{1,1}, b_1) \mid \boldsymbol{\Phi}'])\big)^2 \mid Z(\boldsymbol{\Phi}') > 0\big] \qquad [\boldsymbol{\Delta}''' \sim \mathrm{Po}(d)] \\
&\leq d(d+1)\big(1 + \mathbb{E}\big[\lambda_\varepsilon(\mathbb{P}[\boldsymbol{\sigma}(x_{1,1}) = \operatorname{sign}(x_{1,1}, b_1) \mid \boldsymbol{\Phi}'])^2 \mid Z(\boldsymbol{\Phi}') > 0\big]\big).
\end{aligned} \tag{6.37}$$

Further, Proposition 2.11 implies that for any $\varepsilon > 0$,

$$\begin{aligned}
\mathbb{E}\left[\lambda_\varepsilon(\mathbb{P}[\boldsymbol{\sigma}(x_{1,1}) = \operatorname{sign}(x_{1,1}, b_1) \mid \boldsymbol{\Phi}'])^2 \mid Z(\boldsymbol{\Phi}') > 0\right] &= \mathbb{E}\big[\lambda_\varepsilon(\boldsymbol{\mu}_{\pi_{d,k},1,1})^2\big] + o(1) \\
&\leq \mathbb{E}\big[\log^2 \boldsymbol{\mu}_{\pi_{d,k},1,1}\big] + o(1).
\end{aligned} \tag{6.38}$$

Finally, the assertion follows from (6.37) and (6.38). □

Lemma 6.7. *Assume that* $0 < d < d_{\mathrm{uniq}}(k)$. *For any* $\delta > 0$ *there exists* $\varepsilon_0 > 0$ *such that for all* $\varepsilon_0 > \varepsilon > 0$ *we have*

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

$$\limsup_{n\to\infty}\Bigg|\mathbb{E}\left[\log\frac{Z(\mathbf{\Phi}''')\vee 1}{Z(\mathbf{\Phi}')\vee 1}\right] - \mathbb{E}\left[\lambda_\varepsilon\left(\sum_{s\in\{\pm1\}}\left(\prod_{i=1}^{\mathbf{\Delta}'''}1-\mathbb{1}\{\mathrm{sign}(x_{n+1},b_i)\neq s\}\prod_{j=1}^{k-1}\mathbb{P}[\boldsymbol{\sigma}(x_{i,j})\neq\mathrm{sign}(x_{i,j},b_i)\mid\mathbf{\Phi}']\right)\right)^2 \mid Z(\mathbf{\Phi}')>0\right]\Bigg| < \delta.$$

Proof. Choose small enough $\xi=\xi(d,k,\delta)>\zeta(\xi)>\eta=\eta(\zeta)>\varepsilon_0=\varepsilon_0(\eta)>0$, let $0<\varepsilon<\varepsilon_0$, suppose that $n>n_0(\varepsilon)$ is sufficiently large and let $0<\gamma=\gamma(n)=o(1)$ be a sequence that converges to zero slowly. Let $\mathfrak{E}$ be the event that the following conditions occur.

E1: $Z(\mathbf{\Phi}')>0$.
E2: $\mathbf{\Delta}'''\le\zeta^{-1}$.
E3: $|\mathscr{X}|=(k-1)\mathbf{\Delta}'''$.
E4: $\max_{x\in\mathscr{X},s\in\{\pm1\}}\mathbb{P}[\boldsymbol{\sigma}(x)=s\mid\mathbf{\Phi}']\le 1-\eta$.
E5: $\sum_{\tau\in\{\pm1\}^{\mathscr{X}}}\left|\mathbb{P}[\forall x\in\mathscr{X}:\boldsymbol{\sigma}(x)=\tau(x)\mid\mathbf{\Phi}']-\prod_{x\in\mathscr{X}}\mathbb{P}[\boldsymbol{\sigma}(x)=\tau(x)\mid\mathbf{\Phi}']\right|<\gamma$.

As in the proof of Lemma 6.5 we find that

$$\mathbb{P}[\mathfrak{E}]\ge 1-2\xi+o(1). \tag{6.39}$$

Let

$$L_\varepsilon=\lambda_\varepsilon\left(\sum_{s\in\{\pm1\}}\prod_{i=1}^{\mathbf{\Delta}'''}\left(1-\mathbb{1}\{\mathrm{sign}(x_{n+1},b_i)\neq s\}\prod_{j=1}^{k-1}\mathbb{P}[\boldsymbol{\sigma}(x_{i,j})\neq\mathrm{sign}(x_{i,j},b_i)\mid\mathbf{\Phi}']\right)\right)$$

for brevity. Combining Proposition 2.7, Lemma 6.6 and (6.39) and using Hölder's inequality, we obtain

$$\mathbb{E}\left|(1-\mathbb{1}\mathfrak{E})\log\frac{Z(\mathbf{\Phi}'')}{Z(\mathbf{\Phi}')}\right|\le\delta/3+o(1),\qquad \mathbb{E}\left|(1-\mathbb{1}\mathfrak{E})L_\varepsilon\mid Z(\mathbf{\Phi}')>0\right|\le\delta/3+o(1). \tag{6.40}$$

Hence, we are left to compare $\mathbb{E}\left|\mathbb{1}\mathfrak{E}\cdot\log\frac{Z(\mathbf{\Phi}'')}{Z(\mathbf{\Phi}')}\right|$ and $\mathbb{E}\left|\mathbb{1}\mathfrak{E}\cdot L_\varepsilon\mid Z(\mathbf{\Phi}')>0\right|$. On the event $\mathfrak{E}$ we have $Z(\mathbf{\Phi}'),Z(\mathbf{\Phi}''')>0$. Consequently,

$$\begin{aligned}\log\frac{Z(\mathbf{\Phi}''')\vee 1}{Z(\mathbf{\Phi}')\vee 1}&=\log\frac{Z(\mathbf{\Phi}''')}{Z(\mathbf{\Phi}')}=\log\sum_{\tau\in\{\pm1\}^{\mathscr{X}\cup\{x_{n+1}\}}}\mathbb{1}\left\{\tau\models b_1,\ldots,b_{\mathbf{\Delta}'''}\right\}\mathbb{P}[\forall x\in\mathscr{X}:\boldsymbol{\sigma}(x)=\tau(x)\mid\mathbf{\Phi}']\\&=\log\sum_{\tau\in\{\pm1\}^{\mathscr{X}\cup\{x_{n+1}\}}}\mathbb{1}\left\{\tau\models b_1,\ldots,b_{\mathbf{\Delta}'''}\right\}\\&\quad\prod_{x\in\mathscr{X}}\mathbb{P}[\boldsymbol{\sigma}(x)=\tau(x)\mid\mathbf{\Phi}']+o(1)\qquad\text{[by \textbf{E4}, \textbf{E5}]}\\&=\log\left[\sum_{s\in\{\pm1\}}\prod_{i=1}^{\mathbf{\Delta}'''}\left(1-\mathbb{1}\left\{\mathrm{sign}(x_{n+1},b_i)\neq s\right\}\prod_{j=1}^{k-1}\mathbb{P}\left[\boldsymbol{\sigma}(x_{i,j})\neq\mathrm{sign}(x_{i,j},b_i)\mid\mathbf{\Phi}'\right]\right)\right]\\&\quad+o(1)\qquad\text{[by \textbf{E3}]}.\end{aligned} \tag{6.41}$$

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

Now, **E4** guarantees that

$$\log\left[\sum_{s\in\{\pm1\}}\prod_{i=1}^{\Delta'''}\left(1-\mathbb{1}\left\{\mathrm{sign}(x_{n+1},b_i)\neq s\right\}\prod_{j=1}^{k-1}\mathbb{P}\left[\boldsymbol{\sigma}(x_{i,j})\neq\mathrm{sign}(x_{i,j},b_i)\mid\boldsymbol{\Phi}'\right]\right)\right]=L_\varepsilon. \quad (6.42)$$

Therefore, we combine (6.41) and (6.42) to obtain

$$\mathbb{E}\left|\mathbb{1}\mathfrak{E}\left(\log\frac{Z(\boldsymbol{\Phi}''')\vee 1}{Z(\boldsymbol{\Phi}')\vee 1}-L_\varepsilon\right)\right|<\delta/3+o(1). \quad (6.43)$$

Finally, the assertion follows from (6.40) and (6.43). □

Proof of Proposition 2.6. Following similar steps as in the proof of Proposition 2.5, we see that the assertion follows from Lemma 6.7, Proposition 2.1, Proposition 2.11, and the dominated convergence theorem. □

Proof of Proposition 2.3. Immediate from Fact 2.4, Proposition 2.5 and Proposition 2.6. □

7. Proof of Proposition 2.15

7.1 Proof of Lemma 2.12

The proof is by induction on the height of the tree. The following claim summarises the main step of the induction.

Claim 7.1. *For all $\ell\geq 0$, all variables x of $\mathbb{T}^{(\ell)}$ and all satisfying assignments $\tau\in S(\mathbb{T}^{(\ell)})$ we have*

$$\frac{Z\big(\mathbb{T}_x^{(\ell)},\tau,\boldsymbol{\tau}^+(x)\big)}{Z\big(\mathbb{T}_x^{(\ell)},\tau\big)}\leq\frac{Z\big(\mathbb{T}_x^{(\ell)},\boldsymbol{\tau}^+,\boldsymbol{\tau}^+(x)\big)}{Z\big(\mathbb{T}_x^{(\ell)},\boldsymbol{\tau}^+\big)}. \quad (7.1)$$

Proof. For boundary variables $x\in\partial^{2\ell}\mathfrak{x}$ there is nothing to show because the r.h.s. of (7.1) equals one. Hence, consider a variable $x\in\partial^{2q}\mathfrak{x}$ for some $q<\ell$. If $Z(\mathbb{T}_x^{(\ell)},\tau,\boldsymbol{\tau}^+(x))=0$, then (7.1) is trivially satisfied. Hence, assume that $Z(\mathbb{T}_x^{(\ell)},\tau,\boldsymbol{\tau}^+(x))>0$.

Let $a_1^+,\ldots,a_g^+$ be the children (clauses) of x with $\mathrm{sign}(x,a_i^+)=\boldsymbol{\tau}^+(x)$. Also let $y_{11},\ldots,y_{1(k-1)},\ldots,y_{g1},\ldots,y_{g(k-1)}$ be the children (variables) of $a_1^+,\ldots,a_g^+$. Similarly, let $a_1^-,\ldots,a_h^-$ be the children of x with $\mathrm{sign}(x,a_i^-)=-\boldsymbol{\tau}^+(x)$ and let $z_{11},\ldots,z_{1(k-1)},\ldots,z_{h1},\ldots,z_{h(k-1)}$ be their children. We claim that for all $\tau\in S(\mathbb{T}^{(\ell)})$,

$$Z\big(\mathbb{T}_x^{(\ell)},\tau,\boldsymbol{\tau}^+(x)\big)=\left(\prod_{i=1}^{g}\prod_{t=1}^{k-1}Z\big(\mathbb{T}_{y_{it}}^{(\ell)},\tau\big)\right)\cdot\prod_{j=1}^{h}\left(\prod_{t=1}^{k-1}Z\big(\mathbb{T}_{z_{jt}}^{(\ell)},\tau\big)-\prod_{t=1}^{k-1}Z\big(\mathbb{T}_{z_{jt}}^{(\ell)},\tau,-\boldsymbol{\tau}^+(z_{jt})\big)\right), \quad (7.2)$$

$$Z\big(\mathbb{T}_x^{(\ell)},\tau,-\boldsymbol{\tau}^+(x)\big)=\prod_{i=1}^{g}\left(\prod_{t=1}^{k-1}Z\big(\mathbb{T}_{y_{it}}^{(\ell)},\tau\big)-\prod_{t=1}^{k-1}Z\big(\mathbb{T}_{y_{it}}^{(\ell)},\tau,\boldsymbol{\tau}^+(y_{it})\big)\right)\cdot\left(\prod_{j=1}^{h}\prod_{t=1}^{k-1}Z\big(\mathbb{T}_{z_{jt}}^{(\ell)},\tau\big)\right). \quad (7.3)$$

For setting x to $\boldsymbol{\tau}^+(x)$ satisfies $a_1^+,\ldots,a_g^+$; hence, arbitrary satisfying assignments of the sub-trees $\mathbb{T}_{y_{it}}^{(\ell)}$ can be combined, which explains the first product in (7.2). By contrast, upon assigning x the value $\boldsymbol{\tau}^+(x)$ we need to ensure that each of the clauses $a_1^-,\ldots,a_g^-$ are satisfied by at least one

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

variable other than x. This explains the second factor of (7.2). A similar argument yields (7.3). Dividing (7.3) by (7.2) and invoking the induction hypothesis (for $q+1$), we obtain

$$\begin{aligned}\frac{Z\big(\mathbb{T}_x^{(\ell)},\tau,-\boldsymbol{\tau}^+(x)\big)}{Z\big(\mathbb{T}_x^{(\ell)},\tau,\boldsymbol{\tau}^+(x)\big)} &= \prod_{i=1}^{g}\left(1-\prod_{t=1}^{k-1}\frac{Z\big(\mathbb{T}_{y_{it}}^{(\ell)},\tau,\boldsymbol{\tau}^+(y_{it})\big)}{Z\big(\mathbb{T}_{y_{it}}^{(\ell)},\tau\big)}\right)\cdot\prod_{j=1}^{h}\left(1-\prod_{t=1}^{k-1}\frac{Z\big(\mathbb{T}_{z_{jt}}^{(\ell)},\tau,-\boldsymbol{\tau}^+(z_i)\big)}{Z\big(\mathbb{T}_{z_{jt}}^{(\ell)},\tau\big)}\right)^{-1}\\ &\geq \prod_{i=1}^{g}\left(1-\prod_{t=1}^{k-1}\frac{Z\big(\mathbb{T}_{y_{it}}^{(\ell)},\boldsymbol{\tau}^+,\boldsymbol{\tau}^+(y_{it})\big)}{Z\big(\mathbb{T}_{y_{it}}^{(\ell)},\boldsymbol{\tau}^+\big)}\right)\cdot\prod_{j=1}^{h}\left(1-\prod_{t=1}^{k-1}\frac{Z\big(\mathbb{T}_{z_{jt}}^{(\ell)},\boldsymbol{\tau}^+,-\boldsymbol{\tau}^+(z_i)\big)}{Z\big(\mathbb{T}_{z_{jt}}^{(\ell)},\boldsymbol{\tau}^+\big)}\right)^{-1}\\ &= \frac{Z\big(\mathbb{T}_x^{(\ell)},\boldsymbol{\tau}^+,-\boldsymbol{\tau}^+(x)\big)}{Z\big(\mathbb{T}_x^{(\ell)},\boldsymbol{\tau}^+,\boldsymbol{\tau}^+(x)\big)},\end{aligned}$$

completing the induction. □

Proof of Lemma 2.12. Applying Claim 7.1 to $x=\mathfrak{x}$ completes the proof of Lemma 2.12. □

7.2 *Proof of Lemma 2.14*

We employ the PULP algorithm introduced in Section 2.3 and its analysis on the random tree from Section 3.2. Recall that given an initial set of literals $\mathscr{L}$, PULP returns a superset $\bar{\mathscr{L}}$ with the property that the partial assignment obtained from setting all literals of $\bar{\mathscr{L}}$ to true, leaves no clause with only unsatisfying literals. Let us write $\bar{\mathscr{L}}=\bar{\mathscr{L}}_{x,s}$ for the set returned by PULP algorithm, initialised with the literal set $\mathscr{L}=\{s\cdot x\}$.

Claim 7.2. *Let $0\leq t<\ell$ and assume that $x\in\partial_{\mathbb{T}}^{2t}\mathfrak{x}$, $s\in\{\pm1\}$, satisfy $|\bar{\mathscr{L}}_{x,s}|<\ell-t$. Then for all $\tau\in S(\mathbb{T}^{(\ell)})$*

$$Z\big(\mathbb{T}_x^{(\ell)},\tau\big)\leq 2^{|\bar{\mathscr{L}}_{x,s}|}\cdot Z\big(\mathbb{T}_x^{(\ell)},\tau,s\big). \tag{7.4}$$

Proof. Notice that under our assumption on the size of $\bar{\mathscr{L}}_{x,s}$, the assignment τ does not clash with the one imposed by PULP. The assertion therefore follows immediately from the same argument as in the proof of Lemma 2.8. □

Claim 7.3. *We have* $\lim_{t\to\infty}\mathbb{P}\left[|\partial_{\mathbb{T}}^{2t}\mathfrak{x}|>(200d\cdot(k-1))^t\right]=0$.

Proof. This is an immediate consequence of Lemma 3.2. □

Proof of Lemma 2.14. Assume that $\ell>ct^c$ for a large enough $c=c(d,k)>0$ and that $t>t_0=t_0(d,k)$ is sufficiently large. Then Corollary 3.3 shows that

$$\mathbb{P}\left[|\bar{\mathscr{L}}_{\mathfrak{x},\pm1}|\geq t^c\right]\leq\exp(-t^2). \tag{7.5}$$

Combining Claim 7.3 with (7.5) and using the union bound, we obtain a sequence $\varepsilon_t\to0$ such that

$$\mathbb{P}\left[\forall x\in\partial_{\mathbb{T}}^{2t}\mathfrak{x}:|\bar{\mathscr{L}}_{x,\pm1}|<t^c\right]\geq1-\varepsilon_t. \tag{7.6}$$

If $x\in\partial_{\mathbb{T}}^{2t}\mathfrak{x}$ satisfies $|\bar{\mathscr{L}}_{x,\pm1}|<t^c$ and $\ell>ct^c$, then Claim 7.2 yields that for all $x\in\partial^{2t}\mathfrak{x}$

$$\left|\boldsymbol{\eta}_x^{(\ell)}\right|\leq\log\frac{Z\big(\mathbb{T}_x^{(\ell)},\boldsymbol{\sigma}^+\big)}{Z\big(\mathbb{T}_x^{(\ell)},\boldsymbol{\sigma}^+,+1\big)}+\log\frac{Z\big(\mathbb{T}_x^{(\ell)},\boldsymbol{\sigma}^+\big)}{Z\big(\mathbb{T}_x^{(\ell)},\boldsymbol{\sigma}^+,-1\big)}\leq|\bar{\mathscr{L}}_{x,+1}|+|\bar{\mathscr{L}}_{x,-1}|\leq2t^c. \tag{7.7}$$

The result now follows from (7.6) and (7.7). □

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

7.3 Proof of Proposition 2.15

We focus on the operator $\mathrm{LL}^{\star}_{d,k}$ introduced in Section 2.5. Let $\rho = (\rho_{\bullet}, \rho_{\oplus}, \rho_{\ominus})$, and $\rho' = (\rho'_{\bullet}, \rho'_{\oplus}, \rho'_{\ominus})$ be two arbitrary triplets in $\mathscr{P}(-\infty, \infty] \times \mathscr{P}(0, +\infty] \times \mathscr{P}(-\infty, 0]$, and write $\hat{\rho} = (\hat{\rho}_{\bullet}\hat{\rho}_{\oplus}, \hat{\rho}_{\ominus})$ and $\hat{\rho}' = (\hat{\rho}'_{\bullet}\hat{\rho}'_{\oplus}, \hat{\rho}'_{\ominus})$ for the images $\mathrm{LL}^{\star}_{d,k}(\rho)$ and $\mathrm{LL}^{\star}_{d,k}(\rho')$, respectively. We wish to bound $\mathrm{dist}_d(\hat{\rho}, \hat{\rho}')$ in terms of $\mathrm{dist}_d(\rho, \rho')$.

To this end, we begin with bounding the W_1-distance separately for each of the coordinates $(\hat{\rho}_{\oplus}, \hat{\rho}'_{\oplus})$, $(\hat{\rho}_{\ominus}, \hat{\rho}'_{\ominus})$ and $(\hat{\rho}_{\bullet}, \hat{\rho}'_{\bullet})$. Observe that it is sufficient to consider only $W_1(\hat{\rho}_{\oplus}, \hat{\rho}'_{\oplus})$ and $W_1(\hat{\rho}_{\ominus}, \hat{\rho}'_{\ominus})$, as the triangle inequality implies that $W_1(\hat{\rho}_{\bullet}, \hat{\rho}'_{\bullet}) \le W_1(\hat{\rho}_{\oplus}, \hat{\rho}'_{\oplus}) + W_1(\hat{\rho}_{\ominus}, \hat{\rho}'_{\ominus})$.

To spell out our bounds, we need to introduce some additional notation. Recall that for $i, j \ge 1$ the random variables $\boldsymbol{\eta}_{\bullet,i,j}$, $\boldsymbol{\eta}_{\oplus,i,j}$, $\boldsymbol{\eta}_{\ominus,i,j}$ follow the law of $\rho_{\bullet}$, $\rho_{\oplus}$, $\rho_{\ominus}$, respectively. Similarly, let $\boldsymbol{\eta}'_{\bullet,i,j}$, $\boldsymbol{\eta}'_{\oplus,i,j}$, $\boldsymbol{\eta}'_{\ominus,i,j}$ be random variables with law $\rho'_{\bullet}$, $\rho'_{\oplus}$ and $\rho'_{\ominus}$, respectively. We denote with $\boldsymbol{\eta}^{\wedge}_{\bullet,i,j}$ the random variable $\boldsymbol{\eta}_{\bullet,i,j} \wedge \boldsymbol{\eta}'_{\bullet,i,j}$, and with $\boldsymbol{\eta}^{\vee}_{\bullet,i,j}$ the random variable $\boldsymbol{\eta}_{\bullet,i,j} \vee \boldsymbol{\eta}'_{\bullet,i,j}$. Similarly, we write $\boldsymbol{\eta}^{\wedge}_{\oplus,i,j} = \boldsymbol{\eta}_{\oplus,i,j} \wedge \boldsymbol{\eta}'_{\oplus,i,j}$ and $\boldsymbol{\eta}^{\vee}_{\oplus,i,j} = \boldsymbol{\eta}_{\oplus,i,j} \vee \boldsymbol{\eta}'_{\oplus,i,j}$, and also write $\boldsymbol{\eta}^{\wedge}_{\ominus,i,j} = \boldsymbol{\eta}_{\ominus,i,j} \wedge \boldsymbol{\eta}'_{\ominus,i,j}$ and $\boldsymbol{\eta}^{\vee}_{\ominus,i,j} = \boldsymbol{\eta}_{\ominus,i,j}, \boldsymbol{\eta}'_{\ominus,i,j}$.

Moreover, for a sign $\varepsilon \in \{\pm 1\}$ and a vector $r = (r_{\bullet}, r_{\oplus}, r_{\ominus}, r_{\circ})$ of non-negative integers with $r_{\bullet} + r_{\oplus} + r_{\ominus} + r_{\circ} = k - 1$ and $1 \le i \le r_{\bullet}$, $1 \le j \le r_{\oplus}$, $1 \le \ell \le r_{\ominus}$, we let

$$\begin{aligned}\mathscr{D}^{\bullet}_i(z, r;\varepsilon) &= \left| \frac{\partial}{\partial z} \log \left(1 - \frac{1}{2^{r_{\circ}}} \Gamma\big(\varepsilon\big(\boldsymbol{\eta}_{\bullet,1,1}, \ldots, \boldsymbol{\eta}_{\bullet,1,i-1}, z, \boldsymbol{\eta}'_{\bullet,1,i+1}, \ldots \boldsymbol{\eta}'_{\bullet,1,r_{\bullet}}\big)\big) \Gamma\big(\varepsilon\big(\boldsymbol{\eta}'_{\oplus,1,1}, \ldots, \boldsymbol{\eta}'_{\oplus,1,r_{\oplus}}\big)\big) \right.\right. \\ &\qquad \left.\left. \Gamma\big(\varepsilon\big(\boldsymbol{\eta}'_{\ominus,1,1}, \ldots, \boldsymbol{\eta}'_{\ominus,1,r_{\ominus}}\big)\big)\big) \right) \right|.\end{aligned}$$

Analogously, we define

$$\begin{aligned}\mathscr{D}^{\oplus}_j(z, r;\varepsilon) &= \left| \frac{\partial}{\partial z} \log \left(1 - \frac{1}{2^{r_{\circ}}} \Gamma\big(\varepsilon\big(\boldsymbol{\eta}_{\bullet,1,1}, \ldots \boldsymbol{\eta}_{\bullet,1,r_{\bullet}}\big)\big) \Gamma\big(\varepsilon\big(\boldsymbol{\eta}_{\oplus,1,1}, \ldots, \boldsymbol{\eta}_{\oplus,1,j-1}, z, \boldsymbol{\eta}'_{\oplus,1,j+1}, \ldots, \boldsymbol{\eta}'_{\oplus,1,r_{\oplus}}\big)\big) \right.\right. \\ &\qquad \left.\left. \Gamma\big(\varepsilon\big(\boldsymbol{\eta}'_{\ominus,1,1}, \ldots \boldsymbol{\eta}'_{\ominus,1,r_{\ominus}}\big)\big)\big) \right) \right|,\end{aligned}$$

$$\begin{aligned}\mathscr{D}^{\ominus}_{\ell}(z, r;\varepsilon) &= \left| \frac{\partial}{\partial z} \log \left(1 - \frac{1}{2^{r_{\circ}}} \Gamma\big(\varepsilon\big(\boldsymbol{\eta}_{\bullet,1,1}, \ldots \boldsymbol{\eta}_{\bullet,1,r_{\bullet}}\big)\big) \Gamma\big(\varepsilon\big(\boldsymbol{\eta}_{\oplus,1,1}, \ldots, \boldsymbol{\eta}_{\oplus,1,r_{\oplus}}\big)\big) \right.\right. \\ &\qquad \left.\left. \Gamma\big(\varepsilon\big(\boldsymbol{\eta}_{\ominus,1,1}, \ldots, \boldsymbol{\eta}_{\ominus,1,\ell-1}, z, \boldsymbol{\eta}'_{\ominus,1,\ell+1}, \ldots, \boldsymbol{\eta}'_{\oplus,1,r_{\ominus}}\big)\big)\big) \right) \right|.\end{aligned}$$

With the above notation in place, we are now ready to bound $W_1(\hat{\rho}_{\oplus}, \hat{\rho}'_{\oplus})$. For each of the pairs of distributions $(\rho_{\bullet}, \rho'_{\bullet})$, $(\rho_{\oplus}, \rho'_{\oplus})$, and $(\rho_{\ominus}, \rho'_{\ominus})$, fix an arbitrary coupling among its coordinates.

Lemma 7.4. *$W_1(\hat{\rho}_{\oplus}, \hat{\rho}'_{\oplus})$ is upper bounded by*

$$\begin{aligned}\frac{d/2}{1 - e^{-\frac{d}{2}}} \cdot \mathbb{E}\left[\sum_{i=1}^{r_{\bullet,1}} \int_{\boldsymbol{\eta}^{\wedge}_{\bullet,1,i}}^{\boldsymbol{\eta}^{\vee}_{\bullet,1,i}} \mathscr{D}^{\bullet}_i(w_i, r_1; +1)\mathrm{d}w_i + \sum_{j=1}^{r_{\oplus,1}} \int_{\boldsymbol{\eta}^{\wedge}_{\oplus,1,j}}^{\boldsymbol{\eta}^{\vee}_{\oplus,1,j}} \mathscr{D}^{\oplus}_j(y_j, r_1; +1)\mathrm{d}y_j \right. \\ \left. + \sum_{\ell=1}^{r_{\ominus,1}} \int_{\boldsymbol{\eta}^{\wedge}_{\ominus,1,\ell}}^{\boldsymbol{\eta}^{\vee}_{\ominus,1,\ell}} \mathscr{D}^{\ominus}_{\ell}(z_{\ell}, r_1; +1)\mathrm{d}z_{\ell} \right]. \qquad (7.8)\end{aligned}$$

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

Proof. Let us write $\Xi'_{i,j}(\varepsilon, r)$ for the expression in the r.h.s. of (2.19) where distribution ρ' is used instead of ρ, i..,

$$\Xi'_{i,j}(\varepsilon, r) = 1 - \frac{1}{2^{r_\circ}} \Gamma\left(\varepsilon\left(\boldsymbol{\eta}'_{\bullet,4i+j,1}, \ldots, \boldsymbol{\eta}'_{\bullet,4i+j,r_\bullet}\right)\right) \Gamma\left(\varepsilon\left(\boldsymbol{\eta}'_{\oplus,4i+j,1}, \ldots, \boldsymbol{\eta}'_{\oplus,4i+j,r_\oplus}\right)\right) \Gamma\left(\varepsilon\left(\boldsymbol{\eta}'_{\ominus,4i+j,1}, \ldots, \boldsymbol{\eta}'_{\ominus,4i+j,r_\ominus}\right)\right). \tag{7.9}$$

By identically coupling the number of clauses and the types of the children variables of each clause in $\hat{\rho}_\oplus$, $\hat{\rho}'_\oplus$, we see that by the definition of the W_1 norm,

$$W_1(\hat{\rho}_\oplus, \hat{\rho}'_\oplus) \leq \mathbb{E}\left[\left| - \sum_{i=1}^{d^\star_+} \log \frac{\Xi_{i,3}(+1, r_{4i+3})}{\Xi'_{i,3}(+1, r_{4i+3})} \right|\right].$$

Applying Wald's lemma, we further obtain

$$W_1(\hat{\rho}_\oplus, \hat{\rho}'_\oplus) \leq \frac{d/2}{1 - e^{-d/2}} \cdot \mathbb{E}\left[\left| \log \frac{\Xi_{1,3}(+1, r_7)}{\Xi'_{1,3}(+1, r_7)} \right|\right] = \frac{d/2}{1 - e^{-d/2}} \cdot \mathbb{E}\left[\left| \log \frac{\Xi_{0,1}(+1, r_1)}{\Xi'_{0,1}(+1, r_1)} \right|\right]. \tag{7.10}$$

Let us now focus on the expectation in the r.h.s. of (7.10). Recalling the definition of Ξ in (2.19), and the definition of Ξ' in (7.9), we expand

$$\log \frac{\Xi_{0,1}(+1, r_1)}{\Xi'_{0,1}(+1, r_1)} = \log \frac{1 - 2^{-r_\circ} \cdot \Gamma\left(\boldsymbol{\eta}_{\bullet,1,1}, \ldots, \boldsymbol{\eta}_{\bullet,1,r_{\bullet,1}}\right) \cdot \Gamma\left(\boldsymbol{\eta}_{\oplus,1,1}, \ldots, \boldsymbol{\eta}_{\oplus,1,r_{\oplus,1}}\right) \cdot \Gamma\left(\boldsymbol{\eta}_{\ominus,1,1}, \ldots, \boldsymbol{\eta}_{\ominus,1,r_{\ominus,1}}\right)}{1 - 2^{-r_\circ} \cdot \Gamma\left(\boldsymbol{\eta}'_{\bullet,1,1}, \ldots, \boldsymbol{\eta}'_{\bullet,1,r_{\bullet,1}}\right) \cdot \Gamma\left(\boldsymbol{\eta}'_{\oplus,1,1}, \ldots, \boldsymbol{\eta}'_{\oplus,1,r_{\oplus,1}}\right) \cdot \Gamma\left(\boldsymbol{\eta}'_{\ominus,1,1}, \ldots, \boldsymbol{\eta}'_{\ominus,1,r_{\ominus,1}}\right)}. \tag{7.11}$$

Telescoping over the arguments of the functions Γ in the r.h.s of (7.11), invoking the fundamental theorem of calculus for each term, and applying the triangle inequality we further obtain

$$\left| \log \frac{\Xi_{0,1}(+1, r_1)}{\Xi'_{0,1}(+1, r_1)} \right| \leq \sum_{i=1}^{r_{\bullet,1}} \left| \int_{\boldsymbol{\eta}'_{\bullet,1,i}}^{\boldsymbol{\eta}_{\bullet,1,i}} \mathscr{D}_i^\bullet(w_i, r_1; +1) \mathrm{d}w_i \right| + \sum_{j=1}^{r_{\oplus,1}} \left| \int_{\boldsymbol{\eta}'_{\oplus,1,j}}^{\boldsymbol{\eta}_{\oplus,1,j}} \mathscr{D}_j^\oplus(y_j, r_1; +1) \mathrm{d}y_j \right| + \sum_{\ell=1}^{r_{\ominus,1}} \left| \int_{\boldsymbol{\eta}'_{\ominus,1,\ell}}^{\boldsymbol{\eta}_{\ominus,1,\ell}} \mathscr{D}_\ell^\ominus(z_\ell, r_1; +1) \mathrm{d}z_\ell \right|.$$

Plugging the above into (7.10) gives the result. □

Following the same steps as above, but replacing '+1' with '−1', yields the corresponding bound for $W_1(\hat{\rho}_\ominus, \hat{\rho}'_\ominus)$.

Lemma 7.5. *$W_1(\hat{\rho}_\ominus, \hat{\rho}'_\ominus)$ is upper bounded by*

$$\frac{d/2}{1 - e^{-\frac{d}{2}}} \cdot \mathbb{E}\left[\sum_{i=1}^{r_{\bullet,1}} \int_{\boldsymbol{\eta}^\wedge_{\bullet,1,i}}^{\boldsymbol{\eta}^\vee_{\bullet,1,i}} \mathscr{D}_i^\bullet(w_i, r_1; -1) \mathrm{d}w_i + \sum_{j=1}^{r_{\oplus,1}} \int_{\boldsymbol{\eta}^\wedge_{\oplus,1,j}}^{\boldsymbol{\eta}^\vee_{\oplus,1,j}} \mathscr{D}_j^\oplus(y_j, r_1; -1) \mathrm{d}y_j + \sum_{\ell=1}^{r_{\ominus,1}} \int_{\boldsymbol{\eta}^\wedge_{\ominus,1,\ell}}^{\boldsymbol{\eta}^\vee_{\ominus,1,\ell}} \mathscr{D}_\ell^\ominus(z_\ell, r_1; -1) \mathrm{d}z_\ell \right]. \tag{7.12}$$

Exploiting the signs of the variables with types $\oplus$ and $\ominus$, we obtain the following bounds for each of the $\mathscr{D}$-functions. For $\lambda \in (0, 1]$, we define the real function $\psi_\lambda : [0, 1] \to \mathbb{R}$ as

$$\psi_\lambda(w) = \frac{\lambda \cdot w}{1 - \lambda \cdot w} \cdot (1 - w). \tag{7.13}$$

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

It is easy to check that $\psi_{\lambda'}(w) \le \psi_\lambda(w)$, for every $\lambda' \le \lambda$.

Claim 7.6. *For every* $r = (r_\bullet, r_\oplus, r_\ominus, r_\circ)$, *and* $i \in [r_\bullet]$ *we have that*

$$\mathscr{D}_i^\bullet(w_i, r; +1) \le \psi_{2^{-r_\circ - r_\ominus}}\left(\frac{1+\tanh(w_i/2)}{2}\right) \quad \text{and} \quad \mathscr{D}_i^\bullet(w_i, r; -1) \le \psi_{2^{-r_\circ - r_\oplus}}\left(\frac{1-\tanh(w_i/2)}{2}\right). \tag{7.14}$$

Similarly, we also have that for $j \in [r_\oplus]$,

$$\mathscr{D}_j^\oplus(y_j, r; +1) \le \psi_{2^{-r_\circ - r_\ominus}}\left(\frac{1+\tanh(y_j/2)}{2}\right) \quad \text{and}$$
$$\mathscr{D}_j^\oplus(y_j, r; -1) \le \psi_{2^{-r_\circ - (r_\oplus - 1)}}\left(\frac{1-\tanh(y_j/2)}{2}\right), \tag{7.15}$$

and for $\ell \in [r_\ominus]$

$$\mathscr{D}_\ell^\ominus(z_\ell, r; +1) \le \psi_{2^{-r_\bullet - (r_\ominus - 1)}}\left(\frac{1+\tanh(z_\ell/2)}{2}\right) \quad \text{and}$$
$$\mathscr{D}_\ell^\ominus(z_\ell, r; -1) \le \psi_{2^{-r_\bullet - r_\oplus}}\left(\frac{1-\tanh(z_\ell/2)}{2}\right). \tag{7.16}$$

Proof. We only prove the first inequality of (7.14) as the rest of them follow in a similar manner. A straightforward calculation shows that for $z \in \mathbb{R}^q$, $\varepsilon \in \{\pm 1\}$, and $i \in [q]$ we have

$$\frac{\partial}{\partial z_i}\Gamma(\varepsilon \cdot z) = \varepsilon \cdot \frac{1-\tanh(\varepsilon \cdot z_i/2)}{2} \cdot \Gamma(\varepsilon \cdot z). \tag{7.17}$$

Writing $K = 2^{-r_\circ}\,\Gamma\big(\boldsymbol{\eta}_{\bullet,1,1}, \ldots, \boldsymbol{\eta}_{\bullet,1,i-1}, \boldsymbol{\eta}'_{\bullet,1,i+1}, \ldots \boldsymbol{\eta}'_{\bullet,1,r_\bullet}\big)\,\Gamma\big(\boldsymbol{\eta}'_{\oplus,1,1}, \ldots, \boldsymbol{\eta}'_{\oplus,1,r_\oplus}\big)\Gamma\big(\boldsymbol{\eta}'_{\ominus,1,1}, \ldots, \boldsymbol{\eta}'_{\ominus,1,r_\ominus}\big)$, applying the chain rule, and using (7.17), we see that

$$\mathscr{D}_i^\bullet(w_i, r; +1) = \psi_K\left(\frac{1+\tanh(w_i/2)}{2}\right). \tag{7.18}$$

Using the fact that $\rho'_\ominus$ is supported in $(-\infty, 0]$, and that $\Gamma \le 1$, we obtain $K \le 2^{-r_\circ - r_\ominus}$. The monotonicity of ψ_λ with respect to the parameter λ concludes the proof. □

Using Claim 7.6, and maximising each of the functions ψ_λ appearing in (7.14)–(7.16), we can recover the bounds of [51]. To obtain sharper bounds, a natural idea is to optimise groups of summands, instead of optimising each $\mathscr{D}$-summand of $W_1(\hat{\rho}_\oplus, \hat{\rho}'_\oplus) + W_1(\hat{\rho}_\ominus, \hat{\rho}'_\ominus)$ in isolation. In particular, it is tempting to pair terms of the form $\mathscr{D}(\cdot, -1)$ with corresponding terms of the form $\mathscr{D}(\cdot, +1)$, as Lemma 7.7 suggests.

Lemma 7.7. *Let* $\phi_\lambda : [0,1] \to \mathbb{R}$ *to be the function* $\phi_\lambda(w) = \psi_\lambda(w) + \psi_\lambda(1-w)$. *For every* $\lambda \in (0,1]$, *we have that* $\phi_\lambda(w) \le \phi_\lambda(1/2) = \frac{\lambda/2}{1-\lambda/2}$, *for all* $w \in [0,1]$.

Proof. For $\lambda = 1$, we have that $\psi_\lambda(w) = w$ implying $\phi_\lambda(w) = 1$, and thus, the result holds trivially. Let now $\lambda \in (0,1)$. Differentiating ψ_λ gives

$$\psi'_\lambda(w) = \frac{\lambda^2 w^2 - 2\lambda w + \lambda}{\lambda^2 w^2 - 2\lambda w + 1} = 1 - \frac{1-\lambda}{(1-\lambda w)^2}.$$

Therefore,

$$\phi'_\lambda(w) = 1 - \frac{1-\lambda}{(1-\lambda w)^2} - \left(1 - \frac{1-\lambda}{(1-\lambda(1-w))^2}\right) = \frac{1-\lambda}{(1-\lambda(1-w))^2} - \frac{1-\lambda}{(1-\lambda w)^2}.$$

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

It is straightforward to check that the above expression has only one root at $w = 1/2$, being non-negative for $w \in [0, 1/2)$, and non-positive for $w \in (1/2, 1]$. Therefore, $\phi_\lambda(1/2) = \frac{\lambda/2}{1-\lambda/2}$ is the maximum value of ϕ_λ. □

However, directly applying Lemma 7.7 to $W_1(\hat{\rho}_\oplus, \hat{\rho}'_\oplus) + W_1(\hat{\rho}_\ominus, \hat{\rho}'_\ominus)$ seems hopeless, since in the bounds supplied by Claim 7.6, the parameters of the functions ψ bounding $\mathscr{D}(\cdot, r, +1)$-terms in $W_1\left(\hat{\rho}_\oplus, \hat{\rho}'_\oplus\right)$ are quite different from the parameters of the functions ψ bounding the corresponding $\mathscr{D}(\cdot, r, -1)$-terms in $W_1\left(\hat{\rho}_\ominus, \hat{\rho}'_\ominus\right)$.

The following lemma reveals, a somewhat unexpected, symmetry between $W_1(\hat{\rho}_\oplus, \hat{\rho}'_\oplus)$ and $W_1(\hat{\rho}_\ominus, \hat{\rho}'_\ominus)$, that facilitates our pairing strategy.

Some additional notation is in order. We denote with $\mathscr{R}(k)$ for the set of all vectors $r = (r_\bullet, r_\oplus, r_\ominus, r_\circ)$ of non-negative integer entries which sum to $k-1$. For every $r \in \mathscr{R}(k)$ we use the shorthand

$$P(r) = \frac{(k-1)!}{r_\bullet! r_\oplus! r_\ominus! r_\circ!} \cdot p_\bullet^{r_\bullet} p_\oplus^{r_\oplus} p_\ominus^{r_\ominus} p_\circ^{r_\circ},$$

where $p_\bullet, p_\oplus, p_\ominus, p_\circ$ are the probabilities defined in (2.18). Finally, we define

$$E_\bullet = \sum_{\substack{r \in \mathscr{R}(k) \\ r_\bullet \geq 1}} P(r) \cdot r_\bullet \cdot \mathbb{E}\left[\int_{\eta^\wedge_{\bullet,1,1}}^{\eta^\vee_{\bullet,1,1}} \phi_{2-r_\circ - r_\ominus}\left(\frac{1+\tanh(w/2)}{2}\right) \mathrm{d}w\right], \tag{7.19}$$

$$E_\oplus = \sum_{\substack{r \in \mathscr{R}(k) \\ r_\oplus \geq 1}} P(r) \cdot r_\oplus \cdot \mathbb{E}\left[\int_{\eta^\wedge_{\oplus,1,1}}^{\eta^\vee_{\oplus,1,1}} \phi_{2-r_\circ - r_\ominus}\left(\frac{1+\tanh(y/2)}{2}\right) \mathrm{d}y\right], \tag{7.20}$$

$$E_\ominus = \sum_{\substack{r \in \mathscr{R}(k) \\ r_\ominus \geq 1}} P(r) \cdot r_\ominus \cdot \mathbb{E}\left[\int_{\eta^\wedge_{\ominus,1,1}}^{\eta^\vee_{\ominus,1,1}} \phi_{2-r_\circ - r_\oplus}\left(\frac{1+\tanh(z/2)}{2}\right) \mathrm{d}z\right]. \tag{7.21}$$

Lemma 7.8. *We have that*

$$W_1(\hat{\rho}_\oplus, \hat{\rho}'_\oplus) + W_1(\hat{\rho}_\ominus, \hat{\rho}'_\ominus) \leq \frac{d/2}{1 - e^{-d/2}} \left(E_\bullet + E_\oplus + E_\ominus\right). \tag{7.22}$$

Proof. Expanding the expectation in (7.8) with respect to $r = \left(r_\bullet, r_\oplus, r_\ominus, r_\circ\right)$, and using the shorthand

$$E^\pm_\bullet(r) = \mathbb{E}\left[\int_{\eta^\wedge_{\bullet,1,1}}^{\eta^\vee_{\bullet,1,1}} \mathscr{D}_1^\bullet\left(w, r; \pm 1\right) \mathrm{d}w\right], E^\pm_\oplus(r) = \mathbb{E}\left[\int_{\eta^\wedge_{\oplus,1,1}}^{\eta^\vee_{\oplus,1,1}} \mathscr{D}_1^\oplus\left(y, r; \pm 1\right) \mathrm{d}y\right],$$
$$E^\pm_\ominus(r) = \mathbb{E}\left[\int_{\eta^\wedge_{\ominus,1,1}}^{\eta^\vee_{\ominus,1,1}} \mathscr{D}_1^\ominus\left(z, r; \pm \mathbf{1}\right) \mathrm{d}z\right],$$

we see that

$$W_1(\hat{\rho}_\oplus, \hat{\rho}'_\oplus) \leq \frac{d/2}{1-e^{-d/2}} \left(\sum_{r \in \mathscr{R}(k)} P(r) \cdot r_\bullet \cdot E^+_\bullet(r) + \sum_{r \in \mathscr{R}(k)} P(r) \cdot r_\oplus \cdot E^+_\oplus(r) + \sum_{r \in \mathscr{R}(k)} P(r) \cdot r_\ominus \cdot E^+_\ominus(r)\right)$$
$$= \frac{d/2}{1-e^{-d/2}} \left(\sum_{\substack{r \in \mathscr{R}(k) \\ r_\bullet \geq 1}} P(r) \cdot r_\bullet \cdot E^+_\bullet(r) + \sum_{\substack{r \in \mathscr{R}(k) \\ r_\oplus \geq 1}} P(r) \cdot r_\oplus \cdot E^+_\oplus(r) + \sum_{\substack{r \in \mathscr{R}(k) \\ r_\ominus \geq 1}} P(r) \cdot r_\ominus \cdot E^+_\ominus(r)\right). \tag{7.23}$$

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

In a similar manner, we derive

$$W_1(\hat{\rho}_\ominus, \hat{\rho}'_\ominus) \leq \frac{d/2}{1-e^{-d/2}} \left(\sum_{\substack{r\in\mathscr{R}(k)\\ r_\bullet \geq 1}} P(r)\cdot r_\bullet \cdot E^-_\bullet(r) + \sum_{\substack{r\in\mathscr{R}(k)\\ r_\oplus \geq 1}} P(r)\cdot r_\oplus \cdot E^-_\oplus(r) + \sum_{\substack{r\in\mathscr{R}(k)\\ r_\ominus \geq 1}} P(r)\cdot r_\ominus \cdot E^-_\ominus(r) \right). \tag{7.24}$$

Let us now consider the bound on the sum $W_1(\hat{\rho}_\oplus, \hat{\rho}'_\oplus) + W_1(\hat{\rho}_\ominus, \hat{\rho}'_\ominus)$ obtained by summing (7.23), (7.24). We next group each of the three sums in (7.23) with the corresponding sum in (7.24), carefully pairing their terms. Specifically, for the $\bullet$–sums we match the term of $\sum_r P(r)\cdot r_\bullet \cdot E^+_\bullet(r)$ corresponding to $r=(r_\bullet, r_\oplus, r_\ominus, r_\circ)$ with the term of $\sum_{r'} P(r')\cdot r'_\bullet \cdot E^-_\bullet(r')$ that corresponds to $r'=(r_\bullet, r_\ominus, r_\oplus, r_\circ)$. Since $r\mapsto r'$ is a bijection of $\mathscr{R}(k)\cap\{r{:}r_\bullet \geq 1\}$, and $r'_\bullet = r_\bullet$, and $P(r')=P(r)$ we see that

$$\sum_{\substack{r\in\mathscr{R}(k)\\ r_\bullet \geq 1}} P(r)\cdot r_\bullet \left(E^+_\bullet(r) + E^-_\bullet(r)\right) = \sum_{\substack{r\in\mathscr{R}(k)\\ r_\bullet \geq 1}} P(r)\cdot r_\bullet \left(E^+_\bullet(r) + E^-_\bullet(r')\right). \tag{7.25}$$

Invoking the bounds (7.14) of Claim 7.6, and recalling the definitions of ϕ, $E_\bullet$, we upper bound the r.h.s. of (7.25) by

$$\sum_{\substack{r\in\mathscr{R}(k)\\ r_\bullet \geq 1}} P(r) r_\bullet \left(\mathbb{E}\left[\int_{\eta^\wedge_{\bullet,1,1}}^{\eta^\vee_{\bullet,1,1}} \psi_{2-r_\circ-r_\ominus}\left(\frac{1+\tanh(w/2)}{2}\right) \mathrm{d}w \right] + \mathbb{E}\left[\int_{\eta^\wedge_{\bullet,1,1}}^{\eta^\vee_{\bullet,1,1}} \psi_{2-r_\circ-r_\ominus}\left(\frac{1-\tanh(w/2)}{2}\right) \mathrm{d}w \right] \right) = E_\bullet. \tag{7.26}$$

The matchings between the terms for the $\oplus, \ominus$–sums of (7.23), (7.24) are more delicate. In particular, for the $\oplus$–sum it turns out that we can pull off the same trick as above by pairing the term of $\sum_r P(r)\cdot r_\oplus \cdot E^+_\oplus(r)$ corresponding to the vector $r=(r_\bullet, r_\oplus, r_\ominus, r_\circ)$ with the term of $\sum_{r''} P(r'')\cdot r''_\oplus \cdot E^-_\oplus(r'')$ that corresponds to $r''=(r_\bullet, r_\ominus+1, r_\oplus-1, r_\circ)$. To see this, note that the mapping $r\mapsto r''$ is a bijection of $\mathscr{R}(k)\cap\{r{:}r_\oplus \geq 1\}$, leaving the quantity $P(r)\cdot r_\oplus$ invariant as

$$\begin{aligned} P(r)\cdot r_\oplus &= \frac{(k-1)!}{r_\bullet!(r_\oplus-1)!r_\ominus!r_\circ!}\cdot p_\bullet^{r_\bullet} p_\oplus^{r_\oplus} p_\ominus^{r_\ominus} p_\circ^{r_\circ} = \frac{(k-1)!}{r_\bullet!(r_\ominus+1)!(r_\oplus-1)!r_\circ!}\cdot p_\bullet^{r_\bullet} p_\oplus^{r_\oplus} p_\ominus^{r_\ominus} p_\circ^{r_\circ}\cdot(r_\ominus+1) \\ &= P(r'')\cdot r''_\oplus. \end{aligned}$$

Invoking the bounds (7.15) of Claim 7.6, recalling the definitions of ϕ, $E_\oplus$, and arguing as above, we obtain

$$\sum_{\substack{r\in\mathscr{R}(k)\\ r_\oplus \geq 1}} P(r)\cdot r_\oplus \left(E^+_\oplus(r) + E^-_\oplus(r'')\right) \leq E_\oplus. \tag{7.27}$$

Similarly, using the mapping $r\mapsto r'''$, with $r'''=(r_\bullet, r_\ominus-1, r_\oplus+1, r_\circ)$, and following the same steps as above, we get

$$\sum_{\substack{r\in\mathscr{R}(k)\\ r_\ominus \geq 1}} P(r)\cdot r_\ominus \left(E^+_\ominus(r) + E^-_\ominus(r''')\right) \leq E_\ominus. \tag{7.28}$$

Summing (7.26)–(7.28) concludes the proof. □

In light of the above, we are now ready to finish the proof of Proposition 2.15.

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

Proof of Proposition 2.15. Applying Lemma 7.7 on the function ϕ in the r.h.s. of (7.19) gives

$$\phi_{2^{-r_\circ-r_\ominus}}\left(\frac{1+\tanh(w/2)}{2}\right) \le \frac{2^{-r_\ominus-r_\circ-1}}{1-2^{-r_\ominus-r_\circ-1}} \le \left(\frac{1}{2}\right)^{r_\ominus+r_\circ}. \tag{7.29}$$

Plugging the above into (7.19) and applying the binomial theorem, further yields

$$E_\bullet \le (k-1)\cdot p_\bullet \left(1-\frac{e^{-\frac{d}{2}}}{2}\right)^{k-2} \mathbb{E}\left[|\boldsymbol{\eta}_{\bullet,1,1}-\boldsymbol{\eta}'_{\bullet,1,1}|\right]. \tag{7.30}$$

Working in a similar manner, we obtain

$$\begin{aligned} E_\oplus &\le (k-1)\cdot p_\oplus \left(1-\frac{e^{-\frac{d}{2}}}{2}\right)^{k-2} \mathbb{E}\left[|\boldsymbol{\eta}_{\oplus,1,1}-\boldsymbol{\eta}'_{\oplus,1,1}|\right], \quad \text{and} \\ E_\ominus &\le (k-1)\cdot p_\ominus \left(1-\frac{e^{-\frac{d}{2}}}{2}\right)^{k-2} \mathbb{E}\left[|\boldsymbol{\eta}_{\ominus,1,1}-\boldsymbol{\eta}'_{\ominus,1,1}|\right]. \end{aligned} \tag{7.31}$$

Finally, plugging the bounds (7.30), and (7.31) into (7.22) we see that $W_1(\hat{\rho}_\oplus, \hat{\rho}'_\oplus) + W_1(\hat{\rho}_\ominus, \hat{\rho}'_\ominus)$ is upper bounded by

$$\begin{aligned} \frac{d\cdot(k-1)}{2}\cdot\left(1-\frac{e^{-\frac{d}{2}}}{2}\right)^{k-2} \Big[&\left(1-e^{-\frac{d}{2}}\right)\mathbb{E}\left[|\boldsymbol{\eta}_{\bullet,1,1}-\boldsymbol{\eta}'_{\bullet,1,1}|\right] + e^{-\frac{d}{2}}\mathbb{E}\left[|\boldsymbol{\eta}_{\oplus,1,1}-\boldsymbol{\eta}'_{\oplus,1,1}|\right] \\ &+ e^{-\frac{d}{2}}\mathbb{E}\left[|\boldsymbol{\eta}_{\ominus,1,1}-\boldsymbol{\eta}'_{\ominus,1,1}|\right]\Big]. \end{aligned} \tag{7.32}$$

Recall that we established (7.32) assuming an arbitrary coupling between the coordinates of each pair of distributions $(\rho_\bullet, \rho'_\bullet)$, $(\rho_\oplus, \rho'_\oplus)$, and $(\rho_\ominus, \rho'_\ominus)$. Therefore, the definition of W_1 norm and (7.32), imply the first inequality below, while (7.33) follows by the definition (2.22) of dist_d

$$\begin{aligned} W_1(\hat{\rho}_\oplus, \hat{\rho}'_\oplus) + W_1(\hat{\rho}_\ominus, \hat{\rho}'_\ominus) &\le \frac{d(k-1)}{2}\left(1-\frac{e^{-\frac{d}{2}}}{2}\right)^{k-2} \Big[\left(1-e^{-\frac{d}{2}}\right) W_1(\rho_\bullet, \rho'_\bullet) \\ &\quad + e^{-\frac{d}{2}} W_1(\rho_\oplus, \rho'_\oplus) + e^{-\frac{d}{2}} W_1(\rho_\ominus, \rho'_\ominus)\Big] \\ &\le \frac{d(k-1)}{2}\left(1-\frac{e^{-\frac{d}{2}}}{2}\right)^{k-2} \mathrm{dist}_d(\rho, \rho'). \end{aligned} \tag{7.33}$$

Moreover, as per the triangle inequality we see that

$$W_1(\hat{\rho}_\bullet, \hat{\rho}'_\bullet) \le W_1(\hat{\rho}_\oplus, \hat{\rho}'_\oplus) + W_1(\hat{\rho}_\ominus, \hat{\rho}'_\ominus) \le \frac{d(k-1)}{2}\left(1-\frac{e^{-\frac{d}{2}}}{2}\right)^{k-2} \mathrm{dist}_d(\rho, \rho'). \tag{7.34}$$

Plugging the bounds (7.33) and (7.34) into the expression of $\mathrm{dist}_d\left(\hat{\rho}, \hat{\rho}'\right)$ yields

$$\begin{aligned} \mathrm{dist}_d\left(\hat{\rho}, \hat{\rho}'\right) &= (1-e^{-d/2})\cdot W_1(\hat{\rho}_\bullet, \hat{\rho}'_\bullet) + e^{-d/2}\left(W_1(\hat{\rho}_\oplus, \hat{\rho}'_\oplus) + W_1(\hat{\rho}_\ominus, \hat{\rho}'_\ominus)\right) \\ &\le \frac{d(k-1)}{2}\left(1-\frac{e^{-\frac{d}{2}}}{2}\right)^{k-2} \mathrm{dist}_d(\rho, \rho'). \end{aligned}$$

Recalling the definition of d_{con}, we see that for $d < d_{\mathrm{con}}(k)$, the operator $\mathrm{LL}^\star_{d,k}$ contracts with respect to the metric dist_d, as desired. □

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

7.4 Proof of Proposition 2.13

To get a handle on the $\boldsymbol{\eta}_x^{(\ell)}$ from (2.12), we show that these quantities can be calculated by propagating the extremal boundary condition $\boldsymbol{\sigma}^+$ bottom-up toward the root of the tree. Specifically, we consider the operator

$$\Lambda^+_{\mathbb{T}^{(\ell)}} : (-\infty,\infty]^{V(\mathbb{T}^{(\ell)})} \to (-\infty,\infty]^{V(\mathbb{T}^{(\ell)})}, \quad \eta \mapsto \hat{\eta} = \Lambda^+_{\mathbb{T}^{(\ell)}}(\eta),$$

defined as follows. For all $x \in \partial^{2\ell}\mathfrak{x}$ we set $\hat{\eta}_x = \infty$. Moreover, for a variable $x \in \partial^{2q}\mathfrak{x}$ with $q < \ell$ having children clauses $a_1, \ldots, a_t$, and grandchildren variables $y_{1,1}, \ldots, y_{1,(k-1)}, \ldots, y_{t,1}, \ldots, y_{t,(k-1)}$ we define

$$\hat{\eta}_x = -\sum_{i=1}^{t} \boldsymbol{\tau}^+(x)\mathrm{sign}(x,a_i) \cdot \log\left(1 - \Gamma\left(\boldsymbol{\tau}^+(x)\mathrm{sign}(x,a_i) \cdot (\eta_{y_{i,1}}, \ldots, \eta_{y_{1,(k-1)}})\right)\right). \tag{7.35}$$

It may not be apparent that the sum above is well-defined as $-\infty$ summands may manifest. The following lemma rules out such possibility and shows that the ℓ-fold iteration of $\Lambda^{+\,(\ell)}_{\mathbb{T}^{(\ell)}}$, initiated all-$(+\infty)$ yields $\boldsymbol{\eta}^{(\ell)} = (\boldsymbol{\eta}_x^{(\ell)})_{x \in V(\mathbb{T}^{(\ell)})}$.

Lemma 7.9. *The operator $\Lambda^+_{\mathbb{T}^{(\ell)}}$ is well-defined and $\Lambda^{+\,(t)}_{\mathbb{T}^{(\ell)}}(+\infty,\ldots,+\infty) = \boldsymbol{\eta}^{(\ell)}$ for every $t \geq \ell$.*

Proof. To show that $\Lambda^+_{\mathbb{T}^{(\ell)}}$ is well defined we verify that, in the notation of (7.35), $\hat{\eta}_x \in (-\infty,\infty]$ for all x. Indeed, in the expression on the r.h.s. of (7.35) a $\pm\infty$ summand can arise only from variables $y_{i,j}$ with $\eta_{y_{i,j}} = \infty$. But the definition of $\boldsymbol{\tau}^+$ ensures that such $y_{i,j}$ either render a zero summand if $\boldsymbol{\tau}^+(x)\mathrm{sign}(x,a_i) = -1$, or a $+\infty$ summand if $\boldsymbol{\tau}^+(x)\mathrm{sign}(x,a_i) = 1$. Thus, the sum is well-defined and $\hat{\eta}_x \in (-\infty,\infty]$.

Further, to verify the identity $\boldsymbol{\eta}^{(\ell)} = \Lambda^{+\,(\ell)}_{\mathbb{T}^{(\ell)}}(\infty,\ldots,\infty)$, consider a variable x of $\mathbb{T}^{(\ell)}$. Let $a_1^+, \ldots, a_g^+$ be the children (clauses) of x with $\mathrm{sign}(x,a_i^+) = \boldsymbol{\tau}^+(x)$. Also let $y_{11}, \ldots, y_{1(k-1)}$, $\ldots, y_{g1}, \ldots, y_{g(k-1)}$ be the children of $a_1^+, \ldots, a_g^+$. Similarly, let $a_1^-, \ldots, a_h^-$ be the children of x with $\mathrm{sign}(x,a_i^-) = -\boldsymbol{\tau}^+(x)$ and let $z_{11}, \ldots, z_{1(k-1)}, \ldots, z_{h1}, \ldots, z_{h(k-1)}$ be their children. Then (7.2), and (7.3) yield

$$\begin{aligned}
\boldsymbol{\eta}_x^{(\ell)} &= -\sum_{i=1}^{g} \log\left(1 - \prod_{q=1}^{k-1} \frac{Z\big(\mathbb{T}^{(\ell)}_{y_{iq}}, \boldsymbol{\tau}^+, \boldsymbol{\tau}^+(y_{iq})\big)}{Z\big(\mathbb{T}^{(\ell)}_{y_{iq}}, \boldsymbol{\tau}^+\big)}\right) + \sum_{j=1}^{h} \log\left(1 - \prod_{q=1}^{k-1} \frac{Z\big(\mathbb{T}^{(\ell)}_{z_{jq}}, \boldsymbol{\tau}^+, -\boldsymbol{\tau}^+(z_{jq})\big)}{Z\big(\mathbb{T}^{(\ell)}_{z_{jq}}, \boldsymbol{\tau}^+\big)}\right) \\
&= -\sum_{i=1}^{g} \log\left(1 - \Gamma\left(\mathrm{sign}(x,a_i^+)\boldsymbol{\tau}^+(x) \cdot \left(\boldsymbol{\eta}^{(\ell)}_{y_{i1}}, \ldots, \boldsymbol{\eta}^{(\ell)}_{y_{i(k-1)}}\right)\right)\right) \\
&\quad + \sum_{j=1}^{h} \log\left(1 - \Gamma\left(\mathrm{sign}(x,a_i^-)\boldsymbol{\tau}^+(x) \cdot \left(\boldsymbol{\eta}^{(\ell)}_{z_{j1}}, \ldots, \boldsymbol{\eta}^{(\ell)}_{z_{j(k-1)}}\right)\right)\right).
\end{aligned}$$

The assertion follows because $\mathrm{sign}(x,a_i^+)\boldsymbol{\tau}^+(x) = 1$ and $\mathrm{sign}(x,a_i^-)\boldsymbol{\tau}^+(x) = -1$. □

The next aim is to approximate the ℓ-fold iteration of $\Lambda^+_{\mathbb{T}^{(\ell)}}$, and more specifically the distribution of $\boldsymbol{\eta}^{(\ell)}_{\mathfrak{x}}$, using a non-random operator. To this end, we need to cope with the $\pm\infty$-entries of the vector $\boldsymbol{\eta}^{(\ell)}$. This is addressed by Lemma 2.14, proven in Section 7.2, which provides a bound on $\boldsymbol{\eta}_x^{(\ell)}$ for variables x near the root of the tree.

In the following we continue to write c and $(\varepsilon_t)_t$ for the number and the sequence supplied by Lemma 2.14. Guided by Lemma 2.14 we consider the vector $\boldsymbol{\eta}^{(\ell)}_{\wedge t}$ of truncated log-likelihood

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

ratios

$$\left(\boldsymbol{\eta}^{(\ell)}_{\wedge t}\right)_x = \begin{cases} -2t^c & \text{if } x \in \partial^{2t}\mathfrak{x} \text{ and } \boldsymbol{\eta}^{(\ell)}_x < -2t^c, \\ 2t^c & \text{if } x \in \partial^{2t}\mathfrak{x} \text{ and } \boldsymbol{\eta}^{(\ell)}_x > 2t^c, \\ \boldsymbol{\eta}^{(\ell)}_x & \text{otherwise.} \end{cases} \tag{7.36}$$

Further, let $\boldsymbol{\eta}^{(\ell,t)}$ be the result of t iterations of $\Lambda^{+}_{\mathbb{T}^{(\ell)}}(\cdot)$ starting from $\boldsymbol{\eta}^{(\ell)}_{\wedge t}$. The following corollary is a direct consequence of Lemma 7.9 and Lemma 2.14.

Corollary 7.10. *For any $\ell > ct^c$ we have $\mathbb{P}[\boldsymbol{\eta}^{(\ell,t)}_{\mathfrak{x}} \neq \boldsymbol{\eta}^{(\ell)}_{\mathfrak{x}}] < \varepsilon_t$.*

Proof. Due to Lemma 2.14, the truncation in (7.36) is inconsequential with probability at least $1 - \varepsilon_t$, in which case

$$\boldsymbol{\eta}^{(\ell,t)} = \Lambda^{+\,(t)}_{\mathbb{T}^{(\ell)}}\left(\boldsymbol{\eta}^{(\ell)}_{\wedge t}\right) = \Lambda^{+\,(t)}_{\mathbb{T}^{(\ell)}}(\boldsymbol{\eta}^{(\ell)}) = \Lambda^{+\,(\ell+t)}_{\mathbb{T}^{(\ell)}}(+\infty, \ldots, +\infty) = \boldsymbol{\eta}^{(\ell)},$$

where the last equality follows from Lemma 7.9. □

Recall that we defined the non-random operator $\mathrm{LL}^{\star}_{d,k}$ from (2.17), mimicking $\Lambda^{+}_{\mathbb{T}^{(\ell)}}$. To make the connection between the random operator $\Lambda^{+}_{\mathbb{T}^{(\ell)}}$ and $\mathrm{LL}^{\star}_{d,k}$ precise, we introduce the following concepts. Given a tree formula T we write $V_{\bullet}(T)$, for the set of x variables of T that appear both as positive and negative literals in the sub-tree T_x comprising x and its the progeny. We define $V_{\circ}(T)$, $V_{\oplus}(T)$, and $V_{\ominus}(T)$ similarly. Note that the above sets constitute a partition of $V(T)$. We use $\mathrm{tp}{:}V(T) \to \{\bullet, \oplus, \ominus, \circ\}$ to indicate the part each vertex belongs to. We denote with $\mathbb{T}^{(\ell)}_{\bullet}$ the random Galton-Watson formula $\mathbb{T}$ conditioned on the root satisfying $\mathrm{tp}(\mathfrak{x}) = \bullet$. We define $\mathbb{T}^{(\ell)}_{\oplus}$, and $\mathbb{T}^{(\ell)}_{\ominus}$ analogously. Degenerately, we also write $\mathbb{T}^{(\ell)}_{\circ}$ for the formula comprised by a single variable $\mathfrak{x}$. Let us denote with $\hat{\eta}^{(\ell,t)}_{\bullet}$ the distribution of $(\boldsymbol{\eta}^{(\ell)}_{\wedge t})_{\mathfrak{x}}$ in $\mathbb{T}^{(\ell)}_{\bullet}$. Moreover, let $\bar{\eta}^{(\ell-t)}_{\bullet}$ be the distribution of

$$\boldsymbol{\eta}^{(\ell-t)}_{\mathfrak{x}} \cdot \mathbb{1}\left\{|\boldsymbol{\eta}^{(\ell-t)}_{\mathfrak{x}}| \le 2t^c\right\} + 2t^c \cdot \mathbb{1}\left\{\boldsymbol{\eta}^{(\ell-t)}_{\mathfrak{x}} > 2t^c\right\} - 2t^c \cdot \mathbb{1}\left\{\boldsymbol{\eta}^{(\ell-t)}_{\mathfrak{x}} < -2t^c\right\},$$

i.e., the truncation of $\boldsymbol{\eta}^{(\ell-t)}_{\mathfrak{x}}$ in $\mathbb{T}^{(\ell)}_{\bullet}$. Analogously we define the distributions $\hat{\eta}^{(\ell,t)}_{\oplus}$, $\hat{\eta}^{(\ell,t)}_{\ominus}$, and $\bar{\eta}^{(\ell-t)}_{\oplus}$, $\bar{\eta}^{(\ell-t)}_{\ominus}$. Notice that, degenerately, $\hat{\eta}^{(\ell,t)}_{\circ} = \bar{\eta}^{(\ell-t)}_{\circ} = \delta_0$.

Lemma 7.11. *For $\ell > ct^c$ we have that $\left(\hat{\eta}^{(\ell,t)}_{\bullet}, \hat{\eta}^{(\ell,t)}_{\oplus}, \hat{\eta}^{(\ell,t)}_{\ominus}\right) = \mathrm{LL}^{\star}_{d,k}\left(\bar{\eta}^{(\ell-t)}_{\bullet}, \bar{\eta}^{(\ell-t)}_{\oplus}, \bar{\eta}^{(\ell-t)}_{\ominus}\right)$.*

Proof. We use induction on t. Specifically, let $\nu = \left(\nu_{\bullet}, \nu_{\oplus}, \nu_{\ominus}\right)$ be any triplet in $\mathscr{P}(-\infty, \infty] \times \mathscr{P}[0, +\infty] \times \mathscr{P}[-\infty, 0]$, and $\nu^{(t)} = \mathrm{LL}^{\star(t)}_{d,k}(\nu)$ be the outcome of the t-fold application of $\mathrm{LL}^{\star}_{d,k}$. Moreover, let $(\boldsymbol{\eta}_x)_{x \in V(\mathbb{T}^{(t)})}$ be a vector of independent samples with $\boldsymbol{\eta}_x \sim \nu_{\mathrm{tp}(x)}$. We claim that root value $\boldsymbol{\eta}^{(t)}_{\mathfrak{x}}$ of the random operator $\Lambda^{+\,(t)}_{\mathbb{T}^{(t)}}$, coincides with $\nu_{\mathrm{tp}(\mathfrak{x})}$. Indeed, for $t = 1$ the claim follows readily from the definitions. For the inductive step, we notice that the t-fold application of $\mathrm{LL}^{\star}_{d,k}$ is obtained by applying $\mathrm{LL}^{\star}_{d,k}$ to the $(t-1)$-fold application. Per the induction hypothesis

$$\left(\Lambda^{+\,(t-1)}_{\mathbb{T}^{(t-1)}}(\boldsymbol{\eta}_x)_x\right)_{\mathfrak{x}} \sim \nu^{(t-1)}_{\mathrm{tp}(\mathfrak{x})}. \tag{7.37}$$

Applying $\mathrm{LL}^{\star}_{d,k}$ to $\nu^{(t-1)}$ implies the result as the first layer of $\mathbb{T}^{(t)}$ is independent of the subtrees rooted at the grandchildren $\partial^2\mathfrak{x}$ of the root, which are distributed as i.i.d. copies of $\mathbb{T}^{(t-1)}$. The lemma follows from applying the above identity to $\nu = \left(\bar{\eta}^{(\ell-t)}_{\bullet}, \bar{\eta}^{(\ell-t)}_{\oplus}, \bar{\eta}^{(\ell-t)}_{\ominus}\right)$. □

Refining the definition of the $\mathrm{BP}_{d,k}$ operator in (1.3), we write $\mathrm{BP}^{\bullet}_{d,k}$ for the operator obtained from $\mathrm{BP}_{d,k}$ upon conditioning on $d^+, d^- \ge 1$. Similarly $\mathrm{BP}^{\oplus}_{d,k}$ and $\mathrm{BP}^{\ominus}_{d,k}$ are obtained from $\mathrm{BP}_{d,k}$

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

upon conditioning on $d^+ \geq 1, d^- = 0$, and $d^+ = 0, d^- \geq 1$, respectively. We define

$$\pi^{\bullet}_{d,k} = \mathrm{BP}^{\bullet}_{d,k}\left(\pi_{d,k}\right), \pi^{\oplus}_{d,k} = \mathrm{BP}^{\oplus}_{d,k}\left(\pi_{d,k}\right), \pi^{\ominus}_{d,k} = \mathrm{BP}^{\ominus}_{d,k}\left(\pi_{d,k}\right).$$

Let us write γ, γ^{-1} for the continuous and mutually inverse real functions

$$\gamma : \mathbb{R} \to (0,1), z \mapsto (1 + \tanh(z/2))/2, \qquad \gamma^{-1} : (0,1) \to \mathbb{R}, p \mapsto \log(p/(1-p)). \tag{7.38}$$

Let $\rho^{\bullet}_{d,k} = \gamma^{-1}(\pi^{\bullet}_{d,k})$, and define $\rho^{\oplus}_{d,k}, \rho^{\ominus}_{d,k}$ similarly.

Claim 7.12. *The vector $(\rho^{\bullet}_{d,k}, \rho^{\oplus}_{d,k}, \rho^{\ominus}_{d,k})$ is a fixed point of the operator $\mathrm{LL}^{\star}_{d,k}$.*

Proof. Let $\rho_{d,k} = \gamma^{-1}\left(\pi_{d,k}\right)$. First, we claim that

$$\mathrm{LL}^{\star}_{d,k}\left(\rho_{d,k}, \rho_{d,k}, \rho_{d,k}\right) = \left(\rho^{\bullet}_{d,k}, \rho^{\oplus}_{d,k}, \rho^{\ominus}_{d,k}\right). \tag{7.39}$$

Indeed, since all input distributions are the same, by Proposition 2.1, the two summands in the left term of (2.21) corresponding to $d^{\star}_{+}$ and $d^{\star}_{-}$ are identically distributed, and also identically distributed to the sums that appear in the other two terms. Therefore, (7.39) follows directly from the definitions of $\mathrm{BP}^{\bullet}_{d,k}$, $\mathrm{BP}^{\oplus}_{d,k}$, and $\mathrm{BP}^{\ominus}_{d,k}$. The claim now follows from (7.39), the definition of the operator $\mathrm{LL}^{\star}_{d,k}$, and the law of total probability. □

Let $\rho^{(\ell)}$ be the distribution of the log-likelihood ratio $\boldsymbol{\eta}^{(\ell)}_{\mathfrak{x}}$.

Corollary 7.13. *For $d < d_{\mathrm{con}}(k)$ the sequence $\left(\gamma\left(\rho^{(\ell)}\right)\right)_{\ell}$ converges weakly to $\pi_{d,k}$.*

Proof. The result follows by combining Corollary 7.10, Lemma 7.11, Proposition 2.15, Claim 7.12, and applying the continuous mapping theorem and the law of total probability. □

Proof of Proposition 2.13. Recall that we write $\Lambda^{+(\ell)}_{\mathbb{T}^{(\ell)}}$ for the ℓ-fold iteration of the operator $\Lambda^{+}_{\mathbb{T}}$. Let us write $\boldsymbol{\theta}^{(\ell)}_{\mathfrak{x}} = \left(\Lambda^{+(\ell)}_{\mathbb{T}^{(\ell)}}(0,\ldots,0)\right)_{\mathfrak{x}}$. Using arguments similar to Fact 4.2, we can show that $\boldsymbol{\theta}^{(\ell)}_{\mathfrak{x}}$ is nothing but the distribution of the random variable $\gamma^{-1}(\mathbb{P}\left[\boldsymbol{\tau}^{(\ell)}(\mathfrak{x}) = 1 \mid \mathbb{T}\right])$. Therefore,

$$\mathbb{P}[\boldsymbol{\tau}^{(\ell)}(\mathfrak{x}) = 1 \mid \mathbb{T}] \sim \gamma(\boldsymbol{\theta}^{(\ell)}_{\mathfrak{x}}), \quad \text{and} \quad \mathbb{P}[\boldsymbol{\tau}^{(\ell)}(\mathfrak{x}) = 1 \mid \forall y \in \partial^{2\ell}\mathfrak{x} : \boldsymbol{\tau}^{(\ell)}(y) = \boldsymbol{\tau}^{+}(y), \mathbb{T}] \sim \gamma(\boldsymbol{\eta}^{(\ell)}_{\mathfrak{x}}).$$

Due to Lemma 2.12, $0 \leq \gamma(\boldsymbol{\theta}^{(\ell)}_{\mathfrak{x}}) \leq \gamma(\boldsymbol{\eta}^{(\ell)}_{\mathfrak{x}}) \leq 1$. Moreover, from Lemma 7.11, Proposition 2.15, and Claim 7.12, we see that for $d < d_{\mathrm{con}}(k)$ the sequence $\left(\gamma(\boldsymbol{\theta}^{(\ell)}_{\mathfrak{x}})\right)_{\ell}$ converges weakly to $\pi_{d,k}$. Finally, Corollary 7.13 implies that $\left(\gamma(\boldsymbol{\eta}^{(\ell)}_{\mathfrak{x}})\right)_{\ell}$ also converges weakly to $\pi_{d,k}$, and thus,

$$\lim_{\ell\to\infty} \mathbb{E}\left|\gamma\left(\boldsymbol{\theta}^{(\ell)}_{\mathfrak{x}}\right) - \gamma\left(\boldsymbol{\eta}^{(\ell)}_{\mathfrak{x}}\right)\right| = \lim_{\ell\to\infty}\left|\mathbb{E}\left[\gamma\left(\boldsymbol{\theta}^{(\ell)}_{\mathfrak{x}}\right)\right] - \mathbb{E}\left[\gamma\left(\boldsymbol{\eta}^{(\ell)}_{\mathfrak{x}}\right)\right]\right| = 0,$$

implying the assertion. □

Acknowledgements

We would like to thank the anonymous referees for thoroughly reviewing our paper and for suggesting valuable corrections and improvements.

Amin Coja-Oghlan is supported by DFG CO 646/3, DFG CO 646/5 and DFG CO 646/6. Catherine Greenhill is supported by ARC DP250101611. Vincent Pfenninger is supported by the Austrian Science Fund (FWF) [10.55776 / 16502]. Pavel Zakharov is supported by DFG CO 646/6. Kostas Zampetakis is supported by DFG CO 646/5. For open access, the authors have applied a CC BY public copyright licence to any Author Accepted Manuscript version arising from this submission.

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

References

[1] Abbe, E. and Montanari, A. (2014) On the concentration of the number of solutions of random satisfiability formulas. *Random Struct. Alg.* **45** 362–382.
[2] Achlioptas, D. and Coja-Oghlan, A. (2008) Algorithmic barriers from phase transitions Proc. 49th FOCS. 793–802.
[3] Achlioptas, D., Coja-Oghlan, A., Hahn-Klimroth, M., et al. (2021) The number of satisfying assignments of random 2-SAT formulas. *Random Struct. Alg.* **58** 609–647.
[4] Achlioptas, D., Coja-Oghlan, A. and Ricci-Tersenghi, F. (2011) On the solution-space geometry of random constraint satisfaction problems. *Random Struct. Alg.* **38** 251–268.
[5] Achlioptas, D. and Moore, C. (2006) Random k-SAT: two moments suffice to cross a sharp threshold. *Siam J Comput* **36** 740–762.
[6] Achlioptas, D., Naor, A. and Peres, Y. (2005) Rigorous location of phase transitions in hard optimization problems. *Nature* **435** 759–764.
[7] Achlioptas, D. and Peres, Y. (2004) The threshold for random k-SAT is $2^k ln2 - O(k)$. *J. AMS* **17** 947–973.
[8] Aizenman, M., Sims, R. and Starr, S. (2003) An extended variational principle for the SK spin-glass model. *Phys. Rev. B.* **68** 214403.
[9] Aldous, D. and Steele, J. (2004) The objective method: probabilistic combinatorial optimization and local weak convergence. In *Probability on Discrete Structures*, vol. 110 (H. Kesten, ed), Springer, pp. 1–72.
[10] Alon, N. and Spencer, J. (2016) *The Probabilistic Method*. Wiley.
[11] Bapst, V., Coja-Oghlan, A., Hetterich, S., Raßmann, F. and Vilenchik, D. (2016) The condensation phase transition in random graph coloring. *Commun. Math. Phys.* **341** 543–606.
[12] Bayati, M., Gamarnik, D. and Tetali, P. (2013) Combinatorial approach to the interpolation method and scaling limits in sparse random graphs. *Ann. Probab.* **41** 4080–4115.
[13] Biswas, R., Chen, W. and Sen, A. (2024) On the replica symmetric solution in general diluted spin glasses. *Random Struct. Alg.* **66**(4), e70013.
[14] Boucheron, S., Lugosi, G. and Massart, P. (2013) *Concentration Inequalities: A Nonasymptotic Theory of Independence*. OUP Oxford.
[15] Bresler, G. and Huang, B. (2021) The algorithmic phase transition of random k-sat for low degree polynomials, Proc. 62th FOCS. 298–309.
[16] Broder, A., Frieze, A. and Upfal, E. (1993) On the satisfiability and maximum satisfiability of random 3-CNF formulas, Proc. 4th SODA. 322–330.
[17] Chatterjee, A., Coja-Oghlan, A., Müller, N., et al. (2024) The number of random 2-SAT solutions is asymptotically log-normal. Proc. 28th RANDOM.
[18] Cheeseman, P., Kanefsky, B. and Taylor, W. (1991) Where the *really* hard problems are, Proc. IJCAI. 331–337.
[19] Chen, Z., Galanis, A., Goldberg, L. A., et al. (2024) Fast sampling of satisfying assignments from random k-SAT with applications to connectivity. *SIAM J. Disc. Math.* **38** 2750–2811.
[20] Chen, Z., Lonkar, A., Wang, C., Yang, K. and Yin, Y. (2024) Counting random k-SAT near the satisfiability threshold. arXiv 2411.02980v1, 2024.
[21] Chvátal, V. and Reed, B. (1992) Mick gets some (the odds are on his side), Proc. 33th FOCS. 620–627.
[22] Coja-Oghlan, A. (2010) A better algorithm for random k-SAT. *Siam J Comput.* **39** 2823–2864.
[23] Coja-Oghlan, A. (2017) Belief propagation fails on random formulas. *J Acm.* **63** 49.
[24] Coja-Oghlan, A., Kapetanopoulos, T. and Müller, N. (2020) The replica symmetric phase of random constraint satisfaction problems. *Comb. Probab. Comput.* **29** 346–422.
[25] Coja-Oghlan, A., Krzakala, F., Perkins, W. and Zdeborová, L. (2018) Information-theoretic thresholds from the cavity method. *Adv Math.* **333** 694–795.
[26] Coja-Oghlan, A., Müller, N. and Ravelomanana, J. (2022) Belief propagation on the random k-SAT model. *Ann Appl Probab.* **32** 3718–3796.
[27] Coja-Oghlan, A. and Panagiotou, K. (2016) The asymptotic k-SAT threshold. *Adv Math.* **288** 985–1068.
[28] Coja-Oghlan, A. and Perkins, W. (2018) Belief propagation on replica symmetric random factor graph models. *Ann. Inst. Henri Poincaré D.* **5** 211–249.
[29] Coja-Oghlan, A. and Perkins, W. (2019) Bethe states of random factor graphs. *Commun Math Phys.* **366** 173–201.
[30] Coja-Oghlan, A. and Wormald, N (2018) The number of satisfying assignments of random regular k-SAT formulas. *Comb. Probab. Comput.* **27** 496–530.
[31] Dembo, A. and Montanari, A. (2010) Gibbs measures and phase transitions on sparse random graphs. *Braz. J. Probab. Stat.* **24** 137–211.
[32] Dembo, A. and Montanari, A. (2010) Ising models on locally tree-like graphs. *Ann. Appl. Probab.* **20** 565–592.
[33] Dembo, A., Montanari, A. and Sun, N. (2013) Factor models on locally tree-like graphs. *Ann. Probab.* **41** 4162–4213.
[34] Ding, J., Sly, A. and Sun, N. (2022) Proof of the satisfiability conjecture for large. *Ann. Math.* **196** 1–388.
[35] Dubois, O. and Mandler, J. (2002) The 3-XORSAT threshold, Proc. 43rd FOCS. 769–778.
[36] Franz, S. and Leone, M. (2003) Replica bounds for optimization problems and diluted spin systems. *J. Stat. Phys.* **111** 535–564.
[37] Friedgut, E. (1999) Sharp thresholds of graph properties, and the k-SAT problem. *J. AMS.* **12** 1017–1054.
[38] Frieze, A. and Suen, S. (1996) Analysis of two simple heuristics on a random instance of k-SAT. *J Algorithm.* **20** 312–355.

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press

[39] Galanis, A., Goldberg, L. A., Guo, H. and Yang, K. (2021) Counting solutions to random SAT formulas. *SIAM J. Comput.* **50** 1701–1738.
[40] Georgii, H. O. (1988) *Gibbs Measures and Phase Transitions*. De Gruyter.
[41] Goerdt, A. (1996) A threshold for unsatisfiability. *J. Comput. Syst. Sci.* **53** 469–486.
[42] Guerra, F. (2003) Broken replica symmetry bounds in the mean field spin glass model. *Comm. Math. Phys.* **233** 1–12.
[43] Hajiaghayi, M. and Sorkin, G. (2003) The satisfiability threshold of random 3-SAT is at least 3.52. IBM Research Report RC22942.
[44] He, K., Wu, K. and Yang, K. (2023) Improved bounds for sampling solutions of random SAT formulas, Proc. 34th SODA. 3330–3361.
[45] Hetterich, S. (2016) Analysing Survey Propagation Guided Decimationon Random Formulas. Proc. 43rd ICALP.
[46] Janson, S., Luczak, T. and Ruciński, A. (2000) *Random Graphs*. Wiley.
[47] Kaporis, A. C., Kirousis, L. M. and Lalas, E. G. (2006) The probabilistic analysis of a greedy satisfiability algorithm. *Random Struct. Alg.* **28** 444–480.
[48] 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. National Academy of Sciences 104. 10318–10323.
[49] Lovász, L. (2012) *Large Networks and Graph Limits*. AMS.
[50] Mertens, S., Mézard, M. and Zecchina, R. (2006) Threshold values of random K-SAT from the cavity method. *Random Struct. Alg.* **28** 340–373.
[51] Mézard, M. and Montanari, A. (2009) *Information, Physics and Computation*. Oxford University Press.
[52] Mézard, M., Parisi, G. and Zecchina, R. (2002) Analytic and algorithmic solution of random satisfiability problems. *Science*. **297** 812–815.
[53] Moitra, A. (2019) Approximate counting, the Lovasz local lemma, and inference in graphical models. *J. ACM*. **66** 10.
[54] Molloy, M. (2005) Cores in random hypergraphs and Boolean formulas. *Random Struct. Alg.* **27** 124–135.
[55] Monasson, R. and Zecchina, R. (1996) The entropy of the *k*-satisfiability problem. *Phys. Rev. Lett*. **76** 3881.
[56] Monasson, R. and Zecchina, R. (1997) Statistical mechanics of the random *K*-SAT model. *Phys. Rev. E*. **56** 1357–1370.
[57] Montanari, A. and Shah, D. (2007) Counting good truth assignments of random *k*-SAT formulae, Proc. 18th SODA. 1255–1264.
[58] Panchenko, D. (2013) *The Sherrington-Kirkpatrick Model*. Springer.
[59] Panchenko, D. (2013) Spin glass models from the point of view of spin distributions. *Ann Probab*. **41** 1315–1361.
[60] Panchenko, D. (2014) On the replica symmetric solution of the *K*-sat model. *Electron. J. Probab*. **19** 67.
[61] Panchenko, D. and Talagrand, M. (2004) Bounds for diluted mean-fields spin glass models. *Probab. Theory Relat. Fields*. **130** 319–336.
[62] Sly, A. (2010) Computational transition at the uniqueness threshold, Proc. 51st FOCS. 287–296.
[63] Talagrand, M. (2001) The high temperature case for the random *K*-sat problem. *Probab. Theory Related Fields*. **119** 187–212.
[64] Valiant, L. (1979) The complexity of enumeration and reliability problems. *Siam J Comput*. **8** 410–421.
[65] Wang, C. and Yin, Y. (2024) A sampling Lovasz local lemma for large domain sizes, Proc. 65th FOCS. 129–150.

Cite this article: Chatterjee A, Coja-Oghlan A, Greenhill C, Pfenninger V, Rolvien M, Zakharov P, and Zampetakis K (2026). The random *k*-SAT Gibbs uniqueness threshold revisited. *Combinatorics, Probability and Computing* **35**, 433–485. https://doi.org/10.1017/S0963548326100406

https://doi.org/10.1017/S0963548326100406 Published online by Cambridge University Press