

THE DISCREPANCY OF RANDOM RECTANGULAR MATRICES

DYLAN J. ALTSCHULER AND JONATHAN NILES-WEEDE

ABSTRACT. A recent approach to the Beck–Fiala conjecture, a fundamental problem in combinatorics, has been to understand when random integer matrices have constant discrepancy. We give a complete answer to this question for two natural models: matrices with Bernoulli or Poisson entries. For Poisson matrices, we further characterize the discrepancy for any rectangular aspect ratio. These results give sharp answers to questions of Hoberg and Rothvoß (SODA 2019) and Franks and Saks (*Random Structures Algorithms* 2020). Our main tool is a conditional second moment method combined with Stein’s method of exchangeable pairs. While previous approaches are limited to dense matrices, our techniques allow us to work with matrices of all densities. This may be of independent interest for other sparse random constraint satisfaction problems.

Keywords. Discrepancy, Stein’s method, random constraint satisfaction, sparse random graphs.

1. INTRODUCTION

Given a matrix $A \in \mathbb{R}^{m \times n}$, its *discrepancy* is

$$\text{disc}(A) := \min_{u \in \{-1, +1\}^n} \|Au\|_\infty. \quad (1.1)$$

If A has entries in $\{0, 1\}$, we can view A as the incidence matrix of a collection of m subsets of $[n]$. Then this optimization problem asks to partition $[n]$ into two classes so that the imbalance of each of the m sets is small. Discrepancy measures how well this partitioning can be done. This problem, a natural generalization of graph coloring, has been studied widely in combinatorics and computer science, due to its connections to problems such as integer rounding, set balancing, and metric embeddings [17, 40].

The foundational result in the field is Spencer’s celebrated “Six Standard Deviations Suffice” [39], which states that $\text{disc}(A) \leq 6\sqrt{n}$. By contrast, if $m \asymp n$, a typical vector $u \in \{-1, +1\}^n$ has $\|Au\|_\infty = \Theta(\sqrt{n \log n})$. Spencer’s result reveals the surprising fact that it is possible to find a vector u for which $\|Au\|_\infty$ is much smaller than this typical value.

Spencer’s theorem is unimprovable in general, but it raises the question of which matrices enjoy better bounds—specifically, bounds that are independent of the matrix dimensions. Beck and Fiala [13] showed that if A is the incidence matrix of a t -sparse set-system, i.e. A has binary entries and columns summing to at most t , then $\text{disc}(A) \leq 2t - 1$, independent of the dimensions of the matrix. They further conjectured that this bound is improvable to $\text{disc}(A) = \mathcal{O}(\sqrt{t})$. This conjecture has resisted significant progress: to date, the best dimension-independent bound, due to Bukh, is just $\text{disc}(A) \leq 2t - \log^* t$, where \log^* is the iterated logarithm function [14]. If mild dependence on the dimension is allowed, then the best bound is due to Banaszczyk [8], who showed that $\text{disc}(A) = \mathcal{O}(\sqrt{t \log n})$. Neither approach seems likely to yield a proof of Beck and Fiala’s conjectured bound.

Since the Beck–Fiala conjecture seems beyond the grasp of current techniques, there has been recent interest in understanding randomized versions of the problem. A striking finding of this line of work is that the random setting evinces sharply different behavior in two different regimes: loosely speaking, past results show that a random binary $m \times n$ matrix A has discrepancy $\Theta(\sqrt{n})$ when $n \asymp m$, but when n is significantly larger than m , then a dimension-free bound is possible—in fact, $\text{disc}(A) = 1$ with high probability. This is the smallest possible discrepancy, since any row whose sum is odd must have discrepancy at least 1.

D.J. ALTSCHULER

COURANT INSTITUTE, NEW YORK UNIVERSITY, 251 MERCER STREET, NEW YORK, NY 10012, USA.

J. NILES-WEEDE

COURANT INSTITUTE, NEW YORK UNIVERSITY, 251 MERCER STREET, NEW YORK, NY 10012, USA.

E-mail addresses: dylan.altschuler@courant.nyu.edu, jnw@cims.nyu.edu.

DJA and JNW were supported in part by NSF grant DMS-2015291. DJA was supported in part by NSF Graduate Research Fellowship Program grant DGE-1839302.

Prior work has investigated this phenomenon in a variety of different parameter ranges, but understanding exactly when $\text{disc}(A) \leq 1$ is achievable has remained an open question. If A is an $m \times n$ matrix¹ with independent Bernoulli(p) entries, then it is known that $\text{disc}(A) = \mathcal{O}(1)$ with high probability as long as A is very wide ($n \gg m^2$) and relatively dense ($p \gg n^{-1/2}$). In the special symmetric case where $p = 1/2$, Potukuchi [36] proved that $\text{disc}(A) \leq 1$ as long as $n \geq Cm \log m$ for a sufficiently large constant, but this is not known for any other values of p . Taken together, these works suggest the presence of a threshold above which constant discrepancy is achievable, but they do not give a hint of where this threshold should be.

Our first main result solves this problem by identifying the precise location at which this transition occurs. Strikingly, the threshold is independent of p , and is valid even for $p = p(n)$ varying with n .

Theorem 1. *Let $A \in \{0, 1\}^{m \times n}$ be a random matrix whose entries are i.i.d. Bernoulli random variables with parameter $p := p(n)$. There exists a universal constant $C > 0$ such that if $n \geq Cm \log m$, then $\text{disc}(A) \leq 1$ with high probability.*

The constant in Theorem 1 can be made explicit: under mild assumptions, it suffices to let C be any constant strictly larger than $(2 \log 2)^{-1}$, which is precisely the threshold at which the expected number of low-discrepancy vectors becomes large. It is easy to see that this cannot be improved in general; for example, if $p = 1/2$ and $n = Cm \log m$ for $C < (2 \log 2)^{-1}$, then by Markov's inequality the probability that $\text{disc}(A) = \mathcal{O}(1)$ is exponentially small.

Theorem 1 shows that $\text{disc}(A) \leq 1$ holds when $n \geq Cm \log m$ for a sufficiently large constant C irrespective of the value of p , substantially generalizing Potukuchi's result [36]. As we discuss in more detail below, the sparse regime where $p = o(1)$ evinces fundamentally different behavior from the $p = 1/2$ case, and requires different techniques. The key challenge is that approximations based on the central limit theorem—which are valid for dense matrices—become too inaccurate when A is sparse. We therefore need to develop tools to obtain precise approximations in a regime where the CLT and other classic asymptotic methods break down.

Establishing Theorem 1 requires a quantitative understanding of how dependent the events $\{\|Ax\|_\infty \leq 1\}$ and $\{\|Ay\|_\infty \leq 1\}$ are, for pairs of vectors $x, y \in \{\pm 1\}^n$. More precisely, we need an upper-bound on

$$\mathbb{P}[\|Ay\|_\infty \leq 1 \mid \|Ax\|_\infty \leq 1] - \mathbb{P}[\|Ay\|_\infty \leq 1] \quad (1.2)$$

When A is sufficiently dense, both probabilities in Eq. (1.2) can be individually approximated to sufficient accuracy via classical tools. However, when A is too sparse, this naïve approach no longer succeeds. To handle this difficulty, we instead directly compare the two probabilities using a version of Stein's method called *the method of exchangeable pairs* [12, 18].

Originally developed to prove CLTs [42], Stein's method has proven to be a powerful general tool for establishing limit laws for dependent random variables. Informally, to compare a complicated distribution μ_c on a set \mathcal{X} to a target distribution μ_0 , Stein proposed to find an operator T_0 , acting on functions from $\mathcal{X} \rightarrow \mathbb{R}$, which satisfies the requirement that

$$\mathbb{E}_{\mu_0}[T_0 f] = 0 \quad \forall f : \mathcal{X} \rightarrow \mathbb{R}.$$

Then, so long as it can be shown that $\mathbb{E}_{\mu_c}[T_0 f] \approx 0$ for all f in a set of suitably rich test functions, one can conclude that the distribution μ_c is close to μ_0 . Though proving $\mathbb{E}_{\mu_c}[T_0 f] \approx 0$ by hand can be challenging, it is often possible to find an operator T_c , satisfying $\mathbb{E}_{\mu_c}[T_c f] = 0$ for all $f : \mathcal{X} \rightarrow \mathbb{R}$, such that $\|T_c f - T_0 f\|_\infty$ is small—in this case, we will have

$$|\mathbb{E}_{\mu_c}[T_0 f]| = |\mathbb{E}_{\mu_c}[T_0 f] - \mathbb{E}_{\mu_c}[T_c f]| \leq \|T_c f - T_0 f\|_\infty \approx 0,$$

which is the desired claim. The method of exchangeable pairs gives a simple way of constructing T_0 and T_c from reversible Markov chains with stationary measures μ_0 and μ_c .

To use Stein's method, we view Eq. (1.2) as an expression measuring how different the law of $\|Ay\|_\infty$ is from the law of $\|Ay\|_\infty$ conditioned on $\{\|Ax\|_\infty \leq 1\}$. To compare these two measures, we construct two Markov chains with the measures as stationary distributions, and use the method of exchangeable pairs to find suitable operators T_0 and T_c . By ensuring that the two Markov chains have similar transition probabilities, we can guarantee that $T_0 - T_c$ is easy to control. Though Stein's method is well known in the probability literature, its use in the context of the second moment method appears to be new.

¹For notational simplicity, we restrict to even n in the remainder of our paper.

Theorem 1 is stated for matrices with Bernoulli entries. In addition to studying this model, we also introduce a natural extension, about which we can prove more powerful bounds. We formalize these two models in the following definition.

Definition 1.1 (Bernoulli, Poisson Ensembles). *Let A be an $m \times n$ random matrix with independent and identically distributed entries. If A_{ij} is $\text{Bernoulli}(p)$, we say A is from the (m, n, p) -Bernoulli ensemble. If A_{ij} is $\text{Poisson}(\lambda)$, we say A is from the (m, n, λ) -Poisson ensemble.*

By symmetry, we may always assume in the Bernoulli ensemble that $p \leq 1/2$. It is useful to view both ensembles as the adjacency matrices of random bipartite factor graphs, where the columns of A correspond to the vertices and the rows correspond to the factors. The Bernoulli ensemble corresponds to an Erdős-Rényi model, which is the one common in the recent discrepancy literature. The Poisson ensemble is a natural extension in which multi-edges are allowed.

In the Poisson ensemble, we are able to prove a significant generalization of Theorem 1 by characterizing the behavior of the discrepancy for any rectangular matrix with $m = o(n)$. We first define a convenient set of candidate solutions.

Definition 1.2. *A vector $u \in \{-1, +1\}^n$ is balanced if $\sum_{i=1}^n u_i = 0$. We write \mathcal{B} for the set of balanced vectors. For a random matrix A and any $r \geq 0$, we write Z_r for the random variable equal to the number of $u \in \mathcal{B}$ for which $\|Au\|_\infty \leq r$.*

Our second main theorem characterizes the discrepancy of Poisson matrices: a matrix from the Poisson ensemble has discrepancy at most r so long as $\mathbb{E}[Z_r]$ is large.

Theorem 2. *Let A be drawn from the (m, n, λ) -Poisson ensemble. If $m = o(n)$ and $\log \mathbb{E}[Z_r] = \Theta(n)$, then $\text{disc}(A) \leq r + 1$ with high probability.*

Theorem 2 shows that the prediction based on the annealed entropy $\log \mathbb{E}[Z_r]$ is correct: as soon as low discrepancy solutions exist in expectation, they exist with high probability. Note that a converse statement holds by Markov's inequality: if the expected number of solutions is vanishing, then with high probability there are no solutions. This theorem captures the transition from constant discrepancy to $\Theta(\sqrt{n})$ discrepancy that occurs as n ranges between m and $m \log m$. For example, if $n \geq m$ and $n\lambda = \omega(1)$, then it is straightforward to verify that $\log \mathbb{E}[Z_r] = \Theta(n)$ as long as $r = \Omega(2^{-n/m}\sqrt{n\lambda})$, and Theorem 2 therefore guarantees that $\text{disc}(A) = \mathcal{O}(2^{-n/m}\sqrt{n\lambda} + 1)$ for such matrices with high probability. Apart from the $+1$ term, this is the same discrepancy bound that prior work shows is achievable for Gaussian matrices with i.i.d. $\mathcal{N}(0, \lambda)$ entries [45]. Theorem 2 therefore implies that the Poisson ensemble has similar qualitative behavior to a corresponding Gaussian model, even though our proofs reveal that there are significant technical differences between the two settings.

We also remark that since $\sqrt{m\lambda} = \Omega(2^{-n/m}\sqrt{n\lambda})$, Theorem 2 matches the conjectured Beck-Fiala bound by analogy, where $m\lambda$ is the *average* column sparsity (see Fig. 1, right). Returning to our original motivation of understanding when matrices have constant discrepancy, we have as an easy consequence of Theorem 2 an analogue of Theorem 1 for Poisson matrices.

Corollary 1.1. *For A drawn from the (m, n, λ) -Poisson ensemble with $\lambda \leq \text{poly}(n)$ and $n > Cm \log m$ for some universal constant C , then $\text{disc}(A) = 1$ with high probability.*

The proofs of Theorems 1 and 2 are nonconstructive and leave open the question of whether it is possible to find a vector u achieving $\|Au\|_\infty = 1$ in polynomial time. As suggested by Aubin et al. [7], it is possible to compare our model to a planted version of the discrepancy problem where the matrix A is generated from the Bernoulli ensemble conditioned on a particular vector u having low discrepancy. Though we lack a rigorous proof that this planted model is contiguous to our original model, we conjecture that the geometry of the solution space in the original model is well-captured by its planted counterpart. It can be shown in this planted model that clusters of low-discrepancy solutions are isolated from each other, which provides heuristic evidence for the following conjecture, which we view as an attractive question for future work.

Conjecture 1. *For A drawn from the $(m, n, 1/2)$ -Bernoulli ensemble with $n \geq Cm \log m$, there is no efficient algorithm that finds a constant discrepancy solution with high probability.*

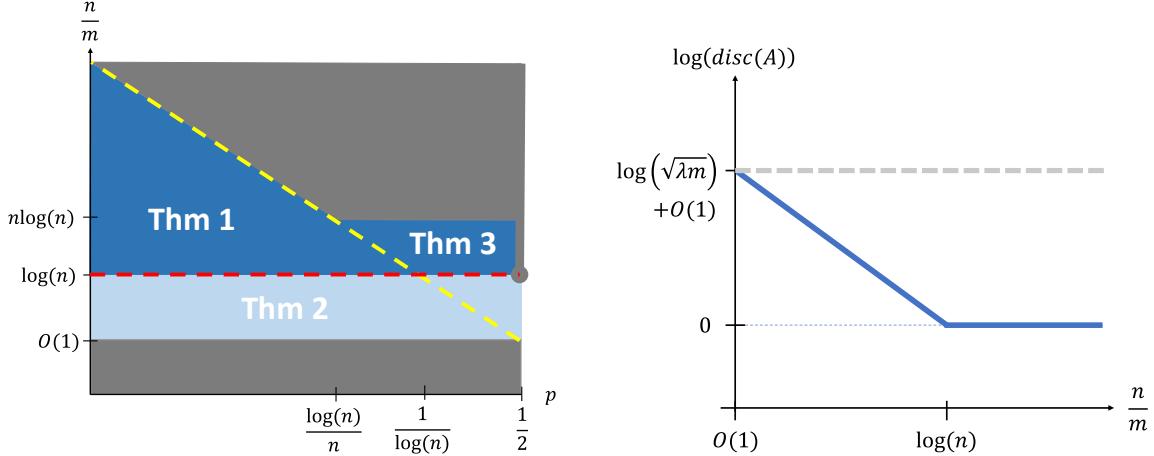


FIGURE 1. **Left:** a diagram of the state-of-the-art (grey) and our contributions (dark blue) for the discrepancy of Bernoulli random matrices. In the case of Poisson matrices (with $\lambda = p$), we contribute the light blue, as well as everything above the light blue. Above the red line, $\text{disc}(A) = 1$ with high probability for either ensemble. Above the yellow line is the “dense” regime. **Right:** the trade-off between aspect-ratio and discrepancy for the Poisson ensemble. The solid blue line is the sharp tradeoff for Poisson matrices given in Theorem 2. The solid grey line is conjectured to be sharp for deterministic λm -sparse matrices (this is the “Beck-Fiala conjecture”), and is a known upper-bound for random λm -sparse matrices [11]. Note the vertical axis is log-scaled.

1.1. Previous Work. A number of recent results study different random matrix models for which $\text{disc}(A) = \mathcal{O}(1)$ with high probability. Ezra and Lovett [21] consider a regular model in which a binary matrix A is chosen uniformly at random conditioned on each column having exactly t ones for some sparsity parameter t . They show when $n > m^t$, then $\text{disc}(A) = \mathcal{O}(1)$ with high probability.

Two independent and concurrent works removed this exponential dependence on t . Franks and Saks [23] consider a fairly general class of matrices and show $\text{disc}(A) \leq 2$ with high probability if $n = \Omega(m^3 \log^2 m)$. Simultaneously, Hoberg and Rothvoß [26] consider A drawn from the (m, n, p) -Bernoulli Ensemble and give the improved bound that $\text{disc}(A) \leq 1$ with high probability if $n = \Omega(m^2 \log m)$ and $mp = \Omega(\log n)$. Shortly after, Potukuchi [36] improved this to $\text{disc}(A) \leq 1$ if $n = \Omega(m \log m)$ in the special symmetric case of $p = 1/2$. Both [23] and [26] used Fourier methods, while [36] used the second moment method. Around the same time as our paper, Macrury et al. [30] showed $\text{disc}(A) \geq \Omega(2^{-n/m} \sqrt{np})$ when $m/(np) \rightarrow 0$, for all $n = \Omega(m)$, via a first moment computation. This is analogous to our Eq. (1.4), but for a wider parameter range.

Bansal and Meka [11] also improved upon [21], except with focus on the Beck-Fiala bound rather than constant discrepancy. They prove that, under a mild growth condition on t , random $m \times n$ binary matrices with t ones per column have $\text{disc}(A) \leq \mathcal{O}(\sqrt{t})$ with high probability. We obtain similar results for Poisson matrices with average column weight t —See the discussion following Theorem 2 comparing the tight rate for random Poisson matrices to the Beck-Fiala bound.

Importantly, all previous work on constant discrepancy on the Bernoulli ensemble requires

$$\lim_{m \rightarrow \infty} \frac{m}{np} \rightarrow 0. \quad (1.3)$$

We call this choice of parameters the *dense* regime. There are *a priori* reasons to expect a non-trivial phase transition when $m/(np) \not\rightarrow 0$; as we show in Theorem 3, this threshold is the point above which the number of optimal solutions to (1.1) no longer concentrates sufficiently well around its expectation. Similar phenomena appear in the analysis of random graphs, whose behavior is very different in the sparse case. Our main technical challenge is proving constant-discrepancy results in the regime where (1.3) does not hold.

A separate line of research has focused on the case where A is a random matrix with independent Gaussian entries [7, 16, 45], showing that $\text{disc}(A) \asymp 2^{-n/m} \sqrt{n}$ when $n = \Omega(m)$; in particular, that $\text{disc}(A) = \mathcal{O}(1)$ once

$n \geq Cm \log m$ for $C > \frac{1}{2 \log 2}$. Discrepancy specialized to the case of iid Gaussian entries and $n = \Theta(m)$ can be seen as a symmetrized version of a famous model in statistical physics known as the binary perceptron model [7, 44], the rigorous understanding of which is an area of active research [19]. In independent and concurrent works, Perkins and Xu [32] and Abbe, Li, and Sly [1] establish the “frozen 1-RSB” geometry of typical solutions in this setting. Further, [32] shows exponential concentration of the number of solutions, while [1] gives an explicit description of the asymptotic distribution of the number of solutions as well as a proof of the “contiguity conjecture” (namely that the planted model and null model are contiguous—see [1, 7] for relevant definitions and discussion).

In the Gaussian case—and more generally, for distributions with sufficiently smooth densities—optimal bounds on the discrepancy can be achieved by a direct application of the second moment method. However, in the sparse, discrete ensembles we consider, the situation is considerably more delicate, and this approach fails. Nevertheless, our results validate the view that, despite being significantly less well behaved than Gaussian matrices, matrices with Bernoulli or Poisson entries also have small discrepancy as soon as $n \geq Cm \log m$.

For square or close-to-square matrices, a variety of efficient algorithms have been discovered matching Spencer’s and Banaszczyk’s bounds [9, 10, 20, 31, 38]. In the Beck–Fiala setting, Potukuchi [36, 37] gives an efficient algorithm achieving $O(\sqrt{t})$ discrepancy for random matrices with t -sparse columns for any $t := t(n)$, as long as $m \geq n$. However, these approaches do not appear to extend to the constant-discrepancy regime when $m = o(n)$. As Hoberg and Rothvoss [26] note, the lack of efficient algorithms for this regime is a common feature of combinatorial problems for which solutions are shown to exist by probabilistic means [29].

The use of the second-moment method for random constraint satisfaction problems was popularized by Achlioptas and Moore [2] and Frieze and Wormald [24]. It has since been successfully applied to a diverse set of problems in theoretical computer science and combinatorics [3–6].

1.2. Heuristics from the first and second moment. Let us first give a heuristic justification for the fact that constant discrepancy is achievable once $n \geq Cm \log m$ for C large enough. Denote by \mathcal{B} the set of balanced vectors and let $Z = Z_1$ be the number of $u \in \mathcal{B}$ for which $\|Au\|_\infty \leq 1$. If $Z > 0$, then $\text{disc}(A) \leq 1$.

Since each of the vectors in \mathcal{B} has an equal probability of satisfying this requirement, fix some $u \in \mathcal{B}$. Then,

$$\mathbb{E}Z = \sum_{u \in \mathcal{B}} \mathbb{P}[\|Au\|_\infty \leq 1] = \binom{n}{n/2} \mathbb{P}[\|Au\|_\infty \leq 1].$$

The entries of Au are independent, and each is a sum of n independent random variables with variance $p(1-p)$, so the local central limit theorem suggests that

$$\mathbb{P}[\|Au\|_\infty \leq 1] = \left(\frac{3}{\sqrt{2\pi np(1-p)}} (1 + o(1)) \right)^m.$$

If $n \gg m$, we therefore expect that

$$\mathbb{E}Z \approx \exp\left(n \log 2 - \frac{m}{2} \log np + o(n)\right) \tag{1.4}$$

So long as $n \geq (1+\varepsilon) \frac{m}{2 \log 2} \log m$, this quantity is exponentially large. In expectation, therefore, $n \geq Cm \log m$ is the right scaling.

Though the annealed entropy $\log \mathbb{E}[Z]$ predicts a threshold at $n \asymp m \log m$, showing that Z is indeed large with high probability past this threshold requires controlling the fluctuations of Z . The classic approach is the so-called *second-moment method*, based on the Paley–Zygmund inequality, which says that a nonnegative integer-valued random variable is positive with high probability as long as $\mathbb{E}[Z^2] = (1 + o(1))\mathbb{E}[Z]^2$ and $\mathbb{E}[Z] > 0$. There is a slight obstruction to naively applying the second moment method to show that the discrepancy of a random matrix is at most 1: the second moment is skewed by the fact that conditioning on $(Au)_i \in \{-1, 0, 1\}$ biases the parity of the i th row of A , because there is only one even number in this set. There are two workarounds: one can ask for $(Au)_i \in \{-2, -1, 1, 2\}$ and prove that $\text{disc}(A) \leq 2$. Or, one can condition on the event that each row of A has even parity, and ask that $Au = 0$. We adopt the second approach and use the following basic construction to extend our result to the unconditioned case.

Lemma 1.1. *Let A be from either the Bernoulli or Poisson ensemble, and let A' be from the same ensemble conditioned on the sum of the entries in each row being even. There exists a coupling of A and A' such that*

$$\text{disc}(A) \leq \text{disc}(A') + 1 \quad a.s.$$

The proof of Lemma 1.1 appears in Appendix B. This conditioning approach was employed by Potukuchi [36] to establish a version of Theorem 1 via the second moment method when $p = 1/2$. We show that this argument can be extended to prove constant discrepancy so long as the matrix satisfies the density requirement (1.3). Moreover, the following theorem shows that, in fact, the density requirement is necessary, and this application of the second moment method provably fails when the matrix is too sparse. Recall Z_r is the number of balanced vectors $u \in \{-1, 1\}^n$ with $\|Au\|_\infty \leq r$.

Theorem 3 (Dense regime). *Let A be drawn from the (m, n, p) -Bernoulli Ensemble. Define P as the event that each row of A sums to an even number. There exists a universal constant $C > 0$ such that for any $n \geq Cm \log m$ and $np = \omega(1)$,*

$$\frac{\mathbb{E}[Z_0^2|P]}{\mathbb{E}[Z_0|P]^2} = \begin{cases} 1 + o(1), & m = o(np) \\ \exp\left\{\Omega\left(\frac{m}{np}\right)\right\}, & m = \Omega(np) \end{cases}$$

Lemma 1.1 immediately yields the following corollary.

Corollary 1.2. *Let A be drawn from the (m, n, p) -Bernoulli ensemble with $n \geq Cm \log m$ and $m = o(np)$. With high probability, $\text{disc}(A) = 1$.*

The constant C is the same constant that appears in Theorem 1, and this suffices to establish Theorem 1 in the dense regime. However, Theorem 3 also shows that this strategy fails when A is sparse. Nevertheless, Theorem 1 maintains that the prediction implied by (1.4) is correct even when the second-moment method fails.

1.3. Our techniques. To prove Theorem 1 in the case where (1.3) does not hold, we employ two strategies. The failure of the second-moment calculation in Theorem 3 stems from the fact that, when $m/np \not\rightarrow 0$, the second moment $\mathbb{E}[Z^2]$ is too sensitive to the sum of the entries in each row of A when A is sparse. We therefore carry out the second-moment method conditional on the weights of each row of A . This technique is common in the literature, and bounding these conditional second-moments still suffices to show that $Z > 0$ with high probability [28].

However, even after conditioning, bounding the second moment requires significant care. For notational simplicity, let us ignore the conditioning argument for now and consider the random variable $Z = Z_1$ counting the $u \in \mathcal{B}$ with $\|Au\|_\infty \leq 1$, as before. Since the entries of Au are i.i.d., we obtain

$$\begin{aligned} \mathbb{E}[Z^2] &= \sum_{u, v \in \mathcal{B}} \mathbb{P}[\|Au\|_\infty \leq 1, \|Av\|_\infty \leq 1] \\ &= \sum_{u, v \in \mathcal{B}} (\mathbb{P}[|(Au)_1| \leq 1, |(Av)_1| \leq 1])^m \\ &= \sum_{u, v \in \mathcal{B}} \mathbb{P}[|(Au)_1| \leq 1 | |(Av)_1| \leq 1]^m \cdot \mathbb{P}[|(Av)_1| \leq 1]^m, \end{aligned}$$

and likewise,

$$\mathbb{E}[Z]^2 = \sum_{u, v \in \mathcal{B}} \mathbb{P}[|(Au)_1| \leq 1]^m \cdot \mathbb{P}[|(Av)_1| \leq 1]^m.$$

To show that $\mathbb{E}[Z^2] = (1 + o(1))\mathbb{E}[Z]^2$, we need to show that for a typical pair $u, v \in \mathcal{B}$, the events $\{|(Au)_1| \leq 1\}$ and $\{|(Av)_1| \leq 1\}$ are approximately independent, so that

$$\mathbb{P}[|(Au)_1| \leq 1 | |(Av)_1| \leq 1]^m \approx \mathbb{P}[|(Au)_1| \leq 1]^m$$

Proving this fact requires approximations on $\mathbb{P}[|(Au)_1| \leq 1 | |(Av)_1| \leq 1]$ which are accurate to $1 + o(m^{-1})$. However, calculating $\mathbb{P}[|(Au)_1| \leq 1 | |(Av)_1| \leq 1]$ explicitly is infeasible; moreover, the local central limit theorem and other classical approximation techniques yield estimates which are accurate only up to a multiplicative factor of $1 + \mathcal{O}\left(\frac{1}{np}\right)$. When $m/np \not\rightarrow 0$, these errors are unacceptably large.

Our second strategy bypasses this difficulty by employing Stein's method. Though this method is well known in the probability literature for its utility in proving limit theorems, to our knowledge the use of this technique combined with the second-moment method is novel. To evaluate $\mathbb{P}[|(Au)_1| \leq 1 | |(Av)_1| \leq 1]^m$, we construct a pair of Markov chains, one of which has stationary distribution given by the law of $(Au)_1$ conditioned on the event $\{|(Av)_1| \leq 1\}$, and the other of which has stationary distribution given by the

law of $(Au)_1$ without conditioning. Stein's method gives a means for comparing these two stationary distributions by inverting a particular functional equation involving the generators of these two chains, which allows us to approximate $\mathbb{P}[|(Au)_1| \leq 1 | (Av)_1 | \leq 1]$ by a simpler, unconditional probability. The resulting approximation has much smaller errors—of order $1 + \mathcal{O}(\frac{1}{n})$ —and this improvement is crucial to obtaining accurate bounds when (1.3) fails.

1.4. Notation. The asymptotic notation $o(\cdot)$, $\mathcal{O}(\cdot)$, $\Omega(\cdot)$, and $\Theta(\cdot)$ refers to the $m \rightarrow \infty$ and therefore $n \rightarrow \infty$ limit. Given a sequence $a = a(m)$ and a nonnegative sequence $b = b(m)$ depending on m , we write $a = \mathcal{O}(b)$ or $b = \Theta(a)$ if $|a(m)| \leq Cb(m)$ for all m sufficiently large. We write $a = o(b)$ if $\lim_{m \rightarrow \infty} \frac{|a(m)|}{b(m)} = 0$. Unless otherwise specified, the implicit constants in these expressions are universal. The phrase “with high probability” means that a sequence of events occurs with probability $1 - o(1)$ in this asymptotic limit. The symbols \wedge and \vee denote min and max respectively. The symbol \log denotes the logarithm base e. We define the binary entropy function H by

$$H(p) := p \log \frac{1}{p} + (1-p) \log \frac{1}{1-p}.$$

1.5. Organization of the remainder of the paper. In Section 2, we formalize the version of the second-moment method that we will employ, and show how to derive Theorem 3. Section 3 introduces Stein's method, and gives the proofs of our central approximation results. The appendices contain additional technical proofs and lemmas.

2. SECOND MOMENT METHOD

The crux of our argument is the second-moment method. Our approach requires two pieces. The first, standard step consists in applying the second-moment method conditionally to ensure that the second moment is not dominated by rare events. We use the following variant of the Paley–Zygmund inequality:

Lemma 2.1 (Conditional Paley–Zygmund [28, Theorem 2.1]). *Let $Z = Z(n)$ be a sequence of nonnegative, integer-valued random variables, and let W be another random variable on the same probability space. If $\mathbb{P}[\mathbb{E}[Z | W] = 0] \rightarrow 0$ and*

$$\frac{\mathbb{E}[Z^2 | W]}{\mathbb{E}[Z | W]^2} \xrightarrow{p} 1,$$

then $Z > 0$ with high probability.

The second step consists of accurately computing the second moment of the conditional distribution, which is the main challenge in our setting. We give a version of the second moment method (similar to Lemma 3 of [2]) tailored for general random constraint satisfaction problems that highlights this aspect. Say a matrix has exchangeable columns if its distribution is invariant under permutations of the columns.

Lemma 2.2 (Second Moment Method for Rectangular CSPs). *Let \mathcal{M} be an ensemble of $m \times n$ matrices with independent rows, exchangeable columns, and $m = o(n)$. Let $A \sim \mathcal{M}$ and fix sets $\mathcal{K}_i \subseteq \mathbb{Z}$ for $i \in [m]$.*

Define $G_i := \{u \in \mathcal{B} : \langle u, A_i \rangle \in \mathcal{K}_i\}$, and let $Z := |\bigcap_i G_i|$ be the number of elements of \mathcal{B} whose inner product with the i th row of A lies in \mathcal{K}_i for all $i \in [m]$. For an arbitrary pair of balanced vectors u and v which agree on βn coordinates, denote

$$\psi_i = \mathbb{P}_{A \sim \mathcal{M}}[u \in G_i], \quad \phi_i(\beta) = \mathbb{P}_{A \sim \mathcal{M}}[u \in G_i, v \in G_i]$$

Suppose that the following conditions hold for n sufficiently large.

- (First Moment:) There exists a positive constant c such that

$$\log \mathbb{E}[Z] > cn \tag{2.1}$$

- (Weak Bound:) For any $\delta \in (0, 1/2)$, there exists a positive constant C_δ such that

$$\phi_i(\beta) \leq C_\delta \psi_i^2 \quad \forall i \in [m], \forall \beta \in [\delta, 1 - \delta] \tag{2.2}$$

- (Strong Bound:) There exists positive universal constants C and ϵ such that

$$\phi_i\left(\frac{1}{2} + x\right) \leq \left(1 + o\left(\frac{1}{m}\right)\right)(1 + Cx^2)\psi_i^2 \quad \forall i \in [m], \forall |x| < \epsilon \tag{2.3}$$

Then the second moment method succeeds: $\mathbb{E}[Z^2] = (1 + o(1))\mathbb{E}[Z]^2$.

Lemma 2.2 is proved in Appendix A.

In the proofs of Theorems 1, 2 and 3, the first-moment bound (2.1) and weak bound (2.2) will follow by simple approximations. In the dense case (Theorem 3), the strong bound is straightforward as well, by applying an Edgeworth expansion for lattice random walks. However, in the context of Theorems 1 and 2, when we no longer assume that $m/np \rightarrow 0$, proving the strong bound (2.3) directly is difficult. While it is easy to show that (2.3) holds with a multiplicative error of $1 + \mathcal{O}\left(\frac{1}{np}\right)$, upgrading the error in (2.3) to $1 + \mathcal{O}\left(\frac{1}{n}\right) = 1 + o\left(\frac{1}{m}\right)$ is the key challenge.

Our main technical idea is to establish (2.3) through Stein's method. By Bayes's rule, (2.3) is equivalent to bounding the difference in the probability mass assigned to the event $u \in G_i$ by the law of u versus the law of u conditioned on the event $v \in G_i$. Stein's method is a powerful tool for proving quantitative comparisons between probability distributions of exactly this type. We survey this approach and show how to derive Theorems 1 and 2 in Section 3. In the remainder of this section, we consider the simpler dense case and use Lemma 2.2 to prove Theorem 3 via a discrete Edgeworth expansion.

The Edgeworth expansion we need for Theorem 3 is for the following simple, lazy random walk:

Definition 2.1. Let $\{X_i\}$ be i.i.d. random variables, each supported on $\{-1, 0, 1\}$ with densities at those points of $p(1-p)$, $p^2 + (1-p)^2$, and $p(1-p)$ respectively. Denote the variance of X_1 as $\sigma^2 := 2p(1-p)$ and the distribution of the lazy random walk $\sum_{i=1}^r X_i$ by $R(r, p)$.

Since X_i can only take three values, its cumulants are easy to compute. The odd cumulants vanish and the first two even cumulants are given by $\kappa_2 = \sigma^2$ and $\kappa_4 = \sigma^2 - 3\sigma^4$. Then, as $r \rightarrow \infty$, Lemma B.5 directly yields the following local central limit theorem

Proposition 2.1. Let $V \sim R(r, p)$. For any constant k :

$$\mathbb{P}[V = k] = \frac{e^{-k^2/(2r\sigma^2)}}{\sqrt{2\pi r\sigma^2}} \left(1 + \frac{(k^4 - 6k^2 + 3)(\sigma^{-2} - 3)}{24r} + \mathcal{O}\left(\frac{1}{r^2\sigma^4}\right) \right)$$

2.1. Proof of Theorem 3. We first show that $\mathbb{E}[Z_0^2|P] = (1 + o(1))\mathbb{E}[Z_0|P]^2$ when $m = o(np)$. Let A be drawn from the (m, n, p) -Bernoulli ensemble conditioned on the event P that each row of A sums to an even number. Define the sets $\mathcal{K}_i := \{0\}$ for all i , and G_i as in Lemma 2.2. Note the rows of A are i.i.d., so we can suppress the subscripts i (e.g. in ψ_i, ϕ_i , and G_i) without ambiguity.

Consider the distribution of a single row of A without conditioning on P . Let $u \in \mathcal{B}$ be a balanced vector (Definition 1.2), and assume without loss of generality that the first $n/2$ coordinates of u are $+1$ and the last $n/2$ are -1 . The number of ones in the first $n/2$ coordinates of A_1 is distributed as $\text{Bin}\left(\frac{n}{2}, p\right)$, as is the number of ones in the last $n/2$ coordinates. Thus, adding these together, $\langle u, A_1 \rangle$ has exactly the distribution of the lazy random walk $R(n/2, p)$.

In this notation, $\mathbb{P}[u \in G] = \mathbb{P}[U = 0]$ where $U \sim R(n/2, p)$. We can now compute the probability that a generic balanced vector u has $u \in G$. Since the event $\{u \in G\}$ contains the event P ,

$$\psi = \mathbb{P}[u \in G|P] = \mathbb{P}[u \in G]/\mathbb{P}[P] = \mathbb{P}[U = 0]/\mathbb{P}[P] \quad (2.4)$$

Before giving an asymptotic expression for ψ , let us derive the corresponding expression for ϕ . Let $\delta \in (0, 1/2)$ be an arbitrary constant. Fix some integer $r \in \{0, \dots, n/2\}$ such that $2r/n \in (\delta, 1 - \delta)$, and consider a pair of balanced vectors $v, w \in \{\pm 1\}^n$ that agree on $2r$ coordinates. Denote the set of indices on which they agree as S . Then, the following events are equal:

$$\begin{aligned} \{v \in G, w \in G\} &= \left\{ \left(\sum_{j \in S} v_j A_{ij} + \sum_{j \in S^c} v_j A_{ij} \right) \in \mathcal{K}, \left(\sum_{j \in S} v_j A_{ij} - \sum_{j \in S^c} v_j A_{ij} \right) \in \mathcal{K} \right\} \\ &= \left\{ \sum_{j \in S} v_j A_{ij} = 0, \sum_{j \in S^c} v_j A_{ij} = 0 \right\} \end{aligned}$$

Note that before we condition on P , the random variables $\sum_{j \in S} v_j A_{ij}$ and $\sum_{j \in S^c} v_j A_{ij}$ are independent with respective distributions $R(r, p)$ and $R\left(\frac{n}{2} - r, p\right)$. So, define two independent random variables V and

V' with $V \sim R(r, p)$ and $V' \sim R(n/2 - r, p)$. Then, since both $\{v \in G\}$ and $\{w \in G\}$ contain the event P ,

$$\phi\left(\frac{2r}{n}\right) = \mathbb{P}[v \in G, w \in G | P] = \mathbb{P}[V = 0, V' = 0] / \mathbb{P}[P] \quad (2.5)$$

Now we give asymptotic expressions for ψ and ϕ by evaluating $\mathbb{P}[P]$ and using the LCLT given in Proposition 2.1. Since we have assumed $2r/n \in (\delta, 1 - \delta)$, we have that $r = \Theta(n)$. This implies that $r\sigma^2 = \Theta(np)$, and since $np = \omega(1)$ we obtain

$$\mathbb{P}[V = 0] = \frac{1}{\sqrt{2\pi r\sigma^2}} \left(1 + \frac{\sigma^{-2} - 3}{8r} + o\left(\frac{1}{np}\right) \right) \quad (2.6)$$

Next, recall P is the event all rows of the matrix A have even sums. Let $P =: \bigcap_{i=1}^m P_i$, where $\{P_i\}$ are the iid events that each row i has an even sum. We also introduce the abbreviation $\beta := 2r/n$, where $\beta \in (\delta, 1 - \delta)$. Returning to the definitions of ϕ and ψ (2.4) and (2.5) and applying (2.6) yields

$$\begin{aligned} \psi^2 &= \frac{1}{\pi n \sigma^2} \left(1 + \frac{\sigma^{-2} - 3}{2n} + o\left(\frac{1}{np}\right) \right) / \mathbb{P}[P]^2 \\ \phi(\beta) &= \frac{1}{\pi n \sigma^2 \sqrt{\beta(1 - \beta)}} \left(1 + \frac{\sigma^{-2} - 3}{4n\beta(1 - \beta)} + o\left(\frac{1}{np}\right) \right) / \mathbb{P}[P] \end{aligned} \quad (2.7)$$

It remains to compute $\mathbb{P}[P_1]$. We claim:

$$\mathbb{P}[P_1] = \frac{1}{2} + o\left(\frac{1}{np}\right) \quad (2.8)$$

To see this, we first note that a simple induction gives the probability that a binomial (n, p) random variable is of even parity:

$$\mathbb{P}[P_1] = \frac{1}{2} + \frac{1}{2}(1 - 2p)^n.$$

Since $p \leq 1/2$, certainly $\mathbb{P}[P_1] \geq 1/2$, and since $np = \omega(1)$, we have

$$(1 - 2p)^n \leq e^{-2np} = o\left(\frac{1}{np}\right).$$

Now, applying (2.8) to (2.7), we obtain

$$\begin{aligned} \psi^2 &= \frac{4}{\pi n \sigma^2} \left(1 + \frac{\sigma^{-2} - 3}{2n} + o\left(\frac{1}{np}\right) \right) \\ \phi(\beta) &= \frac{2}{\pi n \sigma^2 \sqrt{\beta(1 - \beta)}} \left(1 + \frac{\sigma^{-2} - 3}{4n\beta(1 - \beta)} + o\left(\frac{1}{np}\right) \right) \end{aligned} \quad (2.9)$$

In particular, if $m = o(np)$, we have

$$\frac{\phi(\beta)}{\psi^2} = \left(1 + o\left(\frac{1}{m}\right) \right) \frac{1}{2\sqrt{\beta(1 - \beta)}}$$

We now verify the conditions of Lemma 2.2. First, since

$$\mathbb{E}[Z_0 | P] = \binom{n}{n/2} \psi^m = 2^{n(1+o(1))} \left(\frac{2}{\sqrt{\pi n \sigma^2}} \right)^m \left(1 + \mathcal{O}\left(\frac{1}{np}\right) \right)^m = \exp\left(n \log 2 - \frac{m}{2} \log np + o(n)\right),$$

we have that $\log \mathbb{E}[Z_0 | P] > cn$ so long as $n \geq Cm \log m$ for any constant $C > (2 \log 2)^{-1}$. This shows that (2.1) holds. Both (2.2) and (2.3) follow from the fact that for any $\delta \in (0, 1/2)$, there exists a constant C_δ such that

$$\frac{1}{2\sqrt{\beta(1 - \beta)}} \leq 1 + C_\delta(\beta - 1/2)^2 \quad \forall \beta \in [\delta, 1 - \delta],$$

so that

$$\phi\left(\frac{1}{2} + x\right) \leq \left(1 + o\left(\frac{1}{m}\right) \right) (1 + C_\delta x^2) \psi^2 \quad \forall |x| \leq \delta.$$

This proves the strong bound (2.3) and, *a fortiori*, the weak bound (2.2).

So, if $m = o(np)$ and $n \geq Cm \log m$ for $C > (2 \log 2)^{-1}$, then Lemma 2.2 yields $\mathbb{E}[Z^2|P] = (1 + o(1))\mathbb{E}[Z|P]^2$. In particular, $\text{disc}(A) = 0$ with high probability over A from the (m, n, p) -Bernoulli ensemble conditioned on P .

This completes one direction of the theorem. Now assume $m/(np) = \Omega(1)$. In particular, we must have $p = o(1)$, so we may assume $\sigma^{-2} - 3 > \sigma^{-2}/2$ for n sufficiently large. Since $1/(\beta(1 - \beta)) \geq 4$, we obtain from (2.7) that for n sufficiently large and β bounded strictly from 0 and 1,

$$\frac{\phi(\beta)}{\psi^2} > \frac{1}{2\sqrt{\beta(1 - \beta)}} \left(1 + \frac{\sigma^{-2} - 3}{2n}\right) \geq 1 + \frac{\sigma^{-2}}{4n} \quad (2.10)$$

We can compute the second moment the same way as in Lemma 2.2:

$$\begin{aligned} \frac{\mathbb{E}[Z_0^2|P]}{\mathbb{E}[Z_0|P]^2} &= \binom{n}{n/2}^{-1} \sum_{r=0}^{n/2} \binom{n/2}{r}^2 \left(\frac{\phi(\frac{2r}{n})}{\psi^2}\right)^m \\ &\geq \binom{n}{n/2}^{-1} \sum_{\delta n/2 < r < (1-\delta)n/2} \binom{n/2}{r}^2 \left(\frac{\phi(\frac{2r}{n})}{\psi^2}\right)^m \end{aligned}$$

We can uniformly lower bound $\phi(2r/n)/\psi^2$ in this last sum with (2.10). This will leave a sum of squared binomial coefficients. By standard tail bounds on the binomial coefficient (Lemma B.3),

$$\sum_{\delta n/2 < r < (1-\delta)n/2} \binom{n/2}{r}^2 = (1 + o(1)) \sum_{r=0}^{n/2} \binom{n/2}{r}^2 = (1 + o(1)) \binom{n}{n/2}$$

Thus, the second moment is exponentially too large: since $\sigma^2 = \Theta(p)$,

$$\frac{\mathbb{E}[Z_0^2|P]}{\mathbb{E}[Z_0|P]^2} \geq (1 + o(1)) \left(1 + \frac{1}{2n\sigma^2}\right)^m = \exp\left\{\Omega\left(\frac{m}{np}\right)\right\}.$$

□

3. STEIN'S METHOD FOR THE BOUNDING THE SECOND MOMENT

In this section, we prove our main results (Theorems 1 and 2) by using Stein's method to establish the inequality (2.3) for the Bernoulli and Poisson ensembles in the sparse regime.

To describe our approach, we begin by rewriting (2.3) as

$$\mathbb{P}_{A \sim \mathcal{M}}[u \in G_i | v \in G_i] \leq \mathbb{P}_{A \sim \mathcal{M}}[u \in G_i] \cdot (1 + C(\beta - 1/2)^2) \cdot (1 + o(m^{-1})). \quad (3.1)$$

The right side involves the law of $\langle u, A_i \rangle$, and the left side involves the law of this same random variable, conditioned on the event that $\langle v, A_i \rangle$ takes particular values, where v is another balanced vector which agrees with u in βn coordinates. Let us write μ_0 and μ_c for the unconditioned and conditioned distribution, respectively, and write \mathbb{E}_0 and \mathbb{E}_c for the corresponding expectation operators. Our key ingredients are:

- (1) Two operators T_0 and T_c , satisfying

$$\mathbb{E}_0[T_0 f] = \mathbb{E}_c[T_c f] = 0 \quad \forall f.$$

- (2) A function $f_{\mathcal{K}}$ satisfying the equation $T_0 f_{\mathcal{K}} = \mathbb{1}_{\mathcal{K}} - \mu_0(\mathcal{K})$.

- (3) A proof that

$$\mathbb{E}_c[T_0 f_{\mathcal{K}}] \leq \mu_0(\mathcal{K}) \cdot C(\beta - 1/2)^2 \cdot (1 + o(m^{-1})). \quad (3.2)$$

The final inequality (3.2) implies (3.1), since $T_0 f_{\mathcal{K}} = \mathbb{1}_{\mathcal{K}} - \mu_0(\mathcal{K})$ and therefore

$$\mathbb{E}_c[T_0 f_{\mathcal{K}}] = \mu_c(\mathcal{K}) - \mu_0(\mathcal{K}).$$

To prove (3.2), we will use the fact that $\mathbb{E}_c[T_c f_{\mathcal{K}}] = 0$, so that $\mathbb{E}_c[T_0 f_{\mathcal{K}}] = \mathbb{E}_c[(T_0 - T_c)f_{\mathcal{K}}]$. We will therefore define T_0 and T_c in such a way that $(T_0 - T_c)$ is easy to control.

This section is organized as follows: first, we define the Stein operator and give general conditions under which we can invert a functional equation of the form $Tf_{\mathcal{K}} = \mathbb{1}_{\mathcal{K}} - \mu_0(\mathcal{K})$. Next, we prove (2.3) holds for the degree-conditioned Poisson ensemble, completing Theorem 2. Finally, using the same techniques, we will prove (2.3) holds for the Binomial ensemble for zero-discrepancy solutions, yielding Theorem 1.

3.1. Stein Operator. Let (S, S') be some exchangeable random variables taking values in $[w] := \{0, 1, \dots, w\}$, with common distribution μ_0 .

Definition 3.1 (Stein Operator). *Fix a constant z . Define the anti-symmetric operator Ξ and the corresponding Stein Operator T_0 by*

$$(\Xi f)(S, S') := z(f(S')\mathbb{1}_{S' > S} - f(S)\mathbb{1}_{S > S'}), \quad T_0 f(S) := \mathbb{E}_0[\Xi f(S, S')|S] \quad (3.3)$$

Note in particular that $\mathbb{E}_0[T_0 f(S)] = 0$ for any bounded f . Now, we would like that T_0 is invertible for the particular class of f needed to examine indicator functions. The following lemma collects the facts we will use about the inverse of T_0 .

Lemma 3.1 (Stein Operator Inverse). *Let $\{a_i\}_{i=0}^w$ and $\{b_i\}_{i=0}^w$ be some sequences that are strictly decreasing and increasing respectively with $a_w = b_0 = 0$ and all other values strictly positive. Define the probability distribution μ on $[w]$ and the operator T on functions from $[w]$ to \mathbb{R} by*

$$\mu(\{s\}) = \mu(\{0\}) \prod_{i=1}^s \frac{a_{i-1}}{b_i}, \quad T f(s) = a_s f(s+1) - b_s f(s) \quad \forall s \in [w], \quad (3.4)$$

where $\mu(\{0\})$ is uniquely determined by the requirement that μ have total mass 1.

Let Δ be the difference operator given by $\Delta f(s) := f(s+1) - f(s)$ for all $s \in [w]$. For any $t \in \{1, 2, \dots, w-1\}$, there exists a bounded function $f : [w] \rightarrow \mathbb{R}$ with the following properties:

- a. (Existence of an inverse) $T_\mu f(s) = \mathbb{1}(s=t) - \mu(\{t\})$
- b. (Monotonicity) f is non-increasing everywhere except between t and $t+1$, where it is increasing. Furthermore, $f(s)$ is non-positive when $s \leq t$ and non-negative when $s \geq t+1$.
- c. (Uniform control)

$$\sup_{0 \leq s \leq w-1} |\Delta f(s)| = \Delta f(t) \leq \min(a_t^{-1}, b_t^{-1})$$

- d. (L^1 bound on Δf)

$$\sum_{s=0}^w |\Delta f(s)| = \mathcal{O}(|\Delta f(t)|) \quad (3.5)$$

All except the last claim, a trivial corollary of the other three, appear as Lemma 1.1.1 and Lemma 9.2.1 of the monograph of Barbour et al. [12]. For completeness, we include a full proof in Appendix A.

We are now ready to prove our two main theorems.

3.2. Proof of Theorem 2. We will apply Lemma 2.2 to the Poisson ensemble conditioned on having fixed row-sums. Under this conditioning, the law of A will still have independent rows and exchangeable columns, meeting the requirements of Lemma 2.2. We will prove the following theorem, which, when combined with Lemma 2.1 and Lemma 1.1, will yield Theorem 2.

Theorem 4. *Fix some even, non-negative numbers $(w_i)_{i=1}^m \in (2\mathbb{N})^m$ as well as some $\{r_i\}_{i=1}^m \in \mathbb{N}^m$ with $r_i = \mathcal{O}(\sqrt{w_i})$ for all i . Let A be drawn from the (m, n, λ) -Poisson ensemble with $m = o(n)$, and denote by $W = (W_1, \dots, W_m)$ the vector of row-sums:*

$$W_i = \sum_{j=1}^n A_{ij} \quad \forall i \in [m].$$

For each i , define $\mathcal{K}_i = \{0, \pm 1, \dots, \pm r_i\}$, and define G_i and Z as in Lemma 2.2. Then, the law of A conditioned on $\{W = w\}$ satisfies the strong and weak bounds of Lemma 2.2. In particular, if the r_i are such that $\log \mathbb{E}[Z|W = w] > cn$ for a constant $c > 0$, then

$$\mathbb{E}[Z^2|W = w] = (1 + o(1))\mathbb{E}[Z|W = w]^2.$$

Let us begin by proving that Theorem 4 implies Theorem 2. The proof of this technical implication can be skipped on first reading.

Proof of Theorem 2 from Theorem 4. By Lemma 1.1 and the Paley-Zygmund inequality, it suffices to show $\mathbb{E}[Z_r^2|P] = (1 + o(1))\mathbb{E}[Z_r|P]^2$. Our plan is to apply Theorem 4 to establish the assumptions of Lemma 2.2.

We begin by approximating $\mathbb{E}[Z_r|P]$ with a similar argument used to reach (1.4). We may assume that r is an even integer. Fix $r \in 2\mathbb{N}$ and let $u \in \mathcal{B}$ be such that the first half of u is $+1$ and the second half is -1 . If $\lambda n = \omega(1)$, we will apply the classical CLT; if $\lambda n = \mathcal{O}(1)$, we will use the fact that any row is empty with constant probability. Either way, for some universal constant $a_0 > 0$, some function δ with $\lim_{x \rightarrow \infty} \delta(x) = 0$, and sufficiently large n :

$$\begin{aligned} \mathbb{E}[Z_r|P] &= \binom{n}{n/2} \mathbb{P}[\|Au\|_\infty \leq r \mid P] = \binom{n}{n/2} \sum_{t=-r/2}^{r/2} \mathbb{P}\left[\sum_{j=1}^{n/2} A_i - \sum_{j=n/2+1}^n A_i = 2t\right] \\ &= 2^{n+o(n)} \left((1 + \delta(\lambda n)) \frac{a_0(r+1)}{\sqrt{\lambda n}} \wedge 1 \right)^m \\ &= 2^{n+o(n)} \left(\frac{r+1}{\sqrt{\lambda n}} \wedge 1 \right)^m \end{aligned} \quad (3.6)$$

The last equality follows by noting $c^m = 2^{o(n)}$ for any fixed $c > 0$. By assumption, $\mathbb{E}[Z_r|P] > e^{c_0 n}$ for some $c_0 > 0$ for all n sufficiently large. Because $m = o(n)$, (3.6) implies that for any constant $a > 0$, $\mathbb{E}[Z_{a\sqrt{\lambda n}}|P] > e^{c_1 n}$ for a constant $c_1 > 0$ depending only on a . So, for instance, if $r > e^{-1}\sqrt{\lambda n}$ and the assumptions of Theorem 2 are met, redefining $r := e^{-1}\sqrt{\lambda n}$ still satisfies the assumptions of Theorem 2 and provides a better upper-bound on discrepancy. Thus, assume without loss of generality $r \leq e^{-1}\sqrt{\lambda n}$. Now consider the three possible cases:

Case 1: $\lambda n = \omega(\log n)$. Define the set $\mathcal{W} := \{w \in (2\mathbb{N})^m : \forall i |w_i - \lambda n| \leq \sqrt{\lambda n \log n}\}$. By Lemma B.6, $\mathbb{P}[W \in \mathcal{W} \mid P] = 1 - o(1)$. And, if $w \in \mathcal{W}$, then $w_i = \lambda n(1 + o(1))$ for all i . Set $r_i := r$ for all i ; since $r \leq e^{-1}\sqrt{\lambda n} = e^{-1}(1 + o(1))\sqrt{w_i}$, we have $r_i \leq 2e^{-1}\sqrt{w_i}$ for all i . Similar to (3.6), we have by Lemma B.1,

$$\mathbb{E}[Z_r|W = w, P] = \binom{n}{n/2} \prod_{i=1}^m \sum_{t=-r/2}^{r/2} \binom{w_i}{w_i/2 + t} 2^{-w_i} = e^{n \log 2 + o(n)} \prod_{i=1}^m \left(\frac{r+1}{\sqrt{w_i}} \wedge 1 \right) \quad (3.7)$$

Since $w_i = (1 + o(1))\lambda n$, $\mathbb{E}[Z_r|W = w, P] = e^{-\mathcal{O}(m)} \mathbb{E}[Z_r|P] > e^{(c-o(1))n}$. Thus, for all $w \in \mathcal{W}$ and all n sufficiently large, $\mathbb{E}[Z|W = w, P] > e^{c_2 n}$ for some universal constant $c_2 > 0$ that in particular does not depend on w . So, we may apply Theorem 4 with $r_i = r$ for all i to conclude that $\mathbb{E}[Z^2|W = w] = (1 + o(1))\mathbb{E}[Z|W = w]^2$ uniformly over $w \in \mathcal{W}$. Since, conditioned on P , $W \in \mathcal{W}$ with high probability, $\mathbb{E}[Z^2|W, P]/\mathbb{E}[Z|W, P]^2 \rightarrow 1$ in probability. Applying Lemma 2.1, we are done.

Case 2: $m \log(\log(n)) = o(n)$ and $\lambda n = \mathcal{O}(\log n)$. Since we are upper-bounding the discrepancy of A , assume $r = 0$ since no better bound is possible. Trivially $r_i = \mathcal{O}(\sqrt{w_i})$, so (3.7) is available. Define $\mathcal{W} := \{w : \forall i, w_i \leq \log(n)^2\}$. By Lemma B.6, $\mathbb{P}[W \in \mathcal{W} \mid P] = 1 - o(1)$. By (3.7), for all $w \in \mathcal{W}$ and n sufficiently large,

$$\mathbb{E}[Z_0|W = w, P] \geq e^{n \log 2 + o(n)} \prod_{i=1}^m \left(\frac{1}{\log n} \right) > e^{n \log 2 - a_1 m \log(\log(n))}$$

Here a_1 is another positive constant. By assumption, $n = \omega(m \log \log n)$, so in particular obtain $\mathbb{E}[Z_0|W = w, P] > e^{9n \log 2}$ for n large enough. Setting $r_i = 0$ for all i and invoking Theorem 4 for each $w \in \mathcal{W}$ yields $\mathbb{E}[Z^2|W = w] = (1 + o(1))\mathbb{E}[Z|W = w]^2$ uniformly over $w \in \mathcal{W}$. As above, an application of Lemma 2.1 yields the claim.

Case 3: $m \log(\log(n)) = \Omega(n)$ and $\lambda n = \mathcal{O}(\log n)$. We would like to again restrict to $W \in \mathcal{W}$ for some asymptotically full measure set \mathcal{W} , and then apply Theorem 4. However, when λn is small, the row-sums of A do not concentrate well. If we naively set $r_i = r$ for all i , we cannot hope to satisfy the theorem's condition that $r_i = \mathcal{O}(\sqrt{w_i})$ for all i . In order to apply Theorem 4, we focus on a specific (random, W -measurable) subset Z'_r of Z_r . Recall, in the notation of Lemma 2.2:

$$Z_r := \left| \bigcap G_i \right|, \quad G_i := \{u \in \mathcal{B} : \langle u, A_i \rangle \in \mathcal{K}\}, \quad \mathcal{K} := \mathbb{Z} \cap [-r, r]$$

Define the (W -measurable) random variables $r'_i := r \wedge \sqrt{W_i}$ for each i . Then, construct the corresponding random sets:

$$Z'_r := \bigcap G'_i, \quad G'_i := \{u \in \mathcal{B} : |(Au)_i| \in \mathcal{K}'_i\}, \quad \mathcal{K}'_i := \mathbb{Z} \cap [-r'_i, r'_i]$$

Certainly $Z'_r \subset Z_r$, so if $Z'_r \neq \emptyset$ then $\text{disc}(A) \leq r$. Since $r'_i \leq \sqrt{W_i}$ for all i , applying Theorem 4 to the set Z'_r and the law of A conditioned on $\{W = w\} \cap P$ establishes the weak and strong bounds of Lemma 2.2. Only the first-moment condition remains: if we can show $\log \mathbb{E}[Z'_r | W = w] = \Theta(n)$ for all w in some \mathcal{W} with $\mathbb{P}[W \in \mathcal{W} | P] = (1 + o(1))$, then Lemma 2.2 may be applied and $Z'_r \neq 0$ with high probability by Lemma 2.1.

We now turn our attention to establishing the first-moment condition. Observe $W_i = 0$ with probability $\Omega(e^{-\lambda n})$; this provides a trivial lower bound on the probability that $u \in \mathcal{B}$ satisfies $|(Au)_i| \leq r$. If $\lambda m = o(1)$,

$$\mathbb{E}[Z'_r | P] \geq 2^{n-o(n)} e^{-\mathcal{O}(\lambda nm)} = 2^{n-o(n)}$$

This would finish the proof, so assume that $\lambda m = \Omega(1)$, and thus $\lambda n = \omega(1)$. By the superposition property for independent Poisson variables, $W_i \sim \text{Po}(n\lambda)$ for each i , and $\sum W_i \sim \text{Po}(nm\lambda)$. For a sufficiently large constant $C > 0$, set

$$\mathcal{W} := \left\{ w \in (2\mathbb{N})^m : \sum_{i=1}^m w_i \leq \lambda mn + C\sqrt{\lambda mn \log n} \right\}$$

By Lemma B.6, we have $\mathbb{P}[W \in \mathcal{W} | P] = (1 + o(1))$. Note $\sum w_i/m = n\lambda(1 + o(1))$ for $w \in \mathcal{W}$ because $\log(n)/m = o(1)$. We now bound $\mathbb{E}[Z'_r | W = w]$ uniformly over $w \in \mathcal{W}$. Observe that if $w_i = 0$, then $|(Au)_i| = 0 \leq r'_i$ with probability 1. Returning to (3.7) and recalling the definition of r'_i ,

$$\begin{aligned} \min_{w \in \mathcal{W}} \mathbb{E}[Z'_r | W = w, P] &= 2^{n+o(n)} \min_{w \in \mathcal{W}} \prod_{i: \substack{1 \leq i \leq m \\ w_i \neq 0}} \left(\frac{(r \wedge \sqrt{w_i}) + 1}{\sqrt{w_i}} \right) \\ &= 2^{n+o(n)} \min_{w \in \mathcal{W}} \prod_{i: \substack{1 \leq i \leq m \\ w_i \neq 0}} \left(\frac{r}{\sqrt{w_i}} \wedge 1 \right) \end{aligned} \tag{3.8}$$

We claim that if $r < \sqrt{w_j}$, then we may assume $w_j = 0$. Indeed, consider $w \in \mathcal{W}$ with $r > \sqrt{w_j} > 0$ for some j . Define w' as a copy of w with a modification: $w'_j = 0$ and for arbitrary $k \neq j$, let $w'_k = w_k + w_j$. Then, a simple computation yields

$$\prod_{i: \substack{1 \leq i \leq m \\ w_i \neq 0}} \left(\frac{r}{\sqrt{w_i}} \wedge 1 \right) \geq \left(\frac{r}{\sqrt{w'_k}} \wedge 1 \right) \prod_{i: \substack{1 \leq i \leq m \\ w_i \neq 0, i \neq j, i \neq k}} \left(\frac{r}{\sqrt{w_i}} \wedge 1 \right) = \prod_{i: \substack{1 \leq i \leq m \\ w'_i \neq 0}} \left(\frac{r}{\sqrt{w'_i}} \wedge 1 \right)$$

Repeating this argument for each index j with $r > \sqrt{w_j} > 0$, we see $\mathbb{E}[Z'_r | W = w, P]$ is minimized (up to a $2^{o(n)}$ factor), by w with $w_j = 0$ or $\sqrt{w_j} > r$, for all j . Now, fix some positive integers S and t . By arithmetic-geometric mean inequality,

$$\max_{\{w_i\}_{i=1}^t : \sum w_i = S} \prod_{i=1}^t \frac{r}{\sqrt{w_i}} \geq \prod_{i=1}^t \frac{r}{\sqrt{S/t}} = \left(\frac{tr^2}{S} \right)^{t/2} =: f_S(t)$$

We apply this to (3.8) for each $w \in \mathcal{W}$ by letting t denote the number of non-zero indices of w and $S = \sum w_i$. Using our observation about the structure of w_j in the second inequality,

$$\min_{w \in \mathcal{W}} \mathbb{E}[Z'_r | W = w] \geq 2^{n+o(n)} \min_{w \in \mathcal{W}} \prod_{i: \substack{1 \leq i \leq m \\ w_i \neq 0}} \left(\frac{r}{\sqrt{w_i}} \wedge 1 \right) \geq 2^{n+o(n)} \min_{w \in \mathcal{W}} \prod_{i: \substack{1 \leq i \leq m \\ w_i \neq 0}} \frac{r}{\sqrt{w_i}} \geq 2^{n+o(n)} \min_{\substack{w \in \mathcal{W} \\ S := \sum w_i}} \min_{t \in [0, m]} f_S(t)$$

It is a simple calculus exercise that $f_S(t)$ is decreasing for $0 < t < S/(r^2 e)$; increasing for $t > S/r^2 e$; and has a global minimum at $S/(r^2 e)$. Recall $S = (1 + o(1))\lambda mn$ and $r \leq e^{-1}\sqrt{\lambda n}$. Then, $S/(r^2 e) > m$, so the minimum of $f_S(t)$ on $[0, m]$ is at $t = m$. Recalling the expression for $\mathbb{E}[Z_r | P]$ given in (3.6),

$$\min_{w \in \mathcal{W}} \mathbb{E}[Z'_r | W = w, P] \geq 2^{n+o(n)} f_{\lambda mn}(m) = 2^{n+o(n)} \left(\frac{r}{\sqrt{\lambda mn}} \right)^m = \mathbb{E}[Z_r | P] e^{o(n)}$$

Since $m = o(n)$ and $\mathbb{E}[Z_r | P] = e^{\Theta(n)}$, then $\mathbb{E}[Z'_r | W = w] = e^{\Theta(n)}$ uniformly over $w \in \mathcal{W}$ and we are done. \square

Proof of Theorem 4. Our goal is to verify that the weak and strong bounds (2.2) and (2.3) hold. Fix a particular i ; we can henceforth suppress all subscripts i without ambiguity. Then, $\mathcal{K} = \{0, \pm 1, \dots, \pm r\}$ for some $r = \mathcal{O}(\sqrt{w})$. Define the set $\mathcal{K}_w := \{(w+k)/2, k \in \mathcal{K}\}$; while this is a slight overload of notation, it will never be ambiguous because the subscript i is suppressed for the remainder of this theorem. Finally, define the shorthand $\gamma := 1 - \beta$. For the Poisson ensemble conditioned on (w_i) described in the theorem statement, and for ϕ and ψ defined in Lemma 2.2,

$$\begin{aligned}\psi &= 2^{-w} \sum_{k \in \mathcal{K}} \binom{w}{(w+k)/2}, \\ \phi(\beta) &= 2^{-w} \sum_{k, k' \in \mathcal{K}} \binom{w}{\frac{w+k}{2}} \sum_c \binom{\frac{w+k}{2}}{\frac{w+k'}{2} - c} \left(\frac{w-k}{c} \right) \beta^{2c + \frac{k-k'}{2}} \gamma^{w + \frac{k'-k}{2} - 2c}\end{aligned}$$

We first verify the weak bound (2.2). Since we restrict to β bounded strictly away from 0 and 1 by constants, the weak bound will follow easily from standard approximations for binomial coefficients (Lemmas B.1 and B.2). Since $r = \mathcal{O}(\sqrt{w})$ and $|k| \leq r$ for each $k \in \mathcal{K}$, Stirling's approximation yields

$$\psi = 2^{-w} \sum_{k \in \mathcal{K}} \binom{w}{(w+k)/2} = \sum_{k \in \mathcal{K}} \Theta\left(\frac{1}{\sqrt{w}}\right) = \Theta(|\mathcal{K}|w^{-1/2}).$$

By a similar argument, making the substitution $c = \beta w/2 + j$ and using again the fact that $k, k' = \mathcal{O}(\sqrt{w})$,

$$\begin{aligned}\phi(\beta) &:= \mathcal{O}\left(\frac{1}{\sqrt{w}}\right) \sum_{k, k' \in \mathcal{K}} \sum_c \left(\binom{\frac{w+k}{2}}{\frac{w+k'}{2} - c} \beta^{\frac{k-k'}{2} + c} \gamma^{\frac{w+k'}{2} - c} \right) \left(\binom{\frac{w-k}{2}}{c} \beta^c \gamma^{\frac{w-k}{2} - c} \right) \\ &\leq \mathcal{O}\left(\frac{1}{\sqrt{w}}\right) \sum_{k, k' \in \mathcal{K}} \sum_{|j| \leq \mathcal{O}(\sqrt{w \log w})} e^{-\Theta(j^2/(w\beta\gamma))} \frac{1}{w\beta\gamma} \\ &= \mathcal{O}\left(\frac{|\mathcal{K}|^2}{w\beta\gamma}\right)\end{aligned}$$

Thus, if $\beta \in [\delta, 1 - \delta]$, then $\phi(\beta) \leq C_\delta \psi^2$. This completes the weak bound.

The strong bound requires a much finer quantitative estimate of ϕ and ψ when β is very close to 1/2. Standard approximation techniques give a $(1 + \mathcal{O}(w^{-1}))$ multiplicative error. Unless we restrict ourselves to the dense case by making the assumption that $m = o(w)$, this error is far too large. Instead, we will use Stein's method of exchangeable pairs to compute $\phi(\beta)$ in terms of ψ_i for β close to 1/2.

Let us consider two balanced vectors, u and v . Since the entries of row A_i of A are i.i.d. Poisson random variables, if we condition on $\{W = w\}$, then A_i can be constructed by starting with the all-zeroes vector of length n , and then choosing w coordinates uniformly at random (with replacement) to increment. By keeping track of whether u is positive or negative in each chosen coordinate, we see $\langle A_i, u \rangle$ is characterized by a binomial random variable. Similarly, the pair $(\langle u, A_i \rangle, \langle v, A_i \rangle)$ is characterized by a multinomial random variable counting how many outcomes in the construction of A_i correspond to coordinates where u and v are both positive, both negative, or of mixed sign.

This description suggests the following construction. We draw w independent random variables from a categorical distribution with four outcomes, labeled $(+, +)$, $(+, -)$, $(-, -)$, and $(-, +)$, where we assign $\beta/2$ probability to each of the outcomes $(+, +)$ and $(-, -)$ and $\gamma/2$ probability to each of the outcomes $(+, -)$ and $(-, +)$. We view these four outcomes as reflecting the signs of the entry of v and the entry of u corresponding to each selected coordinate.

Let $\sigma = (\sigma_{\ddagger}, \sigma_{\pm}, \sigma_=, \sigma_{\mp})$ be the respective counts of how many outcomes of each type are observed. Then σ has a multinomial distribution. By construction,

$$\langle u, A_i \rangle \stackrel{d}{=} \sigma_{\ddagger} + \sigma_{\mp} - \sigma_{\pm} - \sigma_= = w - 2(\sigma_{\pm} + \sigma_=).$$

To obtain an exchangeable pair, we construct another tuple σ' by selecting one of the w outcomes uniformly at random and resampling it from the original categorical distribution. We call the joint law of (σ, σ') generated by this procedure the *unconditioned distribution*, which we denote by \mathbf{P}_0 .

Next, we consider a different process for generating σ , which reflects the law of $\langle u, A_i \rangle$ when we condition on the value of $\langle v, A_i \rangle$. Concretely, we group the four outcomes of our categorical random variables into two

types depending on their first coordinate: $(+, +)$ and $(+, -)$ are one type and $(-, -)$ and $(-, +)$ are another type. Define the events

$$E_k := \left\{ \sigma_{\ddagger} + \sigma_{\pm} = \frac{w+k}{2} \right\}, \quad E_{\mathcal{K}} := \bigcup_{k \in \mathcal{K}} E_k$$

We draw w independent random variables from the same categorical distribution as above, but we condition on the event $E_{\mathcal{K}}$ that the number of $(+, +)$ and $(+, -)$ outcomes is $(w+k)/2$ for some $k \in \mathcal{K}$. This yields a new distribution on the tuple $\sigma = (\sigma_{\ddagger}, \sigma_{\pm}, \sigma_{=}, \sigma_{\mp})$ of counts. To obtain an exchangeable pair, we can generate another tuple σ' from σ by picking one of the w outcomes uniformly at random and resampling it from the categorical distribution *conditioned on the outcome being of the same type*. We call the joint law of the resulting pair the *conditioned distribution*, which we denote by \mathbf{P}_c .

As above, if we view the outcomes as the signs of the entries of v and u corresponding to each selected coordinate, then under \mathbf{P}_c

$$\sigma_{\ddagger} + \sigma_{\mp} - \sigma_{\pm} - \sigma_{=} = w - 2(\sigma_{\pm} + \sigma_{=}) \stackrel{d}{=} \{ \langle u, A_i \rangle \mid \langle v, A_i \rangle \in \mathcal{K} \}$$

We focus on the quantity $S := S(\sigma) = \sigma_{\pm} + \sigma_{=}$ under the conditioned and unconditioned distributions. We have by construction:

$$\mathbf{P}_0[S \in \mathcal{K}_w] = \psi = \mathbf{P}_0[E_{\mathcal{K}}], \quad \mathbf{P}_c[S \in \mathcal{K}_w]\psi = \phi(\beta) \quad (3.9)$$

Thus, we want to show that the probability that $S \in \mathcal{K}_w$ is close under \mathbf{P}_0 and \mathbf{P}_c . We will use Stein's method to compare these probabilities. As in (3.3), define the Stein operators

$$T_0 f(\sigma) = \mathbb{E}_0[\Xi f(S(\sigma), S(\sigma'))|\sigma], \quad T_c f(\sigma) = \mathbb{E}_c[\Xi f(S(\sigma), S(\sigma'))|\sigma].$$

Writing μ_0 and μ_c for the probability measures on \mathbb{N} induced by S under \mathbf{P}_0 and \mathbf{P}_c , we find a function f for which $T_0 f(\sigma) = \mathbb{1}_{S(\sigma) \in \mathcal{K}_w} - \mu_0(\mathcal{K}_w)$, and then we compute

$$\mathbb{E}_c[(T_c - T_0)f(S)] = -\mathbb{E}_c[T_0 f(S)] = \mu_0(\mathcal{K}_w) - \mu_c(\mathcal{K}_w).$$

Carrying out these constructions by means of Lemma 3.1, we obtain the following result whose proof is deferred to the next section.

Lemma 3.2. *Define Δ as the one-step difference operator, $\Delta f(s) := f(s+1) - f(s)$. There exists a function f satisfying*

$$\mu_0(\mathcal{K}_w) - \mu_c(\mathcal{K}_w) = \mathbb{E}_c \left[(\gamma - \beta) \left(\frac{\sigma_{=} - \sigma_{\pm}}{2} \right) \Delta f(S) \right] \quad (3.10)$$

with the property:

$$\sum_{s=0}^w |\Delta f(s)| = \mathcal{O}(|\mathcal{K}|w^{-1}) \quad (3.11)$$

It remains to bound (3.10). Let $x = \beta - 1/2$. We establish the following proposition.

Proposition 3.1. *Uniformly over s ,*

$$|\mathbb{E}_c[(\sigma_{=} - \sigma_{\pm})\mathbb{1}_{S=s}]| = \mathcal{O}\left(xw^{1/2}\right) e^{-\Theta\left(\frac{1}{w}\right)(\frac{w}{2}-s)^2}.$$

The proof of Proposition 3.1 is the most involved part of the theorem. This task, though technical, is significantly simplified by the fact that it suffices to estimate the quantity in question to *constant* multiplicative error, whereas our original goal required estimating $\phi_i(\beta)$ to error $1 + o(\frac{1}{m})$. Before proving Proposition 3.1,

we first show how it implies the strong bound. Applying Proposition 3.1 to (3.10) yields

$$\begin{aligned}
|\mu_0(\mathcal{K}_w) - \mu_c(\mathcal{K}_w)| &= x |\mathbb{E}_c[(\sigma_+ - \sigma_\pm) \Delta f(S)]| \\
&= x \left| \sum_{s=0}^w \mathbb{E}_c[(\sigma_+ - \sigma_\pm) \mathbb{1}_{S=s}] \Delta f(s) \right| \\
&\leq \mathcal{O}\left(x^2 w^{1/2}\right) \sum_{s=0}^w e^{-\Theta(\frac{1}{w})(\frac{w}{2}-s)^2} |\Delta f(s)| \\
&\leq \mathcal{O}\left(x^2 w^{1/2}\right) \sum_{s=0}^w |\Delta f(s)| \\
&= \mathcal{O}\left(x^2 |\mathcal{K}| w^{-1/2}\right),
\end{aligned}$$

where the last step follows from (3.11).

Since S under \mathbf{P}_0 has distribution $\text{Bin}(w, 1/2)$, the de Moivre–Laplace theorem (Lemma B.1) implies that if $k = \mathcal{O}(\sqrt{w})$, then $\mu_0(\{(w+k)/2\}) = \Theta(\mu_0(\{w/2\}))$. In particular, this holds for all $k \in \mathcal{K}$. Thus, $\psi = \mu_0(\mathcal{K}_w) = \Theta(|\mathcal{K}| w^{-1/2})$. Rearranging yields $|\mathcal{K}| = \Theta(\mu_0(\mathcal{K}_w) \sqrt{w})$. So,

$$|\mu_c(\mathcal{K}_w) - \mu_0(\mathcal{K}_w)| = \mathcal{O}(x^2 \mu_0(\mathcal{K}_w)) = \mathcal{O}(x^2 \psi)$$

The last equality follows from the identity (3.9). Using (3.9) again yields the strong bound (2.3) for some universal positive constant C .

$$\phi\left(\frac{1}{2} + x\right) = \psi \mu_c(\mathcal{K}_w) = (1 + Cx^2)\psi^2$$

We have now shown all the assumptions of Lemma 2.2 are satisfied. Thus the second moment method succeeds for the Poisson ensemble conditioned on $\{W = w\}$, proving Theorem 4. \square

It remains to prove Proposition 3.1, modulo some technical lemmas which we defer to the following section.

Proof of Proposition 3.1. Let $G(k, c) := \mathbf{P}_c[\sigma_+ = c, S = s | E_k] = \mathbf{P}_c[\sigma_+ = c, \sigma_\pm = s - c | E_k]$. Conditioned on E_k , σ_+ and σ_\pm are independent, with σ_\pm having distribution $\text{Bin}((w+k)/2, \gamma)$ and σ_+ having distribution $\text{Bin}((w-k)/2, \beta)$. We can therefore write $G(k, c)$ as a product of binomial densities:

$$G(k, c) = \binom{(w+k)/2}{s-c} \binom{(w-k)/2}{c} \beta^{\frac{w+k}{2}+2c-s} \gamma^{\frac{w-k}{2}-2c+s}$$

Fix s with $0 \leq s \leq w$. Our goal is to prove uniformly for $k \in \mathcal{K}$:

$$\mathbb{E}_c[(\sigma_+ - \sigma_\pm) \mathbb{1}_{S=s} | E_k] = \sum_{c=0}^s (2c-s) G(k, c) = \mathcal{O}\left(x w^{1/2}\right) e^{-\Theta(\frac{1}{w})(\frac{w}{2}-s)^2}, \quad (3.12)$$

The first equality is by definition; the second is the claim that directly yields the proposition after averaging over $k \in \mathcal{K}$. We begin with a crude approximation for G . By standard binomial inequalities:

Lemma 3.3. *Uniformly over c ,*

$$G(k, c) = \mathcal{O}\left(\frac{1}{w}\right) \exp\left\{-\Theta\left(\frac{(s-w/2)^2 + (c-\beta s)^2}{w}\right)\right\}$$

We will employ the following basic fact about Gaussian sums, which follows immediately upon comparison with a Gaussian integral.

Lemma 3.4. *Let $q \geq 0$, $r \in \mathbb{R}$, and $w \geq 1$. For some implicit constant depending only on q ,*

$$\sum_{y \in \mathbb{Z}} |y-r|^q e^{-(y-r)^2/w} = \mathcal{O}\left(w^{(q+1)/2}\right),$$

We consider three cases separately, depending on the size of x .

Case 1: $|x| \geq w^{-1/2}$. Here, Lemma 3.3 suffices. Indeed, also applying Lemma 3.4 yields

$$\begin{aligned} |\mathbb{E}_c[(\sigma_{\pm} - \sigma_{\pm})\mathbb{1}_{S=s}|E_k]| &= \mathcal{O}\left(\frac{1}{w}\right) \sum_{c=0}^s (2c-s) e^{-\Theta\left(\frac{1}{w}\right)((c-\beta s)^2 + (s-w/2)^2)} \\ &= \mathcal{O}\left(\frac{1}{w}\right) \left(2 \sum_{c=0}^s (|xs| + |c - \beta s|) e^{-\Theta\left(\frac{1}{w}\right)(c-\beta s)^2}\right) e^{-\Theta\left(\frac{1}{w}\right)(\frac{w}{2}-s)^2} \\ &= \left[\mathcal{O}\left(xw^{1/2}\right) + \mathcal{O}(1)\right] e^{-\Theta\left(\frac{1}{w}\right)(\frac{w}{2}-s)^2} \\ &= \mathcal{O}\left(xw^{1/2}\right) e^{-\Theta\left(\frac{1}{w}\right)(\frac{w}{2}-s)^2}, \end{aligned}$$

where the last step uses the assumption that $|x| \geq w^{-1/2}$. Averaging over $k \in \mathcal{K}$ completes the proposition in the case $x \geq w^{-1/2}$.

Case 2: $|x| \leq w^{-1}$. For this case, we need to exploit a symmetry. We have assumed that $k \in \mathcal{K}$ implies $-k \in \mathcal{K}$. And, since $S \stackrel{d}{=} w - S$, we also have $\mathbf{P}_c[E_k] = \mathbf{P}_c[E_{-k}]$. Thus:

$$\begin{aligned} |\mathbb{E}_c[(\sigma_{\pm} - \sigma_{\pm})\mathbb{1}_{S=s}]| &\leq \sum_{k \geq 0} |\mathbb{E}_c[(\sigma_{\pm} - \sigma_{\pm})\mathbb{1}_{S=s}|E_k] + \mathbb{E}_c[(\sigma_{\pm} - \sigma_{\pm})\mathbb{1}_{S=s}|E_{-k}]|\mathbf{P}_c[E_k]| \\ &= \sum_{k \geq 0} \mathbf{P}_c[E_k] \sum_{c=0}^s |2c-s| |G(k, c) - G(-k, s-c)| \end{aligned} \tag{3.13}$$

The case when $|x| \leq w^{-1}$ now follows from easy arguments. We have:

$$\begin{aligned} G(k, c) - G(-k, s-c) &= \binom{\frac{w+k}{2}}{s-c} \binom{\frac{w-k}{2}}{c} \left(\beta^{\frac{w-k}{2}+2c-s} \gamma^{\frac{w+k}{2}-2c+s} - \beta^{\frac{w+k}{2}-2c+s} \gamma^{\frac{w-k}{2}+2c-s} \right) \\ &= G(k, c) \left(1 - \left(\frac{\beta}{\gamma} \right)^{-k-4c+2s} \right) \end{aligned}$$

Since $x = \mathcal{O}(w^{-1})$, note that $x(k+4c-2s) = \mathcal{O}(1)$. Hence, observing $\beta/\gamma = 1 + \mathcal{O}(x)$ yields

$$\left(1 - \left(\frac{\beta}{\gamma} \right)^{-k-4c+2s} \right) = \mathcal{O}(x(k+4c-2s)).$$

In total, combining this with Lemma 3.3, we obtain

$$|G(k, c) - G(-k, s-c)| = \mathcal{O}(x(k+4c-2s)) \mathcal{O}\left(\frac{1}{w}\right) \exp\left\{-\Theta\left(\frac{1}{w}\right)[(s-w/2)^2 + (\beta s - c)^2]\right\}.$$

Recall $x := \beta - 1/2$. Since $x = \mathcal{O}(w^{-1})$ by assumption, $|2c-s| \leq |2c-2\beta s| + 1$. Applying Lemma 3.4,

$$\begin{aligned} \sum_{c=0}^s |2c-s| |G(k, c) - G(-k, s-c)| &= \mathcal{O}\left(\frac{x}{w}\right) \sum_{c=0}^s (|k| |2c-s| + |2c-s|^2) e^{-\Theta\left(\frac{1}{w}\right)[(s-w/2)^2 + (\beta s - c)^2]} \\ &= \mathcal{O}\left(\frac{x}{w}\right) \sum_{c=0}^s (1 + |k| |c - \beta s| + |c - \beta s|^2) e^{-\Theta\left(\frac{1}{w}\right)[(s-w/2)^2 + (\beta s - c)^2]} \\ &= \mathcal{O}\left(xw^{1/2}\right) e^{-\Theta\left(\frac{1}{w}\right)(\frac{w}{2}-s)^2} \end{aligned}$$

Since this holds uniformly for $k \in \mathcal{K}$, returning to (3.13) completes the proposition for $x \leq w^{-1}$.

Case 3: $w^{-1} \leq |x| \leq w^{-1/2}$. For this case, we again employ a symmetrized expression. Fix a $k \geq 0$. As in (3.13), we have

$$|\mathbb{E}_c[(\sigma_{\pm} - \sigma_{\pm})\mathbb{1}_{S=s}]| \leq \sum_{k \geq 0} |\mathbb{E}_c[(\sigma_{\pm} - \sigma_{\pm})\mathbb{1}_{S=s}|E_k] + \mathbb{E}_c[(\sigma_{\pm} - \sigma_{\pm})\mathbb{1}_{S=s}|E_{-k}]|\mathbf{P}_c[E_k]|$$

First, note that we can assume that $|s - w/2| = \mathcal{O}(\sqrt{w \log w})$. Indeed, by Lemma 3.3,

$$\begin{aligned} |\mathbb{E}_c[(\sigma_+ - \sigma_\pm) \mathbb{1}_{S=s} | E_k] + \mathbb{E}_c[(\sigma_+ - \sigma_\pm) \mathbb{1}_{S=s} | E_{-k}]| &\leq \sum_{c=0}^s |2c - s| |G(k, c) - G(-k, s - c)| \\ &\leq \mathcal{O}(w) \sum_{c=0}^s G(k, c) + G(-k, s - c) \\ &= \mathcal{O}(w^{1/2}) e^{-\frac{C(s-w/2)^2}{w}} \end{aligned}$$

for some positive constant C ; if $|s - w/2| \geq C' \sqrt{w \log w}$ for a sufficiently large positive constant C' , then

$$e^{-\frac{C(s-w/2)^2}{w}} \leq w^{-1} e^{-\frac{C(s-w/2)^2}{2w}} = \mathcal{O}(x) e^{-\frac{C(s-w/2)^2}{2w}}.$$

Therefore, if $|s - w/2| \geq C' \sqrt{w \log w}$, we already have

$$\sum_{c=0}^s |2c - s| |G(k, c) - G(-k, s - c)| = \mathcal{O}(xw^{1/2}) e^{-\Theta(\frac{1}{w})(\frac{w}{2}-s)^2},$$

which is the desired bound. We therefore assume in what follows that $|s - w/2| = \mathcal{O}(\sqrt{w \log w})$.

For this case, we develop a slightly different symmetrized expression based on (3.13). Fix some $k \geq 0$. Writing $[xs]$ for the nearest integer to xs , we have

$$\begin{aligned} \mathbb{E}_c[(\sigma_+ - \sigma_\pm) \mathbb{1}_{S=s} | E_k] + \mathbb{E}_c[(\sigma_+ - \sigma_\pm) \mathbb{1}_{S=s} | E_{-k}] &= \sum_{c=0}^s (2c - s) G(k, c) + \sum_{c=0}^s (2c - s) G(-k, c) \\ &= \sum_{c \in \mathbb{Z}} (2c - s) G(k, c) + \sum_{c \in \mathbb{Z}} (s + 4[xs] - 2c) G(-k, s + 2[xs] - c) \\ &= \sum_{c \in \mathbb{Z}} (2c - 2\beta s) (G(k, c) - G(-k, s + 2[xs] - c)) \\ &\quad + 2xs \sum_{c \in \mathbb{Z}} (G(k, c) - G(-k, s + 2[xs] - c)) \\ &\quad + 4[xs] \sum_{c \in \mathbb{Z}} G(-k, s + 2[xs] - c). \end{aligned}$$

We first claim that the last two sums are small enough. Indeed, using Lemmas 3.3 and 3.4, we see that both terms are bounded by

$$\mathcal{O}\left(\frac{xs}{w}\right) \sum_{c \in \mathbb{Z}} e^{-\Theta(\frac{1}{w})[(s-w/2)^2 + (c-\beta s)^2]} = \mathcal{O}(xw^{1/2}) e^{-\Theta(\frac{1}{w})(\frac{w}{2}-s)^2},$$

which is of the desired size. Moreover, we further claim that

$$\begin{aligned} \sum_{c \in \mathbb{Z}} (2c - 2\beta s) (G(k, c) - G(-k, s + 2[xs] - c)) &= \sum_{|c-\beta s|=\mathcal{O}(\sqrt{w \log w})} (2c - 2\beta s) (G(k, c) - G(-k, s + 2[xs] - c)) \\ &\quad + \mathcal{O}(xw^{1/2}) e^{-\Theta(\frac{1}{w})(\frac{w}{2}-s)^2}. \end{aligned}$$

This truncation is valid because Lemma 3.3 guarantees that there exists a positive constant C such that the portion of the sum outside the range $|c - \beta s| \leq C\sqrt{w \log w}$ contributes at most

$$w^{-1/2} e^{-\Theta(\frac{1}{w})(\frac{w}{2}-s)^2}$$

to the sum; since we have assumed that $w^{-1} \leq |x|$, this error is also of size $\mathcal{O}(xw^{1/2}) e^{-\Theta(\frac{1}{w})(\frac{w}{2}-s)^2}$.

Combining the above bounds, we obtain that

$$\begin{aligned} |\mathbb{E}_c[(\sigma_+ - \sigma_\pm) \mathbb{1}_{S=s}]| &= \sum_{k \geq 0} \mathbf{P}_c[E_k] \sum_{|c-\beta s|=\mathcal{O}(\sqrt{w \log w})} (2c - 2\beta s) (G(k, c) - G(-k, s + 2[xs] - c)) \\ &\quad + \mathcal{O}(xw^{1/2}) e^{-\Theta(\frac{1}{w})(\frac{w}{2}-s)^2}, \quad (3.14) \end{aligned}$$

so it suffices to obtain an accurate approximation of the first sum under the restriction that both $|s - w/2|$ and $|c - \beta s|$ are $\mathcal{O}(\sqrt{w \log w})$. We use the following refinement of Lemma 3.3.

Lemma 3.5. *Define $l := s - \frac{w}{2}$ and $j := c - \beta s$. Let $|j| = \mathcal{O}(\sqrt{w \log w})$, $|l| = \mathcal{O}(\sqrt{w \log w})$, and $k \in \mathcal{K}$. For $x \leq w^{-1/2}$,*

$$G(k, c) = (1 + E(j, l)) \frac{1}{\pi \beta \gamma w} \exp \left\{ -\frac{l^2 + 2j^2}{\beta \gamma w} \right\} \quad (3.15)$$

where $E(j, l)$ denotes a quantity satisfying

$$E(j, l) = \mathcal{O} \left(\frac{1}{w^{1/2}} \right) + \mathcal{O} \left(\frac{(|j| \vee |l|)^3}{w^2} \right)$$

Continuing to write l and j as in the statement of Lemma 3.5, we obtain

$$(G(k, c) - G(-k, s + 2[xs] - c)) = (1 + E(j, l)) \frac{1}{\pi \beta \gamma w} \exp \left\{ -\frac{l^2 + 2j^2}{\beta \gamma w} \right\} - (1 + E(-\tilde{j}, l)) \frac{1}{\pi \beta \gamma w} \exp \left\{ -\frac{l^2 + 2\tilde{j}^2}{\beta \gamma w} \right\},$$

where

$$\tilde{j} := s + 2[xs] - c - \beta s = j + 2([xs] - xs).$$

Since $|\tilde{j} - j| \leq 2$, we have in particular that $E(-\tilde{j}, l) = \mathcal{O}(E(j, l))$, as well as:

$$\exp \left\{ -\frac{l^2 + 2\tilde{j}^2}{\beta \gamma w} \right\} = \left(1 + \mathcal{O} \left(\frac{|j|}{w} \right) \right) \exp \left\{ -\frac{l^2 + 2j^2}{\beta \gamma w} \right\} = (1 + \mathcal{O}(E(j, l))) \exp \left\{ -\frac{l^2 + 2j^2}{\beta \gamma w} \right\}$$

We obtain

$$|(G(k, c) - G(-k, s + 2[xs] - c))| = \mathcal{O}(E(j, l)) \frac{1}{\pi \beta \gamma w} \exp \left\{ -\frac{l^2 + 2j^2}{\beta \gamma w} \right\} \quad (3.16)$$

By Lemma 3.4,

$$\begin{aligned} \sum_{c \in \mathbb{Z}} |c - \beta s| E(j, l) \frac{1}{\pi \beta \gamma w} e^{-\frac{l^2 + 2j^2}{\beta \gamma w}} &= \mathcal{O} \left(\frac{1}{w} \right) e^{-\frac{l^2}{\beta \gamma w}} \sum_{c \in \mathbb{Z}} \left(\frac{|c - \beta s|}{w^{1/2}} + \frac{|c - \beta s|^4}{w^2} + \frac{|c - \beta s| |l|^3}{w^2} \right) e^{-\Theta \left(\frac{(c - \beta s)^2}{w} \right)} \\ &= \left(\mathcal{O} \left(\frac{1}{w^{1/2}} \right) + \mathcal{O} \left(\frac{|l|^3}{w^2} \right) \right) e^{-\frac{l^2}{\beta \gamma w}} \\ &= \mathcal{O} \left(\frac{1}{w^{1/2}} \right) e^{-\frac{l^2}{2\beta \gamma w}}, \end{aligned}$$

where the last step uses that $\mathcal{O} \left(\frac{|l|^3}{w^2} \right) e^{-\frac{l^2}{\beta \gamma w}} = \mathcal{O} \left(\frac{1}{w^{1/2}} \right) e^{-\frac{l^2}{2\beta \gamma w}}$.

Since $w^{-1/2} = \mathcal{O}(xw^{1/2})$, combining this calculation with (3.16) yields

$$\sum_{c \in \mathbb{Z}} |c - \beta s| |(G(k, c) - G(-k, s + 2[xs] - c))| = \mathcal{O}(xw^{1/2}) e^{-\Theta \left(\frac{1}{w} \right) \left(\frac{w}{2} - s \right)^2},$$

and combining this fact with (3.14) finishes the proof of Proposition 3.1. \square

3.3. Proofs of lemmas.

Proof of lemma 3.2. We let $S = S(\sigma)$ and $S' = S(\sigma')$, so that (S, S') is an exchangeable pair under both \mathbf{P}_0 and \mathbf{P}_c . Define a_S and b_S ,

$$\mathbf{P}_0[S' > S \mid S] = \frac{w - S}{2w} =: \frac{a_S}{w}, \quad \mathbf{P}_0[S' < S \mid S] = \frac{S}{2w} =: \frac{b_S}{w}. \quad (3.17)$$

On the other hand, conditioning on E_k for some $k \in \mathcal{K}$:

$$\begin{aligned} \mathbf{P}_c[S' < S \mid \sigma_{\pm}, \sigma_{\mp}, E_k] &= \frac{\sigma_{\pm} \beta + \sigma_{\mp} \gamma}{w}, \\ \mathbf{P}_c[S' > S \mid \sigma_{\pm}, \sigma_{\mp}, E_k] &= \frac{(w/2 - k/2 - \sigma_{\mp}) \beta + (w/2 + k/2 - \sigma_{\pm}) \gamma}{w} \end{aligned} \quad (3.18)$$

Consider the space of functions on $[w]$. We define a skew-symmetric operator Ξ on such functions by

$$(\Xi f)(s, s') = w(f(s') \mathbb{1}_{s'=s+1} - f(s) \mathbb{1}_{s=s'+1})$$

We then define the Stein operators

$$T_0 f(\sigma) := \mathbb{E}_0[(\Xi f)(S, S') \mid \sigma] \quad (3.19)$$

$$T_c f(\sigma) := \mathbb{E}_c[(\Xi f)(S, S') \mid \sigma]. \quad (3.20)$$

Explicitly, by (3.17),

$$T_0 f(\sigma) = a_S f(S+1) - b_S f(S).$$

In particular, $T_0 f$ depends on σ only through $S(\sigma)$, so that T_0 agrees with the operator T on functions on $[w]$ defined by

$$T f(s) = a_s f(s+1) - b_s f(s) \quad \forall s \in [w]. \quad (3.21)$$

Using the tower property of conditional expectation (i.e. the identity $\mathbb{E}[\mathbb{E}[X|Y]] = \mathbb{E}[X]$) to condition over E_k for each $k \in \mathcal{K}$ via (3.18),

$$\begin{aligned} T_c f(\sigma) - T_0 f(\sigma) &= (\gamma - \beta) \sum_{k \in \mathcal{K}} \left(\left(\frac{\sigma_- - \sigma_\pm}{2} \right) \Delta f(S) + \frac{k}{2} f(S+1) \right) \mathbf{P}_c[E_k | E_{\mathcal{K}}] \\ &= (\gamma - \beta) \left(\left(\frac{\sigma_- - \sigma_\pm}{2} \right) \Delta f(S) + f(S+1) \sum_{k \in \mathcal{K}, k \geq 0} \frac{k}{2} (\mathbf{P}_c[E_k | E_{\mathcal{K}}] - \mathbf{P}_c[E_{-k} | E_{\mathcal{K}}]) \right) \end{aligned}$$

Marginally, $\sigma_\pm + \sigma_\pm \sim \text{Bin}(w, 1/2)$, so $\sigma_\pm + \sigma_\pm \stackrel{d}{=} w - (\sigma_\pm + \sigma_\pm)$ and therefore $\mathbf{P}_c[E_k | E_{\mathcal{K}}] = \mathbf{P}_c[E_{-k} | E_{\mathcal{K}}]$ for all $k \in \mathcal{K}$. So, the summation in the last equation cancels to zero and we obtain

$$T_c f(\sigma) - T_0 f(\sigma) = (\gamma - \beta) \left(\frac{\sigma_- - \sigma_\pm}{2} \right) \Delta f(S) \quad (3.22)$$

Since $S(\sigma)$ under μ_0 is an unbiased Binomial random variable with w trials, it is easy to check that taking $\mu := \mu_0(S)$, a_s and b_s as given in (3.17), and T as in (3.21), satisfies condition (3.4). Thus, for each $k \in \mathcal{K}$, we may apply Lemma 3.1 with $t := (w+k)/2$ to obtain a function f_k . Defining the superposition $f := \sum_{k \in \mathcal{K}} f_k$, we have from the linearity of T_0 ,

$$T_0 f(\sigma) = \mathbb{1}_{S \in \mathcal{K}} - \mu_0(\mathcal{K}_w)$$

Taking expectations of both sides of (3.22) with this choice of f establishes the first desired claim:

$$\mathbb{E}_c \left[(\gamma - \beta) \left(\frac{\sigma_- - \sigma_\pm}{2} \right) \Delta f(S) \right] = \mu_0(\mathcal{K}_w) - \mu_c(\mathcal{K}_w)$$

We now turn to bounding $\sum_{s=0}^k |\Delta f(s)|$. Fix some $k \in \mathcal{K}$ and let $t := (w+k)/2$. Then $k = \mathcal{O}(\sqrt{w})$, and so $|t \vee (w-t)| = \Theta(w)$. By Lemma 3.1, part c, we have

$$\Delta f_k(t) \leq \min(a_t^{-1}, b_t^{-1}) = \min(t^{-1}, (w-t)^{-1}) = \mathcal{O}\left(\frac{1}{w}\right)$$

Lemma 3.1, part d, then implies $\sum_s |\Delta f_k(s)| = \mathcal{O}(w^{-1})$. Summing over $k \in \mathcal{K}$ yields the second desired result, completing the lemma. □

$$\sum_s |\Delta f(s)| \leq \sum_s \sum_{k \in \mathcal{K}} |\Delta f_k(s)| = \mathcal{O}(|\mathcal{K}|w^{-1})$$

Proof of lemma 3.3: Write $l = s - w/2$, and let $j = c - \beta s$. By standard tail bounds for binomial random variables (Lemma B.3), we have

$$\begin{aligned} \binom{(w+k)/2}{\gamma s - j} \beta^{\frac{w+k}{2} - \gamma s + j} \gamma^{\gamma s - j} &= \mathcal{O}\left(\frac{1}{\sqrt{w}}\right) \exp\left\{-\frac{2(\gamma l - j - \gamma k/2)^2}{w+k}\right\} \\ \binom{(w-k)/2}{\beta s + j} \beta^{\beta s + j} \gamma^{\frac{w-k}{2} - \beta s - j} &= \mathcal{O}\left(\frac{1}{\sqrt{w}}\right) \exp\left\{-\frac{2(\beta l + j + \beta k/2)^2}{w-k}\right\} \end{aligned}$$

Multiplying these two inequalities together and using the fact that $k = \mathcal{O}(w^{1/2})$ yields

$$G(k, \beta s + j) = \mathcal{O}\left(\frac{1}{w}\right) \exp\left\{-\frac{(l+xk)^2 + (xl+2j+k)^2}{2w}\right\}$$

There are three cases to check. If $|l|, |j| = \mathcal{O}(w^{1/2})$, then the quantity in the exponent is of constant order, so the entire expression is $\mathcal{O}(\frac{1}{w})$, which agrees with the desired bound. If $|l| = \omega(w^{1/2})$ and $|l| \geq |j|$, then $(l + xk)^2 = \Theta(l^2) = \Theta(l^2 + j^2)$, which yields the desired bound. If $|j| = \omega(w^{1/2})$ and $|j| \geq |l|$, then $(xl + 2j + k)^2 = \Theta(j^2) = \Theta(l^2 + j^2)$ as well. All together, we have

$$G(k, \beta s + j) = \mathcal{O}\left(\frac{1}{w}\right) \exp\left\{-\Theta\left(\frac{l^2 + j^2}{w}\right)\right\},$$

as claimed. \square

Proof of Lemma 3.5. Rewrite the definition of G as the product of two binomial densities:

$$G(k, \beta s + j) := \binom{\frac{w+k}{2}}{\gamma s - j} \beta^{\frac{w+k}{2} - \gamma s + j} \gamma^{\gamma s - j} \binom{\frac{w-k}{2}}{\beta s + j} \beta^{\beta s + j} \gamma^{\frac{w-k}{2} - \beta s - j}$$

Let

$$\zeta := \frac{\gamma s - j}{(w+k)/2}.$$

Then, $\zeta - \gamma = \mathcal{O}(|j| \vee |l|)/w = o(1)$. Indeed,

$$\zeta - \gamma = \left(\frac{l}{w} - \frac{2(j+xl)}{w}\right) \left(1 + \mathcal{O}\left(\frac{k}{w}\right)\right) \quad (3.23)$$

and this is $o(1)$ since we have assumed that j and l are both $\mathcal{O}(\sqrt{w \log w})$.

Applying Lemma B.1 and recalling that $k = \mathcal{O}(w^{1/2})$, we obtain

$$\begin{aligned} \binom{\frac{w+k}{2}}{\gamma s - j} \beta^{\frac{w+k}{2} - \gamma s + j} \gamma^{\gamma s - j} &= \left(1 + \mathcal{O}\left(\frac{1}{w}\right)\right) \frac{1}{\sqrt{\pi(w+k)\beta\gamma}} e^{-\frac{w(\zeta-\gamma)^2}{4\beta\gamma} + \mathcal{O}(k(\zeta-\gamma)^2) + \mathcal{O}(\zeta-\gamma) + \mathcal{O}(w(\zeta-\gamma)^3)} \\ &= \left(1 + \mathcal{O}\left(\frac{1}{w}\right) + \mathcal{O}\left(\frac{|j| \vee |l|}{w}\right) + \mathcal{O}\left(\frac{(|j| \vee |l|)^3}{w^2}\right)\right) \frac{1}{\sqrt{\pi(w+k)\beta\gamma}} e^{-\frac{w(\zeta-\gamma)^2}{4\beta\gamma}}. \end{aligned} \quad (3.24)$$

Similarly, if we let

$$\eta := \frac{\beta s + j}{(w-k)/2},$$

then an identical computation shows

$$\eta - \beta = \left(\frac{l}{w} + \frac{2(j+xl)}{w}\right) \left(1 + \mathcal{O}\left(\frac{k}{w}\right)\right) \quad (3.25)$$

and

$$\binom{\frac{w-k}{2}}{\beta s + j} \beta^{\beta s + j} \gamma^{\frac{w-k}{2} - \beta s - j} = \left(1 + \mathcal{O}\left(\frac{1}{w}\right) + \mathcal{O}\left(\frac{|j| \vee |l|}{w}\right) + \mathcal{O}\left(\frac{(|j| \vee |l|)^3}{w^2}\right)\right) \frac{1}{\sqrt{\pi(w-k)\beta\gamma}} e^{-\frac{w(\eta-\beta)^2}{4\beta\gamma}} \quad (3.26)$$

We would like to multiply (3.24) and (3.26) together in order to derive an approximation for G . First, the product of the polynomial prefactors is

$$\frac{1}{\pi\beta\gamma\sqrt{w^2 - k^2}} = \left(1 + \mathcal{O}\left(\frac{k^2}{w^2}\right)\right) \frac{1}{\pi\beta\gamma} = \left(1 + \mathcal{O}\left(\frac{1}{w}\right)\right) \frac{1}{\pi w \beta \gamma}.$$

Combining (3.23) and (3.25) and using the fact that $k = \mathcal{O}(\sqrt{k})$ and $x \leq w^{-1/2}$ yields

$$(\zeta - \gamma)^2 + (\eta - \beta)^2 = \left(1 + \mathcal{O}\left(\frac{k}{w}\right)\right) \left(\frac{2l^2}{w^2} + \frac{8(j+xl)^2}{w^2}\right) = \left(1 + \mathcal{O}\left(w^{-1/2}\right)\right) \left(\frac{2l^2}{w^2} + \frac{8j^2}{w^2}\right);$$

therefore, the product of the exponential terms is

$$\exp\left\{-\frac{2l^2 + 4j^2}{2\beta\gamma w} \left(1 + \mathcal{O}\left(w^{-1/2}\right)\right)\right\} = \left(1 + \mathcal{O}\left(\frac{(|j| \vee |l|)^2}{w^{3/2}}\right)\right) e^{-\frac{2l^2 + 4j^2}{2\beta\gamma w}}.$$

We conclude that

$$\begin{aligned} G(k, \beta t + j) &= \left(1 + \mathcal{O}\left(\frac{1}{w}\right) + \mathcal{O}\left(\frac{|j| \vee |l|}{w}\right) + \mathcal{O}\left(\frac{(|j| \vee |l|)^3}{w^2}\right) + \mathcal{O}\left(\frac{(|j| \vee |l|)^2}{w^{3/2}}\right)\right) \frac{1}{\pi w \beta \gamma} e^{-\frac{2l^2 + 4j^2}{2\beta\gamma w}} \\ &= \left(1 + \mathcal{O}\left(\frac{1}{w^{-1/2}}\right) + \mathcal{O}\left(\frac{(|j| \vee |l|)^3}{w^2}\right)\right) \frac{1}{\pi w \beta \gamma} e^{-\frac{2l^2 + 4j^2}{2\beta\gamma w}}, \end{aligned}$$

as claimed. \square

3.4. Proof of Theorem 1. The proof will closely follow the proof of Theorem 2. Again by Lemma 2.1 and Lemma 1.1, the following suffices:

Theorem 5. *Let A be from the (m, n, p) -Bernoulli ensemble with $n \geq Cm \log m$, where $C > (2 \log 2)^{-1}$, and let P be the event that each row of A sums to an even number. Let W be the vector of row weights. Then, for any such $w \in (2\mathbb{N})^m$ with $w_i < .49n$ for all i ,*

$$\mathbb{E}[Z_0^2 | W = w] = (1 + o(1)) \mathbb{E}[Z_0 | W = w]^2$$

To obtain Theorem 1, let A be from the (m, n, p) -Bernoulli ensemble, conditioned on the event P that each row of A sums to an even number. By assumption $p \leq 1/2$. And, Theorem 3 already established that $\text{disc}(A) = 0$ with high probability if $m = o(np)$, so we may assume that $p \leq .48$. Then, with high probability, $W_i < .49n$ for all i simultaneously. In other words the collection of events $\{W = w\}$ for $w \in (2\mathbb{N})^m$ with $w_i < .49n$ satisfy

$$\sum_w \mathbb{P}[W = w | P] = 1 - o(1).$$

Therefore, by Lemma 2.1, proving Theorem 5 will imply that $\text{disc}(A) = 0$ with high probability. Finally, removing the conditioning on P by Lemma 1.1 proves the claim.

We will check the three conditions of Lemma 2.2 for $\mathcal{K} = \{0\}$ when A is from the Bernoulli ensemble conditioned on $W = w$. Say u and v are balanced $\{-1, +1\}^n$ vectors agreeing on βn coordinates, and again suppress the subscript i whenever not ambiguous, e.g. $w = w_i$, $\phi = \phi_i$, and $\psi = \psi_i$. Recalling $\gamma = 1 - \beta$, we have

$$\psi := \binom{n/2}{w/2}^2 \binom{n}{w}^{-1}, \quad \phi(\beta) = \sum_{t=0}^{w/2} \binom{\beta n/2}{t}^2 \binom{\gamma n/2}{w/2-t}^2 \binom{n}{w}^{-1}$$

By Stirling's formula (Lemma B.2), $\psi = \Theta(w^{-1/2}) = \Omega(n^{-1/2})$. Now consider $\phi(\beta)$ with $\beta \in [\delta, 1 - \delta]$ for some universal constant $\delta > 0$. Since we have assumed that $w_i \leq .49n$ for all i , the binomial coefficients in the definition of ϕ are all non-zero by taking δ sufficiently small. For a sufficiently large constant M , standard hypergeometric tail bounds (Lemma B.4) yield

$$\begin{aligned} \phi(\beta) &= \left(1 + \mathcal{O}\left(\frac{1}{w}\right)\right) \sum_{|t - \beta w/2| \leq M \sqrt{w \log w}} \binom{\beta n/2}{t}^2 \binom{\gamma n/2}{w/2-t}^2 \binom{n}{w}^{-1} \\ &= \left(1 + \mathcal{O}\left(\frac{1}{w}\right)\right) \binom{n/2}{w/2}^2 \binom{n}{w}^{-1} \sum_{|t - \beta w/2| \leq M \sqrt{w \log w}} \binom{\beta n/2}{t}^2 \binom{\gamma n/2}{w/2-t}^2 \binom{n/2}{w/2}^{-2} \\ &= \mathcal{O}\left(\frac{1}{w^{-3/2}}\right) \sum_{t \in \mathbb{Z}} e^{-\Theta\left(\frac{t - \beta w/2)^2}{w}\right)} \\ &= \mathcal{O}(w^{-1}). \end{aligned}$$

where the last step uses Lemma 3.4. Since

$$\mathbb{E}[Z_0 | W = w] = \binom{n}{n/2} \prod_{i=1}^n \psi_i = \binom{n}{n/2} \Omega(n^{-1/2})^m \geq \exp\left\{n \log 2 - (1 + o(1)) \frac{m}{2} \log n\right\},$$

the assumption that $n \geq Cm \log m$ for $C > (2 \log 2)^{-1}$ implies that the first-moment condition (2.1) holds. Similarly, the weak bound (2.2) holds: for some implicit constants that depend only on δ ,

$$\phi(\beta) = \mathcal{O}(w^{-1}) \leq C_\delta \psi_i^2.$$

It remains to establish the strong inequality (2.3). As in the proof of Theorem 2, Stirling's approximation only allows us to compute ϕ and ψ up to an error of $(1 + \mathcal{O}(w^{-1}))$, which we cannot afford. We again use Stein's method to circumvent this challenge.

We adopt a modified version of the construction used in Theorem 2. If we condition on the event that $W_i = w$, we can generate A_i by choosing w coordinates uniformly *without replacement* from $\{1, \dots, n\}$ and setting the corresponding coordinates of A_i to 1. Given two balanced vectors u and v , it again suffices to track how many of the chosen coordinates correspond to entries of u and v which are both positive, both negative, or of mixed sign.

Consider a urn of n balls labeled with $(+, +)$, $(+, -)$, $(-, +)$ or $(-, -)$, where we assign $\beta n/2$ to the labels $(+, +)$ and $(-, -)$ and $\gamma n/2$ to $(+, -)$ and $(-, +)$. We select w balls independently without replacement from this urn, and view each outcome as reflecting the signs of the entry of v and entry of u corresponding to the selected coordinate. As before, we define the vector $\sigma = (\sigma_{\ddagger}, \sigma_{\pm}, \sigma_{=}, \sigma_{\mp})$ of counts, and construct an exchangeable copy of σ' by choosing one of the w selected balls uniformly at random and swapping it with a random ball in the urn. The joint law of (σ, σ') is the *unconditioned distribution*, \mathbf{P}_0 .

For the conditioned distribution, we again divide the balls into two types— $\{(+, +), (+, -)\}$ on the one hand, $\{(-, -), (-, +)\}$ on the other—and consider drawing as above w balls without replacement from the urn, but conditioned on the event that exactly $w/2$ balls of each type are chosen. We obtain a different distribution on count vectors σ ; to construct an exchangeable pair, we generate another vector σ' by choosing one of the w selected balls uniformly at random and swapping it with a random ball in the urn of the same type. This induces a joint law on (σ, σ') under which these variables are again exchangeable, which we call the *conditioned distribution*, \mathbf{P}_c .

We again focus on $S = S(\sigma) = \sigma_{\pm} + \sigma_{=}$. Adopting the same notation as in the proof of Theorem 2 (note that now $\mathcal{K} = \{0\}$), we have

$$\mathbf{P}_0[S = w/2] = \psi = \mathbf{P}_0[E_{\mathcal{K}}], \quad \mathbf{P}_c[S = w/2]\psi = \phi(\beta) \quad (3.27)$$

We also have the following analogue of Lemma 3.1.

Lemma 3.6. *There is a function f satisfying the identity*

$$|\mu_c(w/2) - \mu_0(w/2)| = \left| \mathbb{E}_c \left[\left((\sigma_{=} - \sigma_{\pm})^2 - n(\sigma_{=} - \sigma_{\pm}) \left(\beta - \frac{1}{2} \right) \right) \Delta f(S) \right] \right| \quad (3.28)$$

as well as:

$$\sum_{s=0}^w |\Delta f(S)| = \mathcal{O}\left(\frac{1}{nw}\right) \quad (3.29)$$

The proofs of this and all succeeding technical lemmas are deferred to the next section. We turn to bounding (3.28). Just as in Theorem 2, the main technical difficulty of this theorem is to bound $\mathbb{E}_C[(\sigma_{=} - \sigma_{\pm})\mathbb{1}_{S=s}]$. The analogue of Proposition 3.1 is the following estimate.

Proposition 3.2. *Uniformly over s ,*

$$|\mathbb{E}_c[(\sigma_{=} - \sigma_{\pm})\mathbb{1}_{S=s}]| = \mathcal{O}\left(xw^{1/2}\right) e^{-\Theta\left(\frac{1}{w}\right)\left(\frac{w}{2}-s\right)^2}.$$

We need a similar result to control the remaining part of (3.28). The proof will follow trivially from the techniques developed in Proposition 3.2.

Proposition 3.3. *Uniformly over all $s \in [w]$,*

$$\mathbb{E}_c[(\sigma_{=} - \sigma_{\pm})^2\mathbb{1}_{S=s}] = \left(\mathcal{O}\left(x^2w^{3/2}\right) + \mathcal{O}\left(w^{1/2}\right) \right) e^{-\Theta\left(\frac{1}{w}\right)\left(\frac{w}{2}-s\right)^2}.$$

Let us first show that together these propositions imply the strong bound before proving them. Combining Proposition 3.2 and Eq. (3.29), we obtain

$$(xn)\mathbb{E}_C[(\sigma_{=} - \sigma_{\pm})\Delta f(S)] = \mathcal{O}\left(x^2w^{1/2}n\right) \sum_{s=0}^w e^{-\Theta\left(\frac{1}{w}\right)\left(\frac{w}{2}-s\right)^2} |\Delta f(s)| = \mathcal{O}\left(x^2w^{-1/2}\right),$$

Similarly, using Proposition 3.3 and recalling $n = \Omega(w)$ yields

$$\begin{aligned}\mathbb{E}_C[(\sigma_- - \sigma_\pm)^2 \Delta f(S)] &= \left(\mathcal{O}\left(x^2 w^{3/2}\right) + \mathcal{O}\left(w^{1/2}\right) \right) \sum_{s=0}^w e^{-\Theta\left(\frac{1}{w}\right)\left(\frac{w}{2}-s\right)^2} |\Delta f(s)| \\ &= \mathcal{O}\left(x^2 w^{-1/2}\right) + \mathcal{O}\left(n^{-1} w^{-1/2}\right)\end{aligned}$$

Applying these to Eq. (3.28), we obtain:

$$|\mu_0(w/2) - \mu_c(w/2)| \leq \mathcal{O}\left(x^2 w^{-1/2}\right) + \mathcal{O}\left(n^{-1} w^{-1/2}\right)$$

The desired strong bound (2.3) then follows from the identities (3.27). Indeed, under the unconditioned law \mathbf{P}_0 , the variables σ_\pm and σ_- are independent hypergeometric variables

$$\sigma_\pm \sim H\left(\frac{w}{2}; \frac{\gamma n}{2}, \frac{n}{2}\right), \quad \sigma_- \sim H\left(\frac{w}{2}; \frac{\beta n}{2}, \frac{n}{2}\right)$$

So, recalling $\psi = \mu_0(w/2) = \Theta(w^{-1/2})$,

$$\phi\left(\frac{1}{2} + x\right) = \psi \mu_c(w/2) = \left(1 + \mathcal{O}\left(\frac{1}{n}\right)\right) \psi^2 (1 + \mathcal{O}(x^2)).$$

This verifies the assumptions of Lemma 2.2 and proves the claim. We conclude by proving Proposition 3.2 and Proposition 3.3.

Proof of Proposition 3.2. We proceed identically to the proof of Proposition 3.1, except with the simplification of only considering $k = 0$.

Let $F(c) := \mathbf{P}_c[\sigma_- = c, S = s] = \mathbf{P}_c[\sigma_- = c, \sigma_\pm = s - c]$. We can write this explicitly as a product of hypergeometric densities:

$$F(c) = \binom{\beta n/2}{c} \binom{\gamma n/2}{w/2 - c} \binom{\beta n/2}{w/2 - (s - c)} \binom{\gamma n/2}{s - c} \binom{n/2}{w/2}^{-2}$$

Fix s with $0 \leq s \leq w$. We aim to show the inequality

$$\mathbb{E}_c[(\sigma_- - \sigma_\pm) \mathbb{1}_{S=s}] =: \sum_{c=0}^s (2c - s) F(c) \leq \mathcal{O}\left(x w^{1/2}\right) e^{-\Theta\left(\frac{1}{w}\right)\left(\frac{w}{2}-s\right)^2} \quad (3.30)$$

We begin by giving a crude tail bound on F . Using Lemma B.4 yields

Lemma 3.7. *For all c ,*

$$F(c) = \mathcal{O}\left(\frac{1}{w}\right) \exp\left\{-\Theta\left(\frac{(s - w/2)^2 + (c - \beta s)^2}{w}\right)\right\}$$

As in Proposition 3.1, we consider three cases for the size of $|x|$.

Case 1: $|x| \geq w^{-1/2}$. Just as in Proposition 3.1, Lemma 3.7 is already enough when $|x| \geq w^{-1/2}$. The calculation is identical, but with Lemma 3.7 in place of Lemma 3.3, so we omit the details.

Case 2: $|x| \leq w^{-1}$. For the remaining cases, we again exploit the symmetry $c \rightarrow s - c$ just as in Proposition 3.1.

$$\begin{aligned}\mathbb{E}_c[(\sigma_- - \sigma_\pm) \mathbb{1}_{S=s}] &= \frac{1}{2} \left[\sum_{c=0}^s (2c - s) F(c) + \sum_{c=0}^s (2c - s) F(s - c) \right] \\ &= \frac{1}{2} \sum_{c=0}^s (2c - s) [F(c) - F(s - c)]\end{aligned} \quad (3.31)$$

If $x = 0$, then (3.31) is identically zero, proving the claim. So assume that $x \neq 0$, which implies that $|xn|$ is a positive integer. If $x > 0$, then expanding the definition of F yields

$$\begin{aligned} \frac{F(c)}{F(s-c)} &= \frac{\binom{\beta n/2}{c} \binom{\gamma n/2}{w/2-c} \binom{\beta n/2}{w/2-(s-c)} \binom{\gamma n/2}{s-c}}{\binom{\gamma n/2}{c} \binom{\beta n/2}{w/2-c} \binom{\gamma n/2}{w/2-(s-c)} \binom{\beta n/2}{s-c}} \\ &= \prod_{r=-xn+1}^{xn} \left(\frac{\frac{n}{4} - \frac{w}{2} + c + r}{\frac{n}{4} - c + r} \right) \left(\frac{\frac{n}{4} - s + c + r}{\frac{n}{4} - (\frac{w}{2} - s + c) + r} \right) \\ &= \prod_{r=-xn+1}^{xn} \left(1 + \mathcal{O}\left(\frac{2c - w/2}{n}\right) \right) \left(1 + \mathcal{O}\left(\frac{w/2 + 2c - 2s}{n}\right) \right) \\ &= \left(1 + \mathcal{O}\left(\frac{2c - w/2}{n}\right) + \mathcal{O}\left(\frac{w/2 - 2(s-c)}{n}\right) \right)^{\Theta(xn)} \end{aligned}$$

If $x < 0$, then the numerator and denominator of the product in the second line are flipped, but the third line is reached unchanged. So, regardless of the sign of x , this computation holds. Then, recalling $|x| \leq w^{-1}$, we are justified in using a first-order Taylor expansion:

$$\frac{F(c)}{F(s-c)} = 1 + \mathcal{O}(x(2c - w/2)) + \mathcal{O}(x(w/2 - 2(s-c))) = 1 + \mathcal{O}(x^2 s) + \mathcal{O}(x(c - \beta s)) + \mathcal{O}(x(w/2 - s)),$$

where the second equality follows from the fact that

$$|2c - w/2| \leq |w/2 - s| + 2|c - \beta s| + 2|xs|,$$

and the analogous bound for $|w/2 - 2(s-c)|$. Since $|x| \leq w^{-1}$, we have $x^2 s = \mathcal{O}(x)$ and $|2c - s| \leq 2|c - \beta s| + 1$. Applying Lemmas 3.4 and 3.7, we obtain

$$\begin{aligned} \sum_{c=0}^s (2c - s)(F(c) - F(s-c)) &= \mathcal{O}\left(\frac{x}{w}\right) \sum_{c=0}^s (\mathcal{O}((c - \beta s)^2) + \mathcal{O}((w/2 - s)^2)) e^{-\Theta\left(\frac{(c - \beta s)^2 + (s - w/2)^2}{w}\right)} \\ &= \left(\mathcal{O}\left(xw^{1/2}\right) + \mathcal{O}\left(xw^{-1/2}(s - w/2)^2\right) \right) e^{-\Theta\left(\frac{1}{w}\right)(\frac{w}{2} - s)^2} \\ &= \mathcal{O}\left(xw^{1/2}\right) e^{-\Theta\left(\frac{1}{w}\right)(\frac{w}{2} - s)^2}, \end{aligned}$$

where the last step uses the fact that, for any positive constant C ,

$$y^2 e^{-Cy^2/w} = \mathcal{O}(w) e^{-Cy^2/2w}.$$

This finishes the proof of the second case.

Case 3: $w^{-1} \leq |x| \leq w^{-1/2}$. Again our plan is to derive a careful estimate for $F(c)$ near $c = \beta s$ via a Taylor expansion of Stirling's formula. Because of the tail estimate Lemma 3.7, as in the proof of Proposition 3.1 we may again restrict to the region where $|c - \beta s| = \mathcal{O}(\sqrt{s \log s})$ and $|s - w/2| = \mathcal{O}(\sqrt{s \log s})$.

Following precisely the steps of the proof of Case 3 in Proposition 3.1, using the tail estimate Lemma 3.7 in place of Lemma 3.3, we obtain that

$$\mathbb{E}_c[(\sigma_+ - \sigma_\pm) \mathbb{1}_{S=s}] = \sum_{|c - \beta s| = \mathcal{O}(\sqrt{w \log w})} (2c - 2\beta s)(F(c) - F(s + 2[xs] - c)) + \mathcal{O}\left(xw^{1/2}\right) e^{-\Theta\left(\frac{1}{w}\right)(\frac{w}{2} - s)^2}, \quad (3.32)$$

The only difference from Proposition 3.1 is that we are dealing with a product F of hypergeometric distributions rather than the product G of binomial distributions. However, F and G are conveniently related because they describe hypergeometric and binomial distributions respectively with the same mean and number of trials. Denote $G(c) := G(0, c)$. An explicit expansion of the hypergeometric density [35] yields:

$$F(c) = G(c) R_\beta(c) R_\gamma(s - c), \quad (3.33)$$

$$R_\beta(c) := \frac{\prod_{t=1}^{c-1} \left(1 - \frac{t}{\beta n/2}\right) \prod_{t=1}^{w/2-c-1} \left(1 - \frac{t}{\gamma n/2}\right)}{\prod_{t=1}^{w/2-1} \left(1 - \frac{t}{n/2}\right)}$$

Let us write $\tilde{s} = s + 2[xs]$. Since $|G(c) - G(\tilde{s} - c)|$ was already bounded in Proposition 3.1, it suffices to bound $R_\beta(c)R_\gamma(s - c) - R_\beta(\tilde{s} - c)R_\gamma(s - \tilde{s} + c)$.

We temporarily define an unorthodox convention for the product that will save us much case work. If $a < b$, then define $\prod_{i=a}^b x_i = x_a x_{a+1} \dots x_b$ as usual. However, if $a > b$, we define $\prod_{i=a}^b x_i = \prod_{i=b+1}^{a-1} x_i^{-1}$. Then, expanding the definition of R ,

$$\begin{aligned} \frac{R_\beta(c)}{R_\gamma(s - \tilde{s} + c)} &= \frac{\prod_{t=1}^{c-1} \left(1 - \frac{t}{\beta n/2}\right) \prod_{t=1}^{w/2-c-1} \left(1 - \frac{t}{\gamma n/2}\right)}{\prod_{t=1}^{s-\tilde{s}+c-1} \left(1 - \frac{t}{\gamma n/2}\right) \prod_{t=1}^{w/2-s+\tilde{s}-c-1} \left(1 - \frac{t}{\beta n/2}\right)} \\ &= \prod_{t=1}^{s-\tilde{s}+c-1} \left(\frac{1 - \frac{t}{\beta n/2}}{1 - \frac{t}{\gamma n/2}}\right) \prod_{t=s-\tilde{s}+c}^{c-1} \left(1 - \frac{t}{\beta n/2}\right) \prod_{t=1}^{w/2-c-1} \left(\frac{1 - \frac{t}{\gamma n/2}}{1 - \frac{t}{\beta n/2}}\right) \prod_{t=w/2-c}^{w/2-s+\tilde{s}-c-1} \left(1 - \frac{t}{\gamma n/2}\right) \end{aligned}$$

A similar expansion yields:

$$\frac{R_\gamma(s - c)}{R_\beta(\tilde{s} - c)} = \prod_{t=1}^{\tilde{s}-c-1} \left(\frac{1 - \frac{t}{\gamma n/2}}{1 - \frac{t}{\beta n/2}}\right) \prod_{t=\tilde{s}-c}^{s-c-1} \left(1 - \frac{t}{\gamma n/2}\right) \prod_{t=1}^{w/2-\tilde{s}+c-1} \left(\frac{1 - \frac{t}{\beta n/2}}{1 - \frac{t}{\gamma n/2}}\right) \prod_{t=w/2-\tilde{s}+c}^{w/2-s+c-1} \left(1 - \frac{t}{\beta n/2}\right)$$

In total, our goal is to bound how far the following expression is from one.

$$\begin{aligned} \frac{R_\beta(c)R_\gamma(s - c)}{R_\gamma(s - \tilde{s} + c)R_\beta(\tilde{s} - c)} &= \prod_{t=\tilde{s}-c}^{s-\tilde{s}+c-1} \left(\frac{1 - \frac{t}{\beta n/2}}{1 - \frac{t}{\gamma n/2}}\right) \prod_{t=w/2-\tilde{s}+c}^{w/2-c-1} \left(\frac{1 - \frac{t}{\gamma n/2}}{1 - \frac{t}{\beta n/2}}\right) \prod_{t=s-\tilde{s}+c}^{c-1} \left(1 - \frac{t}{\beta n/2}\right) \\ &\quad \times \prod_{t=w/2-c}^{w/2-s+\tilde{s}-c-1} \left(1 - \frac{t}{\gamma n/2}\right) \prod_{t=\tilde{s}-c}^{s-c-1} \left(1 - \frac{t}{\gamma n/2}\right) \prod_{t=w/2-\tilde{s}+c}^{w/2-s+c-1} \left(1 - \frac{t}{\beta n/2}\right) \end{aligned} \quad (3.34)$$

We bound this expression in several parts. Throughout, we will repeatedly use the arithmetic fact that $(xw)^2 + |xw(c - \beta s)| = \mathcal{O}((xw)^2 + (c - \beta s)^2)$. Observe that $2c - \tilde{s} + s - \tilde{s} = \mathcal{O}(xs) + \mathcal{O}(c - \beta s)$. Noting $xt/n = \mathcal{O}(w^{1/2}/n) = o(1)$, we may apply a first-order Taylor expansion:

$$\begin{aligned} \prod_{t=\tilde{s}-c}^{s-\tilde{s}+c-1} \left(\frac{1 - \frac{t}{\beta n/2}}{1 - \frac{t}{\gamma n/2}}\right) \prod_{t=w/2-\tilde{s}+c}^{w/2-c-1} \left(\frac{1 - \frac{t}{\gamma n/2}}{1 - \frac{t}{\beta n/2}}\right) &= \prod_{t=\tilde{s}-c}^{s-\tilde{s}+c-1} \left(1 + \mathcal{O}\left(\frac{xt}{n}\right)\right) \prod_{t=w/2-\tilde{s}+c}^{w/2-c-1} \left(1 + \mathcal{O}\left(\frac{xt}{n}\right)\right) \\ &= \mathcal{O}\left(\exp\left\{\mathcal{O}\left(\frac{xw(|xs| + |c - \beta s|)}{n}\right)\right\}\right) \\ &= 1 + \mathcal{O}\left(\frac{x^2 w^2 + (c - \beta s)^2}{n}\right) \exp\left\{\frac{xw(c - \beta s)}{n}\right\} \end{aligned} \quad (3.35)$$

In the last line, we used that $x^2 w^2/n = \mathcal{O}(1)$ as well as the inequality $e^y \leq 1 + ye^y$ for all $y \in \mathbb{R}$. Next, by our product notation convention,

$$\begin{aligned} \prod_{t=s-\tilde{s}+c}^{c-1} \left(1 - \frac{t}{\beta n/2}\right) \prod_{t=\tilde{s}-c}^{s-c-1} \left(1 - \frac{t}{\gamma n/2}\right) &= \prod_{t=s-\tilde{s}+c}^{c-1} \left(\frac{1 - \frac{t}{\beta n/2}}{1 - \frac{t+2c-\tilde{s}}{\gamma n/2}}\right) \\ &= \prod_{t=s-c}^{c-1} \left(1 + \mathcal{O}\left(\frac{xt}{n}\right) + \mathcal{O}\left(\frac{2c - \tilde{s}}{n}\right)\right) \end{aligned}$$

Both $xt/n = o(1)$ and $(2c - s)/n = o(1)$, so we may apply a first-order Taylor expansion again. Since $|s - \tilde{s}| = \mathcal{O}(xw)$,

$$\begin{aligned} \prod_{t=s-\tilde{s}+c}^{c-1} \left(1 - \frac{t}{\beta n/2}\right) \prod_{t=\tilde{s}-c}^{s-c-1} \left(1 - \frac{t}{\gamma n/2}\right) &= 1 + \mathcal{O}\left(\frac{xw(2c - s)}{n}\right) + \mathcal{O}\left(\frac{(2c - s)^2}{n}\right) \\ &= 1 + \mathcal{O}\left(\frac{x^2 w^2 + (c - \beta s)^2}{n}\right) \end{aligned} \quad (3.36)$$

The last line follows from observing $2c - s = 2(c - \beta s + xs)$, and hence $2c - s = \mathcal{O}(c - \beta s) + \mathcal{O}(xw)$. An identical computation also yields control over the remaining products in (3.34).

$$\prod_{t=w/2-c}^{w/2-s+\tilde{s}-c-1} \left(1 - \frac{t}{\gamma n/2}\right) \prod_{t=w/2-\tilde{s}+c}^{w/2-s+c-1} \left(1 - \frac{t}{\beta n/2}\right) = 1 + \mathcal{O}\left(\frac{x^2 w^2 + (c - \beta s)^2}{n}\right) \quad (3.37)$$

Hence the contribution of (3.35) dominates those of (3.36) and (3.37). Returning to (3.34),

$$\frac{R_\beta(c)R_\gamma(\tilde{s} - c)}{R_\beta(s - c)R_\gamma(c + s - \tilde{s})} = 1 + \mathcal{O}\left(\frac{x^2 w^2 + (c - \beta s)^2}{n}\right) \exp\left\{\frac{xw(c - \beta s)}{n}\right\} \quad (3.38)$$

This bound is sufficient to control R . Next, recall that by Lemma 3.5 from Proposition 3.1,

$$|G(c) - G(\tilde{s} - c)| \leq G(c)E(j, l),$$

where we have set $l = s - w/2$ and $j = c - \beta s$. Returning to (3.33) with this fact and (3.38), and then using the tail bound on F given in Lemma 3.7, we have:

$$\begin{aligned} |F(c) - F(\tilde{s} - c)| &\leq F(c) \left(E(j, l) + \mathcal{O}\left(\frac{x^2 w^2 + j^2}{n}\right) e^{\frac{xw j}{n}} \right) \\ &= \mathcal{O}\left(\frac{1}{w}\right) \exp\left\{-\Theta\left(\frac{j^2 + l^2}{w}\right)\right\} \left(E(j, l) + \mathcal{O}\left(\frac{x^2 w^2 + j^2}{n}\right) e^{\frac{xw j}{n}} \right) \end{aligned}$$

Apart from the $\frac{x^2 w^2 + j^2}{n} e^{\frac{xw j}{n}}$ term, this is same bound obtained as equation (3.16) in Proposition 3.1. From here, the proof in Proposition 3.1 may be followed exactly to see:

$$\begin{aligned} \sum_c (2c - s) |F(c) - F(s - c)| &= \mathcal{O}(xw^{1/2}) e^{-\Theta\left(\frac{l^2}{w}\right)} \\ &\quad + \sum_c \mathcal{O}\left(\frac{(|xs| + |j|)(x^2 w^2 + j^2)}{nw}\right) e^{\mathcal{O}\left(\frac{xw j}{n}\right)} e^{-\Theta\left(\frac{j^2 + l^2}{w}\right)} \end{aligned} \quad (3.39)$$

Let us simplify the exponents. Recalling $x = \mathcal{O}(w^{-1/2})$ and $n = \Omega(w)$, we have $xwj/n = \mathcal{O}(j/\sqrt{w})$. Fixing $b > 0$ large enough, if $j \geq b\sqrt{w}$, then $e^{\mathcal{O}\left(\frac{xw j}{n}\right)} e^{-\Theta\left(\frac{j^2 + l^2}{w}\right)} = e^{-\Theta\left(\frac{j^2 + l^2}{w}\right)}$. And, if $j \leq b\sqrt{w}$, then $e^{\mathcal{O}\left(\frac{xw j}{n}\right)} = \mathcal{O}(1)$. So, in total, we have by Lemma 3.4,

$$\begin{aligned} \sum_c \mathcal{O}\left(\frac{(|xs| + |j|)(x^2 w^2 + |xwj|)}{nw}\right) e^{\mathcal{O}\left(\frac{xw j}{n}\right)} e^{-\Theta\left(\frac{j^2 + l^2}{w}\right)} &= \sum_c \mathcal{O}\left(\frac{(|xs| + |j|)(x^2 w^2 + |xwj|)}{nw}\right) e^{-\Theta\left(\frac{j^2 + l^2}{w}\right)} \\ &= \mathcal{O}(xw^{1/2}) e^{-\Theta\left(\frac{l^2}{w}\right)} \end{aligned}$$

Combining this with (3.39) and (3.32), we are done. This completes the proposition for all possible x . \square

Proof of Proposition 3.3. Our goal is to bound:

$$\mathbb{E}_c[(\sigma_\pm - \sigma_=)^2 \mathbb{1}_{S=s}] = \sum_{c=0}^s (2c - s)^2 F(c) = 4 \sum_{c=0}^s (c - \beta s + xs)^2 F(c) = \mathcal{O}\left(\sum_{c=0}^s (x^2 s^2 + (c - \beta s)^2) F(c)\right)$$

We can directly apply the crude tailbound Lemma 3.7 on F to conclude. By Lemma 3.4,

$$\begin{aligned} \mathbb{E}_c[(\sigma_\pm - \sigma_=)^2 \mathbb{1}_{S=s}] &\leq \mathcal{O}\left(\frac{1}{w}\right) e^{-\Theta\left(\frac{1}{w}\right)(\frac{w}{2} - s)^2} \sum_{c=0}^s (x^2 s^2 + (c - \beta s)^2) e^{-\Theta\left(\frac{(c - \beta s)^2}{w}\right)} \\ &= \left(\mathcal{O}\left(x^2 w^{3/2}\right) + \mathcal{O}\left(w^{1/2}\right)\right) e^{-\Theta\left(\frac{1}{w}\right)(\frac{w}{2} - s)^2}. \end{aligned}$$

\square

3.5. Proofs of lemmas.

Proof of Lemma 3.6. The proof is almost identical to that of Lemma 3.2. Let $S := S(\sigma)$ and $S' := S(\sigma')$. A simple computation yields:

$$\begin{aligned}\mathbf{P}_0[S' > S \mid S] &= \frac{1}{w(n-w)} \left(\frac{wn}{2} - S \left(w + \frac{n}{2} \right) + S^2 \right) =: \frac{a_S}{w(n-w)} \\ \mathbf{P}_0[S' < S \mid S] &= \frac{1}{w(n-w)} \left(S \left(\frac{n}{2} - w \right) + S^2 \right) =: \frac{b_S}{w(n-w)}\end{aligned}\tag{3.40}$$

Similarly,

$$\begin{aligned}\mathbf{P}_c[S > S' \mid \sigma_{\pm}, \sigma_=] &= \frac{2}{w(n-w)} \left((\sigma_{\pm})^2 + (\sigma_=)^2 - (\sigma_{\pm} + \sigma_=) \frac{w}{2} + \sigma_= \left(\frac{n}{2} - r \right) + \sigma_{\pm} r \right) \\ \mathbf{P}_c[S < S' \mid \sigma_{\pm}, \sigma_=] &= \frac{2}{w(n-w)} \left((\sigma_{\pm})^2 + (\sigma_=)^2 + \frac{wn}{4} - (\sigma_{\pm} + \sigma_=) \frac{w}{2} - \sigma_{\pm} \left(\frac{n}{2} - r \right) - \sigma_= r \right)\end{aligned}$$

Recall the definitions of the skew-symmetric operator Ξ and T_0 given in (3.3). Then

$$T_0 f(S) = a_S f(S+1) - b_S f(S)\tag{3.41}$$

Recall $\Delta f(S)$ is the one-step forward difference operator. For any function f that depends only on S ,

$$T_c f(\sigma) - T_0 f(\sigma) = \left[(\sigma_{\pm} - \sigma_=)^2 - n(\sigma_= - \sigma_{\pm}) \left(\beta - \frac{1}{2} \right) \right] \Delta f(S)\tag{3.42}$$

We aim to find an f such that $T_0 f = \mathbb{1}(S = w/2) - \mu_0(w/2)$. Since $S(\sigma)$ under μ_0 is a hypergeometric random variable with w trials, success population $n/2$ and total population n , it is easy to check that taking $\mu := \mu_0(S)$, a_S and b_S as given in (3.40), and T as in (3.41) satisfies (3.4). Thus, we may apply Lemma 3.1 with $t := w/2$ to obtain such an f . Taking the expectation of both sides of (3.42) with this choice of f yields the first desired claim, (3.28).

Next, we have by definition of a and b , as well as the assumption that $n - w = \Theta(n)$,

$$\Delta f(w/2) \leq \min(a_{w/2}^{-1}, b_{w/2}^{-1}) = \frac{4}{(n-w)w} = \mathcal{O}\left(\frac{1}{nw}\right)$$

The second desired claim, (3.29), then directly follows from the third and fourth parts of Lemma 3.1:

$$\sum_s |\Delta f(w/2)| = \mathcal{O}\left(\frac{1}{nw}\right)$$

□

Proof of Lemma 3.7. Except for two changes, the proof is the same as that of Lemma 3.3. First, we now have the simplification $k = 0$. Second, we use Lemma B.4 in place of Lemma B.3. □

APPENDIX A. OMITTED PROOFS

A.1. Proof of Lemma 1.1. Let μ be either the Poisson or Binomial distribution on \mathbb{Z} . We will construct a variable $X \sim \mu$ as well as another random variable X' satisfying $|X - X'| \leq 1$ almost surely, such that X' has the same distribution as X conditioned on the event that X is even. If we can construct such a coupling (X, X') , then we can couple a matrix A (from either the Bernoulli or Poisson ensemble) with A' from the same ensemble conditioned on P (the event of even row parities), such that each row of A and A' differ in at most one entry.

To construct this coupling, consider a matrix A drawn from either the Poisson or Bernoulli ensemble and write X_i , $i = 1, \dots, m$ for its row sums. For each i , use the above coupling to construct an X'_i with $|X_i - X'_i| \leq 1$ and X'_i even. For A drawn from the Bernoulli ensemble, copy the rows of A to A' with the following modification: if $X'_i = X_i - 1$, flip a one to a zero uniformly at random from row i . If $X'_i = X_i + 1$, then flip a zero to a one uniformly at random. For A drawn from the Poisson ensemble: if $X'_i = X_i - 1$, decrement a non-zero entry of the row i , with entry A_{ij} picked with probability proportional to A_{ij} . If $X'_i = X_i + 1$, increment an entry of row i uniformly at random.

It remains to construct the desired coupling (X, X') . We use a construction due to Pinelis [34]. Denote by μ' the distribution induced by conditioning μ on being even, and define the shorthand $\mu_i := \mu(i)$, $\mu'_i := \mu'(i)$. The coupling is defined as follows:

$$\begin{aligned}\mathbb{P}[X = 2j, X' = 2j] &= \mu_{2j} \\ \mathbb{P}[X = 2j - 1, X' = 2j] &= t_{2j} \\ \mathbb{P}[X = 2j - 1, X' = 2(j - 1)] &= \mu_{2j-1} - t_{2j}\end{aligned}$$

We claim that we can define t to satisfy the following conditions for all $j \in \mathbb{Z} \cap [-1, \infty]$:

$$\mu'_{2j} = \mu_{2j} + t_{2j} + \mu_{2j+1} - t_{2j+2} \quad (\text{A.1})$$

$$0 \leq t_{2j} \leq \mu_{2j-1} \quad (\text{A.2})$$

If such a t exists, then (X, X') is clearly the desired coupling, so let us construct it. First, let Q be the probability that $X \sim \mu$ is even. This probability is positive if μ is a Poisson distribution. In the Binomial case, since we have assumed that $p \leq 1/2$ in the definition of the Bernoulli ensemble, we also have $Q > 0$. If $Q = 1$, then X is even almost surely and there is nothing to show, so we assume in what follows that $Q \in (0, 1)$.

Let Q be the probability that $X \sim \mu$ is even; then $\mu_{2j}/Q = \mu'_{2j}$. Since both the binomial and Poisson distribution are supported on nonnegative integers, we will set $t_j = 0$ for all $j \leq -2$. Now we construct t to exactly satisfy (A.1).

$$t_{2j+2} = t_{2j} + \mu_{2j} \left(1 - \frac{1}{Q}\right) + \mu_{2j+1}$$

It remains to check that this construction satisfies (A.2). We check the two inequalities individually. First,

$$t_{2j} = \sum_{i=-2}^{j-1} t_{2(i+1)} - t_{2i} = \sum_{i=-1}^{j-1} \mu_{2j} \left(1 - \frac{1}{Q}\right) + \mu_{2j+1}$$

From the expression we have just derived for t , it is clear that $t_{2j} \geq t_{2(j-1)}$ if and only if $\frac{\mu_{2j-1}}{\mu_{2j-2}} > \frac{1}{Q} - 1$. For both the Poisson and Binomial cases, the sequence μ is log-concave:

$$\mu_j^2 \geq \mu_{j-1} \mu_{j+1} \quad \forall j \in \mathbb{Z},$$

which implies that the sequence of ratios $\frac{\mu_{2j-1}}{\mu_{2j-2}}$ is decreasing. Therefore, the sequence t_{2j} changes from increasing to decreasing at most once. Since $t_{-1} = 0 = t_\infty$ and $t_0 > 0$, this implies that t can never be negative. Similarly, we can write a telescoping sum to check the other condition:

$$t_{2j} - \mu_{2j-1} = \sum_{i=-1}^{j-1} \mu_{2j} \left(1 - \frac{1}{Q}\right) + \mu_{2j-1}$$

By identical reasoning, this expression is always non-negative, so indeed (A.2) is satisfied. This completes the lemma. \square

A.2. Laplace's Method: Proof of Lemma 2.2. By assumption, the columns of $A \sim \mathcal{M}$ are exchangeable, so ψ_i and ϕ_i do not depend on particular choices of u and v . We also assumed the rows of A are independent, so the expectation of Z is given by

$$\mathbb{E}[Z]^2 = \left(\sum_{u \in \mathcal{B}} \mathbb{P} \left[u \in \bigcap_{i=1}^m G_i \right] \right)^2 = \binom{n}{n/2}^2 \prod_{i=1}^m \psi_i^2 \quad (\text{A.3})$$

Turning to the second moment, the number of pairs of balanced vectors that are equal in $2r$ indices and different in the remaining $n/2 - 2r$ indices is:

$$\binom{n}{n/2} \binom{n/2}{r}^2$$

The first coefficient gives the number of ways to pick a balanced vector u ; then, looking at the $n/2$ coordinates each where u is $+1$ and -1 we must pick r coordinates from each for v to agree with u and have v differ on the remaining coordinates.

We will break the second moment calculation into three regimes. Consider $r \in [n]$ and define $\beta = 2r/n$. Let $\delta \in (0, 1/2)$ be small enough that $2H(2\delta) \leq c$, where c is the constant appearing in the first-moment bound (2.1). Let $I_1 := \{r : \beta \in (1/2 - \epsilon, 1/2 + \epsilon)\}$ be the central region, $I_2 := \{r : \beta \in [\delta, 1 - \delta] \setminus I_1\}$ be an annulus, and $I_3 := [n] \setminus (I_1 \cup I_2)$ be the remainder. Then:

$$\mathbb{E}[Z^2] = \binom{n}{n/2} \sum_{r=0}^{n/2} \binom{n/2}{r}^2 \prod_{i=1}^m \phi_i\left(\frac{2r}{n}\right) = \binom{n}{n/2} \sum_{r \in I_1 \cup I_2 \cup I_3} \binom{n/2}{r}^2 \prod_{i=1}^m \phi_i\left(\frac{2r}{n}\right)$$

We begin by showing the annular region I_2 is negligible. Recall $\beta := 2r/n$ and let $x := \beta - 1/2$. Using the weak bound (2.2) and standard tail bounds on binomial coefficients (Lemma B.3) yields

$$\begin{aligned} \binom{n}{n/2} \sum_{I_2} \binom{n/2}{\beta n/2}^2 \prod_{i=1}^m \phi_i(\beta) &\leq \binom{n}{n/2} \sum_{I_2} \binom{n/2}{n/4}^2 e^{-\Theta(x^2 n)} C_\delta^m \prod_{i=1}^m \psi_i^2 \\ &\leq \mathbb{E}[Z]^2 \exp\{-\Theta(n) + \Theta(m)\} \\ &= \mathbb{E}[Z]^2 \exp\{-\Theta(n)\} \end{aligned} \tag{A.4}$$

The second line follows from the fact that β is bounded away from $1/2$ on I_2 , and the third line uses $m = o(n)$. Now, we show the outermost region is also negligible. Using the trivial bound that $\phi_i(\beta) \leq \psi_i$ for all β ,

$$\begin{aligned} \binom{n}{n/2} \sum_{I_3} \binom{n/2}{r}^2 \prod_{i=1}^m \phi_i\left(\frac{2r}{n}\right) &\leq n \mathbb{E}[Z] \binom{n/2}{\delta n}^2 \\ &\leq \mathbb{E}[Z] e^{(1+o(1))nH(2\delta)} \\ &\leq \mathbb{E}[Z]^2 e^{-(1+o(1))nH(2\delta)}, \end{aligned} \tag{A.5}$$

where the final inequality uses the assumption that $\mathbb{E}[Z] \geq e^{cn} \geq e^{2nH(2\delta)}$. Hence the contribution of I_3 is also exponentially small.

So, I_1 is the only region with meaningful contribution. The sum over I_1 can be understood through Laplace's method: away from $r = n/4$, the term $\binom{n/2}{r}$ decays exponentially in n and $\prod_{i=1}^m \phi_i$ grows at most exponentially in m . In order for the central term when $r = n/4$ to agree with $\mathbb{E}[Z]^2$ to first order, we therefore need that $\phi_i(1/2) = (1 + o(m^{-1}))\psi_i^2$, and to ensure that the central term is the dominant one we need to show that $\phi_i(1/2 + x)$ grows at most quadratically in a window of constant radius around $x = 0$. This is the purpose of the inequality (2.3). Using (A.4) and (A.5) to ignore I_2 and I_3 , and then applying (2.3) to bound ϕ in I_1 ,

$$\begin{aligned} \mathbb{E}[Z^2] &\leq (1 + o(1)) \binom{n}{n/2} \sum_{r \in I_1} \binom{n/2}{r}^2 \prod_{i=1}^m \left(1 + o\left(\frac{1}{m}\right)\right) \left(1 + C\left(\frac{2r}{n} - \frac{1}{2}\right)^2\right) \psi_i^2 \\ &\leq (1 + o(1)) \binom{n}{n/2}^{-1} \mathbb{E}[Z]^2 \sum_{I_1} \binom{n/2}{r}^2 \exp\left(mC\left(\frac{2r}{n} - \frac{1}{2}\right)^2\right) \end{aligned} \tag{A.6}$$

Recall $m = o(n)$ and $I_1 = \{r : |r - n/4| < \epsilon n\}$. Using Stirling's approximation (Lemma B.2), we obtain

$$\begin{aligned} \sum_{|r-n/4| \leq \epsilon n} \binom{n/2}{r}^2 e^{mC\left(\frac{2r}{n} - \frac{1}{2}\right)^2} &= (1 + o(1)) \sum_{|r-n/4| \leq \sqrt{n \log n}} \binom{n/2}{r}^2 e^{mC\left(\frac{2r}{n} - \frac{1}{2}\right)^2} \\ &= (1 + o(1)) \frac{2^{n+2}}{n\pi} \sum_{|r-n/4| \leq \sqrt{n \log n}} e^{-\frac{8(r-n/4)^2}{n} + mC\left(\frac{2r}{n} - \frac{1}{2}\right)^2} \\ &= (1 + o(1)) \frac{2^{n+2}}{n\pi} \sum_{j \in \mathbb{Z}} e^{-\frac{8j^2}{n}(1+o(\frac{m}{n}))} \\ &= \left(1 + o(1) + o\left(\frac{m}{n}\right)\right) \sqrt{\frac{2}{\pi n}} 2^n \sum_{j \in \mathbb{Z}} e^{-\frac{n\pi^2 j^2}{8}(1+o(\frac{m}{n}))}, \end{aligned}$$

where in the last line we have used the functional equation for the Jacobi theta function [43, Theorem 3.2]. Since $\sqrt{\frac{2}{\pi n}} 2^n = (1 + o(1)) \binom{n}{n/2}$ and $\sum_{j \in \mathbb{Z}} e^{-\frac{n\pi^2 j^2}{8}(1+O(\frac{m}{n}))} = 1 + o(1)$, we have shown

$$\sum_{|r-n/4| \leq \epsilon n} \binom{n/2}{r}^2 e^{mC(\frac{2r}{n}-\frac{1}{2})^2} = (1 + o(1)) \binom{n}{n/2}.$$

Returning to (A.6) and applying this bound for I_1 , the lemma is established: $\mathbb{E}[Z^2] \leq (1 + o(1))\mathbb{E}[Z]^2$. \square

A.3. Proof of Lemma 3.1. Denote by U_s the integers $\{0, \dots, s\}$ and use the shorthand $\mu_i := \mu(\{i\})$. Define f by

$$f(0) = 0, \quad f(s) := \frac{\mu_t}{b_s \mu_s} (\mathbb{1}_{t < s} - \mu(U_{s-1})) \quad \forall s \in \{1, \dots, w\}. \quad (\text{A.7})$$

Then, since $\mu_{s+1} = \mu_s a_s / b_{s+1}$, we have for all $s \in [w]$,

$$\begin{aligned} a_s f(s+1) - b_s f(s) &= \mu_t \left(\frac{a_s}{b_{s+1} \mu_{s+1}} \mathbb{1}_{t < s+1} - \frac{1}{\mu_s} \mathbb{1}_{t < s} - \frac{a_s}{b_{s+1} \mu_{s+1}} \mu(U_s) + \frac{1}{\mu_s} \mu(U_{s-1}) \right) \\ &= \frac{\mu_t}{\mu_s} [(\mathbb{1}_{t < s+1} - \mathbb{1}_{t < s}) - (\mu(U_s) - \mu(U_{s-1}))] \\ &= \mathbb{1}_{s=t} - \mu_t \end{aligned}$$

This establishes claim (a). Now we prove (b). Note that $a_s, b_s, \mu_s \geq 0$ for all s . So, if $s \leq t$, then $f(s) = -\frac{\mu_t}{b_s \mu_s} \mu(U_{s-1}) \leq 0$. And, if $s \geq t+1$, then $f(s) = \frac{\mu_t}{b_s \mu_s} (1 - \mu(U_{s-1})) \geq 0$ since $\mu \leq 1$. This establishes the sign of f . Next is monotonicity. Since $t \geq 1$, we have $f(1) \leq 0 = f(0)$, so f is non-increasing from $f(0)$ to $f(1)$. Now, say $2 \leq s \leq t$. We have:

$$f(s-1) - f(s) = \mu_t \left(\frac{\mu(U_{s-1})}{b_s \mu_s} - \frac{\mu(U_{s-2})}{b_{s-1} \mu_{s-1}} \right) \geq \frac{\mu_t}{b_s \mu_s} \left(\sum_{i=1}^{s-1} \mu_i \left(1 - \frac{b_s \mu_s \mu_{i-1}}{b_{s-1} \mu_{s-1} \mu_i} \right) \right) \geq 0$$

The last inequality follows from $i \leq s-1 < t$ and the fact that a is decreasing and b is increasing:

$$\frac{b_s \mu_s \mu_{i-1}}{b_{s-1} \mu_{s-1} \mu_i} = \frac{a_{s-1} b_i}{a_{i-1} b_{s-1}} \leq 1$$

The argument for $s \geq t+1$ is almost identical and yields the reverse inequality. Now, we prove (c). By (b), we already know $\sup |\Delta f(s)| \leq \Delta f(t)$. So, we just need to show that $\Delta f(t) \leq \min(a_t^{-1}, b_t^{-1})$.

$$\Delta f(t) = a_t^{-1} \sum_{i=t+1}^w \mu_i + b_t^{-1} \sum_{i=0}^{t-1} \mu_i = a_t^{-1} \left(\sum_{i=t+1}^w \mu_i + \sum_{i=0}^{t-1} \frac{\mu_{i+1} b_{i+1} a_t}{a_i b_t} \right) \leq a_t^{-1} \left(\sum_{i=1}^w \mu_i \right)$$

The last inequality again uses the monotonicity of a and b and the fact μ sums to 1. The proof that $\Delta f(t) \leq b_t^{-1}$ is essentially identical so we omit it.

Finally, we turn to claim (d) of the lemma. We need three facts that we have already established: $\Delta f(s) < 0$ for all $s \neq t$; $f(s) \leq 0$ for $s \leq t$; and $f(s) \geq 0$ for $s > t$. Combining the first two facts, $\sum_{s \leq t-1} |\Delta f(s)|$ telescopes and is bounded above by $|f(t)|$. Similarly, combining the second two facts, $\sum_{s \geq t+1} |\Delta f(s)| \leq |f(t+1)|$. Since $|f(t+1)| + |f(t)| = f(t+1) - f(t) = |\Delta f(t)|$,

$$\begin{aligned} \sum_s |\Delta f(s)| &= \sum_{s \leq t-1} |\Delta f(s)| + |\Delta f(t)| + \sum_{t+1 \leq s} |\Delta f(s)| \leq |f(t)| + |\Delta f(t)| + |f(t+1)| \\ &= 2|\Delta f(t)|. \end{aligned}$$

\square

APPENDIX B. LOCAL LIMIT THEOREMS

We collect several approximations which we use throughout.

Lemma B.1 (De Moivre–Laplace [22, Theorem VII.3.1]). *Let $p \in [\delta, 1 - \delta]$ for a constant $\delta \in (0, 1/2)$, and write $q = 1 - p$. Let $k = r - np$. For any $r \in \{0, \dots, n\}$, it holds*

$$\binom{n}{r} p^r (1-p)^{n-r} = \left(1 + \mathcal{O}\left(\frac{1}{n}\right)\right) \frac{1}{\sqrt{2\pi np(1-p)}} e^{-\frac{k^2}{2np(1-p)} + \mathcal{O}\left(\frac{k}{n}\right) + \mathcal{O}\left(\frac{k^3}{n^2}\right)}.$$

Specializing to $p = 1/2$ yields the following simplified bound.

Lemma B.2 (Stirling’s approximation for Binomial coefficients [41, Equation (5.43)]). *If $|r - n/2| = o(n^{-2/3})$, then*

$$\binom{n}{r} = (1 + o(1)) \sqrt{\frac{2}{\pi n}} 2^n e^{-2(r-n/2)^2/n}.$$

We also have a coarser estimate valid for all r .

Lemma B.3 (Gaussian tails for Binomial). *In the same setting as Lemma B.1, for any $r \in \{0, \dots, n\}$,*

$$\binom{n}{r} p^r (1-p)^{n-r} = \mathcal{O}\left(\frac{1}{\sqrt{n}}\right) e^{-\frac{k^2}{4np(1-p)}}.$$

Proof. When $k = o(n^{2/3})$, this bound follows from Lemma B.1. When $k = \Omega(n^{2/3})$, we employ Hoeffding’s bound [27]; letting $X \sim \text{Bin}(n, p)$,

$$\binom{n}{r} p^r (1-p)^{n-r} \leq \mathbb{P}[|X - \mathbb{E}[X]| \geq k] \leq 2e^{-\frac{k^2}{2np(1-p)}} = \mathcal{O}\left(\frac{1}{\sqrt{n}}\right) e^{-\frac{k^2}{4np(1-p)}},$$

where the last step uses $e^{-k^2/n} = \mathcal{O}\left(\frac{1}{\sqrt{n}}\right)$ if $k = \Omega(n^{2/3})$. \square

Lemma B.4 (Gaussian tail for Hypergeometric distribution [25, Theorem 2, (i)]). *Consider a hypergeometric random variable $X \sim H(w; pn, n)$ with w trials, n total population, and np population successes. There is some universal constant K such that for any $\lambda > 0$,*

$$\mathbb{P}[\sqrt{w}|X - pw| = \lambda] \leq \frac{K}{\sqrt{n}} \exp\left\{-\frac{2\lambda^2}{1 - w/n} - \left(\frac{1}{4} + \frac{1}{3} \left(\frac{w}{n-w}\right)^3\right) \frac{\lambda^4}{n}\right\}$$

In particular, if $n - w = \Theta(n)$,

$$\binom{np}{k} \binom{n(1-p)}{w-k} \binom{n}{w}^{-1} = \mathcal{O}\left(\frac{1}{\sqrt{w}}\right) e^{-\Theta(\frac{1}{w})(k-wp)^2}$$

Lemma B.5 (Edgeworth Series for lattice sums [33, Theorem 2]). *Let $\{X_i\}_{i \in \mathbb{N}}$ be independent identically distributed random variables with $X_1 \in \{-1, 0, 1\}$, $\mathbb{E}[X_1] = 0$, and $\mathbb{E}[X_1^2] = \sigma^2$. Denote by κ_r the r -th cumulant of X_1 , and denote by ϕ the density of standard unit Gaussian. Define the quantities*

$$P_n(N) = \mathbb{P}\left[\sum_{j=1}^n X_j = N\right], \quad x = \frac{N}{\sigma\sqrt{n}}$$

Then, there exists a collection of polynomials $\{q_{3\nu}(x)\}_{\nu \in \mathbb{N}}$, each of degree 3ν with coefficients depending only on the moments of X_1 up to order $\nu + 2$ (inclusive), satisfying

$$\left| \sigma\sqrt{n}P_n(N) - \phi(x) - \sum_{\nu=1}^{k-2} \frac{q_{3\nu}(x)\phi(x)}{n^{\nu/2}} \right| = \frac{o(1)}{(1 + |x|^k)n^{(k-2)/2}}$$

In particular, $q_{3\nu} = 0$ for any odd ν , and the first non-zero q is q_6 given by:

$$q_6(x)\phi(x) = -\left(\frac{1}{24}\lambda_4\phi^{(4)}(x) + \frac{1}{72}\lambda_3^2\phi^{(6)}(x)\right)$$

where $\lambda_r = \frac{\kappa_r}{\sigma^r}$ and $\phi^{(r)}$ is the r th derivative of the Gaussian density.

Lemma B.6 (Exponential upper-tail for Poisson distribution [15, Theorem 1]). *If S has Poisson distribution with mean λ , then*

$$\mathbb{P}[|S - \lambda| > x] \leq 2e^{-\frac{x^2}{2(\lambda+x)}}.$$

REFERENCES

- [1] ABBE, E., LI, S., AND SLY, A. Proof of the contiguity conjecture and lognormal limit for the symmetric perceptron, 2021.
- [2] ACHLIOPAS, D., AND MOORE, C. The asymptotic order of the random k -SAT threshold. In *The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings.* (2002), IEEE, pp. 779–788.
- [3] ACHLIOPAS, D., AND MOORE, C. Random k -SAT: two moments suffice to cross a sharp threshold. *SIAM J. Comput.* 36, 3 (2006), 740–762.
- [4] ACHLIOPAS, D., AND NAOR, A. The two possible values of the chromatic number of a random graph. *Ann. of Math.* (2) 162, 3 (2005), 1335–1351.
- [5] ACHLIOPAS, D., NAOR, A., AND PERES, Y. On the maximum satisfiability of random formulas. *J. ACM* 54, 2 (2007), Art. 10, 21.
- [6] ACHLIOPAS, D., AND PERES, Y. The threshold for random k -SAT is $2^k \log 2 - O(k)$. *J. Amer. Math. Soc.* 17, 4 (2004), 947–973.
- [7] AUBIN, B., PERKINS, W., AND ZDEBOROVÁ, L. Storage capacity in symmetric binary perceptrons. *J. Phys. A* 52, 29 (2019), 294003, 32.
- [8] BANASZCZYK, W. Balancing vectors and gaussian measures of n -dimensional convex bodies. *Random Structures & Algorithms* 12, 4 (1998), 351–360.
- [9] BANSAL, N., DADUSH, D., AND GARG, S. An algorithm for Komlós conjecture matching Banaszczyk's bound. *SIAM J. Comput.* 48, 2 (2019), 534–553.
- [10] BANSAL, N., DADUSH, D., GARG, S., AND LOVETT, S. The Gram-Schmidt walk: a cure for the Banaszczyk blues. In *STOC'18—Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing* (2018), ACM, New York, pp. 587–597.
- [11] BANSAL, N., AND MEKA, R. On the discrepancy of random low degree set systems. *Random Structures Algorithms* 57, 3 (2020), 695–705.
- [12] BARBOUR, A. D., HOLST, L., AND JANSON, S. *Poisson approximation*, vol. 2 of *Oxford Studies in Probability*. The Clarendon Press, Oxford University Press, New York, 1992. Oxford Science Publications.
- [13] BECK, J., AND FIALA, T. “Integer-making” theorems. *Discrete Applied Mathematics* 3, 1 (1981), 1–8.
- [14] BUKH, B. An improvement of the Beck-Fiala theorem. *Combin. Probab. Comput.* 25, 3 (2016), 380–398.
- [15] CANONNE, C. L. A short note on poisson tail bounds. Retrieved from the website: <http://www.cs.columbia.edu/~ccanonne> (2017).
- [16] CHANDRASEKARAN, K., AND VEMPALA, S. S. Integer feasibility of random polytopes. In *ITCS'14—Proceedings of the 2014 Conference on Innovations in Theoretical Computer Science* (2014), ACM, New York, pp. 449–458.
- [17] CHAZELLE, B. *The discrepancy method*. Cambridge University Press, Cambridge, 2000. Randomness and complexity.
- [18] DIACONIS, P., AND HOLMES, S. *Stein's method: expository lectures and applications*, vol. 46 of *Institute of Mathematical Statistics Lecture Notes—Monograph Series*. Institute of Mathematical Statistics, Beachwood, OH, 2004. Papers from the Workshop on Stein's Method held at Stanford University, Stanford, CA, 1998.
- [19] DING, J., AND SUN, N. Capacity lower bound for the Ising perceptron. In *STOC'19—Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing* (2019), ACM, New York, pp. 816–827.
- [20] ELDAN, R., AND SINGH, M. Efficient algorithms for discrepancy minimization in convex sets. *Random Structures Algorithms* 53, 2 (2018), 289–307.
- [21] EZRA, E., AND LOVETT, S. On the Beck-Fiala conjecture for random set systems. *Random Structures Algorithms* 54, 4 (2019), 665–675.
- [22] FELLER, W. *An introduction to probability theory and its applications. Vol. I*. Third edition. John Wiley & Sons, Inc., New York-London-Sydney, 1968.
- [23] FRANKS, C., AND SAKS, M. On the discrepancy of random matrices with many columns. *Random Structures Algorithms* 57, 1 (2020), 64–96.
- [24] FRIEZE, A., AND WORMALD, N. C. Random k -SAT: a tight threshold for moderately growing k . *Combinatorica* 25, 3 (2005), 297–305.
- [25] GREENE, E., AND WELLNER, J. A. Exponential bounds for the hypergeometric distribution. *Bernoulli* 23, 3 (2017), 1911–1950.
- [26] HOBERG, R., AND ROTHVOSS, T. A Fourier-analytic approach for the discrepancy of random set systems. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms* (2019), SIAM, Philadelphia, PA, pp. 2547–2556.
- [27] HOEFFDING, W. Probability inequalities for sums of bounded random variables. In *The collected works of Wassily Hoeffding*. Springer-Verlag, New York, 1994, pp. 409–426.
- [28] JANSON, S. The second moment method, conditioning and approximation. In *Random discrete structures (Minneapolis, MN, 1993)*, vol. 76 of *IMA Vol. Math. Appl.* Springer, New York, 1996, pp. 175–183.
- [29] KUPERBERG, G., LOVETT, S., AND PELED, R. Probabilistic existence of rigid combinatorial structures. In *STOC'12—Proceedings of the 2012 ACM Symposium on Theory of Computing* (2012), ACM, New York, pp. 1091–1105.
- [30] MACRURY, C., MASAŘÍK, T., PAI, L., AND PÉREZ-GIMÉNEZ, X. The phase transition of discrepancy in random hypergraphs, 2021.

- [31] NIKOLOV, A., AND TALWAR, K. Approximating hereditary discrepancy via small width ellipsoids. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms* (2015), SIAM, Philadelphia, PA, pp. 324–336.
- [32] PERKINS, W., AND XU, C. Frozen 1-rsb structure of the symmetric ising perceptron, 2021.
- [33] PETROV, V. V. On local limit theorems for sums of independent random variables. *Theory of Probability & Its Applications* 9, 2 (1964), 312–320.
- [34] PINELIS, I. Coupling a binomial - parity conditioning. Mathematics Overflow. URL: <https://mathoverflow.net/q/379062> (version: 2020-12-18).
- [35] PINSKY, M. A. The normal approximation to the hypergeometric distribution.
- [36] POTUKUCHI, A. Discrepancy in random hypergraph models. *arXiv preprint arXiv:1811.01491* (2018).
- [37] POTUKUCHI, A. A spectral bound on hypergraph discrepancy. *arXiv preprint arXiv:1907.04117* (2019).
- [38] ROTHVOSS, T. Constructive discrepancy minimization for convex sets. *SIAM J. Comput.* 46, 1 (2017), 224–234.
- [39] SPENCER, J. Six standard deviations suffice. *Trans. Amer. Math. Soc.* 289, 2 (1985), 679–706.
- [40] SPENCER, J. *Ten lectures on the probabilistic method*, second ed., vol. 64 of *CBMS-NSF Regional Conference Series in Applied Mathematics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1994.
- [41] SPENCER, J. *Asymptopia*, vol. 71 of *Student Mathematical Library*. American Mathematical Society, Providence, RI, 2014. With Laura Florescu.
- [42] STEIN, C. A bound for the error in the normal approximation to the distribution of a sum of dependent random variables. In *Proceedings of the Sixth Berkeley Symposium on Mathematical Statistics and Probability (Univ. California, Berkeley, Calif., 1970/1971), Vol. II: Probability theory* (1972), pp. 583–602.
- [43] STEIN, E. M., AND SHAKARCHI, R. *Fourier analysis*, vol. 1 of *Princeton Lectures in Analysis*. Princeton University Press, Princeton, NJ, 2003. An introduction.
- [44] TALAGRAND, M. Intersecting random half cubes. *Random Structures Algorithms* 15, 3-4 (1999), 436–449. Statistical physics methods in discrete probability, combinatorics, and theoretical computer science (Princeton, NJ, 1997).
- [45] TURNER, P., MEKA, R., AND RIGOLLET, P. Balancing Gaussian vectors in high dimension. In *Conference on Learning Theory* (2020), pp. 3455–3486.