

SIZE BIASED COUPLINGS AND THE SPECTRAL GAP FOR RANDOM REGULAR GRAPHS

BY NICHOLAS COOK^{1,*}, LARRY GOLDSTEIN^{2,†} AND TOBIAS JOHNSON^{3,†}

Stanford University and University of Southern California[†]*

Let λ be the second largest eigenvalue in absolute value of a uniform random d -regular graph on n vertices. It was famously conjectured by Alon and proved by Friedman that if d is fixed independent of n , then $\lambda = 2\sqrt{d-1} + o(1)$ with high probability. In the present work, we show that $\lambda = O(\sqrt{d})$ continues to hold with high probability as long as $d = O(n^{2/3})$, making progress toward a conjecture of Vu that the bound holds for all $1 \leq d \leq n/2$. Prior to this work the best result was obtained by Broder, Frieze, Suen and Upfal (1999) using the configuration model, which hits a barrier at $d = o(n^{1/2})$. We are able to go beyond this barrier by proving concentration of measure results directly for the uniform distribution on d -regular graphs. These come as consequences of advances we make in the theory of concentration by size biased couplings. Specifically, we obtain Bennett-type tail estimates for random variables admitting certain unbounded size biased couplings.

CONTENTS

| | |
|---------------------------------------------------------------------------------------------|-----|
| 1. Introduction | 73 |
| 1.1. Organization of the paper | 76 |
| 1.2. Notation, definitions and facts | 76 |
| 2. Spectral concentration from measure concentration | 77 |
| 2.1. The uniform tails property | 77 |
| 2.2. High level proofs of the spectral gap for the uniform and permutation models | 79 |
| 3. Concentration by size biased couplings | 82 |
| 3.1. Introduction to size biased couplings | 82 |
| 3.2. New concentration results by size biased couplings | 84 |
| 3.3. Examples | 86 |
| 3.4. Proofs | 88 |
| 4. Size biased couplings for random regular graphs | 94 |
| 5. Concentration for random regular graphs | 99 |
| 5.1. Concentration for the permutation model | 99 |
| 5.2. Uniform tails property for the uniform model | 102 |

Received January 2016; revised February 2017.

¹Supported by NSF Grant DMS-12-66164.

²Supported by NSA Grant H98230-15-1-0250.

³Supported by NSF Grant DMS-14-01479.

MSC2010 subject classifications. Primary 05C80; secondary 60B20, 60E15.

Key words and phrases. Second eigenvalue, random regular graph, Alon’s conjecture, size biased coupling, Stein’s method, concentration.

6. The Kahn–Szemerédi argument 108
 6.1. Light couples 111
 6.2. Heavy couples 113
 6.3. The ε -net and Proof of Proposition 2.4 120
 Acknowledgments 122
 References 123

1. Introduction. Let A be the adjacency matrix of a d -regular graph (i.e., a graph where every vertex has exactly d neighbors), and let $\lambda_1(A) \geq \dots \geq \lambda_n(A)$ be the eigenvalues of A . The trivial eigenvalue $\lambda_1(A)$ is always equal to d ; the second eigenvalue $\lambda_2(A)$, on the other hand, has been the focus of much study over the last thirty years. Alon and Milman demonstrated a close connection between a graph’s second eigenvalue and its expansion properties [2]. Expander graphs were seen to be extraordinarily useful for a range of applications in computer science and beyond (see [37, 43] for good surveys). Alon and Boppana proved a lower bound on $\lambda_2(A)$, showing it to be at least $2\sqrt{d-1}(1 - O(1/\log^2 n))$ [1, 46]. Alon conjectured in [1] that if A is the adjacency matrix of a *random* d -regular graph, the eigenvalue $\lambda_2(A)$ is at most $2\sqrt{d-1} + o(1)$ with probability tending to 1.

Now, take A to be the adjacency matrix of a random graph chosen uniformly from all d -regular graphs on n vertices with no loops or parallel edges, which from now on we call a *uniform random d -regular simple graph on n vertices*. Let $\lambda(A) = \max(\lambda_2(A), -\lambda_n(A))$. After pioneering work by Broder and Shamir [14], Kahn and Szemerédi [28] and Friedman [26], Friedman proved Alon’s conjecture in [27], showing that for any fixed $d \geq 3$ and $\epsilon > 0$,

$$\lim_{n \rightarrow \infty} \mathbf{P}[\lambda(A) \leq 2\sqrt{d-1} + \epsilon] = 1.$$

Also see [12] for a simpler proof of this result.

This result is about sparse graphs; the number of vertices n must be very large compared to d to obtain information about $\lambda(A)$. It is natural to ask about $\lambda(A)$ when both n and d are large. In [15], it is shown that if $d = o(\sqrt{n})$, then $\lambda(A) = O(\sqrt{d})$ with probability tending to 1 as $n \rightarrow \infty$. Vu has conjectured that this holds for all $1 \leq d \leq n/2$; see (1) below for the more precise version of this conjecture made in print. Our main result extends this bound to the range $d = O(n^{2/3})$.

THEOREM 1.1. *Let A be the adjacency matrix of a uniform random d -regular simple graph on n vertices. Let $\lambda_1(A) \geq \dots \geq \lambda_n(A)$ be the eigenvalues of A , and let $\lambda(A) = \max(\lambda_2(A), -\lambda_n(A))$. For any $C_0, K > 0$, there exists $\alpha > 0$ depending only on C_0, K such that if $1 \leq d \leq C_0 n^{2/3}$, then*

$$\mathbf{P}[\lambda(A) \leq \alpha\sqrt{d}] \geq 1 - n^{-K}$$

for all n sufficiently large depending on K .

REMARK 1.2. The proof shows we can take $\alpha = 459,652 + 229,452K + \max(30C_0^{3/2}, 768)$ and $n \geq 7 + K^2$, though we do not attempt to optimize these constants; see Remark 2.7.

REMARK 1.3. The complement of a uniform random d -regular simple graph with adjacency matrix A is a uniform random $(n - d - 1)$ -regular simple graph whose adjacency matrix B has entries $1 - A_{ij}$ for $i \neq j$ and 0 on the diagonal. The nontrivial eigenvalues of B are $-1 - \lambda_i(A)$ for $i = 2, \dots, n$, which implies that $|\lambda(A) - \lambda(B)| \leq 1$. As a consequence, if $\lambda(A) = O(\sqrt{d})$ for $1 \leq d \leq n/2$, then $\lambda(A) = O(\sqrt{d})$ for the full range $1 \leq d \leq n - 1$.

Previous arguments to bound $\lambda(A)$ all proceeded by first establishing the bound for a different distribution on random regular multigraphs (in which loops and multiple edges are permitted), and then transferring the bound to the uniform distribution on d -regular simple graphs by some comparison procedure. [15] work with random regular multigraphs drawn from the *configuration model* (see [54] for a description). A key property of this model is that it gives the same probability to every simple d -regular graph on n vertices. This makes it possible to prove that properties hold with high probability for the uniform model by showing that the probability of failure in the configuration model tends to zero faster than the probability of being simple. When $d \gg n^{1/2}$, the probability that the configuration model being simple decays faster than exponentially in n . Estimating the spectral gap in this way would then require proving that the probabilities of relevant events in the configuration model decay at this rate, and we are unaware of any methods to do so.

Other past work, including [27] and [28], used the *permutation model*, a random $2d$ -regular multigraph whose adjacency matrix is the sum of d independent uniform random permutation matrices and their transposes. It is proven in [34] that the permutation and configuration models are *contiguous*, allowing a second eigenvalue bound to be transferred from the permutation model to the configuration model, from which it can be transferred to the uniform model. Both of these transferences require d to be fixed independently of n .

The reason for using the permutation or configuration model rather than working directly with the uniform distribution on simple d -regular graphs is that the distributions have more independence or martingale structure. In particular, this approach gives access to standard concentration estimates, which play a key role in the approach introduced by Kahn and Szemerédi for the permutation model [28], and adapted for the configuration model in [15]. The argument, which borrows ideas from geometric functional analysis, is explained in more detail in Section 6.

In contrast to previous works, to prove Theorem 1.1 we work directly with the uniform distribution on d -regular simple graphs. A key obstacle is the lack of concentration estimates for this setting. We obtain these using a method based on *size*

biased couplings, developed initially in [30]. These techniques are an offshoot of Stein’s method for distributional approximation; see Section 3 for further discussion. The theory developed in [30] and improved in [3] can show that a nonnegative random variable X is concentrated if there exists a bounded size biased coupling for X (all of these terms are explained in Section 3). These results are analogues of an inequality of Hoeffding [36], Theorem 1, line (2.1), for sums of independent random variables, in which the bound is in terms of the mean of the sum. To make size biasing work in our situation, we extend the theory developed in [3, 30] in two ways. First, we relax the condition that the coupling be bounded. Second, we prove an analogue of Bennett’s inequality [10], equation (8b), in which the concentration bound for a sum is given in terms of its variance rather than its mean.

We apply this theory in several ways besides proving our main result, Theorem 1.1. In Theorem 5.3, we give an edge discrepancy result for random regular graphs, showing concentration for the number of edges between two given sets. We also apply size bias couplings to prove Theorem 2.6, yielding second eigenvalue bounds for distributions of random regular graphs constructed from independent random permutations. The case where the permutations are chosen uniformly, often called the *permutation model*, was considered in [14], [28], [26], [27] and elsewhere, and a second eigenvalue bound for the permutation model was previously proven in [23], Theorem 24. We also consider the case that the permutations have (not necessarily identical) distributions invariant under conjugation and supported on permutations without fixed points. A more graph theoretic interpretation is as follows. For even d , take $d/2$ independent random 2-regular loopless graphs, with no conditions on their distributions except for being invariant under relabeling of vertices. Superimpose these graphs to make a random d -regular graph. We show that this graph obeys a second eigenvalue bound, with no assumption on the distributions of the individual 2-regular graphs. We expect our size bias coupling results to be applicable beyond random regular graphs as well.

Very recently (after the submission of this paper), Tikhomirov and Youssef proved that with high probability $\lambda(A) = O(\sqrt{d})$ for the range $n^\varepsilon \leq d \leq n/2$, which extends Theorem 1.1 to the full range $d \leq n/2$ [50]. Their proof first reduces the question to bounding the second singular value of the adjacency matrix of a random *directed* graph with fixed degree sequence. They tackle this by the Kahn–Szemerédi method, using an inequality for linear forms that they prove with Freedman’s martingale inequality.

One open problem beyond this is to determine the correct constant in the $O(\sqrt{d})$ bound. Vu conjectures that

$$(1) \quad \lambda(A) = (2 + o(1))\sqrt{d(1 - d/n)}$$

whenever $d \leq n/2$ tends to infinity with n [53]. There is also the problem of determining the *fluctuations* of $\lambda(A)$ about this asymptotic value. Numerical simulations for small values of d suggest that after centering and rescaling, $\lambda(A)$

asymptotically follows the $\beta = 1$ Tracy–Widom law [45], which also describes the asymptotic fluctuations of the largest eigenvalue of matrices from the Gaussian Orthogonal Ensemble (GOE) [51]. It was recently shown in [8] that fluctuations of eigenvalues of A in the bulk of the spectrum [i.e., eigenvalues λ_i with $\varepsilon n \leq i \leq (1 - \varepsilon)n$ for some arbitrary fixed $\varepsilon > 0$] asymptotically match those of the GOE, assuming $n^\varepsilon \leq d \leq n^{2/3-\varepsilon}$ for arbitrary fixed $\varepsilon > 0$. It may be possible to extend their approach to establish the universal fluctuations of eigenvalues at the spectral edge for this range of d .

The argument in [8] relied on the *local semicircle law* for uniform random regular graphs, which was established in [9] for the range $\log^4 n \leq d \leq n^{2/3}/\log^{4/3} n$, improving on earlier results in [24, 52]. It follows from the local semicircle law that with high probability, the bulk of the spectrum of A is confined to the scale $O(\sqrt{d})$. Theorem 1.1 complements their result by saying that with high probability the *entire* spectrum, with the exception of the Perron–Frobenius eigenvalue $\lambda_1 = d$, lives at this scale. It would follow from Vu’s conjecture (1) that $\lambda_2(A)$ and $\lambda_n(A)$ “stick to the bulk,” that is, after rescaling by \sqrt{d} these eigenvalues converge to the edge of the support $[-2, 2]$ of the limiting spectral distribution. Interestingly, the limitation to $d \ll n^{2/3}$ in [9] appears to be for reasons similar to our constraint $d = O(n^{2/3})$ in Theorem 1.1, stemming from their use of double switchings (as described in Section 4).

1.1. Organization of the paper. The idea of the proof of Theorem 1.1 is to prove concentration results for random regular graphs by size biasing, and then to apply the Kahn–Szemerédi argument to derive eigenvalue bounds from these concentration inequalities. Section 2 presents this argument at a high level: Proposition 2.3 gives the concentration result and Proposition 2.4 translates it into eigenvalue bounds, with proofs deferred to later in the paper. The proof of Theorem 1.1 appears in Section 2.2 and is a simple application of these two propositions. This section also includes Theorem 2.6, which gives eigenvalue bounds for distributions of random regular graphs derived from independent random permutations.

Section 3, which is entirely self-contained, develops the theory of size biased couplings for concentration. For the permutation models, it is easy to form size bias couplings that let us apply the results of Section 3. For the uniform model, we construct the necessary couplings in Section 4 using a combinatorial technique called *switchings*. We then apply size biasing in Section 5 to establish Proposition 2.3, a concentration bound for general linear functions of the adjacency matrices of random regular graphs. We specialize this to prove an edge discrepancy bound in Theorem 5.3. Section 6 presents Kahn and Szemerédi’s argument to prove Proposition 2.4, deducing a second eigenvalue bound given a concentration bound like Proposition 2.3.

1.2. Notation, definitions and facts. The *degree* of a vertex in a graph is the number of edges incident to it, or in a weighted graph, the sum of all edge weights

incident to it. A loop in a graph contributes its weight twice to the degree of the vertex. A graph is d -regular if every vertex has degree d . When considering d -regular graphs on n vertices, we always assume that nd is even. We also assume that $n \geq 5$ to avoid some pathologies. A graph is *simple* if it contains no loops or parallel edges.

For an adjacency matrix A , we define the set $\mathcal{N}_A(v)$ to be the neighbors of v in the graph corresponding to A ; when it is clear which graph we are referring to, we omit the A . We define $\overline{\mathcal{N}}_A(v)$ as the vertices which are neither neighbors of v nor v itself in the graph corresponding to A . For $S, T \subseteq [n]$ and an adjacency matrix A , define the edge count

$$(2) \quad e_A(S, T) = \sum_{u \in S} \sum_{v \in T} A_{uv}.$$

Note that this can count the same edge twice if $S \cap T \neq \emptyset$.

By the invariance of the law of a uniform random d -regular simple graph on n vertices under the swapping of vertex labels, the neighbors of v in such a graph form a set of d vertices sampled uniformly from $[n] \setminus \{v\}$, where $[n]$ denotes the set of integers $\{1, \dots, n\}$. Thus, the probability of a given edge uv appearing in the graph is $d/(n-1)$.

2. Spectral concentration from measure concentration. The main result of the present work is to extend the bound $O(\sqrt{d})$ on the second eigenvalue of a random d -regular graph to the uniform model with $d = O(n^{2/3})$. Our argument follows a streamlined version of the Kahn–Szemerédi approach, with all of the necessary concentration estimates unified into an assumption that we call the “uniform tails property” (Definition 2.1 below), which gives uniform tail bounds for linear functions of the adjacency matrix. This property is shown to hold in different models of random regular graphs in Proposition 2.3. In Section 2.2, we state a technical result, Proposition 2.4, which gives a bound on $\lambda(A)$ holding with high probability for any random regular multigraph satisfying the uniform tails property. Based on these results, whose proofs appear later in the paper, we prove Theorem 1.1, the second eigenvalue bound for the uniform model, as well as Theorem 2.6, for permutation models.

2.1. The uniform tails property. We will prove high probability bounds of the optimal order $O(\sqrt{d})$ for random regular graph models satisfying the following concentration property. As is common in the literature on concentration of measure, we phrase our tail bounds in terms of the function

$$(3) \quad h(x) = (1+x) \log(1+x) - x, \quad \text{for } x \geq -1.$$

An $n \times n$ matrix Q is associated to a linear function f_Q of the entries of a matrix M as follows:

$$(4) \quad f_Q(M) = \sum_{u,v=1}^n Q_{uv} M_{uv}.$$

When M is symmetric, we lose no generality in restricting to symmetric matrices Q .

DEFINITION 2.1 (Uniform tails property). Let M be a random symmetric $n \times n$ matrix with nonnegative entries. With f_Q as in (4), write

$$(5) \quad \mu := \mathbf{E}f_Q(M) = f_Q(\mathbf{E}M) \quad \text{and} \quad \tilde{\sigma}^2 := f_{Q \circ Q}(\mathbf{E}M) = \sum_{u,v=1}^n Q_{uv}^2 \mathbf{E}M_{uv},$$

where \circ denotes the Hadamard (entrywise) matrix product. Say that M satisfies the *uniform tails property* $\text{UTP}(c_0, \gamma_0)$ with $c_0 > 0$, $\gamma_0 \geq 0$, if the following holds: for any $a, t > 0$ and for any $n \times n$ symmetric matrix Q with entries $Q_{uv} \in [0, a]$ for all $u, v \in [n]$,

$$(6) \quad \begin{aligned} & \mathbf{P}[f_Q(M) \geq (1 + \gamma_0)\mu + t], \\ & \mathbf{P}[f_Q(M) \leq (1 - \gamma_0)\mu - t] \leq \exp\left(-c_0 \frac{\tilde{\sigma}^2}{a^2} h\left(\frac{at}{\tilde{\sigma}^2}\right)\right). \end{aligned}$$

We will say that M satisfies the *uniform upper tail property* $\text{UUTP}(c_0, \gamma_0)$ if the above bound holds for the first quantity on the left-hand side, with no assumption on the lower tail.

REMARK 2.2. From the bound $h(x) \geq \frac{x^2}{2(1+x/3)}$ for $x \geq 0$, the bound (6) implies

$$(7) \quad \mathbf{P}[|f_Q(M) - \mu| \geq \gamma_0\mu + t] \leq 2 \exp\left(-\frac{c_0 t^2}{2(\tilde{\sigma}^2 + \frac{1}{3}at)}\right).$$

However, (6) is superior for large t —a fact we will use to establish a key graph regularity property (see Lemma 6.4).

The uniform tails property is closely related to extensive work in the literature on *Hoeffding's combinatorial statistic*, defined as $f_Q(P)$ with P a uniform random $n \times n$ permutation matrix and Q a fixed $n \times n$ matrix with bounded entries. See Remark 5.2 for a lengthier discussion.

In addition to proving the uniform tails property for uniform random regular graphs, we will show it holds for various random regular graphs derived from random permutations. In these models, for $d \geq 2$ even, we let $P_1, \dots, P_{d/2}$ be independent random $n \times n$ permutation matrices, and we put $A = \sum_{k=1}^{d/2} (P_k + P_k^\top)$. Note that A is the adjacency matrix of a *multigraph*, with loops and parallel edges.

First, we will consider the case when the permutations are uniform over the symmetric group on n elements. This is frequently called the *permutation model* and is considered in [23, 27, 28] and elsewhere. We will call it the *uniform permutation model* here. Note that we then have $\mathbf{E}A_{uv} = d/n$ for all $u, v \in [n]$, while

for the uniform model we have $\mathbf{E}A_{uv} = d/(n-1)$ for $u \neq v$, giving rise to slightly different values of the quantities μ and $\tilde{\sigma}^2$ in Definition 2.1.

Next, we prove the uniform tails property for graphs derived from permutations π with distribution constant on conjugacy class, that is, that satisfy

$$\sigma^{-1}\pi\sigma =_d \pi \quad \text{for all permutations } \sigma,$$

and are fixed-point-free, that is, $\pi(u) \neq u$ for all u a.s. The permutation matrices $P_1, \dots, P_{d/2}$ are independently created from random permutations satisfying this property, but they need not be identically distributed. As for the uniform model, $\mathbf{E}A_{uv} = d/(n-1)$ for $u \neq v$. One example of such a graph model is given when the permutations are uniformly distributed over the set of all fixed-point-free involutions. The graphs thus produced have edges given by independently choosing $d/2$ uniformly random matchings of the vertices, and then forming two parallel edges between each pair of matched vertices. Dividing the resulting adjacency matrix by two, we have the random regular graph model $\mathcal{I}_{n,d}$ considered in [27], Theorem 1.3. Another example is when the permutations are uniformly distributed permutations with one long cycle, the graph model considered in [27], Theorem 1.2.

The following propositions state that the uniform tails property holds with appropriate c_0, γ_0 for the various random regular multigraph models we consider. Parts (a) and (b) are proved in Section 5.1, and part (c) is deduced from a stronger result, Theorem 5.1, in Section 5.2.

PROPOSITION 2.3. *Let A be the adjacency matrix of a random d -regular multigraph on n vertices.*

(a) *If A is drawn from the uniform permutation model, then it satisfies $\text{UTP}(\frac{1}{4}, 0)$.*

(b) *If A is drawn from the permutation model with distribution constant on conjugacy class and fixed-point-free, then it satisfies $\text{UTP}(\frac{1}{8}, 0)$.*

(c) *If A is drawn from the uniform model, then it satisfies $\text{UTP}(c_0, \gamma_0)$ with*

$$(8) \quad c_0 = \frac{1}{6} \left(1 - \frac{d+1}{n-1} \right), \quad \gamma_0 = \frac{d+1}{n-d-2}.$$

2.2. High level proofs of the spectral gap for the uniform and permutation models. The following proposition shows that $\lambda(A) = O(\sqrt{d})$ with high probability for a wide class of distributions on random d -regular multigraphs satisfying the uniform tails property for suitable $c_0 > 0, \gamma_0 \geq 0$. The setup is sufficiently general to cover all the graph models we consider here; hence, in combination with Proposition 2.3 it yields control of $\lambda(A)$. The assumptions also cover any random regular multigraph whose expected adjacency matrix has uniformly bounded entries and is close in the Hilbert–Schmidt norm to a constant matrix. Recall that the Hilbert–Schmidt norm of a matrix B is given by $\|B\|_{\text{HS}} = (\sum_{u,v} B_{uv}^2)^{1/2}$. We let $\mathbf{1} = (1, \dots, 1)^T \in \mathbb{R}^n$ denote the all-ones vector.

PROPOSITION 2.4 (Spectral concentration from measure concentration). *Let A be the adjacency matrix of a random d -regular multigraph on n vertices. Assume that the following hold for some constants $c_0 > 0$, $a_1 \geq 1$, $a_2, a_3 \geq 0$:*

1. $\mathbf{E}A_{uv} \leq a_1 \frac{d}{n}$ for all $u, v \in [n]$;
2. $\|\mathbf{E}A - \frac{d}{n} \mathbf{1}\mathbf{1}^\top\|_{\text{HS}} \leq a_2 \sqrt{d}$;
3. A has $\text{UTP}(c_0, a_3/\sqrt{d})$.

Then for all $K > 0$ and some $\alpha > 0$ sufficiently large depending on K, c_0, a_1, a_2, a_3 ,

$$(9) \quad \mathbf{P}[\lambda(A) \geq \alpha\sqrt{d}] \leq n^{-K} + 4e^{-n}.$$

REMARK 2.5. The above proposition deduces an upper tail bound on the spectral gap from uniform tail bounds on functionals of the form (4). In the other direction, since $\lambda(A)$ is the supremum over $x \in \langle \mathbf{1} \rangle^\perp \setminus \{0\}$ of the Rayleigh quotients $x^\top Ax / |x^\top Ax|$, on the event that $\lambda(A) \leq \alpha\sqrt{d}$ we have the uniform bound $f_{xx^\top}(A) \leq \alpha\sqrt{d}$ for unit vectors $x \in \mathbf{S}^{n-1} \cap \langle \mathbf{1} \rangle^\perp$, where here and in the sequel $\langle y \rangle$ denotes the span of a vector y . (Taking x to be constant on a set of vertices S gives the well-known *expander mixing lemma* [37], Lemma 2.5.) Thus, (9) implies a uniform polynomial upper tail for the random variables $f_Q(A)$ for the case that Q is a rank-1 projection. Of course, the advantage of (9) over the uniform tails property is that we obtain a single high probability event on which all rank-1 functionals $f_Q(A)$ are bounded.

The proof of this proposition is deferred to Section 6. Combining Proposition 2.4 and Proposition 2.3, we deduce the following results.

THEOREM 2.6 (Spectral gap for the permutation models). *Consider either of the following two distributions for A :*

(a) *For all $n \geq 5$ and all even $d \geq 2$, let $A = \sum_{k=1}^{d/2} (P_k + P_k^\top)$ be a random d -regular multigraph from the uniform permutation model.*

(b) *For all $n \geq 5$ and all even $2 \leq d \leq C_0 n$ for some constant C_0 , let $A = \sum_{k=1}^{d/2} (P_k + P_k^\top)$ be a random d -regular multigraph from the permutation model with permutations constant on conjugacy class and fixed-point-free.*

In both cases, for any $K > 0$, there is a constant α sufficiently large, depending only on K and in the second case also on C_0 , such that

$$\mathbf{P}[\lambda(A) \geq \alpha\sqrt{d}] \leq n^{-K} + 4e^{-n}.$$

PROOF. For case (a), for each $u, v \in [n]$ we have $\mathbf{E}A_{uv} = d/n$. Together with Proposition 2.3(a), this means we can apply Proposition 2.4 with $a_1 = 1$ and $a_2 = a_3 = 0$, and the result follows.

For case (b), we have $\mathbf{E}A_{uv} = d/(n-1)$ for all $u \neq v$ and $\mathbf{E}A_{uu} = 0$ for all $u \in [n]$, so may take $a_1 = 2$, say. Then we can compute

$$\left\| \mathbf{E}A - \frac{d}{n} \mathbf{1}\mathbf{1}^\top \right\|_{\text{HS}} = \frac{d}{\sqrt{n-1}},$$

and set $a_2 = \sqrt{2C_0}$. By Proposition 2.3(b), we may take $a_3 = 0$, completing the proof. \square

Case (a) of Theorem 2.6 was previously shown in [23]. The proof there also uses the Kahn–Szemerédi approach, with the necessary concentration proven via martingale methods and by direct evaluation of the moment generating function. Our Stein’s method machinery makes the proof much simpler: contrast our proof of Proposition 2.3(a) with those of Theorem 26 and Lemma 30 in [23].

It is natural to ask about relaxing the condition in Theorem 2.6(b) that the permutations be fixed-point-free. It seems possible to remove this condition entirely from Proposition 2.3(b), at the cost of a significantly more complicated construction of the size bias coupling. On the other hand, if the number of fixed points is of larger order than n/\sqrt{d} , then the resulting matrix models have a larger spectral gap than $O(\sqrt{d})$ from the terms along the diagonal. We expect that it is possible to prove a version of Theorem 2.6(b) for permutations having $O(n/\sqrt{d})$ fixed points, though we have not pursued it here.

We prove Theorem 1.1 along the same lines as Theorem 2.6, combining Proposition 2.4 and Proposition 2.3.

PROOF OF THEOREM 1.1. Following the proof of Theorem 2.6, for the first two conditions in Proposition 2.4 we can take $a_1 = 2$ and $a_2 = 1$. By Proposition 2.3(c), A has UTP(c_0, γ_0) with the parameters (8). Now let $C_0, K > 0$, and assume $1 \leq d \leq C_0 n^{2/3}$. From Remark 1.3, we may also assume $d \leq n/2$. Applying these bounds on d , for all n sufficiently large we have

$$(10) \quad \gamma_0 = \frac{d+1}{n-d-2} \leq \frac{10d}{n} \leq \frac{10C_0^{3/2}}{\sqrt{d}}.$$

(The first inequality holds for all $n \geq 7$ and $1 \leq d \leq n/2$.) Note that this is where we used the assumption that $d \leq C_0 n^{2/3}$. See Remark 6.10 for more on why we require $d = O(n^{2/3})$.

Hence, we may apply Proposition 2.4 with $a_3 = 2C_0^{3/2}$. We can also shrink c_0 to some constant independent of n (say $1/12$). Now having fixed the parameters c_0, a_1, a_2 as constants, from Proposition 2.4 applied with $K+1$ in the role of K , we may take α sufficiently large depending only on C_0, K such that $\lambda(A) \leq \alpha\sqrt{d}$ except with probability at most $n^{-K-1} + 4e^{-n}$. The result follows from this. \square

REMARK 2.7. To get the explicit values of α , we refer to Remark 6.9 for the explicit value of α in Proposition 2.4, assuming for now that $d \leq n/2$. First, we evaluate (80). Note that $\gamma_0 \leq 10d/n \leq 10$. Thus, (80) gives

$$\begin{aligned} \alpha_0 &\leq 16 + 32(2)(1 + e^2(11)^2) + 128(12)(11)(K + 5)(1 + e^{-2}) \\ &\leq 153,214 + 76,484K. \end{aligned}$$

(Strictly speaking, we are applying Proposition 2.4 with $a_3 = 10d^{3/2}/n$ rather than the larger value $a_3 = 10C_0^{3/2}$ here.) From (79), we then get

$$\begin{aligned} \alpha &\leq 3(\alpha_0 + 3) + \max(30C_0^{3/2}, 768) \\ &\leq 459,651 + 229,452K + \max(30C_0^{3/2}, 768). \end{aligned}$$

Choosing n large enough requires us to have $n \geq 7$ and $1/n + 4n^K e^{-n} \leq 1$, and these conditions hold when $n \geq 7 + K^2$, for instance. As a consequence of Remark 1.3, increasing α by one is more than enough to drop the requirement that $d \leq n/2$, leading to the constants stated in Remark 1.2. As one might suspect after seeing these bounds, we have not made an effort to optimize these constants.

3. Concentration by size biased couplings.

3.1. *Introduction to size biased couplings.* If X is a nonnegative random variable with finite mean $\mu > 0$, we say that X^s has the X -size biased distribution if

$$\mathbf{E}[Xf(X)] = \mu\mathbf{E}[f(X^s)]$$

for all functions f such that the left-hand side above exists. The law $\mathcal{L}(X^s)$ always exists for such X , as can be seen by (equivalently) specifying the distribution ν^s of X^s as the one with Radon–Nikodym derivative $d\nu^s/d\nu = x/\mu$, where ν is the distribution of X . Many appearances of the size biased distribution in probability and statistics, some quite unexpected, are reviewed in [4].

For such an X , we say the pair of random variables (X, X^s) defined on a common space is a size biased coupling for X when X^s has the X -size biased distribution. Couplings of this sort were used throughout the history of Stein’s method (see [49], pp. 89–90, [5] and [6]), though the connection to size biasing was not made explicit until [33]. See [18] or [48] for surveys of Stein’s method including size biased coupling.

Proving concentration using couplings borrowed from Stein’s method began with the work of [47], and, absent the Stein equation tying the analysis to a particular distribution, in [16]. Using Stein’s classical exchangeable pair, [16] and [17] show concentration for Hoeffding’s combinatorial statistic, in the Curie–Weiss and Ising models, and for the number of triangles in the Erdős–Rényi random graph.

Similar techniques are also used in [21] to show concentration for statistics of random regular digraphs.

We say that a size biased coupling (X, X^s) is bounded when there exists a constant c such that $X^s \leq X + c$ almost surely. It is shown in [3] that the existence of such a coupling implies that X is concentrated, an improvement of a result in [30], where the idea originated. We will present concentration bounds that generalize the results in [3], relaxing the boundedness assumption and giving a Bennett-type inequality (see the following section for the details of what this means). Previous work for concentration by unbounded size biased couplings was limited to [31], with a construction particular to the example treated, and dependent on a negative association property holding. There was no previous Bennett-type inequality by size biasing, though [32] gives one by the related method of zero biasing; see Remark 5.2.

At the heart of nearly all applications of size biasing is a construction of a coupling for a sum $X = \sum_{i=1}^n X_i$, as first outlined in [33], Lemma 2.1. We follow the treatment in [4], Section 2.3. Suppose that ν is the distribution of a random vector (X_1, \dots, X_n) with nonnegative entries each with positive mean. We say that the distribution $\nu^{(i)}$ defined by its Radon–Nikodym derivative

$$\frac{d\nu^{(i)}}{d\nu}(x_1, \dots, x_n) = \frac{x_i}{\mathbf{E}X_i}$$

has the distribution of (X_1, \dots, X_n) size biased by X_i . One can think of $\nu^{(i)}$ as the distribution of the random vector formed by size biasing X_i and then giving the vector of other entries its distribution conditional on the new value of X_i .

LEMMA 3.1. *Let X_1, \dots, X_n be nonnegative random variables with positive means, and let $X = \sum_{i=1}^n X_i$. For each i , let $(X_1^{(i)}, \dots, X_n^{(i)})$ have the distribution of (X_1, \dots, X_n) size biased by X_i . Independent of everything else, choose an index I with $\mathbf{P}[I = i] = \mathbf{E}X_i / \mathbf{E}X$. Then $X^s = \sum_{i=1}^n X_i^{(I)}$ has the size biased distribution of X .*

This reduces the problem of forming a size biased coupling for X to forming couplings of (X_1, \dots, X_n) with $(X_1^{(i)}, \dots, X_n^{(i)})$ for each i . We demonstrate now how to do this when X_1, \dots, X_n are independent, but it is often possible to do even when they are not.

EXAMPLE 3.2 (Size biased couplings for independent sums). Suppose $X = \sum_{i=1}^n X_i$ with the summands independent. Let $\mu = \mathbf{E}X$ and $\mu_i = \mathbf{E}X_i$. Let $X_i^{(i)}$ have the X_i -size biased distribution, and make it independent of all other random variables. For $i \neq j$, let $X_j^{(i)} = X_j$. By the independence of the random variables, $(X_1^{(i)}, \dots, X_n^{(i)})$ has the distribution of (X_1, \dots, X_n) size biased by X_i . With I and

X^s as in Lemma 3.1, we have a size biased coupling (X, X^s) . Note that X^s can be expressed as

$$X^s = X - X_I + X_I^{(I)}.$$

In our applications of size biasing in Section 5, we will have $X_i = a_i F_i$, where F_i is a nontrivial indicator and $a_i \geq 0$. In this case, the X_i -size biased transform is a_i , and the distribution of $(X_1^{(i)}, \dots, X_n^{(i)})$ can be described by specifying that $X_i^{(i)} = a_i$ and $(X_j^{(i)})_{j \neq i}$ is distributed as $(X_j)_{j \neq i}$ conditional on $F_i = 1$.

3.2. *New concentration results by size biased couplings.* Throughout this section, X is a nonnegative random variable with nonzero, finite mean μ . We say the size biased coupling (X, X^s) is (c, p) -bounded for the upper tail if

$$(11) \quad \text{for any } x, \quad \mathbf{P}[X^s \leq X + c \mid X^s \geq x] \geq p,$$

and (c, p) -bounded for the lower tail if

$$(12) \quad \text{for any } x, \quad \mathbf{P}[X^s \leq X + c \mid X \leq x] \geq p.$$

The probabilities in (11) and (12) conditional on null events may be defined arbitrarily. In Theorems 3.3 and 3.4, we recall the definition:

$$(13) \quad h(x) = (1 + x) \log(1 + x) - x, \quad x \geq -1,$$

which satisfies

$$(14) \quad \begin{aligned} h(x) &\geq \frac{x^2}{2(1+x/3)} && \text{for all } x \geq 0, \quad \text{and} \\ h(x) &\geq x^2/2 && \text{for } -1 \leq x \leq 0; \end{aligned}$$

see the second and first inequalities of Exercise 2.8 of [13], respectively.

THEOREM 3.3. (a) *If X admits a (c, p) -bounded size biased coupling for the upper tail, then for all $x \geq 0$,*

$$(15) \quad \mathbf{P}\left[X - \frac{\mu}{p} \geq x\right] \leq \exp\left(-\frac{\mu}{cp} h\left(\frac{px}{\mu}\right)\right) \leq \exp\left(-\frac{x^2}{2c(x/3 + \mu/p)}\right).$$

(b) *If X admits a (c, p) -bounded size biased coupling for the lower tail, then for all $0 \leq x < p\mu$,*

$$(16) \quad \mathbf{P}[X - p\mu \leq -x] \leq \exp\left(-\frac{p\mu}{c} h\left(-\frac{x}{p\mu}\right)\right) \leq \exp\left(-\frac{x^2}{2pc\mu}\right).$$

The special case $p = 1$ yields Theorem 1.3 and Corollary 1.1 from [3], with the second inequality in (15) a slight improvement to (12) of the latter, through the use of (14) in place of [3], Lemma 4.2.

Theorem 3.3 is an analogue of a bound for sums of independent variables due to Hoeffding [36], Theorem 1, line (2.2), with tails that incorporate the mean μ as well as an L^∞ bound on the summands (from Example 3.2 we see this role is played by c in the above theorem). Hence, for sums of independent nonnegative variables whose expectation is small in comparison to their L^∞ norms, Theorem 3.3 provides better estimates than Azuma–Hoeffding-type (or “bounded differences”) inequalities, such as [36], line (2.3) (the bound that is widely referred to as “Hoeffding’s inequality”). See also Section 4 of [7] for a fuller comparison of concentration results obtained by bounded size bias couplings to those obtained via more classical means.

To prove concentration of the light couples in the Kahn–Szemerédi argument (see Section 6), we will need tail bounds incorporating the *variance* rather than the mean. For sums of independent variables, such bounds are provided⁴ by Bernstein’s inequality [11] or Bennett’s inequality [10], equation (8b). In previous applications of the Kahn–Szemerédi argument, [28] and [15] used ad hoc arguments working directly with the moment generating function, and [42] and [23] used Freedman’s inequality, the martingale version of Bennett’s inequality. We instead develop the following Bennett-type inequality by size biased coupling. Let x^+ denote $\max(0, x)$ in the following theorem.

THEOREM 3.4. *Let (X, X^s) be a size biased coupling with $\mathbf{E}X = \mu$, and let \mathcal{B} be an event on which $X^s - X \leq c$. Let $D = (X^s - X)^+$, and suppose that $\mathbf{E}[D\mathbf{1}_{\mathcal{B}} | X] \leq \tau^2/\mu$ a.s.:*

(a) *If $\mathbf{P}[\mathcal{B} | X^s] \geq p$ a.s., then for $x \geq 0$*

$$(17) \quad \begin{aligned} \mathbf{P}\left[X - \frac{\mu}{p} \geq x\right] &\leq \exp\left(-\frac{\tau^2}{pc^2}h\left(\frac{pcx}{\tau^2}\right)\right) \\ &\leq \exp\left(-\frac{x^2}{2c(x/3 + \tau^2/cp)}\right). \end{aligned}$$

(b) *If $\mathbf{P}[\mathcal{B} | X] \geq p$ a.s., then for $0 \leq x \leq p\mu$*

$$(18) \quad \begin{aligned} \mathbf{P}[X - p\mu \leq -x] &\leq \exp\left(-\frac{\tau^2}{c^2}h\left(\frac{cx}{\tau^2}\right)\right) \\ &\leq \exp\left(-\frac{x^2}{2c(x/3 + \tau^2/c)}\right). \end{aligned}$$

⁴As explained in [10], Bernstein’s work [11] was originally published in Russian and went largely unnoticed in the English-speaking world. Bennett himself was unable to access Bernstein’s original paper, but proved a strengthened form of Bernstein’s result in [10].

We use the notation τ^2 to suggest that it plays the role of the variance in Bennett's inequality. In our applications of Theorem 3.4 in this paper, τ^2 is indeed on the same order as $\text{Var } X$; see Example 3.5 for a simple example.

We compare Theorems 3.3 and 3.4, assuming $c = 1$ by rescaling if necessary. Note that by taking $\mathcal{B} = \{X^s \leq X + 1\}$ in the former, we have $\mathbf{E}[D\mathbf{1}_{\mathcal{B}} | X] \leq 1$, and hence one may set $\tau^2 = \mu$. Doing so, the upper bound (17) of Theorem 3.4 recovers (15) of Theorem 3.3 when $\mathbf{P}[X^s - X \leq 1 | X^s] \geq p$. For the lower tail, one can easily verify that

$$\exp(-\mu h(x/\mu)) \leq \exp(-p\mu h(-x/p\mu)) \quad \text{for all } 0 \leq x < p\mu,$$

showing the left tail bound of Theorem 3.3 superior to that of Theorem 3.4 in the absence of a better bound on $\mathbf{E}[D\mathbf{1}_{\mathcal{B}} | X]$.

3.3. Examples. We now give some examples to give a sense of what can be done with Theorems 3.3 and 3.4. Example 3.5 applies Theorem 3.4 to recover a weakened form of Bennett's inequality for independent summands. Our theorems prove concentration around a shifted mean; Examples 3.6 and 3.7 demonstrate that this is unavoidable. In Example 3.8, we give a simple application of Theorem 3.3 to Erdős–Rényi graphs to show that our theory has applications beyond the ones we give in Section 5.

EXAMPLE 3.5 (Weakened form of Bennett's inequality). Suppose $X = \sum_{i=1}^n X_i$ with the summands independent and contained in $[0, 1]$. Let $\mu = \mathbf{E}X$ and $\mu_i = \mathbf{E}X_i$. Let X_i^s have the size biased distribution of X_i , and make it independent of all other random variables. Choose $I \in [n]$ independently of all else, taking $\mathbf{P}[I = i] = \mu_i/\mu$. As in Example 3.2, the pair (X, X^s) is a size biased coupling with $X^s = X - X_I + X_I^s$.

Since X_i^s has the same support as X_i , we have $X^s \leq X + 1$. In applying Theorem 3.4, we can then take the event \mathcal{B} to be the entire probability space, and obtain

$$\begin{aligned} \mathbf{E}[(X^s - X)^+ | X] &= \mathbf{E}[(X_I^s - X_I)^+ | X] \\ &\leq \mathbf{E}[X_I^s | X] \\ &= \mathbf{E}[X_I^s] = \frac{1}{\mu} \sum_{i=1}^n \mu_i \mathbf{E}X_i^s. \end{aligned}$$

From the definition of the size biased transform, $\mathbf{E}X_i^s = \mathbf{E}X_i^2/\mu_i$. Thus,

$$\mathbf{E}[(X^s - X)^+ | X] \leq \frac{1}{\mu} \sum_{i=1}^n \mathbf{E}X_i^2.$$

We then apply Theorem 3.4 with $c = 1$, $p = 1$, and $\tau^2 = \sum_{i=1}^n \mathbf{E}X_i^2$ to show that

$$\mathbf{P}[X - \mu \geq t], \mathbf{P}[X - \mu \leq -t] \leq \exp\left(-\tau^2 h\left(\frac{t}{\tau^2}\right)\right),$$

which would be Bennett's inequality if τ^2 were $\text{Var } X$ rather than the larger $\sum_{i=1}^n \mathbf{E}X_i^2$ (see [13], Section 2.7).

When applied with $p < 1$, Theorems 3.3 and 3.4 show concentration of X not around its mean μ , but rather around μ/p for the upper tail and $p\mu$ for the lower tail. The following two examples demonstrate that this behavior may reflect the true nature of X , thus showing these theorems to be unimprovable in this sense.

EXAMPLE 3.6 (Upper tail concentration around μ/p). Let $Z \sim \text{Poi}(\lambda)$ and $B \sim \text{Bernoulli}(1/2)$ be independent, and define $X = BZ$. We show $(X, Z + 1)$ is a size biased coupling. By a well-known property of the Poisson distribution (e.g., see (6) of [4]), $Z + 1$ has the Z -size biased distribution. Mixing a distribution with the measure δ_0 does not change its size biased transform (see Lemma 2.6 of [4]). Thus, $Z + 1$ also has the size biased distribution of X , and the coupling $(X, Z + 1)$ is $(1, 1/2)$ -bounded for the upper tail. Theorem 3.3 then shows exponential decay for the upper tail of X starting at $\mu/p = 2\mu = \lambda$, reflecting its actual behavior.

EXAMPLE 3.7 (Lower tail concentration around $p\mu$). Let $N > 1$ and let X_1, \dots, X_n be i.i.d. with distribution

$$X_i = \begin{cases} 0 & \text{with probability } 1/2 - \epsilon, \\ 1 & \text{with probability } 1/2, \\ N & \text{with probability } \epsilon, \end{cases}$$

where $\epsilon = 1/(2N)$. As $\mathbf{E}X_i = 1$, for $i = 1, \dots, n$ the variables

$$X_i^s = \begin{cases} 1 & \text{with probability } 1/2, \\ N & \text{with probability } 1/2 \end{cases}$$

have the X_i -size biased distribution. Let X_1^s, \dots, X_n^s be independent of each other and of X_1, \dots, X_n and set $X = X_1 + \dots + X_n$. Then by Lemma 3.1, choosing I uniformly from $\{1, \dots, n\}$, independent of all other variables, we obtain a size biased coupling (X, X^s) by defining

$$X^s = X - X_I + X_I^s.$$

This coupling is $(1, 1/2)$ -bounded for the lower tail. Theorem 3.3 shows concentration starting at $p\mathbf{E}X = n/2$. When N is large, X is nearly distributed as $n(n, 1/2)$, so this is the correct behavior.

The next example gives a lower tail bound for the number of isolated vertices in an Erdős-Rényi graph. The bound is inferior to the one given in [31], but we can get it with very little work.

EXAMPLE 3.8. Let G be a random graph on n vertices with each edge included independently with probability p . Let X be the number of isolated vertices in G . To form a size biased coupling, select a random vertex V from $\{1, \dots, n\}$ independently of G , and form G^s by deleting from G all edges incident to V . Let X^s be the number of isolated vertices in G^s . As G^s is distributed as G conditional on vertex V being isolated, X^s has the X -size biased distribution by Lemma 3.1.

Call a vertex a leaf if it has degree one. In any (deterministic) graph, we claim that at most $1/3$ of the vertices are connected to two or more leaves. To see this, let l be the number of leaves and m the number of vertices connected to two or more leaves. The claim then follows from the observation that $l \geq 2m$.

Thus, conditional on G , there is at most a $1/3$ chance that V is connected to two or more leaves. Deleting the edges incident to V isolates V as well as any neighboring leaves, giving us

$$\mathbf{P}[X^s - X \leq 2 \mid G] \geq 2/3.$$

Since X is measurable with respect to G , the coupling is $(2, 2/3)$ -bounded for the lower tail, and Theorem 3.3 gives the bound

$$(19) \quad \mathbf{P}\left[X - \frac{2\mu}{3} \leq -t\right] \leq \exp\left(-\frac{\mu}{3}h\left(-\frac{3t}{2\mu}\right)\right) \leq \exp\left(-\frac{3t^2}{8\mu}\right)$$

with $\mu = \mathbf{E}X$.

A variation on this argument shows that the coupling is $(k, k/(k+1))$ -bounded for the lower tail. Applying this fact with larger values of k yields a concentration bound around a quantity closer to the true mean than in (19), but with a worse constant in the exponent.

3.4. *Proofs.* We start with a modified version of [3], Lemma 2.1.

LEMMA 3.9. *If X admits a (c, p) -bounded size biased coupling for the upper tail, then*

$$(20) \quad \forall x > 0, \quad \mathbf{P}[X \geq x] \leq \frac{\mu}{px} \mathbf{P}[X \geq x - c]$$

and if X admits a (c, p) -bounded size biased coupling for the lower tail, then

$$(21) \quad \forall x, \quad \mathbf{P}[X \leq x] \leq \frac{x+c}{p\mu} \mathbf{P}[X \leq x+c].$$

PROOF. For (X, X^s) , the upper tail coupling,

$$px \mathbf{P}[X \geq x] = px \mathbf{E}\mathbf{1}_{\{X \geq x\}} \leq p \mathbf{E}[X \mathbf{1}_{\{X \geq x\}}] = p\mu \mathbf{P}[X^s \geq x].$$

If $\mathbf{P}[X^s \geq x] = 0$, then $\mathbf{P}[X \geq x] = 0$, since the support of X contains the support of X^s . Thus, in this case (20) holds trivially. If $\mathbf{P}[X^s \geq x] > 0$, then we apply (11) to get

$$\begin{aligned} px\mathbf{P}[X \geq x] &\leq \mu\mathbf{P}[X^s \leq X + c \mid X^s \geq x]\mathbf{P}[X^s \geq x] \\ &= \mu\mathbf{P}[X^s \leq X + c \text{ and } X^s \geq x] \\ &\leq \mu\mathbf{P}[X \geq x - c]. \end{aligned}$$

The proof for the lower tail follows by a similar modification of [3], Lemma 2.1. \square

Inequality (20) corresponds to (14) of [3] with μ replaced by μ/p , and inequality (21) corresponds to (15) of [3], Lemma 2.1, with μ replaced by $p\mu$. As iteration of the bounds (14) and (15) results in [3], Theorems 1.1 and 1.2, respectively, Lemma 3.9 implies that the bounds of these theorems hold more generally with this replacement. In particular, replacing the functions $u(x, \mu, c)$ and $l(x, \mu, c)$ by $u(x, \mu/p, c)$ and $l(x, p\mu, c)$ respectively, inequalities (3) and (4) of [3], Theorem 1.1, hold over the ranges $x \geq \mu/p$ and $0 \leq x \leq p\mu$, with k as given in (1) with the mean μ replaced by μ/p and $p\mu$, under the upper and lower tail conditions (11) and (12), respectively. Likewise, under the upper and lower tail conditions (11) and (12), [3], Theorem 1.2, holds with all occurrences of the mean μ replaced by μ/p and $p\mu$ in (7) and (8), with equalities holding if and only if $x - \mu/p$ and $x - p\mu$ are integers, respectively.

Theorem 3.3 generalizes [3], Theorem 1.3 and Corollary 1.1, by these same replacements. As those results are not shown there as a direct consequence of (14) and (15), we provide separate arguments, beginning by applying Lemma 3.9 to prove that (11) implies that the moment generating function $M(\beta) = \mathbf{E}e^{\beta X}$ of X is finite. The following proof is essentially the same as that of [3], Corollary 2.1, with μ replaced by μ/p in the upper tail inequality, and using a bound on the upper tail directly rather than bounding that tail using the upper bound product function $u(x, a, c)$.

PROPOSITION 3.10. *If X admits a (c, p) -bounded size bias coupling for the upper tail for some $p > 0$, then the moment generating function $M(\beta)$ is finite for all β .*

PROOF. As $X \geq 0$, the claim is clearly true for $\beta \leq 0$. Let $\beta > 0$ and $x_0 \geq 2\mu e^{\beta c}/p$. As in [3], Corollary 2.1, the idea is that beyond x_0 , for every increase by c , the tail of the distribution of X decreases in probability by enough to make $M(\beta)$ finite. More precisely, by (20), for $x \geq x_0$,

$$\mathbf{P}[X \geq x + c] \leq \frac{\mu}{p(x + c)}\mathbf{P}[X \geq x] \leq \frac{1}{2}e^{-\beta c}\mathbf{P}[X \geq x].$$

By iterating this bound, $\mathbf{P}[X \geq x + ic] \leq 2^{-i} e^{-i\beta c}$. Applying this inequality, we have

$$\begin{aligned}
 M(\beta) &= \mathbf{E}e^{\beta X} \\
 &\leq e^{\beta x} \mathbf{P}[X < x] + \sum_{i=0}^{\infty} e^{\beta(x+(i+1)c)} \mathbf{P}[x + ic \leq X < x + (i+1)c] \\
 &\leq e^{\beta x} \mathbf{P}[X < x] + \sum_{i=0}^{\infty} e^{\beta(x+(i+1)c)} \mathbf{P}[X \geq x + ic] \\
 &\leq e^{\beta x} \mathbf{P}[X < x] + \sum_{i=0}^{\infty} e^{\beta(x+c)} 2^{-i} < \infty.
 \end{aligned}$$

□

LEMMA 3.11. *If X admits a (c, p) -bounded size bias coupling for the upper tail, then*

$$(22) \quad M(\beta) \leq \exp\left[\frac{\mu}{pc}(e^{\beta c} - 1)\right]$$

for all $\beta \geq 0$.

If X admits (c, p) -bounded size bias coupling for the lower tail, then

$$(23) \quad M(\beta) \leq \exp\left[\frac{p\mu}{c}(e^{\beta c} - 1)\right]$$

for all $\beta \leq 0$.

PROOF. Let (X, X^s) be a (c, p) -bounded size biased coupling for the upper tail, and let $\beta \geq 0$. We will bound $M'(\beta)$ in terms of $M(\beta)$. It follows from the finiteness of $M(\beta)$ for all β proved in Proposition 3.10 that $\mu \mathbf{E}e^{\beta X^s} = \mathbf{E}[X e^{\beta X}] = M'(\beta)$. Using $\beta \geq 0$, we have

$$e^{\beta X} = e^{\beta(X^s - (X^s - X))} \geq e^{\beta(X^s - (X^s - X))} \mathbf{1}_{X^s \leq X+c} \geq e^{\beta X^s - c} \mathbf{1}_{X^s \leq X+c},$$

whence

$$\begin{aligned}
 M(\beta) &= \mathbf{E}e^{\beta X} \\
 &\geq \mathbf{E}[e^{\beta X^s - c} \mathbf{1}_{X^s \leq X+c}] \\
 (24) \quad &= \mathbf{E} \int_0^{\infty} \mathbf{1}\{x \leq e^{\beta(X^s - c)} \text{ and } X^s \leq X+c\} dx \\
 &= \int_0^{\infty} \mathbf{P}[x \leq e^{\beta(X^s - c)} \text{ and } X^s \leq X+c] dx.
 \end{aligned}$$

As a consequence of (11),

$$\mathbf{P}[x \leq e^{\beta(X^s - c)} \text{ and } X^s \leq X+c] \geq p \mathbf{P}[x \leq e^{\beta(X^s - c)}].$$

Applying this inequality to (24) gives

$$M(\beta) \geq p \int_0^\infty \mathbf{P}[x \leq e^{\beta(X^s - c)}] dx = p \mathbf{E} e^{\beta(X^s - c)} = \frac{p M'(\beta)}{\mu e^{\beta c}}.$$

Thus,

$$(\log M)'(\beta) = \frac{M'(\beta)}{M(\beta)} \leq \frac{\mu e^{\beta c}}{p},$$

and integrating we obtain

$$\log M(\beta) = \log M(\beta) - \log M(0) \leq \int_0^\beta \frac{\mu e^{cu}}{p} du = \frac{\mu}{pc} (e^{\beta c} - 1).$$

Exponentiating proves (22).

Next, let (X, X^s) be a (c, p) -bounded size bias coupling for the lower tail, and let $\beta \leq 0$. Note that $M(\beta)$ is now finite simply because $\beta \leq 0$, and again $M'(\beta) = \mu \mathbf{E} e^{\beta X^s}$. Now using $e^{\beta X^s} \geq e^{\beta(X+c)} \mathbf{1}_{X^s \leq X+c}$, we obtain

$$\begin{aligned} \frac{M'(\beta)}{\mu} &= \mathbf{E} e^{\beta X^s} \\ &\geq \mathbf{E}[e^{\beta(X+c)} \mathbf{1}_{X^s \leq X+c}] \\ &= \mathbf{E} \int_0^\infty \mathbf{1}\{x \leq e^{\beta(X+c)} \text{ and } X^s \leq X+c\} dx \\ &= \int_0^\infty \mathbf{P}[x \leq e^{\beta(X+c)} \text{ and } X^s \leq X+c] dx. \end{aligned}$$

By (12),

$$\frac{M'(\beta)}{\mu} \geq p \int_0^\infty \mathbf{P}[x \leq e^{\beta(X+c)}] dx = p \mathbf{E} e^{\beta(X+c)} = p e^{\beta c} M(\beta).$$

Therefore,

$$(\log M)'(\beta) \geq p \mu e^{\beta c}$$

and

$$\log M(\beta) = - \int_\beta^0 (\log M)'(u) du \leq \int_\beta^0 -p \mu e^{cu} du = \frac{p \mu}{c} (e^{\beta c} - 1). \quad \square$$

PROOF OF THEOREM 3.3. If X admits a (c, p) -bounded size bias coupling for the upper tail, then by Markov's inequality and Lemma 3.11,

$$\begin{aligned} \mathbf{P}[X - \mu/p \geq x] &= \mathbf{P}[e^{\beta X} \geq e^{\beta(x + \mu/p)}] \\ &\leq e^{-\beta(x + \mu/p)} M(\beta) \\ &\leq \exp\left[\frac{\mu}{pc} (e^{\beta c} - 1) - \beta(x + \mu/p)\right] \end{aligned}$$

for $\beta \geq 0$. Setting $\beta = \log(px/\mu + 1)/c$, which is nonnegative for $x \geq 0$, yields the first inequality in (15). The second inequality in (15) now follows from the first inequality in (14).

To prove (16), for any $\beta \leq 0$,

$$\begin{aligned} \mathbf{P}[X - p\mu \leq -x] &= \mathbf{P}[e^{\beta X} \geq e^{\beta(-x+p\mu)}] \\ &\leq M(\beta)e^{\beta(x-p\mu)} \\ &\leq \exp\left[\frac{p\mu}{c}(e^{\beta c} - 1) + \beta(x - p\mu)\right]. \end{aligned}$$

Setting $\beta = \log(-x/p\mu + 1)/c$, which is nonpositive for $0 \leq x < p\mu$, yields the first inequality in (16). The second inequality in (16) now follows from the second inequality in (14). \square

Next, we turn toward the proof of Theorem 3.4, beginning with the following simple lemma.

LEMMA 3.12. *If $0 \leq y \leq 1$, then for all $x \in \mathbb{R}$,*

$$(25) \quad e^{xy} \leq 1 + (e^x - 1)y$$

and

$$(26) \quad e^{-xy} \geq 1 - (e^x - 1)y.$$

PROOF. The function $f(u) = u^y$ for $u \geq 0$ is concave, and hence it lies below its tangent line at $u = 1$, showing that

$$u^y \leq 1 + (u - 1)y.$$

Substituting $u = e^x$ shows (25).

To prove (26), the function $g(u) = u^{-y}$ is convex, and hence lies above its tangent line at $u = 1$, and the same argument completes the proof. \square

PROOF OF THEOREM 3.4. We start with the upper tail bound, assuming for now that $c = 1$. As $\{X^s \leq X + c\} \supseteq \mathcal{B}$, the hypothesis of (a) implies (11), hence the moment generating function $M(\beta) = \mathbf{E}e^{\beta X}$ of X is finite for all β by Proposition 3.10. Assume $\beta \geq 0$. Applying $\mathbf{P}[\mathcal{B} \mid X^s] \geq p$, we have

$$\mathbf{E}[e^{\beta X^s} \mathbf{1}_{\mathcal{B}}] \geq p\mathbf{E}[e^{\beta X^s}] = \frac{p}{\mu}\mathbf{E}[Xe^{\beta X}] = \frac{p}{\mu}M'(\beta),$$

since by finiteness of the moment generating function we can differentiate inside the expectation. Rewriting this inequality and using the definition of D , we have

$$M'(\beta) \leq \frac{\mu}{p}\mathbf{E}[e^{\beta X^s} \mathbf{1}_{\mathcal{B}}] \leq \frac{\mu}{p}\mathbf{E}[e^{\beta D} e^{\beta X} \mathbf{1}_{\mathcal{B}}].$$

Since $0 \leq D \leq 1$ on \mathcal{B} , we can apply Lemma 3.12 to conclude that

$$\begin{aligned} \mathbf{E}[e^{\beta D} \mathbf{1}_{\mathcal{B}} | X] &\leq \mathbf{E}[(1 + (e^{\beta} - 1)D) \mathbf{1}_{\mathcal{B}} | X] \\ &= 1 + (e^{\beta} - 1) \mathbf{E}[D \mathbf{1}_{\mathcal{B}} | X] \\ &\leq 1 + \frac{\tau^2}{\mu} (e^{\beta} - 1). \end{aligned}$$

Thus,

$$M'(\beta) \leq \frac{1}{p} (\mu + \tau^2 (e^{\beta} - 1)) M(\beta)$$

and

$$\begin{aligned} \log M(\beta) &= \int_0^{\beta} (\log M)'(u) du \\ &\leq \int_0^{\beta} \frac{1}{p} (\mu + \tau^2 (e^u - 1)) du \\ &= \frac{1}{p} (\mu\beta + \tau^2 (e^{\beta} - 1 - \beta)). \end{aligned}$$

By Markov's inequality,

$$\mathbf{P}[X - \mu/p \geq x] \leq M(\beta) e^{-\beta(x + \mu/p)} \leq \exp\left(\frac{\tau^2}{p} (e^{\beta} - 1 - \beta) - \beta x\right).$$

Substituting $\beta = \log(1 + px/\tau^2)$, which is nonnegative for $x \geq 0$, yields

$$(27) \quad \mathbf{P}[X - \mu/p \geq x] \leq \exp\left[-\frac{\tau^2}{p} h\left(\frac{px}{\tau^2}\right)\right].$$

Now, we consider the general case $c > 0$. We obtain the first inequality in (17) by rescaling and applying (27):

$$\mathbf{P}[X - \mu/p \geq x] = \mathbf{P}[X/c - \mu/pc \geq x/c] \leq \exp\left[-\frac{\tau^2}{pc^2} h\left(\frac{px/c}{\tau^2/c^2}\right)\right],$$

noting that we must replace τ by τ/c when applying (27) to X/c . The second inequality now follows by the first inequality in (14).

Next, we prove the lower tail bound, again assuming $c = 1$. Using that the moment generating function $M(-\beta)$ exists for all $\beta \geq 0$, we have

$$\begin{aligned} M'(-\beta) &= \mu \mathbf{E} e^{-\beta X^s} \geq \mu \mathbf{E}[e^{-\beta X^s} \mathbf{1}_{\mathcal{B}}] \\ &= \mu \mathbf{E}[e^{-\beta(X^s - X)} e^{-\beta X} \mathbf{1}_{\mathcal{B}}] \\ &\geq \mu \mathbf{E}[e^{-\beta D} e^{-\beta X} \mathbf{1}_{\mathcal{B}}]. \end{aligned}$$

Since $0 \leq D \leq 1$ on \mathcal{B} , we can apply Lemma 3.12 to obtain the bound

$$\begin{aligned} \mathbf{E}[e^{-\beta D} \mathbf{1}_{\mathcal{B}} | X] &\geq \mathbf{E}[(1 - (e^\beta - 1)D) \mathbf{1}_{\mathcal{B}} | X] \\ &= \mathbf{P}[\mathcal{B} | X] - (e^\beta - 1) \mathbf{E}[D \mathbf{1}_{\mathcal{B}} | X] \\ &\geq p - \frac{\tau^2}{\mu} (e^\beta - 1). \end{aligned}$$

We then have

$$M'(-\beta) \geq (p\mu - \tau^2(e^\beta - 1))M(-\beta),$$

and arguing as for the upper tail leads to

$$\log M(-\beta) \leq \tau^2(e^\beta - 1 - \beta) - \mu p \beta.$$

Applying Markov's inequality and setting $\beta = \log(1 + x/\tau^2)$, which is nonnegative for $x \geq 0$, gives

$$\mathbf{P}[X - p\mu \leq -x] \leq M(-\beta)e^{-\beta(x - \mu p)} = \exp\left[-\tau^2 h\left(\frac{x}{\tau^2}\right)\right],$$

and scaling by $c > 0$ as before now yields the first inequality of (18). The second inequality now follows by the second inequality of (14). \square

4. Size biased couplings for random regular graphs. Suppose that A is the adjacency matrix of a random regular graph. In this section, we construct size biased couplings for linear combinations of the entries of A with positive coefficients. Statistics of this form include the number of edges between two given sets of vertices, and the positive part of a truncated quadratic form, as described in Section 6. To construct a size biased coupling for any such statistic, it is enough to give a coupling between A and $A^{(uv)}$, which we define to have the distribution of A conditional on $A_{uv} = 1$. The size biased coupling can then be defined as a mixture of $A^{(uv)}$ for different choices of (u, v) , following the standard recipe for a size biased coupling given in Lemma 3.1.

To make the coupling between A and $A^{(uv)}$, we will use switchings, which are local manipulations of a graph that preserve regularity; see [54], Section 2.4, for an introduction. The most natural thing to do to form the coupling is to apply a switching to A at random out of the ones that yield a graph containing uv . This creates a matrix whose distribution is slightly off from what we want. We then tweak the coupling to get the right distribution, taking care that most of the time, A and $A^{(uv)}$ still differ from each other by a switching.

Switchings, Stein's method and concentration have bumped into each other in a variety of ways in the past. In the configuration model, switchings give easy proofs of concentration by martingale arguments [54], Theorem 2.19. In the uniform model, switchings have been applied to prove tail bounds by ad hoc arguments;

for some examples, see [54], Section 2.4, [44], Theorem 4, and [15], Lemma 16. In [9], switchings are combined with a nonstandard martingale argument to prove concentration of the resolvent of the adjacency matrix of a random regular graph. In [21], switchings were used to define an exchangeable pair in order to apply [16] to prove concentration in random digraphs. Switchings and exchangeable pairs also met in [39], where they were used for Poisson approximation. Janson observed that switchings produce “approximate” couplings of graphs conditioned to have certain edges [38], Remark 5.6. In this section, we essentially make these approximate couplings exact in order to construct size biased couplings.

To make switchings work to achieve our goals, we will view things from a more combinatorial perspective. First, we recast the problem of constructing a coupling as constructing a bipartite graph. We call a bipartite graph *biregular* if all vertices within each vertex class have the same degree, recalling that the degree of a vertex in a weighted graph is the sum of the weights of the edges incident to the vertex.

LEMMA 4.1. *Suppose that \mathfrak{G} is a biregular weighted bipartite graph on vertex sets U and V . Let X be uniformly distributed on U , and let X' be given by walking from X along an edge with probability proportionate to its weight. Then X' is uniformly distributed on V .*

PROOF. Let every vertex in U have degree d and every vertex in V have degree e . Let $w(u, v)$ be the weight of the edge from u to v or 0 if there is none. Since every vertex in U has degree d ,

$$\mathbf{P}[X' = v \mid X = u] = \frac{w(u, v)}{d}.$$

Thus,

$$\mathbf{P}[X' = v] = \sum_{u \in U} \mathbf{P}[X' = v \mid X = u] \mathbf{P}[X = u] = \frac{1}{|U|} \sum_{u \in U} \frac{w(u, v)}{d},$$

and since every vertex in V has degree e , this is $e/d|U| = 1/|V|$. \square

Thus, our goal in this section will be to construct a biregular bipartite graph \mathfrak{G} on the vertex sets \mathcal{G} and \mathcal{G}_{uv} , where \mathcal{G} is the set of adjacency matrices of simple d -regular graphs on n vertices, and \mathcal{G}_{uv} is the subset of \mathcal{G} of matrices with uv entry equal to 1. Roughly speaking, the goal is for the edges of \mathfrak{G} to have as their endpoints graphs that are as similar to each other as possible.

We now define our switchings, which in the combinatorics literature are sometimes called double switchings. See Figure 1 for a pictorial depiction of what we now formally define.

DEFINITION 4.2. Let A be the adjacency matrix of a simple regular graph. Suppose that $A_{v_2 v_3} = A_{v_4 v_5} = A_{v_6 v_1} = 1$ and $A_{v_1 v_2} = A_{v_3 v_4} = A_{v_5 v_6} = 0$, and




FIG. 1. A solid line means an edge between two vertices, and a dotted line means that the two vertices are nonequal and are not connected. The action of replacing the subgraph indicated by the left diagram by the subgraph indicated by the right diagram is a switching at (v_1, \dots, v_6) .

that $v_1 \neq v_2$, $v_3 \neq v_4$, and $v_5 \neq v_6$. Note that we do not assume that all vertices v_1, \dots, v_6 are distinct. Then (v_1, \dots, v_6) is a *valid switching* for A , and we define the application of the switching to A as the adjacency matrix of the graph with edges $v_1 v_2$, $v_3 v_4$ and $v_5 v_6$ added and $v_2 v_3$, $v_4 v_5$ and $v_6 v_1$ deleted.

It is not obvious that a valid switching (v_1, \dots, v_6) preserves regularity if v_1, \dots, v_6 are not all distinct. To see that it does, consider the vertex v_1 . We will show that its degree is unchanged by the switching. Identical arguments apply to the other vertices. By the definition of valid switching, v_1 cannot equal v_2 or v_6 , since it is connected to v_6 and assumed nonequal to v_2 . It cannot equal v_3 , since $A_{v_3 v_2} = 1$ but $A_{v_1 v_2} = 0$, and in the same way it cannot be v_5 . If $v_1 \neq v_4$, then $v_1 v_2$ and $v_1 v_6$ are the only edges incident to v_1 , and its degree is unchanged when $v_1 v_2$ is added and $v_1 v_6$ is deleted. If $v_1 = v_4$, then similar arguments show that v_2, v_3, v_5, v_6 are distinct. Then the switching adds $v_1 v_6$ and $v_1 v_5$ and deletes $v_1 v_3$ and $v_1 v_2$, again leaving the degree of v_1 unchanged.

LEMMA 4.3. For a given adjacency matrix A , let let $s_{uv}(A)$ be the number of valid switchings of the form $(u, v, \cdot, \cdot, \cdot, \cdot)$, and let $t_{uv}(A)$ be the number of valid switchings of the form $(u, \cdot, \cdot, \cdot, \cdot, v)$. For $u \neq v$ with $A_{uv} = 0$,

$$(28) \quad d^3(n - 2d - 2) \leq s_{uv}(A) \leq d^3(n - d - 1)$$

and for $u \neq v$ with $A_{uv} = 1$,

$$(29) \quad d^2(n - d - 1)(n - 2d - 2) \leq t_{uv}(A) \leq d^2(n - d - 1)^2.$$

PROOF. We start by bounding $s_{v_1 v_2}(A)$. Consider the $d^3(n - d - 1)$ tuples $(v_1, v_2, v_3, \dots, v_6)$ given by choosing $v_6 \in \mathcal{N}(v_1)$ and $v_3 \in \mathcal{N}(v_2)$, then $v_5 \in \overline{\mathcal{N}}(v_6)$, and finally $v_4 \in \mathcal{N}(v_5)$ (Figure 1 is very helpful here). This is an upper bound for $s_{v_1 v_2}(A)$. For the lower bound, let K be the number of these tuples that do not allow for a switching, so that

$$s_{v_1 v_2}(A) = d^3(n - d - 1) - K.$$




FIG. 2. A tuple (v_1, \dots, v_6) counted by K coincides with one of the two subgraphs pictured above, with solid lines denoting edges and dotted lines denoting that the endpoints are neither equal nor neighbors. For a given choice of v_1 and v_2 , there are at most d^3 subgraphs of the first kind and d^4 of the second kind.

Now, we bound K from above (see Figure 2). A tuple chosen as above allows for a switching if and only if $v_3 \in \overline{\mathcal{N}}(v_4)$. The number of these tuples where $v_3 = v_4$ is at most d^3 , since in this case $v_3 \in \mathcal{N}(v_2)$, $v_5 \in \mathcal{N}(v_3)$, and $v_6 \in \mathcal{N}(v_1)$, making for d^3 choices total. Similarly, the number of these tuples where $v_3 \in \mathcal{N}(v_4)$ is at most d^4 . Thus, $K \leq d^4 + d^3$, and $s_{ab}(A) \geq d^3(n - 2d - 2)$.

The bound for $t_{v_1 v_6}(A)$ is essentially the same. Consider the tuples (v_1, \dots, v_6) given by choosing $v_2 \in \overline{\mathcal{N}}(v_1)$, then $v_3 \in \mathcal{N}(v_2)$, then $v_5 \in \overline{\mathcal{N}}(v_6)$, and last $v_4 \in \mathcal{N}(v_5)$. There are at most $d^2(n - d - 1)^2$ of these, giving an upper bound for $t_{v_1 v_6}(A)$. For the lower bound, let L be the number of these tuples that are not valid switchings. A tuple fails to be a valid switching if v_3 and v_4 are equal or are neighbors, and we obtain a bound $L \leq (n - d - 1)(d^2 + d^3)$ by counting as in the first case. Thus,

$$\begin{aligned} t_{v_1 v_6}(A) &\geq d^2(n - d - 1)^2 - (n - d - 1)(d^2 + d^3) \\ &= d^2(n - d - 1)(n - 2d - 2). \end{aligned} \quad \square$$

REMARK 4.4. Switchings in which two rather than three edges are added and deleted are known as simple or single switchings. They have been used to analyze regular graphs, though they are typically less effective than double switchings, as mentioned in [54], Section 2.4. The problem is that in the equivalent of (28) for simple switchings, no lower bound is possible. There is no further improvement for us to be found in higher order switchings, however.

LEMMA 4.5. Fix two distinct vertices $u, v \in [n]$. Make a bipartite graph \mathfrak{G}_0 with weighted edges on two vertex classes \mathcal{G} and \mathcal{G}_{uv} by forming edges as follows:

- If $A \in \mathcal{G}$ has $A_{uv} = 0$, then form an edge of weight 1 between A and every element of \mathcal{G}_{uv} that is the result of applying a valid switching of the form $(u, v, \cdot, \cdot, \cdot, \cdot)$.
- If $A \in \mathcal{G}$ has $A_{uv} = 1$, then form an edge of weight $d^3(n - d - 1)$ between A and its identical copy in \mathcal{G}_{uv} .

In \mathfrak{G}_0 , every element of \mathcal{G} has degree between $d^3(n - 2d - 2)$ and $d^3(n - d - 1)$, and every element of \mathcal{G}_{uv} has degree between $d^2(n - d - 1)(n - d - 2)$ and $d^2(n - d - 1)(n - 1)$. Furthermore, \mathfrak{G}_0 can be embedded in a biregular bipartite graph \mathfrak{G} on the same vertex sets, with vertices in \mathcal{G} having degree $d^3(n - d - 1)$ and in \mathcal{G}_{uv} having degree $d^2(n - d - 1)(n - 1)$.

PROOF. We start with the claims about \mathfrak{G}_0 . For any $A \in \mathcal{G}$ with $A_{uv} = 0$, the bound $s_{uv}(A) \leq d^3(n - d - 1)$ from Lemma 4.3 shows that the degree of A in \mathfrak{G}_0 is between $d^3(n - 2d - 2)$ and $d^3(n - d - 1)$. If $A_{uv} = 1$, then A has exactly one incident edge of weight $d^3(n - d - 1)$ in \mathfrak{G}_0 .

If A' is the result of applying a switching $(u, v, w_1, w_2, w_3, w_4)$ to A , then A is the result of applying a switching (u, w_4, w_3, w_2, v) to A' . Thus, $A' \in \mathcal{G}_{uv}$ has $t_{uv}(A')$ incident edges of weight 1, as well as one extra edge of weight $d^3(n - d - 1)$ to its identical copy in \mathcal{G} . The bounds on the degree of A' then follow from the bounds on t_{uv} in Lemma 4.3. This proves all the claims about \mathfrak{G}_0 .

To form \mathfrak{G} , we start with \mathfrak{G}_0 and add edges as follows. Go through the vertices of \mathcal{G} , and for each vertex with degree less than $d^3(n - d - 1)$, arbitrarily make edges from the vertex to vertices in \mathcal{G}_{uv} of weight less than $d^2(n - d - 1)(n - 1)$. Continue this procedure until either all vertices in \mathcal{G} have degree $d^3(n - d - 1)$ or all vertices in \mathcal{G}_{uv} have degree $d^2(n - d - 1)(n - 1)$. We claim that in fact, both are true when the procedure is done. Since the probability of a random regular graph containing edge uv is $d/(n - 1)$, it holds that $|\mathcal{G}_{uv}|/|\mathcal{G}| = d/(n - 1)$. We can count the total edge weight in the graph when the procedure has terminated by summing the degrees of all vertices in \mathcal{G} , or by summing the degrees of all vertices in \mathcal{G}_{uv} . If all degrees in \mathcal{G} are $d^3(n - d - 1)$ and all degrees in \mathcal{G}_{uv} are at most $d^2(n - d - 1)(n - 1)$, then

$$|\mathcal{G}|d^3(n - d - 1) \leq |\mathcal{G}_{uv}|d^2(n - d - 1)(n - 1) = |\mathcal{G}|d^3(n - 1),$$

and so all vertices in \mathcal{G}_{uv} must have degree exactly $d^2(n - d - 1)(n - 1)$. In the same way, if all degrees in \mathcal{G}_{uv} are $d^2(n - d - 1)(n - 1)$, then all degrees in \mathcal{G} must be exactly $d^3(n - d - 1)$. Thus, we have embedded \mathfrak{G}_0 in a biregular bipartite graph \mathfrak{G} as desired. \square

This lemma together with Lemma 4.1 yields a coupling of $(A, A^{(uv)})$ satisfying

$$\begin{aligned} & \mathbf{P}[A \text{ and } A^{(uv)} \text{ are identical or differ by a switching} \mid A^{(uv)}] \\ (30) \quad & \geq 1 - \frac{d + 1}{n - 1}, \end{aligned}$$

and

$$\begin{aligned} & \mathbf{P}[A \text{ and } A^{(uv)} \text{ are identical or differ by a switching} \mid A] \\ (31) \quad & \geq 1 - \frac{d + 1}{n - d - 1} \end{aligned}$$

which can be used to construct size biased couplings for linear sums of A bounded both for the upper and lower tail. This immediately gives tail bounds for any statistic $f(A) = \sum_{u \neq v} a_{uv} A_{uv}$ with $0 \leq a_{uv} \leq c$, since by choosing (U, V) with $\mathbf{P}[(U, V) = (u, v)]$ in proportion to a_{uv} , we obtain a size biased coupling $(f(A), f(A^{(UV)}))$ by Lemma 3.1. For the full details, see Section 5.2, where we carry this out.

5. Concentration for random regular graphs. In this section, we prove Proposition 2.3, establishing the uniform tails property for all the models of random regular graphs we consider. We also prove a concentration result for the edge count $e_A(S, T)$ in the uniform model in Theorem 5.3. Results like this bounding the *edge discrepancy* for random regular graphs have often been of interest; see the expander mixing lemma [37], Lemma 2.5, and [41], Lemma 4.1, for example.

5.1. *Concentration for the permutation model.* Recall that in our permutation models, an adjacency matrix A is given as the symmetrized sum of $d/2$ independent random permutation matrices, for some even d . A more graph theoretic description of the model is as follows. Let $\pi_1, \dots, \pi_{d/2}$ be independent random permutations of $[n]$. Then A is the adjacency matrix of the graph formed by making an edge between i and j for every (i, j, l) such that $\pi_l(i) = j$. Equivalently,

$$(32) \quad A_{ij} = \sum_{l=1}^{d/2} (\mathbf{1}_{\{\pi_l(i)=j\}} + \mathbf{1}_{\{\pi_l(j)=i\}})$$

for $i, j \in [n]$. Note that the graph allows for loops and parallel edges, and that a loop contributes to the adjacency matrix twice. We now show that when the distribution of the permutations is uniform over the symmetric group or is constant on conjugacy classes with no fixed points, the matrix A has the uniform tails property, which we recall from Definition 2.1. Proposition 2.4 then implies that the second eigenvalue of A is $O(\sqrt{d})$ with probability tending to 1. For uniform permutations, this result was previously shown in [23], Theorem 24, and it is included here to highlight that our concentration proofs by size biasing are simpler than previous martingale-based proofs such as [23], Theorem 26.

PROOF OF PROPOSITION 2.3, PARTS (A) AND (B). Fix a symmetric matrix Q and a as in Definition 2.1, and let $\pi_1, \dots, \pi_{d/2}$ be the random permutations defining A . By the symmetry of Q and A , we can view $f_Q(A)$ as

$$(33) \quad f_Q(A) = 2 \sum_{u,v=1}^n \sum_{l=1}^{d/2} Q_{uv} \mathbf{1}_{\{\pi_l(u)=v\}}.$$

First, we consider the case where the common permutation distribution is uniform. We show how to couple π_l with a random permutation $\pi_l^{(uv)}$ distributed as

π_l conditional on $\pi_l(u) = v$. Let τ be the transposition swapping $\pi_l(u)$ and v [or the identity if $\pi_l(u) = v$], and define $\pi_l^{(uv)} = \tau \circ \pi_l$. It is straightforward to check that $\pi_l^{(uv)}$ is distributed as a uniformly random permutation conditioned to map u to v .

Choose (U, V) from $[n] \times [n]$ with $\mathbf{P}[(U, V) = (u, v)]$ proportional to Q_{uv} , and choose L uniformly from $\{1, \dots, d/2\}$, independently of each other and of A . Define A' as we defined A , but with $\pi_L^{(UV)}$ substituting for π_L . This gives us a size biased coupling $(f_Q(A), f_Q(A'))$ by Lemma 3.1. Let $U' = \pi_L^{-1}(V)$ and $V' = \pi_L(U)$. Applying (33), we then have

$$f_Q(A') - f_Q(A) = 2(Q_{UV} + Q_{U'V'} - Q_{UV'} - Q_{U'V}) \leq 2(Q_{UV} + Q_{U'V'}).$$

This shows that $f_Q(A') - f_Q(A) \leq 4a$. With $D = (f_Q(A') - f_Q(A))^+$ and $\mathcal{F} = \{\pi_1, \dots, \pi_{d/2}\}$, we have

$$\begin{aligned} \mathbf{E}[D \mid \mathcal{F}] &\leq 2\mathbf{E}[Q_{UV} + Q_{U'V'} \mid \mathcal{F}] \\ (34) \quad &= \frac{2}{\sum_{u,v=1}^n Q_{uv}} \sum_{u,v=1}^n Q_{uv} \left(Q_{uv} + \frac{2}{d} \sum_{l=1}^{d/2} Q_{\pi_l^{-1}(v)\pi_l(u)} \right) \\ &= \frac{2d}{n\mu} \left(\sum_{u,v=1}^n Q_{uv}^2 + \frac{2}{d} \sum_{l=1}^{d/2} \sum_{u,v=1}^n Q_{uv} Q_{\pi_l^{-1}(v)\pi_l(u)} \right). \end{aligned}$$

Applying the Cauchy–Schwarz inequality,

$$\begin{aligned} \sum_{u,v=1}^n Q_{uv} Q_{\pi_l^{-1}(v)\pi_l(u)} &\leq \left(\sum_{u,v=1}^n Q_{uv}^2 \right)^{1/2} \left(\sum_{u,v=1}^n Q_{\pi_l^{-1}(v)\pi_l(u)}^2 \right)^{1/2} \\ &= \left(\sum_{u,v=1}^n Q_{uv}^2 \right)^{1/2} \left(\sum_{u,v=1}^n Q_{uv}^2 \right)^{1/2} = \sum_{u,v=1}^n Q_{uv}^2. \end{aligned}$$

Substitution into (34) yields

$$\mathbf{E}[D \mid \mathcal{F}] \leq \frac{4d}{n\mu} \sum_{u,v=1}^n Q_{uv}^2 = \frac{4\tilde{\sigma}^2}{\mu}.$$

As A is \mathcal{F} -measurable, the same bound holds for $\mathbf{E}[D \mid A]$. Now apply Theorem 3.4 with $\tau^2 = 4\tilde{\sigma}^2$, $c = 4a$ and $p = 1$ to complete the proof for the uniform permutation case.

Next, let $\pi_l, l = 1, \dots, d/2$ be independent random permutations with distributions constant on conjugacy class and having no fixed points. Lack of fixed points implies that the matrix A has zeros all along its diagonal, and we may therefore assume without loss of generality that $Q_{uu} = 0$. By [18], Section 6.1.2, we have

$$\mathbf{P}[\pi(u) = v] = \frac{1}{n-1} \quad \text{for all } u \neq v,$$




FIG. 3. On the left is π_L and on the right π_L^{UV} , assuming that vertices W, U, V', W', U' and V are distinct.

hence

$$\mu := \mathbf{E}f_Q(A) = \frac{d}{n-1} \sum_{u,v=1}^n Q_{uv}.$$

Sample (U, V) and L as in the uniform case, noting here that since $Q_{uu} = 0$ for all u we have $U \neq V$ a.s. With τ the identity if $\pi_L(U) = V$ and otherwise the transposition $(U, \pi_L^{-1}(V))$, one may check that the permutation $\pi_L^{UV} = \tau \circ \pi \circ \tau$ has the distribution of π_L conditional on $\pi_L(U) = V$, and that therefore, $f_Q(A')$ has the size biased distribution of $f_Q(A)$, where A' is defined as A , but with π_L^{UV} replacing π_L . Let $U' = \pi_L^{-1}(V)$, $V' = \pi_L(U)$ and $W = \pi_L^{-1}(U)$, $W' = \pi_L^{-2}(V)$. See Figure 3 for depictions of these vertices in π_L and π_L^{UV} .

There are two cases we need to consider. In the first case, $V' = U'$, which forces $U = W'$ and puts us in the situation shown in Figure 4. Consulting the figure and applying (33),

$$\begin{aligned} f_Q(A') - f_Q(A) &= 2(Q_{WU'} + Q_{U'U} + Q_{UV} - Q_{WU} - Q_{UU'} - Q_{U'V}) \\ &= 2(Q_{WU'} + Q_{UV} - Q_{WU} - Q_{U'V}) \\ &\leq 2(Q_{WU'} + Q_{UV}) \leq 2(Q_{W'U} + Q_{WU'} + Q_{UV} + Q_{U'V}). \end{aligned}$$

In the other case, we claim that $\{V, V', W, W'\} \cap \{U, U'\} = \emptyset$. Indeed, since π_L has no fixed points, $V' \neq U$, $W \neq U$, $W' \neq U'$ and $V \neq U'$. Since we are not




FIG. 4. On the left, π_L and on the right, π_L^{UV} , in the case where $V' = U'$ and $U = W'$.

in the first case, $V' \neq U'$ and $W' \neq U$. From the way we selected them, $V \neq U$. Since $W = \pi_L^{-1}(U)$ and $U' = \pi_L^{-1}(V)$, we have $W \neq U'$. This confirms the claim. Since τ swaps U and U' (or does nothing if $U = U'$), it leaves V, V', W and W' fixed, giving

$$\begin{aligned}\pi_L^{UV}(W') &= U, & \pi_L^{UV}(W) &= U', \\ \pi_L^{UV}(U) &= V, & \pi_L^{UV}(U') &= V'.\end{aligned}$$

Thus, the only positive terms on the right-hand side of

$$f_Q(A') - f_Q(A) = 2 \sum_{u,v=1}^n \sum_{l=1}^{d/2} Q_{uv} (\mathbf{1}_{\{\pi_l^{UV}(u)=v\}} - \mathbf{1}_{\{\pi_l(u)=v\}})$$

occur when $(u, v) \in \{(W', U), (W, U'), (U, V), (U', V')\}$. We therefore have

$$f_Q(A) - f_Q(A') \leq 2(Q_{W'U} + Q_{WU'} + Q_{UV} + Q_{U'V'}).$$

In both cases, then we have

$$f_Q(A') - f_Q(A) \leq 2(Q_{W'U} + Q_{WU'} + Q_{UV} + Q_{U'V'}) \leq 8a,$$

and following the same argument as for uniform permutations yields

$$\mathbf{E}[D | A] \leq \frac{8}{\sum_{u,v} Q_{uv}} \sum_{u,v} Q_{u,v}^2 = \frac{8\tilde{\sigma}^2}{\mu}.$$

The proof is completed by applying Theorem 3.4 with $\tau^2 = 8\tilde{\sigma}^2$, $c = 8a$ and $p = 1$. \square

5.2. Uniform tails property for the uniform model. Our proof of the uniform tails property for the model where a graph is chosen uniformly from all random d -regular simple graph on n vertices will be similar to the proof for the permutation model in the previous section. The main difference is that here our size biased coupling will take more work to construct and will not be bounded with probability 1. We note that when A is the adjacency matrix of a uniform random regular graph, $A_{uu} = 0$ for $u \in [n]$.

THEOREM 5.1. *Let A be the adjacency matrix of a uniform random simple d -regular graph on n vertices. Let Q be an $n \times n$ symmetric matrix with all entries in $[0, a]$, and let $f_Q(A) = \sum_{u,v} Q_{uv} A_{uv}$. Let $\mu = \mathbf{E} f_Q(A) = \frac{d}{n-1} \sum_{u \neq v} Q_{uv}$ and let $\tilde{\sigma}^2 = \frac{d}{n-1} \sum_{u \neq v} Q_{uv}^2$. Then, with h as given in (13), for all $t \geq 0$,*

$$\begin{aligned}(35) \quad \mathbf{P}\left[f_Q(A) - \frac{\mu}{p} \geq t\right] &\leq \exp\left(-\frac{\tilde{\sigma}^2}{6pa^2} h\left(\frac{pat}{\tilde{\sigma}^2}\right)\right) \\ &\leq \exp\left(\frac{t^2}{12a(t/3 + \tilde{\sigma}^2/ap)}\right)\end{aligned}$$

with $p = 1 - (d + 1)/(n - 1)$, and

$$\begin{aligned}
 \mathbf{P}[f_Q(A) - p'\mu \leq -t] &\leq \exp\left(-\frac{\tilde{\sigma}^2}{6a^2}h\left(\frac{at}{\tilde{\sigma}^2}\right)\right) \\
 &\leq \exp\left(-\frac{t^2}{12a(t/3 + \tilde{\sigma}^2/a)}\right)
 \end{aligned}
 \tag{36}$$

with $p' = 1 - (d + 1)/(n - d - 1)$.

PROOF. We now construct a size biased coupling using the tools we developed in Section 4. Let $A^{(v_1 v_2)}$ be the matrix obtained by walking randomly in the bipartite graph \mathfrak{G} , constructed in Lemma 4.5, from A along an edge chosen with probability proportional to its weight. By Lemma 4.1, the matrix $A^{(v_1 v_2)}$ is distributed as A conditioned on $A_{v_1 v_2} = 1$. Independently of A , choose $(V_1, V_2) = (v_1, v_2)$ with probability proportional to $Q_{v_1 v_2}$ for all $v_1 \neq v_2$, and set $A' = A^{(V_1 V_2)}$. By Lemma 3.1, the pair $(f_Q(A), f_Q(A'))$ is a size biased coupling. Define \mathcal{B} as the event that the edge traversed in \mathfrak{G} from A to $A^{(V_1 V_2)}$ belongs to \mathfrak{G}_0 . By (30) and (31), $\mathbf{P}[\mathcal{B} \mid A'] \geq p$, and $P[\mathcal{B} \mid A] \geq p'$.

Let $\mathcal{S}(A, v_1, v_2)$ consist of all tuples (v_3, \dots, v_6) such that (v_1, \dots, v_6) is a valid switching. Note that if $A_{v_1 v_2} = 1$, then $\mathcal{S}(A, v_1, v_2)$ is the empty set. For $(v_3, \dots, v_6) \in \mathcal{S}(A, v_1, v_2)$, let $A(v_1, \dots, v_6)$ denote A after application of the switching (v_1, \dots, v_6) . Looking back at Lemma 4.5, we can describe the coupling of A and A' as follows. Conditional on A, V_1 and V_2 and assuming $A_{V_1 V_2} = 0$, the matrix A' takes the value $A(V_1, V_2, v_3, \dots, v_6)$ with probability $1/d^3(n - d - 1)$ for each $(v_3, \dots, v_6) \in \mathcal{S}(A, V_1, V_2)$, and these events make up the set \mathcal{B} . The matrix A' can take other values as well, if $|\mathcal{S}(A, V_1, V_2)|$ is strictly smaller than $d^3(n - d - 1)$, in which case \mathcal{B} does not hold.

In view of Figure 1, we have

$$\begin{aligned}
 f_Q(A(v_1, \dots, v_6)) - f_Q(A) &= 2(Q_{v_1 v_2} + Q_{v_3 v_4} + Q_{v_5 v_6} - Q_{v_2 v_3} - Q_{v_4 v_5} - Q_{v_6 v_1}) \\
 &\leq 2(Q_{v_1 v_2} + Q_{v_3 v_4} + Q_{v_5 v_6}),
 \end{aligned}$$

the factor of 2 arising because addition or deletion of edge uv adds or removes both terms Q_{uv} and Q_{vu} . This shows that $f_Q(A') - f_Q(A) \leq 6a$ on the event \mathcal{B} .

Let $\overline{\mathcal{S}}(A, v_1, v_2)$ denote the set of tuples (v_3, \dots, v_6) with $v_3 \in \mathcal{N}(v_2)$, $v_4 \in \overline{\mathcal{N}}(v_3)$, $v_5 \in \mathcal{N}(v_4)$ and $v_6 \in \mathcal{N}(v_1)$. Recalling that $\overline{\mathcal{N}}(v)$ is the set of $n - d - 1$ vertices not equal to v or the neighbors of v , we see that $\overline{\mathcal{S}}(A, v_1, v_2)$ has size $d^3(n - d - 1)$, and that it contains $\mathcal{S}(A, v_1, v_2)$. Letting $D = (f_Q(A') - f_Q(A))^+$,

we have

$$\begin{aligned}
& \mathbf{E}[D\mathbf{1}_B \mid A, V_1, V_2] \\
&= \frac{1}{d^3(n-d-1)} \\
&\quad \times \sum_{(v_3, \dots, v_6) \in \mathcal{S}(A, V_1, V_2)} (f_Q(A(V_1, V_2, v_3, \dots, v_6)) - f_Q(A))^+ \\
&\leq \frac{2}{d^3(n-d-1)} \sum_{(v_3, \dots, v_6) \in \bar{\mathcal{S}}(A, V_1, V_2)} (Q_{V_1 V_2} + Q_{v_3 v_4} + Q_{v_5 v_6}).
\end{aligned}$$

Recalling the distribution of (V_1, V_2) and observing that $\sum_{u \neq v} Q_{uv} = (n-1)\mu/d$,

$$\begin{aligned}
& \mathbf{E}[D\mathbf{1}_B \mid A] \\
&\leq \sum_{v_1 \neq v_2} \frac{Q_{v_1 v_2}}{\sum_{u \neq v} Q_{uv}} \\
(37) \quad &\quad \times \left(\frac{2}{d^3(n-d-1)} \sum_{(v_3, \dots, v_6) \in \bar{\mathcal{S}}(A, v_1, v_2)} (Q_{v_1 v_2} + Q_{v_3 v_4} + Q_{v_5 v_6}) \right) \\
&= \frac{2}{(n-1)(n-d-1)d^2\mu} \\
&\quad \times \sum_{\substack{v_1 \neq v_2 \\ (v_3, \dots, v_6) \in \bar{\mathcal{S}}(A, v_1, v_2)}} (Q_{v_1 v_2}^2 + Q_{v_1 v_2} Q_{v_3 v_4} + Q_{v_1 v_2} Q_{v_5 v_6}).
\end{aligned}$$

We now consider each term of this sum. For the first one,

$$\begin{aligned}
(38) \quad & \sum_{\substack{v_1 \neq v_2 \\ (v_3, \dots, v_6) \in \bar{\mathcal{S}}(A, v_1, v_2)}} Q_{v_1 v_2}^2 = d^3(n-d-1) \sum_{v_1 \neq v_2} Q_{v_1 v_2}^2 \\
&= (n-1)(n-d-1)d^2\tilde{\sigma}^2.
\end{aligned}$$

For the next term, we apply the Cauchy–Schwarz inequality in an argument similar to what we used in the proof of parts (a) and (b) of Proposition 2.3:

$$\begin{aligned}
& \sum_{\substack{v_1 \neq v_2 \\ (v_3, \dots, v_6) \in \bar{\mathcal{S}}(A, v_1, v_2)}} Q_{v_1 v_2} Q_{v_3 v_4} \\
&\leq \left(\sum_{\substack{v_1 \neq v_2 \\ (v_3, \dots, v_6) \in \bar{\mathcal{S}}(A, v_1, v_2)}} Q_{v_1 v_2}^2 \right)^{1/2} \left(\sum_{\substack{v_1 \neq v_2 \\ (v_3, \dots, v_6) \in \bar{\mathcal{S}}(A, v_1, v_2)}} Q_{v_3 v_4}^2 \right)^{1/2}.
\end{aligned}$$

The first factor on the right-hand side was evaluated in (38). For the second one, observe that for a given $v_3 \neq v_4$, there are $d^3(n-d-1)$ tuples (v_1, v_2, v_5, v_6) such that $(v_3, \dots, v_6) \in \overline{\mathcal{S}}(A, v_1, v_2)$, giving

$$\begin{aligned} \sum_{\substack{v_1 \neq v_2 \\ (v_3, \dots, v_6) \in \overline{\mathcal{S}}(A, v_1, v_2)}} Q_{v_3 v_4}^2 &= d^3(n-d-1) \sum_{v_3 \neq v_4} Q_{v_3 v_4}^2 \\ &= (n-1)(n-d-1)d^2\tilde{\sigma}^2. \end{aligned}$$

Thus,

$$\sum_{\substack{v_1 \neq v_2 \\ (v_3, \dots, v_6) \in \overline{\mathcal{S}}(A, v_1, v_2)}} Q_{v_1 v_2} Q_{v_3 v_4} \leq (n-1)(n-d-1)d^2\tilde{\sigma}^2.$$

The same bound holds for the final term in (37). Thus, we have

$$\mathbf{E}[D\mathbf{1}_B \mid A] \leq \frac{6\tilde{\sigma}^2}{\mu}.$$

Theorem 3.4 now proves (35) and (36). \square

Now we deduce part (c) of Proposition 2.3 from Theorem 5.1.

PROOF OF PROPOSITION 2.3, PART (C). We start with an elementary estimate: for any $p \in [0, 1]$ and $x \geq 0$,

$$(39) \quad p^{-1}h(px) \geq ph(x).$$

Indeed, for fixed $p \in [0, 1]$, note that by concavity of $x \mapsto (1+x)^p$,

$$1 + px \geq (1+x)^p$$

for all $x \geq 0$. Taking logarithms and integrating the inequality gives

$$\begin{aligned} h(px) &= \int_0^x \frac{d}{dt} h(pt) dt \\ &= \int_0^x p \log(1+pt) dt \\ &\geq p^2 \int_0^x \log(1+t) dt \\ &= p^2 h(x) \end{aligned}$$

as desired.

Recall

$$(40) \quad c_0 = \frac{1}{6} \left(1 - \frac{d+1}{n-1} \right) = \frac{p}{6}, \quad \gamma_0 = \frac{d+1}{n-d-2} = \frac{1}{p} - 1$$

with p as in Theorem 5.1. Let Q be an $n \times n$ symmetric matrix with entries in $[0, a]$, as in Definition 2.1. By Theorem 5.1, for all $t \geq 0$,

$$\begin{aligned} \mathbf{P}[f_Q(A) - \mu \geq \gamma_0 \mu + t] &= \mathbf{P}\left[f_Q(A) - \frac{\mu}{p} \geq t\right] \\ &\leq \exp\left(-\frac{\tilde{\sigma}^2}{6pa^2} h\left(\frac{pat}{\tilde{\sigma}^2}\right)\right) \\ &\leq \exp\left(-c_0 \frac{\tilde{\sigma}^2}{a^2} h\left(\frac{at}{\tilde{\sigma}^2}\right)\right), \end{aligned}$$

where in the last step we applied (39). Similarly,

$$\begin{aligned} \mathbf{P}[f_Q(A) - \mu \leq -(\gamma_0 \mu + t)] &= \mathbf{P}[f_Q(A) - p'\mu \leq -((\gamma_0 - 1 + p')\mu + t)] \\ &\leq \mathbf{P}[f_Q(A) - p'\mu \leq -t] \\ &\leq \exp\left(-\frac{\tilde{\sigma}^2}{6a^2} h\left(\frac{at}{\tilde{\sigma}^2}\right)\right) \\ &\leq \exp\left(-c_0 \frac{\tilde{\sigma}^2}{a^2} h\left(\frac{at}{\tilde{\sigma}^2}\right)\right), \end{aligned}$$

where in the second line we used that $1 - p' = (d + 1)/(n - d - 1)$, which we see from (40) is (slightly) smaller than γ_0 . \square

REMARK 5.2. Proposition 2.3 on the statistic $f_Q(A)$ can be seen as extensions of results on $f_Q(P)$ where P is a random permutation matrix. This is Hoeffding's combinatorial statistic, as studied in [35]. Concentration for this statistic was achieved using exchangeable pairs by [16], who showed, with $\mu = \mathbf{E}f_Q(P)$, that

$$\mathbf{P}(|f_Q(P) - \mu| \geq t) \leq 2 \exp\left(-\frac{t^2}{4\mu + 2t}\right) \quad \text{for all } t \geq 0$$

when $Q_{uv} \in [0, 1]$. Under these same conditions, using zero biasing [32] obtained the Bennett-type inequality:

$$\mathbf{P}(|f_Q(P) - \mu| \geq t) \leq 2 \exp\left(-\frac{t^2}{2\sigma^2 + 16t}\right) \quad \text{for all } t \geq 0,$$

where $\sigma^2 = \text{Var}(f_Q(P))$, as well as Bennett-type bounds whose tails decay asymptotically at the faster ‘‘Poisson’’ rate $\exp(-\Omega(t \log t))$, as do the bounds given in Proposition 2.3.

In some applications, ours among them, concentration bounds that depend on the variance are preferable to those depending on the mean. In our case, however, the variance proxy $\tilde{\sigma}^2$ in Definition 2.1 suffices. For the permutation model, it seems likely that the zero bias method can be applied to yield a concentration

bound for $f_Q(A)$ depending on the true variance. For the uniform model, it appears difficult to create a zero bias coupling for $f_Q(A)$, but it appears possible to construct an *approximate* zero bias coupling at the expense of some additional complexity.

Since the edge counts $e_A(S, T)$ can be expressed as $f_Q(A) = \sum_{u,v} A_{uv} Q_{uv}$ with

$$Q_{uv} = \frac{1}{2}(\mathbf{1}_{\{u \in S, v \in T\}} + \mathbf{1}_{\{v \in S, u \in T\}}),$$

concentration for $e_A(S, T)$ follows as a corollary of Theorem 5.1. With a bit of extra effort, we can improve the constant in the tail bound. Since edge discrepancy concentration is of independent interest, we make the effort and give the better result.

THEOREM 5.3. *Let A be the adjacency matrix of a uniformly random d -regular graph on n vertices, and let $S, T \subseteq [n]$. Define*

$$\mu = \mathbf{E}e_A(S, T) = \frac{(|S||T| - |S \cap T|)d}{n - 1}.$$

(a) *For any $t \geq 1$,*

$$(41) \quad \mathbf{P}\left[e_A(S, T) \geq \frac{t\mu}{p}\right] \leq \exp\left(-\frac{\mu}{2p}h(t - 1)\right) \leq \exp\left(-\frac{3\mu(t - 1)^2}{4p(2 + t)}\right)$$

where $p = 1 - (d + 1)/(n - 1)$.

(b) *For any $0 < t \leq 1$,*

$$(42) \quad \mathbf{P}[e_A(S, T) \leq t p \mu] \leq \exp\left(-\frac{p\mu}{2}h(t - 1)\right) \leq \exp\left(-\frac{p\mu(1 - t)^2}{4}\right),$$

where $p = 1 - (d + 1)/(n - d - 1)$.

PROOF. Recall that

$$e_A(S, T) = \sum_{\substack{u \in S \\ v \in T}} A_{uv}.$$

Take \mathfrak{G} from Lemma 4.5, and form a coupling $(A, A^{(uv)})$ by defining $A^{(uv)}$ to be the result of walking from A along an edge in \mathfrak{G} chosen with probability proportionate to its weight. By Lemma 4.1, the matrix $A^{(uv)}$ is distributed as A conditional on $A_{uv} = 1$. Choosing U uniformly from S and V uniformly from T , independent of each other and of A , and setting $A' = A^{(UV)}$, by Lemma 3.1 we obtain a size biased coupling $(e_A(S, T), e_{A'}(S, T))$.

We claim that if A and A' differ by a switching, then $e_{A'}(S, T) \leq e_A(S, T) + 2$. Suppose the switching adds v_1v_2 , v_3v_4 , and v_5v_6 and deletes v_2v_3 , v_4v_5 and v_6v_1 . Considering indices modulo 6 and referring to Figure 1, let

$$I_i = \mathbf{1}_{\{v_i \in S, v_{i+1} \in T\}}, \quad J_i = \mathbf{1}_{\{v_i \in T, v_{i+1} \in S\}}.$$

Then

$$\begin{aligned} e_{A'}(S, T) - e_A(S, T) \\ = (I_1 + I_3 + I_5 - J_2 - J_4 - J_6) + (J_1 + J_3 + J_5 - I_2 - I_4 - J_6). \end{aligned}$$

If $I_i = I_{i+2} = 1$, then $J_{i+1} = 1$. From this observation, one can work out that the first term is at most 1, and by the same argument the second term is also at most 1.

By (30) and (31), the coupling is then $(2, 1 - (d + 1)/(n - 1))$ -bounded for the upper tail, and $(2, 1 - (d + 1)/(n - d - 1))$ -bounded for the lower tail. Theorem 3.3 then proves (41) and (42). \square

REMARK 5.4. Similar results were established for random regular *digraphs* by the first author in [21] using Chatterjee's exchangeable pairs approach [16], another variant of Stein's method. This approach would likely give effective bounds when d is on the same order as n , in which case the bounds given by Theorem 5.3 start to break down. For instance, if $d = n/2$, then $p \approx 1/2$, and the upper tail bound (41) becomes effective only starting at $2\mu \approx |S||T|$, a trivial upper bound. Similarly, as d rises to $n/2$, the factor p' approaches zero, and the lower bound (42) breaks down as well.

6. The Kahn–Szemerédi argument. In [28], Kahn and Szemerédi introduced a general approach for bounding the second eigenvalue of a random regular graph, which they used to show that the second eigenvalue of a random graph from the permutation model is $O(\sqrt{d})$ with high probability as $n \rightarrow \infty$ with d fixed. The disadvantage of their approach as compared to the *trace method* used by Friedman [27] and Broder–Shamir [14] is that it is incapable of capturing the correct constant in front of \sqrt{d} . However, it is more flexible in some ways than the trace method: it has been adapted to establish bounds on the spectral gap for several other random graph models (see, for instance, [15, 19, 20, 25, 29, 40, 42]), and it can be applied when d grows with n , as observed in [15].

We now describe how the argument will go for us. For now, we let A denote the adjacency matrix of a random d -regular graph without specifying the distribution further. Recall our notation $\lambda(A) = \max(\lambda_2(A), -\lambda_n(A))$ for the largest (in magnitude) nontrivial eigenvalue. Alternatively, $\lambda(A) = s_2(A)$, the second-largest singular value [recall that $\lambda_1(A) = s_1(A) = d$].

The Kahn–Szemerédi approach stems from the Courant–Fischer variational formula:

$$(43) \quad \lambda(A) = \sup_{x \in \mathbf{S}_0^{n-1}} |x^\top Ax|,$$

where \mathbf{S}^{n-1} is the unit sphere in \mathbb{R}^n and

$$\mathbf{S}_0^{n-1} := \left\{ x \in \mathbf{S}^{n-1} : \sum_{i=1}^n x_i = 0 \right\} = \mathbf{S}^{n-1} \cap \langle \mathbf{1} \rangle^\perp,$$

which follows from the fact that $\mathbf{1} = (1, \dots, 1)$ is the eigenvector corresponding to $\lambda_1(A) = d$. Broadly speaking, the approach is to bound the supremum by first demonstrating concentration results for random variables $x^\top Ax$ for a fixed vector x . (Kahn and Szemerédi actually considered $x^\top My$ for various choices of x, y with M a nonsymmetrized version of A , but it makes little difference to the argument.) A short continuity argument shows that to control the supremum in (43), it suffices to control $x^\top Ax$ for all x in a suitable net of \mathbf{S}_0^{n-1} of cardinality C^n for some constant $C > 0$ (see Section 6.3). Toward applying a union bound over such a net, one might seek bounds on $|x^\top Ax|$ of order $O(\sqrt{d})$ holding with probability $1 - O(e^{-C'n})$ for some $C' > 0$ sufficiently large depending on C . It turns out that this is impossible, at least when d is fixed as n grows, since a $O(\sqrt{d})$ eigenvalue bound is only expected to hold with probability approaching one polynomially in this case [indeed, in the permutation model it is not hard to see that the graph is *disconnected* with probability $\Omega(n^{-c})$ for some $c > 0$ depending on d]. However, Kahn and Szemerédi gave a modification of this argument that works.

We motivate their argument by first considering a simpler problem: to show that $|x^\top Bx| = O(\sqrt{n})$ with high probability when B is the adjacency matrix of an Erdős–Rényi graph with expected density $p = d/n$ and $x \in \mathbf{S}_0^{n-1}$. It easily follows from Hoeffding’s inequality that for a fixed unit vector x and any $t \geq 0$,

$$(44) \quad \mathbf{P}[|x^\top Bx - \mathbf{E}x^\top Bx| \geq t] \leq 2 \exp\left(-\frac{ct^2}{\sum_{u,v=1}^n |x_u x_v|^2}\right) = 2 \exp(-ct^2)$$

for some absolute constant $c > 0$. Moreover, if $x \in \mathbf{S}_0^{n-1}$ we have $\mathbf{E}x^\top Bx = 0$, and we conclude that $x^\top Bx = O(\sqrt{n})$ except with probability $O(e^{-C'n})$, where we can take the constant $C' > 0$ as large as we please. Combined with a union bound over the net described above, and taking C' sufficiently large depending on C we deduce that

$$(45) \quad \sup_{x \in \mathbf{S}_0^{n-1}} |x^\top Bx| = O(\sqrt{n})$$

except with exponentially small probability.

There are two difficulties one encounters in trying to extend this argument to random d -regular graphs. This first is that Hoeffding’s inequality is unavailable as the entries of A are not independent. In Kahn and Szemerédi’s proof for the permutation model, a martingale argument was used instead. In the present work, we use size biased couplings for the uniform model, through the uniform tails property (Definition 2.1).

The second barrier is that the bound (45) is not of the desired order $O(\sqrt{d})$. This stems from the appearance of the L^∞ bound $|x_u x_v|$ on the summands of $x^\top Bx$ that appears in the denominator of the exponential in Hoeffding's inequality (44). We would like to substitute this with an L^2 bound, which has size on the order of the density p of B (and can be shown to have order d/n for adjacency matrices A with hypotheses as in Proposition 2.4). Such a substitute is provided by concentration inequalities of Bennett-type, which for A would give bounds of the form

$$(46) \quad \mathbf{P}[|x^\top Ax - \mathbf{E}x^\top Ax| \geq t] \leq 2 \exp\left(-\frac{ct^2}{(\sum_{u,v=1}^n |x_u x_v|^2 \mathbf{E}A_{uv}^2) + t \max_{u,v} |x_u x_v|}\right).$$

The first term in the denominator of the exponent is order $O(d/n)$. Substituting $C\sqrt{d}$ for t , we need the term $\max_{u,v} |x_u x_v|$ to be of size $O(\sqrt{d}/n)$ in order that the bound decay exponentially in n .

This motivates a key step in Kahn and Szemerédi's argument, which is to split the sum $\sum_{u,v} x_u x_v A_{uv}$ into two pieces. For fixed $x \in \mathbf{S}_0^{n-1}$, we define the *light* and *heavy* couples of vertices, respectively, by

$$(47) \quad \mathcal{L}(x) = \{(u, v) \in [n]^2 : |x_u x_v| \leq \sqrt{d}/n\} \quad \text{and} \quad \mathcal{H}(x) = [n]^2 \setminus \mathcal{L}(x),$$

using the terminology from [25]. We then use the decomposition

$$(48) \quad x^\top Ax = \sum_{(u,v) \in [n]^2} x_u x_v A_{uv} = \sum_{(u,v) \in \mathcal{L}(x)} x_u x_v A_{uv} + \sum_{(u,v) \in \mathcal{H}(x)} x_u x_v A_{uv}.$$

We can express this in the notation of (4) as

$$x^\top Ax = f_{xx^\top}(A) = f_{L(x)}(A) + f_{H(x)}(A),$$

where $L(x)$ is the matrix with entries

$$[L(x)]_{uv} = \begin{cases} x_u x_v & (u, v) \in \mathcal{L}(x), \\ 0 & \text{otherwise,} \end{cases}$$

and $H(x) = xx^\top - L(x)$.

The goal is now to show that $f_{L(x)}(A)$ and $f_{H(x)}(A)$ are each of size $O(\sqrt{d})$ with high probability. The light couples contribution $f_{L(x)}(A)$ can be handled by a bound of the form (46) (which we have thanks to the uniform tails property) together with a union bound over a discretization of the sphere, as outlined above for the Erdős-Rényi case.

The contribution of heavy couples $f_{H(x)}(A)$ does not enjoy sufficient concentration to beat the cardinality of a net of the sphere. Here, the key idea is to prove that a *discrepancy property* holds with high probability for the associated random regular graph. This essentially means that the edge counts

$$(49) \quad e_A(S, T) = \sum_{u \in S, v \in T} A_{uv} = \mathbf{1}_S^\top A \mathbf{1}_T$$

are not much larger than their expectation, uniformly over choices of $S, T \subset [n]$ (here $\mathbf{1}_S \in \{0, 1\}^n$ denotes the vector with j th component equal to 1 if $j \in S$ and 0 otherwise). This is accomplished using tail estimates for the random variables $e_A(S, T)$. One then shows that conditional on the event that the discrepancy property holds, the contribution $f_{H(x)}(A)$ of the heavy couples to the sum (48) is $O(\sqrt{d})$ with probability 1.

We point out that concentration estimates play a crucial role in both parts of the argument above, though in different guises: in the light couples argument it is for the random variables $f_{L(x)}(A)$ with $x \in \mathbf{S}_0^{n-1}$, while in the heavy couples argument it is for the random variables $e_A(S, T)$ with $S, T \subset [n]$. In our implementation of the Kahn–Szemerédi argument below, the necessary concentration bounds both follow from the uniform tails property (Definition 2.1).

The remainder of this section establishes Proposition 2.4 and is organized as follows. We bound the contribution of the light couples in Section 6.1 and the heavy couples in Section 6.2. Proposition 2.4 follows easily from these two sections; we give the final proof in Section 6.3. We do all of this without reference to a specific graph model. Instead, we assume the uniform tails property. Proposition 2.4 is then applicable to any graph model where this is shown to hold.

6.1. *Light couples.* In this section, we establish Lemma 6.2, which says that the uniform tails property implies that $f_{L(x)}(A)$ is $O(\sqrt{d})$ with overwhelming probability for any particular vector $x \in \mathbf{S}_0^{n-1}$. The uniform tails property was tailored for exactly this purpose, so it is just matter of working out the details. The work of extending this bound from a single vector to a supremum over the entire sphere \mathbf{S}_0^{n-1} occurs in Section 6.3.

LEMMA 6.1 (Expected contribution of light couples). *Let A be the adjacency matrix of a random d -regular multigraph on n vertices satisfying the conditions of Proposition 2.4. Then for any fixed $x \in \mathbf{S}_0^{n-1}$, $|\mathbf{E}f_{L(x)}(A)| \leq (a_1 + a_2)\sqrt{d}$, with a_1, a_2 as in Proposition 2.4.*

PROOF. Fix $x \in \mathbf{S}_0^{n-1}$. From the decomposition (48),

$$\begin{aligned} |\mathbf{E}f_{L(x)}(A)| &\leq |\mathbf{E}x^\top Ax| + |\mathbf{E}f_{H(x)}(A)| \\ &\leq \left| x^\top \left(\mathbf{E}A - \frac{d}{n} \mathbf{1}\mathbf{1}^\top \right) x \right| + a_1 \frac{d}{n} \sum_{(u,v) \in \mathcal{H}(x)} |x_u x_v| \\ &\leq \left\| \mathbf{E}A - \frac{d}{n} \mathbf{1}\mathbf{1}^\top \right\|_{\text{HS}} + a_1 \frac{d}{n} \sum_{u,v=1}^n \frac{|x_u x_v|^2}{\sqrt{d}/n} \\ &\leq a_2 \sqrt{d} + a_1 \sqrt{d}, \end{aligned}$$

where in the second line we have used $x \perp \mathbf{1}$, and in the third line applied the Cauchy–Schwarz inequality to the first term. \square

LEMMA 6.2. *Let A be the adjacency matrix of a random d -regular multigraph on n vertices satisfying the conditions of Proposition 2.4. Then for any $x \in \mathbf{S}_0^{n-1}$ and $\beta \geq 4a_1a_3$,*

$$(50) \quad \mathbf{P}[|f_{L(x)}(A)| \geq (\beta + a_1 + a_2)\sqrt{d}] \leq 4 \exp\left(-\frac{c_0\beta^2n}{32(a_1 + \frac{\beta}{12})}\right).$$

PROOF. Applying Lemma 6.1,

$$(51) \quad \begin{aligned} \mathbf{P}[|f_{L(x)}(A)| \geq (\beta + a_1 + a_2)\sqrt{d}] \\ \leq \mathbf{P}[|f_{L(x)}(A) - \mathbf{E}f_{L(x)}(A)| \geq \beta\sqrt{d}]. \end{aligned}$$

Splitting $L(x) = L^+(x) - L^-(x)$ into positive and negative parts, by a union bound the right-hand side of (51) is bounded by

$$(52) \quad \begin{aligned} \mathbf{P}[|f_{L^+(x)}(A) - \mathbf{E}f_{L^+(x)}(A)| \geq (\beta/2)\sqrt{d}] \\ + \mathbf{P}[|f_{L^-(x)}(A) - \mathbf{E}f_{L^-(x)}(A)| \geq (\beta/2)\sqrt{d}]. \end{aligned}$$

Considering the first term, abbreviate $\mu := \mathbf{E}f_{L^+(x)}(A)$. Note that by Cauchy–Schwarz and the assumption that $\mathbf{E}A_{uv} \leq a_1 \frac{d}{n}$,

$$(53) \quad \mu \leq a_1 \frac{d}{n} \sum_{u,v=1}^n |x_u x_v| \leq a_1 d \left(\sum_{u,v=1}^n |x_u x_v|^2 \right)^{1/2} = a_1 d.$$

From (47), each entry of the matrix $L^+(x)$ lies in $[0, \sqrt{d}/n]$. Moreover, again using our first assumption in Proposition 2.4,

$$\tilde{\sigma}^2 := f_{L^+(x) \circ L^+(x)}(\mathbf{E}A) \leq \sum_{u,v=1}^n |x_u x_v|^2 \mathbf{E}A_{uv} \leq a_1 \frac{d}{n},$$

where we use the notation of Definition 2.1 with $Q = L^+(x)$. Recall that we are assuming that A has UTP(c_0, γ_0) for $\gamma_0 = a_3/\sqrt{d}$. Applying (7),

$$\begin{aligned} \mathbf{P}[|f_{L^+(x)}(A) - \mu| \geq (\beta/2)\sqrt{d}] \\ \leq \mathbf{P}[|f_{L^+(x)}(A) - \mu| \geq \gamma_0\mu - \gamma_0a_1d + (\beta/2)\sqrt{d}] \\ \leq 2 \exp\left(-\frac{c_0(\frac{\beta}{2}\sqrt{d} - \gamma_0a_1d)^2}{2a_1\frac{d}{n} + \frac{2}{3}\frac{\sqrt{d}}{n}(\frac{\beta}{2}\sqrt{d} - \gamma_0a_1d)}\right). \end{aligned}$$

Recall that $\gamma_0a_1d = a_1a_3\sqrt{d}$. Hence, if $\beta \geq 4a_1a_3$, then since $t \mapsto t^2/(a + bt)$ is

nondecreasing on $[0, \infty)$ for $a, b > 0$, we conclude the bound

$$\begin{aligned} \mathbf{P}[|f_{L+(x)}(A) - \mu| \geq (\beta/2)\sqrt{d}] & \\ & \leq 2 \exp\left(-\frac{c_0\beta^2 d}{32(a_1\frac{d}{n} + \frac{\beta}{12}\frac{d}{n})}\right) \\ & = 2 \exp\left(-\frac{c_0\beta^2 n}{32(a_1 + \frac{\beta}{12})}\right). \end{aligned}$$

The same bound holds for the second term in (52), which combined with (51) proves the lemma. \square

6.2. *Heavy couples.* In this section, we define a *discrepancy property* for a matrix. For an adjacency matrix, the discrepancy property essentially says that the number of edges between any two sets of vertices is not too much larger than its expectation. Lemma 6.4 shows that the uniform upper tail property (see Definition 2.1) implies that the discrepancy property holds except with polynomially small probability. Lemma 6.6 then shows that if the discrepancy property holds for A , then deterministically the heavy couples give a small contribution to $x^T Ax$ for any vector x .

DEFINITION 6.3 (Discrepancy property). Let M be an $n \times n$ matrix with non-negative entries. For $S, T \subset [n]$, recall that

$$e_M(S, T) := \sum_{u \in S} \sum_{v \in T} M_{uv}.$$

We say that M has the *discrepancy property with parameters* $\delta \in (0, 1)$, $\kappa_1 > 1$, $\kappa_2 \geq 0$, or $\text{DP}(\delta, \kappa_1, \kappa_2)$, if for all nonempty $S, T \subset [n]$ at least one of the following hold:

1. $\frac{e_M(S, T)}{\delta|S||T|} \leq \kappa_1$;
2. $e_M(S, T) \log \frac{e_M(S, T)}{\delta|S||T|} \leq \kappa_2(|S| \vee |T|) \log \frac{en}{|S| \vee |T|}$.

The following lemma shows that if a symmetric matrix A has the uniform upper tail property with parameters $c_0 > 0$, $\gamma_0 \geq 0$, the discrepancy property holds with high probability for some κ_1, κ_2 depending on c_0, γ_0 .

LEMMA 6.4 (UUTP \Rightarrow DP holds with high probability). Let M be an $n \times n$ symmetric random matrix with nonnegative entries. Assume that for some $\delta \in (0, 1)$, $\mathbf{E}M_{uv} \leq \delta$ for all $u, v \in [n]$ and that M has $\text{UUTP}(c_0, \gamma_0)$ for some $c_0 > 0$ and $\gamma_0 \geq 0$. Then for any $K > 0$, $\text{DP}(\delta, \kappa_1, \kappa_2)$ holds for M with probability at least $1 - n^{-K}$ with

$$(54) \quad \kappa_1(\gamma_0) = e^2(1 + \gamma_0)^2, \quad \kappa_2(c_0, \gamma_0, K) = \frac{2}{c_0}(1 + \gamma_0)(K + 4).$$

REMARK 6.5 (Smaller deviations for edge counts). The above lemma controls *large deviations* of edge counts $e_M(S, T)$ for random matrices with the uniform tails property. One can also use the uniform tails property (or Theorem 5.3 in particular for the uniform random regular graph) to obtain tighter control of $e_M(S, T)$ around its expectation, uniformly over all *sufficiently large* sets S, T . Control of this type was used in [22] to show that adjacency matrices of random d -regular digraphs with $\min(d, n - d) \geq C \log^2 n$ are invertible with high probability.

PROOF. For $S, T \subset [n]$, we write

$$\mu(S, T) := \mathbf{E}e_M(S, T) \leq \delta|S||T|.$$

Fix $K > 0$. Put $\gamma_1 = e^2(1 + \gamma_0)^2 - 1$, and for $S, T \subset [n]$, let $\gamma = \gamma(S, T, n) = \max(\gamma^*, \gamma_1)$, where γ^* is the unique solution for x in $[\gamma_0, \infty)$ to

$$(55) \quad c_0h(x - \gamma_0)\mu(S, T) = (K + 4)(|S| \vee |T|) \log\left(\frac{en}{|S| \vee |T|}\right).$$

We can recast $e_M(S, T)$ in the notation of (4) as $f_Q(M)$ with $Q = \frac{1}{2}(\mathbf{1}_S \mathbf{1}_T^T + \mathbf{1}_T \mathbf{1}_S^T)$, where $\mathbf{1}_S \in \{0, 1\}^n$ denotes the vector with j th component equal to 1 if $j \in S$ and 0, otherwise. Taking $a = 1$ in Definition 2.1 and applying our assumption that M has UUTP(c_0, γ_0), then for any $S, T \subset [n]$ and any $\gamma > \gamma_0$,

$$\mathbf{P}[e_M(S, T) \geq (1 + \gamma)\mu(S, T)] \leq \exp(-c_0h(\gamma - \gamma_0)\mu(S, T)).$$

By a union bound, for any $s, t \in [n]$,

$$(56) \quad \begin{aligned} & \mathbf{P}[\exists S, T \subset [n] : |S| = s, |T| = t, e_M(S, T) \geq (1 + \gamma)\mu(S, T)] \\ & \leq \sum_{S \in \binom{[n]}{s}} \sum_{T \in \binom{[n]}{t}} \exp(-c_0h(\gamma - \gamma_0)\mu(S, T)) \\ & \leq \binom{n}{s} \binom{n}{t} \exp\left(- (K + 4)(s \vee t) \log\left(\frac{en}{s \vee t}\right)\right) \\ & \leq \exp\left(- (K + 2)(s \vee t) \log \frac{en}{s \vee t}\right), \end{aligned}$$

where in the last line we used the bound $\binom{n}{k} \leq (ne/k)^k$ along with the fact that $x \mapsto x \log(e/x)$ is increasing on $[0, 1]$. Applying this fact again, we can bound (56) by its value when $s \vee t = 1$, which is $(ne)^{-K-2} \leq n^{-K-2}$. Now by a union bound over the n^2 choices of $s, t \in [n]$, we have that with probability at least $1 - n^{-K}$,

$$(57) \quad \forall S, T \subset [n], \quad e_M(S, T) \leq (1 + \gamma)\mu(S, T).$$

If S, T are such that $\gamma(S, T, n) = \gamma_1$, then on the event that (57) holds,

$$(58) \quad e_M(S, T) \leq (1 + \gamma_1)\mu(S, T) \leq e^2(1 + \gamma_0)^2\delta|S||T|,$$

putting us in case (1) of the discrepancy property with $\kappa_1 = e^2(1 + \gamma_0)^2$. Otherwise, on the event (57), we have

$$(59) \quad e_M(S, T) \leq (1 + \gamma^*)\mu(S, T),$$

and consequently,

$$(60) \quad \begin{aligned} & c_0 \frac{h(\gamma^* - \gamma_0)}{1 + \gamma^*} e_M(S, T) \\ & \leq c_0 h(\gamma^* - \gamma_0) \mu(S, T) \\ & = (K + 4)(|S| \vee |T|) \log\left(\frac{en}{|S| \vee |T|}\right) \end{aligned}$$

by the definition of γ^* . Note that when $\gamma^* \geq \gamma_1 = e^2(1 + \gamma_0)^2 - 1$,

$$(61) \quad \log(1 + \gamma^*) \geq 2 + 2\log(1 + \gamma_0).$$

Hence, we can lower bound

$$\begin{aligned} & \frac{h(\gamma^* - \gamma_0)}{1 + \gamma^*} \\ & = \frac{1 + \gamma^* - \gamma_0}{1 + \gamma^*} \log(1 + \gamma^* - \gamma_0) - \frac{\gamma^* - \gamma_0}{1 + \gamma^*} \\ & = \frac{1 + \gamma^* - \gamma_0}{1 + \gamma^*} \left[\log(1 + \gamma^*) - \log\left(\frac{1 + \gamma^*}{1 + \gamma^* - \gamma_0}\right) - \frac{\gamma^* - \gamma_0}{1 + \gamma^* - \gamma_0} \right] \\ & \geq \frac{1}{1 + \gamma_0} (\log(1 + \gamma^*) - \log(1 + \gamma_0) - 1) \\ & \geq \frac{1}{2(1 + \gamma_0)} \log(1 + \gamma^*) \\ & \geq \frac{1}{2(1 + \gamma_0)} \log \frac{e_M(S, T)}{\mu(S, T)}, \end{aligned}$$

where we used (61) in the fourth line and (59) in the fifth. Combined with (60), we conclude that when $\gamma^* \geq \gamma_1$,

$$(62) \quad e_M(S, T) \log \frac{e_M(S, T)}{\mu(S, T)} \leq \frac{2}{c_0} (1 + \gamma_0)(K + 4)(|S| \vee |T|) \log \frac{en}{|S| \vee |T|}.$$

Finally, note that the left-hand side can only decrease if we replace $\mu(S, T)$ by its upper bound $\delta|S||T|$, putting us in case (2) of the discrepancy property, with $\kappa_2 = 2(1 + \gamma_0)(K + 4)/c_0$ as claimed. \square

The following deterministic lemma shows that when the discrepancy property holds, the heavy couples contribution $f_{H(x)}(A)$ to $x^\top Ax$ is of order $O(\sqrt{d})$, as desired.

LEMMA 6.6 (DP \Rightarrow heavy couples are small). *Let M be a nonnegative symmetric $n \times n$ matrix with all row and column sums bounded by d . Suppose that M has DP($\delta, \kappa_1, \kappa_2$) with $\delta = Cd/n$, for some $C > 0$, $\kappa_1 > 1$, $\kappa_2 \geq 0$. Then for any $x \in \mathbf{S}^{n-1}$,*

$$(63) \quad f_{H(x)}(M) \leq \alpha_0 \sqrt{d},$$

where

$$(64) \quad \alpha_0 = \alpha_0(C, \kappa_1, \kappa_2) = 16 + 32C(1 + \kappa_1) + 64\kappa_2 \left(1 + \frac{2}{\kappa_1 \log \kappa_1} \right).$$

REMARK 6.7. The same argument can be applied to control the heavy couples contribution to bilinear expressions $x^T M y$ for general nonsymmetric matrices M , as was done in [28] for the case that M is a sum of d i.i.d. permutation matrices.

PROOF OF LEMMA 6.6. Fix $x \in \mathbf{S}^{n-1}$. For $i \geq 1$, let

$$S_i = \left\{ u \in [n] : |x_u| \in \frac{1}{\sqrt{n}} [2^{i-1}, 2^i] \right\}.$$

Note that S_i is empty for $i > \log_2 \sqrt{n} + 1$. For any $(u, v) \in \mathcal{H}(x) \cap (S_i \times S_j)$ we have

$$(65) \quad \frac{\sqrt{d}}{n} \leq |x_u x_v| \leq \frac{2^{i+j}}{n}.$$

Thus,

$$(66) \quad |f_{H(x)}(M)| \leq \sum_{(i,j): 2^{i+j} \geq d} \frac{2^{i+j}}{n} e_M(S_i, S_j).$$

For notational convenience, we would like to assume that $|S_i| \geq |S_j| > 0$. Thus, we define

$$\mathcal{I} := \{(i, j) : 2^{i+j} \geq \sqrt{d}, |S_i| \geq |S_j| > 0\}.$$

Since the summands in (66) are symmetric in i and j ,

$$(67) \quad |f_{H(x)}(M)| \leq 2 \sum_{(i,j) \in \mathcal{I}} \frac{2^{i+j}}{n} e_M(S_i, S_j),$$

with inequality only because pairs (i, j) with $|S_i| = |S_j|$ are counted twice. For $(i, j) \in \mathcal{I}$, denote the discrepancy ratio of the pair (S_i, S_j) by

$$r_{ij} = \frac{e_M(S_i, S_j)}{\delta |S_i| |S_j|}.$$

Define also the quantities

$$(68) \quad \alpha_i := \frac{2^{2i}}{n} |S_i|$$

and

$$(69) \quad s_{ij} := \frac{\sqrt{d}}{2^{i+j}} r_{ij}.$$

In terms of these, the bound (67) becomes

$$(70) \quad \begin{aligned} |f_{H(x)}(M)| &\leq 2 \sum_{(i,j) \in \mathcal{I}} \frac{2^{i+j}}{n} \delta |S_i| |S_j| r_{ij} \\ &= 2C\sqrt{d} \sum_{(i,j) \in \mathcal{I}} \alpha_i \alpha_j \frac{\sqrt{d}}{2^{i+j}} r_{ij} \\ &= 2C\sqrt{d} \sum_{(i,j) \in \mathcal{I}} \alpha_i \alpha_j s_{ij}. \end{aligned}$$

Note that for $(i, j) \in \mathcal{I}$, $s_{ij} \leq r_{ij}$. Note also that

$$(71) \quad \sum_{i \geq 1} \alpha_i = 4 \sum_{i \geq 1} |S_i| \frac{2^{2i-2}}{n} \leq 4 \sum_{i \geq 1} \sum_{u \in S_i} x_u^2 \leq 4.$$

From (70), our aim is to show

$$g(M) := \sum_{(i,j) \in \mathcal{I}} \alpha_i \alpha_j s_{ij} = O(1).$$

We now list our a priori bounds on s_{ij} and r_{ij} . By the assumption that all column sums of M are bounded by d , we have the easy bound

$$e_M(S_i, S_j) \leq d |S_j|$$

giving

$$(72) \quad r_{ij} \leq \frac{d |S_j|}{\delta |S_i| |S_j|} = \frac{n}{C |S_i|} = \frac{2^{2i}}{C \alpha_i}.$$

Now by our assumption that $\text{DP}(\delta, \kappa_1, \kappa_2)$ holds, we have that for all $i, j \geq 1$, either

$$(73) \quad r_{ij} \leq \kappa_1$$

or

$$(74) \quad r_{ij} \log r_{ij} \leq \frac{\kappa_2}{\delta} \frac{1}{|S_j|} \log \frac{en}{|S_i|} \leq \frac{\kappa_2}{\delta n} \frac{2^{2j}}{\alpha_j} \log \frac{2^{2(i+1)}}{\alpha_i},$$

where we have written $2^{2(i+1)}$ rather than 2^{2i} inside the logarithm to absorb the factor e .

In addition to \mathcal{I} , we define the following five sets of pairs (i, j) :

$$\mathcal{I}_1 := \{(i, j) \in \mathcal{I} : s_{ij} \leq \kappa_1\},$$

$$\mathcal{I}_2 := \left\{ (i, j) \in \mathcal{I} : 2^i \leq \frac{2^j}{\sqrt{d}} \right\},$$

$$\mathcal{I}_3 := \left\{ (i, j) \in \mathcal{I} : r_{ij} > \left(\frac{2^{2(i+1)}}{\alpha_i} \right)^{1/4} \right\} \setminus (\mathcal{I}_1 \cup \mathcal{I}_2),$$

$$\mathcal{I}_4 := \left\{ (i, j) \in \mathcal{I} : \frac{1}{\alpha_i} \leq 2^{2(i+1)} \right\} \setminus (\mathcal{I}_1 \cup \mathcal{I}_3),$$

$$\mathcal{I}_5 := \mathcal{I} \setminus (\mathcal{I}_3 \cup \mathcal{I}_4).$$

For $1 \leq k \leq 5$, write

$$g_k(M) = \sum_{(i,j) \in \mathcal{I}_k} \alpha_i \alpha_j s_{ij}.$$

Note that $g(M) \leq \sum_{k=1}^5 g_k(M)$. It remains to show that $g_k = O_{\kappa_1, \kappa_2}(1)$ for each $1 \leq k \leq 5$, which we do in the following five claims.

CLAIM 1. $g_1(M) \leq 16\kappa_1$.

PROOF. Using (71),

$$g_1(M) \leq \kappa_1 \sum_{(i,j) \in \mathcal{I}_1} \alpha_i \alpha_j \leq \kappa_1 \sum_{i \geq 1} \alpha_i \sum_{j \geq 1} \alpha_j \leq 16\kappa_1. \quad \square$$

CLAIM 2. $g_2(M) \leq 8/C$.

PROOF. Here, we use the crude bound (72):

$$\begin{aligned} g_2(M) &= \sum_{(i,j) \in \mathcal{I}_2} \alpha_i \alpha_j \frac{\sqrt{d}}{2^{i+j}} r_{ij} \\ &\leq \sum_{(i,j) \in \mathcal{I}_2} \alpha_i \alpha_j \frac{\sqrt{d}}{2^{i+j}} \frac{2^{2i}}{C\alpha_i} \\ &\leq C^{-1} \sum_{j \geq 1} \alpha_j 2^{-j} \sum_{i: (i,j) \in \mathcal{I}_2} 2^i \sqrt{d}. \end{aligned}$$

As the inner sum is geometric with all terms bounded by 2^j , it is bounded by 2^{j+1} . This gives

$$g_2(M) \leq \frac{2}{C} \sum_{j \geq 1} \alpha_j \leq 8/C. \quad \square$$

CLAIM 3. $g_3(M) \leq 32\kappa_2/C$.

PROOF. First, note that for any $(i, j) \in \mathcal{I} \setminus \mathcal{I}_1$, by (65),

$$r_{ij} \geq \frac{\sqrt{d}}{2^{i+j}} r_{ij} = s_{ij} > \kappa_1.$$

It follows that (74) holds, which gives

$$r_{ij} \leq \frac{\kappa_2}{\delta n} \frac{2^{2j}}{\alpha_j} \frac{\log \frac{2^{2(i+1)}}{\alpha_i}}{\log r_{ij}},$$

and so multiplying through by $\alpha_j \sqrt{d}/2^{i+j}$,

$$(75) \quad \alpha_j s_{ij} \leq \kappa_2 \frac{2^j}{2^i \sqrt{C\delta n}} \frac{\log \frac{2^{2(i+1)}}{\alpha_i}}{\log r_{ij}}.$$

Now the assumption $r_{ij} > (2^{2(i+1)}/\alpha_i)^{1/4}$ gives that the ratio of logarithms is bounded by 4. Hence,

$$\alpha_j s_{ij} \leq 4\kappa_2 \frac{2^j}{2^i \sqrt{C\delta n}} = 4\kappa_2 \frac{2^j}{2^i C \sqrt{d}}.$$

Now

$$g_3(M) \leq \frac{4\kappa_2}{C} \sum_{i \geq 1} \alpha_i 2^{-i} \sum_{j: (i,j) \in \mathcal{I}_2^c} \frac{2^j}{\sqrt{d}} \leq \frac{4\kappa_2}{C} \sum_{i \geq 1} \alpha_i 2^{-i} 2^{i+1} \leq \frac{32\kappa_2}{C},$$

where in the second inequality we used that the inner sum is geometric with every term bounded by 2^i (by the restriction to \mathcal{I}_2^c). \square

CLAIM 4. $g_4(M) \leq \frac{64\kappa_2}{C\kappa_1 \log \kappa_1}$.

PROOF. As in the proof of Claim 3, inequality (75) also holds here since we are summing over $(i, j) \notin \mathcal{I}_1$. Now, by virtue of summing over \mathcal{I}_4 , we have $\frac{1}{\alpha_i} \leq 2^{2(i+1)}$, and hence $\log \frac{2^{2(i+1)}}{\alpha_i} \leq \log 2^{4(i+1)}$. Since $\kappa_1 < s_{ij} \leq r_{ij}$ on $\mathcal{I} \setminus \mathcal{I}_1$, $\log r_{ij} > \log \kappa_1$, so (75) gives

$$\alpha_j s_{ij} \leq \frac{\kappa_2}{\log \kappa_1} \frac{2^j}{\sqrt{C\delta n}} \frac{\log 2^{4(i+1)}}{2^i} \leq \frac{\kappa_2 \log 16}{\log \kappa_1} \frac{2^j}{C \sqrt{d}},$$

where in the second bound we crudely bounded $i + 1 \leq 2^i$. For any $(i, j) \in \mathcal{I}_4 \setminus (\mathcal{I}_3 \cup \mathcal{I}_1)$,

$$\kappa_1 < s_{ij} = \frac{\sqrt{d}}{2^{i+j}} r_{ij} \leq \frac{\sqrt{d}}{2^{i+j}} \left(\frac{2^{2(i+1)}}{\alpha_i} \right)^{1/4} \leq \frac{\sqrt{d}}{2^{i+j}} (2^{4(i+1)})^{1/4} = \frac{\sqrt{d}}{2^{j-1}}.$$

Hence, $2^j/\sqrt{d} < 2/\kappa_1$ for any such (i, j) , so by summing over j first we conclude by similar reasoning as in the previous proof that

$$g_4(M) \leq \frac{4\kappa_2 \log 16}{C\kappa_1 \log \kappa_1} \sum_{i \geq 1} \alpha_i \leq \frac{64\kappa_2}{C\kappa_1 \log \kappa_1}. \quad \square$$

CLAIM 5. $g_5(M) \leq 16$.

PROOF. Now we will sum over i first. Using that $(i, j) \notin \mathcal{I}_3$ for the first inequality, and that $\alpha_i \leq 4$ and $(i, j) \notin \mathcal{I}_4$ for the last, we obtain

$$\alpha_i s_{ij} = \alpha_i \frac{\sqrt{d}}{2^{i+j}} r_{ij} \leq \alpha_i \frac{\sqrt{d}}{2^{i+j}} \left(\frac{2^{2(i+1)}}{\alpha_i} \right)^{1/4} = \alpha_i^{1/2} \frac{\sqrt{d}}{2^{i+j}} (\alpha_i 2^{2(i+1)})^{1/4} \leq 2 \frac{\sqrt{d}}{2^{i+j}}.$$

Summing first the geometric series in i (and noting that all terms in the inner sum are bounded by 1 from the restriction to \mathcal{I}), we have

$$g_5(M) \leq 2 \sum_{j \geq 1} \alpha_j \sum_{i: (i,j) \in \mathcal{I}} \frac{\sqrt{d}}{2^{i+j}} \leq 4 \sum_{j \geq 1} \alpha_j \leq 16. \quad \square$$

All together Claims 1–5 give

$$g(M) = \sum_{(i,j) \in \mathcal{I}} \alpha_i \alpha_j s_{ij} \leq 16\kappa_1 + \frac{8}{C} + \frac{32\kappa_2}{C} \left(1 + \frac{2}{\kappa_1 \log \kappa_1} \right) + 16.$$

Together with (70), this gives the desired result. \square

6.3. *The ε -net and Proof of Proposition 2.4.* Now, we will prove Proposition 2.4 by combining the bounds on the light and heavy couples and applying a union bound over a discretization of \mathbf{S}_0^{n-1} . To achieve this goal, we need the following standard lemma. Recall that for a set $E \subset \mathbb{R}^n$ and $\varepsilon > 0$, a subset $\mathcal{N}_\varepsilon \subset E$ is an ε -net of E if every element of E is within Euclidean distance ε of some element of \mathcal{N}_ε .

LEMMA 6.8 (ε -net). *Let $E \subset \mathbf{S}^{n-1}$ be a subset of the unit sphere, and let $\varepsilon > 0$. There is an ε -net of E of cardinality at most $(1 + 2/\varepsilon)^n$.*

PROOF. Let $\mathcal{N}_\varepsilon \subset E$ be a maximal (under set inclusion) ε -separated set in E . Observe that \mathcal{N}_ε is an ε -net of E . Indeed, if there exists $x \in E$ such that x is distance at least ε from every element of \mathcal{N}_ε , then $\mathcal{N}_\varepsilon \cup \{x\}$ is still ε -separated, contradicting maximality.

Now we bound the cardinality of \mathcal{N}_ε by a volumetric argument. Observe that $(\mathcal{N}_\varepsilon)_{\varepsilon/2}$ —the $\varepsilon/2$ neighborhood of \mathcal{N}_ε —is a disjoint union of balls of radius $\varepsilon/2$. Hence, its volume is $|\mathcal{N}_\varepsilon| \times v_n(\varepsilon/2)^n$, where v_n is the volume of the unit ball in \mathbb{R}^n .

On the other hand, $(\mathcal{N}_\varepsilon)_{\varepsilon/2}$ is contained in $B(0, 1 + \varepsilon/2)$, the volume of which is $v_n(1 + \varepsilon/2)^n$. The claim follows by monotonicity of Lebesgue measure under set inclusion. \square

PROOF OF PROPOSITION 2.4. Let $K > 0$, and denote $\delta = a_1 d/n$, $\gamma_0 = a_3/\sqrt{d}$. By our assumption of $\text{UTP}(c_0, \gamma_0)$ and Lemma 6.4, there are constants $\kappa_1, \kappa_2 > 0$ depending on c_0, a_3, K such that A has $\text{DP}(\delta, \kappa_1, \kappa_2)$ except on an event of probability at most n^{-K} . Hence, letting \mathcal{G} denote the event that $\text{DP}(\delta, \kappa_1, \kappa_2)$ holds, it suffices to show

$$(76) \quad \mathbf{P}(\mathcal{G} \cap \{\lambda(A) \geq \alpha\sqrt{d}\}) \leq 4e^{-n}$$

for α sufficiently large depending on K, c_0, a_1, a_2, a_3 . Let $\varepsilon > 0$ to be fixed later, and let \mathcal{N} be an ε -net of \mathbf{S}_0^{n-1} of size at most $(1 + 2/\varepsilon)^n$ (which exists by Lemma 6.8). By the variational formula (43), continuity of $x \mapsto x^\top Ax$ and the compactness of \mathbf{S}_0^{n-1} , there exists $\tilde{x} \in \mathbf{S}_0^{n-1}$ such that $\lambda(A) = \tilde{x}^\top A\tilde{x}$. Let $x \in \mathcal{N}$ such that $\|x - \tilde{x}\| \leq \varepsilon$. We have

$$\begin{aligned} \lambda(A) &\leq |x^\top Ax| + 2|(x - \tilde{x})^\top Ax| + |(x - \tilde{x})^\top A(x - \tilde{x})| \\ &\leq |x^\top Ax| + (2\varepsilon + \varepsilon^2)\lambda(A), \end{aligned}$$

where in the second line we rescaled $x - \tilde{x}$ to lie in \mathbf{S}_0^{n-1} , and applied the variational formula (43). Taking $\varepsilon = 1/4$, upon rearranging we have

$$(77) \quad \lambda(A) \leq 3|x^\top Ax|$$

(say). Note that with this choice of ε we have $|\mathcal{N}| \leq 9^n$. We have shown that on the event $\{\lambda(A) \geq \alpha\sqrt{d}\}$ there exists $x \in \mathcal{N}$ such that $|x^\top Ax| \geq (\alpha/3)\sqrt{d}$. Hence,

$$\begin{aligned} \mathbf{P}[\mathcal{G} \cap \{\lambda(A) \geq \alpha\sqrt{d}\}] &\leq \sum_{x \in \mathcal{N}} \mathbf{P}[\mathcal{G} \cap \{|x^\top Ax| \geq (\alpha/3)\sqrt{d}\}] \\ (78) \quad &\leq \sum_{x \in \mathcal{N}} \mathbf{P}[\mathcal{G} \cap \{|f_{L(x)}(A)| \geq (\alpha/3)\sqrt{d} - |f_{H(x)}(A)|\}] \\ &\leq \sum_{x \in \mathcal{N}} \mathbf{P}[|f_{L(x)}(A)| \geq (\alpha/3 - \alpha_0)\sqrt{d}], \end{aligned}$$

where in the second line we applied the decomposition (48), and in the third line we applied Lemma 6.6 (taking the constant C there to be a_1) in view of our restriction to \mathcal{G} . Let $\beta = \alpha/3 - \alpha_0 - a_1 - a_2$, and apply Lemma 6.2 and a union bound to show

$$\begin{aligned} \mathbf{P}[\mathcal{G} \cap \{\lambda(A) \geq \alpha\sqrt{d}\}] &\leq 4|\mathcal{N}| \exp\left(-\frac{c_0\beta^2 n}{32(a_1 + \frac{\beta}{12})}\right) \\ &\leq 4(9^n) \exp\left(-\frac{c_0\beta^2 n}{32(a_1 + \frac{\beta}{12})}\right). \end{aligned}$$

Taking α large enough establishes (76), proving the proposition. \square

REMARK 6.9. We now determine just how large α must be in Proposition 2.4. If we take $\beta \geq \max(12a_1, 64/(3c_0))$, then

$$\frac{\beta^2}{a_1 + \frac{\beta}{12}} \geq 6\beta \geq 128/c_0 \geq 32(1 + \log 9)/c_0,$$

and we obtain (76). Together with the assumption $\beta \geq 4a_1a_3$ required by Lemma 6.2, this means we can take

$$(79) \quad \alpha = 3(\alpha_0 + a_1 + a_2) + \max(36a_1, 12a_1a_3, 64/c_0).$$

Further unraveling the constants by looking back at Lemmas 6.4 and 6.6, we have

$$(80) \quad \begin{aligned} \alpha_0 &= 16 + 32a_1(1 + e^2(1 + \gamma_0)^2) \\ &+ \frac{128}{c_0}(1 + \gamma_0)(K + 4) \left(1 + \frac{1}{e^2(1 + \gamma_0)^2(1 + \log(1 + \gamma_0))} \right), \end{aligned}$$

where $\gamma_0 = a_3/\sqrt{d}$.

REMARK 6.10. The restriction $d = O(n^{2/3})$ in Theorem 1.1 arises as follows. The idea of the uniform tails property is that it allows us to show that for a vector $x \in \mathbb{R}^n$ with $|x| = 1$ and $\sum_u x_u = 0$,

$$X - \mathbf{E}X := \sum_{u,v: 0 \leq x_u x_v \leq \frac{\sqrt{d}}{n}} x_u x_v (A_{uv} - \mathbf{E}A_{uv}) = O(\sqrt{d})$$

with high probability. The random variable X has mean on the order of d (this is because the sum is restricted to positive $x_u x_v$). Our construction of a size biased coupling for the uniform model is bounded with probability $1 - O(d/n)$, and Section 3 then gives concentration for X around its mean multiplied by a factor of $1 + O(d/n)$, which introduces a shift of size $O(d^2/n)$. This needs to be $O(\sqrt{d})$ for the argument to work, leading to the condition $d = O(n^{2/3})$.

It is interesting to note that this barrier also appears in a recent result of Bauerschmidt, Knowles and Yau on the local semicircular law for the uniform random d -regular graph with d growing to infinity with n [9]. They also employ double switchings, though in a different manner from the present work, and their analysis requires taking $d = o(n^{2/3})$ (see [9] for a more precise quantitative statement).

Acknowledgments. Larry Goldstein and Tobias Johnson thank the Institute for Mathematics and its Applications in Minneapolis, MN, where some of this research was carried out.

REFERENCES

- [1] ALON, N. (1986). Eigenvalues and expanders. *Combinatorica* **6** 83–96. Theory of computing (Singer Island, Fla., 1984). [MR0875835](#)
- [2] ALON, N. and MILMAN, V. D. (1985). λ_1 , isoperimetric inequalities for graphs, and super-concentrators. *J. Combin. Theory Ser. B* **38** 73–88. [MR0782626](#)
- [3] ARRATIA, R. and BAXENDALE, P. (2015). Bounded size bias coupling: A Gamma function bound, and universal Dickman-function behavior. *Probab. Theory Related Fields* **162** 411–429. [MR3383333](#)
- [4] ARRATIA, R., GOLDSTEIN, L. and KOCHMAN, F. (2015). Size bias for one and all. ArXiv preprint. Available at [arXiv:1308.2729](#).
- [5] BALDI, P., RINOTT, Y. and STEIN, C. (1989). A normal approximation for the number of local maxima of a random function on a graph. In *Probability, Statistics, and Mathematics* 59–81. Academic Press, Boston, MA. [MR1031278](#)
- [6] BARBOUR, A. D., HOLST, L. and JANSON, S. (1992). *Poisson Approximation. Oxford Studies in Probability* **2**. Oxford Univ. Press, New York. [MR1163825](#)
- [7] BARTROFF, J., GOLDSTEIN, L. and IŞLAK, Ü. (2014). Bounded size biased couplings, log concave distributions and concentration of measure for occupancy models. ArXiv preprint. Available at [arXiv:1402.6769](#).
- [8] BAUERSCHMIDT, R., HUANG, J., KNOWLES, A. and YAU, H.-T. (2015). Bulk eigenvalue statistics for random regular graphs. ArXiv preprint. Available at [arXiv:1505.06700](#).
- [9] BAUERSCHMIDT, R., KNOWLES, A. and YAU, H.-T. (2015). Local semicircle law for random regular graphs. ArXiv preprint. Available at [arXiv:1503.08702](#).
- [10] BENNETT, G. (1962). Probability inequalities for the sum of independent random variables. *J. Amer. Statist. Assoc.* **57** 33–45.
- [11] BERNSTEIN, S. (1924). On a modification of Chebyshev’s inequality and of the error formula of Laplace. *Ann. Sci. Inst. Sav. Ukraine, Sect. Math.* **1** 38–49.
- [12] BORDENAVE, C. (2015). A new proof of Friedman’s second eigenvalue theorem and its extension to random lifts. ArXiv preprint. Available at [arXiv:1502.04482](#).
- [13] BOUCHERON, S., LUGOSI, G. and MASSART, P. (2013). *Concentration Inequalities. A Nonasymptotic Theory of Independence*. Oxford Univ. Press, Oxford. [MR3185193](#)
- [14] BRODER, A. and SHAMIR, E. (1987). On the second eigenvalue of random regular graphs. In *28th Annual Symposium on Foundations of Computer Science (Los Angeles)* 286–294. IEEE Comput. Soc. Press, Washington, DC.
- [15] BRODER, A. Z., FRIEZE, A. M., SUEN, S. and UPFAL, E. (1999). Optimal construction of edge-disjoint paths in random graphs. *SIAM J. Comput.* **28** 541–573. [MR1634360](#)
- [16] CHATTERJEE, S. (2007). Stein’s method for concentration inequalities. *Probab. Theory Related Fields* **138** 305–321. [MR2288072](#)
- [17] CHATTERJEE, S. and DEY, P. S. (2010). Applications of Stein’s method for concentration inequalities. *Ann. Probab.* **38** 2443–2485. [MR2683635](#)
- [18] CHEN, L. H. Y., GOLDSTEIN, L. and SHAO, Q.-M. (2011). *Normal Approximation by Stein’s Method. Probability and Its Applications (New York)*. Springer, Heidelberg. [MR2732624](#)
- [19] CHUNG, F., LU, L. and VU, V. (2003). Spectra of random graphs with given expected degrees. *Proc. Natl. Acad. Sci. USA* **100** 6313–6318. [MR1982145](#)
- [20] COJA-OGHLAN, A. and LANKA, A. (2009). The spectral gap of random graphs with given expected degrees. *Electron. J. Combin.* **16** Research Paper 138. [MR2577306](#)
- [21] COOK, N. A. (2017). Discrepancy properties for random regular digraphs. *Random Structures Algorithms* **50** 23–58. [MR3583025](#)
- [22] COOK, N. A. (2017). On the singularity of adjacency matrices for random regular digraphs. *Probab. Theory Related Fields* **167** 143–200. [MR3602844](#)

- [23] DUMITRIU, I., JOHNSON, T., PAL, S. and PAQUETTE, E. (2013). Functional limit theorems for random regular graphs. *Probab. Theory Related Fields* **156** 921–975. [MR3078290](#)
- [24] DUMITRIU, I. and PAL, S. (2012). Sparse regular random graphs: Spectral density and eigenvectors. *Ann. Probab.* **40** 2197–2235. [MR3025715](#)
- [25] FEIGE, U. and OFEK, E. (2005). Spectral techniques applied to sparse random graphs. *Random Structures Algorithms* **27** 251–275. [MR2155709](#)
- [26] FRIEDMAN, J. (1991). On the second eigenvalue and random walks in random d -regular graphs. *Combinatorica* **11** 331–362. [MR1137767](#)
- [27] FRIEDMAN, J. (2008). A proof of Alon’s second eigenvalue conjecture and related problems. *Mem. Amer. Math. Soc.* **195** viii+100. [MR2437174](#)
- [28] FRIEDMAN, J., KAHN, J. and SZEMERÉDI, E. (1989). On the second eigenvalue of random regular graphs. In *Proceedings of the Twenty-First Annual ACM Symposium on Theory of Computing, STOC’89* 587–598. ACM, New York.
- [29] FRIEDMAN, J. and WIGDERSON, A. (1995). On the second eigenvalue of hypergraphs. *Combinatorica* **15** 43–65. [MR1325271](#)
- [30] GHOSH, S. and GOLDSTEIN, L. (2011). Concentration of measures via size-biased couplings. *Probab. Theory Related Fields* **149** 271–278. [MR2773032](#)
- [31] GHOSH, S., GOLDSTEIN, L. and RAIČ, M. (2011). Concentration of measure for the number of isolated vertices in the Erdős–Rényi random graph by size bias couplings. *Statist. Probab. Lett.* **81** 1565–1570. [MR2832913](#)
- [32] GOLDSTEIN, L. and IŞLAK, Ü. (2014). Concentration inequalities via zero bias couplings. *Statist. Probab. Lett.* **86** 17–23. [MR3162712](#)
- [33] GOLDSTEIN, L. and RINOTT, Y. (1996). Multivariate normal approximations by Stein’s method and size bias couplings. *J. Appl. Probab.* **33** 1–17. [MR1371949](#)
- [34] GREENHILL, C., JANSON, S., KIM, J. H. and WORMALD, N. C. (2002). Permutation pseudographs and contiguity. *Combin. Probab. Comput.* **11** 273–298. [MR1909503](#)
- [35] HOEFFDING, W. (1951). A combinatorial central limit theorem. *Ann. Math. Stat.* **22** 558–566. [MR0044058](#)
- [36] HOEFFDING, W. (1963). Probability inequalities for sums of bounded random variables. *J. Amer. Statist. Assoc.* **58** 13–30. [MR0144363](#)
- [37] HOORY, S., LINIAL, N. and WIGDERSON, A. (2006). Expander graphs and their applications. *Bull. Amer. Math. Soc. (N.S.)* **43** 439–561. [MR2247919](#)
- [38] JANSON, S. (2009). The probability that a random multigraph is simple. *Combin. Probab. Comput.* **18** 205–225. [MR2497380](#)
- [39] JOHNSON, T. (2015). Exchangeable pairs, switchings, and random regular graphs. *Electron. J. Combin.* **22** Paper 1.33. [MR3315475](#)
- [40] KESHAVAN, R. H., MONTANARI, A. and OH, S. (2010). Matrix completion from a few entries. *IEEE Trans. Inform. Theory* **56** 2980–2998. [MR2683452](#)
- [41] KRIVELEVICH, M., SUDAKOV, B., VU, V. H. and WORMALD, N. C. (2001). Random regular graphs of high degree. *Random Structures Algorithms* **18** 346–363. [MR1839497](#)
- [42] LUBETZKY, E., SUDAKOV, B. and VU, V. (2011). Spectra of lifted Ramanujan graphs. *Adv. Math.* **227** 1612–1645. [MR2799807](#)
- [43] LUBOTZKY, A. (2012). Expander graphs in pure and applied mathematics. *Bull. Amer. Math. Soc. (N.S.)* **49** 113–162. [MR2869010](#)
- [44] MCKAY, B. D., WORMALD, N. C. and WYSOCKA, B. (2004). Short cycles in random regular graphs. *Electron. J. Combin.* **11** Research Paper 66. [MR2097332](#)
- [45] MILLER, S. J., NOVIKOFF, T. and SABELLI, A. (2008). The distribution of the largest nontrivial eigenvalues in families of random regular graphs. *Exp. Math.* **17** 231–244. [MR2433888](#)
- [46] NILLI, A. (1991). On the second eigenvalue of a graph. *Discrete Math.* **91** 207–210. [MR1124768](#)

- [47] RAIČ, M. (2007). CLT-related large deviation bounds based on Stein's method. *Adv. in Appl. Probab.* **39** 731–752. [MR2357379](#)
- [48] ROSS, N. (2011). Fundamentals of Stein's method. *Probab. Surv.* **8** 210–293. [MR2861132](#)
- [49] STEIN, C. (1986). *Approximate Computation of Expectations*. Institute of Mathematical Statistics Lecture Notes—Monograph Series **7**. IMS, Hayward, CA. [MR0882007](#)
- [50] TIKHOMIROV, K. and YOUSSEF, P. (2016). The spectral gap of dense random regular graphs. ArXiv preprint. Available at [arXiv:1610.01765](#).
- [51] TRACY, C. A. and WIDOM, H. (1994). Level-spacing distributions and the Airy kernel. *Comm. Math. Phys.* **159** 151–174. [MR1257246](#)
- [52] TRAN, L. V., VU, V. H. and WANG, K. (2013). Sparse random graphs: Eigenvalues and eigenvectors. *Random Structures Algorithms* **42** 110–134. [MR2999215](#)
- [53] VU, V. (2008). Random discrete matrices. In *Horizons of Combinatorics*. Bolyai Soc. Math. Stud. **17** 257–280. Springer, Berlin. [MR2432537](#)
- [54] WORMALD, N. C. (1999). Models of random regular graphs. In *Surveys in Combinatorics, 1999 (Canterbury)*. London Mathematical Society Lecture Note Series **267** 239–298. Cambridge Univ. Press, Cambridge. [MR1725006](#)

N. COOK
DEPARTMENT OF MATHEMATICS
STANFORD UNIVERSITY
450 SERRA MALL, BLDG. 380
STANFORD, CALIFORNIA 94305
USA
E-MAIL: nickcook@stanford.edu

L. GOLDSTEIN
T. JOHNSON
DEPARTMENT OF MATHEMATICS
UNIVERSITY OF SOUTHERN CALIFORNIA
3620 S. VERMONT AVE., KAP 104
LOS ANGELES, CALIFORNIA 90089
USA
E-MAIL: larry@usc.edu
tobias.johnson@nyu.edu