Decoding from Pooled Data: Sharp Information-Theoretic Bounds

Ahmed El Alaoui, Aaditya Ramdas, Florent Krzakala, Lenka Zdeborova, Michael I. Jordan

Decoding from Pooled Data: Sharp Information-Theoretic Bounds

Ahmed El Alaoui* Aaditya Ramdas*† Florent Krzakala‡ Lenka Zdeborová§ Michael I. Jordan*†

Abstract

Consider a population consisting of n individuals, each of whom has one of d types (e.g. their blood type, in which case $d = 4$). We are allowed to query this database by specifying a subset of the population, and in response we observe a noiseless histogram (a d-dimensional vector of counts) of types of the pooled individuals. This measurement model arises in practical situations such as pooling of genetic data and may also be motivated by privacy considerations. We are interested in the number of queries one needs to unambiguously determine the type of each individual. In this paper, we study this information-theoretic question under the random, dense setting where in each query, a random subset of individuals of size proportional to n is chosen. This makes the problem a particular example of a random constraint satisfaction problem (CSP) with a "planted" solution. We establish almost matching upper and lower bounds on the minimum number of queries m such that there is no solution other than the planted one with probability tending to 1 as $n \to \infty$. Our proof relies on the computation of the exact "annealed free energy" of this model in the thermodynamic limit, which corresponds to the exponential rate of decay of the expected number of solution to this planted CSP. As a by-product of the analysis, we show an identity of independent interest relating the Gaussian integral over the space of Eulerian flows of a graph to its spanning tree polynomial.

1 Introduction

Constraint satisfaction problems (CSPs) have been the object of intense study in recent years in probability theory, computer science, information theory and statistical physics. For certain families of CSPs, a deep understanding has begun to emerge regarding the number of solutions as a function of problem size, as well as the algorithmic feasibility of finding solutions when they exist (see e.g. [COMV09, COF14, COHH16, COP16, DSS15, DSS16, SSZ16]...) Consider in particular a *planted* random constraint satisfaction problem with n variables that take their values in the discrete set $\{1, \cdots, d\}$, with $d \geq 2$, and with m clauses drawn uniformly at random under the constraint that they are all satisfied by a pre-specified assignment, which is referred to as *the planted solution*. It is of interest to determine the properties of the set of all solutions of CSP as n and m grow to infinity at a some relative rate. Two questions are of particular importance: *(1) how large should m be so that the planted solution is the unique solution to CSP?* and *(2) given that it is unique, how large should m be so that it is recoverable by a "tractable" algorithm?* Significant progress has been made on these questions, often initiated by insights from statistical physics and followed by a growing body of rigorous mathematical investigation. The emerging

*Department of Electrical Engineering and Computer Sciences, UC Berkeley, CA.

†Department of Statistics, UC Berkeley, CA.

‡Laboratoire de Physique Statistique, CNRS, PSL Universités & Ecole Normale Supérieure, Sorbonne Universités et Université Pierre & Marie Curie, Paris, France.

§Institut de Physique Théorique, CNRS, CEA, Université Paris-Saclay, Gif-sur-Yvette, France.

picture is that in many planted CSPs, when n is sufficiently large, all solutions become highly correlated with the planted one when $m > \kappa_{\mathsf{IT}} \cdot n$, for some "Information-Theoretic" (IT) constant $\kappa_{\mathsf{IT}} > 0$. Furthermore, one of these highly correlated solutions becomes typically recoverable by a random walk or a Belief Propagation (BP)-inspired algorithm when $m > \kappa_{\mathsf{BP}} \cdot n$ for some $\kappa_{\mathsf{BP}} > \kappa_{\mathsf{IT}}$ [COMV09, KZ09, KMZ12, COF14]. Interestingly, it is known in many problems, at least heuristically, that these algorithms fail when $\kappa_{\mathsf{IT}} < m/n < \kappa_{\mathsf{BP}}$, and a tractable algorithm that succeeds in this regime is still lacking [ACO08, CO09, ZK15, COHH16]. In other words, there is a non-trivial regime $m/n \in (\kappa_{\mathsf{IT}}, \kappa_{\mathsf{BP}})$ where an essentially unique solution exists, but is hard to recover. In the picture we described, the uniqueness and recoverability thresholds differ by a constant factor; that is, they both happen at the same scale where m is proportional to n. In other examples, notably a class of planted CSPs related to XORSAT, the gap between the information theoretic and the putative algorithmic threshold is much larger [FPV15].

The systematic study of this phenomenon has mostly dealt with planted CSPs with Boolean variables (the n variables are $\{0,1\}$-valued and $d = 2$). Among CSPs with larger variable domain, the random Graph Coloring problem with d colors is a prominent prototype. Here the size of the gap is constant.

In this paper we consider a naturally motivated random CSP with arbitrary but fixed domain size d, which we call the Histogram Query Problem (HQP), where we present evidence that the gap between existence of a unique solution and its tractable recovery could be as large as $\log n$. More precisely, in this paper, we undertake a detailed information-theoretic analysis which shows that in HQP, the planted solution becomes unique at as soon as $m > \gamma^* n/\log n$ with high probability as $n \to \infty$ for a constant $\gamma^* > 0$. In a sequel paper, we consider the algorithmic aspect of the problem and provide a BP-based algorithm that recovers the planted assignment if $m \geq \kappa^* \cdot n$ for a specific threshold κ^* and fails otherwise. This indicates the existence of a logarithmic gap between the information-theoretic and algorithmic thresholds.

1.1 Problem and motivation

The setting Let $\{\boldsymbol{h}_a\}_{1 \leq a \leq m}$ be a collection of d-dimensional arrays with non-negative integer entries. For an assignment $\tau : \{1, \cdots, n\} \mapsto \{1, \cdots, d\}$ of the n variables, and given a realization of m random subsets $S_a \subset \{1, \cdots, n\}$, the constraints of the HQP are given by $\boldsymbol{h}_a = \boldsymbol{h}_a(\tau)$ for all $1 \leq a \leq m$ with

$$\boldsymbol{h}_a(\tau) := \left(\left|\tau^{-1}(1) \cap S_a\right|, \cdots, \left|\tau^{-1}(d) \cap S_a\right|\right) \in \mathbb{Z}_+^d.$$

We let $\tau^* : \{1, \cdots, n\} \mapsto \{1, \cdots, d\}$ be a planted assignment; i.e., we set $\boldsymbol{h}_a := \boldsymbol{h}_a(\tau^*)$ for all a for some realization of the sets $\{S_a\}$, and consider the problem of recovering the map τ^* given the observation of the arrays $\{\boldsymbol{h}_a\}_{1 \leq a \leq m}$.

This problem can be viewed informally as that of decoding a discrete high-dimensional signal consisting of categorical variables from a set of measurements formed by pooling together the variables belonging to a subset of the signal. It is useful to think of the n variables as each describing the type or category of an individual in a population of size n, where each individual has exactly one type among d. For instance the categories may represent blood types or some other discrete feature such as ethnicity or age group. Then, the observation $\boldsymbol{h}_a$ is the histogram of types of a subpopulation S_a. We let $\boldsymbol{\pi} = \frac{1}{n}\left(\left|\tau^{*-1}(1)\right|, \cdots, \left|\tau^{*-1}(d)\right|\right)$ denote the vector of proportions of assigned values; i.e., the empirical distribution of categories.

We consider here a model in which each variable participates in a given constraint independently and with probability $\alpha \in (0,1)$. Thus, the sets $\{S_a\}_{1 \leq a \leq m}$ are independent draws of a random set S where $\Pr(i \in S) = \alpha$ independently for each $i \in \{1, \ldots, n\}$. We are thus in the "dense regime"

where $\mathbb{E}[|S|] = \alpha n$; i.e., the number of variables participating in each constraint (the degree of each factor in the CSP) is linear in n.

Motivation This model is inspired by practical problems in which a data analyst can only assay certain summary statistics involving a moderate or large number of participants. This may be done for privacy reasons, or it may be inherent in the data-collection process (see e.g. [SBC+02, HLB+01]). For example, in DNA assays, the pooling of allele measurements across multiple strands of DNA is necessary given the impracticality of separately analyzing individual strands. Thus the data consists of a frequency spectrum of alleles; a "histogram" in our language. In the privacy-related situation, one may take the viewpoint of an attacker whose goal is to gain a granular knowledge of the database from coarse measurements, or that of a guard who wishes to prevent this scenario from happening. It is then natural to ask how many histogram queries it takes to exactly determine the category of each individual.

Related problems Note that the case $d = 2$ of HQP can be seen as a compressed sensing problem with a binary sensing matrix and binary signal. While the bulk of the literature in the field of compressed sensing is devoted to the case in which both the signal of interest and the sensing matrix are real-valued, the binary case has also been considered, notably in relation to Code Division Multiple Access (CDMA) [Zig04, Tan02], and Group Testing [DH06, MT11]: in the latter, one observes the logical "OR" of subsets of the entries of the signal. In the case of categorical variables with $d \geq 3$ categories, it is natural to consider measurements consisting of histograms of the categories in the pooled sub-population. In the literature on compressed sensing one commonly considers the setting where the sensing matrices have i.i.d. entries with finite second moment, and the signal has an arbitrary empirical distribution of its entries. It has been established that, under the scaling $m = \kappa n$, whereas the success of message-passing algorithms requires $\kappa > \kappa_{\mathsf{BP}}$ [BLM15], the information-theoretic threshold is $\kappa_{\mathsf{IT}} = 0$ in the discrete signal case [WV09, DJM13], indicating that uniqueness of the solution happens at a finer scale $m = o(n)$. Here we consider the HQP with arbitrary d, for which the exact scaling for investigating uniqueness is $m = \gamma \frac{n}{\log n}$ with finite $\gamma > 0$, and provide tight bounds on the information-theoretic threshold.

Prior work on HQP The study of this problem for generic values of d was initiated in [WHLC16] in the two settings where the sets $\{S_a\}$ are deterministic and random. They showed in both these cases with a simple counting argument that under the condition that $\boldsymbol{\pi}$ is the uniform distribution, if $m < \frac{\log d}{d-1} \frac{n}{\log n}$ then the set of collected histograms does not uniquely determine the planted assignment τ^* (with high probability in the random case). On the other hand, for the deterministic setting, they provided a querying strategy that recovers τ^* provided that $m > c_0 \frac{n}{\log n}$, where c_0 is an absolute constant independent of d. For the random setting and under the condition that the sets S_a are of average size $n/2$, they proved via a first moment bound that $m > c_1 \frac{n}{\log n}$ with c_1 also constant and independent of d, suffices to uniquely determine τ^*, although no algorithm was proposed in this setting.

In the above results, there is a gap that is both information-theoretic and algorithmic depending on the dimension d between the upper and lower bounds. Intuitively, the upper bounds should also depend on d since the decoding problem becomes easier (or at least, it is no harder) for large d, for the simple reason that if it is possible to determine the categories of the population for $d = 2$, then one can proceed by dichotomy for larger d by merging the d groups into two super-groups, identifying which individuals belong to each of the two super-groups, and then recurse. We attempt to fill the information-theoretic gap in the random setting by providing tighter upper and lower bounds on

the number of queries m necessary and sufficient to uniquely determine the planted assignment τ^* with high probability, which depend on the dimension d and $\boldsymbol{\pi}$ along with explicit constants. In a sequel paper, we consider the algorithmic aspect of the problem and provide a Belief Propagation-based algorithm that recovers the planted assignment if $m \geq \kappa^*(\boldsymbol{\pi}, d) \cdot n$ for a specific threshold $\kappa^*(\boldsymbol{\pi}, d)$ and fails otherwise, indicating the putative existence of a statistical-computational gap in the random setting.

1.2 Main result

Let Δ^{d-1} be the $d-1$-dimensional simplex and $H(\boldsymbol{x}) = -\sum_{r=1}^{d} x_r \log x_r$ for $\boldsymbol{x} \in \Delta^{d-1}$ be the Shannon entropy function. We write $\tau \sim \boldsymbol{\pi}$ to indicate that τ is a random assignment drawn from the uniform distribution over maps $\tau : \{1, \cdots, n\} \mapsto \{1, \cdots, d\}$ such that $\frac{1}{n}\left(\left|\tau^{-1}(1)\right|, \cdots, \left|\tau^{-1}(d)\right|\right) = \boldsymbol{\pi}$.

Theorem 1. *For $n \geq 2$ integer, $m = \gamma \frac{n}{\log n}$, $\gamma > 0$, $\alpha \in (0,1)$, and $\boldsymbol{\pi} \in \Delta^{d-1}$ with entries bounded away from 0 and 1. Let $\mathcal{E}$ be the event that τ^* is* not *the unique satisfying assignment to* HQP*:*

$$\mathcal{E} = \left\{\exists \tau \in \{1, \cdots, d\}^n \ : \ \tau \neq \tau^*, \ \boldsymbol{h}_a(\tau) = \boldsymbol{h}_a(\tau^*) \ \forall a \in \{1, \cdots, m\}\right\}.$$

(i) *If*

$$\gamma < \gamma_{low} := \frac{H(\boldsymbol{\pi})}{d-1},$$

then

$$\lim_{n \to \infty} \mathbb{E}_{\tau^* \sim \boldsymbol{\pi}} \Pr\left(\mathcal{E}\right) = 1.$$

(ii) *On the other hand, let $\boldsymbol{\pi}_{[\cdot]}$ be the vector of order statistics of $\boldsymbol{\pi}$: $\pi_{[1]} \geq \pi_{[2]} \geq \cdots \geq \pi_{[d]}$. For $1 \leq k \leq d-1$, let $\boldsymbol{\pi}^{(k)} \in \Delta^{k-1}$ be defined as $\pi_1^{(k)} = \sum_{r=1}^{d-k+1} \pi_{[r]}$ and $\pi_l^{(k)} = \pi_{[d-k+l]}$ for all $2 \leq l \leq k$ (if $k \geq 2$). If*

$$\gamma > \gamma_{up} := 2 \max_{1 \leq k \leq d-1} \frac{H(\boldsymbol{\pi}) - H(\boldsymbol{\pi}^{(k)})}{d-k},$$

then

$$\lim_{n \to \infty} \mathbb{E}_{\tau^* \sim \boldsymbol{\pi}} \Pr\left(\mathcal{E}\right) = 0.$$

Remarks and special cases:

- For $d = 2$, $\gamma_{\text{up}} = 2H(\boldsymbol{\pi}) = 2\gamma_{\text{low}}$.
- If $\boldsymbol{\pi} = (\frac{1}{d}, \cdots, \frac{1}{d})$, or more generally, if $\boldsymbol{\pi}$ is such that $k = 1$ maximizes the expression defining γ_{up} then $\gamma_{\text{up}} = 2\frac{H(\boldsymbol{\pi})}{d-1} = 2\gamma_{\text{low}}$.
- The resulting bounds do not depend on α as long as it is fixed and bounded away from 0 and 1. Its contribution in the problem is sub-dominant and vanishes as $n \to \infty$ under the scaling considered here.
- The number k in the expression of γ_{up} can be interpreted as the number of connected components of a graph on d vertices that depends on the overlap structure of the two assignments τ and τ^*, and induces "maximum confusion" between them. This will become clear in latter sections.

The proof of the above Theorem occupies the rest of the manuscript.

1.3 Main ideas of the proof

Our main contribution is the second part of Theorem 1, which establishes an upper bound on the uniqueness threshold of the random CSP with histogram constraints HQP. The proof uses the first moment method to upper bound the probability of existence of a non-planted solution. Since we are in a planted model, the analysis of the first moment ends up bearing many similarities with a second moment computation in a purely random (non-planted) model. Although second moment computations often require approximations, for the HQP it turns out that we are able to compute the exact annealed free energy of the model in the thermodynamic limit. That is, letting $\mathcal{Z}$ be the number of solutions of the CSP, we show that the limit

$$\mathfrak{F}(\gamma) := \lim_{n\to\infty} \frac{1}{n} \log \mathbb{E}\left[\mathcal{Z} - 1\right]$$

exists and we compute its value exactly. Then the value of the threshold γ_{up} is obtained by locating the first point at which $\mathfrak{F}$ becomes negative:

$$\gamma_{\text{up}} = \inf\left\{\gamma > 0 \ : \ \mathfrak{F}(\gamma) < 0\right\}.$$

Together with the fact that $\mathfrak{F}$ is a monotone function, which will become clear once $\mathfrak{F}$ is computed, it is clear that for any $\gamma > \gamma_{\text{up}}$, $\mathbb{E}[\mathcal{Z} - 1]$ decays exponentially with n when the latter is sufficiently large.

This general strategy has been successfully pursued for a range of CSPs, such as K-SAT, NAE-SAT, and Independent Set, most of which are Boolean. For larger domain sizes, in order to carry out the second moment method one needs fine control of the overlap structure between the planted and a candidate solution. This control is at the core of the difficulty that arises in any second moment computation. To obtain such control, researchers have often imposed additional assumptions, at a cost of a weakening of the resulting bounds. For example, existing proofs for Graph Coloring and similar problems assume certain balancedness conditions (the overlap matrix needs to be close to doubly stochastic.) without which the annealed free energy cannot be computed [AN05, AM04, COEH16, BCOH$^+$16, BMNN16]; this yields results that fall somewhat short of the bounds that the second moment method could achieve in principle [DMO12]. In the present problem, due its rich combinatorial structure, we are able to obtain unconditional control of the overlap structure, for any domain size d, and compute the exact annealed free energy.

Concretely, computing the function $\mathfrak{F}$ requires tight control of the "collision probability" of two non-equal assignments τ_1 and τ_2. This is the probability that the random histograms $\boldsymbol{h}(\tau_1) = \left(\left|{\tau_1}^{-1}(1) \cap S\right|, \cdots, \left|{\tau_1}^{-1}(d) \cap S\right|\right)$ and $\boldsymbol{h}(\tau_2) = \left(\left|{\tau_2}^{-1}(1) \cap S\right|, \cdots, \left|{\tau_2}^{-1}(d) \cap S\right|\right)$ generated from a random draw of a pool S coincide. The collision probability roughly measures the correlation strength between the two assignments. Specifically, we will be interested in the collision probabilities of the pairs (τ^*, τ) where τ^* is the planted assignment and τ is any candidate assignment. Its decay reveals how long an assignment τ "survives" as a satisfying assignment to HQP as $n \to \infty$. The study of these collision probabilities requires the evaluation of certain Gaussian integrals over the space of *Eulerian flows* of a weighted graph on d vertices that is defined based on the overlap structure of τ and τ^*. We prove a family of identities that relate these integrals to some combinatorial polynomials in the weights of the graph: the spanning tree and spanning forest polynomials. We believe that these identities are of independent interest beyond the problem studied in this paper. Once these collision probabilities are controlled, the computation of $\mathfrak{F}(\gamma)$ per se requires the analysis of a certain sequence of optimization problems. We show that the sequence of maximum values converges to a finite limit that yields the value of the annealed free energy.

On the other hand, the proof of the first part of Theorem 1 is straightforward—it is an extension of a standard counting argument used in [ZKMZ13] and [WHLC16]. The argument goes as follows: if

m is too small then the number of possible histograms one could potentially observe is exponentially smaller than the number of assignments of n variables that agree with $\boldsymbol{\pi}$. Therefore when the planted assignment τ^* is drawn at random, there will exist at least one $\tau \neq \tau^*$ that satisfies the constraints of the CSP with overwhelming probability. We begin with this argument in the next section and then turn to the more challenging computation of the upper bound.

2 Proof of Theorem 1

Notation We denote vectors in $\mathbb{R}^d$ in bold lower case letters, e.g., $\boldsymbol{x}$, and matrices in $\mathbb{R}^{d\times d}$ will be written in bold lower case underlined letters, e.g., $\underline{\boldsymbol{x}}$. We denote the coordinates of such vectors and matrices as x_r and x_{rs} respectively. Matrices that act either as linear operators on the space $\mathbb{R}^{d\times d}$ or that are functions of elements in this space are written in bold upper case letters, e.g., $\boldsymbol{M}\underline{\boldsymbol{x}}$ and $\boldsymbol{L}(\underline{\boldsymbol{x}})$, for $\underline{\boldsymbol{x}} \in \mathbb{R}^{d\times d}$. These choices will be clear from the context. We may write $\underline{\boldsymbol{x}}/\underline{\boldsymbol{y}}$ to indicate coordinate-wise division. Additionally, for two $d \times d$ matrices $\underline{\boldsymbol{a}}, \underline{\boldsymbol{b}} \in \mathbb{R}^{d\times d}$, $\underline{\boldsymbol{a}} \odot \underline{\boldsymbol{b}} \in \mathbb{R}^{d\times d}$ is their Hadamard product. We let $\mathbf{1} \in \mathbb{R}^d$ be the all-ones vector.

2.1 The first part of Theorem 1: the lower bound

Let $m = \gamma \frac{n}{\log n}$ with $\gamma > 0$. The number of potential histograms one could possibly observe in a single query with pool size $|S| = k$ is $f(k,d) := \binom{d+k-1}{d-1} \leq (k+1)^{d-1}$. Since the queries are independent, the number of collections of histograms $\{\boldsymbol{h}_a\}_{1\leq a\leq m}$ one could potentially observe in m queries is $\prod_{a=1}^m f(|S_a|, d)$. On the other hand, the number of possible assignments $\tau : \{1, \cdots, n\} \mapsto \{1, \cdots, d\}$ satisfying the constraint $\boldsymbol{\pi} = \frac{1}{n}\left(|\tau^{*-1}(1)|, \cdots, |\tau^{*-1}(d)|\right)$ is $\binom{n}{n\boldsymbol{\pi}} = \binom{n}{n\pi_1, \cdots, n\pi_d} \geq C(\boldsymbol{\pi}) n^{-(d-1)/2} \exp(H(\boldsymbol{\pi})n)$, for some constant $C(\boldsymbol{\pi}) > 0$ depending on $\boldsymbol{\pi}$.

Now, the probability that τ^* is the unique satisfying assignment of the CSP with constraints given by the random histograms $\{\boldsymbol{h}_a(\tau^*)\}_{1\leq a\leq m}$, averaged over the random choice of $\tau^* \sim \boldsymbol{\pi}$, is

$$\begin{aligned}
&\mathbb{E}_{\tau^*\sim\boldsymbol{\pi}} \mathbb{E}_{\{S_a\}}\Big[\mathbb{1}\{\forall \tau \in \{1,\cdots,d\}^n \ : \ \boldsymbol{h}_a(\tau) = \boldsymbol{h}_a(\tau^*)\ \forall a \in \{1,\cdots,m\} \implies \tau = \tau^*\}\Big] \\
&\qquad \leq \binom{n}{n\boldsymbol{\pi}}^{-1} \cdot \mathbb{E}_S\left[f(|S|,d)\right]^m \\
&\qquad \leq \binom{n}{n\boldsymbol{\pi}}^{-1} \cdot \mathbb{E}_S\left[(|S|+1)^{d-1}\right]^m \\
&\qquad \leq C(\boldsymbol{\pi})\ n^{(d-1)/2} \cdot \exp\Big(-H(\boldsymbol{\pi})n\Big) \cdot (n+1)^{m(d-1)} \\
&\qquad \leq C(\boldsymbol{\pi})\ n^{(d-1)/2} \cdot \exp\Big((\gamma(d-1) - H(\boldsymbol{\pi}))n\Big).
\end{aligned}$$

If $\gamma < \gamma_{\text{low}}$ the last quantity tends to 0 as $n \to \infty$. This concludes the proof of the first assertion of the theorem.

2.2 The second part of Theorem 1 : the upper bound

We use a first moment method to show that when γ is greater than γ_{up}, the only assignment satisfying HQP is τ^* with high probability. Let $\mathcal{Z}$ be the number of satisfying assignments to HQP:

$$\mathcal{Z} := \Big|\{\tau \in \{1,\cdots,d\}^n \ : \ \boldsymbol{h}_a(\tau) = \boldsymbol{h}_a(\tau^*)\ \forall a \in \{1,\cdots,m\}\}\Big|. \tag{1}$$

The planted assignment τ^* is obviously a solution, so we always have $\mathcal{Z} \geq 1$. Recall the definition of the annealed free energy

$$\mathfrak{F}(\gamma) := \lim_{n\to\infty} \frac{1}{n} \log \mathbb{E}\left[\mathcal{Z} - 1\right]. \tag{2}$$

Also, recall that for $1 \leq k \leq d-1$, $\boldsymbol{\pi}^{(k)} \in \Delta^{k-1}$ be defined as $\pi_1^{(k)} = \sum_{r=1}^{d-k+1} \pi_{[r]}$ and $\pi_l^{(k)} = \pi_{[d-k+l]}$ for all $2 \leq l \leq k$ (if $k \geq 2$).

Theorem 2. *Let $m = \gamma \frac{n}{\log n}$ with $\gamma > 0$. The limit (2) exists for all $\gamma > 0$ and its value is*

$$\mathfrak{F}(\gamma) = \max_{1\leq k\leq d-1} \left\{ H(\boldsymbol{\pi}) - H(\boldsymbol{\pi}^{(k)}) - \frac{\gamma}{2}(d-k) \right\}. \tag{3}$$

We can deduce from Theorem 2 the smallest value of γ past which $\mathfrak{F}(\gamma)$ becomes negative. In particular, we see that $\mathfrak{F}$ is a decreasing function of γ that crosses the horizontal axis at

$$\gamma_{\text{up}} = 2 \max_{1\leq k\leq d-1} \frac{H(\boldsymbol{\pi}) - H(\boldsymbol{\pi}^{(k)})}{d-k}.$$

From this result it is easy to prove the second assertion of Theorem 1. By averaging over τ^* and applying Markov's inequality, we have:

$$\mathbb{E}_{\tau^*\sim\boldsymbol{\pi}} \Pr\left(\exists \tau \in \{1,\cdots,d\}^n \ : \ \tau \neq \tau^*, \boldsymbol{h}_a(\tau) = \boldsymbol{h}_a(\tau^*)\ \forall a \in \{1,\cdots,m\}\right) = \mathbb{E}_{\tau^*\sim\boldsymbol{\pi}} \Pr\left(\mathcal{Z} \geq 2\right) \leq \mathbb{E}[\mathcal{Z}-1].$$

For $\gamma > \gamma_{\text{up}}$, it is clear that $\mathfrak{F}(\gamma) < 0$. Let $0 < \epsilon < |\mathfrak{F}(\gamma)|/2$; then there is an integer $n_0(\epsilon) \geq 0$ such that for all $n \geq n_0(\epsilon)$,

$$\begin{aligned}
\mathbb{E}_{\tau^*\sim\boldsymbol{\pi}} \Pr\left(\exists \tau \in \{1,\cdots,d\}^n \ : \ \tau \neq \tau^*, \boldsymbol{h}_a(\tau) = \boldsymbol{h}_a(\tau^*)\ \forall a \in \{1,\cdots,m\}\right) &\leq \exp n\ (\mathfrak{F}(\gamma) + \epsilon), \\
&\leq \exp n\ \mathfrak{F}(\gamma)/2, \\
&\xrightarrow[n\to\infty]{} 0.
\end{aligned}$$

Now it remains to prove Theorem 2, and this represents the main technical thrust of our paper.

2.3 Collisions, overlaps, and the first moment

Preliminaries We begin by presenting the main quantities to be analyzed in our application of the first moment method. We have

$$\begin{aligned}
\mathbb{E}_{\tau^*\sim\boldsymbol{\pi}} \mathbb{E}_{\{S_a\}}[\mathcal{Z}-1] &= \mathbb{E}_{\tau^*\sim\boldsymbol{\pi}} \left[\sum_{\substack{\tau\in\{1,\cdots,d\}^n \\ \tau\neq\tau^*}} \Pr\left(\boldsymbol{h}_a(\tau) = \boldsymbol{h}_a(\tau^*)\ \forall a \in \{1,\cdots,m\}\right) \right] \\
&= (d^n - 1) \Pr_{\tau,\tau^*,\{S_a\}} \left(\boldsymbol{h}_a(\tau) = \boldsymbol{h}_a(\tau^*)\ \forall a \in \{1,\cdots,m\}\right),
\end{aligned}$$

where $\tau^* \sim \boldsymbol{\pi}$, $\tau \sim \text{Unif}(\{1,\cdots,d\}^n \backslash \{\tau^*\})$. By conditional independence,

$$\begin{aligned}
\Pr_{\tau,\tau^*,\{S_a\}} \left(\boldsymbol{h}_a(\tau) = \boldsymbol{h}_a(\tau^*)\ \forall a \in \{1,\cdots,m\}\right) &= \mathbb{E}_{\tau,\tau^*} \left[\Pr_{\{S_a\}} \left(\boldsymbol{h}_a(\tau) = \boldsymbol{h}_a(\tau^*)\ \forall a \in \{1,\cdots,m\}\right) \right] \\
&= \mathbb{E}_{\tau,\tau^*} \left[\Pr_{S} \left(\boldsymbol{h}(\tau) = \boldsymbol{h}(\tau^*)\right)^m \right].
\end{aligned}$$

Next, we write the *collision probability*, $\Pr_S(\boldsymbol{h}(\tau) = \boldsymbol{h}(\tau^*))$, for fixed τ and τ^* in a convenient form. Let us first define the *overlap matrix*, $\underline{\boldsymbol{\mu}}(\tau, \tau^*) = (\mu_{rs})_{1 \le r,s \le d} \in \mathbb{Z}_+^{d \times d}$, of τ and τ^*, by

$$\mu_{rs} = \left| \tau^{-1}(r) \cap \tau^{*-1}(s) \right| \quad \text{for all } r, s = 1, \cdots, d. \tag{4}$$

Remark that $\boldsymbol{h}(\tau) = \boldsymbol{h}(\tau^*)$ if and only if $\left|S \cap \tau^{-1}(r)\right| = \left|S \cap \tau^{*-1}(r)\right|$ for all $r \in \{1, \cdots, d\}$. Since the collection of sets $\{\tau^{-1}(r)\}_{1 \le r \le d}$ forms a partition of $\{1, \cdots, n\}$, and similarly with τ^*, we have the following equality of events

$$\{\boldsymbol{h}(\tau) = \boldsymbol{h}(\tau^*)\} = \left\{ \sum_{s=1}^{d} \left| S \cap \tau^{-1}(r) \cap \tau^{*-1}(s) \right| = \sum_{s=1}^{d} \left| S \cap \tau^{-1}(s) \cap \tau^{*-1}(r) \right|, \ \forall r \in \{1, \cdots, d\} \right\}.$$

Therefore, the probability that two assignments τ and τ^* collide on a random pool S—meaning that their histograms formed on the pool S coincide—is

$$\Pr_S(\boldsymbol{h}(\tau) = \boldsymbol{h}(\tau^*)) = \sum_{\underline{\boldsymbol{\nu}}} \left(\prod_{r,s=1}^{d} \binom{\mu_{rs}}{\nu_{rs}} \alpha^{\nu_{rs}} (1-\alpha)^{\mu_{rs} - \nu_{rs}} \right) \mathbb{1}\left\{ \sum_{s=1}^{d} \nu_{rs} = \sum_{s=1}^{d} \nu_{sr} \ , \ \forall r \in [d] \right\}, \tag{5}$$

where the outer sum is over all arrays of integer numbers $\underline{\boldsymbol{\nu}} = (\nu_{rs})_{1 \le r,s \le d}$ such that $0 \le \nu_{rs} \le \mu_{rs}$ for all r, s. We see from the above expression that the collision probability of τ and τ^* only depends on the overlap matrix $\underline{\boldsymbol{\mu}}(\tau, \tau^*)$. We henceforth denote the probability in equation (5) by $q(\underline{\boldsymbol{\mu}})$, where we dropped the dependency on τ and τ^*. Remark that $\tau = \tau^*$ if and only if their overlap matrix $\underline{\boldsymbol{\mu}}$ is diagonal. Thus, we can rewrite the expected number of solutions as

$$\mathbb{E}[\mathcal{Z} - 1] = \binom{n}{n\boldsymbol{\pi}}^{-1} \cdot \sum_{\underline{\boldsymbol{\mu}}} \binom{n}{\underline{\boldsymbol{\mu}}} q(\underline{\boldsymbol{\mu}})^m \ \mathbb{1}\left\{ \sum_{r=1}^{d} \mu_{rs} = n\pi_s, \ s \in \{1, \cdots, d\} \right\}, \tag{6}$$

where the sum is over all non-diagonal arrays $\underline{\boldsymbol{\mu}} = (\mu_{rs})_{1 \le r,s \le d}$ with non-negative integer entries that sum to n, and $\binom{n}{\underline{\boldsymbol{\mu}}} = \frac{n!}{\prod_{r,s} \mu_{rs}!}$.

The rest of the proof From here, the proof of Theorem 2 roughly breaks into three parts:

(*i*) One needs to have tight asymptotic control on the collision probability $q(\underline{\boldsymbol{\mu}})$ when any subset of the entries of $\underline{\boldsymbol{\mu}}$ becomes large. This will be achieved via the Laplace method (see, e.g., [DB70]). The outcome of this analysis is an asymptotic estimate that exhibits two different speeds of decay, polynomial or exponential, depending on the "balancedness" of $\underline{\boldsymbol{\mu}}$ as its entries become large. This notion of balancedness, namely that $\underline{\boldsymbol{\mu}}$ must have equal row- and column-sums[1], is specific to the histogram setting and departs from the usual "double stochasticity" that arises in other more classical problems such as Graph Coloring, and Community Detection under the stochastic block model [AN05, AM04, COEH16, BCOH$^+$16, BMNN16]. As we will explain in the next section, configurations (τ, τ^*) with an unbalanced overlap matrix have an exponentially decaying collision probability, i.e., they exhibit weak correlation, and disappear very early on as $n \to \infty$ under the scaling $m = \gamma \frac{n}{\log n}$. On the other hand, those configurations with balanced overlap exhibit a slow decay of correlation: their collision probability decays only polynomially, and these are the last surviving configurations in expression (6) as $n \to \infty$.

(*ii*) Understanding the above-mentioned polynomial decay of $q(\underline{\boldsymbol{\mu}})$ requires the evaluation of a multivariate Gaussian integral (which is a product of the above analysis) over the space of constraints

[1]These are exactly the constraints on $\underline{\boldsymbol{\nu}}$ showing up in (5).

of the array $\underline{\boldsymbol{\nu}}$ in (5); the latter being the space of *Eulerian flows* on the graph on d vertices whose edges are weighted by the (large) entries of $\underline{\boldsymbol{\mu}}$. We show that this integral, properly normalized, evaluates to *the inverse square root of the spanning tree (or forest) polynomial* of this graph. This identity seems to be new, to the best of our knowledge, and may be of independent interest. We therefore provide two very different proofs of it, each highlighting different combinatorial aspects.

(*iii*) Lastly, armed with these estimates, we show the existence of, and compute the exact value of, the annealed free energy of the model in the thermodynamic limit, thereby completing the proof of Theorem 2. This last part requires the analysis of a certain optimization problem involving an entropy term and an "energy" term accounting for the correlations discussed above. Here we can exactly characterize the maximizing configurations for large n, and this allows the computation of the value of $\mathfrak{F}(\gamma)$. We note once more that this situation contrasts with the more traditional case of Graph Coloring, where we lack a rigorous understanding of the maximizing configurations of the second moment, except when certain additional constraints are imposed on their overlap matrix.

3 Bounding the collision probabilities

Here we provide tight asymptotic bounds on the collision probabilities $q(\underline{\boldsymbol{\mu}})$ defined in (5). Consider the following subspace of $\mathbb{R}^{d\times d}$, which will play a key role in the analysis:

$$\mathcal{F} := \left\{ \underline{\boldsymbol{x}} \in \mathbb{R}^{d\times d} \ : \ \sum_{s=1}^{d} x_{rs} = \sum_{s=1}^{d} x_{sr} \ , \ \forall r \in \{1,\cdots,d\} \right\}. \tag{7}$$

This is a linear subspace of dimension $(d-1)^2 + d$ in $\mathbb{R}^{d\times d}$. For $p, q \in (0,1)$, let $D(p \parallel q) = p\log(p/q) + (1-p)\log((1-p)/(1-q))$ be the Kullback-Leibler divergence. Let $G = (V, E)$ be an undirected graph on d vertices where we allow up to two parallel edges between each pair of vertices, i.e., $V = \{1,\cdots,d\}$, and $E \subseteq \{(r,s) \ : \ r,s \in V,\ r \neq s\}$. For $\underline{\boldsymbol{\nu}}, \underline{\boldsymbol{\mu}} \in \mathbb{R}_+^{d\times d}$, $\underline{\boldsymbol{x}} \in [0,1]^{d\times d}$ let

$$\varphi_{\underline{\boldsymbol{\mu}}}(\underline{\boldsymbol{x}}) := \sum_{(r,s)\in E} \mu_{rs} D(x_{rs} \parallel \alpha). \tag{8}$$

and recalling that $\odot$ represents the Hadamard product, we let

$$\vartheta(\underline{\boldsymbol{\nu}}, \underline{\boldsymbol{\mu}}) := \min_{\substack{\underline{\boldsymbol{x}}\in[0,1]^{d\times d} \\ \boldsymbol{M}_G(\underline{\boldsymbol{x}}\odot\underline{\boldsymbol{\mu}},\underline{\boldsymbol{\nu}})\in\mathcal{F}}} \sum_{(r,s)\in E} \mu_{rs} D(x_{rs} \parallel \alpha), \tag{9}$$

where for two $d\times d$ matrices $\underline{\boldsymbol{a}}, \underline{\boldsymbol{b}}$, $\boldsymbol{M}_G(\underline{\boldsymbol{a}}, \underline{\boldsymbol{b}})$ is the $d\times d$ matrix with entries a_{rs} if $(r,s) \in E$ and b_{rs} otherwise. By strong duality (see, e.g., [BV04, Roc70]), the function (9) can be written in the more transparent form

$$\begin{aligned}\vartheta(\underline{\boldsymbol{\nu}}, \underline{\boldsymbol{\mu}}) &= \sup_{\boldsymbol{\lambda}\in\mathbb{R}^d} \left\{ \sum_{(r,s)\notin E} \nu_{rs}(\lambda_r - \lambda_s) + \sum_{(r,s)\in E} \mu_{rs} \log\left(\frac{e^{\lambda_r-\lambda_s}}{\alpha + (1-\alpha)e^{\lambda_r-\lambda_s}} \right) \right\}, \\ &= \phi^*_{\underline{\boldsymbol{\mu}}}(\underline{\boldsymbol{\nu}}\mathbf{1} - \underline{\boldsymbol{\nu}}^\intercal\mathbf{1}),\end{aligned}$$

where $\phi^*_{\underline{\boldsymbol{\mu}}}$ is the Legendre-Fenchel transform of the (convex) function

$$\phi_{\underline{\boldsymbol{\mu}}}(\boldsymbol{\lambda}) := -\sum_{(r,s)\in E} \mu_{rs} \log\left(\frac{e^{\lambda_r-\lambda_s}}{\alpha + (1-\alpha)e^{\lambda_r-\lambda_s}} \right).$$

We may note that since $\phi^*_{\underline{\boldsymbol{\mu}}}$ is convex on $\mathbb{R}^d$, ϑ is a continuous function of its first argument. Before we state our bounds on the collision probability, we recall the following concept from algebraic graph theory. Define *the spanning tree polynomial* of G as

$$T_G(\underline{\boldsymbol{z}}) := \frac{1}{\mathsf{nst}(G)} \sum_T \prod_{(r,s)\in T} z_{rs},$$

for $\underline{\boldsymbol{z}} \in \mathbb{R}_+^{d\times d}$, where the sum is over all spanning trees of G, and $\mathsf{nst}(G)$ is the number of spanning trees of G. In cases where G is not connected, we define the following polynomial

$$P_G := \prod_{l=1}^{\mathsf{ncc}(G)} T_{G_l},$$

where G_l is the lth connected component of G, and we denote by $\mathsf{ncc}(G)$ the number of connected components of G. This polynomial may be interpreted as the generating polynomial of *spanning forests* of G having exactly $\mathsf{ncc}(G)$ trees. The polynomials T_G and P_G are multi-affine, homogenous of degree $d-1$ for T_G (when G is connected) and $d-\mathsf{ncc}(G)$ for P_G, and do not depend on the diagonal entries $\{z_{rr} : 1 \le r \le d\}$. Furthermore, letting $z_{rs} = 1$ for all $r \neq s$, we have $P_G(\underline{\boldsymbol{z}}) = T_G(\underline{\boldsymbol{z}}) = 1$. We now provide tight asymptotic bounds on the collision probability $q(\underline{\boldsymbol{\mu}})$ when a subset E of the entries of $\underline{\boldsymbol{\mu}}$ become large.

Theorem 3. *Let $G = (V, E)$ with $V = \{1, \cdots, d\}$, $E = \{(r, s) \in V^2 \ : \ r \neq s\}$, and $\epsilon \in (0, 1)$. There exist two constants $0 < c_u < c_l$ depending on ϵ, d and α such that for all n sufficiently large, and all $\underline{\boldsymbol{\mu}} \in \{0, \cdots, n\}^{d\times d}$ with $\mu_{rs} \ge \epsilon n$ if and only if $(r, s) \in E$, we have*

$$c_l \frac{e^{-\vartheta_l(\underline{\boldsymbol{\mu}})}}{P_G(\underline{\boldsymbol{\mu}})^{1/2}} \le q(\underline{\boldsymbol{\mu}}) \le c_u \frac{e^{-\vartheta_u(\underline{\boldsymbol{\mu}})}}{P_G(\underline{\boldsymbol{\mu}})^{1/2}}.$$

with

$$\vartheta_u(\underline{\boldsymbol{\mu}}) = \inf_{\underline{\boldsymbol{\nu}}} \{\vartheta(\underline{\boldsymbol{\nu}}, \underline{\boldsymbol{\mu}}) : 0 \le \nu_{rs} \le \mu_{rs} \ \forall (r, s) \notin E\},$$

and

$$\vartheta_l(\underline{\boldsymbol{\mu}}) = \sup_{\underline{\boldsymbol{\nu}}} \{\vartheta(\underline{\boldsymbol{\nu}}, \underline{\boldsymbol{\mu}}) : 0 \le \nu_{rs} \le \mu_{rs} \ \forall (r, s) \notin E\}.$$

Let us now expand on the above result and derive some special cases and corollaries. First, we see that the collision probabilities can decay at two different speeds—polynomial or exponential—in the entries of the overlap matrix $\underline{\boldsymbol{\mu}}$, depending on whether $\vartheta_u(\underline{\boldsymbol{\mu}})$ (and/or $\vartheta_l(\underline{\boldsymbol{\mu}})$) is zero or strictly negative. Second, the apparent gap in the exponential decay of $q(\underline{\boldsymbol{\mu}})$ in the above characterization is artificial; one can make ϑ_u and ϑ_l equal by taking $\mu_{rs} = 0$ for all $(r, s) \notin E$. Alternatively, they could be made arbitrarily close to each other under an appropriate limit: Assume for simplicity that $\mu_{rs} = nw_{rs} > 0$ for all $(r, s) \in E$ for some $\underline{\boldsymbol{w}} \in [0, 1]^{d\times d}$. We have

$$\vartheta(\underline{\boldsymbol{\nu}}, \underline{\boldsymbol{\mu}}) = n\vartheta(\underline{\boldsymbol{\nu}}/n, \underline{\boldsymbol{w}}).$$

For $(r, s) \notin E$, we have $\mu_{rs} < \epsilon n$, therefore

$$\vartheta_u(\underline{\boldsymbol{\mu}})/n \le \inf_{\underline{\boldsymbol{x}}} \{\vartheta(\underline{\boldsymbol{x}}, \underline{\boldsymbol{w}}) : 0 \le x_{rs} \le \epsilon \ \forall (r, s) \notin E\} \underset{\epsilon\to 0}{\longrightarrow} \vartheta(\mathbf{0}, \underline{\boldsymbol{w}}).$$

The last step is justified by the continuity of $\vartheta(\cdot, \underline{\boldsymbol{w}})$. The same argument holds for $v_l(\underline{\boldsymbol{\mu}})$. Denoting the limiting function under this operation as $\vartheta(\underline{\boldsymbol{w}})$, we obtain:

$$\vartheta(\underline{\boldsymbol{w}}) = \sup_{\boldsymbol{\lambda}\in\mathbb{R}^d} \sum_{(r,s)\in E} w_{rs} \log\left(\frac{e^{\lambda_r-\lambda_s}}{\alpha+(1-\alpha)e^{\lambda_r-\lambda_s}}\right) = \min_{\substack{\underline{\boldsymbol{x}}\in[0,1]^{d\times d}\\ \underline{\boldsymbol{w}}\odot\underline{\boldsymbol{x}}\in\mathcal{F}}} \varphi_{\underline{\boldsymbol{w}}}(\underline{\boldsymbol{x}}).$$

The function ϑ can be seen as the exponential rate of decay of $q(\underline{\boldsymbol{\mu}})$. The reason ϑ_u and ϑ_l cannot (in general) be replaced by ϑ in Theorem 3 is that all control on the constants c_u and c_l is lost when $\epsilon \to 0$. Next, we identify the cases where this exponential decay is non-vacuous.

Lemma 4. *Let $\alpha \in (0,1)$, and $\underline{\boldsymbol{\mu}} \in \mathbb{R}_+^{d\times d}$. We have*

(i) $\vartheta(\underline{\boldsymbol{\mu}}) = 0$ if and only if $\underline{\boldsymbol{\mu}} \in \mathcal{F}$,

(ii) $\vartheta_u(\underline{\boldsymbol{\mu}}) = 0$ if and only if $\boldsymbol{M}_G(\alpha\underline{\boldsymbol{\mu}}, \underline{\boldsymbol{\nu}}) \in \mathcal{F}$ for some $\underline{\boldsymbol{\nu}} \in \mathbb{R}_+^{d\times d}$ such that $0 \le \nu_{rs} \le \mu_{rs}$ for all $(r,s) \notin E$.

Now we specialize Theorem 3 to the case where the entries of the overlap matrix are either zero or grow proportionally to n. From Theorem 3 and Lemma 4, we deduce a key corollary on the convergence of the properly rescaled logarithm of the collision probabilities.

Corollary 5. *Given a graph $G = (V, E)$, let $\underline{\boldsymbol{w}} \in [0,1]^{d\times d}$ be such that $w_{rs} > 0$ if and only if $(r,s) \in E$. If $\underline{\boldsymbol{w}} \in \mathcal{F}$ then*

$$\lim_{n\to\infty} \frac{\log q(n\underline{\boldsymbol{w}})}{\log n} = -\frac{d - \mathsf{ncc}(G)}{2}.$$

Otherwise if $\underline{\boldsymbol{w}} \notin \mathcal{F}$, then

$$\lim_{n\to\infty} \frac{\log q(n\underline{\boldsymbol{w}})}{n} = -\vartheta(\underline{\boldsymbol{w}}).$$

We see that the assignments τ such that $\underline{\boldsymbol{\mu}}(\tau, \tau^*) \in \mathcal{F}$ exhibit a much stronger correlation to τ^* than those for which this overlap matrix does not belong to $\mathcal{F}$, and will hence survive much longer as $n \to \infty$.

Proof of Lemma 4. Let $\underline{\boldsymbol{\mu}}, \underline{\boldsymbol{\nu}} \in \mathbb{R}_+^{d\times d}$ with $\underline{\boldsymbol{\mu}} \neq \underline{\boldsymbol{0}}$. Let $\alpha \in (0,1)$, and let $G = (V, E)$ denote a graph on d vertices. The function $\varphi_{\underline{\boldsymbol{\mu}}}$ defined in (8) is strictly convex on the support of $\underline{\boldsymbol{\mu}}$, i.e., on the subspace induced by the non-zero coordinates of $\underline{\boldsymbol{\mu}}$, so it admits a unique minimizer on the closed convex set $\{\underline{\boldsymbol{x}} \in [0,1]^{d\times d} \ : \ \boldsymbol{M}_G(\underline{\boldsymbol{x}}^* \odot \underline{\boldsymbol{\mu}}, \underline{\boldsymbol{\nu}}) \in \mathcal{F}\}$ intersected with that subspace. Let $\underline{\boldsymbol{x}}^*$ be this minimizer. By differentiating the associated Lagrangian, the entries of $\underline{\boldsymbol{x}}^*$ admit the expressions

$$x^*_{rs} = \frac{\alpha}{\alpha + (1-\alpha)e^{\lambda_r-\lambda_s}},$$

for all $(r,s) \in E$ (recall that $\mu_{rs} > 0$ for all such (r,s)), and where the vector $\boldsymbol{\lambda} \in \mathbb{R}^d$ is the unique solution up to global shifts of the system of equations

$$\sum_{s:(r,s)\in E} \frac{\alpha\mu_{rs}}{\alpha+(1-\alpha)e^{\lambda_r-\lambda_s}} + \sum_{s:(r,s)\notin E} \nu_{rs} = \sum_{s:(r,s)\in E} \frac{\alpha\mu_{sr}}{\alpha+(1-\alpha)e^{\lambda_s-\lambda_r}} + \sum_{s:(r,s)\notin E} \nu_{sr} \quad \forall r \in \{1,\cdots,d\}. \tag{10}$$

The claims of the lemma follow directly from the system of equations (10) and the fact that the non-negative function $\varphi_{\underline{\boldsymbol{\mu}}}$ vanishes if and only if $x^*_{rs} = \alpha$ for all $(r,s) \in E$: to show (i), we take $\underline{\boldsymbol{\nu}} = \underline{\boldsymbol{0}}$. It is clear from the equations that $\underline{\boldsymbol{\mu}} \in \mathcal{F}$ if and only if $\boldsymbol{\lambda} = c\boldsymbol{1}$, $c \in \mathbb{R}$, is a solution to the

above equations; and this is equivalent to $x^*_{rs} = \alpha$ whenever $\mu_{rs} > 0$. This is in turn equivalent to $\vartheta(\underline{\boldsymbol{\mu}}) = \varphi_{\underline{\boldsymbol{\mu}}}(\underline{\boldsymbol{x}}^*) = 0$. The same strategy is employed to show (ii), in conjunction with the continuity of the function $\underline{\boldsymbol{\nu}} \mapsto \vartheta(\underline{\boldsymbol{\nu}}, \underline{\boldsymbol{\mu}})$ over a compact domain (the infimum defining ϑ_u is attained). ∎

Proof of Corollary 5. Fix $G = (V, E)$, let $\underline{\boldsymbol{w}} \in (0,1)^{d\times d}$ with $w_{rs} > 0$ if and only if $(r, s) \in E$, and let n be an integer. For simplicity, assume that for $n\underline{\boldsymbol{w}}$ is an array of integer entries. The non-integer part introduces easily manageable error terms. Applying Theorem 3 with $\epsilon = \min_{(r,s)\in E} w_{rs}$, we have for n large

$$c_l P_G(n\underline{\boldsymbol{w}})^{-1/2} \exp -\vartheta_l(n\underline{\boldsymbol{w}}) \leq q(n\underline{\boldsymbol{w}}) \leq c_u P_G(n\underline{\boldsymbol{w}})^{-1/2} \exp -\vartheta_u(n\underline{\boldsymbol{w}}).$$

Moreover, since $w_{rs} = 0$ for $(r, s) \notin E$, we have

$$\vartheta_u(n\underline{\boldsymbol{w}}) = \vartheta_l(n\underline{\boldsymbol{w}}) = n\vartheta(\underline{\boldsymbol{w}}).$$

On the other hand, by homogeneity of the polynomial P_G, $P_G(n\underline{\boldsymbol{w}}) = n^{d-\mathsf{ncc}(G)} P_G(\underline{\boldsymbol{w}})$. Applying Lemma 4 yields the desired result: If $\underline{\boldsymbol{w}} \in \mathcal{F}$ then

$$\lim_{n\to\infty} \frac{\log q(n\underline{\boldsymbol{w}})}{\log n} = -\frac{d - \mathsf{ncc}(G)}{2}.$$

Otherwise,

$$\lim_{n\to\infty} \frac{\log q(n\underline{\boldsymbol{w}})}{n} = -\vartheta(\underline{\boldsymbol{w}}).$$

∎

3.1 A Gaussian integral

One important step in proving Theorem 3 (specifically for obtaining the polynomial decay part of $q(\underline{\boldsymbol{\mu}})$) is the following identity relating the Gaussian integral on a linear space $\mathcal{F}(G)$ defined based on a graph G to the spanning tree/forest polynomial of G. We denote by K_d the complete graph on d vertices where every pair of distinct vertices is connected by *two parallel edges*.

Proposition 6. *Let $G = (V, E)$ be a graph on d vertices, where self-loops and up to two parallel edges are allowed: $V = \{1, \cdots, d\}$, $E \subseteq V \times V$. Further, let*

$$\mathcal{F}(G) = \Big\{ \underline{\boldsymbol{x}} \in \mathcal{F} \ : \ x_{rs} = 0 \text{ for } (r, s) \notin E \Big\}.$$

For any array of positive real numbers $(w_{rs})_{(r,s)\in E}$, we have

$$\int_{\mathcal{F}(G)} e^{-\sum_{rs} x_{rs}^2/2w_{rs}} \, \mathrm{d}\underline{\boldsymbol{x}} = \left((2\pi)^{\dim(\mathcal{F}(G))} \ \frac{\prod_{r,s} w_{rs}}{P_G(\underline{\boldsymbol{w}})} \right)^{1/2}.$$

In the case where G is the complete graph K_d, $\mathcal{F}(G) = \mathcal{F}$, $\dim(\mathcal{F}) = (d-1)^2 + d$, and $P_G = T_G = (2^{d-1}d^{d-2})^{-1} \sum_T \prod_{(r,s)\in T} w_{rs}$ where the sum is over all spanning trees of K_d. The pre-factor in the last expression comes from Cayley's formula for the number of spanning trees of the complete graph. We will show that it suffices to prove Proposition 6 in the case where $G = K_d$ in order to establish it for any graph G. We were not able to locate this identity in the literature. To illuminate the combinatorial mechanisms behind it, we provide what appear to be two very different proofs of it. A first "direct" and purely combinatorial proof views $\mathcal{F}(G)$ as the space of *Eulerian flows* of the graph G. A second, slightly indirect proof which is mainly analytic, and relates the above Gaussian integral to the characteristic polynomial of the Laplacian matrix of G then invokes the Principal Minors Matrix Tree theorem (see, e.g., [Cha82]).

4 Computing the annealed free energy

In this section we establish the existence of $\mathfrak{F}(\gamma)$, and compute its value for all $\gamma > 0$. For $1 \le k \le d$ let $\mathcal{D}_k$ denote the set of binary matrices $\boldsymbol{X} \in \{0,1\}^{k\times d}$ such that each column of $\boldsymbol{X}$ contains *exactly* one non-zero entry and each row contains *at least* one non-zero entry. The elements of $\mathcal{D}_k$ represent partitions of the set $\{1, \cdots, d\}$ into k non-empty subsets.

Proposition 7. *Let $m = \gamma \frac{n}{\log n}$ with $\gamma > 0$ fixed for all $n \ge 2$. We have*

$$\mathfrak{F}(\gamma) = \max_{1\le k\le d-1} \left\{ H(\boldsymbol{\pi}) - \min_{\boldsymbol{X}\in\mathcal{D}_k} H(\boldsymbol{X}\boldsymbol{\pi}) - \frac{\gamma}{2}(d-k) \right\}.$$

Moreover, the inner minimization problem in the above expression can be solved explicitly:

Lemma 8. *Let $\boldsymbol{\pi}_{[\cdot]}$ be a permutation of the vector $\boldsymbol{\pi}$ such that $\pi_{[1]} \ge \pi_{[2]} \ge \cdots \ge \pi_{[d]}$. And for $1 \le k \le d-1$, let $\boldsymbol{\pi}^{(k)} \in \Delta^{k-1}$ defined as $\pi^{(k)}_1 = \sum_{r=1}^{d-k+1} \pi_{[r]}$ and $\pi^{(k)}_l = \pi_{[d-k+l]}$ for all $2 \le l \le k$ (if $k \ge 2$). Then*

$$\min_{\boldsymbol{X}\in\mathcal{D}_k} H(\boldsymbol{X}\boldsymbol{\pi}) = H(\boldsymbol{\pi}^{(k)}).$$

Theorem 2 follows from Proposition 7 and Lemma 8. We begin with the proof of the latter and devote the next subsection to the lengthier proof of the former.

Proof of Lemma 8. We start with an arbitrary partition of $\boldsymbol{\pi}$ into k groups, and define a sequence of operations on the set of k-partitions of $\boldsymbol{\pi}$ that strictly decreases $H(\underline{\boldsymbol{X}}\boldsymbol{\pi})$ at each step, and, irrespective of the starting point, always converges to $\boldsymbol{\pi}^{(k)}$. Starting with an arbitrary k-partition, write down the groups from left to right in decreasing order of total weight of each group. Initially, every group is marked *incomplete*. Then we perform the following operations:

1. Start with the rightmost incomplete group.
2. If it has more than one element, transfer the largest element to the leftmost group. This strictly decreases the entropy, since the heaviest group gets heavier and the lightest group gets lighter. Repeat this step until the rightmost group has exactly one element, and then move to the next step.
3. Consider this (now singleton) group. If there is no element to its left that is lighter than it, mark the group as complete. Else, swap this element with the lightest element to its left, and then mark it complete. Then go back to step 1.

■

4.1 Proof of Proposition 7

Let $m = \gamma \frac{n}{\log n}$. Recall from equation (6) that

$$\mathbb{E}[\mathcal{Z} - 1] = \binom{n}{n\boldsymbol{\pi}}^{-1} \cdot \sum_{\underline{\boldsymbol{\mu}}} \binom{n}{\underline{\boldsymbol{\mu}}} q(\underline{\boldsymbol{\mu}})^m \, \mathbb{1}\left\{\underline{\boldsymbol{\mu}}^\intercal \mathbf{1} = n\boldsymbol{\pi}\right\},$$

where the sum is over all arrays $\underline{\boldsymbol{\mu}} \in \mathbb{Z}_+^{d\times d}$ such that $\mathbf{1}^\intercal \underline{\boldsymbol{\mu}} \mathbf{1} = n$, $1 \le \sum_{r\ne s} \mu_{rs}$. Since the sum defining $\mathbb{E}[\mathcal{Z} - 1]$ is larger than its maximum term and smaller than the maximum term times

$(n+1)^{d^2}$, we only need to understand the convergence of the sequence

$$\mathfrak{F}_n := \frac{1}{n}\log\left(\max_{\substack{\underline{\boldsymbol{\mu}}\in\{0,\cdots,n\}^{d\times d}\\ \text{non-diagonal}}}\binom{n}{\underline{\boldsymbol{\mu}}}q(\underline{\boldsymbol{\mu}})^m\ \mathbb{1}\left\{\underline{\boldsymbol{\mu}}^\intercal\mathbf{1}=n\boldsymbol{\pi}\right\}\right)$$
$$=\max\left\{\frac{1}{n}\log\binom{n}{\underline{\boldsymbol{\mu}}}+\gamma\frac{\log q(\underline{\boldsymbol{\mu}})}{\log n}\ :\ \underline{\boldsymbol{\mu}}\in\{0,\cdots,n\}^{d\times d},\sum_{r\neq s}\mu_{rs}\geq 1,\underline{\boldsymbol{\mu}}^\intercal\mathbf{1}=n\boldsymbol{\pi}\right\}.$$

If this sequence converges, we would have

$$\mathfrak{F}(\gamma)=-H(\boldsymbol{\pi})+\lim_{n\to\infty}\mathfrak{F}_n, \tag{11}$$

since $\frac{1}{n}\log\binom{n}{n\boldsymbol{\pi}}\to H(\boldsymbol{\pi})$ by Stirling's formula. Next, we show that the above limit indeed exists. Let

$$\psi_n(\underline{\boldsymbol{w}}):=\frac{1}{n}\log\binom{n}{n\underline{\boldsymbol{w}}}+\gamma\frac{\log q(n\underline{\boldsymbol{w}})}{\log n}. \tag{12}$$

By Corollary 5, the function

$$\psi(\underline{\boldsymbol{w}}):=\begin{cases}H(\underline{\boldsymbol{w}})-\frac{\gamma}{2}(d-\mathsf{ncc}(\underline{\boldsymbol{w}})) & \text{if } \underline{\boldsymbol{w}}\in\mathcal{F},\\ -\infty & \text{otherwise,}\end{cases} \tag{13}$$

is the point-wise limit of the sequence of functions $\{\psi_n\}_{n\geq 2}$ on $\Delta^{d\times d-1}$. Next, we use the following lemma which states that any non-diagonal sequence of maximizers $\{\underline{\boldsymbol{\mu}}^{(n)}\}_n$ of ψ_n is such that $\sum_{r\neq s}\mu_{rs}^{(n)}$ grows proportionally to n.

Lemma 9. *For all $n\geq 2$, let*

$$\underline{\boldsymbol{\mu}}^{(n)}\in\arg\max\left\{\psi_n(\underline{\boldsymbol{\mu}}/n)\ :\ \underline{\boldsymbol{\mu}}\in\{0,\cdots,n\}^{d\times d},\ 1\leq\sum_{r\neq s}\mu_{rs}\leq n,\ \underline{\boldsymbol{\mu}}^\intercal\mathbf{1}=n\boldsymbol{\pi}\right\}.$$

It holds that

$$\liminf_{n\to\infty}\ \frac{\sum_{r\neq s}\mu_{rs}^{(n)}}{n}>0.$$

By Lemma 9, which we prove at the end of the current argument, we can safely restrict the set of candidate maximizers to those $\underline{\boldsymbol{\mu}}$ such that $\sum_{r\neq s}\mu_{rs}\geq c_0 n$ for some fixed but small $c_0>0$. From here, and by a change of variables $\underline{\boldsymbol{\mu}}=n\underline{\boldsymbol{w}}$, mere point-wise convergence suffices to interchange lim inf and sup:

$$\liminf_{n\to\infty}\mathfrak{F}_n\geq\liminf_{n\to\infty}\ \sup\left\{\psi_n(\underline{\boldsymbol{w}})\ :\ \underline{\boldsymbol{w}}\in\{i/n:0\leq i\leq n\}^{d\times d},\ c_0\leq\sum_{r\neq s}w_{rs}\leq 1,\ \underline{\boldsymbol{w}}^\intercal\mathbf{1}=\boldsymbol{\pi}\right\}$$
$$\geq\sup\left\{\psi(\underline{\boldsymbol{w}})\ :\ \underline{\boldsymbol{w}}\in[0,1]^{d\times d}\cap\mathcal{F},\ c_0\leq\sum_{r\neq s}w_{rs}\leq 1,\ \underline{\boldsymbol{w}}^\intercal\mathbf{1}=\boldsymbol{\pi}\right\}. \tag{14}$$

Now we present a matching upper bound for $\limsup \mathfrak{F}_n$. For $\epsilon > 0$, let $G_n = (\{1, \cdots, d\}, E_n)$ be defined such that $(r,s) \in E_n$ if and only if $w_{rs}^{(n)} \geq \epsilon$. Let $(G_l)_{l=1}^k$ denote the connected components of the graph G_n, $k = \mathsf{ncc}(G_n)$. Also, for $\underline{\boldsymbol{w}}$ an array for positive entries, let $\mathsf{ncc}^\epsilon(\underline{\boldsymbol{w}})$ denote the number of connected components of the graph $G(\underline{\boldsymbol{w}}, \epsilon) = (V, E(\underline{\boldsymbol{w}}, \epsilon))$, $V = \{1, \cdots, d\}$, $E(\underline{\boldsymbol{w}}, \epsilon) = \{(r,s) \ : \ r \neq s, \ w_{rs} > \epsilon\}$, and let

$$\vartheta^\epsilon(\underline{\boldsymbol{w}}) := \inf_{\underline{\boldsymbol{x}}} \{\vartheta(\underline{\boldsymbol{x}}, \underline{\boldsymbol{w}}) : 0 \leq x_{rs} \leq \epsilon \ \forall (r,s) \notin E(\underline{\boldsymbol{w}}, \epsilon)\}.$$

We will also write $\mathsf{ncc}(\underline{\boldsymbol{w}})$ for $\mathsf{ncc}^0(\underline{\boldsymbol{w}})$. Let $\underline{\boldsymbol{w}}^{(n)} = \underline{\boldsymbol{\mu}}^{(n)}/n$ for all $n \geq 2$, where $\underline{\boldsymbol{\mu}}^{(n)}$ is defined in Lemma 9. By Theorem 3, we have for n sufficiently large

$$q(n\underline{\boldsymbol{w}}^{(n)}) \leq c_u(\epsilon, d, \alpha) P_{G_n}(n\underline{\boldsymbol{w}}^{(n)})^{-1/2} \exp -\vartheta^\epsilon(n\underline{\boldsymbol{w}}^{(n)}).$$

Since $w_{rs}^{(n)} \geq \epsilon$ of all the edges (r,s) of G_n^ϵ, $\prod_l T_{G_l}(\underline{\boldsymbol{w}}^{(n)})$ is bounded below by ϵ^d *independently of* n. Therefore, for n sufficiently large,

$$\begin{aligned}
\psi_n(\underline{\boldsymbol{w}}^{(n)}) &= \frac{1}{n} \log \binom{n}{n\underline{\boldsymbol{w}}^{(n)}} + \gamma \frac{\log q(n\underline{\boldsymbol{w}}^{(n)})}{\log n} \\
&\leq \frac{1}{n} \log \binom{n}{n\underline{\boldsymbol{w}}^{(n)}} - \frac{\gamma}{2}(d - \mathsf{ncc}^\epsilon(\underline{\boldsymbol{w}}^{(n)})) - \frac{\gamma n}{\log n} \vartheta^\epsilon(\underline{\boldsymbol{w}}^{(n)}) + \mathcal{O}\left(\frac{\log c_u(\epsilon, d, \alpha) + d \log(1/\epsilon)}{\log n}\right) \\
&\leq \sup \left\{ H(\underline{\boldsymbol{w}}) - \frac{\gamma}{2}(d - \mathsf{ncc}^\epsilon(\underline{\boldsymbol{w}})) - \frac{\gamma n}{\log n} \vartheta^\epsilon(\underline{\boldsymbol{w}}) \ : \ \underline{\boldsymbol{w}} \in [0,1]^{d \times d}, c_0 \leq \sum_{r \neq s} w_{rs} \leq 1, \ \underline{\boldsymbol{w}}^\intercal \mathbf{1} = \boldsymbol{\pi} \right\} \\
&\quad + \mathcal{O}\left(\frac{\log c_u(\epsilon, d, \alpha) + d \log(1/\epsilon)}{\log n}\right),
\end{aligned}$$

where the last inequality is obtained by Stirling's formula and taking a supremum over all $\underline{\boldsymbol{w}}$. By Lemma 4, $\vartheta^\epsilon(\underline{\boldsymbol{w}}) = 0$ if and only if $\boldsymbol{M}_G(\alpha \underline{\boldsymbol{w}}, \underline{\boldsymbol{x}}) \in \mathcal{F}$ for some $\underline{\boldsymbol{x}} \in [0,1]^{d \times d}$ such that $0 \leq x_{rs} \leq \epsilon$ for all $(r,s) \notin E$, $G = (V, E)$ being the graph whose edges are $(r,s) : w_{rs} \geq \epsilon$. This constrains the supremum to be achieved in the space of such $\underline{\boldsymbol{w}}$ for n sufficiently large. Moreover, this condition implies in particular that

$$\|\underline{\boldsymbol{w}}\mathbf{1} - \underline{\boldsymbol{w}}^\intercal \mathbf{1}\|_{\ell_\infty} \leq 2d\alpha^{-1}\epsilon,$$

where $\|\cdot\|_{\ell_\infty}$ is the ℓ_∞ norm of a vector in $\mathbb{R}^d$. Consequently, this yields the following upper bound as $n \to \infty$,

$$\limsup_{n \to \infty} \mathfrak{F}_n \leq \sup \left\{ H(\underline{\boldsymbol{w}}) - \frac{\gamma}{2}(d - \mathsf{ncc}^\epsilon(\underline{\boldsymbol{w}})) \ : \ \begin{array}{c} \underline{\boldsymbol{w}} \in [0,1]^{d \times d}, \|\underline{\boldsymbol{w}}\mathbf{1} - \underline{\boldsymbol{w}}^\intercal \mathbf{1}\|_{\ell_\infty} \leq 2d\alpha^{-1}\epsilon, \\ c_0 \leq \sum_{r \neq s} w_{rs} \leq 1, \ \underline{\boldsymbol{w}}^\intercal \mathbf{1} = \boldsymbol{\pi} \end{array} \right\}, \quad (15)$$

for all $\epsilon > 0$. Next, we argue that as $\epsilon \to 0$, the right-hand side of the above inequality converges to

$$\sup \left\{ H(\underline{\boldsymbol{w}}) - \frac{\gamma}{2}(d - \mathsf{ncc}(\underline{\boldsymbol{w}})) \ : \ \underline{\boldsymbol{w}} \in [0,1]^{d \times d} \cap \mathcal{F}, c_0 \leq \sum_{r \neq s} w_{rs} \leq 1, \ \underline{\boldsymbol{w}}^\intercal \mathbf{1} = \boldsymbol{\pi} \right\},$$

thereby establishing the existence of the limit $\lim \mathfrak{F}_n$ along with its precise value. Since the function $\epsilon \to \mathsf{ncc}^\epsilon(\underline{\boldsymbol{w}})$ is non-decreasing for any fixed $\underline{\boldsymbol{w}}$, the limit of the right-hand side of (15) as $\epsilon \to 0$ exists by monotone convergence. The limit can be decomposed as

$$\begin{aligned}
&\lim_{\epsilon \to 0} \ \sup \left\{ H(\underline{\boldsymbol{w}}) - \frac{\gamma}{2}(d - \mathsf{ncc}^\epsilon(\underline{\boldsymbol{w}})) \ : \ \begin{array}{c} \underline{\boldsymbol{w}} \in [0,1]^{d \times d}, \|\underline{\boldsymbol{w}}\mathbf{1} - \underline{\boldsymbol{w}}^\intercal \mathbf{1}\|_{\ell_\infty} \leq 2d\alpha^{-1}\epsilon, \\ c_0 \leq \sum_{r \neq s} w_{rs} \leq 1, \ \underline{\boldsymbol{w}}^\intercal \mathbf{1} = \boldsymbol{\pi} \end{array} \right\} \\
&= \max_{1 \leq k \leq d} \max_{\{V_l\}_{l=1}^k} \lim_{\epsilon \to 0} \ \sup \left\{ H(\underline{\boldsymbol{w}}) - \frac{\gamma}{2}(d - k) \ : \ \begin{array}{c} \underline{\boldsymbol{w}} \in [0,1]^{d \times d}, \ \|\underline{\boldsymbol{w}}\mathbf{1} - \underline{\boldsymbol{w}}^\intercal \mathbf{1}\|_{\ell_\infty} \leq 2d\alpha^{-1}\epsilon, \\ w_{rs} \leq \epsilon \ \forall (r,s) \in V_l \times V_{l'}, \ l \neq l', \\ G_l(\underline{\boldsymbol{w}}) \text{ is connected } \forall l, \ c_0 \leq \sum_{r \neq s} w_{rs} \leq 1, \ \underline{\boldsymbol{w}}^\intercal \mathbf{1} = \boldsymbol{\pi} \end{array} \right\},
\end{aligned}$$

where $\{V_l\}_{l=1}^k$ ranges over a partitions of the set $\{1,\cdots,d\}$ with k non-empty subsets, and $G_l(\underline{\boldsymbol{w}}) = (V_l, \{(r,s) \in V_l \times V_l \ : \ w_{rs} > \epsilon\})$ for all $1 \leq l \leq k$. Letting $\epsilon < c_0$, the range of the outer-most maximum becomes $1 \leq k \leq d-1$. By concavity of the entropy, the constraint that the graphs $G_l(\underline{\boldsymbol{w}})$ must be connected can be safely removed from the maximization problem without changing its maximum value since it will be automatically satisfied. Thus, the inner-most optimization problem is that of a continuous function on a closed and bounded domain that shrinks with ϵ. Its value is therefore a continuous function of ϵ. Hence, by sending ϵ to 0, in conjunction with the lower bound (14), we conclude that

$$\lim_{n\to\infty} \mathfrak{F}_n = \sup\left\{ \psi(\underline{\boldsymbol{w}}) \ : \ \underline{\boldsymbol{w}} \in [0,1]^{d\times d},\ c_0 \leq \sum_{r\neq s} w_{rs} \leq 1,\ \underline{\boldsymbol{w}}\mathbf{1} = \underline{\boldsymbol{w}}^\intercal \mathbf{1} = \boldsymbol{\pi} \right\}. \tag{16}$$

As a final step, we make the above expression a bit more explicit. As argued previously, the supremum in (16) can be decomposed such that one first takes the maximum of $\psi(\underline{\boldsymbol{w}})$ over all $\underline{\boldsymbol{w}}$ such that $w_{rs} = 0$ for all $(r,s) \in V_l \times V_{l'},\ l \neq l'$ where $\{V_l\}_{1\leq l\leq k}$ is a fixed partition of $\{1,\cdots,d\}$ into non-empty subsets, then maximize over all such partitions, then over all $1 \leq k \leq d-1$. The first optimization problem has a value

$$\sup\left\{ H(\underline{\boldsymbol{w}}) - \frac{\gamma}{2}(d-k) \ : \ \underline{\boldsymbol{w}} \in [0,1]^{d\times d},\ w_{rs} = 0,\ (r,s) \in V_l \times V_{l'},\ l \neq l',\ \underline{\boldsymbol{w}}\mathbf{1} = \underline{\boldsymbol{w}}^\intercal \mathbf{1} = \boldsymbol{\pi} \right\},$$

where the constraint $c_0 \leq \sum_{r\neq s} w_{rs} \leq 1$ is not active for c_0 small enough, hence can be removed. Let $\underline{\boldsymbol{w}}$ be in the above constraint set. Then $H(\underline{\boldsymbol{w}}) = -\sum_{l=1}^k \sum_{(r,s)\in V_l\times V_l} w_{rs} \log w_{rs}$, and this is maximized at

$$w^*_{rs} = \begin{cases} (\pi_r \pi_s)/\sum_{r'\in V_l} \pi_{r'} & \text{if } (r,s) \in V_l \times V_l,\ l \in \{1,\cdots,k\}, \\ 0 & \text{otherwise,} \end{cases} \tag{17}$$

with maximum value

$$\begin{aligned} H(\underline{\boldsymbol{w}}^*) &= 2H(\boldsymbol{\pi}) + \sum_{l=1}^k \left(\sum_{r\in V_l} \pi_r \right) \log \left(\sum_{r\in V_l} \pi_r \right), \\ &= 2H(\boldsymbol{\pi}) - H(\boldsymbol{X}\boldsymbol{\pi}), \end{aligned} \tag{18}$$

where $\boldsymbol{X} \in \{0,1\}^{k\times d}$, $X_{l,r} = 1$ if and only if $r \in V_l$. Note that $\mathcal{D}_k$ is the set of all such matrices (each one corresponding to a partition $\{V_l\}$ of $\{1,\cdots,d\}$). Finally, by maximizing over all possible partitions, and using (11) we get

$$\mathfrak{F}(\gamma) = \max_{1\leq k\leq d-1} \left\{ H(\boldsymbol{\pi}) - \min_{\boldsymbol{X}\in\mathcal{D}_k} H(\boldsymbol{X}\boldsymbol{\pi}) - \frac{\gamma}{2}(d-k) \right\}.$$

This completes the proof of Proposition 7, except for the proof Lemma 9, which we provide below.

Proof of Lemma 9. Let

$$\underline{\boldsymbol{\mu}}^{(n)} \in \arg\max \left\{ \psi_n(\underline{\boldsymbol{\mu}}/n) \ : \ \underline{\boldsymbol{\mu}} \in \{0,\cdots,n\}^{d\times d},\ 1 \leq \sum_{r\neq s} \mu_{rs},\ \underline{\boldsymbol{\mu}}^\intercal \mathbf{1} = n\boldsymbol{\pi} \right\}.$$

We show that

$$\liminf_{n\to\infty}\ n^{-1}\sum_{r\neq s}\mu_{rs}^{(n)} > 0.$$

Let us first show that

$$\frac{(\log n)^3}{n}\sum_{r\neq s}\mu_{rs}^{(n)} \longrightarrow \infty,$$

and then remove the logarithmic factor. We proceed by contradiction, by showing that if the above statement is not true, then the expected number of non-planted solutions $\mathbb{E}[\mathcal{Z}-1]$ vanishes as $n\to\infty$ for any $\gamma>0$, which contradicts our lower bound of Theorem 1. We have

$$\mathbb{E}\left[\mathcal{Z}-1\right] \le \binom{n}{n\boldsymbol{\pi}}^{-1}\cdot(n+1)^{d^2}\cdot\binom{n}{\underline{\boldsymbol{\mu}}^{(n)}}\cdot q_{\max}^{\gamma n/\log n},$$

with $q_{\max}=\max\left\{q(\underline{\boldsymbol{\mu}})\ :\ 1\le\sum_{r\neq s}\mu_{rs}\right\}<1$. Moreover,

$$\binom{n}{\underline{\boldsymbol{\mu}}^{(n)}} = \binom{n}{n\boldsymbol{\pi}}\prod_{r=1}^{d}\frac{(n\pi_r)!}{\prod_{s\neq r}\mu_{sr}!(n\pi_r-\sum_{s\neq r}\mu_{sr})!} \le \binom{n}{n\boldsymbol{\pi}}\prod_{r=1}^{d}(n\pi_r)^{\sum_{s\neq r}\mu_{sr}}.$$

If $\sum_{r\neq s}\mu_{rs}^{(n)}\le Cn/(\log n)^3$ for some constant $C>0$, then

$$\mathbb{E}\left[\mathcal{Z}-1\right] \le (n+1)^{d^2}\cdot n^{Cn/(\log n)^3}\cdot q_{\max}^{\gamma n/\log n} \underset{n\to\infty}{\longrightarrow} 0,$$

for all $\gamma>0$, and this contradicts the fact that below γ_{low} there are exponentially many distinct satisfying assignments.

Now let us assume that $\frac{(\log n)^3}{n}\sum_{r\neq s}\mu_{rs}^{(n)}\to\infty$ but $\liminf n^{-1}\sum_{r\neq s}\mu_{rs}^{(n)}=0$. We proceed by contradiction once more, and construct a sequence of points that have a higher objective value than $\underline{\boldsymbol{\mu}}^{(n)}$. Instead of working with convergent subsequences, we may as well assume that $\{\underline{\boldsymbol{\mu}}^{(n)}\}$ is convergent. Let

$$E_n=\left\{(r,s)\ :\ r\neq s,\ \mu_{rs}^{(n)}>\epsilon\sum_{r\neq s}\mu_{rs}^{(n)}\right\},\quad\text{and}\quad E_\infty=\left\{(r,s)\ :\ r\neq s,\ \liminf_{n\to\infty}\frac{\mu_{rs}^{(n)}}{\sum_{r\neq s}\mu_{rs}^{(n)}}>0\right\},$$

for all n and some $\epsilon>0$ sufficiently small. Let $k_n=\mathsf{ncc}(G_n)$ be the number of connected components of the graph $G_n=(\{1,\cdots,d\},E_n)$, and similarly, let $k_\infty=\mathsf{ncc}(G_\infty)$ with $G_\infty=(\{1,\cdots,d\},E_\infty)$. Observe that E_∞ and E_n are both non-empty sets, hence $k_\infty,k_n\le d-1$ for all n.

Now we consider an arbitrary partition of the set of vertices $\{1,\cdots,d\}$ into k_∞ subsets $\{V_l\}_{1\le l\le k_\infty}$, and let G be the graph on d vertices with edge set $\cup_{l=1}^{k_\infty}V_l\times V_l$; i.e., G is the union of k_∞ *complete* connected components. Finally, let $\underline{\boldsymbol{v}}^{(n)}:=n\underline{\boldsymbol{w}}$ for all n, with

$$w_{rs}=\begin{cases}(\pi_r\pi_s)/\sum_{r'\in V_l}\pi_{r'} & \text{if } (r,s)\in V_l\times V_l,\ l\in\{1,\cdots,k_\infty\},\\ 0 & \text{otherwise,}\end{cases}$$

Recall that this construction provides one of the candidate maximizers of the annealed free energy (see (17)). Observe that $\underline{\boldsymbol{v}}^{(n)}$ satisfies all the constraints satisfied by $\underline{\boldsymbol{\mu}}^{(n)}$, and additionally, $\underline{\boldsymbol{v}}^{(n)}\in\mathcal{F}$. Therefore, by Corollary 5, we have

$$\psi_n(\underline{\boldsymbol{v}}^{(n)}/n)=H(\underline{\boldsymbol{w}})-\frac{\gamma}{2}(d-k_\infty)+o_n(1).$$

Recall that the function ψ_n is defined in (12). On the other hand, to study the asymptotics of $\psi_n(\underline{\boldsymbol{\mu}}^{(n)}/n)$, we apply Theorem 3 with n replaced by $\sum_{r\neq s}\mu_{rs}^{(n)}$ (which grows to infinity), and we get

$$\psi_n(\underline{\boldsymbol{\mu}}^{(n)}/n) \leq H(\boldsymbol{\pi}) - \frac{\gamma}{2}(d-k_n)\left(1-3\frac{\log\log n}{\log n}\right) - \frac{\vartheta_u(\underline{\boldsymbol{\mu}}^{(n)})}{\log n} + o_n(1).$$

The term in the right-hand side follows from Stirling's formula and the fact that $\mu_{rs}^{(n)}/n \to 0$ for all $r \neq s$. The second term follows from the fact that

$$P_{G_n}(\underline{\boldsymbol{\mu}}^{(n)}) \geq \left(\epsilon\sum_{r\neq s}\mu_{rs}^{(n)}\right)^{d-k_n} \gg \left(\frac{n}{(\log n)^3}\right)^{d-k_n}.$$

Next, we argue based on these estimates that $\psi_n(\underline{\boldsymbol{v}}^{(n)}/n) > \psi_n(\underline{\boldsymbol{\mu}}^{(n)}/n)$ for all n large enough. First, the term involving ϑ_u in the upper bound on $\psi_n(\underline{\boldsymbol{\mu}}^{(n)}/n)$ can be dropped since it is always non-negative. By direct computation (we already showed this in (18)), we have

$$H(\underline{\boldsymbol{w}}) - H(\boldsymbol{\pi}) = H(\boldsymbol{\pi}) - H(\boldsymbol{p}),$$

with $\boldsymbol{p} \in \Delta^{k_\infty - 1}$ with $p_l = \sum_{r\in V_l}\pi_r$ for all $1 \leq l \leq k_\infty$. We show that the right-hand side of this equality is strictly positive:

$$\begin{aligned}
H(\boldsymbol{\pi}) - H(\boldsymbol{p}) &= -\sum_{r=1}^{d}\pi_r\log\pi_r + \sum_{l=1}^{k_\infty}\left(\sum_{r\in V_l}\pi_r\right)\log\left(\sum_{r\in V_l}\pi_r\right)\\
&= -\sum_{l=1}^{k_\infty}\sum_{r\in V_l}\pi_r\log\left(\frac{\pi_r}{p_l}\right)\\
&= -\sum_{l=1}^{k_\infty}p_l\sum_{r\in V_l}\frac{\pi_r}{p_l}\log\left(\frac{\pi_r}{p_l}\right)\\
&\geq -\sum_{l=1}^{k_\infty}p_l\log\left(\frac{\sum_{r\in V_l}\pi_r^2}{p_l^2}\right),\\
&\geq 0.
\end{aligned}$$

We used Jensen's inequality on the concave function $x \mapsto \log x$, and the fact that $\sum_{r\in V_l}\pi_r^2 \leq p_l\sum_{r\in V_l}\pi_r = p_l^2$ for all l. Moreover, since all coordinates of $\boldsymbol{\pi}$ are strictly positive, equality holds if and only if $\pi_r = p_l$ for all l and $r \in V_l$ which implies that the partition must be trivial; i.e., $k_\infty = d$. Recall that this does not happen since E_∞ is non-empty.

On the other hand, by setting ϵ sufficiently small (smaller than all the limits in the definition of E_∞), any edge in E_∞ will eventually (and permanently from then on) be in E_n. Therefore the number of connected components of G_n does not exceed that of G_∞: $k_n \leq k_\infty$ for n sufficiently large. We conclude that $\psi_n(\underline{\boldsymbol{v}}^{(n)}/n) > \psi_n(\underline{\boldsymbol{\mu}}^{(n)}/n)$ for all n large enough. Therefore $\underline{\boldsymbol{\mu}}^{(n)}$ is not always a maximizer of ψ_n, and this leads to a contradiction. ∎

5 Proof of Theorem 3

Our proof is based on the method of Laplace from asymptotic analysis: when the entries of $\underline{\boldsymbol{\mu}}$ are large, the sum defining $q(\underline{\boldsymbol{\mu}})$ is dominated by its largest term corrected by a sub-exponential term

which is represented by a Gaussian integral (see, e.g., [DB70] for the univariate case). Since we are in a multivariate situation, the asymptotics of q depend on which subset of the entries of $\underline{\boldsymbol{\mu}}$ are large. Our approach is inspired by [AN05]. We recall that for $\underline{\boldsymbol{\mu}} \in \mathbb{Z}_+^{d\times d}$,

$$q(\underline{\boldsymbol{\mu}}) = \sum_{\substack{\underline{\boldsymbol{\nu}}\in\mathbb{Z}_+^{d\times d}\cap\mathcal{F} \\ 0\le\nu_{rs}\le\mu_{rs}}} \left(\prod_{r,s=1}^{d} \binom{\mu_{rs}}{\nu_{rs}} \alpha^{\nu_{rs}}(1-\alpha)^{\mu_{rs}-\nu_{rs}} \right).$$

Let $G=(V,E)$ with $V=\{1,\cdots,d\}$ and $E=\{(r,s)\in V^2 \ : \ r\neq s\}$. The graph G will be used to store information about which entries of $\underline{\boldsymbol{\mu}}$ are going to infinity linearly in n, and which entries are not. We can split the sum defining q into a double sum, one involving the large terms (A in subsequent notation), and the rest:

$$q(\underline{\boldsymbol{\mu}}) = \sum_{\substack{0\le\nu'_{rs}\le\mu_{rs} \\ (r,s)\notin E}} \prod_{(r,s)\notin E} \binom{\mu_{rs}}{\nu'_{rs}} \alpha^{\nu'_{rs}}(1-\alpha)^{\mu_{rs}-\nu'_{rs}} A(\underline{\boldsymbol{\nu}}', \underline{\boldsymbol{\mu}}),$$

with

$$A(\underline{\boldsymbol{\nu}}', \underline{\boldsymbol{\mu}}) = \sum_{\substack{0\le\nu_{rs}\le\mu_{rs} \\ (r,s)\in E}} \prod_{(r,s)\in E} \binom{\mu_{rs}}{\nu_{rs}} \alpha^{\nu_{rs}}(1-\alpha)^{\mu_{rs}-\nu_{rs}} \mathbb{1}\left\{ \boldsymbol{M}_G(\underline{\boldsymbol{\nu}}, \underline{\boldsymbol{\nu}}') \in \mathcal{F} \right\},$$

where for two $d\times d$ matrices $\underline{\boldsymbol{a}}, \underline{\boldsymbol{b}}$, $\boldsymbol{M}_G(\underline{\boldsymbol{a}}, \underline{\boldsymbol{b}})$ is the $d\times d$ matrix with entries a_{rs} if $(r,s)\in E$ and b_{rs} otherwise. The quantity A will be approximated using the Laplace method. Recall from the expressions (8) and (9) that

$$\varphi_{\underline{\boldsymbol{\mu}}}(\underline{\boldsymbol{x}}) = \sum_{(r,s)\in E} \mu_{rs} D(x_{rs} \parallel \alpha),$$

and

$$\vartheta(\underline{\boldsymbol{\nu}}, \underline{\boldsymbol{\mu}}) = \min_{\substack{\underline{\boldsymbol{x}}\in[0,1]^{d\times d} \\ \boldsymbol{M}_G(\underline{\boldsymbol{x}}\odot\underline{\boldsymbol{\mu}}, \underline{\boldsymbol{\nu}})\in\mathcal{F}}} \varphi_{\underline{\boldsymbol{\mu}}}(\underline{\boldsymbol{x}}).$$

Let $\underline{\boldsymbol{x}}^*(\underline{\boldsymbol{\nu}}, \underline{\boldsymbol{\mu}})$ be the optimal solution of the above optimization problem.

Before stating our asymptotic approximation result for A, we state an important lemma on the boundedness of the entries of $\underline{\boldsymbol{x}}^*(\underline{\boldsymbol{\nu}}, \underline{\boldsymbol{\mu}})$, where the bounds depend only on ϵ and α.

Lemma 10. *Let G be fixed as above, $\alpha\in(0,1)$ and $\epsilon\in(0,1)$. There exist two constants $0<c_l\le c_u<1$ depending only on d, α and ϵ such that the following is true: For all integers $n\ge 1$, and $\underline{\boldsymbol{\mu}}\in\{0,\cdots,n\}^{d\times d}$ such that $\mu_{rs}\ge\epsilon n$ iff $(r,s)\in E$. For all $\underline{\boldsymbol{\nu}}'\in\{0,\cdots,n\}^{\bar{E}}$ such that $0\le\nu'_{rs}\le\mu_{rs}$ for all $(r,s)\notin E$, we have*

$$c_l \le \min_{(r,s)\in E} x^*_{rs} \le \max_{(r,s)\in E} x^*_{rs} \le c_u.$$

Therefore, the entries of $\underline{\boldsymbol{x}}^*$ can effectively be treated as constants throughout the rest of the proof. Now we state our asymptotic estimate for A.

Proposition 11. *Let G be fixed as above, and $\epsilon>0$. For all n sufficiently large, all $\underline{\boldsymbol{\mu}}\in\{0,\cdots,n\}^{d\times d}$ with $\mu_{rs}\ge\epsilon n$ iff $(r,s)\in E$, and all $\underline{\boldsymbol{\nu}}'\in\{0,\cdots,n\}^{\bar{E}}$ such that $0\le\nu'_{rs}\le\mu_{rs}$ for all $(r,s)\notin E$, we have*

$$A(\underline{\boldsymbol{\nu}}', \underline{\boldsymbol{\mu}}) \asymp_{G,d,\epsilon,\alpha} \frac{e^{-\vartheta(\underline{\boldsymbol{\nu}}', \underline{\boldsymbol{\mu}})}}{P_G(\underline{\boldsymbol{\mu}})^{1/2}}.$$

Here, the symbol " $\asymp_{G,d,\epsilon,\alpha}$ " means that the ratio is upper- and lower-bounded by constants depending only on G, d, ϵ and α.

By the above proposition, we have

$$q(\underline{\boldsymbol{\mu}}) \quad \asymp_{G,d,\epsilon,\alpha} \sum_{\substack{\underline{\boldsymbol{\nu}} \in \mathbb{Z}_+^{\bar{E}} \\ 0 \le \nu_{rs} \le \mu_{rs}}} \left(\prod_{(r,s) \notin E} \binom{\mu_{rs}}{\nu_{rs}} \alpha^{\nu_{rs}} (1-\alpha)^{\mu_{rs}-\nu_{rs}} \right) \frac{e^{-\vartheta(\underline{\boldsymbol{\nu}},\underline{\boldsymbol{\mu}})}}{P_G(\underline{\boldsymbol{\mu}})^{1/2}}.$$

The estimate above (ignoring the term $P_G(\underline{\boldsymbol{\mu}})$) can be interpreted as the expected value of the function $e^{-\vartheta(\underline{\boldsymbol{\nu}},\underline{\boldsymbol{\mu}})}$ under the law of the random variable $\underline{\boldsymbol{\nu}}$ where each entry ν_{rs} for $(r,s) \notin E$ is independently binomial with parameters α and μ_{rs}. From here, the bounds claimed in Theorem 3 follow immediately.

Proof of Proposition 11. We will show that

$$A(\underline{\boldsymbol{\nu}}', \underline{\boldsymbol{\mu}}) \quad \asymp_{G,d,\epsilon,\alpha} \quad e^{-\vartheta(\underline{\boldsymbol{\nu}}',\underline{\boldsymbol{\mu}})} \prod_{(r,s)\in E} \mu_{rs}^{-1/2} \int_{\mathcal{F}(G)} e^{-\sum_{(r,s)\in E} z_{rs}^2/2\mu_{rs}} \, \mathrm{d}\underline{\boldsymbol{z}}.$$

Then the result follows by applying Proposition 6 to evaluate the Gaussian integral. We proceed by showing the upper and lower bounds separately.

The upper bound For $\underline{\boldsymbol{\nu}}' \in \mathbb{Z}_+^{\bar{E}}, \underline{\boldsymbol{\mu}} \in \mathbb{Z}_+^{d\times d}$ fixed and some parameter $C(\underline{\boldsymbol{\mu}}) > 0$ to be adjusted, let

$$\Omega = \left\{ \underline{\boldsymbol{\nu}} \in \mathbb{Z}_+^E \ : \ \boldsymbol{M}_G(\underline{\boldsymbol{\nu}}, \underline{\boldsymbol{\nu}}') \in \mathcal{F}, \ 0 \le \nu_{rs} \le \mu_{rs}, \ \sum_{(r,s)\in E} \frac{(\nu_{rs} - x_{rs}^* \mu_{rs})^2}{x_{rs}^*(1-x_{rs}^*)\mu_{rs}} \le C(\underline{\boldsymbol{\mu}})^2 \right\}.$$

For $\underline{\boldsymbol{\nu}} \in \mathbb{Z}_+^E$, we let $\underline{\boldsymbol{x}} \in [0,1]^E$ defined by $x_{rs} = \nu_{rs}/\mu_{rs}$ for all $(r,s) \in E$. We upper bound the binomial coefficients $\binom{\mu_{rs}}{\nu_{rs}}$ based on whether $\underline{\boldsymbol{\nu}}$ is in Ω or not:

- If $\underline{\boldsymbol{\nu}} \in \Omega$ we use the upper bound $\binom{\mu_{rs}}{\nu_{rs}} \le (2\pi\mu_{rs}x_{rs}(1-x_{rs}))^{-1/2} \exp \mu_{rs} H(x_{rs})$.
- Otherwise, we use the upper bound $\binom{\mu_{rs}}{\nu_{rs}} \le 3\sqrt{\mu_{rs}} \exp \mu_{rs} H(x_{rs})$.

Here, $H(x_{rs}) = -x_{rs}\log x_{rs} - (1-x_{rs})\log(1-x_{rs})$. Thus, the summand in $A(\underline{\boldsymbol{\nu}}', \underline{\boldsymbol{\mu}})$ is bounded by

$$\prod_{(r,s)\in E} (2\pi\mu_{rs}x_{rs}(1-x_{rs}))^{-1/2} \exp \mu_{rs} D(x_{rs} \parallel \alpha) = \prod_{(r,s)\in E} (2\pi\mu_{rs}x_{rs}(1-x_{rs}))^{-1/2} \exp(-\varphi_{\underline{\boldsymbol{\mu}}}(\underline{\boldsymbol{x}}))$$

if $\underline{\boldsymbol{\nu}} \in \Omega$, and

$$\prod_{(r,s)\in E} 3\mu_{rs}^{1/2} \exp(-\varphi_{\underline{\boldsymbol{\mu}}}(\underline{\boldsymbol{x}}))$$

if not. The function $\varphi_{\underline{\boldsymbol{\mu}}}$ is smooth, and we have $\frac{\mathrm{d}\varphi_{\underline{\boldsymbol{\mu}}}}{\mathrm{d}x_{rs}}(\underline{\boldsymbol{x}}) = \mu_{rs} \log\left(\frac{x_{rs}(1-\alpha)}{\alpha(1-x_{rs})}\right)$, and $\frac{\mathrm{d}^2\varphi_{\underline{\boldsymbol{\mu}}}}{\mathrm{d}x_{rs}^2}(\underline{\boldsymbol{x}}) = \frac{\mu_{rs}}{x_{rs}(1-x_{rs})} \ge 0$. Therefore, by convexity,

$$\varphi_{\underline{\boldsymbol{\mu}}}(\underline{\boldsymbol{x}}) \ge \varphi_{\underline{\boldsymbol{\mu}}}(\underline{\boldsymbol{x}}^*) + \frac{1}{2} \sum_{(r,s)\in E}^{d} \frac{\mu_{rs}}{x_{rs}^*(1-x_{rs}^*)} (x_{rs} - x_{rs}^*)^2.$$

Let

$$\ell_{\underline{\boldsymbol{\mu}}}(\underline{\boldsymbol{\nu}}) = \sum_{(r,s)\in E} \frac{(\nu_{rs} - \mu_{rs}x^*_{rs})^2}{x^*_{rs}(1-x^*_{rs})\mu_{rs}}. \tag{19}$$

Based on Lemma 10, all the entries of $\underline{\boldsymbol{x}}^*$ will be treated as constants. If $\underline{\boldsymbol{\nu}} \in \Omega$ then $\frac{(\nu_{rs}-\mu_{rs}x^*_{rs})^2}{x^*_{rs}(1-x^*_{rs})\mu_{rs}} \le \ell_{\underline{\boldsymbol{\mu}}}(\underline{\boldsymbol{\nu}}) \le C(\underline{\boldsymbol{\mu}})^2$, and

$$\nu_{rs} \in \left[\mu_{rs}x^*_{rs} \pm C(\underline{\boldsymbol{\mu}})\sqrt{x^*_{rs}(1-x^*_{rs})\mu_{rs}}\right].$$

Now let us assume that $C(\underline{\boldsymbol{\mu}}) = o(\mu_{rs}^{1/2})$ for all $(r,s) \in E$. Then we have

$$\frac{\nu_{rs}}{\mu_{rs}}(1-\frac{\nu_{rs}}{\mu_{rs}}) \ge x^*_{rs}(1-x^*_{rs}) - o_n(1).$$

If $\underline{\boldsymbol{\nu}} \notin \Omega$ then $\ell(\underline{\boldsymbol{\nu}}) \ge C(\underline{\boldsymbol{\mu}})^2$, therefore A is bounded by

$$\left(\prod_{(r,s)\in E}(2\pi\mu_{rs}x^*_{rs}(1-x^*_{rs}))^{-1/2}\sum_{\underline{\boldsymbol{\nu}}\in\Omega} e^{-\ell_{\underline{\boldsymbol{\mu}}}(\underline{\boldsymbol{\nu}})/2} + \prod_{(r,s)\in E} 3\mu_{rs}^{1/2} \sum_{\substack{\underline{\boldsymbol{\nu}}\notin\Omega \\ 0\le\nu_{rs}\le\mu_{rs}}} e^{-C(\underline{\boldsymbol{\mu}})^2/2}\right)\cdot\exp(-\varphi_{\underline{\boldsymbol{\mu}}}(\underline{\boldsymbol{x}}^*)). \tag{20}$$

The second term in the sum above is bounded by $c^{d^2}(\prod_{(r,s)\in E}\mu_{rs}^{3/2})e^{-C(\underline{\boldsymbol{\mu}})^2/2}$ for some constant $c > 0$. Taking $C(\underline{\boldsymbol{\mu}})^2 = 5\log\prod_{(r,s)\in E}\mu_{rs}$, this term is $c^{d^2}\prod_{(r,s)\in E}\mu_{rs}^{-1} = \mathcal{O}(n^{-|E|})$. Moreover, for all $(r,s) \in E$, $\mu_{rs} > \epsilon n$, therefore

$$C(\underline{\boldsymbol{\mu}})^2 \le 4d^2\log n \ll \mu_{rs};$$

this choice satisfies the condition $C(\underline{\boldsymbol{\mu}}) = o(\mu_{rs}^{1/2})$ for $(r,s) \in E$. On the other hand, let

$$S(\underline{\boldsymbol{\mu}}) = \sum_{\underline{\boldsymbol{\nu}}\in\mathbb{Z}^E} \mathbb{1}\{\boldsymbol{M}_G(\underline{\boldsymbol{\nu}},\underline{\boldsymbol{\nu}}') \in \mathcal{F}\}\exp(-\ell_{\underline{\boldsymbol{\mu}}}(\underline{\boldsymbol{\nu}})/2), \tag{21}$$

with $\ell_{\underline{\boldsymbol{\mu}}}$ defined in (19). We ignored the dependence of S on $\underline{\boldsymbol{\nu}}'$ in the notation on purpose: this dependence is inessential. The first sum in (20) is upper bounded by $S(\underline{\boldsymbol{\mu}})$. Therefore

$$A(\underline{\boldsymbol{\nu}}',\underline{\boldsymbol{\mu}}) \le c_u \prod_{(r,s)\in E}\mu_{rs}^{-1/2}\ (S(\underline{\boldsymbol{\mu}}) + \varepsilon_n)\ e^{-\vartheta(\underline{\boldsymbol{\nu}}',\underline{\boldsymbol{\mu}})}, \tag{22}$$

for some c_u depending on ϵ and d, and $\epsilon_n \to 0$ as $n \to \infty$. We now turn our attention to the lower bound, deferring the analysis of the Gaussian sum $S(\underline{\boldsymbol{\mu}})$ to a subsequent paragraph.

The lower bound We have

$$A(\underline{\boldsymbol{\nu}}',\underline{\boldsymbol{\mu}}) \ge \sum_{\underline{\boldsymbol{\nu}}\in\Omega}\left(\prod_{(r,s)\in E}\binom{\mu_{rs}}{\nu_{rs}}\alpha^{\nu_{rs}}(1-\alpha)^{\mu_{rs}-\nu_{rs}}\right).$$

Using $\binom{\mu_{rs}}{\nu_{rs}} \ge (8\pi\nu_{rs}(1-\nu_{rs}/\mu_{rs}))^{-1/2}e^{H(\nu_{rs}/\mu_{rs})}$ for all $(r,s) \in E$, we get

$$A(\underline{\boldsymbol{\nu}}',\underline{\boldsymbol{\mu}}) \ge \sum_{\underline{\boldsymbol{\nu}}\in\Omega}\left(\prod_{(r,s)\in E} 8\pi\nu_{rs}(1-\frac{\nu_{rs}}{\mu_{rs}})\right)^{-1/2}\cdot\exp\left(-\varphi_{\underline{\boldsymbol{\mu}}}\left(\underline{\boldsymbol{\nu}}/\underline{\boldsymbol{\mu}}\right)\right).$$

For $\underline{\boldsymbol{\nu}} \in \Omega$, we have $\frac{\nu_{rs}}{\mu_{rs}}(1-\frac{\nu_{rs}}{\mu_{rs}}) \le x^*_{rs}(1-x^*_{rs})+o_n(1)$ for all r,s, and since $\varphi_{\underline{\boldsymbol{\mu}}}$ is a smooth function, a Taylor expansion yields

$$\varphi_{\underline{\boldsymbol{\mu}}}(\underline{\boldsymbol{\nu}}/\underline{\boldsymbol{\mu}}) = \varphi_{\underline{\boldsymbol{\mu}}}(\underline{\boldsymbol{x}}^*) + \frac{1}{2}\sum_{(r,s)\in E}\frac{(\nu_{rs}-x^*_{rs}\mu_{rs})^2}{x^*_{rs}(1-x^*_{rs})\mu_{rs}} + o_n(1).$$

Therefore,

$$\begin{aligned} A(\underline{\boldsymbol{\nu}}', \underline{\boldsymbol{\mu}}) &\ge c_l e^{-\varphi_{\underline{\boldsymbol{\mu}}}(\underline{\boldsymbol{x}}^*)} \prod_{(r,s)\in E}\mu_{r,s}^{-1/2}\cdot\sum_{\underline{\boldsymbol{\nu}}\in\Omega}\exp\left(-\frac{1}{2}\sum_{rs}\frac{(\nu_{rs}-x^*_{rs}\mu_{rs})^2}{x^*_{rs}(1-x^*_{rs})\mu_{rs}}\right) \\ &= c_l e^{-\vartheta(\underline{\boldsymbol{\nu}}',\underline{\boldsymbol{\mu}})}\prod_{(r,s)\in E}\mu_{r,s}^{-1/2}\cdot\left(S(\underline{\boldsymbol{\mu}})-\varepsilon_n\right), \end{aligned}$$

where $c_l = c_l(\epsilon, d)$, $S(\underline{\boldsymbol{\mu}})$ is defined in (21) and

$$\varepsilon_n := \prod_{(r,s)\in E}(\mu_{rs}+1)e^{-C(\underline{\boldsymbol{\mu}})^2/2} + \sum_{\substack{\underline{\boldsymbol{\nu}}\in\mathbb{Z}^E_+ \\ \nu_{rs}\ge\mu_{rs}+1}}\exp -\ell_{\underline{\boldsymbol{\mu}}}(\underline{\boldsymbol{\nu}})/2 + \sum_{\underline{\boldsymbol{\nu}}\in\mathbb{Z}^E_-}\exp -\ell_{\underline{\boldsymbol{\mu}}}(\underline{\boldsymbol{\nu}})/2.$$

We take $C(\underline{\boldsymbol{\mu}})^2 = 4\log\prod_{(r,s)\in E}\mu_{rs}$. This makes the first term in ε_n bounded by $\prod_{r,s}\mu_{rs}^{-1} = \mathcal{O}(n^{-|E|})$. On the other hand, the remaining tail sums are easily bounded by the tail probability function of a normal random variable (i.e., the error function):

$$\begin{aligned} \sum_{\substack{\underline{\boldsymbol{\nu}}\in\mathbb{Z}^E_+ \\ \nu_{rs}\ge\mu_{rs}+1}}\exp -\ell_{\underline{\boldsymbol{\mu}}}(\underline{\boldsymbol{\nu}})/2 &\le \prod_{(r,s)\in E}\mu_{rs}^{1/2}\,\mathsf{erfc}\left(\sqrt{\frac{x^*_{rs}}{1-x^*_{rs}}}\mu_{rs}\right), \\ \sum_{\underline{\boldsymbol{\nu}}\in\mathbb{Z}^E_-}\exp -\ell_{\underline{\boldsymbol{\mu}}}(\underline{\boldsymbol{\nu}})/2 &\le \prod_{(r,s)\in E}\mu_{rs}^{1/2}\,\mathsf{erfc}\left(\sqrt{\frac{1-x^*_{rs}}{x^*_{rs}}}\mu_{rs} - \frac{1}{\sqrt{x^*_{rs}(1-x^*_{rs})\mu_{rs}}}\right), \end{aligned}$$

with $\mathsf{erfc}(x) = \int_x^\infty e^{-t^2/2}dt$. Since $\mathsf{erfc}(x) \le e^{-x^2/2}/x$ for all $x > 0$, these two terms decay in a sub-Gaussian way in n.

Bounding the Gaussian sum. Here we approximate S by a continuous Gaussian integral. We prove that

$$S(\underline{\boldsymbol{\mu}}) \asymp \int_{\mathcal{F}(G)}\exp\left(-\sum_{(r,s)\in E}\frac{\mu_{rs}^{-1}}{2x^*_{rs}(1-x^*_{rs})}z_{rs}^2\right)\mathrm{d}\underline{\boldsymbol{z}},$$

where the symbol " $\asymp$ " hides constants depending on G, ϵ, d and α as $n\to\infty$. For $\underline{\boldsymbol{\nu}} \in \mathcal{F}(G)$ an array of integer numbers such that $0 \le \nu_{rs} \le \mu_{rs}$, let $T(\underline{\boldsymbol{\nu}}) = \underline{\boldsymbol{\nu}} + \mathcal{C}\cap\mathcal{F}(G)$ where $\mathcal{C} = [-1/2, 1/2]^E$. The sum is understood in the Minkowski sense. $T(\underline{\boldsymbol{\nu}})$ is a "tile" of side 1 centered around $\underline{\boldsymbol{\nu}}$. Two crucial facts are (i) : $T(\underline{\boldsymbol{\nu}})$ and $T(\underline{\boldsymbol{\nu}}')$ are of disjoint interiors when $\underline{\boldsymbol{\nu}} \ne \underline{\boldsymbol{\nu}}'$ and (ii) : $T(\underline{\boldsymbol{\nu}}) \subset \mathcal{F}(G)$. Now for a fixed $\underline{\boldsymbol{\nu}}$, let $\underline{\boldsymbol{z}} \in T(\underline{\boldsymbol{\nu}})$. For $r,s\in E$, we have $\nu_{rs}-1/2 \le z_{rs} \le \nu_{rs}+1/2$ and $\frac{\nu_{rs}-1/2}{\mu_{rs}} - x^*_{rs} \le \frac{z_{rs}}{\mu_{rs}} - x^*_{rs} \le \frac{\nu_{rs}+1/2}{\mu_{rs}} - x^*_{rs}$. Thus

$$\left(\frac{z_{rs}}{\mu_{rs}} - x^*_{rs}\right)^2 \le \max\left\{\left(\frac{\nu_{rs}-1/2}{\mu_{rs}} - x^*_{rs}\right)^2, \left(\frac{\nu_{rs}+1/2}{\mu_{rs}} - x^*_{rs}\right)^2\right\}.$$

Using the fact $\max\{(x-1/2)^2,(x+1/2)^2\} \le 2x^2+1$ for all $x \in \mathbb{R}$, we get

$$\exp\left(-\sum_{(r,s)\in E}\frac{\mu_{rs}^{-1}}{4x_{rs}^*(1-x_{rs}^*)}\left(2\left(\nu_{rs}-\mu_{rs}x_{rs}^*\right)^2+1\right)\right) \le \exp\left(-\sum_{(r,s)\in E}\frac{\mu_{rs}^{-1}}{4x_{rs}^*(1-x_{rs}^*)}\left(z_{rs}-\mu_{rs}x_{rs}^*\right)^2\right).$$

By integrating both sides of the above inequality on $T(\underline{\boldsymbol{\nu}})$ in the variable $\underline{\boldsymbol{z}}$, and summing over all $\underline{\boldsymbol{\nu}}$ with integer entries such that $\boldsymbol{M}_G(\underline{\boldsymbol{\nu}},\underline{\boldsymbol{\nu}}') \in \mathcal{F}$, we get

$$\begin{aligned}\mathrm{vol}(\mathcal{C}\cap\mathcal{F}(G))e^{-\sum_{(r,s)\in E}\frac{\mu_{rs}^{-1}}{4x_{rs}^*(1-x_{rs}^*)}}\ S(\underline{\boldsymbol{\mu}}) &\le \sum_{\substack{\underline{\boldsymbol{\nu}}\in\mathbb{Z}^E\\ \boldsymbol{M}_G(\underline{\boldsymbol{\nu}},\underline{\boldsymbol{\nu}}')\in\mathcal{F}}}\int_{T(\underline{\boldsymbol{\nu}})}\exp\left(-\sum_{(r,s)\in E}\frac{\mu_{rs}^{-1}\left(z_{rs}-\mu_{rs}x_{rs}^*\right)^2}{4x_{rs}^*(1-x_{rs}^*)}\right)\mathrm{d}\underline{\boldsymbol{z}},\\ &= \sum_{\substack{\underline{\boldsymbol{\nu}}\in\mathbb{Z}^E\\ \boldsymbol{M}_G(\underline{\boldsymbol{\nu}},\underline{\boldsymbol{\nu}}')\in\mathcal{F}}}\int_{T(\underline{\boldsymbol{\nu}}-\underline{\boldsymbol{x}}^*\odot\underline{\boldsymbol{\mu}})}\exp\left(-\sum_{(r,s)\in E}\frac{\mu_{rs}^{-1}}{4x_{rs}^*(1-x_{rs}^*)}z_{rs}^2\right)\mathrm{d}\underline{\boldsymbol{z}},\end{aligned}$$

where vol is the volume according to the $\dim(\mathcal{F}(G))$-dimensional Lebesgue measure. Since $\boldsymbol{M}_G(\underline{\boldsymbol{x}}^*\odot\underline{\boldsymbol{\mu}},\underline{\boldsymbol{\nu}}') \in \mathcal{F}$, we have $\underline{\boldsymbol{\nu}}-\underline{\boldsymbol{x}}^*\odot\underline{\boldsymbol{\mu}} \in \mathcal{F}$ for all $\underline{\boldsymbol{\nu}}$ we are summing over. Moreover, since the tiles $T(\underline{\boldsymbol{\nu}})$ are of mutually disjoint interiors, and given that their union is in $\mathcal{F}(G)$, the left-hand side is upper bounded by (there is actually equality)

$$\int_{\mathcal{F}(G)}\exp\left(-\sum_{(r,s)\in E}\frac{\mu_{rs}^{-1}}{4x_{rs}^*(1-x_{rs}^*)}z_{rs}^2\right)\mathrm{d}\underline{\boldsymbol{z}}.$$

Here, to get sharper constants, one could apply a theorem by Vaaler [Vaa79] which states that the volume of any linear subspace intersected with the cube $\mathcal{C}$ is at least 1; i.e., $\mathrm{vol}(\mathcal{C}\cap\mathcal{F}(G)) \ge 1$. This yields

$$S(\underline{\boldsymbol{\mu}}) \le e^{\sum_{(r,s)\in E}\frac{\mu_{rs}^{-1}}{4x_{rs}^*(1-x_{rs}^*)}}\cdot\int_{\mathcal{F}(G)}\exp\left(-\sum_{(r,s)\in E}\frac{\mu_{rs}^{-1}}{4x_{rs}^*(1-x_{rs}^*)}z_{rs}^2\right)\mathrm{d}\underline{\boldsymbol{z}}.$$

As for the reverse inequality, slightly more care is needed in constructing the approximation. For a given $\underline{\boldsymbol{\nu}}$, let $\Omega^+ = \{(r,s) \ : \ \nu_{rs} \ge x_{rs}^*\mu_{rs}+1/2\}$ and $\Omega^- = \{(r,s) \ : \ \nu_{rs} \le x_{rs}^*\mu_{rs}-1/2\}$. For $\underline{\boldsymbol{z}} \in T(\underline{\boldsymbol{\nu}})$, we have $(z_{rs}-x_{rs}^*\mu_{rs})^2 \ge (\nu_{rs}-x_{rs}^*\mu_{rs}-1/2)^2$ if $(r,s)\in\Omega^+$ and $(z_{rs}-x_{rs}^*\mu_{rs})^2 \ge (\nu_{rs}-x_{rs}^*\mu_{rs}+1/2)^2$ if $(r,s)\in\Omega^-$. Otherwise, for $(r,s)\notin\Omega^+\cup\Omega^-$, we have $|\nu_{rs}-x_{rs}^*\mu_{rs}|<1/2$ and $\left|(z_{rs}-x_{rs}^*\mu_{rs})^2-(\nu_{rs}-x_{rs}^*\mu_{rs})^2\right| < 1/2(1+1/2) = 3/4$. Therefore

$$\begin{aligned}&\sum_{(r,s)\in\Omega^+}\frac{(\nu_{rs}-x_{rs}^*\mu_{rs}+1/2)^2}{\mu_{rs}x_{rs}^*(1-x_{rs}^*)}+\sum_{(r,s)\in\Omega^-}\frac{(\nu_{rs}-x_{rs}^*\mu_{rs}-1/2)^2}{\mu_{rs}x_{rs}^*(1-x_{rs}^*)}+\sum_{(r,s)\notin\Omega^+\cup\Omega^-}\frac{(\nu_{rs}-x_{rs}^*\mu_{rs})^2}{\mu_{rs}x_{rs}^*(1-x_{rs}^*)}\\ &\le \sum_{(r,s)\in E}\frac{(z_{rs}-x_{rs}^*\mu_{rs})^2}{\mu_{rs}x_{rs}^*(1-x_{rs}^*)}+\sum_{(r,s)\in E}\frac{3\mu_{rs}^{-1}}{4x_{rs}^*(1-x_{rs}^*)}.\end{aligned}$$

On the other hand, $(\nu_{rs}-x_{rs}^*\mu_{rs})^2 \le (\nu_{rs}-x_{rs}^*\mu_{rs}+1/2)^2$ when $(r,s)\in\Omega^+$ and $(\nu_{rs}-x_{rs}^*\mu_{rs})^2 \le (\nu_{rs}-x_{rs}^*\mu_{rs}-1/2)^2$ when $(r,s)\in\Omega^-$. After integrating on $T(\underline{\boldsymbol{\nu}})$ and summing over all $\underline{\boldsymbol{\nu}} \in \mathbb{Z}^E$ such that $\underline{\boldsymbol{\nu}}-\underline{\boldsymbol{x}}^*\odot\underline{\boldsymbol{\mu}} \in \mathcal{F}$, we obtain:

$$\mathrm{vol}(\mathcal{C}\cap\mathcal{F}(G))S(\underline{\boldsymbol{\mu}}) \ge e^{-\sum_{(r,s)\in E}\frac{3\mu_{rs}^{-1}}{8x_{rs}^*(1-x_{rs}^*)}}\cdot\sum_{\substack{\underline{\boldsymbol{\nu}}\in\mathbb{Z}^E\\ \underline{\boldsymbol{\nu}}-\underline{\boldsymbol{x}}^*\odot\underline{\boldsymbol{\mu}}\in\mathcal{F}}}\int_{T(\underline{\boldsymbol{\nu}}-\underline{\boldsymbol{x}}^*\odot\underline{\boldsymbol{\mu}})}\exp\left(-\sum_{(r,s)\in E}\frac{\mu_{rs}^{-1}}{2x_{rs}^*(1-x_{rs}^*)}z_{rs}^2\right)\mathrm{d}\underline{\boldsymbol{z}},$$

and the last sum is equal to

$$\sum_{\underline{\boldsymbol{\nu}}\in(\mathbb{Z}^E+\underline{\boldsymbol{x}}^*\odot\underline{\boldsymbol{\mu}})\cap\mathcal{F}(G)}\int_{T(\underline{\boldsymbol{\nu}})}\exp\left(-\sum_{(r,s)\in E}\frac{\mu_{rs}^{-1}}{2x_{rs}^*(1-x_{rs}^*)}z_{rs}^2\right)\mathrm{d}\underline{\boldsymbol{z}}=\int_{\mathcal{F}(G)}\exp\left(-\sum_{(r,s)\in E}\frac{\mu_{rs}^{-1}}{2x_{rs}^*(1-x_{rs}^*)}z_{rs}^2\right)\mathrm{d}\underline{\boldsymbol{z}}.$$

Finally,

$$S(\underline{\boldsymbol{\mu}})\geq c(G,d)e^{-\sum_{(r,s)\in E}\frac{3\mu_{rs}^{-1}}{2x_{rs}^*(1-x_{rs}^*)}}\cdot\int_{\mathcal{F}(G)}\exp\left(-\sum_{(r,s)\in E}\frac{\mu_{rs}^{-1}}{2x_{rs}^*(1-x_{rs}^*)}z_{rs}^2\right)\mathrm{d}\underline{\boldsymbol{z}}.$$

■

Proof of Lemma 10. Recall that $\underline{\boldsymbol{x}}^*$ is the unique minimizer of the function

$$\varphi_{\underline{\boldsymbol{\mu}}}=\sum_{(r,s)\in E}\mu_{rs}D(x_{rs}\parallel\alpha)$$

on $[0,1]^{d\times d}$ subject to $\boldsymbol{M}_G(\underline{\boldsymbol{x}}^*\odot\underline{\boldsymbol{\mu}},\underline{\boldsymbol{\nu}})\in\mathcal{F}$. Recall also that the entries of $\underline{\boldsymbol{x}}^*$ admit the expressions

$$x_{rs}^*=\frac{\alpha}{\alpha+(1-\alpha)e^{\lambda_r^*-\lambda_s^*}},$$

for all $(r,s)\in E$. The vector $\boldsymbol{\lambda}^*\in\mathbb{R}^d$ is the unique solution up to global shifts to the dual optimization problem (strong duality holds here [BV04, Roc70])

$$\sup_{\boldsymbol{\lambda}\in\mathbb{R}^d}\left\{\sum_{(r,s)\notin E}\nu_{rs}(\lambda_r-\lambda_s)+\sum_{(r,s)\in E}\mu_{rs}\log\left(\frac{e^{\lambda_r-\lambda_s}}{\alpha+(1-\alpha)e^{\lambda_r-\lambda_s}}\right)\right\}. \tag{23}$$

Our claim reduces to the boundedness of the differences $|\lambda_r^*-\lambda_s^*|$ for all $(r,s)\in E$ independently of $n,\underline{\boldsymbol{\mu}},\underline{\boldsymbol{\nu}}$ and r,s. We will shortly prove the following inequality

$$\sum_{(r,s)\in E}\mu_{rs}(\lambda_r^*-\lambda_s^*)^2\leq\kappa(\alpha)\sum_{(r,s)\in E}\mu_{rs}, \tag{24}$$

where $\kappa(\alpha)=\frac{1}{\alpha^2}+\frac{1}{(1-\alpha)^2}$. Assuming the above is true, by the Cauchy-Schwarz inequality, we would have

$$\sum_{(r,s)\in E}|\lambda_r^*-\lambda_s^*|\leq\left(\sum_{(r,s)\in E}\mu_{rs}^{-1}\right)^{1/2}\left(\kappa(\alpha)\sum_{(r,s)\in E}\mu_{rs}\right)^{1/2}\leq d^2(\kappa(\alpha)/\epsilon)^{1/2},$$

since $\epsilon n\leq\mu_{rs}\leq n$ for all $(r,s)\in E$. We would then be done. Now, the inequality (24) follows from convexity considerations. We let ϕ be the function being maximized in (23). By concavity of ϕ, we have

$$\phi(\boldsymbol{\lambda}^*)-\phi(\mathbf{0})\leq\boldsymbol{\lambda}^{*\intercal}\nabla\phi(\mathbf{0})+\frac{1}{2}\boldsymbol{\lambda}^{*\intercal}\nabla^2\phi(\mathbf{0})\boldsymbol{\lambda}^*. \tag{25}$$

The gradient and the Hessian of ϕ are

$$[\nabla\phi(\boldsymbol{\lambda})]_r=\sum_{s:(r,s)\in E}\frac{\alpha\mu_{rs}}{\alpha+(1-\alpha)e^{\lambda_r-\lambda_s}}-\frac{\alpha\mu_{sr}}{\alpha+(1-\alpha)e^{\lambda_s-\lambda_r}}+\sum_{s:(r,s)\notin E}\nu_{rs}-\nu_{sr},\quad r\in\{1,\cdots,d\},$$

$$\nabla^2\phi(\boldsymbol{\lambda}) = -\alpha(1-\alpha)\sum_{(r,s)\in E} w_{rs}(\boldsymbol{\lambda})(\boldsymbol{e}_r-\boldsymbol{e}_s)(\boldsymbol{e}_r-\boldsymbol{e}_s)^\intercal,$$

with

$$w_{rs}(\boldsymbol{\lambda}) = \frac{\mu_{rs}e^{\lambda_r-\lambda_s}}{(\alpha+(1-\alpha)e^{\lambda_r-\lambda_s})^2} + \frac{\mu_{sr}e^{\lambda_s-\lambda_r}}{(\alpha+(1-\alpha)e^{\lambda_s-\lambda_r})^2},$$

and $\boldsymbol{e}_1,\cdots,\boldsymbol{e}_d$ being the standard unit vectors in $\mathbb{R}^d$. The concavity inequality (25) becomes

$$\phi(\boldsymbol{\lambda}^*) \le \alpha\sum_{(r,s)\in E}\mu_{rs}(\lambda_r^*-\lambda_s^*) + \sum_{(r,s)\notin E}\nu_{rs}(\lambda_r^*-\lambda_s^*) - \alpha(1-\alpha)\sum_{(r,s)\in E}\mu_{rs}(\lambda_r^*-\lambda_s^*)^2.$$

Substituting in the expression of $\phi(\boldsymbol{\lambda}^*)$, the term $\sum_{(r,s)\notin E}\nu_{rs}(\lambda_r^*-\lambda_s^*)$ cancels out on both sides and we get

$$\sum_{(r,s)\in E}\mu_{rs}\log\left(\frac{e^{\lambda_r^*-\lambda_s^*}}{\alpha+(1-\alpha)e^{\lambda_r^*-\lambda_s^*}}\right) \le \alpha\sum_{(r,s)\in E}\mu_{rs}(\lambda_r^*-\lambda_s^*) - \alpha(1-\alpha)\sum_{(r,s)\in E}\mu_{rs}(\lambda_r^*-\lambda_s^*)^2,$$

which can be written as

$$\sum_{(r,s)\in E}\mu_{rs}\left(\alpha(1-\alpha)(\lambda_r^*-\lambda_s^*)^2 + (1-\alpha)(\lambda_r^*-\lambda_s^*) - \log\left(\alpha+(1-\alpha)e^{\lambda_r^*-\lambda_s^*}\right)\right) \le 0. \tag{26}$$

Now we approximate the logarithm by the positive part: $\log(\alpha+(1-\alpha)e^x) \le x_+ = \max\{0,x\}$ for all $x\in\mathbb{R}$ and $\alpha\in(0,1)$, so that we almost get a quadratic polynomial inequality. We make this a genuine quadratic inequality by applying the additional approximation that for all $x\in\mathbb{R}$ and $\alpha\in(0,1)$:

$$\alpha(1-\alpha)x^2 + (1-\alpha)x - x_+ \ge \frac{\alpha(1-\alpha)}{2}x^2 - \frac{1-\alpha}{2\alpha} - \frac{\alpha}{2(1-\alpha)}.$$

This is easy to check by verifying that the discriminants of the resulting quadratics (one for $x\ge 0$ and one for $x<0$) are negative. Now, inequality (26) implies

$$\frac{\alpha(1-\alpha)}{2}\sum_{(r,s)\in E}\mu_{rs}(\lambda_r^*-\lambda_s^*)^2 \le \left(\frac{1-\alpha}{2\alpha}+\frac{\alpha}{2(1-\alpha)}\right)\sum_{(r,s)\in E}\mu_{rs}.$$

In other words,

$$\sum_{(r,s)\in E}\mu_{rs}(\lambda_r^*-\lambda_s^*)^2 \le \kappa(\alpha)\sum_{(r,s)\in E}\mu_{rs}.$$

■

6 Two proofs of Proposition 6

We first reduce the proof to the case where $G=K_d$ by a limiting argument. Let $G=(V,E)$ be a graph on d vertices. If G is not connected then the constraints defining the space $\mathcal{F}(G)$ decouple across the connected components of G and so does the integrand $\exp -\frac{1}{2}\sum_{(r,s)\in E}x_{rs}^2/w_{rs}$, therefore the Gaussian integral factors across the connected components of G. Hence, we may assume that G is connected. Now, if

$$\int_{\mathcal{F}} e^{-\frac{1}{2}\sum_{rs}x_{rs}^2/w_{rs}}\,\mathrm{d}\underline{\boldsymbol{x}} = (2\pi)^{((d-1)^2+d)/2}\left(\frac{\prod_{r,s}w_{rs}}{T(\underline{\boldsymbol{w}})}\right)^{1/2},$$

for all $\underline{\boldsymbol{w}} \in \mathbb{R}_+^{d\times d}$ where $T = T_{K_d}$, then taking a limit $w_{rs} \to 0$ for all $(r,s) \notin E$, we get

$$\frac{1}{\left(\prod_{(r,s)\notin E} w_{rs}\right)^{1/2}} \int_{\mathcal{F}} e^{-\frac{1}{2}\sum_{rs} x_{rs}^2/w_{rs}} \, \mathrm{d}\underline{\boldsymbol{x}} \longrightarrow c(G) \int_{\mathcal{F}(G)} e^{-\frac{1}{2}\sum_{(r,s)\in E} x_{rs}^2/w_{rs}} \, \mathrm{d}\underline{\boldsymbol{x}},$$

where $c(G) > 0$ is a constant that only depends on G. On the other hand

$$T(\underline{\boldsymbol{w}}) \longrightarrow \frac{\mathsf{nst}(G)}{2^{d-1}d^{d-2}} \, T_G(\underline{\boldsymbol{w}}).$$

Therefore

$$c(G) \int_{\mathcal{F}(G)} e^{-\frac{1}{2}\sum_{(r,s)\in E} x_{rs}^2/w_{rs}} \, \mathrm{d}\underline{\boldsymbol{x}} = (2\pi)^{((d-1)^2+d)/2} \left(\frac{2^{d-1}d^{d-2}}{\mathsf{nst}(G)} \frac{\prod_{(r,s)\in E} w_{rs}}{T_G(\underline{\boldsymbol{w}})}\right)^{1/2}.$$

Now we set $w_{rs} = 1$ for all $(r,s) \in E$ to clear out the constants. Since $\int_{\mathcal{F}(G)} e^{-\frac{1}{2}\sum_{(r,s)\in E} x_{rs}^2} \, \mathrm{d}\underline{\boldsymbol{x}} = (2\pi)^{\dim(\mathcal{F}(G))/2}$, we get

$$\int_{\mathcal{F}(G)} e^{-\frac{1}{2}\sum_{(r,s)\in E} x_{rs}^2/w_{rs}} \, \mathrm{d}\underline{\boldsymbol{x}} = (2\pi)^{\dim(\mathcal{F}(G))/2} \left(\frac{\prod_{(r,s)\in E} w_{rs}}{T_G(\underline{\boldsymbol{w}})}\right)^{1/2}.$$

Now it remains to prove the proposition for the complete graph.

6.1 A combinatorial proof

We proceed by adopting a combinatorial view on the structure of the space $\mathcal{F}$. This will lead us to consider a very special basis of $\mathcal{F}$ in which the computations become tractable. (Background on the concepts used in this construction can be found in [Big97].) We first orient K_d in such a way that every pair of distinct vertices is connected by two parallel edges pointing in opposite directions. There are $d(d-1)$ (oriented) edges in total. Then, the subgraphs whose edges are weighted by an array $\underline{\boldsymbol{x}} \in \mathcal{F}$ are called *Eulerian*: the total weight of the incoming edges is equal to that of the outgoing edges on each vertex. An important property of Eulerian graphs is that they can be decomposed into a superposition of cycles. In particular, fix a spanning tree T^* of K_d (the tree uses only one edge, if any, between each pair of vertices, and ignores their orientation). Every edge $e \notin T^*$ can be identified with the oriented cycle C_e in the graph which consists of the oriented edge e and the unique path between the endpoints of e in the tree T^* (where the direction of the edges on the path are flipped if necessary). Let $\boldsymbol{\chi}_e \in \{0, \pm 1\}^{d(d-1)}$ be the indicator vector of the cycle C_e[2]. Since a cycle is Eulerian, the vector $\boldsymbol{\chi}_e$—when folded into a $d \times d$ matrix—belongs to $\mathcal{F}$. Furthermore, the family $\{\boldsymbol{\chi}_e \ : \ e \notin T^*\}$ is linearly independent since a cycle C_e contains at least one edge—namely e—that is not contained in any other cycle $C_{e'}$, $e' \neq e$. There are exactly $d(d-1) - (d-1) = (d-1)^2$ off-tree edges in K_d, and this number coincides with the dimension of $\mathcal{F}$. Therefore $\mathcal{B} = \{\boldsymbol{\chi}_e \ : \ e \notin T^*\}$ is a basis of $\mathcal{F}$, that we henceforth call a *fundamental cycle basis* of $\mathcal{F}$.

Let $\boldsymbol{P} \in \{0, \pm 1\}^{(d-1)^2 \times d(d-1)}$ be the matrix where the rows are indexed by the off-tree edges of the graph, and whose eth row is equal to $\boldsymbol{\chi}_e$. The matrix $\boldsymbol{P}$ can be regarded as the *cycle-edge incidence matrix* of the graph K_d: an entry (e, e') is non-zero if and only if $e' \in C_e$.

[2]Each non-zero entry in the vector corresponds to an edge present in the cycle, and the non-zero value is $+1$ if the cycle flows along the orientation of that edge, and -1 if the flow is in the opposite direction. In particular, the eth coordinate of $\boldsymbol{\chi}_e$ is always $+1$.

Let $\boldsymbol{M} \in \mathbb{R}^{d(d-1)\times d(d-1)}$ be the diagonal matrix with entries w_{rs}, $r \neq s$ on the diagonal. Then by a change of variables

$$\begin{aligned}\int_{\mathcal{F}} e^{-\sum_{rs} x_{rs}^2/2w_{rs}} \mathrm{d}\underline{\boldsymbol{x}} &= Det(\boldsymbol{P}\boldsymbol{P}^\intercal)^{1/2} \int_{\mathbb{R}^{(d-1)^2}} e^{-\boldsymbol{z}^\intercal(\boldsymbol{P}\boldsymbol{M}^{-1}\boldsymbol{P}^\intercal)\boldsymbol{z}/2} \mathrm{d}\boldsymbol{z} \\ &= (2\pi)^{(d-1)^2/2}\ Det(\boldsymbol{P}\boldsymbol{P}^\intercal)^{1/2} Det(\boldsymbol{P}\boldsymbol{M}^{-1}\boldsymbol{P}^\intercal)^{-1/2}.\end{aligned}$$

Now it remains to show that $Det(\boldsymbol{P}\boldsymbol{M}^{-1}\boldsymbol{P}^\intercal) = \sum_T \prod_{(r,s)\notin T} w_{rs}^{-1}$ where the sum is over all spanning trees of K_d. This will finish the proof since we would then have $Det(\boldsymbol{P}\boldsymbol{P}^\intercal) = \mathsf{nst}(K_d) = 2^{d-1}d^{d-2}$ by Cayley's formula on the number of spanning trees in the complete graph.

We expand the determinant using the Cauchy-Binet formula. Let $\boldsymbol{D} = \boldsymbol{M}^{-1/2}$, and let E be the set of edges in K_d. For a matrix $\boldsymbol{A}$ of size $n \times m$, $I \subseteq \{1, \cdots, n\}, J \subseteq \{1, \cdots, m\}$, we denote by $\boldsymbol{A}[I, J]$ the matrix of size $|I| \times |J|$ whose rows and columns are indexed by I and J respectively. If $I = \{1, \cdots, n\}$, then we write $\boldsymbol{A}[\ :\ ,\ J]$, and likewise for the column indices. Then, we have

$$Det(\boldsymbol{P}\boldsymbol{M}^{-1}\boldsymbol{P}^\intercal) = \sum_{\substack{S\subseteq E \\ |S|=(d-1)^2}} Det(\boldsymbol{P}\boldsymbol{D}[\ :\ ,\ S\ ])^2. \tag{27}$$

Now we use the following key lemma that we prove later.

Lemma 12. *Assuming the diagonal entries of the (diagonal) matrix $\boldsymbol{D}$ are positive, the matrix $\boldsymbol{P}\boldsymbol{D}[\ :\ ,\ S\ ]$ is singular if and only if the graph spanned by the complement $\overline{S} = E\backslash S$ of S in K_d contains a cycle.*

Since there are exactly $(d-1)$ edges left unchosen by S, this lemma implies that they must form a spanning tree in order for the corresponding term to contribute to the sum in identity (27). Hence

$$Det(\boldsymbol{P}\boldsymbol{M}^{-1}\boldsymbol{P}^\intercal) = \sum_{T\ :\ \text{spanning tree}} Det(\boldsymbol{P}\boldsymbol{D}[\ :\ ,\ \overline{T}\ ])^2.$$

Fix a spanning tree T of K_d. Observe that if $T = T^*$ then the edges that generate the cycles in the fundamental cycle basis $\mathcal{B}$ are exactly the ones that are selected in $\overline{T}$. In other words, each row and each column of $\boldsymbol{P}\boldsymbol{D}[\ :\ ,\ \overline{T}\ ]$ contain exactly one non-zero entry, (i.e., $\boldsymbol{P}[\ :\ ,\ \overline{T}\ ]$ is a permutation matrix), hence $Det(\boldsymbol{P}\boldsymbol{D}[\ :\ ,\ \overline{T}\ ]) = \pm\prod_{(r,s)\notin T} w_{rs}^{-1/2}$. If $T \neq T^*$ then we split the set of edges in $\overline{T}$ into those that belong to T^* and those that do not. Each column in $\boldsymbol{P}\boldsymbol{D}[\ :\ ,\ \overline{T}\ ]$ corresponding to an edge in $\overline{T} \cap \overline{T^*}$ contains only one non-zero entry (since this edge is contained in only one cycle in $\mathcal{B}$). Therefore all such edges (columns) along with the corresponding cycles (rows of the non-zero entry) can be successively eliminated from the determinant, yielding

$$Det(\boldsymbol{P}\boldsymbol{D}[\ :\ ,\ \overline{T}\ ]) = \pm\left(\prod_{(r,s)\in\overline{T}\cap\overline{T^*}} w_{rs}^{-1/2}\right) \cdot Det\left(\boldsymbol{P}\boldsymbol{D}[\ T\cap\overline{T^*}\ ,\ \overline{T}\cap T^*\ ]\right). \tag{28}$$

Notice that this operation has drastically reduced the size of the problem; the common size k of the sets $T \cap \overline{T^*}$ and $\overline{T} \cap T^*$ is anywhere between 0 and $d-1$ at most. Now we will show that

$$Det\left(\boldsymbol{P}\boldsymbol{D}[\ T\cap\overline{T^*}\ ,\ \overline{T}\cap T^*\ ]\right) = \pm\prod_{(r,s)\in\overline{T}\cap T^*} w_{rs}^{-1/2},$$

using a peeling argument slightly more delicate than the one previously applied. Observe that $\boldsymbol{P}\boldsymbol{D}[\ T\cap\overline{T^*}\ ,\ \overline{T}\cap T^*\ ]$ is the $k \times k$ cycle-edge incidence matrix with k edges $T \cap \overline{T^*}$ indexing the

rows and k edges in $\overline{T} \cap T^*$ indexing the columns, such that a row indexed by e indicates the edges $e' \in \overline{T} \cap T^*$ that participate in the cycle C_e.

So far, the spanning tree T^* was arbitrary. To continue, we choose T^* to be the *star tree* rooted at vertex 1 (see Figure 1, left). This choice simplifies the combinatorial argument to come, because the fundamental cycle basis $\mathcal{B}$ is now composed of triangles rooted at vertex 1. Crucially, this is where the assumption $G = K_d$ is needed; to ensure the existence of a star spanning tree. Figure 1 (right) illustrates the remaining edges after the first elimination procedure.

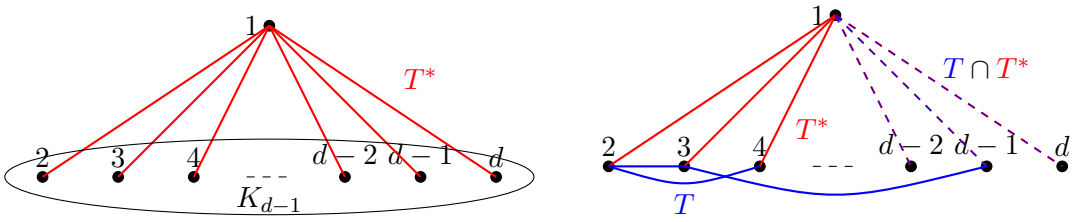

Figure 1. Left: the graph K_d where the star tree T^* is highlighted in red. Right: remaining edges in red and blue after the first elimination procedure (violet edges were removed).

Since T is a tree, by Lemma 9, each row and column of the matrix $\boldsymbol{PD}[\ T \cap \overline{T^*}\ ,\ \overline{T} \cap T^*\ ]$ contains at least one non-zero entry. Furthermore, T^* being the star graph, each cycle $C_e \in \mathcal{B}$ is a triangle rooted at vertex 1, thus each row of the above matrix contains at most two non-zero entries. This is simply because one of the three edges that compose the triangle C_e—namely e—is not selected by the set $\overline{T} \cap T^*$ that indexes the columns of the matrix. See Figure 1, right (any blue edge has at most two adjacent red edges).

Furthermore, if all the rows contain exactly two non-zero entries then by the pigeonhole principle (since $|T \cap \overline{T^*}| = |\overline{T} \cap T^*|$), there will exist three edges in $T \cap \overline{T^*}$ that form a cycle C (see Figure 2, left). However, we assumed that T is a tree so this cannot happen. Therefore there must exist at least one row in the matrix with exactly one non-zero entry (i.e., there must exist an edge $e \in T \cap \overline{T^*}$ such that $C_e = \{e, e_1, e_2\} \in \mathcal{B}$ with $e_1 \in \overline{T} \cap T^*$ and $e_2 \in T \cap T^*$). See Figure 2.

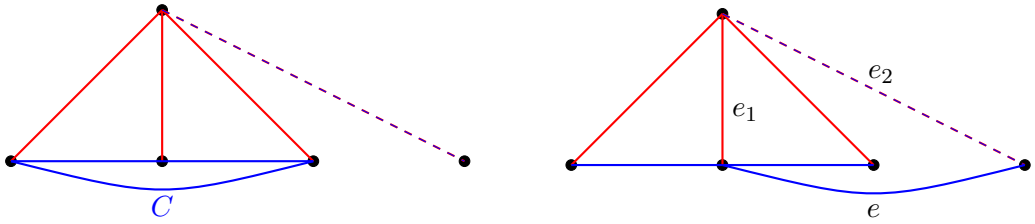

Figure 2. Left: an impossible situation where there remains a cycle C where no edge was eliminated in the first step. Right: a logical situation where there exist a fundamental cycle $C_e = (e, e_1, e_2)$ with one edge in T^* only, one edge in T only, and one edge in their intersection.

Hence, we can eliminate this row and its corresponding column from the determinant. This corresponds to eliminating (dashing) the edges e and e_1 in the right figure above. Applying this argument iteratively allows us to peel all the edges and the cycles they belong to (see Figure 3), so that we obtain

$$Det\left(\boldsymbol{PD}[\ T \cap \overline{T^*}\ ,\ \overline{T} \cap T^*\ ]\right) = \pm \prod_{(r,s) \in \overline{T} \cap T^*} w_{rs}^{-1/2}.$$

This completes the proof.

Proof of Lemma 12. Since we assumed the entries of the diagonal matrix $\boldsymbol{D}$ are strictly positive, we assume without loss of generality that $\boldsymbol{D}$ is the identity matrix. Assume now that

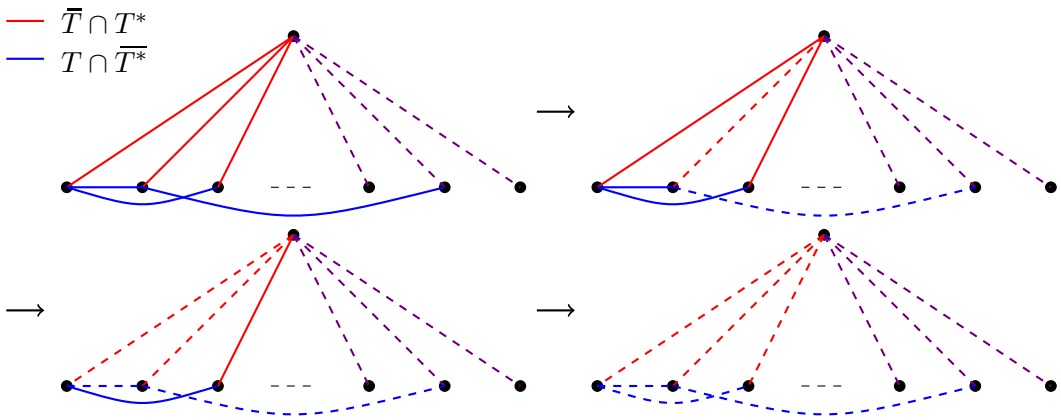

Figure 3. An illustration of the peeling process. "Wedges" with one edge in T only and the other in T^* only are eliminated successively until no edges remain. Violet edges were eliminated in the first step.

the complement of S contains a cycle whose indicator vector is $\boldsymbol{\xi} \in \{0, \pm 1\}^{d(d-1)}$. Since $\mathcal{B}$ is a fundamental cycle basis, there exists $\boldsymbol{x} \in \mathbb{R}^{(d-1)^2} \backslash \{0\}$ such that $\boldsymbol{\xi} = \sum_{e \notin T^*} x_e \boldsymbol{\chi}_e = \boldsymbol{P}^\intercal \boldsymbol{x}$. Since S selects no edges in the cycle indicated by $\boldsymbol{\xi}$, it is clear that $\boldsymbol{x}^\intercal \boldsymbol{P}[\ :\ ,\ S\ ] = 0$, and this settles one direction. As for the other direction, let $\boldsymbol{x} \in \mathbb{R}^{(d-1)^2} \backslash \{0\}$ lie in the null space of $(\boldsymbol{P}[\ :\ ,\ S\ ])^\intercal$. The vector $\boldsymbol{P}^\intercal \boldsymbol{x}$ indicates the weights of a Eulerian subgraph in K_d (this vector belong to $\mathcal{F}$ when written in the form of a $d \times d$ matrix). The condition $(\boldsymbol{P}[\ :\ ,\ S\ ])^\intercal \boldsymbol{x} = 0$ implies that this Eulerian subgraph involves no edges from S. In particular, any cycle from this subgraph (there always exists one) is in the complement of E. This completes the proof. ∎

6.2 An analytic proof

This proof contrasts with the previous purely combinatorial approach in that it is mainly analytic. The approach relies on an interpolation argument that involves expressing the Gaussian integral over $\mathcal{F}$ as the *limit* of another parameterized Gaussian integral, when the parameter tends to zero. This latter integral can on the other hand be written in closed form, by relating it to the characteristic polynomial of a Laplacian matrix. Then the Principal Minors Matrix-Tree Theorem is invoked to finish the argument. This final step is the only place where combinatorics appear. (This proof approach was suggested to us by Andrea Sportiello.) Incidentally, this proof can be carried out with an arbitrary graph G; there is no need to reduce to the complete case. For $\delta > 0$ let

$$I(\delta) = \frac{1}{(2\pi\delta^2)^{(d-1)/2}} \int_{\mathbb{R}^{d \times d}} e^{-\frac{1}{2}\sum_{rs} x_{rs}^2 / w_{rs}}\ e^{-\frac{1}{2\delta^2}\|(\underline{\boldsymbol{x}} - \underline{\boldsymbol{x}}^\intercal)\mathbf{1}\|_{\ell_2}^2}\ \mathrm{d}\underline{\boldsymbol{x}}.$$

The additional Gaussian term in $I(\delta)$ gradually concentrates the mass of the integral on $\mathcal{F}$ as δ becomes small, and we have the following limiting statement:

Lemma 13. *We have*

$$\lim_{\delta \to 0} I(\delta) = c_d \int_{\mathcal{F}} e^{-\frac{1}{2}\sum_{rs} x_{rs}^2 / 2w_{rs}}\ \mathrm{d}\underline{\boldsymbol{x}},$$

with

$$c_d = \frac{1}{(2\pi)^{(d-1)/2}} \int_{\mathcal{F}^\perp} e^{-2\|\underline{\boldsymbol{z}}\mathbf{1}\|_{\ell_2}^2}\ \mathrm{d}\underline{\boldsymbol{z}} = (2d)^{-(d-1)/2}.$$

On the other hand, a straightforward computation allows us to write $I(\delta)$ in closed form:

Lemma 14. *Let $G = (V, E)$ be a weighted graph with $V = \{1, \cdots, d\}$, $E = \{(r, s) \in V \times V, r \neq s\}$ where the edges are weighted by the array $\underline{\boldsymbol{w}} \in \mathbb{R}_+^{d\times d}$. Let $\boldsymbol{L}(\underline{\boldsymbol{w}}) \in \mathbb{R}^{d\times d}$ be the Laplacian matrix of G. For all $\delta > 0$, it holds that*

$$I(\delta) = (2\pi)^{((d-1)^2+d)/2} \left(\prod_{r,s} w_{rs}\right)^{1/2} \frac{\delta}{Det\left(\delta^2 \boldsymbol{I} + \boldsymbol{L}(\underline{\boldsymbol{w}})\right)^{1/2}}.$$

Now, by the Principal Minors Matrix-Tree Theorem (see, e.g., [Cha82]), the characteristic polynomial of the Laplacian matrix of a graph admits the following expansion

$$Det\left(x\boldsymbol{I} + \boldsymbol{L}(\underline{\boldsymbol{w}})\right) = \sum_F x^{|\mathrm{roots}(F)|} \prod_{(r,s)\in F} w_{rs},$$

where the sum is over all rooted spanning forests F of the graph. We finish the argument by taking a limit in δ:

$$\delta^{2(d-1)} Det\left(\boldsymbol{I} + \delta^{-2}\boldsymbol{L}(\underline{\boldsymbol{w}})\right) = \delta^{-2} Det\left(\delta^2\boldsymbol{I} + \boldsymbol{L}(\underline{\boldsymbol{w}})\right) \xrightarrow[\delta\to 0]{} d\sum_T \prod_{(r,s)\in T} w_{rs} = (2d)^{d-1} T(\underline{\boldsymbol{w}}),$$

since the above limit singles out the rooted spanning forests with exactly one root—i.e., rooted spanning trees—from the characteristic polynomial, and there are d ways of choosing the root of a spanning tree. This exactly leads to the desired identity

$$\int_{\mathcal{F}} e^{-\frac{1}{2}\sum_{rs} x_{rs}^2/2w_{rs}}\, \mathrm{d}\underline{\boldsymbol{x}} = (2\pi)^{((d-1)^2+d)/2} \left(\frac{\prod_{r,s} w_{rs}}{T(\underline{\boldsymbol{w}})}\right)^{1/2}.$$

Proof of Lemma 13. We decompose $\mathbb{R}^{d\times d}$ into the direct sum $\mathcal{F} \oplus \mathcal{F}^\perp$. It is easy to see that $\mathcal{F}^\perp = \{\underline{\boldsymbol{z}} = \boldsymbol{\lambda}\mathbf{1}^\intercal - \mathbf{1}\boldsymbol{\lambda}^\intercal,\ \boldsymbol{\lambda} \in \mathbb{R}^d\}$ which is a $(d-1)$-dimensional space. For $\underline{\boldsymbol{x}} \in \mathbb{R}^{d\times d}$, let $\underline{\boldsymbol{y}} \in \mathbb{R}^{d\times d}$ be its orthogonal projection on $\mathcal{F}$, and $\underline{\boldsymbol{z}} = \underline{\boldsymbol{x}} - \underline{\boldsymbol{y}}$. Therefore $(\underline{\boldsymbol{x}} - \underline{\boldsymbol{x}}^\intercal)\mathbf{1} = (\underline{\boldsymbol{z}} - \underline{\boldsymbol{z}}^\intercal)\mathbf{1} = 2\underline{\boldsymbol{z}}\mathbf{1} = 2(d\boldsymbol{\lambda} - (\mathbf{1}^\intercal\boldsymbol{\lambda})\mathbf{1})$. For $\delta > 0$, we have

$$I(\delta) = \frac{1}{(2\pi\delta^2)^{(d-1)/2}} \int_{F\times F^\perp} e^{-\frac{1}{2}\sum_{r,s}(y_{rs}+z_{rs})^2/w_{rs}}\ e^{-\frac{2}{\delta^2}\|\underline{\boldsymbol{z}}\mathbf{1}\|_{\ell_2}^2}\ \mathrm{d}\underline{\boldsymbol{y}}\mathrm{d}\underline{\boldsymbol{z}}.$$

We make the change of variables $\underline{\boldsymbol{z}}' = \underline{\boldsymbol{z}}/\delta$:

$$I(\delta) = \frac{1}{(2\pi)^{(d-1)/2}} \int_{\mathcal{F}\times\mathcal{F}^\perp} e^{-\frac{1}{2}\sum_{r,s}(y_{rs}+\delta z'_{rs})^2/w_{rs}}\ e^{-2\|\underline{\boldsymbol{z}}'\mathbf{1}\|_{\ell_2}^2}\ \mathrm{d}\underline{\boldsymbol{y}}\mathrm{d}\underline{\boldsymbol{z}}'.$$

By dominated convergence,

$$\begin{aligned}\lim_{\delta\to 0} I(\delta) &= \frac{1}{(2\pi)^{(d-1)/2}} \int_{\mathcal{F}\times\mathcal{F}^\perp} e^{-\frac{1}{2}\sum_{r,s} y_{rs}^2/w_{rs}}\ e^{-2\|\underline{\boldsymbol{z}}\mathbf{1}\|_{\ell_2}^2}\ \mathrm{d}\underline{\boldsymbol{y}}\mathrm{d}\underline{\boldsymbol{z}} \\ &= \frac{1}{(2\pi)^{(d-1)/2}} \int_{\mathcal{F}} e^{-\frac{1}{2}\sum_{r,s} y_{rs}^2/w_{rs}}\mathrm{d}\underline{\boldsymbol{y}} \int_{\mathcal{F}^\perp} e^{-2\|\underline{\boldsymbol{z}}\mathbf{1}\|_{\ell_2}^2}\mathrm{d}\underline{\boldsymbol{z}}.\end{aligned}$$

Moreover,

$$\int_{\mathcal{F}^\perp} e^{-2\|\underline{\boldsymbol{z}}\mathbf{1}\|_{\ell_2}^2}\mathrm{d}\underline{\boldsymbol{z}} = (2d)^{(d-1)/2} \int_{\{\boldsymbol{\lambda}\in\mathbb{R}^d, \mathbf{1}^\intercal\boldsymbol{\lambda}=0\}} e^{-2d^2\|\boldsymbol{\lambda}\|_{\ell_2}^2}\ \mathrm{d}\boldsymbol{\lambda} = (2\pi)^{(d-1)/2}(2d)^{-(d-1)/2},$$

where the pre-factor in the first equality comes from the fact that $\|\underline{\boldsymbol{z}}\|_F^2 = 2d\,\|\boldsymbol{\lambda}\|_{\ell_2}^2$ for $\underline{\boldsymbol{z}} = \boldsymbol{\lambda}\mathbf{1}^\intercal - \mathbf{1}\boldsymbol{\lambda}^\intercal,\ \boldsymbol{\lambda} \in \mathbb{R}^d,\ \mathbf{1}^\intercal\boldsymbol{\lambda} = 0$. ∎

Proof of Lemma 14. Let $\delta > 0$. We linearize the quadratic term $\|(\underline{\boldsymbol{x}} - \underline{\boldsymbol{x}}^\intercal)\mathbf{1}\|_{\ell_2}^2$ in $I(\delta)$ by writing the corresponding Gaussian as the Fourier transform of another Gaussian: $\forall \underline{\boldsymbol{x}} \in \mathbb{R}^{d\times d}$,

$$e^{-\frac{1}{2\delta^2}\|(\underline{\boldsymbol{x}}-\underline{\boldsymbol{x}}^\intercal)\mathbf{1}\|_{\ell_2}^2} = \frac{1}{(2\pi)^{d/2}} \int_{\mathbb{R}^d} e^{-\mathbf{i}\delta^{-1}\boldsymbol{y}^\intercal(\underline{\boldsymbol{x}}-\underline{\boldsymbol{x}}^\intercal)\mathbf{1} - \frac{1}{2}\|\boldsymbol{y}\|_{\ell_2}^2}\ \mathrm{d}\boldsymbol{y},$$

where $\mathbf{i}^2 = -1$. Then

$$I(\delta) = \frac{1}{(2\pi\delta^2)^{(d-1)/2}} \frac{1}{(2\pi)^{d/2}} \int_{\mathbb{R}^{d\times d}} \int_{\mathbb{R}^d} e^{-\frac{1}{2}\sum_{rs} x_{rs}^2/w_{rs}}\ e^{-\mathbf{i}\delta^{-1}\boldsymbol{y}^\intercal(\underline{\boldsymbol{x}}-\underline{\boldsymbol{x}}^\intercal)\mathbf{1} - \frac{1}{2}\|\boldsymbol{y}\|_{\ell_2}^2}\ \mathrm{d}\underline{\boldsymbol{x}}\mathrm{d}\boldsymbol{y}.$$

We complete the square involving x_{rs} in the exponentiated expression:

$$-\frac{1}{2}\sum_{r,s} x_{rs}^2/w_{rs} - \mathbf{i}\delta^{-1}\boldsymbol{y}^\intercal(\underline{\boldsymbol{x}} - \underline{\boldsymbol{x}}^\intercal)\mathbf{1} = -\frac{1}{2}\sum_{r,s} \frac{1}{w_{rs}} \left(\left(x_{rs} + \mathbf{i}\frac{w_{rs}}{\delta}(y_r - y_s) \right)^2 + \frac{w_{rs}^2}{\delta^2}(y_r - y_s)^2 \right).$$

Then by Fubini's theorem,

$$I(\delta) = \frac{1}{(2\pi\delta^2)^{(d-1)/2}} \frac{1}{(2\pi)^{d/2}} \int_{\mathbb{R}^d} e^{-\frac{1}{2}\|\boldsymbol{y}\|_{\ell_2}^2 - \frac{1}{2}\sum_{rs}\frac{w_{rs}}{\delta^2}(y_r-y_s)^2} \int_{\mathbb{R}^{d\times d}} e^{-\frac{1}{2}\sum_{rs}\frac{1}{w_{rs}}\left(x_{rs} + \mathbf{i}\frac{w_{rs}}{\delta}(y_r-y_s)\right)^2}\ \mathrm{d}\underline{\boldsymbol{x}}\mathrm{d}\boldsymbol{y}.$$

The inner integral evaluates to $\left(\prod_{r,s} 2\pi w_{rs}\right)^{1/2}$. Hence

$$\begin{aligned} I(\delta) &= \frac{(2\pi)^{(d-1)^2/2}}{\delta^{d-1}} \left(\prod_{r,s} w_{rs}\right)^{1/2} \int_{\mathbb{R}^d} e^{-\frac{1}{2}\|\boldsymbol{y}\|_{\ell_2}^2 - \frac{1}{2}\sum_{r,s}\frac{w_{rs}}{\delta^2}(y_r-y_s)^2}\ \mathrm{d}\boldsymbol{y} \\ &= \frac{(2\pi)^{((d-1)^2+d)/2}}{\delta^{d-1}} \left(\prod_{r,s} w_{rs}\right)^{1/2} Det\left(\boldsymbol{I} + \delta^{-2}\boldsymbol{L}(\underline{\boldsymbol{w}})\right)^{-1/2}, \end{aligned}$$

where $\boldsymbol{L}(\underline{\boldsymbol{w}}) \in \mathbb{R}^{d\times d}$ is the Laplacian matrix of the weighted graph G. ∎

7 Discussion

Our main result, Theorem 1, leaves a gap of essentially a factor of two between γ_{low} and γ_{up}. This is a limitation of the methods employed. In particular, the upper bound is likely to be loose due to the lack of concentration of the random variable $\mathcal{Z}$ about its mean, and this translates to the possibility of existence of a non-trivial interval inside $[\gamma_{\mathrm{low}}, \gamma_{\mathrm{up}}]$ where $\mathcal{Z}$ is typically close to 1 while its expectation is exponentially large. A sharper bound could be obtained by computing $\mathbb{E}[|\mathcal{Z} - 1|^{1/n}]$, or even, and perhaps less ambitiously, $\mathbb{E}[|\mathcal{Z} - 1|^\beta]$ for some $0 < \beta < 1$. The first quantity would correspond to the free energy of the model in the limit; the quantity $|\mathcal{Z} - 1|^{1/n}$ is believed to concentrate for large n, so taking its logarithm before or after averaging would lead to the same outcome.

In a different vein, the "sparse" regime where the sets S_a are of constant size k (exactly or on average) could also be of interest. Here, the relevant scaling is one where m is proportional to n. The lower bound argument could be easily extended and yields a bound of $\frac{H(\boldsymbol{\pi})}{(d-1)\log k}$. As for the upper bound, one could in principle follow the same first moment strategy, but our analysis breaks in a quite serious fashion, in that none of our asymptotic estimates hold true in this regime.

Acknowledgments We benefited from insightful conversations with many people. We thank I-Hsiang Wang for sharing the slides of his talk at ITA. We thank Cris Moore for bringing the work of Achlioptas and Naor [AN05] to our attention. We thank Ross Boczar for his input using Mathematica in the trial-and-error process that led to the discovery of Proposition 6. We thank Andrea Sportiello for suggesting the interpolation method that is used in our second proof of Proposition 6. We thank Nikhil Srivastava for bringing the survey [Big97] to our attention. Part of this work was performed during the spring of 2016 when FK and LF were visiting the Simons Institute for the Theory of Computing at UC Berkeley. FK acknowledges funding from the EU (FP/2007-2013/ERC grant agreement 307087-SPARCS). MJ acknowledges the support of the Mathematical Data Science program of the Office of Naval Research under grant number N00014-15-1-2670.

References

[ACO08] Dimitris Achlioptas and Amin Coja-Oghlan. Algorithmic barriers from phase transitions. In *Foundations of Computer Science, 2008. FOCS'08. IEEE 49th Annual IEEE Symposium on*, pages 793–802. IEEE, 2008.

[AM04] Dimitris Achlioptas and Cristopher Moore. The chromatic number of random regular graphs. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 219–228. Springer, 2004.

[AN05] Dimitris Achlioptas and Assaf Naor. The two possible values of the chromatic number of a random graph. *Annals of Mathematics*, 162(3):1335–1351, 2005.

[BCOH$^+$16] Victor Bapst, Amin Coja-Oghlan, Samuel Hetterich, Felicia Raßmann, and Dan Vilenchik. The condensation phase transition in random graph coloring. *Communications in Mathematical Physics*, 341(2):543–606, 2016.

[Big97] Norman Biggs. Algebraic potential theory on graphs. *Bulletin of the London Mathematical Society*, 29(6):641–682, 1997.

[BLM15] Mohsen Bayati, Marc Lelarge, and Andrea Montanari. Universality in polytope phase transitions and message passing algorithms. *Annals of Applied Probability*, 25(2):753–822, 2015.

[BMNN16] Jess Banks, Cristopher Moore, Joe Neeman, and Praneeth Netrapalli. Information-theoretic thresholds for community detection in sparse networks. In *29th Annual Conference on Learning Theory*, volume 49, pages 383–416, 2016.

[BV04] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, Cambridge, UK, 2004.

[Cha82] Seth Chaiken. A combinatorial proof of the all minors matrix tree theorem. *SIAM Journal on Algebraic Discrete Methods*, 3(3):319–329, 1982.

[CO09] Amin Coja-Oghlan. Random constraint satisfaction problems. *arXiv preprint arXiv:0911.2322*, 2009.

[COEH16] Amin Coja-Oghlan, Charilaos Efthymiou, and Samuel Hetterich. On the chromatic number of random regular graphs. *Journal of Combinatorial Theory, Series B*, 116:367–439, 2016.

[COF14] Amin Coja-Oghlan and Alan Frieze. Analyzing walksat on random formulas. *SIAM Journal on Computing*, 43(4):1456–1485, 2014.

[COHH16] Amin Coja-Oghlan, Amir Haqshenas, and Samuel Hetterich. Walksat stalls well below the satisfiability threshold. *arXiv preprint arXiv:1608.00346*, 2016.

[COMV09] Amin Coja-Oghlan, Elchanan Mossel, and Dan Vilenchik. A spectral approach to analysing belief propagation for 3-colouring. *Combinatorics, Probability and Computing*, 18(6):881–912, 2009.

[COP16] Amin Coja-Oghlan and Will Perkins. Belief propagation on replica symmetric random factor graph models. *arXiv preprint arXiv:1603.08191*, 2016.

[DB70] Nicolaas Govert De Bruijn. *Asymptotic Methods in Analysis*. Dover Publications, 1970.

[DH06] Dingzhu Du and Frank Hwang. *Pooling Designs and Nonadaptive Group Testing: Important Tools for DNA Sequencing*, volume 18. World Scientific Publishing Company, 2006.

[DJM13] David L Donoho, Adel Javanmard, and Andrea Montanari. Information-theoretically optimal compressed sensing via spatial coupling and approximate message passing. *IEEE Transactions on Information Theory*, 59(11):7434–7464, 2013.

[DMO12] Varsha Dani, Cristopher Moore, and Anna Olson. Tight bounds on the threshold for permuted k-colorability. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 505–516. Springer, 2012.

[DSS15] Jian Ding, Allan Sly, and Nike Sun. Proof of the satisfiability conjecture for large k. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing*, pages 59–68. ACM, 2015.

[DSS16] Jian Ding, Allan Sly, and Nike Sun. Satisfiability threshold for random regular nae-sat. *Communications in Mathematical Physics*, 341(2):435–489, 2016.

[FPV15] Vitaly Feldman, Will Perkins, and Santosh Vempala. On the complexity of random satisfiability problems with planted solutions. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing*, pages 77–86. ACM, 2015.

[HLB+01] M. Heo, R. L. Leibel, B. B. Boyer, W. K. Chung, M. Koulu, M. K. Karvonen, U. Pesonen, A. Rissanen, M. Laakso, M. I. J. Uusitupa, Y. Chagnon, C. Bouchard, P. A. Donohoue, T. L. Burns, A. R. Shuldiner, K. Silver, R. E. Andersen, O. Pedersen, S. Echwald, T. I. A. Sørensen, P. Behn, M. A. Permutt, K. B. Jacobs, R. C. Elston, D. J. Hoffman, and D. B. Allison. Pooling analysis of genetic data: The association of leptin receptor (LEPR) polymorphisms with variables related to human adiposity. *Genetics*, 159(3):1163–1178, 2001.

[KMZ12] Florent Krzakala, Marc Mézard, and Lenka Zdeborová. Reweighted belief propagation and quiet planting for random K-SAT. *arXiv preprint arXiv:1203.5521*, 2012.

[KZ09] Florent Krzakala and Lenka Zdeborová. Hiding quiet solutions in random constraint satisfaction problems. *Physical Review Letters*, 102(23):238701, 2009.

[MT11] Marc Mézard and Cristina Toninelli. Group testing with random pools: Optimal two-stage algorithms. *IEEE Transactions on Information Theory*, 57(3):1736–1745, 2011.

[Roc70] R. Tyrrell Rockafellar. *Convex Analysis*. Princeton University Press, Princeton, 1970.

[SBC+02] Pak Sham, Joel S Bader, Ian Craig, Michael O'Donovan, and Michael Owen. DNA pooling: a tool for large-scale association studies. *Nature Reviews Genetics*, 3(11):862–871, 2002.

[SSZ16] Allan Sly, Nike Sun, and Yumeng Zhang. The number of solutions for random regular nae-sat. *arXiv preprint arXiv:1604.08546*, 2016.

[Tan02] Toshiyuki Tanaka. A statistical-mechanics approach to large-system analysis of CDMA multiuser detectors. *IEEE Transactions on Information theory*, 48(11):2888–2910, 2002.

[Vaa79] Jeffrey D. Vaaler. A geometric inequality with applications to linear forms. *Pacific Journal of Mathematics*, 83(2):543–553, 1979.

[WHLC16] I-Hsiang Wang, Shao-Lun Huang, Kuan-Yun Lee, and Kwang-Cheng Chen. Data extraction via histogram and arithmetic mean queries: Fundamental limits and algorithms. In *2016 IEEE International Symposium on Information Theory (ISIT)*, pages 1386–1390. IEEE, 2016.

[WV09] Yihong Wu and Sergio Verdú. Fundamental limits of almost lossless analog compression. In *2009 IEEE International Symposium on Information Theory*, pages 359–363. IEEE, 2009.

[Zig04] Kamil Sh. Zigangirov. *Theory of Code Division Multiple Access Communication*, volume 6. John Wiley & Sons, 2004.

[ZK15] Lenka Zdeborová and Florent Krzakala. Statistical physics of inference: Thresholds and algorithms. *arXiv preprint arXiv:1511.02476*, 2015.

[ZKMZ13] Pan Zhang, Florent Krzakala, Marc Mézard, and Lenka Zdeborová. Non-adaptive pooling strategies for detection of rare faulty items. In *2013 IEEE International Conference on Communications Workshops (ICC)*, pages 1409–1414. IEEE, 2013.