

Hypergraph coloring up to condensation

Peter Ayre

School of Mathematics and Statistics
UNSW Sydney
NSW 2052, Australia
peter.ayre@unsw.edu.au

Amin Coja-Oghlan*

Mathematics Institute
Goethe University
Frankfurt 60325, Germany
acoghlan@math.uni-frankfurt.de

Catherine Greenhill[†]

School of Mathematics and Statistics
UNSW Sydney
NSW 2052, Australia
c.greenhill@unsw.edu.au

April 13, 2018

Abstract

Improving a result of Dyer, Frieze and Greenhill [Journal of Combinatorial Theory, Series B, 2015], we determine the q -colorability threshold in random k -uniform hypergraphs up to an additive error of $\ln 2 + \varepsilon_q$, where $\lim_{q \rightarrow \infty} \varepsilon_q = 0$. The new lower bound on the threshold matches the “condensation phase transition” predicted by statistical physics considerations [Krzakala et al., PNAS 2007].

Mathematics Subject Classification: 05C80 (primary), 05C15 (secondary)

1 Introduction

Recent work on random constraint satisfaction problems has focused either on the case of binary variables and k -ary constraints (e.g., random k -SAT) or on the case of k -ary variables and binary constraints (e.g., random graph coloring) for some $k \geq 3$. In these two cases substantial progress has been made over the past few years. For instance, the k -SAT threshold has been identified precisely for large enough k [12]. Moreover, in the random hypergraph 2-coloring problem (or equivalently the k -NAESAT problem) the threshold is known up to an error term that tends to 0 rapidly in terms of the size k of the edges [11]. In addition, the best current upper and lower bounds on the k -colorability threshold of the Erdős-Rényi random graph are within a small additive constant [9]. By comparison, little is known about problems in which both the arity of the constraints and the domain of the variables have size greater than two. Although it has been asserted that the techniques developed in recent work should carry over [9], this claim has hardly been put to the test.

The present paper deals with one of the most natural examples of a problem with k -ary constraints and q -ary variables with $q, k \geq 3$, namely q -colorability of random k -uniform hypergraphs. Let $[m]$ denote the set $\{1, \dots, m\}$ for any positive integer m . To be precise, by a q -coloring of $G = (V, E)$ we mean a map $\sigma : V \rightarrow [q]$ such that $|\sigma(e)| > 1$ for all $e \in E$, i.e., no edge is monochromatic. The chromatic number of G is the least q for which a q -coloring exists. The random hypergraph model that we consider is the most natural one, i.e., $\mathcal{G} \in \mathcal{G}(n, k, m)$

*The research leading to these results has received funding from the European Research Council under the European Union’s Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement n. 278857-PTCC.

[†]Research supported by the Australian Research Council Discovery Project DP140101519.

is a (simple) k -uniform hypergraph on the vertex set $[n] := \{1, 2, 3, \dots, n\}$ with a set of precisely m edges chosen uniformly at random.

For every $q \geq 2, k \geq 3$ there exists a (non-uniform) sharp threshold $c_{q,k} = c_{q,k}(n)$ for q -colorability [17]. That is, if $m = m(n)$ is a sequence such that for some fixed $\varepsilon > 0$ we have $m(n) < (1 - \varepsilon)nc_{q,k}(n)$, then $\mathcal{G}(n, k, m)$ is q -colorable w.h.p., whereas w.h.p. the random hypergraph fails to be q -colorable if $m(n) > (1 + \varepsilon)nc_{q,k}(n)$. The best prior bounds on this threshold, obtained by Dyer, Frieze and Greenhill [13, Remark 2.1, (82)], read

$$(q^{k-1} - 1) \ln q - 1 - \varepsilon_{q,k} \leq \liminf_{n \rightarrow \infty} c_{q,k}(n) \leq \limsup_{n \rightarrow \infty} c_{q,k}(n) \leq (q^{k-1} - 1/2) \ln q, \quad (1.1)$$

where $\lim_{q \rightarrow \infty} \varepsilon_{q,k} = 0$ for any fixed $k \geq 3$. Thus, the upper and the lower bound differ by an additive $\frac{1}{2} \ln q + 1 + \varepsilon_{q,k}$, a term that diverges in the limit of large q . The main result of this paper provides an improved lower bound that is within an additive $\ln 2$ of the upper bound from (1.1), in the large- q limit.

Theorem 1.1. *For each $k \geq 3$ there is a number $q_0 = q_0(k) > 0$ such that for all $q > q_0$ we have*

$$\liminf_{n \rightarrow \infty} c_{q,k}(n) \geq (q^{k-1} - 1/2) \ln q - \ln 2 - 1.01 \ln q / q.$$

The proof of Theorem 1.1 is based on the second moment method. So is [13], which generalises the second moment argument of Achlioptas and Naor [4] from graphs to hypergraphs. The result of Achlioptas and Naor was recently improved by Coja-Oghlan and Vilenchik [9], and in this paper we generalise the argument from that paper to hypergraphs. While numerous details need adjusting, the basic proof strategy that we pursue is similar to the one suggested in [9]. In particular, the improvement over [13] results from studying the second moment of a subtly chosen random variable. While the random variable considered in [13] is just the number of (balanced) q -colorings of the random hypergraph, here we use a random variable that is inspired by ideas from statistical mechanics; we will give a more detailed outline in Section 3 below. Thus, the present paper shows that, indeed, with a fair number of careful modifications the method from [9] can be generalised to hypergraphs.

Notation. We assume throughout that the number of vertices, n , is sufficiently large for our estimates to hold. We also assume that the number of colors, q , exceeds some large enough constant $q_0 = q_0(k)$. But of course q, k are always assumed to remain fixed as $n \rightarrow \infty$.

We use the O -notation to refer to the limit $n \rightarrow \infty$. For example, $f(n) = O(g(n))$ means that there exists some $C > 0, n_0 > 0$ such that for all $n > n_0$ we have $|f(n)| \leq C \cdot |g(n)|$. In addition, $o(\cdot), \Omega(\cdot), \Theta(\cdot)$ take their usual definitions, except that we assume the expression $\Omega(n)$ is positive (for sufficiently large n) whenever we write $\exp(-\Omega(n))$. We write $f(n) \sim g(n)$ if $\lim_{n \rightarrow \infty} f(n)/g(n) = 1$.

When discussing estimates that hold in the limit of large q we will make this explicit by adding the subscript q to the asymptotic notation. Therefore, $f(q) = O_q(g(q))$ means that there exists positive constants C, q_0 such that for all $q > q_0$ we have $|f(q)| \leq C \cdot |g(q)|$. Furthermore, we will write $f(q) = \tilde{O}_q(g(q))$ to indicate that there exists positive C, q_0 such that for all $q > q_0$ we have $|f(q)| \leq (\ln q)^C \cdot |g(q)|$.

2 Related work

The quest for the chromatic number of random graphs (i.e., $\mathcal{G}(n, 2, m)$) goes back to the seminal 1960 paper of Erdős and Rényi in which they established the “giant component” phase transition [14]. But it took almost thirty years until a celebrated paper of Bollobás [7] determined the asymptotic value of the chromatic number of dense random graphs. His proof used martingale tail bounds, which were introduced to combinatorics by Shamir and Spencer [26] to investigate the concentration of the chromatic number. Building upon ideas of Matula [25], Łuczak [23] determined the asymptotic value of the chromatic number of the Erdős-Rényi random graph in the case that $m = m(n)$ satisfies $m/n \rightarrow \infty$. However, the results from [7, 23] only determine the chromatic number up to a *multiplicative* error of $1 + o(1)$ as $n \rightarrow \infty$, and the resulting error term exceeds the width within which

the chromatic number is known to be concentrated. Indeed, in the case that $m = m(n) \leq n^{3/2 - \Omega(1)}$ it is known that the chromatic number of the random graph is concentrated on two subsequent integers [6, 24]. In the sparse case $m = O(n)$ the precise values of these two integers are implied by the current bounds on the q -colorability threshold [4, 8, 9].

The 2-colorability problem in random hypergraphs, which is essentially equivalent to the random k -NAESAT problem, has also been studied. Achlioptas and Moore [2, 3] showed that the 2-colorability threshold can be approximated within a small additive constant via the second moment method. Furthermore, Coja-Oghlan and Zdeborová [10] established the existence of a further phase transition apart from the threshold for 2-colorability, the “condensation phase transition”. The name derives from an intriguing connection to the statistical mechanics of glasses [19, 21]. Moreover, the argument of Coja-Oghlan and Panagiotou [11] determines the 2-colorability threshold in k -uniform random hypergraphs up to an additive error term ε_k that tends to 0 exponentially as a function of k .

Prior to the aforementioned work of Dyer, Frieze and Greenhill [13] the q -colorability problem in hypergraphs was studied by Krivelevich and Sudakov [20], who also considered other possible notions of colorings. Their results are of a similar nature to Łuczak’s [23] in the case of graphs. That is, they determine the value of the chromatic number up to a multiplicative $1 + o(1)$ factor, with $o(1)$ hiding a term that vanishes as $m/n \rightarrow \infty$. The same is true of the results of Kupavskii and Shabanov [22], which partly improve upon [20]. However, the bounds on the q -colorability threshold that can be read out of [20, 22] are less precise than those obtained in [13] (upon which Theorem 1.1 improves).

3 Outline

Throughout, we assume that n is sufficiently large for our error estimates to hold, and that $q > q_0$. Further, we assume that $m = \lceil cn \rceil$ and for ease of notation will often write cn rather than $\lceil cn \rceil$.

The second moment method. The second moment method has become the mainstay for lower-bounding satisfiability thresholds [2, 5, 15].

Suppose that we can construct a non-negative random variable Z on $\mathcal{G}(n, k, cn)$ such that the event $Z(G) > 0$ implies q -colorability, and such that

$$\mathbf{E}[Z^2] = O(\mathbf{E}[Z]^2) \quad \text{as } n \rightarrow \infty. \quad (3.1)$$

Then the Paley-Zygmund inequality implies that

$$\liminf_{n \rightarrow \infty} \mathbf{P}[Z > 0] \geq \liminf_{n \rightarrow \infty} \frac{\mathbf{E}[Z]^2}{\mathbf{E}[Z^2]} > 0. \quad (3.2)$$

Combining (3.2) with the sharp threshold result from [17], which establishes the existence of a sharp threshold sequence $c_{q,k}(n)$, yields $\liminf_{n \rightarrow \infty} c_{q,k}(n) \geq c$. Hence, the second moment method can be summarised as follows.

Fact 3.1. *If there is a non-negative random variable Z on $\mathcal{G}(n, k, cn)$ such that $Z > 0$ implies q -colorability and (3.1) holds, then $\liminf_{n \rightarrow \infty} c_{q,k}(n) \geq c$.*

Thus, our task is to exhibit a random variable Z on $\mathcal{G}(n, k, cn)$ that satisfies (3.1) for as large a value of c as possible.

Balanced colorings. Certainly the most natural choice for Z seems to be the number Z_q of q -colorings of the random hypergraph. Clearly, $Z_q \geq 0$ and $Z_q(G) > 0$ only if G is q -colorable. However, technically Z_q is a bit unwieldy. Therefore, following Achlioptas and Naor [4], Dyer, Frieze and Greenhill [13] considered a slightly modified random variable. Namely, let us call a map $\sigma : [n] \rightarrow [q]$ *balanced* if $|\sigma^{-1}(j) - n/q| \leq \sqrt{n}$ for all $j \in [q]$ and let $Z_{q,\text{bal}}$ be the number of balanced q -colorings of \mathcal{G} .

Lemma 3.2 ([13]). *For any $q, k \geq 3$ and any $c > 0$ we have*

$$\mathbf{E}[Z_{q,\text{bal}}] = \Theta \left[\left(q(1 - q^{1-k})^c \right)^n \right] \quad (3.3)$$

and

$$\lim_{n \rightarrow \infty} \frac{1}{n} \ln \mathbf{E}[Z_q] = \lim_{n \rightarrow \infty} \frac{1}{n} \ln \mathbf{E}[Z_{q,\text{bal}}] = \ln q + c \ln(1 - q^{1-k}). \quad (3.4)$$

Proof. Calculations similar to the following ones were performed in [13]; we repeat them here to keep the paper self-contained. Given a balanced map $\sigma : [n] \rightarrow [q]$, let $\alpha_i = |\sigma^{-1}(i)|/n$ for $i \in [q]$ and define $\alpha = (\alpha_1, \dots, \alpha_q)$. Stirling's formula yields

$$\mathbf{P}[\sigma \text{ is a proper } q\text{-coloring of } \mathcal{G}] = \frac{\binom{n}{k} - \sum_{i=1}^q \binom{\alpha_i n}{k}}{c^n} \binom{n}{cn}^{-1} = \Theta \left(\exp \left[cn \ln \left(1 - \sum_{i \in [q]} \alpha_i^k \right) \right] \right), \quad (3.5)$$

cf. [13, equation (8)]. Let $\bar{\alpha} = (1/q, \dots, 1/q)$ denote the uniform distribution on $[q]$. The gradient of the function $f : (x_1, \dots, x_k) \mapsto \ln(1 - \sum_{i \in [q]} x_i^k)$ at the point $\bar{\alpha}$ is simply the vector $\nabla f(\bar{\alpha})$ with every entry equal to $k(1 - q^{k-1})^{-1}$. Consequently, because $\sum_{i \in [q]} (\alpha_i - 1/q) = 0$, expanding f to the second order around $\bar{\alpha}$ yields

$$\ln \left(1 - \sum_{i \in [q]} \alpha_i^k \right) = f(\alpha) = f(\bar{\alpha}) + \nabla f(\bar{\alpha})(\alpha - \bar{\alpha}) + O(\|\alpha - \bar{\alpha}\|_2^2) = \ln(1 - q^{1-k}) + O(\|\alpha - \bar{\alpha}\|_2^2). \quad (3.6)$$

Since σ is balanced, we have the bound $\|\alpha - \bar{\alpha}\|_2^2 = O(1/n)$. Therefore, combining (3.5) and (3.6), we obtain

$$\mathbf{P}[\sigma \text{ is a proper } q\text{-coloring of } \mathcal{G}] = \Theta \left((1 - q^{1-k})^{cn} \right), \quad (3.7)$$

uniformly for all balanced σ . Finally, the number of balanced maps corresponding to a given α is $\binom{n}{\alpha_1 n, \dots, \alpha_q n} = \Theta(n^{(1-q)/2}) q^n$, by Stirling's formula, and the number of choices for the vector α is $\Theta(n^{(q-1)/2})$. Hence the total number of balanced maps σ is $\Theta(q^n)$. Combining this with (3.7) implies (3.3).

Next, as observed in the proof of [13, Lemma 2.1], the probability that a map $\sigma : [n] \rightarrow [q]$ is a q -colouring of \mathcal{G} is maximised when σ is perfectly balanced, and this probability equals $O(1)(1 - q^{1-k})^{cn}$. (Here the $O(1)$ factor is needed only when q does not divide n .) Hence, by linearity of expectation,

$$\mathbf{E}[Z_{q,\text{bal}}] \leq \mathbf{E}[Z_q] = O(1) \left(q(1 - q^{1-k})^c \right)^n,$$

which differs from (3.3) by at most a constant factor. This implies (3.4), completing the proof. \square

It is easily verified that the r.h.s. of (3.4) is positive if $c < (q^{k-1} - \frac{1}{2}) \ln q - \ln 2$. Hence, for such c , both $\mathbf{E}[Z_q]$ and $\mathbf{E}[Z_{q,\text{bal}}]$ are exponential in n . They differ only in their sub-exponential terms. Consequently, we do not give anything away by confining ourselves to balanced colorings only. In the following we will see why neither Z_q nor $Z_{q,\text{bal}}$ is a good random variable to work with and why neither can be used to prove Theorem 1.1. What we learn will guide us towards constructing a better random variable.

While working out the first moment of $Z_{q,\text{bal}}$ (i.e., the proof of Lemma 3.2) is pretty straightforward, getting a handle on the second moment is not quite so easy. Of course, the second moment of $Z_{q,\text{bal}}$ is nothing but the expected number of *pairs* of balanced q -colorings. Moreover, the probability that two maps $\sigma, \tau : [n] \rightarrow [q]$ simultaneously happen to be q -colorings of \mathcal{G} will depend on how "similar" σ, τ are. To gauge similarity, define the *overlap* of σ, τ as the $q \times q$ -matrix $a(\sigma, \tau) = (a_{ij}(\sigma, \tau))_{i,j \in [q]}$ with entries

$$a_{ij}(\sigma, \tau) = n^{-1} |\sigma^{-1}(i) \cap \tau^{-1}(j)|.$$

In words, $a_{ij}(\sigma, \tau)$ is the probability that a random vertex $v \in [n]$ has color i under σ and color j under τ . Then we can cast the second moment in terms of the overlap as follows. Let $\mathcal{R} = \mathcal{R}_{n,q}$ be the set of all overlaps $a(\sigma, \tau)$ of balanced $\sigma, \tau : [n] \rightarrow [q]$. Though the results of the next lemma can be found in [13], for completeness we provide a brief proof here.

Lemma 3.3 ([13]). Let $\|a\|_k = \left[\sum_{i,j \in [q]} a_{ij}^k \right]^{1/k}$ be the ℓ_k -norm and define

$$H(a) = - \sum_{i,j \in [q]} a_{ij} \ln a_{ij}, \quad E(a) = E_{q,c,k}(a) = c \ln \left[1 - 2q^{1-k} + \|a\|_k^k \right].$$

Let $F(a) = H(a) + E(a)$ and suppose that $a^* \in \mathcal{R}$ satisfies $F(a^*) = \max_{a \in \mathcal{R}} F(a)$. Then

$$\mathbf{E}[Z_{q,\text{bal}}^2] = \exp[nF(a^*) + o(n)]. \quad (3.8)$$

Next, let ξ be a positive constant and suppose that $\mathcal{A} \subseteq \mathcal{R}$ has the following property: $a_{ij} \geq \xi$ for all $a \in \mathcal{A}$ and all $i, j \in [q]$. Then

$$\mathbf{E}[Z_{q,\text{bal}}^2 \cdot \mathbf{1}_{\mathcal{A}}] = \Theta(n^{(1-q^2)/2}) \sum_{a \in \mathcal{A}} \exp[nF(a)]. \quad (3.9)$$

Proof. First, observe that for a given $a \in \mathcal{R}$, the number of σ, τ with overlap a is given by the multinomial coefficient

$$\binom{n}{a_{11}n, a_{12}n, \dots, a_{qq}n}.$$

Next, fix balanced maps σ, τ with overlap a . By inclusion-exclusion, the probability that a random edge chosen uniformly out of all $\binom{n}{k}$ possible edges is monochromatic under either σ or τ equals

$$\begin{aligned} & \binom{n}{k}^{-1} \left[\sum_{i \in [q]} \left[\binom{\sum_{j \in [q]} a_{ij}n}{k} + \binom{\sum_{j \in [q]} a_{ji}n}{k} \right] - \sum_{i,j \in [q]} \binom{a_{ij}n}{k} \right] \\ &= \sum_{i \in [q]} \left[\left(\sum_{j \in [q]} a_{ij} \right)^k + \left(\sum_{j \in [q]} a_{ji} \right)^k \right] - \sum_{i,j \in [q]} a_{ij}^k + O\left(\frac{1}{n}\right). \end{aligned} \quad (3.10)$$

To simplify this expression we observe that since σ, τ are balanced,

$$\begin{aligned} \sum_{i \in [q]} \left(\sum_{j \in [q]} a_{ij} \right)^k &= \sum_{i \in [q]} \left(q^{-1} - \left(q^{-1} - \sum_{j \in [q]} a_{ij} \right) \right)^k \\ &= q^{1-k} - kq^{1-k} \left(1 - \sum_{i,j \in [q]} a_{ij} \right) + O\left(\sum_{i \in [q]} \left(\frac{1}{q} - \sum_{j \in [q]} a_{ij} \right)^2 \right) \\ &= q^{1-k} + O(1/n), \end{aligned}$$

because $\sum_{i,j \in [q]} a_{ij} = 1$ and $|q^{-1} - \sum_{j \in [q]} a_{ij}| = O(n^{-1/2})$. Of course, the same steps apply to $\sum_{i \in [q]} \left(\sum_{j \in [q]} a_{ji} \right)^k$. Hence, since σ and τ are balanced, (3.10) can be written as

$$2q^{1-k} - \sum_{i,j=1}^q a_{ij}^k + O\left(\frac{1}{n}\right).$$

Therefore

$$\mathbf{E}[Z_{q,\text{bal}}^2] \sim \sum_{a \in \mathcal{R}} \binom{n}{a_{11}n, a_{12}n, \dots, a_{qq}n} \exp(nE(a) + O(1)). \quad (3.11)$$

Let $b \vee 1$ denote $\max\{b, 1\}$. We give upper and lower bounds on the multinomial coefficient by applying Stirling's formula in the form

$$b! = \sqrt{2\pi(b \vee 1)} \left(\frac{b}{e} \right)^b \left[1 + O\left(\frac{1}{b+1} \right) \right],$$

which holds for all nonnegative integers b . This gives

$$\binom{n}{a_{11}n, a_{12}n, \dots, a_{qq}n} \sim (2\pi n)^{(1-q^2)/2} \exp[nH(a)] \prod_{i,j \in [q]} (a_{ij} \vee 1/n)^{-1/2}. \quad (3.12)$$

Since $1/n \leq a_{ij} \vee 1/n \leq 2/q$ for all $i, j \in [q]$, and since each row and column sum equals $1/q + o(1)$, the product over $i, j \in [q]$ in (3.12) is always bounded below by a constant and (easily) bounded above by $O(n^{(q^2-1)/2})$. Therefore

$$\Omega(n^{(1-q^2)/2}) \exp[nH(a)] \leq \binom{n}{a_{11}n, a_{12}n, \dots, a_{qq}n} = O(1) \exp[nH(a)].$$

Combining the above leads to

$$\Omega(n^{(1-q^2)/2}) \sum_{a \in \mathcal{R}} \exp[nF(a)] \leq \mathbf{E}[Z_{q,\text{bal}}^2] \leq O(1) \sum_{a \in \mathcal{R}} \exp[nF(a)].$$

Taking just the term corresponding to $a = a^*$ in the lower bound gives the lower bound of (3.8), and the upper bound follows using the fact that $|\mathcal{R}| \leq n^{q^2}$.

Next, observe that if $a \in \mathcal{A}$ then $\prod_{i,j \in [q]} (a_{ij} \vee 1/n)^{-1/2} = \Theta(1)$. Substituting this into (3.12) and restricting the sum in (3.11) to \mathcal{A} completes the proof of (3.9). \square

Let $\mathcal{D} \subseteq \mathbb{R}^{q^2}$ be the polytope comprising of all $a = (a_{ij})_{i,j \in [q]}$ such that

$$\sum_{j \in [q]} a_{ij} = \sum_{j \in [q]} a_{ji} = 1/q \quad \text{for all } i \in [q], \quad a_{ij} \geq 0 \quad \text{for all } i, j \in [q].$$

Then \mathcal{D} is the Birkhoff polytope, scaled by a constant factor, and $\mathcal{R} \cap \mathcal{D}$ is dense in \mathcal{D} as $n \rightarrow \infty$. Therefore, (3.8) yields

$$\lim_{n \rightarrow \infty} \frac{1}{n} \ln \mathbf{E}[Z_{q,\text{bal}}^2] = \max_{a \in \mathcal{D}} F(a).$$

Further, evaluating the function $F(a)$ from Lemma 3.3 at the ‘‘flat’’ overlap $\bar{a} = (\bar{a}_{ij})$ with $\bar{a}_{ij} = q^{-2}$ for all $i, j \in [q]$, we find

$$F(\bar{a}) = 2 \left[\ln q + c \ln(1 - q^{1-k}) \right]. \quad (3.13)$$

This term is precisely twice the exponential order of the first moment from (3.4). Consequently, the second moment bound $\mathbf{E}[Z_{q,\text{bal}}^2] = O(\mathbf{E}[Z_{q,\text{bal}}]^2)$ can hold *only if*

$$F(\bar{a}) = \max_{a \in \mathcal{D}} F(a). \quad (3.14)$$

In fact, the Laplace method applied along the lines of [16, Theorem 2.3] shows that the condition (3.14) is both necessary and sufficient for the success of the second moment method. In summary, the second moment argument reduces to the analytic problem of maximising the function F over the polytope \mathcal{D} .

A relaxation. This maximisation problem is anything but straightforward. Following [4], Dyer, Frieze and Greenhill [13] consider a relaxation. Namely, instead of optimising F over \mathcal{D} , they consider the (substantially) bigger domain \mathcal{S} of all $a = (a_{ij})_{i,j \in [q]}$ such that $\sum_{j=1}^q a_{ij} = 1/q$ for all $i \in [q]$ and $a_{ij} \geq 0$ for all $i, j \in [q]$, dropping the constraint that the ‘‘column sums’’ $\sum_j a_{ji}$ equal $1/q$. Note that \mathcal{S} is the set of singly (row) stochastic matrices, scaled by a constant factor. Clearly, $\max_{a \in \mathcal{D}} F(a) \leq \max_{a \in \mathcal{S}} F(a)$. Furthermore, Dyer, Frieze and Greenhill solve the latter maximisation problem precisely by generalising the techniques from [4], requiring rather lengthy technical arguments. The result is that for c up to the lower bound in (1.1) we indeed have $\max_{a \in \mathcal{S}} F(a) = F(\bar{a})$.

But this method does not work up to the density promised by Theorem 1.1. There are two obstacles. First, not far beyond the lower bound in (1.1) the maximum of F over \mathcal{S} is attained at a point $a' \in \mathcal{S} \setminus \mathcal{D}$, i.e., $F(a') > F(\bar{a})$.

Thus, relaxing \mathcal{D} to the larger domain \mathcal{S} gives too much away. Second, there exists a constant $\gamma > \ln 2$ such that for $c = (q^{k-1} - 1/2) \ln q - \gamma$, the value of F attained at

$$a_{\text{stable}} = (q^{-1} - q^{-k}) \text{id} + q^{-k}(q-1)^{-1}(q^2 \bar{a} - \text{id}) \in \mathcal{D}$$

is strictly greater than $F(\bar{a})$. (Note that every entry of the matrix $q^2 \bar{a}$ equals 1.) Consequently, even if we could solve the analytic problem of maximising F over the actual domain \mathcal{D} it would be insufficient to prove Theorem 1.1.

Tame colorings. The above discussion shows that it is impossible to prove Theorem 1.1 via the second moment method applied to $Z_{q,\text{bal}}$. A similar problem occurs in the case of random graphs ($k = 2$), see [9]. To remedy this problem in the hypergraph case we will generalise the strategy from [9].

The key idea is to introduce a random variable $Z_{q,\text{tame}}$ that takes the typical geometry of the set $B(\mathcal{G})$ of all balanced q -colorings of $\mathcal{G} \in \mathcal{G}(n, k, cn)$ into account, such that $0 \leq Z_{q,\text{tame}} \leq Z_{q,\text{bal}}$. According to predictions based on non-rigorous physics considerations [21], the set $B(\mathcal{G})$ has a geometry that is very different from that of a random subset of the cube $[q]^{[n]}$ of the same size. More precisely, for almost all k -uniform hypergraphs G with cn edges, the set $B(G)$ decomposes into well-separated ‘‘clusters’’ which each contains an exponential number of colorings. However, the fraction of colorings that any single cluster contains is only an exponentially small fraction of the total number of q -colorings of G . Furthermore, while it is possible to walk inside the set $B(G)$ from any coloring to any other colouring in the same cluster by only changing the colors of $O(\ln n)$ vertices at a time, it is impossible to get from one cluster to another without changing the colors of $\Omega(n)$ vertices in a single step. Now, the basic idea is to let $Z_{q,\text{tame}} = Z_{q,\text{bal}} \cdot \mathbf{1}\{\mathcal{T}\}$, where \mathcal{T} is the event that the geometry of the set $B(G)$ has the aforementioned properties.

To make this rigorous, we define the *cluster* of a q -coloring σ of a hypergraph G as the set

$$\mathcal{C}(G, \sigma) = \left\{ \tau \in B(G) : \min_{i \in [q]} a_{ii}(\sigma, \tau) > q^{-1} (1.01/k)^{1/(k-1)} \right\}.$$

In words, $\mathcal{C}(G, \sigma)$ contains all balanced q -colorings τ of G where, for each color i , at least a $(1.01/k)^{1/(k-1)}$ fraction of all vertices colored i under σ retain color i under τ . Call a q -coloring σ of G *separable* if

$$\forall \tau \in B(G), \forall i, j \in [q], a_{ij}(\sigma, \tau) \notin (q^{-1} (1.01/k)^{1/(k-1)}, q^{-1} (1 - \kappa)) \quad \text{where } \kappa = q^{1-k} \ln^{20} q. \quad (3.15)$$

Definition 3.4. A q -coloring σ of the (fixed) hypergraph G is tame if

$$\mathbf{T1}: \sigma \text{ is balanced}, \quad \mathbf{T2}: \sigma \text{ is separable}, \quad \mathbf{T3}: |\mathcal{C}(G, \sigma)| \leq \mathbf{E}[Z_{q,\text{bal}}].$$

Definition 3.4 generalises the concept of ‘‘tame graph colorings’’ from [9, Definition 2.3].

The set of tame colorings of a given hypergraph G decomposes into well-separated clusters. Indeed, the separability condition ensures that the clusters of two tame colorings σ, τ of G are either disjoint or identical. Furthermore, **T3** ensures that no cluster size exceeds the expected number of balanced colorings, i.e., the clusters are ‘‘small’’. This will allow us to control the contribution to the second moment from colourings which lie in the same cluster (see Lemma 5.4). Furthermore, if σ, τ are tame colorings then the overlap $a(\sigma, \tau)$ cannot equal the matrix a_{stable} defined above, as this matrix fails **T2**. So, restricting attention to tame colourings excludes the matrix a_{stable} .

Let $Z_{q,\text{tame}}$ be the number of tame q -colorings of $\mathcal{G}(n, k, cn)$. With the right random variable in place, our task boils down to calculating the first and the second moment. In Section 4 we will prove that the first moment of $Z_{q,\text{tame}}$ is asymptotically equal to the first moment of $Z_{q,\text{bal}}$. For the following two propositions we assume that

$$(q^{k-1} - 1/2) \ln q - 2 \leq c \leq (q^{k-1} - 1/2) \ln q - \ln 2 - 1.01 \ln q / q.$$

That is, we consider values of c which lie between the standard second-moment lower bound (on the q -colorability threshold $c_{q,k}$) and the one we prove here.

Proposition 3.5. *There is a number $q_0 > 0$ such that for all $q > q_0$ we have $\mathbf{E}[Z_{q,\text{tame}}] \sim \mathbf{E}[Z_{q,\text{bal}}]$.*

Further, in Section 5 we establish the following bound on the second moment.

Proposition 3.6. *There is a number $q_0 > 0$ such that for all $q > q_0$, if $\mathbf{E}[Z_{q,\text{tame}}] \sim \mathbf{E}[Z_{q,\text{bal}}]$ then*

$$\mathbf{E}[Z_{q,\text{tame}}^2] = O(\mathbf{E}[Z_{q,\text{bal}}]^2) = O(\mathbf{E}[Z_{q,\text{tame}}]^2).$$

Thus, while moving to tame colorings has no discernible effect on the first moment, Proposition 3.6 shows that the impact on the second moment is dramatic. Indeed, the matrix a_{stable} shows that $\mathbf{E}[Z_{q,\text{bal}}^2] \geq \exp(\Omega(n))\mathbf{E}[Z_{q,\text{bal}}]^2$ for c near the bound in Theorem 1.1, while $\mathbf{E}[Z_{q,\text{tame}}^2] = O(\mathbf{E}[Z_{q,\text{bal}}]^2)$ for all c up to $(q^{k-1} - 1/2)\ln q - \ln 2 - 1.01 \ln q/q$. Then Theorem 1.1 follows from applying Fact 3.1 to $Z_{q,\text{tame}}$, by Propositions 3.5 and 3.6.

Finally, the obvious question is whether the approach taken in this work can be pushed further to actually obtain tight upper and lower bounds on the q -colorability threshold. However, it follows from the proof of Propositions 3.5 that the answer is “no”. More specifically, in Section 4.4 we prove the following.

Corollary 3.7. *For any $k \geq 3$ there exists a sequence $(\varepsilon_q)_{q \geq 3}$ such that $\lim_{q \rightarrow \infty} \varepsilon_q = 0$ and such that the following is true: For any c such that*

$$(q^{k-1} - 1/2)\ln q - \ln 2 + \varepsilon_q < c < (q^{k-1} - 1/2)\ln q$$

there exists $\delta > 0$ such that

$$\lim_{n \rightarrow \infty} \mathbf{P}[Z_q < \exp(-\delta n)\mathbf{E}[Z_q]] = 1. \quad (3.16)$$

Now, assume for a contradiction that there is a random variable $0 \leq Z \leq Z_q$ with the following properties. First, $Z(G) > 0$ only if G is q -colorable. Second, $\ln \mathbf{E}[Z] \sim \ln \mathbf{E}[Z_q]$ (cf. Lemma 3.2 and Proposition 3.5). Third, $\mathbf{E}[Z^2] = O(\mathbf{E}[Z]^2)$. Then the Paley-Zygmund inequality implies that

$$\lim_{\delta \rightarrow 0} \lim_{n \rightarrow \infty} \mathbf{P}[Z_q \geq \exp(-\delta n)\mathbf{E}[Z_q]] \geq \lim_{\delta \rightarrow 0} \lim_{n \rightarrow \infty} \mathbf{P}[Z \geq \exp(-\delta n)\mathbf{E}[Z]] > 0,$$

in contradiction to (3.16). Corollary 3.7 is in line with the physics prediction that the actual q -colorability threshold is preceded by another phase transition called *condensation* [21], beyond which w.h.p. $Z_q \leq \exp(-\Omega(n))\mathbf{E}[Z_q]$. In particular, the lower bound of Theorem 1.1 matches this “condensation threshold” up to an error term that tends to 0 in the limit of large q .

4 The first moment

Throughout this section, unless specified otherwise we take $\sigma, \tau : [n] \rightarrow [q]$ as balanced maps, and assume that

$$(q^{k-1} - 1/2)\ln q - 2 \leq c \leq (q^{k-1} - 1/2)\ln q - \ln 2 - 1.01 \ln q/q. \quad (4.1)$$

We frequently make use of the Chernoff bound.

Lemma 4.1. *([18, Theorem 2.1]) Let $\phi(x) = (1+x)\ln(1+x) - x$. Let X be a binomial random variable with mean $\mu > 0$. Then for any $t > 0$*

$$\mathbf{P}[X > \mu + t] \leq \exp\{-\mu\phi(t/\mu)\}, \quad \mathbf{P}[X < \mu - t] \leq \exp\{-\mu\phi(-t/\mu)\}.$$

In particular, for any $t > 1$ we have $\mathbf{P}[X > t\mu] \leq \exp\{-t\mu \ln(t/e)\}$.

4.1 The planted model

The aim in this section is to establish Proposition 3.5, the lower bound on the expected number of tame colorings. Let $\sigma : [n] \rightarrow [q]$ be a (fixed) balanced map that assigns each vertex a color. It suffices to prove that

$\mathbf{P}[\sigma \text{ is a tame coloring of } \mathcal{G} | \sigma \text{ is a coloring of } \mathcal{G}] = 1 - o(1)$. Furthermore, the conditional distribution of \mathcal{G} given that σ is a coloring admits an easy explicit description: the conditional random hypergraph simply consists of m random edges chosen uniformly out of all edges that are not monochromatic under σ .

It will however be convenient to work with a slightly different distribution. Let $\mathcal{G}_\sigma \in \mathcal{G}(n, k, cn, \sigma)$ be the hypergraph on $[n]$ obtained by including every edge that is not monochromatic under σ with probability p , independently, where

$$p = \frac{cn}{\binom{n}{k} - \prod_{j=1}^q \binom{|\sigma^{-1}(j)|}{k}} \sim \frac{ck! \cdot (1 + O(1/n))}{n^{k-1}(1 - q^{1-k})} = O(n^{1-k}). \quad (4.2)$$

Observe that the expected number of edges equals cn . We call $\mathcal{G}(n, k, cn, \sigma)$ the *planted coloring model*.

Lemma 4.2. *Let $\sigma : [n] \rightarrow [q]$ be a fixed balanced map. For any event \mathcal{E} we have*

$$\mathbf{P}[\mathcal{G} \in \mathcal{E} | \sigma \text{ is a coloring of } \mathcal{G}] \leq O(\sqrt{n}) \mathbf{P}[\mathcal{G}_\sigma \in \mathcal{E}].$$

Proof. By Stirling's formula, the probability that \mathcal{G}_σ has precisely m edges is $\Theta(n^{-1/2})$. If this event occurs then the conditional distributions of \mathcal{G}_σ and of \mathcal{G} coincide. \square

Hence, we are left to show that the probability that σ fails to be tame in \mathcal{G}_σ is $o(n^{-1/2})$. Indeed, in Sections 4.2 and 4.3 we will establish the following two statements. In both cases the proofs are by careful generalisation of the arguments from [9] to the hypergraph case.

Lemma 4.3. *With probability $1 - \exp(-\Omega(n))$ the planted coloring σ is separable in $\mathcal{G}(n, k, cn, \sigma)$.*

Lemma 4.4. *With probability $1 - o(n^{-1/2})$ we have $|\mathcal{E}(\mathcal{G}_\sigma, \sigma)| \leq \mathbf{E}[Z_{q,\text{bal}}]$.*

Proposition 3.5 is immediate from Lemmas 4.2–4.4.

Much of the analysis in this section will involve random variables defined using the following edge counts. For sets $X_1, X_2, X_3 \subset [n]$ and $\alpha \in [k]$, we let $m_\alpha(X_1, X_2, X_3)$ be the number of edges e of \mathcal{G}_σ such that there exists $x \in X_1$ and $v_1, \dots, v_\alpha \in X_2$ distinct from one another and from x , such that $x, v_1, \dots, v_\alpha \in e$ and $e \setminus \{x, v_1, \dots, v_\alpha\} \subseteq X_3$. If $\alpha = k - 1$ then we write $m_{k-1}(X_1, X_2)$ instead of $m_{k-1}(X_1, X_2, X_3)$, since X_3 has no effect in this case. For ease of notation, if $X_1 = \{v\}$ we simply write $m_\alpha(v, X_2, X_3)$, or $m_{k-1}(v, X_2)$. We set $V_i = \sigma^{-1}(i)$ to ease the notational burden. The following lemmas will be useful later. Recall that $\kappa = q^{1-k} \ln^{20} q$, as in the definition of separability.

Lemma 4.5. *For all sets $A, B \subseteq [n]$ such that $|A|, |B| \leq \frac{nk}{e}$ we have $m_1(A, B, V_i) < 20k(|A| + |B|)$ with probability $1 - O(1/n)$.*

Proof. Fix sets A, B such that $|A| = a, |B| = b$. Using (4.1) and (4.2), there exists a constant $\lambda \leq \kappa^{-2/3}$ such that

$$\mathbb{P}[m_1(A, B, V_i) > 20k(|A| + |B|)] \leq \binom{ab \binom{|V_i|}{k-2}}{20k(a+b)} p^{20k(a+b)} \leq \left(\frac{epab|V_i|^{k-2}}{20k(a+b)(k-2)!} \right)^{20k(a+b)} \leq \left(n^{-1}\lambda \cdot \frac{ab}{a+b} \right)^{20k(a+b)}.$$

Summing over all choices for A, B , it follows from the union bound that the probability that any such pair of sets exists is at most

$$\sum_{a,b=1}^{nk/e} \binom{n}{a} \binom{n}{b} \left(n^{-1}\lambda \cdot \frac{ab}{a+b} \right)^{20k(a+b)} \leq \sum_{a,b=1}^{nk/e} \left(\frac{ne}{a} \right)^a \left(\frac{ne}{b} \right)^b \left(n^{-1}\lambda \cdot \frac{ab}{a+b} \right)^{20k(a+b)}.$$

Now

$$\left(\frac{ne}{a} \right)^a \left(\frac{ne}{b} \right)^b \left(n^{-1}\lambda \cdot \frac{ab}{a+b} \right)^{k(a+b)} = \left[\frac{ne}{a} \left(n^{-1}\lambda \cdot \frac{ab}{a+b} \right)^k \right]^a \left[\frac{ne}{b} \left(n^{-1}\lambda \cdot \frac{ab}{a+b} \right)^k \right]^b$$

$$\begin{aligned}
&\leq \left[\frac{ne}{a} (n^{-1}\lambda \cdot a)^k \right]^a \left[\frac{ne}{b} (n^{-1}\lambda \cdot b)^k \right]^b \\
&\leq \left[e\lambda^k \cdot (\kappa/e)^{k-1} \right]^a \left[e\lambda^k \cdot (\kappa/e)^{k-1} \right]^b < 1.
\end{aligned}$$

Therefore

$$\sum_{a,b=1}^{n\kappa/e} \binom{n}{a} \binom{n}{b} \left(n^{-1}\lambda \cdot \frac{ab}{a+b} \right)^{20k(a+b)} \leq \sum_{a,b=1}^{n\kappa/e} \left(n^{-1}\lambda \cdot \frac{ab}{a+b} \right)^{19k(a+b)} \leq n\kappa \sum_{a=1}^{n\kappa/e} (n^{-1}\lambda)^{19ka} = O(1/n)$$

where the last equality follows since the summand is decreasing in a when $a \leq \frac{n\kappa}{e}$. \square

Lemma 4.6. *For every $S \subseteq V$, define*

$$B_S := \{v \in V : m_1(v, S \cap V_j, V_j) > 0 \text{ for some } j \neq \sigma(v)\}.$$

With probability $1 - \exp\{-\Omega(n)\}$, every set S of size $|S| \leq nq^{-9k}$ has $|B_S| \leq nq^{-6k}$.

Proof. Fix a subset $S \subseteq V$ of size at most nq^{-9k} and take $v \in V \setminus S$, and some $j \neq \sigma(v)$. Now $m_1(v, S \cap V_j, V_j)$ is stochastically dominated by $\text{Bin}\left(|S| \binom{|V_j|}{k-2}, p\right)$. Therefore, the union bound in conjunction with (4.1), (4.2) gives

$$\begin{aligned}
\mathbf{P}[m_1(v, S \cap V_j, V_j) > 0 \text{ for some } j \neq \sigma(v)] &\leq 1 - q \cdot \exp\left\{-p \cdot |S| \cdot \binom{\max_{j \neq \sigma(v)} |V_j|}{k-2}\right\} \\
&\leq 1 - q \cdot \exp\left\{-(1 + o(1)) \cdot \frac{k! \cdot q^{k-1} \ln q}{n^{k-1}(1-q^{k-1})} \cdot \frac{(n/q)^{k-2}}{(k-2)!} \cdot \frac{n}{q^{9k}}\right\} \leq q^{-8k}.
\end{aligned}$$

With B_S defined above, it follows that $|B_S|$ is stochastically dominated by $\text{Bin}(n, q^{-8k})$, and so from the Chernoff bound (see Lemma 4.1), we have

$$\mathbf{P}[|B_S| > nq^{-7k}] \leq \exp\{-nq^{-7k} \ln(q^k/e)\} \leq \exp\{-nq^{-7k}\}.$$

Finally, for $\alpha \leq q^{-9k}$ let X_α be the number of sets S of size αn such that $|B_S| \geq nq^{-7k}$. Then

$$\mathbf{P}[X_\alpha > 0] \leq \mathbf{P}[X_{q^{-9k}} > 0] \leq \binom{n}{q^{-9k}n} \cdot \exp\{-nq^{-7k}\} \leq \exp\left\{-n\left(q^{-9k}(q^{2k} - qk \ln q - 1)\right)\right\}.$$

Therefore for sufficiently large q we have $\mathbf{P}[X_\alpha > 0] \leq \exp\{-\Omega(n)\}$. The claim follows from taking the union bound over all $\alpha \leq q^{-9k}$ such that αn is an integer: the number of terms in the summation is linear and so is absorbed by the exponential small probability. \square

4.2 Separability: proof of Lemma 4.3

Let $\tau : [n] \rightarrow [q]$ be a balanced map which is not separable: that is, for which there exist $i, j \in [q]$ such that (3.15) is violated. Of course, we may assume without loss that $i = j = 1$. We aim to show that τ is unlikely to be a coloring of \mathcal{G}_σ . Clearly, if τ is a coloring of \mathcal{G}_σ then $\tau^{-1}(1)$ is an independent set of size about n/q that has a rather substantial intersection with the independent set $\sigma^{-1}(1)$. Here, as for graphs, an independent set is a set of vertices which contains no edge. The following lemma rules this constellation out for a wide range of intersection sizes.

Lemma 4.7. *With probability $1 - \exp(-\Omega(n))$ the hypergraph \mathcal{G}_σ has no independent set I of order $(1+o(1))\frac{n}{q}$ such that*

$$n^{-1}|I \cap \sigma^{-1}(1)| \in \left(q^{-1}(1.01/k)^{1/(k-1)}, q^{-1}(1 - q^{(1.01-k)/2})\right).$$

Proof. Suppose that I is an independent set with $|I| = \frac{n}{q}(1 + o(1))$ such that $S = I \cap \sigma^{-1}(1)$ contains $|S| = \frac{sn}{q}$ vertices, for some $s \in (0, 1)$. Then the set

$$V_0(S) := \{v \in V \setminus \sigma^{-1}(1) : m_{k-1}(v, S) = 0\}$$

contains $I \setminus S$. Observe that

$$\mathbf{P}[m_{k-1}(v, S) = 0] = \exp\left\{-p \binom{|S|}{k-1}\right\} \cdot (1 + O(1/n)) = \exp\left\{-\frac{kc(s/q)^{k-1}}{(1-q^{1-k})}\right\} \cdot (1 + O(1/n)) \leq 2q^{-ks^{k-1}}.$$

Let $n_0(S) := |V_0(S)|$, and observe that in order for I to exist, the inequality $n_0(S) > (1-s+o(1))\frac{n}{q}$ must hold. Thus it suffices to prove that when $(1.01/k)^{1/(k-1)} < s < 1 - q^{(1.01-k)/2}$, with probability $1 - \exp(-\Omega(n))$ there is no subset $S \subseteq \sigma^{-1}(1)$ of size sn/q with $n_0(S) > (1-s+o(1))\frac{n}{q}$.

Since $n_0(S)$ is stochastically dominated by $\text{Bin}(|V \setminus \sigma^{-1}(1)|, 2q^{-ks^{k-1}})$, we have by the Chernoff bound (see Lemma 4.1) that

$$\mathbf{P}\left[n_0(S) > (1-s+o(1))\frac{n}{q}\right] \leq \exp\left\{-(1-s+o(1))\frac{n}{q} \ln\left(\frac{1-s}{2q^{1-ks^{k-1}}e}\right)\right\}.$$

The number of choices for a subset S of $\sigma^{-1}(1)$ of size sn/q equals

$$\binom{n/q(1+o(1))}{sn/q} \leq \left(\frac{e}{1-s}\right)^{(1-s+o(1))n/q} = \exp\left\{(1-s+o(1))\frac{n}{q}(1-\ln(1-s))\right\},$$

as established in [9, equation (A.5)]. Hence, by the union bound over S , the probability that such a subset S exists with the desired lower bound on $n_0(S)$ is at most

$$\exp\left\{-(1-s+o(1))\frac{n}{q} \left(\ln\left(\frac{1-s}{2q^{1-ks^{k-1}}e}\right) - 1 + \ln(1-s)\right)\right\} = \exp\left\{(1-s+o(1))\frac{n}{q} \ln\left(\frac{2e^2}{q^{ks^{k-1}-1}(1-s)^2}\right)\right\}.$$

This probability tends to zero if and only if

$$\frac{\sqrt{2}e}{q^{(1-ks^{k-1})/2}} < 1-s. \quad (4.3)$$

By convexity, the exponential function on the l.h.s. intersects the linear function on the r.h.s. at most twice, and between these two points of intersection the linear function is largest. For sufficiently large q , explicit calculation shows that the values $s = (1.01/k)^{1/(k-1)}$ and $s = 1 - q^{(1.01-k)/2}$ satisfy (4.3). Therefore, for fixed s such that $(1.01/k)^{1/(k-1)} \leq s \leq 1 - q^{(1.01-k)/2}$, the probability that the set S exists is bounded by $\exp(-\Omega(n))$. Finally, as there are only linearly many such choices for s that make sn/q an integer, this completes the proof. \square

Lemma 4.7 does not quite cover the entire interval of intersections required by (3.15). To rule out the remaining subinterval $(q^{-1}(1 - q^{(1-k)/2}), q^{-1}(1 - \kappa))$ we use an expansion argument. The starting point is the observation that most vertices that have color 1 under τ but not under σ are likely to occur in a good number of edges in which *all* the $k-1$ other vertices are colored 1 under σ . We have not attempted to optimise the constants in this lemma.

Lemma 4.8. *Let $\tau : [n] \rightarrow [q]$ be a balanced map such that $a_{11}(\sigma, \tau) \in (q^{-1}(1 - q^{(1.01-k)/2}), q^{-1}(1 - \kappa))$. With probability $1 - \exp(-\Omega(n))$, the random hypergraph $\mathcal{G}_\sigma \in \mathcal{G}(n, k, cn, \sigma)$ has the following properties:*

1. *The set $Y := \{v \in V \setminus \sigma^{-1}(1) : m_{k-1}(v, \sigma^{-1}(1)) < 15\}$ has size at most $nk/(3q)$.*
2. *The set $U := \tau^{-1}(1) \setminus (\sigma^{-1}(1) \cup Y)$ satisfies $m_1(U, \sigma^{-1}(1) \setminus \tau^{-1}(1), \sigma^{-1}(1)) \leq 5|\sigma^{-1}(1) \setminus \tau^{-1}(1)|$.*

Proof. By assumption, $|\sigma^{-1}(1) \cap \tau^{-1}(1)| = \frac{sn}{q}$ where $s \in (1 - q^{(1.01-k)/2}, 1 - \kappa)$. Fix $v \in V \setminus V_1$. We know that

$$m_{k-1}(v, V_1) \sim \text{Bin}\left(\binom{|V_1|}{k-1}, p\right).$$

Therefore

$$\mathbf{P}[m_{k-1}(v, V_1) < 15] \leq \sum_{j=0}^{14} \binom{\binom{|V_1|}{k-1}}{j} p^j (1-p)^{\binom{|V_1|}{k-1}-j} \leq (1-p)^{\binom{|V_1|}{k-1}-14} \sum_{j=0}^{14} \frac{\left(\binom{|V_1|}{k-1} p\right)^j}{j!}.$$

Combining (4.2) with the lower bound from (4.1) shows that $\binom{|V_1|}{k-1} p > k \ln q$, which in turn implies that

$$\mathbf{P}[m_{k-1}(v, V_1) < 15] \leq 3(k \ln q)^{14} q^{-k}.$$

As the event $\{m_{k-1}(v, V_1) < 15\}$ occurs independently for all $v \in V \setminus V_1$, the total number Y of such vertices is stochastically dominated by $\text{Bin}(n(1-1/q), 3(k \ln q)^{14} q^{-k})$. Therefore $\mathbf{E}[Y] \leq n \cdot 3(k \ln q)^{14} q^{-k}$. Finally, by the Chernoff bound (see Lemma 4.1) and using the definition of κ from (3.15),

$$\mathbf{P}[Y > n\kappa/(3q)] \leq \exp\{-n\kappa/(3q)\} = \exp\{-\Omega(n)\}$$

and so the proof of (i) is complete.

For notational convenience, we write $R = \sigma^{-1}(1) \setminus \tau^{-1}(1)$ and $T = \tau^{-1}(1) \setminus \sigma^{-1}(1)$. Observe that $m_1(U, R, V_1)$ is stochastically dominated by $m_1(T, R, V_1)$, since U is a subset of T . Furthermore, $m_1(T, R, V_1)$ is stochastically dominated by $\text{Bin}\left(|T||R| \binom{|V_1|}{k-2}, p\right)$. Therefore since q is large with respect to k and $c < q^{k-1} \ln q$, it follows that

$$\binom{|V_1|}{k-2} p \leq (1+o(1)) \cdot \left(\frac{n}{q}\right)^{k-2} \cdot \frac{ck(k-1)}{n^{k-1}(1-q^{1-k})} \leq (1+o(1)) \cdot n^{-1} q \ln q \cdot \frac{k(k-1)}{(1-q^{1-k})} \leq n^{-1} k^2 q \ln q,$$

and so

$$\begin{aligned} \mathbf{E}[m_1(T, R, V_1)] &\leq |T||R| \binom{|V_1|}{k-2} p \leq (1+o(1)) |R| \cdot \frac{n(1-s)}{q} \cdot n^{-1} k^2 q \ln q \\ &= (1+o(1)) |R| \cdot (1-s) k^2 \ln q. \end{aligned}$$

Finally, as $\kappa \leq 1 - s \leq q^{(1.01-k)/2}$, part (ii) follows from the Chernoff Bound. \square

Proof of Lemma 4.3. Suppose that τ is a balanced map such that $a_{11}(\sigma, \tau) > q^{-1}(1.01/k)^{1/(k-1)}$. By Lemma 4.7, we may assume that $a_{11}(\sigma, \tau) > q^{-1}(1 - q^{1.01-k/2})$. With U, Y as in Lemma 4.8 we have that

$$15|U| \leq m_1(U, R, V_1) \leq 5|\sigma^{-1}(1) \setminus \tau^{-1}(1)|,$$

and so $|U| \leq \frac{1}{3}|\sigma^{-1}(1) \setminus \tau^{-1}(1)| \sim \frac{n}{3q} - \frac{1}{3}|\sigma^{-1}(1) \cap \tau^{-1}(1)|$. Since τ is balanced, we have

$$\frac{n}{q} \sim |\tau^{-1}(1)| \leq |\sigma^{-1}(1) \cap \tau^{-1}(1)| + |U| + |Y|.$$

Substituting our bound on $|U|$ from above, and using Lemma 4.8, implies that $na_{11}(\sigma, \tau) = |\sigma^{-1}(1) \cap \tau^{-1}(1)| > n(1-\kappa)/q$, as required. (The failure probability $e^{-\Omega(n)}$ equals the sum of the failure probabilities from Lemma 4.7 and Lemma 4.8.) \square

4.3 The cluster size: proof of Lemma 4.4

To upper bound the cluster size we will exhibit a large “core” of vertices of \mathcal{G}_σ that are difficult to recolor. More specifically, the core will consist of vertices v such that for every color $i \neq \sigma(v)$ there are several edges e containing v such that $e \setminus \{v\} \subset V_i$ and such that all vertices of e belong to the core. Therefore, if we attempt to change the color of v to $i \neq \sigma(v)$, then it will be necessary to recolor several other vertices of the core. In other words, recoloring a single vertex in the core leads to an avalanche that will stop only once at least $nq^{-1}(1.01/k)^{1/(k-1)}$ vertices in some color class have been recolored. Hence, the outcome is a coloring that does not belong to $\mathcal{C}(\mathcal{G}_\sigma, \sigma)$.

Formally, given a fixed balanced map σ and fixed hypergraph G , the core V_{core} of G is defined as the largest subset $V' \subseteq [n]$ of vertices such that

$$m_{k-1}(v, V_i \cap V') \geq 100k \text{ for all } v \in V' \text{ and all } i \neq \sigma(v).$$

The core is well-defined; for if V', V'' are sets with the property, then so is $V' \cup V''$.

Lemma 4.9. *With probability $1 - o(n^{-1/2})$ the random hypergraph \mathcal{G}_σ has the following two properties:*

- (i) *The core of \mathcal{G}_σ contains at least $(1 - q^{1-k} \ln^{500k} q)n$ vertices.*
- (ii) *If τ is a balanced coloring of \mathcal{G}_σ such that $\tau(v) \neq \sigma(v)$ for some v in the core, then $\tau \notin \mathcal{C}(\mathcal{G}_\sigma, \sigma)$.*

We proceed to prove Lemma 4.9. To estimate the size of the core we consider the following process:

CR1 For $i, j \in [q]$ and $i \neq j$, let $W_{ij} = \{v \in V_i : m_{k-1}(v, V_j) < 300k\}$, $W_i = \cup_{j:j \neq i} W_{ij}$, $W = \cup_i W_i$.

CR2 For $i \neq j$ let $U_{ij} = \{v \in V_i : m_1(v, W_j, V_j) > 100k\}$, and $U = \cup_{i \neq j} U_{ij}$.

CR3 Set $Z^{(0)} = U$ and repeat the following for $\ell \in \mathbb{N}$,

- if there is a $v \in V \setminus Z^{(\ell)}$ such that $m_1(v, Z^{(\ell)}, V_j) > 100k$ for some $j \neq \sigma(v)$ then take one such v and let $Z^{(\ell+1)} = Z^{(\ell)} \cup \{v\}$;
- otherwise, set $Z^{(\ell+1)} = Z^{(\ell)}$.

Let $Z = \cup_{\ell \geq 0} Z^{(\ell)}$ be the final set resulting from **CR3**.

Claim 4.10. *The set $V \setminus (W \cup Z)$ is contained within the core.*

Proof. For a contradiction, let $v \in V \setminus (W \cup Z)$. Since $W_j \subseteq V_j$, any edge counted by $m_{k-1}(v, V_j)$ that does not contribute to $m_1(v, W_j, V_j)$ must have empty intersection with W_j . Since $m_{k-1}(v, V_j) \geq 300k$ but $m_1(v, W_j, V_j) \leq 100k$, we must have that $m_{k-1}(v, V_j \setminus W_j) \geq 200k$. Similarly, since $v \notin Z$ we have

$$m_1(v, Z \cap V_j, V_j) \leq m_1(v, Z, V_j) \leq 100k,$$

and therefore $m_{k-1}(v, V_j \setminus (W_j \cup Z)) \geq 100k$. Furthermore, this statement holds for all $j \neq \sigma(v)$ and all $v \in V \setminus (W \cup Z)$. It follows that the entire set $V \setminus (W \cup Z)$ may be added to the core, which contradicts maximality unless $V \setminus (W \cup Z) \subseteq V_{\text{core}}$, as required. \square

We now bound the size of W, U and Z .

Claim 4.11. *Define the function $Q(q, k) = q^{-k-1} \ln^{400k} q$. With probability at least $1 - \exp\{-\Omega(n)\}$ we have $|W_{ij}| \leq n \cdot Q(q, k)$ for all distinct $i, j \in [q]$.*

Proof. Fix $v \in V_i$. Due to the independence of edges in $\mathcal{G}(n, k, cn, \sigma)$ we know that $m_{k-1}(v, V_j)$ is distributed binomially with mean $\binom{|V_j|}{k-1} p(1 + o(1)) \geq k \ln q + O_q(q^{-1})$. It follows from Lemma 4.1 that $\mathbf{P}(v \in W_{ij}) \leq \frac{q}{3} \cdot Q(q, k)$ for $v \in V_i$ and sufficiently large q . Therefore $\mathbf{E}[|W_{ij}|] \leq \frac{n}{3} \cdot Q(q, k)$. Finally, since $|W_{ij}|$ is distributed binomially, a straightforward application of the Chernoff bound shows that $\mathbf{P}[|W_{ij}| \geq n \cdot Q(q, k)] \leq \exp\{-n \cdot Q(q, k) \ln(3/e)\} = \exp\{-\Omega(n)\}$. \square

Claim 4.12. *We have $|U| \leq n/q^{10k}$ with probability at least $1 - \exp\{-\Omega(n)\}$.*

Proof. Fix $v \in V_i$. The quantity $m_1(v, W_j, V_j)$ is stochastically dominated by $\text{Bin}\left(|W_j| \binom{|V_j|}{k-2}, p\right)$. Hence, with $Q(q, k)$ as previously, we know that

$$\mathbf{E}\left[m_1(v, W_j, V_j) \mid |W_j| \leq n \cdot qQ(q, k)\right] \leq nk p \binom{|V_j|}{k-2} \cdot qQ(q, k) = \tilde{O}_q(q^{1-k}).$$

Applying the Chernoff bound gives $\mathbf{P}[v \in U_{ij} \mid |W_j| \leq n \cdot qQ(q, k)] \leq \tilde{O}_q(q^{-19k})$. Then $|U_{ij}|$, conditional on the event $|W_j| \leq n \cdot qQ(q, k)$, is stochastically dominated by a binomial random variable with mean $n \cdot O_q(q^{-15k})$. The Chernoff bound implies that

$$\mathbf{P}\left[|U_{ij}| > nq^{-10k} \mid |W_j| \leq n \cdot qQ(q, k)\right] \leq \exp\{-\Omega(n)\}.$$

The result follows by Claim 4.11. \square

Claim 4.13. *We have $|Z| \leq n/q^{9k}$ with probability at least $1 - \exp\{-\Omega(n)\}$.*

Proof. Claim 4.12 tells us that $|U| \leq n/q^{10k}$ with probability $1 - \exp\{-\Omega(n)\}$. We will condition on this event. Suppose that $|Z \setminus U| \geq i^* = n/q^{10k}$ and consider the set $Z^{(i^*)}$ obtained after i^* steps of **CR3**. The construction of Z implies that there exists $100k|Z^{(i^*)} \setminus U|$ vertex-edge pairs (v, e) such that $e \cap Z \geq 2$ and $e \setminus \{v\} \subseteq V_j$ for some $j \in [q]$. Since each edge may appear in at most k vertex-edge pairs, this implies that there are at least $100|Z^{(i^*)} \setminus U|$ such edges. Therefore, there are at least $100i^* = 100n/q^{10k}$ edges e such that $e \cap Z^{(i^*)} \geq 2$ and $e \setminus \{v\} \subseteq V_j$ for some $j \in [q], v \in e$, despite the set $Z^{(i^*)}$ only being of size at most $2n/q^{10k}$. We prove that with high probability, no such set can exist.

Let $\alpha = q^{-10k}$ and let $T \subset [n]$ be a set of $|T| = \alpha n$ vertices. Let m_T^* be the number of edges e such that $e \cap T \geq 2$ and $e \setminus \{v\} \subseteq V_j$ for some $j \in [q], v \in V$. We know that m_T^* is stochastically dominated by $\text{Bin}\left(2\alpha \binom{n}{2} \binom{n/q}{k-2}, p\right)$, and so we may observe by the Chernoff bound that

$$\mathbf{P}\left[m_T^* \geq 100 \cdot \alpha n\right] \leq \exp\{100\alpha n \ln \alpha\}.$$

If we let N be the number of sets T of size $|T| = \alpha n$ such that $m_T^* \geq 100 \cdot \alpha n$, then

$$\mathbf{P}[N > 0] \leq \binom{n}{\alpha n} \exp\{100\alpha n \ln \alpha\} \leq \exp\left\{-\left(\alpha \ln \alpha + (1-\alpha) \ln(1-\alpha) - 100\alpha \ln \alpha\right)n\right\} = \exp\{-\Omega(n)\}.$$

The final bound holds since α is constant with respect to n and $\alpha \in (0, 1)$. Therefore with probability $1 - \exp\{-\Omega(n)\}$ we have $|Z \setminus U| \leq n/q^{10k}$, which implies the claim. \square

Lemma 4.9 (i) then follows from Claims 4.10–4.13.

To establish (ii) we say that a vertex v is j -blocked if there is an edge $e \ni v$ such that $e \setminus \{v\}$ is contained in the core and $e \setminus \{v\} \subset V_j$. We say that a vertex v is σ -complete if it is j -blocked for all $j \neq \sigma(v)$. Note that, as with vertices inside the core, recoloring any σ -complete vertex will set off a coloring avalanche.

Claim 4.14. *With probability $1 - O(1/n)$ the random graph \mathcal{G}_σ has the following property:*

$$\text{if } \tau \in \mathcal{C}(\mathcal{G}_\sigma, \sigma) \text{ then for all } \sigma\text{-complete vertices } v \text{ we have } \sigma(v) = \tau(v).$$

Proof. Note that it suffices to prove that $\sigma(v) = \tau(v)$ for all v in the core, since this implies the result for all σ -complete vertices outside the core as well, by definition of σ -complete.

Recalling Lemma 4.3, we may assume that σ is separable in $\mathcal{G}(n, k, cn, \sigma)$. For $i \in [q]$, let

$$\Delta_i^+ = \{v \in V_{\text{core}} : \tau(v) = i \neq \sigma(v)\}, \quad \Delta_i^- = \{v \in V_{\text{core}} : \tau(v) \neq i = \sigma(v)\}.$$

Then

$$\sum_{i=1}^q |\Delta_i^+| = |\{v \in V_{\text{core}} : \sigma(v) \neq \tau(v)\}| = \sum_{i=1}^q |\Delta_i^-|. \quad (4.4)$$

Since σ is separable and $\tau \in \mathcal{C}(\mathcal{G}_\sigma, \sigma)$ we have $\max_{i \in [q]} |\Delta_i^+| \leq \frac{n}{q} \kappa(1 + o(1))$ and $\max_{i \in [q]} |\Delta_i^-| \leq \frac{n}{q} \kappa(1 + o(1))$. If we can show that $\{v \in V_{\text{core}} : \sigma(v) \neq \tau(v)\} = \emptyset$ then $\sigma(v) = \tau(v)$ for all σ -complete vertices.

Take $v \in \Delta_i^+$. Since $v \in V_{\text{core}}$ we know that $m_{k-1}(v, V_i) \geq 100k$. Further, since τ is a coloring, we must have that $m_1(v, \Delta_i^-, V_i) \geq 100k$. By Lemma 4.5 we know that with probability $1 - O(1/n)$,

$$m_1(\Delta_i^+, \Delta_i^-, V_i) \leq 20k(|\Delta_i^+| + |\Delta_i^-|) \quad \text{for all } i \in [q]. \quad (4.5)$$

Observe that if (4.5) holds then for all $i \in [q]$,

$$100k|\Delta_i^+| \leq m_1(\Delta_i^+, \Delta_i^-, V_i) \leq 20k(|\Delta_i^-| + |\Delta_i^+|).$$

But this implies that $4|\Delta_i^+| \leq |\Delta_i^-|$ for all $i \in [q]$, which contradicts (4.4) unless $\Delta_i^+ = \Delta_i^- = \emptyset$ for all $i \in [q]$. Therefore we conclude that with probability $1 - O(1/n)$, for all v in the core we have $\sigma(v) = \tau(v)$, completing the proof. \square

Lemma 4.9 (ii) is immediate from Claim 4.14. The core size guaranteed by Lemma 4.9 is not quite big enough to deduce a good bound on the cluster size (due to the polylogarithmic factor). Recall that a vertex v is j -blocked if it is contained in an edge e such that $e \setminus \{v\}$ is contained in the core and $e \setminus \{v\} \subset V_j$. Further, we say that v is α -free if there are at least $\alpha + 1$ colors j (including $\sigma(v)$) such that v fails to be j -blocked. A careful study of how the vertices outside the core connect to those inside yields the following.

Lemma 4.15. *With probability $1 - \exp\{-\Omega(n)\}$ there exists a set A_W of vertices such that there are at most $nq^{1-k}(1 + O_q(q^{-2}))$ vertices outside of A_W which are 1-free and there are at most $nq^{-k}(1 + O_q(q^{-1}))$ vertices which belong to A_W or are 2-free.*

We proceed to prove Lemma 4.15. Let

$$A_i = \{v \in V \setminus V_i : m_{k-1}(v, V_i) = 0\},$$

and define

$$\begin{aligned} A_0 &= \bigcup_{i \in [q]} A_i, & A_{00} &= \bigcup_{i \neq j} (A_i \cap A_j), & A_Z &= \{v \in V : m_1(v, Z \cap V_i, V_i) > 0 \text{ for some } i \neq \sigma(v)\}, \\ A_W &= \{v \in V : m_{k-1}(v, V_i \setminus W_i) = 0 \text{ for some } i \neq \sigma(v)\} \setminus A_0. \end{aligned}$$

We claim that if v is 1-free then $v \in A_0 \cup A_Z \cup A_W$, and if v is 2-free then $v \in A_{00} \cup A_Z \cup A_W$. To see that this is the case, note that if v is 1-free then there is some $i \neq \sigma(v)$ such that there is no edge $e \ni v$ with $e \setminus \{v\} \subseteq V_i \cap V_{\text{core}}$. For a contradiction, suppose that $v \notin A_0 \cup A_W \cup A_Z$. Then there must be an edge $e' \ni v$ such that

$$e' \setminus \{v\} \subseteq V_i \setminus (W_i \cup Z) \quad \text{for some } i \neq \sigma(v).$$

However, we know from Claim 4.10 that $V_i \setminus (W_i \cup Z) \subseteq V_i \cap V_{\text{core}}$, giving the desired contradiction. The case for 2-free vertices is similar. Suppose that v is 2-free and $v \notin A_{00} \cup A_Z \cup A_W$. Since v is also 1-free, must have $v \in (A_0 \setminus A_{00}) \cap (A_Z \cup A_W)^c$. Since $v \in A_0 \setminus A_{00}$, there exists $i \neq \sigma(v)$ so that $m_{k-1}(v, V_j) > 0$ for all $j \notin \{\sigma(v), i\}$. That is to say, we have $v \in (A_j \cup A_Z \cup A_W)^c$ for all $j \notin \{\sigma(v), i\}$. This means that for all $j \notin \{\sigma(v), i\}$ there exists an edge $e' \ni v$ such that $e' \setminus \{v\} \subseteq V_j \setminus (W_j \cup Z)$. As above, we conclude that $e' \subseteq V_j \cap V_{\text{core}}$ and so v is j -blocked for all $j \notin \{\sigma(v), i\}$. Therefore v is not 2-free and so we have a contradiction.

Thus, to prove Lemma 4.15 it suffices to bound the size of the sets A_0, A_{00}, A_W, A_Z .

Claim 4.16. *We have $|A_0| \leq n/q^{k-1}$ and $|A_{00}| \leq n/q^{2k-2}$ with probability at least $1 - \exp\{-\Omega(n)\}$.*

Proof. Take $v \in V_j$, and $i \neq j$. Now $\mathbf{P}[m_{k-1}(v, V_i) = 0] < \exp\{-k \ln q\}$, and hence $\mathbf{P}[v \in A_0] \leq (q-1)q^{-k}$. It follows that $\mathbf{E}[|A_0|] := \mu < n \cdot (q-1)q^{-k}$. Since $\mathbf{P}[m_{k-1}(v, V_i) = 0] > \exp\{-(k+1) \ln q\}$ we must have that $\mu > nq^{-k}$ and so by the Chernoff bound

$$\mathbf{P}\left[|A_0| > n/q^{k-1}\right] \leq \exp\left\{-nq^{-k} \left[\frac{q}{q-1} \ln\left(\frac{q}{q-1}\right) - \frac{1}{q-1}\right]\right\} = \exp\{-\Omega(n)\}$$

as desired. Further, the argument for A_{00} follows quickly after noting that the edge sets of $m_{k-1}(v, V_i)$ and $m_{k-1}(v, V_j)$ are independent for $i \neq j$. \square

Claim 4.17. *We have $|A_Z| \leq n/q^{6k}$ with probability at least $1 - \exp\{-\Omega(n)\}$.*

Proof. The proof follows immediately from Lemma 4.6 and Claim 4.13. \square

Claim 4.18. *We have $|A_W| \leq n/q^k$ with probability at least $1 - \exp\{-\Omega(n)\}$.*

Proof. Fix $i \neq j$ and $v \in V_i$. We seek to compute the following probability:

$$\mathbf{P}\left[m_{k-1}(v, V_j \setminus W_j) = 0 \text{ and } m_{k-1}(v, V_j) > 0\right] = \mathbf{P}\left[m_{k-1}(v, V_j \setminus W_j) = 0\right] \cdot \mathbf{P}\left[m_{k-1}(v, W_j) > 0\right].$$

Since $V_j \setminus W_j \subseteq V_j$, we know from the calculations in Claim 4.16 that $\mathbf{P}\left[m_{k-1}(v, V_j \setminus W_j) = 0\right] \leq q^{-k}$. Further, define the event \mathcal{E} that $|W_j| \leq nq \cdot Q(q, k)$ where $Q(q, k) = q^{-k-1} \ln^{400k} q$ as in Claim 4.11. We know that $m_{k-1}(v, W_j) | \mathcal{E}$ is stochastically dominated by $\text{Bin}\left(\binom{nq \cdot Q(q, k)}{k-1}, p\right)$ and so for sufficiently large q , we have

$$\mathbf{P}\left[m_{k-1}(v, W_j) > 0 | \mathcal{E}\right] \leq p \binom{nq \cdot Q(q, k)}{k-1} \leq \frac{p(nqQ(q, k))^{k-1}}{(k-1)!} \leq 2ck(qQ(q, k))^{k-1} \leq q^{-k(k-2)}.$$

Since Claim 4.11 implies that \mathcal{E} only fails with exponentially small probability, it follows that

$$\mathbf{P}\left[m_{k-1}(v, V_j \setminus W_j) = 0 \text{ and } m_{k-1}(v, V_j) > 0\right] \leq q^{-k} \left(q^{-k(k-2)} \cdot \mathbf{P}[\mathcal{E}] + \mathbf{P}[\neg \mathcal{E}]\right) = q^{-k^2+k} + \exp\{-\Omega(n)\}.$$

Taking the union bound over $j \in [q] \setminus \{i\}$ shows that for $v \in V_i$,

$$\mathbf{P}[v \in A_W] \leq q^{-k^2+k+1} + \exp\{-\Omega(n)\} < q^{-k-1}.$$

Therefore $|A_W|$ is stochastically dominated by $\text{Bin}(n, q^{-k-1})$, and applying Lemma 4.1 completes the proof. \square

Thus Lemma 4.15 follows from Claims 4.16-4.18.

Proof of Lemma 4.4. Assume that the properties described in Claim 4.14 and Lemma 4.15 both hold, noting that this is an event with probability $1 - o(n^{-1/2})$. The remainder of the proof is deterministic.

Since we have assumed that Claim 4.14 succeeds, for all σ -complete v and all $\tau \in \mathcal{C}(\mathcal{G}_\sigma, \sigma)$ we have $\tau(v) = \sigma(v)$. Let F_x be the set of x -free vertices. Next, by our assumption that Lemma 4.15 succeeds, we have

$$|F_1 \setminus A_W| \leq \frac{n}{q^{k-1}} + n \cdot O_q\left(q^{-k-1}\right), \quad |F_2 \cup A_W| = \frac{n}{q^k} + n \cdot O_q\left(q^{-k-1}\right).$$

For any $v \in F_x \setminus F_{x+1}$ there are at most $x+1$ choices for the color of v . Since $F_{x+1} \subseteq F_x$ it follows that

$$|\mathcal{C}(\mathcal{G}_\sigma, \sigma)| \leq 2^{|F_1 \setminus F_2|} 3^{|F_2 \setminus F_3|} \dots q^{|F_q|} \leq 2^{|F_1 \setminus A_W|} \cdot q^{|F_2 \cup A_W|},$$

and so

$$\frac{1}{n} \ln |\mathcal{C}(\mathcal{G}_\sigma, \sigma)| \leq \frac{\ln 2}{q^{k-1}} + \frac{\ln q}{q^k} + \tilde{O}_q(q^{-k-1}).$$

Furthermore, since $c \leq (q^{k-1} - 1/2) \ln q - \ln 2 - \frac{1.01 \ln q}{q}$, we have

$$\frac{1}{n} \ln \mathbf{E}[Z_{q, \text{bal}}] \geq \ln q + c \ln(1 - q^{1-k}) = \frac{\ln 2}{q^{k-1}} + \frac{1.01 \ln q}{q^k} + \tilde{O}_q(q^{-k-1}).$$

These bounds imply that $|\mathcal{C}(\mathcal{G}_\sigma, \sigma)| \leq \mathbf{E}[Z_{q, \text{bal}}]$, completing the proof. \square

4.4 Proof of Corollary 3.7

Here we assume that

$$(q^{k-1} - 1/2) \ln q - \ln 2 + 1/\ln q < c < (q^{k-1} - 1/2) \ln q.$$

The proof of Corollary 3.7 is similar to the proof of [9, Proposition 2.1]. The starting point is the following observation, which is reminiscent of the “planting trick” from [1]. Call $\sigma : [n] \rightarrow [q]$ ε -balanced for some $\varepsilon > 0$ if $\max_{i \in [q]} |\sigma^{-1}(i)| - n/q < \varepsilon n$.

Claim 4.19. *Suppose there exist $\varepsilon, \varepsilon' > 0$ and a sequence $(\mathcal{E}_n)_n$ of events such that for large n and all ε -balanced $\sigma : [n] \rightarrow [q]$ we have*

$$\mathbf{P}[\mathcal{G}_\sigma \in \mathcal{E}_n] \leq \exp(-\varepsilon' n), \quad (4.6)$$

$$\lim_{n \rightarrow \infty} \mathbf{P}[\mathcal{G} \in \mathcal{E}_n] = 1. \quad (4.7)$$

Then there exists $\delta > 0$ such that w.h.p. $Z_q(\mathcal{G}) \leq \exp(-\delta n) \mathbf{E}[Z_q(\mathcal{G})]$.

Proof. Let $Z'_q(G)$ be the number of ε -balanced q -colorings of G . By [13, proof of Lemma 2.1] there exists $\alpha > 0$ such that

$$\mathbf{E}[Z_q(\mathcal{G}) - Z'_q(\mathcal{G})] \leq \exp(-\alpha n) \mathbf{E}[Z_q(\mathcal{G})]. \quad (4.8)$$

Further, let $Z''_q(\mathcal{G}) = Z'_q(\mathcal{G}) \mathbf{1}\{\mathcal{G} \in \mathcal{E}_n\}$. Combining Lemma 4.2 and (4.6) shows that

$$\mathbf{E}[Z''_q(\mathcal{G})] \leq \exp(-\varepsilon' n/2) \mathbf{E}[Z'_q(\mathcal{G})]. \quad (4.9)$$

Moreover, let $\mathcal{A}_n = \{Z_q(\mathcal{G}) \geq \exp(-\delta n) \mathbf{E}[Z_q(\mathcal{G})]\}$ for a small enough $\delta > 0$. Combining (4.8) and (4.9), we obtain

$$\begin{aligned} \exp(-\delta n) \mathbf{E}[Z_q(\mathcal{G})] \mathbf{P}[\mathcal{G} \in \mathcal{A}_n \cap \mathcal{E}_n] &\leq \mathbf{E}[Z_q(\mathcal{G}) \mathbf{1}\{\mathcal{G} \in \mathcal{A}_n \cap \mathcal{E}_n\}] \\ &\leq \mathbf{E}[Z''_q(\mathcal{G})] + \mathbf{E}[Z_q(\mathcal{G}) - Z'_q(\mathcal{G})] \leq (\exp(-\varepsilon' n/2) + \exp(-\alpha n)) \mathbf{E}[Z_q(\mathcal{G})]. \end{aligned}$$

Hence, choosing $\delta > 0$ small enough and recalling (4.7), we obtain $\mathbf{P}[\mathcal{A}_n] = o(1)$. \square

Thus, we are left to exhibit a sequence of events as in Claim 4.19. Given a map $\tau : [n] \rightarrow [q]$ and a hypergraph G on $[n]$ let $E_\tau(G)$ be the number of monochromatic edges of G under τ . Further, for $\beta > 0$ let

$$Z_{q,\beta}(G) = \sum_{\tau} \exp(-\beta E_\tau(G)),$$

where the sum ranges over all $\tau : [n] \rightarrow [q]$. The function $\beta \mapsto Z_{q,\beta}(G)$ can be viewed as the partition function of a hypergraph variant of the “Potts antiferromagnet” from statistical physics. We consider this random variable because it is concentrated in the following sense.

Claim 4.20. *For any $\varepsilon > 0$ there is $\delta > 0$ such that for any $\sigma : [n] \rightarrow [q]$ we have*

$$\mathbf{P}[|\ln Z_{q,\beta}(\mathcal{G}) - \mathbf{E} \ln Z_{q,\beta}(\mathcal{G})| > \varepsilon n] < \exp(-\delta n), \quad \mathbf{P}[|\ln Z_{q,\beta}(\mathcal{G}_\sigma) - \mathbf{E} \ln Z_{q,\beta}(\mathcal{G}_\sigma)| > \varepsilon n] < \exp(-\delta n).$$

Proof. Either adding or removing a single edge alters the value of $\ln Z_{q,\beta}$ by at most β . Therefore, the assertion follows from a standard application of Azuma’s inequality. \square

Additionally, we have the following estimate of $\mathbf{E} \ln Z_{q,\beta}(\mathcal{G}_\sigma)$.

Claim 4.21. *There is $\delta > 0$ such that for all $\beta > 0$ and all δ -balanced σ we have $\mathbf{E} \ln Z_{q,\beta}(\mathcal{G}_\sigma) > \delta n + \ln \mathbf{E}[Z_q(\mathcal{G})]$.*

Proof. We are going to show that for a small enough $\delta > 0$ we have w.h.p.

$$n^{-1} \ln Z_q(\mathcal{G}_\sigma) \geq q^{1-k} \ln 2 + \tilde{O}_q(q^{-k}). \quad (4.10)$$

Since $Z_{q,\beta}(\mathcal{G}_\sigma) \geq Z_q(\mathcal{G}_\sigma)$ for all β and because (3.4) implies that

$$n^{-1} \ln \mathbf{E}[Z_q(\mathcal{G})] \leq q^{1-k} (\ln 2 - (2 \ln q)^{-1}) + \tilde{O}_q(q^{-k}),$$

the claim follows from (4.10).

To prove (4.10), we let F_{ij} be the set of vertices $v \in V_i$ such that $m_{k-1}(v, V_j) = 0$. Further, let F'_{ij} be the set of all $v \in F_{ij}$ such that $m_{k-1}(v, V_h) = 0$ for some $h \in [q] \setminus \{i, j\}$. Due to the independence of the edges, $|F_{ij}|, |F'_{ij}|$ are binomial random variables. The expected sizes of these sets satisfy

$$\mathbf{E}|F_{ij}| = (q^{-k-1} + \tilde{O}_q(q^{-k-2}))n, \quad \mathbf{E}|F'_{ij}| = \tilde{O}_q(q^{-k-2})n.$$

Hence, the Chernoff bound implies that with probability $1 - \exp(-\Omega(n))$, for all i, j we have

$$|F_{ij}| = (q^{-k-1} + \tilde{O}_q(q^{-k-2}))n, \quad |F'_{ij}| = \tilde{O}_q(q^{-k-2})n. \quad (4.11)$$

Let $F_\star = \bigcup_{i \neq j} F_{ij} \setminus F'_{ij}$. Further, for every vertex $v \in F_\star$ let $\sigma_\star(v) \in [q]$ be the (unique) color such that $v \in F_{\sigma(v)\sigma_\star(v)}$. Further, let E_\star be the set of edges e of \mathcal{G}_σ such that there exist $v, w \in e \cap F_\star$ such that

$$\sigma(e \setminus \{v, w\}) \subset \{\sigma(v), \sigma(w), \sigma_\star(v), \sigma_\star(w)\}.$$

The random variable $|E_\star|$ is stochastically dominated by a binomial random variable $\text{Bin}(cn, p_0)$ where

$$p_0 = \frac{2|F_\star|^2}{n^2} (4q^{-1})^{k-2}.$$

So by (4.11),

$$\mathbf{E}|E_\star| = 2(q^{1-k} + \tilde{O}_q(q^{-k}))^2 (4q^{-1})^{k-2} cn + \exp(-\Omega(n)) = \tilde{O}_q(q^{1-k})n.$$

Then, the Chernoff bound, we find that with probability $1 - \exp(-\Omega(n))$,

$$|E_\star| = \tilde{O}_q(q^{-k})n. \quad (4.12)$$

Now, let F_0 be the set of all vertices $v \in F_\star$ that do not occur in any $e \in E_\star$. Then by construction any map $\tau : [n] \rightarrow [q]$ such that $\tau(v) \in \{\sigma(v), \sigma_\star(v)\}$ for all $v \in F_0$ and $\tau(v) = \sigma(v)$ for all $v \notin F_0$ is a q -coloring of \mathcal{G}_σ . Furthermore, there are $2^{|F_0|}$ such τ and (4.11), (4.12) entail that $|F_0| \geq (q^{1-k} + \tilde{O}_q(q^{-k}))n$ w.h.p., whence (4.10) follows. \square

By comparison, $\ln \mathbf{E}[Z_{q,\beta}(\mathcal{G})]$ is upper-bounded as follows.

Claim 4.22. *For any $\delta > 0$ there is $\beta_0 > 0$ such that for all $\beta > \beta_0$ we have $\ln \mathbf{E}[Z_{q,\beta}(\mathcal{G})] \leq \delta n + \ln \mathbf{E}[Z_q(\mathcal{G})]$.*

Proof. Using (3.4) and the fact that monochromatic edges are least likely when τ is balanced, we obtain

$$\frac{1}{n} \ln \mathbf{E}[Z_q(\mathcal{G})] = \ln q + c \ln(1 - q^{1-k}) + o(1), \quad \frac{1}{n} \ln \mathbf{E}[Z_{q,\beta}(\mathcal{G})] \leq \ln q + c \ln(1 - q^{1-k}(1 - \exp(-\beta))).$$

Making β sufficiently large and taking logarithms, we obtain the assertion. \square

Finally, we know from Claims 4.21–4.22 and Jensen's inequality that there exists $\delta > 0$ such that $\mathbf{E} \ln Z_{q,\beta}(\mathcal{G}) + \delta n \leq \mathbf{E} \ln Z_{q,\beta}(\mathcal{G}_\sigma)$. However, Claim 4.20 implies that both $\ln Z_{q,\beta}(\mathcal{G})$ and $\ln Z_{q,\beta}(\mathcal{G}_\sigma)$ are close to their expectations. Therefore Corollary 3.7 follows by applying Claim 4.19 to the event

$$\mathcal{E}_n = \{G : |\ln Z_{q,\beta}(G) - \mathbf{E} \ln Z_{q,\beta}(\mathcal{G})| > \varepsilon n\}.$$

5 The second moment

In this section we prove Proposition 3.6. We keep the notation and the assumptions of Section 3 and Section 4.

5.1 Overview

We reduce the problem of estimating $\mathbf{E}[Z_{q,\text{tame}}^2]$ to that of optimising the function $F(a)$ from Lemma 3.3 over a certain domain $\mathcal{D}_{\text{tame}}$. Due to the additional constraints imposed by the ‘‘tame’’ condition, this domain $\mathcal{D}_{\text{tame}}$ is a relatively small subset of \mathcal{D} , which was the domain of optimisation for (3.14). In the end, $\max_{a \in \mathcal{D}_{\text{tame}}} F(a)$ will be seen to be significantly smaller than $\max_{a \in \mathcal{D}} F(a)$, and additionally, the problem of maximising F over $\mathcal{D}_{\text{tame}}$ technically less demanding.

To define $\mathcal{D}_{\text{tame}}$ formally, call $a \in \mathcal{D}$ *separable* if $a_{ij} \notin (q^{-1}(1.01/k)^{1/(k-1)}, q^{-1}(1-\kappa))$ for all $i, j \in [q]$ (cf. (3.15)). Additionally, we say that $a \in \mathcal{D}$ is *s-stable* if there are precisely s pairs (i, j) such that $a_{ij} > q^{-1}(1.01/k)^{1/(k-1)}$. We denote by \mathcal{D}_s the set of all s -stable $a \in \mathcal{D}$, and by $\mathcal{D}_{[q-1]} = \cup_{s < q} \mathcal{D}_s$. Geometrically, each \mathcal{D}_s is close to a $(q-s)$ -dimensional face of the Birkhoff polytope, for if a has entries greater than $(1.01/kq)^{1/(k-1)}$ then by separability these entries are in fact at least $(1-\kappa)/q$ (with $\kappa = \ln^{20} q/q^{k-1}$). Finally, let $\mathcal{D}_{\text{tame}}$ be the (compact) set of all $a \in \mathcal{D}$ that are separable and s -stable for some $0 \leq s < q$.

Lemma 5.1. *If $F(a) < F(\bar{a})$ for all $a \in \mathcal{D}_{\text{tame}} \setminus \{\bar{a}\}$ then $\mathbf{E}[Z_{q,\text{tame}}^2] = O(\mathbf{E}[Z_{q,\text{bal}}]^2)$.*

The proof of Lemma 5.1 is by a standard application of the Laplace method. We defer the details to Section 5.2.

In order to prove that $\max_{a \in \mathcal{D}_{\text{tame}}} F(a) = F(\bar{a})$, we observe that the set $\mathcal{D}_{\text{tame}}$ naturally decomposes into a number of disjoint subsets. Namely, let $\mathcal{D}_{s,\text{tame}}$ be the set of all s -stable $a \in \mathcal{D}_{\text{tame}}$ for $0 \leq s < q$. We will argue that for $1 \leq s < q$ the maximum of F over $\mathcal{D}_{s,\text{tame}}$ is not much greater than the function value attained at certain canonical points $\bar{a}(s)$ with entries

$$\bar{a}_{ij}(s) = q^{-1} \mathbf{1}\{i = j\} \mathbf{1}\{i \leq s\} + (q(q-s))^{-1} \mathbf{1}\{i > s\} \mathbf{1}\{j > s\}. \quad (5.1)$$

Hence, $\bar{a}(s)$ is a block-diagonal matrix. The upper-left block is the $s \times s$ identity matrix, divided by q , and the lower-right block is the $(q-s) \times (q-s)$ matrix with all entries equal to $(q(q-s))^{-1}$. Clearly, $\bar{a}(s) \in \mathcal{D}_{s,\text{tame}}$.

The following statement, which we prove in Section 5.3, is the heart of the second moment analysis.

Lemma 5.2. *We have $F(a) < F(\bar{a})$ for all $a \in \mathcal{D}_{\text{tame}} \setminus \{\bar{a}\}$.*

Proposition 3.6 follows immediately from Lemma 5.1 and Lemma 5.2.

5.2 The Laplace method: proof of Lemma 5.1

We seek to show that there exists some positive constant $C(q)$ such that

$$\mathbf{E}[Z_{q,\text{tame}}^2] \leq C(q) \cdot \mathbf{E}[Z_{q,\text{bal}}]^2. \quad (5.2)$$

The expected value of $Z_{q,\text{tame}}^2$ can be written as a sum over pairs of tame colourings. Define

$$\mathcal{E} = \{a \in \mathcal{R} \cap \mathcal{D}_{\text{tame}} : \|a - \bar{a}\|_k^k < \eta(q)\}.$$

We split $Z_{q,\text{tame}}^2$ into three components as follows:

$$Z_{q,\text{tame}}^2 = Z_{q,\text{tame}}^2 \cdot \mathbf{1}_{\mathcal{E}} + Z_{q,\text{tame}}^2 \cdot \mathbf{1}_{\mathcal{D}_{[q-1]} \setminus \mathcal{E}} + Z_{q,\text{tame}}^2 \cdot \mathbf{1}_{\mathcal{D}_q}.$$

First we estimate the contribution of the first summand above by performing a Taylor expansion of F around \bar{a} .

Lemma 5.3. *There exists $C(q)$ and $\eta(q)$ such that with we have*

$$\mathbf{E}[Z_{q,\text{tame}}^2 \cdot \mathbf{1}_{\mathcal{E}}] \leq C(q) \cdot \mathbf{E}[Z_{q,\text{bal}}]^2$$

Proof. We may parametrise $\mathcal{R} \cap \mathcal{D}_{\text{tame}}$ as follows: disregard the (q, q) entry and consider each matrix a as a $q^2 - 1$ dimensional vector. Let

$$\mathcal{L} : [0, 1/q]^{q^2-1} \longrightarrow [0, 1/q]^{q^2}, \quad a_{ij} \mapsto \begin{cases} a_{ij} & \text{if } (i, j) \neq (q, q), \\ 1 - \sum_{(i,j) \neq (q,q)} a_{ij} & \text{otherwise.} \end{cases}$$

We compute the Hessian of $F \circ \mathcal{L} = H \circ \mathcal{L} + E \circ \mathcal{L}$. For $(i, j) \neq (s, t)$ we have

$$\frac{\partial}{\partial a_{ij}} (H \circ \mathcal{L}(a)) \Big|_{a=\bar{a}} = 0, \quad \frac{\partial^2}{\partial a_{ij}^2} (H \circ \mathcal{L}(a)) \Big|_{a=\bar{a}} = -2q^2, \quad \frac{\partial^2}{\partial a_{ij} \partial a_{st}} (H \circ \mathcal{L}(a)) \Big|_{a=\bar{a}} = -q^2.$$

Further

$$\frac{\partial}{\partial a_{ij}} \|\mathcal{L}(a)\|_k^k \Big|_{a=\bar{a}} = 0, \quad \frac{\partial^2}{\partial a_{ij}^2} \|\mathcal{L}(a)\|_k^k \Big|_{a=\bar{a}} = \frac{2k(k-1)}{q^{2k-4}}, \quad \frac{\partial^2}{\partial a_{ij} \partial a_{st}} \|\mathcal{L}(a)\|_k^k \Big|_{a=\bar{a}} = \frac{k(k-1)}{q^{2k-4}},$$

and so

$$\begin{aligned} \frac{\partial}{\partial a_{ij}} (E \circ \mathcal{L}(a)) \Big|_{a=\bar{a}} &= 0, & \frac{\partial^2}{\partial a_{ij}^2} (E \circ \mathcal{L}(a)) \Big|_{a=\bar{a}} &= \frac{2ck(k-1)}{q^{2k-4}(1-q^{1-k})^2}, \\ \frac{\partial^2}{\partial a_{ij} \partial a_{st}} (E \circ \mathcal{L}(a)) \Big|_{a=\bar{a}} &= \frac{ck(k-1)}{q^{2k-4}(1-q^{1-k})^2}. \end{aligned}$$

Thus, we have that the first derivative of $F \circ \mathcal{L}$ vanishes at \bar{a} , and that the Hessian is

$$D^2(F \circ \mathcal{L}(a)) \Big|_{a=\bar{a}} = -q^2 \left(1 - \frac{2ck(k-1)}{q^{2(k-1)}(1-q^{1-k})^2} \right) (\text{id} + \mathbf{1})$$

where $\mathbf{1}$ is the matrix with all all entries equal to one, and id is the identity matrix. As id is positive definite, $\mathbf{1}$ is positive semidefinite and $c < q^{k-1} \ln q$ we have that the Hessian is negative definite at \bar{a} . Further, it follows from continuity that there exists some $\tilde{\eta}, \tilde{\xi}$ independent of n such that the largest eigenvalue of $D^2(F \circ \mathcal{L})$ is smaller than $-\tilde{\xi}$ for all points $\|a - \bar{a}\|_2 < \tilde{\eta}$. Since \mathcal{L} is linear, there exists some positive η , independent of n , such that for all a such that $\|a - \bar{a}\|_2 < \eta$ we have $\|\mathcal{L}^{-1} - \bar{a}\|_2 < \tilde{\eta}$. Taylor's theorem then implies that there is some positive ξ , independent of n , such that

$$F \circ \mathcal{L}(a) \leq F(\bar{a}) - \xi \sum_{(i,j) \neq (q,q)} (a_{ij} - q^{-2})^2 \quad \text{for all } a : \|a - \bar{a}\|_2 < \eta.$$

As \mathcal{E} satisfies the conditions required for the event \mathcal{A} in Lemma 3.3, we may apply (3.9) with $\mathcal{A} = \mathcal{E}$ to obtain

$$\begin{aligned} \mathbf{E}[Z_{q,\text{tame}}^2 \cdot \mathbf{1}_{\mathcal{E}}] &= \exp\{nF(\bar{a})\} \cdot O(n^{(1-q^2)/2}) \cdot \sum_{a \in \mathcal{E}} \exp\left\{-\xi n \sum_{(i,j) \neq (q,q)} (a_{ij} - q^{-2})^2\right\} \\ &\leq \exp\{nF(\bar{a})\} \cdot O(n^{(1-q^2)/2}) \cdot \int_{\mathbb{R}^{q^2-1}} \exp\left\{-\xi n \sum_{(i,j) \neq (q,q)} (z_{ij} - q^{-2})^2\right\} dz_{ij} \\ &\leq \exp\{nF(\bar{a})\} \cdot O(n^{(1-q^2)/2}) \cdot \left[\int_{\infty}^{\infty} \exp\{-\xi n z^2\} dz \right]^{q^2-1} \leq C(q) \cdot \mathbf{E}[Z_{q,\text{bal}}]^2 \end{aligned}$$

for some constant $C(q)$ depending only on q . Here the final inequality follows from (3.3). \square

There are two remaining cases to consider, namely $a \in \mathcal{D}_{[q-1]} \setminus \mathcal{E}$ and $a \in \mathcal{D}_q$. We begin with the latter.

Lemma 5.4. *There exists a constant $C(q) > 0$ such that*

$$\mathbf{E}[Z_{q,tame}^2 \cdot \mathbf{1}_{\mathcal{D}_q}] \leq C(q) \cdot \mathbf{E}[Z_{q,bal}]^2$$

Proof. Recall that $a(\sigma, \tau) \in \mathcal{D}_q$ if and only if there is a permutation of the colours of τ such that the resulting colouring is in $\mathcal{C}(\mathcal{G}, \sigma)$. Therefore

$$\begin{aligned} \mathbf{E}[Z_{q,tame}^2 \cdot \mathbf{1}_{\mathcal{D}_q}] &= \sum_{a(\sigma, \tau) \in \mathcal{D}_q} \mathbf{P}[\sigma, \tau \text{ are tame colourings}] \\ &\leq q! \sum_{\substack{\text{balanced} \\ \sigma: [n] \rightarrow [q]}} \mathbf{E}[|\mathcal{C}(\mathcal{G}, \sigma)| \mid \sigma \text{ is a tame colouring}] \cdot \mathbf{P}[\sigma \text{ is a tame colouring}] \\ &\leq q! \cdot \mathbf{E}[Z_{q,bal}] \sum_{\substack{\text{balanced} \\ \sigma: [n] \rightarrow [q]}} \mathbf{P}[\sigma \text{ is a tame colouring}] \leq q! \cdot \mathbf{E}[Z_{q,bal}]^2, \end{aligned} \quad [\text{by T3}],$$

as desired. □

Lemma 5.5. *If $F(a) < F(\bar{a})$ for all $a \in \mathcal{D}_{tame} \setminus \{\bar{a}\}$ then we have*

$$\mathbf{E}[Z_{q,tame}^2 \cdot \mathbf{1}_{\mathcal{D}_{[q-1]} \setminus \mathcal{E}}] \leq \mathbf{E}[Z_{q,bal}]^2.$$

Proof. We take η as in Lemma 5.3 and set

$$\mathcal{E}' = \{a \in \mathcal{R} \cap \mathcal{D}_{tame} : \|a - \bar{a}\|_2 \geq \eta\}.$$

As \mathcal{E}' is compact, the assumption that $F(a) < F(\bar{a})$ for all $a \in \mathcal{E}'$ additionally implies that there exists some γ such that $\max_{a \in \mathcal{E}'} F(a) < F(\bar{a}) - \gamma$. Then it follows from Lemma 3.2 and (3.13) that

$$\begin{aligned} \mathbf{E}[Z_{q,tame}^2 \cdot \mathbf{1}_{\mathcal{D}_{[q-1]} \setminus \mathcal{E}}] &\leq |\mathcal{E}'| \exp\{n(F(\bar{a}) - \gamma)\} \leq n^{q^2} \exp\{n(F(\bar{a}) - \gamma)\} \\ &\leq \exp\{n(F(\bar{a}) - \gamma/2)\} \leq \mathbf{E}[Z_{q,bal}]^2 \cdot \exp\{-n\gamma/3\} \leq \mathbf{E}[Z_{q,bal}]^2, \end{aligned}$$

as desired. □

Finally, (5.2) follows from combining Lemmas 5.3-5.5.

5.3 The maximisation problem: proof of Lemma 5.2

Throughout this subsection it is sufficient to assume that c equals the upper bound of (4.1), that is,

$$c = (q^{k-1} - 1/2) \ln q - \ln 2 - 1.01 q / \ln q.$$

To see this, suppose that Lemma 5.2 is true with this value of c . Then \bar{a} is the unique maximum of F on \mathcal{D}_{tame} . Now F is the sum of the concave function H and the convex function E , which attain their maximum, respectively minimum, at \bar{a} . Further, since H is independent of c and E is a linear multiple of c , decreasing the value of c only makes the maximum of F at \bar{a} more pronounced.

5.3.1 The strategy

The proof is based on the local variation technique developed in [9]. Roughly speaking, for each $0 < s < q$ we will argue that for any arbitrary $a \in \mathcal{D}_s$, we can move slightly toward a nicer matrix while increasing F . The new matrix that we produce is then regular enough that we may perform calculations and compare it the point $\bar{a}(s)$ whose first s diagonal entries are $1/q$, and whose (i, j) -entries are equal to $(q(q-s))^{-1}$ for $i, j > s$. As it turns out, $\bar{a}(s)$ comes

close enough to maximising F over \mathcal{D}_s (up to a negligible error term in each case). The final step is then to show that $F(\tilde{a}(s))$ is strictly less than $F(\tilde{a})$.

Let us take a moment to collect some results that will be used throughout the remainder of this section. In particular, it may come as no surprise that in a local variations argument we make extensive use of derivatives. Taking partials of F we have

$$\left(\frac{\partial}{\partial a_{ix}} - \frac{\partial}{\partial a_{iy}} \right) F(a) = \ln \frac{a_{iy}}{a_{ix}} + \frac{ck(a_{ix}^{k-1} - a_{iy}^{k-1})}{1 - 2/q^{k-1} + \|a\|_k^k}, \quad i, x, y \in [q]. \quad (5.3)$$

This represents the change in F when we increase a_{ix} at the expense of a_{iy} (see Lemma 5.7, which describes when the above quantity is positive). Further, we will often tackle the changes in entropy and energy separately. We need the following elementary inequalities (cf. [9, Corollary 4.10]). As usual, the entropy of a vector $b \in [0, 1]^q$ is defined by $H(b) = -\sum_{i \in [q]} b_i \ln b_i$.

Fact 5.6. *Let $qb \in [0, 1]^q$ be such that $\sum_{i=1}^q b_i = 1/q$, and define*

$$h : [0, 1] \rightarrow \mathbb{R}, \quad z \mapsto -z \ln z - (1-z) \ln(1-z).$$

Then

- (i) for $J \subseteq [q]$ and $r = \sum_{i \in J} qb_i$ we have $H(b) \leq h(r) + r \ln |J| + (1-r) \ln(q-|J|)$, and
- (ii) for $J \subseteq \{2, \dots, q\}$ with $0 < |J| < q-1$ and $r = \sum_{i \in J} qb_i$, if $qb_1 < 1$ then

$$H(b) \leq h(qb_1) + (1-qb_1)h(r/(1-qa_1)) + r \ln |J| + (1-r-qa_1) \ln(q-|J|-1).$$

The following lemma is the main tool to carry out the local variations argument. Recall that \mathcal{S} is the set of all matrices $a = (a_{ij})_{i,j \in [q]}$ with entries $a_{ij} \geq 0$ such that $\sum_j a_{ij} = 1/q$ for all i .

Lemma 5.7. *Suppose $a \in \mathcal{S}$. If $i \in [q]$ and $\emptyset \neq J \subseteq [q]$ are such that for some number $3 \ln \ln q / \ln q \leq \mu \leq 1$ we have*

$$|J| \geq q^\mu \quad \text{and} \quad \max_{j \in J} a_{ij}^{k-1} < \frac{0.995}{kq^{k-1}} (\mu - \ln \ln q / \ln q), \quad (5.4)$$

then the matrix $\tilde{a} \in \mathcal{S}$ obtained from a by setting

$$\tilde{a}_{xy} = \mathbf{1}\{(x, y) \notin \{i\} \times J\} a_{xy} + \frac{\mathbf{1}\{(x, y) \in \{i\} \times J\}}{|J|} \sum_{j \in J} a_{ij}$$

is such that $F(a) \leq F(\tilde{a})$. In fact, the inequality is strict unless $a = \tilde{a}$.

Proof. Take $i \in [q]$, $J \subset [q]$ as described and $x, y \in J$ such that

$$a_{ix}^{k-1} = \min_{j \in J} a_{ij}^{k-1} < a_{iy}^{k-1} < \frac{0.995}{kq^{k-1}} (\mu - \ln \ln q / \ln q).$$

We will show that (5.3) is positive for the range of a_{ix} and a_{iy} that we have at hand. It will be convenient to make the substitution $\delta_{xy} = a_{iy}^{k-1} - a_{ix}^{k-1} > 0$ and instead consider whether

$$\begin{aligned} (k-1) \left(\frac{\partial}{\partial a_{ix}} - \frac{\partial}{\partial a_{iy}} \right) F(a) &= \ln \left(\left(\frac{a_{iy}}{a_{ix}} \right)^{k-1} \right) - \frac{ck(k-1)(a_{iy}^{k-1} - a_{ix}^{k-1})}{1 - 2/q^{k-1} + \|a\|_k^k} \\ &= \ln \left(1 + \frac{\delta_{xy}}{a_{ix}^{k-1}} \right) - \frac{ck(k-1)\delta_{xy}}{1 - 2/q^{k-1} + \|a\|_k^k} =: \Delta(\delta_{xy}) > 0. \end{aligned} \quad (5.5)$$

After noting that $\Delta(0) = 0$, it follows from the concavity of Δ that if $\delta^* > 0$ satisfies (5.5) then so does δ_{xy} for all $0 < \delta_{xy} < \delta^*$. Therefore we take

$$\delta^* = \frac{0.999}{kq^{k-1}} (\mu - \ln \ln q / \ln q) > \max_{x,y \in J} \delta_{xy} = \max_{x,y \in J} |a_{iy}^{k-1} - a_{ix}^{k-1}|,$$

and observe that $a_{ix} \leq \frac{1}{|J|} \sum_{j \in J} a_{ij} \leq \frac{1}{q|J|}$. After taking the exponential of (5.5), we have

$$\begin{aligned} \exp \left\{ \frac{ck(k-1)\delta^*}{1-2/q^{k-1} + \|a\|_k^k} \right\} &< \exp \{ (k-1) \ln q (\mu - \ln \ln q / \ln q) \} = (q^\mu / \ln q)^{k-1} \leq (|J| / \ln q)^{k-1} \\ &\leq (qa_{ix} \ln q)^{1-k} \leq 1 + \frac{1.99 \ln \ln q}{kq^{k-1} a_{ix}^{k-1} \ln q} \leq 1 + \frac{0.995}{kq^{k-1}} (\mu - \ln \ln q / \ln q) \cdot \frac{1}{a_{ix}^{k-1}} < 1 + \delta^* / a_{ix}^{k-1}, \end{aligned}$$

as required. \square

In other words, if we take a row i and a set J of not too few columns such that the largest entry a_{ij} , $j \in J$, is not too big, then the function value does not drop if we replace all entries a_{ij} , $j \in J$, by their average. Thus, Lemma 5.7 can be used to “flatten” parts of the matrix a without reducing the function value.

In what follows we will use Lemma 5.7 to show that Lemma 5.2 holds for each $0 \leq s < q$ separately. Formally, we set out to show that:

Claim 5.8. For all $a \in \mathcal{D}_{0,tame} \setminus \{\bar{a}\}$ we have $F(a) < F(\bar{a})$.

Claim 5.9. Suppose that $1 \leq s \leq q^{0.999}$. Then for all $a \in \mathcal{D}_{s,tame}$ we have $F(a) < F(\bar{a})$.

Claim 5.10. Suppose that $q^{0.999} < s < q - q^{0.49}$. Then for all $a \in \mathcal{D}_{s,tame}$ we have $F(a) < F(\bar{a})$.

Claim 5.11. Suppose that $q - q^{0.49} \leq s < q$. Then for all $a \in \mathcal{D}_{s,tame}$ we have $F(a) < F(\bar{a})$.

Lemma 5.2 is then immediate from Claims 5.8–5.11.

The general strategy will be to compare $a \in \mathcal{D}_{s,tame}$ to the overlap $\bar{a}(s)$ defined in (5.1) and finally to the central overlap above. To this end, we observe that

$$H(\bar{a}) = 2 \ln q, \quad \text{and} \quad E(\bar{a}) = -2 \ln q + \frac{2 \ln 2}{q^{k-1}} + o(q^{1-k}), \quad (5.6)$$

and if $s < q$ then

$$H(\bar{a}(s)) = \frac{s}{q} \ln q + \frac{q-s}{q} \ln(q(q-s)), \quad E(\bar{a}(s)) < -2 \ln q + \frac{s}{q} \ln q + \tilde{O}_q(q^{1-k}) \quad \text{and thus} \quad (5.7)$$

$$\begin{aligned} F(\bar{a}(s)) &= \frac{s}{q} \ln q + \frac{q-s}{q} \ln(q(q-s)) + c \ln \left(1 + \frac{s-2q}{q^k} + \frac{(q-s)^2}{q^k(q-s)^k} \right) \\ &< \ln q + \frac{q-s}{q} \ln(q-s) + o_q(q^{1-k}) \\ &\quad + \left[(q^{k-1} - 1/2) \ln q - \ln 2 \right] \cdot \left[\left(\frac{s-2q}{q^k} + \frac{(q-s)^2}{q^k(q-s)^k} \right) - \frac{1}{2} \left(\frac{s-2q}{q^k} + \frac{(q-s)^2}{q^k(q-s)^k} \right)^2 \right] \\ &= (1-s/q) \ln(1-s/q) + \frac{2 \ln 2}{q^{k-1}} - \frac{s \ln 2}{q^k} - \frac{s \ln q}{2q^k} + \frac{\ln q}{q^{k-1}} + \frac{\ln q}{q^{k-1}(1-s/q)^{k-2}} \\ &\quad - \frac{q^{k-1} \ln q}{2} \left[\left(\frac{s-2q}{q^k} + \frac{(q-s)^2}{q^k(q-s)^k} \right)^2 \right] + o_q(q^{1-k}). \end{aligned} \quad (5.8)$$

5.3.2 Proof of Claim 5.8

We begin with the following consequence of Lemma 5.7.

Claim 5.12. *Suppose that $a \in \mathcal{S}$ has an entry $a_{ij} \in [1.02/(qk), q^{-1}(1.01/k)^{1/(k-1)}]$. Then the matrix $a' \in \mathcal{S}$ with entries*

$$a'_{xy} = \mathbf{1}\{x \neq i\}a_{xy} + \mathbf{1}\{x = i\}q^{-2} \quad (x, y \in [q])$$

satisfies $F(a') > F(a)$.

Proof. Without loss of generality we may assume that a maximises $F(a)$ over the set $a \in \mathcal{S}$ with respect to $a_{11} \in [1.02/(qk), q^{-1}(1.01/k)^{1/(k-1)}]$. If we apply Lemma 5.7 to the set $J = [q] \setminus \{1\}$ with $\mu = \ln(q-1)/\ln q$ then the maximality of $F(a)$ implies that $a_{1j} = (1 - qa_{11})/(q(q-1))$ for $j \geq 2$. Let a_1 denote the first row of a . Because a' is obtained from a by replacing the first row by (q^{-2}, \dots, q^{-2}) , the change in entropy comes to

$$H(a') - H(a) = q^{-1} \ln q - H(qa_1) \geq q^{-1} (\ln q - \ln 2 - (1 - 1.02/k) \ln q) \geq q^{-1} ((1.02 \ln q)/k - \ln 2). \quad (5.9)$$

Furthermore,

$$\|a\|_k^k - \|a'\|_k^k = a_{11}^k - q^{1-2k} + (q-1) \left[\frac{1 - qa_{11}}{q(q-1)} \right]^k \leq q^{-k} (1.01/k)^{k/(k-1)} + 4q^{1-2k}. \quad (5.10)$$

The derivative of the function E from Lemma 3.3 satisfies

$$\frac{\partial E(a)}{\partial \|a\|_k^k} = \frac{c}{1 - 2q^{1-k} + \|a\|_k^k} \leq 1.001q^{k-1} \ln q. \quad (5.11)$$

Hence, (5.10) implies that $E(a) - E(a') \leq 1.02k^{-k/(k-1)}q^{-1} \ln q$. Combining this bound with (5.9) and assuming that $q \geq q_0$ for a large enough constant q_0 , we find $F(a') - F(a) = H(a') - H(a) + E(a') - E(a) > 0$. \square

Proof of Claim 5.8. The set $\mathcal{D}_{0,\text{tame}}$ is compact. Therefore, the continuous function F attains a maximum at some point $a \in \mathcal{D}_{0,\text{tame}}$. Assume for contradiction that $a \neq \bar{a}$. Then we will construct a sequence of matrices $a[i]$, $i \in [q]$, such that $a[0] = a$, $a[q] = \bar{a}$, with $F(a[i+1]) \geq F(a[i])$ for all $i < q$ and $F(a[0]) \neq F(a[q])$, clearly arriving at a contradiction to the maximality of $F(a)$. Specifically, let $a[0] = a$ and obtain $a[i]$ from $a[i-1]$ by letting

$$a_{xy}[i] = \mathbf{1}\{x \neq i\}a_{xy}[i-1] + \mathbf{1}\{x = i\}q^{-2} \quad \text{for } i, x, y \in [q].$$

This construction ensures that $a[q] = \bar{a}$. To show that $F(a[i+1]) \geq F(a[i])$ we consider two cases.

Case 1: $\max_{j \in [q]} a_{ij} \leq 1.02/(qk)$. We apply Lemma 5.7 with $J = [q]$ and $\mu = 1$. Since $a_{ij} \leq 1.02/(qk)$, the assumption (5.4) is satisfied. Consequently, $F(a[i]) \geq F(a[i-1])$, with equality if and only if $a[i] = a[i-1]$.

Case 2: $\max_{j \in [q]} a_{ij} > 1.02/(qk)$. Claim 5.12 shows that $F(a[i]) > F(a[i-1])$.

Finally, since $a \neq \bar{a}$ we have $a[i] \neq a[i-1]$ for some $i \in [q]$, whence $F(\bar{a}) = F(a[q]) > F(a[0]) = F(a)$. Note that although we may temporarily leave $\mathcal{D}_{0,\text{tame}}$ during this process, we are guaranteed to return to $\bar{a} \in \mathcal{D}_{0,\text{tame}}$. \square

5.3.3 Proof of Claim 5.9

The strategy of this proof is to compare an arbitrary element of \mathcal{D}_s to a matrix that is more evenly distributed (using Lemma 5.7), which we then compare to the barycentre of the face of \mathcal{D} (i.e. $\bar{a}(s)$) and finally, to which we compare \bar{a} . Let $1 \leq s \leq q^{0.999}$ and take $a \in \mathcal{D}_s$. It follows from Corollary 5.12 and the definition of separability that we may

assume $qa_{ii} \geq 1 - \kappa$ for $i \leq s$ with $\kappa = \ln^{20} q / q^{k-1}$, and further, that we may also assume $qa_{ij} < 1.02/k$ for all $i \neq j \leq s$ and $s < i, j \leq q$. Let $q\hat{a}$ be the singly-stochastic matrix with entries

$$\hat{a}_{ij} = \begin{cases} a_{ij} & \text{if } i \in [q], j \leq s, \\ \frac{1}{q-s} \sum_{\ell > s} a_{i\ell} & \text{if } i \in [q], j > s. \end{cases}$$

Since $q-s = q(1 - o_q(1))$ we may apply Lemma 5.7 to $J = [q] \setminus [s]$ for any $i \in [q]$. It follows that $F(a) \leq F(\hat{a})$. We will now compare $F(\hat{a})$ and $F(\bar{a}(s))$. To this end we must first estimate $F(\hat{a})$. We start with the entropy term. As \hat{a} is stochastic and $q\hat{a}_{ii} \geq 1 - \kappa$ for $i \leq s$, we find that

$$r_i = q \sum_{i \neq j} \hat{a}_{ij} = 1 - qa_{ii} \leq \kappa, \quad \text{for } i \leq s.$$

Further, if we set $r_i = q \sum_{j=1}^s \hat{a}_{ij}$ for $i > s$ then it follows from the fact that qa is doubly-stochastic that

$$\sum_{i > s} r_i = q \sum_{i > s} \sum_{j=1}^s \hat{a}_{ij} = q \sum_{i > s} \sum_{j=1}^s a_{ij} \leq \kappa s, \quad \text{for } i > s.$$

Let \hat{a}_i denote the i th row of \hat{a} . We know from Fact 5.6 that

$$H(q\hat{a}_i) \leq h(r_i) + r_i \ln(q-1) \leq h(\kappa) + \kappa \ln q \quad \text{for } i \leq s,$$

and

$$H(q\hat{a}_i) \leq h(r_i) + r_i \ln s + (1 - r_i) \ln(q-s) \leq h(r_i) + r_i \ln s + \ln(q-s), \quad \text{for } i > s.$$

Since h is concave, it follows that

$$\sum_{i > s} H(q\hat{a}_i) \leq (q-s) \ln(q-s) + \sum_{i > s} (h(r_i) - r_i \ln s) \leq (q-s) \ln(q-s) + qh\left(\frac{\kappa s}{q}\right) + \kappa s \ln s.$$

Therefore

$$\begin{aligned} H(\hat{a}) &= \ln q + \frac{1}{q} \sum_{i=1}^q H(q\hat{a}_i) \leq \ln q + \frac{s}{q} (h(\kappa) + \kappa \ln q) + \frac{q-s}{q} \ln(q-s) + h\left(\frac{\kappa s}{q}\right) + \frac{\kappa s}{q} \ln s \\ &\leq \ln q + \frac{q-s}{q} \ln(q-s) + o_q(q^{1-k}) \quad \text{[as } s \leq q^{0.999} \text{ and } h(\kappa s/q) = \tilde{O}_q(q^{1-k})\text{]} \\ &= H(q\bar{a}(s)) + o_q(q^{1-k}) \quad \text{[by (5.7)].} \end{aligned} \tag{5.12}$$

Next we deal with estimation of the energy term. It will be convenient to break down the problem as follows:

$$\|\hat{a}\|_k^k = \sum_{i \leq s} \sum_{j \leq q} \hat{a}_{ij}^k + \sum_{i > s} \sum_{j > s} \hat{a}_{ij}^k + \sum_{i > s} \sum_{j \leq s} \hat{a}_{ij}^k.$$

As the k -norm is maximised when summands are as unequal as possible, we have $\|\hat{a}_i\|_k^k \leq q^{1-k}$ for $i \leq s$. Further, by the same logic we have

$$\sum_{i > s} \sum_{j > s} \hat{a}_{ij}^k = (q-s)^2 \left(\frac{1}{q-s} \sum_{\ell > s} a_{i\ell} \right)^k \leq (q-s)^{2-k} q^{-k},$$

and

$$\sum_{i > s} \sum_{j \leq s} \hat{a}_{ij}^k \leq \left(\sum_{i > s} \sum_{j \leq s} \hat{a}_{ij} \right)^k \leq \left(\frac{\kappa s}{q} \right)^k.$$

As $s/q \leq q^{-0.001}$ we know

$$\sum_{i>s} \sum_{j \leq s} \hat{a}_{ij}^k \leq \left(\frac{\kappa s}{q}\right)^k \leq q^{k(1-k)} q^{-\gamma},$$

for some $\gamma > 0$. If we combine the above results then we have shown that

$$\|\hat{a}\|_k^k \leq s q^{1-k} + (q-s)^{2-k} q^{-k} + q^{k(1-k)} q^{-\gamma} = \|\bar{a}(s)\|_k^k + q^{k(1-k)} q^{-\gamma},$$

and so

$$E(\hat{a}) - E(\bar{a}(s)) \leq \frac{\partial E(a)}{\partial \|a\|_k^k} (\|\hat{a}\|_k^k - \|\bar{a}(s)\|_k^k) \leq q^{-(1-k)^2} q^{-\gamma} \ln q (1 + o_q(1/q)) = o_q(q^{1-k}). \quad (5.13)$$

Therefore it follows from (5.12) and (5.13) that

$$F(a) \leq F(\hat{a}) \leq F(\bar{a}(s)) + o_q(q^{1-k}).$$

Recalling that $s/q \leq q^{-0.001}$, it follows from (5.8) that

$$\begin{aligned} F(a) &\leq F(\hat{a}) \leq F(\bar{a}(s)) + o_q(q^{1-k}) \\ &= (1-s/q) \ln(1-s/q) + \frac{2\ln 2}{q^{k-1}} - \frac{s \ln 2}{q^k} - \frac{s \ln q}{2q^k} + \frac{\ln q}{q^{k-1}} + \frac{\ln q}{q^{k-1}(1-s/q)^{k-2}} \\ &\quad - \frac{q^{k-1} \ln q}{2} \left[\left(\frac{s-2q}{q^k} + \frac{(q-s)^2}{q^k(q-s)^k} \right)^2 \right] + o_q(q^{1-k}) \\ &= (1-s/q) \ln(1-s/q) + \frac{2\ln 2}{q^{k-1}} + \frac{\ln q}{q^{k-1}} + \frac{\ln q}{q^{k-1}(1-s/q)^{k-2}} \\ &\quad - \frac{q^{k-1} \ln q}{2} \left(\frac{s-2q}{q^k} \right)^2 + o_q(q^{1-k}) \\ &= -\frac{s}{q} (1-s/q) + \frac{2\ln 2}{q^{k-1}} + \frac{\ln q}{q^{k-1}} + \frac{\ln q}{q^{k-1}(1-s/q)^{k-2}} - \frac{2\ln q}{q^{k-1}} + o_q(q^{1-k}) \\ &\leq (1-s/q) \ln(1-s/q) + \frac{2\ln 2}{q^{k-1}} + o_q(q^{1-k}) = F(\bar{a}) - \frac{s}{q} (1-s/q) + o_q(q^{1-k}). \end{aligned}$$

As the $\frac{s}{q}(1-s/q)$ is decreasing in s , we have shown that $F(a) < F(\bar{a}) - 1/q + 1/q^2 + o_q(q^{1-k})$. This implies our original assertion. \square

5.3.4 Proof of Claim 5.10

Let $q^{0.999} < s < q - q^{0.49}$ and take $a \in \mathcal{D}_s$. As before, we may assume $qa_{ii} \geq 1 - \kappa$ for $i \leq s$, and $qa_{ij} < 1.02/k$ for all $i \neq j \leq s$ and $s < i, j \leq q$. Let $q\hat{a}$ be the singly-stochastic matrix with entries

$$\hat{a}_{ij} = \begin{cases} a_{ij} & \text{if } i = j \in [s], \\ \frac{1}{s-1} \sum_{\ell \in [s] \setminus \{i\}} a_{i\ell} & \text{if } i, j \leq s, i \neq j, \\ \frac{1}{q-s} \sum_{\ell > s} a_{i\ell} & \text{if } j > s, \\ \frac{1}{s} \sum_{\ell \leq s} a_{i\ell} & \text{if } j \leq s < i. \end{cases}$$

Since $s, q-s > q^{0.49}$ we may apply Lemma 5.7 to $J = [q] \setminus [s]$ and $J' = [s] \setminus \{i\}$ for any $i \in [q]$. It follows that $F(a) \leq F(\hat{a})$. To estimate $F(\hat{a})$ we will now define

$$r_i = q \sum_{j>s} a_{ij} = q \sum_{j>s} \hat{a}_{ij} \text{ for } i \leq s, \text{ and } r_i = q \sum_{j \leq s} a_{ij} = q \sum_{j \leq s} \hat{a}_{ij} \text{ for } i > s,$$

and since qa is doubly stochastic,

$$r = \sum_{i>s} r_i = \sum_{i\leq s} r_i \leq \sum_{i\leq s} 1 - qa_{ii} \leq \kappa s.$$

Further, we also set

$$t_i = q \sum_{j \in [s] \setminus \{i\}} \hat{a}_{ij} = q \sum_{j \in [s] \setminus \{i\}} a_{ij} \leq 1 - qa_{ii} \leq \kappa \quad \text{for } i \leq s. \quad (5.14)$$

As before we will now estimate the entropy term and the energy term separately. Again, we let \hat{a}_i denote the i th row of \hat{a} . We know from Fact 5.6 (ii) that

$$\begin{aligned} H(q\hat{a}_i) &\leq h(qa_{ii}) + (1 - qa_{ii})h(t_i/(1 - qa_{ii})) + t_i \ln(s-1) + (1 - t_i - qa_{ii}) \ln(q-s) \\ &\leq h(qa_{ii}) + (1 - qa_{ii})h(r_i/(1 - qa_{ii})) + t_i \ln s + r_i \ln(q-s) \\ &= -qa_{ii} \ln(qa_{ii}) - t_i \ln t_i - r_i \ln r_i + t_i \ln s + r_i \ln(q-s) \leq h(t_i) + t_i \ln s + h(r_i) + r_i \ln(q-s), \quad i \leq s, \end{aligned}$$

where the last line follows as the function $g : x \mapsto -(1-x) \ln(1-x)$ is decreasing with $g'(x) \leq 1$ for small x . If we set $\tilde{H} = \frac{1}{q} \sum_{i \leq s} (h(t_i) + t_i \ln s)$ then it follows from the concavity of h that

$$\frac{1}{q} \sum_{i \leq s} H(q\hat{a}_i) \leq \tilde{H} + \frac{s}{q} h(r/s) + \frac{r}{q} \ln(q-s).$$

Furthermore, by Fact 5.6 (i) and the concavity of h , we have

$$\frac{1}{q} \sum_{i>s} H(q\hat{a}_i) \leq \frac{q-s}{q} h(r/(q-s)) + \frac{r}{q} \ln s + \frac{q-s-r}{q} \ln(q-s).$$

Combining these results, it follows that

$$H(\hat{a}) \leq \ln q + \tilde{H} + \left(\frac{s}{q} h(r/s) + \frac{r}{q} \ln(q-s) \right) + \left(\frac{q-s}{q} h(r/(q-s)) + (r/q) \ln s \right) + \frac{q-s-r}{q} \ln(q-s).$$

Further as $h(x) \leq x(1 - \ln x)$, we have

$$\begin{aligned} H(\hat{a}) - \tilde{H} &\leq \ln q + \frac{r}{q} [2 - 2 \ln r + 2 \ln s + \ln(q-s)] + \frac{q-s}{q} \ln(q-s) \\ &\leq \ln q + \frac{r}{q} (2 + 3 \ln q) + \frac{q-s}{q} \ln(q-s) + O_q(1/q), \quad [\text{as } -z \ln z \leq 1 \text{ for } z \geq 0] \\ &= 2 \ln q + \frac{r}{q} (2 + 3 \ln q) + (1 - s/q) \ln(1 - s/q) - \frac{s \ln q}{q} + O_q(1/q). \end{aligned} \quad (5.15)$$

Since $s < q$, we obtain

$$\tilde{H} - \frac{2 \ln q}{q} \sum_{i \leq s} t_i = \frac{1}{q} \sum_{i \leq s} (h(t_i) + t_i (\ln s - 2 \ln q)) \leq \frac{1}{q} \sum_{i \leq s} (h(t_i) - t_i \ln q) \leq \frac{1}{q}, \quad (5.16)$$

where the last inequality follows from noting that $\max_{x \in [0,1]} h(x) - x \ln q \leq 1/q$. Thus, by combining (5.15) and (5.16) we have shown

$$H(\hat{a}) \leq 2 \ln q + \frac{r}{q} (2 + 3 \ln q) + (1 - s/q) \ln(1 - s/q) - \frac{s \ln q}{q} + \frac{2 \ln q}{q} \sum_{i \leq s} t_i + O_q(1/q). \quad (5.17)$$

Next, we move on to estimating the energy term. As before we firstly estimate $\|a\|_k^k$, then we apply a bound for $\partial E / \partial \|a\|_k^k$ in order to approximate $E(\hat{a}) - E(\bar{a}(s))$. Firstly note from (5.14) that for $i \leq s$, we have

$$\hat{a}_{ii}^k \leq q^{-k} (1 - t_i)^k = \frac{1}{q^k} - \frac{kt_i}{q^k} + o_q(q^{1-2k}), \quad \text{and}$$

$$\sum_{j \in [s] \setminus \{i\}} \hat{a}_{ij}^k = (s-1) \left(\frac{t_i/q}{s-1} \right)^k \leq \frac{(\kappa/q)^k}{(s-1)^{k-1}} \leq (\kappa/q)^k.$$

Moreover, since $q\hat{a}$ is stochastic and $q\hat{a}_{ii} \geq 1 - \kappa$ if $i \leq s$, we have

$$\sum_{j \in [q] \setminus [s]} \hat{a}_{ij}^k \leq (\kappa/q)^k, \quad \text{for } i \leq s.$$

Combining the above equations yields

$$\sum_{i \leq s} \|\hat{a}_i\|_k^k \leq sq^{-k} - \frac{k}{q^k} \sum_{i \leq s} t_i + o_q(q^{1-2k}).$$

Since $qa_{ii} \geq 1 - \kappa$ for $i \leq s$ we have $qa_{ij} \leq \kappa$ for $j \leq s < i$. By construction, this implies that $q\hat{a}_{ij} \leq \kappa$ for $j \leq s < i$. Furthermore, we have that

$$\sum_{i > s} \sum_{j \leq s} \hat{a}_{ij}^k \leq \frac{\kappa^k s}{q^k} \quad \text{and} \quad \sum_{i > s} \sum_{j > s} \hat{a}_{ij} = (q-s)^2 \left(\frac{\sum_{j > s} a_{ij}}{q-s} \right)^k \leq q^{-k} (q-s)^{2-k}.$$

We have shown that

$$\|\hat{a}\|_k^k \leq sq^{-k} + q^{-k} (q-s)^{2-k} - \frac{k}{q^k} \sum_{i \leq s} t_i + o_q(q^{1-2k}) = \|\bar{a}(s)\|_k^k - \frac{k}{q^k} \sum_{i \leq s} t_i + o_q(q^{1-2k}),$$

and so from (5.7) and (5.11), we have

$$\begin{aligned} E(\hat{a}) &= E(\bar{a}(s)) + \frac{\partial E(a)}{\partial \|a\|_k^k} \cdot (\|\hat{a}\|_k^k - \|\bar{a}(s)\|_k^k) \\ &\leq E(\bar{a}(s)) - q^{k-1} \ln q (1 + o_q(1/q)) \cdot \left(\frac{k}{q^k} \sum_{i \leq s} t_i + o_q(q^{1-2k}) \right) \\ &\leq E(\bar{a}(s)) - \frac{k \ln q}{q} \sum_{i \leq s} t_i + o_q(q^{1-k}) \\ &= -2 \ln q + \frac{s}{q} \ln q - \frac{k \ln q}{q} \sum_{i \leq s} t_i + O_q(1/q). \end{aligned} \tag{5.18}$$

Finally then, it follows from (5.17) and (5.18) that

$$\begin{aligned} F(a) \leq F(\hat{a}) &\leq \frac{r}{q} (2 + 3 \ln q) + (1 - s/q) \ln(1 - s/q) + \frac{(2-k) \ln q}{q} \sum_{i \leq s} t_i + O_q(1/q) \\ &= (1 - s/q) \ln(1 - s/q) + O_q(1/q) \leq -\frac{s}{q} (1 - s/q) + O_q(1/q). \end{aligned}$$

Fortunately, our assumption $q^{0.999} < s < q - q^{0.49}$ ensures that $F(a) < 0 < F(\bar{a})$. □

5.3.5 Proof of Claim 5.11

Let $q - \sqrt{q} \leq s \leq q - 1$ and take $a \in \mathcal{D}_s$. As before we may assume $qa_{ii} \geq 1 - \kappa$ for $i \in [s]$, and $qa_{ij} < 1.02/k$ for all $i \neq j \leq s$ and $s < i, j \leq q$. Let $r_i = q \sum_{j \neq i} a_{ij}$. As qa is doubly-stochastic and $qa_{ii} \geq 1 - \kappa$ for $i \leq s$, we have

$$r = \sum_{i \leq s} r_i = q \sum_{i \leq s} \sum_{j \neq i} a_{ij} = \sum_{i \leq s} 1 - qa_{ii} \leq \kappa s.$$

Further, we let

$$t_i = \sum_{j > s} qa_{ij}, \quad \text{and} \quad t = \sum_{i \leq s} t_i.$$

Since qa is doubly-stochastic we have

$$t = \sum_{i \leq s} \sum_{j > s} qa_{ij} = \sum_{i > s} \sum_{j \leq s} qa_{ij}. \quad (5.19)$$

The strategy of this proof is to compare $F(a)$ to $F(q^{-1}\text{id})$ where id is the $q \times q$ identity matrix. To this end we firstly estimate the entropy of a . Again, let a_i denote the i th row of a . We now set $\bar{H} = \frac{1}{q} \sum_{i \leq s} h(qa_{ii})$ and as before apply Fact 5.6 (ii) and the concavity of h to observe that

$$\begin{aligned} \frac{1}{q} \sum_{i \leq s} H(qa_i) &\leq \frac{1}{q} \sum_{i \leq s} h(qa_{ii}) + r_i h(t_i/r_i) + t_i \ln(q-s) + (r_i - t_i) \ln s \\ &\leq \bar{H} + \frac{r}{q} h(t/r) + \frac{t}{q} \ln(q-s) + \frac{r-t}{q} \ln s \\ &\leq \bar{H} + \frac{t}{q} (1 - \ln t + \ln r) + \frac{t}{q} \ln(q-s) + \frac{r-t}{q} \ln s \quad [\text{as } h(z) \leq z(1 - \ln z)]. \end{aligned}$$

As $-z \ln z \leq 1$ for $z > 0$, we have that $-t \ln t \leq 1$. Furthermore, as qa is doubly-stochastic we have that $t \leq q-s$, and so

$$\frac{t}{q} (1 - \ln t + \ln r) \leq \frac{q-s}{q} \cdot (1 + \tilde{O}(1/q)).$$

Therefore

$$\frac{1}{q} \sum_{i \leq s} H(qa_i) \leq \bar{H} + \frac{t}{q} \ln(q-s) + \frac{r-t}{q} \ln s + \frac{q-s}{q} (1 + \tilde{O}_q(1/q)). \quad (5.20)$$

We now move to estimating $H(qa_i)$ for $i > s$. As is by now routine, we apply Fact 5.6 (i) along with the concavity of h and (5.19) to conclude that

$$\begin{aligned} \frac{1}{q} \sum_{i > s} H(qa_i) &\leq \frac{1}{q} \sum_{i > s} \left[h\left(\sum_{j \leq s} qa_{ij}\right) + \sum_{j \leq s} qa_{ij} \ln(s) + \left(1 - \sum_{j \leq s} qa_{ij}\right) \ln(q-s) \right] \\ &\leq \frac{q-s}{q} h\left(\frac{t}{q-s}\right) + \frac{t}{q} \ln s + \frac{q-s-t}{q} \ln(q-s) \\ &\leq \frac{q-s}{q} \ln 2 + \frac{t}{q} \ln s + \frac{q-s-t}{q} \ln(q-s) \quad [\text{as } h(z) \leq \ln 2 \text{ for all } z]. \end{aligned} \quad (5.21)$$

Finally then, we have from (5.20) and (5.21) that

$$\begin{aligned} H(a) &= \ln q + \bar{H} + \frac{1}{q} \sum_{i \leq q} H(qa_i) \\ &\leq \ln q + \frac{r}{q} \ln s + \frac{q-s}{q} \ln 2 + \frac{q-s}{q} \ln(q-s) + \frac{q-s}{q} (1 + \tilde{O}(1/q)) \\ &\leq \ln q + \bar{H} + \frac{r}{q} \ln s + \frac{q-s}{q} \ln 2 + \frac{q-s}{2q} \ln q + \frac{q-s}{q} (1 + \tilde{O}(1/q)), \quad [\text{as } q-s \leq \sqrt{q}]. \end{aligned}$$

Moving on to the energy term, we firstly estimate $\|a\|_k^k$. As the norm is maximised when the summands are widely distributed, we have

$$\sum_{i \leq s} \|a_i\|_k^k \leq \kappa^k s + \sum_{i \leq s} a_{ii}^k = \sum_{i \leq s} a_{ii}^k + o(1/q^{k+1}).$$

A similar argument then applies to the remaining $q-s$ rows. Recalling Corollary 5.12, it follows that

$$\sum_{i > s} \|a_i\|_k^k \leq (q-s) \left(\frac{1.02}{qk}\right)^k.$$

Therefore,

$$\|a\|_k^k \leq (q-s) \left(\frac{1.02}{qk} \right)^k + \sum_{i \leq s} a_{ii}^k + o(1/q^{k+1}),$$

and so

$$\|a\|_k^k - \|q^{-1}\text{id}\|_k^k \leq \frac{q-s}{q^k} \left[\left(\frac{1.02}{k} \right)^k - 1 \right] + \frac{1}{q^k} \sum_{i \leq s} [(qa_{ii})^k - 1] + o(1/q^{k+1}).$$

Finally then, we have from (5.11) that

$$\begin{aligned} E(a) - E(q^{-1}\text{id}) &= \frac{\partial E(a)}{\partial a} (\|a\|_k^k - \|q^{-1}\text{id}\|_k^k) \\ &\leq \ln q \cdot \frac{q-s}{q} \left[\left(\frac{1.02}{k} \right)^k - 1 \right] + \frac{\ln q}{q} \sum_{i \leq s} [(qa_{ii})^k - 1] + o_q(\ln q/q). \end{aligned}$$

If we combine our estimates for the entropy and energy terms, we have shown that

$$\begin{aligned} F(a) - F(q^{-1}\text{id}) &\leq \bar{H} + \frac{r}{q} \ln s + \frac{\ln q}{q} \sum_{i \leq s} [(qa_{ii})^k - 1] + o_q(\ln q/q) \\ &\quad + \ln q \cdot \frac{q-s}{q} \left[\left(\frac{1.02}{k} \right)^k + \frac{2 + \ln 2}{\ln q} - 1/2 \right]. \end{aligned}$$

Since $\max_{0 < z < 1} h(z) - z \ln q \leq 1/q$ and $h(z) = h(1-z)$, if we set $\rho_{ii} = 1 - qa_{ii}$ then

$$\begin{aligned} \bar{H} + \frac{r}{q} \ln s + \frac{\ln q}{q} \sum_{i \leq s} [(qa_{ii})^k - 1] &\leq \frac{1}{q} \sum_{i \leq s} [h(\rho_{ii}) + \rho_{ii} \ln q + ((1 - \rho_{ii})^k - 1) \ln q] \\ &\leq \frac{1}{q} \sum_{i \leq s} [h(\rho_{ii}) + \rho_{ii} \ln q - k\rho_{ii} \ln q] + O(1/q) \leq \frac{1}{q} \sum_{i \leq s} [h(\rho_{ii}) - \rho_{ii} \ln q] + O(1/q) = O(1/q). \end{aligned}$$

Finally then

$$F(a) \leq F(q^{-1}\text{id}) - \ln q/3q + o_q(\ln q/q) \leq F(q^{-1}\text{id}) = \frac{1}{2}F(\bar{a}).$$

□

Acknowledgement

We would like to thank the referees for their many helpful comments which have improved the paper.

References

- [1] D. Achlioptas, A. Coja-Oghlan: Algorithmic barriers from phase transitions. Proc. 49th FOCS (2008) 793–802.
- [2] D. Achlioptas, C. Moore: Random k -SAT: two moments suffice to cross a sharp threshold. SIAM Journal on Computing **36** (2006) 740–762.
- [3] D. Achlioptas, C. Moore: On the 2-colorability of random hypergraphs. Proc. 6th RANDOM (2002) 78–90.
- [4] D. Achlioptas, A. Naor: The two possible values of the chromatic number of a random graph. Annals of Mathematics **162** (2005), 1333–1349.

- [5] D. Achlioptas, A. Naor, Y. Peres: Rigorous location of phase transitions in hard optimization problems. *Nature* **435** (2005) 759–764.
- [6] N. Alon, M. Krivelevich: The concentration of the chromatic number of random graphs. *Combinatorica* **17** (1997) 303–313
- [7] B. Bollobás: The chromatic number of random graphs. *Combinatorica* **8** (1988) 49–55
- [8] A. Coja-Oghlan: Upper-bounding the k -colorability threshold by counting covers. *Electronic Journal of Combinatorics* **20** (2013) P32.
- [9] A. Coja-Oghlan, D. Vilenchik: Chasing the k -colorability threshold. *Proc. 54th FOCS* (2013) 380–389.
- [10] A. Coja-Oghlan, L. Zdeborová: The condensation transition in random hypergraph 2-coloring. *Proc. 23rd SODA* (2012) 241–250.
- [11] A. Coja-Oghlan, K. Panagiotou: Catching the k -NAESAT threshold. *Proc. 44th STOC* (2012) 899–908.
- [12] J. Ding, A. Sly, N. Sun: Proof of the satisfiability conjecture for large k . *Proc. 47th STOC* (2015) 59–68.
- [13] M. Dyer, A. Frieze, C. Greenhill: On the chromatic number of a random hypergraph. *Journal of Combinatorial Theory, Series B.* **113** (2015), 68-122
- [14] P. Erdős, A. Rényi: On the evolution of random graphs. *Magyar Tud. Akad. Mat. Kutato Int. Kozl.* **5** (1960) 17–61.
- [15] A. Frieze, N. Wormald: Random k -Sat: a tight threshold for moderately growing k . *Combinatorica* **25** (2005) 297–305.
- [16] C. Greenhill, S. Janson, A. Ruciński: On the number of perfect matchings in random lifts. *Combinatorics, Probability and Computing*, 19 (2010), pp. 791–817.
- [17] H. Hatami, M. Molloy: Sharp thresholds for constraint satisfaction problems and homomorphisms. *Random Structures and Algorithms* **33** (2008) 310–332.
- [18] S. Janson, T. Łuczak, A. Ruciński: *Random Graphs*, Wiley, New York, 2000.
- [19] W. Kauzmann: The nature of the glassy state and the behavior of liquids at low temperatures. *Chem. Rev.* **43** (1948) 219–256.
- [20] M. Krivelevich, B. Sudakov: The chromatic numbers of random hypergraphs. *Random Structures and Algorithms* **12** (1998) 381–403.
- [21] F. Krzakala, A. Montanari, F. Ricci-Tersenghi, G. Semerjian, L. Zdeborova: Gibbs states and the set of solutions of random constraint satisfaction problems. *Proc. National Academy of Sciences* **104** (2007) 10318–10323.
- [22] A. Kupavskii, D. Shabanov: Colorings of partial Steiner systems and their applications, *Journal of Mathematical Sciences* **206** (2015) 511–538.
- [23] T. Łuczak: The chromatic number of random graphs. *Combinatorica* **11** (1991) 45–54
- [24] T. Łuczak: A note on the sharp concentration of the chromatic number of random graphs. *Combinatorica* **11** (1991) 295–297
- [25] D. Matula: Expose-and-merge exploration and the chromatic number of a random graph. *Combinatorica* **7** (1987) 275–284.
- [26] E. Shamir, J. Spencer: Sharp concentration of the chromatic number of random graphs $G(n, p)$. *Combinatorica* **7** (1987) 121–129