

SHARP NONASYMPTOTIC BOUNDS ON THE NORM OF RANDOM MATRICES WITH INDEPENDENT ENTRIES

BY AFONSO S. BANDEIRA¹ AND RAMON VAN HANDEL²

Princeton University

We obtain nonasymptotic bounds on the spectral norm of random matrices with independent entries that improve significantly on earlier results. If X is the $n \times n$ symmetric matrix with $X_{ij} \sim N(0, b_{ij}^2)$, we show that

$$\mathbf{E}\|X\| \lesssim \max_i \sqrt{\sum_j b_{ij}^2} + \max_{ij} |b_{ij}| \sqrt{\log n}.$$

This bound is optimal in the sense that a matching lower bound holds under mild assumptions, and the constants are sufficiently sharp that we can often capture the precise edge of the spectrum. Analogous results are obtained for rectangular matrices and for more general sub-Gaussian or heavy-tailed distributions of the entries, and we derive tail bounds in addition to bounds on the expected norm. The proofs are based on a combination of the moment method and geometric functional analysis techniques. As an application, we show that our bounds immediately yield the correct phase transition behavior of the spectral edge of random band matrices and of sparse Wigner matrices. We also recover a result of Seginer on the norm of Rademacher matrices.

1. Introduction. Understanding the behavior of the spectral norm of random matrices is a fundamental problem in probability theory, as well as a problem of considerable importance in many modern applications. If X is a Wigner matrix, that is, a symmetric $n \times n$ matrix whose entries are i.i.d. with unit variance, then a classical result in random matrix theory [1, 3, 9, 23] shows that $\|X\|/\sqrt{n} \rightarrow 2$ under mild moment assumptions (as is expected

Received August 2014; revised March 2015.

¹Supported by AFOSR Grant FA9550-12-1-0317.

²Supported in part by NSF Grant CAREER-DMS-1148711 and by the ARO through PECASE award W911NF-14-1-0094.

AMS 2000 subject classifications. Primary 60B20; secondary 46B09, 60F10.

Key words and phrases. Random matrices, spectral norm, nonasymptotic bounds, tail inequalities.

This is an electronic reprint of the original article published by the Institute of Mathematical Statistics in *The Annals of Probability*, 2016, Vol. 44, No. 4, 2479–2506. This reprint differs from the original in pagination and typographic detail.

from the well-known fact that the empirical spectral density converges to the semicircle law supported in $[-2, 2]$). The corresponding result for rectangular matrices with i.i.d. entries is even older [10]. More recently, there has been considerable interest in structured random matrices where the entries are no longer identically distributed. As the combinatorial methods that are used for this purpose typically exploit the specific structure of the entries, precise asymptotic results on the spectral norm of structured matrices must generally be obtained on a case-by-case basis; see, for example, [20, 21].

In order to gain a deeper understanding of the spectral norm of structured matrices, it is natural to ask whether one can find a unifying principle that captures at least the correct scale of the norm in a general setting, that is, in the absence of specific structural assumptions. This question is most naturally phrased in a nonasymptotic setting: can we obtain upper and lower bounds on $\|X\|$, in terms of natural parameters that capture the structure of X , that differ only by universal constants? Nonasymptotic bounds on the norm of a random matrix have long been developed in a different area of probability that arises from problems in geometric functional analysis, and have had a significant impact on various areas of pure and applied mathematics [8, 18, 25]. Unfortunately, as we will shortly see, the best known general results along these lines fail to capture the correct scale of the spectral norm of structured matrices except in extreme cases.

In this paper, we investigate the norm of random matrices with independent entries. Consider for concreteness the case of Gaussian matrices (our main results will extend to more general distributions of the entries). Let X be the $n \times n$ symmetric random matrix with entries $X_{ij} = g_{ij}b_{ij}$, where $\{g_{ij} : i \geq j\}$ are independent standard Gaussian random variables and $\{b_{ij} : i \geq j\}$ are given scalars.

Perhaps the most useful known nonasymptotic bound on the spectral norm $\|X\|$ can be obtained as a consequence of the noncommutative Khintchine inequality of Lust-Piquard and Pisier [16], or alternatively (in a much more elementary fashion) from the “matrix concentration” method that has been widely developed in recent years [15, 24]. This yields the following inequality in our setting:

$$\mathbf{E}\|X\| \lesssim \sigma \sqrt{\log n} \quad \text{with } \sigma := \max_i \sqrt{\sum_j b_{ij}^2}.$$

Unfortunately, this inequality already fails to be sharp in the simplest case of Wigner matrices: here $\sigma = \sqrt{n}$, so that the resulting bound $\mathbf{E}\|X\| \lesssim \sqrt{n \log n}$ falls short of the correct scaling $\mathbf{E}\|X\| \sim \sqrt{n}$. On the other hand, the logarithmic factor in this bound is necessary: if X is the diagonal matrix with independent standard Gaussian entries, then $\sigma = 1$ and $\mathbf{E}\|X\| \sim \sqrt{\log n}$. We therefore conclude that while the noncommutative Khintchine bound is

sharp in extreme cases, it fails to capture the structure of the matrix X in a satisfactory manner.

A different bound on $\|X\|$ can be obtained by a method due to Gordon (see [8]) that exploits Slepian’s comparison lemma for Gaussian processes, or alternatively from a simple ε -net argument [23, 25]. This yields the following inequality:

$$\mathbf{E}\|X\| \lesssim \sigma_* \sqrt{n} \quad \text{with } \sigma_* := \max_{ij} |b_{ij}|.$$

While the parameter σ_* that appears in this bound is often much smaller than σ , the dimensional scaling of this bound is much worse than in the noncommutative Khintchine bound. In particular, while this bound captures the correct \sqrt{n} rate for Wigner matrices, it is vastly suboptimal in almost every other situation (e.g., in the diagonal matrix example considered above).

Further nonasymptotic bounds on $\|X\|$ have been obtained in the present setting by Latała [12] and by Riemer and Schütt [17]. In most examples, these bounds provide even worse rates than the noncommutative Khintchine bound. Seginer [19] obtained a slight improvement on the noncommutative Khintchine bound that is specific to the special case where the random matrix has uniformly bounded entries; see Section 4.2 below. None of these results provides a sharp understanding of the scale of the spectral norm for general structured matrices.

The present paper develops a new family of nonasymptotic bounds on the spectral norm of structured random matrices that prove to be optimal in a surprisingly general setting. Our main bounds are of the form

$$\mathbf{E}\|X\| \lesssim \sigma + \sigma_* \sqrt{\log n},$$

which provides a sort of interpolation between the two bounds discussed above. For example, the following is one of the main results of this paper.

THEOREM 1.1. *Let X be the $n \times n$ symmetric matrix with $X_{ij} = g_{ij} b_{ij}$, where $\{g_{ij} : i \geq j\}$ are i.i.d. $\sim N(0, 1)$ and $\{b_{ij} : i \geq j\}$ are given scalars. Then*

$$\mathbf{E}\|X\| \leq (1 + \varepsilon) \left\{ 2\sigma + \frac{6}{\sqrt{\log(1 + \varepsilon)}} \sigma_* \sqrt{\log n} \right\}$$

for any $0 < \varepsilon \leq 1/2$, where σ, σ_* are as defined above.

Let us emphasize two important features of this result:

- It is almost trivial to obtain a matching lower bound of the form

$$\mathbf{E}\|X\| \gtrsim \sigma + \sigma_* \sqrt{\log n}$$

that holds as long as the coefficients b_{ij} are not too inhomogeneous (Section 3.5). This means that Theorem 1.1 captures the optimal scaling of the expected norm $\mathbf{E}\|X\|$ under surprisingly minimal structural assumptions.

- In the case of Wigner matrices, Theorem 1.1 yields a bound of the form

$$\mathbf{E}\|X\| \leq (1 + \varepsilon)2\sqrt{n} + o(\sqrt{n})$$

for arbitrarily small $\varepsilon > 0$. Thus Theorem 1.1 not only captures the correct scaling of the spectral norm, but even recovers the precise asymptotic behavior $\|X\|/\sqrt{n} \rightarrow 2$ as $n \rightarrow \infty$. This feature of Theorem 1.1 makes it possible to effortlessly prove nontrivial results, such as the precise phase transition behavior of the spectral edge of random band matrices (Section 4.1), that would be distinctly nontrivial to obtain by classical combinatorial methods.

In view of these observations, it seems that Theorem 1.1 is essentially the optimal result of its kind: there is little hope to accurately capture inhomogeneous models where Theorem 1.1 is not sharp in terms of simple parameters such as σ, σ_* ; see Remarks 3.8 and 3.16. On the other hand, we can now understand the previous bounds as extreme cases of Theorem 1.1. The non-commutative Khintchine bound matches Theorem 1.1 when $\sigma/\sigma_* \lesssim 1$: this case is minimal as $\sigma/\sigma_* \geq 1$. Gordon's bound matches Theorem 1.1 when $\sigma/\sigma_* \gtrsim \sqrt{n}$: this case is maximal as $\sigma/\sigma_* \leq \sqrt{n}$. In intermediate regimes, Theorem 1.1 yields a strictly better scaling.

While we have formulated the specific result of Theorem 1.1 for concreteness, our methods are not restricted to this particular setting. Once a complete proof of Theorem 1.1 has been given in Section 2, we will develop various extensions and complements in Section 3. These results are developed both for their independent interest, and in view of their potential utility in applications to other problems. In Section 3.1, we prove a sharp analogue of Theorem 1.1 in the setting of rectangular matrices. In Section 3.2, we develop versions of our main results when the entries are not necessarily Gaussian: for sub-Gaussian variables we obtain very similar results to the Gaussian case, while the scaling in our main results must be modified in the case of heavy-tailed entries. In Section 3.3, we develop variants of our main results where the explicit-dimensional dependence is replaced by a certain notion of effective dimension. In Section 3.4, we obtain sharp inequalities for the tail probabilities of the spectral norm $\|X\|$ rather than for its expectation. Finally, we obtain in Section 3.5 lower bounds on the spectral norm of Gaussian matrices that match our upper bounds under rather mild assumptions.

In Section 4, we develop two applications that illustrate the power of our main results. In Section 4.1, we investigate a phase transition phenomenon for the spectral edge of random band matrices and more general sparse Wigner matrices. Our main results effortlessly provide a precise understanding of this transition, which sharpens earlier results that were obtained by much more delicate combinatorial methods in [4, 11, 21]. In Section 4.2,

we investigate the setting of Rademacher random matrices with entries $X_{ij} = \varepsilon_{ij} b_{ij}$, where ε_{ij} are independent Rademacher (symmetric Bernoulli) variables. Here we recover a result of Seginer [19] with a much simpler proof, and develop insight into how such bounds can be improved.

One of the nice features of Theorem 1.1 is that its proof explains very clearly why the result is true. Once the idea has been understood, the technical details prove to be of minimal difficulty, which suggests that the “right” approach has been found. Let us briefly illustrate the idea behind the proof in the special case where the coefficients b_{ij} take only the values $\{0, 1\}$ (this setting guided our intuition, though the ultimate proof is no more difficult in the general setting). We can then interpret the matrix of coefficients (b_{ij}) as the adjacency matrix of a graph G on n points, and we have $\sigma_* = 1$ and $\sigma = \sqrt{k}$ where k is the maximal degree of G .

Following a classical idea in random matrix theory, we use the fact that the spectral norm $\|X\|$ is comparable to the quantity $\text{Tr}[X^p]^{1/p}$ for $p \sim \log n$. If one writes out the expression for $\mathbf{E} \text{Tr}[X^p]$ in terms of the coefficients, it is readily seen that controlling this quantity requires us to count the number of cycles in G for which every edge is visited an even number of times. One might expect that the graph G of degree k that possesses the most such cycles is the complete graph on k points. If this were the case, then one could control $\mathbf{E} \text{Tr}[X^p]$ by $\mathbf{E} \text{Tr}[Y^p]$ where Y is a Wigner matrix of dimension k . This intuition is almost, but not entirely correct: while a k -clique typically possesses more distinct topologies of cycles, each cycle of a given topology can typically be embedded in more ways in a regular graph on n points than in a k -clique. Careful bookkeeping shows that the latter can be accounted for by choosing a slightly larger Wigner matrix of dimension $k + p$. We therefore obtain a comparison theorem between the spectral norm of X and the spectral norm of a $(k + p)$ -dimensional Wigner matrix, which is of the desired order $\sqrt{k + p} \sim \sqrt{k} + \sqrt{\log n}$ for $p \sim \log n$. We can now conclude by using standard ideas from probability in Banach spaces to obtain sharp nonasymptotic bounds on the norm of the resulting Wigner matrix, avoiding entirely any combinatorial complications. (A purely combinatorial approach would be nontrivial as very high moments of Wigner matrices can appear in this argument.)

We conclude the [Introduction](#) by noting that both the noncommutative Khintchine inequality and Gordon’s bound can be formulated in a more general context beyond the case of independent entries. Whether the conclusion of Theorem 1.1 extends to this situation is a natural question of considerable interest.

Notation. Let us clarify a few notational conventions that will be used throughout the paper. In the sequel, $a \lesssim b$ means that $a \leq Cb$ for a universal constant C . (If C depends on other quantities, this will be indicated

explicitly.) If $a \lesssim b$ and $b \lesssim a$, we write $a \asymp b$. We write $a \wedge b := \min(a, b)$ and $a \vee b := \max(a, b)$, and we denote by $[n] := \{1, \dots, n\}$. Finally, we occasionally write $\|\xi\|_p := \mathbf{E}[\xi^p]^{1/p}$.

2. Proof of Theorem 1.1. The main idea behind the proof of Theorem 1.1 is the following comparison theorem.

PROPOSITION 2.1. *Let Y_r be the $r \times r$ symmetric matrix such that $\{(Y_r)_{ij} : i \geq j\}$ are independent $N(0, 1)$ random variables, and suppose that $\sigma_* \leq 1$. Then*

$$\mathbf{E} \operatorname{Tr}[X^{2p}] \leq \frac{n}{\lceil \sigma^2 \rceil + p} \mathbf{E} \operatorname{Tr}[Y_{\lceil \sigma^2 \rceil + p}^{2p}] \quad \text{for every } p \in \mathbb{N}.$$

Let us begin by completing the proof of Theorem 1.1 given this result. We need the following lemma, which is a variation on standard ideas; cf. [8].

LEMMA 2.2. *Let Y_r be the $r \times r$ symmetric matrix such that $\{(Y_r)_{ij} : i \geq j\}$ are independent $N(0, 1)$ random variables. Then for every $p \geq 2$*

$$\mathbf{E}[\|Y_r\|^{2p}]^{1/2p} \leq 2\sqrt{r} + 2\sqrt{2p}.$$

PROOF. We begin by noting that

$$\|Y_r\| = \lambda_+ \vee \lambda_-, \quad \lambda_+ := \sup_{v \in S} \langle v, Y_r v \rangle, \quad \lambda_- = - \inf_{v \in S} \langle v, Y_r v \rangle,$$

where S is the unit sphere in \mathbb{R}^r . We are therefore interested in the supremum of the Gaussian process $\{\langle v, Y_r v \rangle\}_{v \in S}$, whose natural distance can be estimated as

$$\mathbf{E}|\langle v, Y_r v \rangle - \langle w, Y_r w \rangle|^2 \leq 2 \sum_{i,j} \{v_i v_j - w_i w_j\}^2 \leq 4\|v - w\|^2$$

[using $1 - x^2 \leq 2(1 - x)$ for $x \leq 1$]. The right-hand side of this expression is the natural distance of the Gaussian process $\{2\langle v, g \rangle\}_{v \in S}$, where g is the standard Gaussian vector in \mathbb{R}^r . Therefore, Slepian's lemma [7], Theorem 13.3, implies

$$\mathbf{E}\lambda_+ = \mathbf{E} \sup_{v \in S} \langle v, Y_r v \rangle \leq 2\mathbf{E} \sup_{v \in S} \langle v, g \rangle = 2\mathbf{E}\|g\| \leq 2\sqrt{r}.$$

Moreover, note that λ_+ and λ_- have the same distribution (as evidently Y_r and $-Y_r$ have the same distribution). Therefore, using the triangle inequality for $\|\cdot\|_{2p}$,

$$\begin{aligned} \mathbf{E}[\|Y_r\|^{2p}]^{1/2p} &= \|\lambda_+ \vee \lambda_-\|_{2p} \\ &\leq \mathbf{E}\lambda_+ + \|\lambda_+ \vee \lambda_- - \mathbf{E}\lambda_+\|_{2p} \\ &= \mathbf{E}\lambda_+ + \|(\lambda_+ - \mathbf{E}\lambda_+) \vee (\lambda_- - \mathbf{E}\lambda_-)\|_{2p}. \end{aligned}$$

It follows from Gaussian concentration [7], Theorems 5.8 and 2.1, that

$$\mathbf{E}[(\lambda_+ - \mathbf{E}\lambda_+)^{2p} \vee (\lambda_- - \mathbf{E}\lambda_-)^{2p}] \leq p!4^{p+1} \leq (2\sqrt{2p})^{2p}$$

for $p \geq 2$. Putting together the above estimates completes the proof. \square

PROOF OF THEOREM 1.1. We can clearly assume without loss of generality that the matrix X is normalized such that $\sigma_* = 1$. For $p \geq 2$, we can estimate

$$\begin{aligned} \mathbf{E}\|X\| &\leq \mathbf{E}[\mathrm{Tr}[X^{2p}]^{1/2p}] \\ &\leq n^{1/2p} \mathbf{E}[\|Y_{\lceil \sigma^2 \rceil + p}\|^{2p}]^{1/2p} \\ &\leq n^{1/2p} \{2\sqrt{\lceil \sigma^2 \rceil + p} + 2\sqrt{2p}\} \end{aligned}$$

by Proposition 2.1 and Lemma 2.2, where we use $\mathrm{Tr}[Y_r^{2p}] \leq r\|Y_r\|^{2p}$. This yields

$$\begin{aligned} \mathbf{E}\|X\| &\leq e^{1/2\alpha} \{2\sqrt{\lceil \sigma^2 \rceil + \lceil \alpha \log n \rceil} + 2\sqrt{2\lceil \alpha \log n \rceil}\} \\ &\leq e^{1/2\alpha} \{2\sigma + 2\sqrt{\alpha \log n + 2} + 2\sqrt{2\alpha \log n + 2}\} \end{aligned}$$

for the choice $p = \lceil \alpha \log n \rceil$. If $n \geq 2$ and $\alpha \geq 1$, then $2 \leq 3 \log 2 \leq 3\alpha \log n$, so

$$\mathbf{E}\|X\| \leq e^{1/2\alpha} \{2\sigma + 6\sqrt{2\alpha \log n}\}.$$

Defining $e^{1/2\alpha} = 1 + \varepsilon$ and noting that $\varepsilon \leq 1/2$ implies $\alpha \geq 1$ yields the result provided that $n \geq 2$ and $p \geq 2$. The remaining cases are easily dealt with separately. The result holds trivially in the case $n = 1$. On the other hand, the case $p = 1$ can only occur when $\alpha \log n \leq 1$ and thus $n \leq 2$. In this case we can estimate directly $\mathbf{E}\|X\| \leq \sqrt{n(\lceil \sigma^2 \rceil + p)} \leq \sigma\sqrt{2} + 2$ using Proposition 2.1. \square

REMARK 2.3. Note that we use the moment method only to prove the comparison theorem of Proposition 2.1; as will be seen below, this requires only trivial combinatorics. All the usual combinatorial difficulties of random matrix theory are circumvented by Lemma 2.2, which exploits the theory of Gaussian processes. After the first version of this paper was posted, we learned from Shahar Mendelson that a related idea has been used in [2] for a different purpose.

REMARK 2.4. The constant 6 in the second term in Theorem 1.1 arises from crude rounding in our proof. While this constant can be somewhat improved for large n , our proof cannot yield a sharp constant here: it can be

verified in the example of the diagonal matrix $b_{ij} = \mathbf{1}_{i=j}$ that the constant $\sqrt{2}$ in the precise asymptotic $\mathbf{E}\|X\| \sim \sqrt{2\log n}$ cannot be recovered from our general proof. We therefore do not insist on optimizing this constant, but rather state the convenient bound in Theorem 1.1 which holds for any n . In contrast to the constant in the second term, it was shown in the [Introduction](#) that the constant in the first term is sharp.

We now turn to the proof of Proposition 2.1. Let us begin by recalling some standard observations. The quantity $\mathbf{E}\text{Tr}[X^{2p}]$ can be expanded as

$$\mathbf{E}\text{Tr}[X^{2p}] = \sum_{u_1, \dots, u_{2p} \in [n]} b_{u_1 u_2} b_{u_2 u_3} \cdots b_{u_{2p} u_1} \mathbf{E}[g_{u_1 u_2} g_{u_2 u_3} \cdots g_{u_{2p} u_1}].$$

Let $G_n = ([n], E_n)$ be the complete graph on n points, that is, $E_n = \{\{u, u'\} : u, u' \in [n]\}$. (Note that we have included self-loops.) We will identify any $\mathbf{u} = (u_1, \dots, u_{2p}) \in [n]^{2p}$ with a cycle $u_1 \rightarrow u_2 \rightarrow \cdots \rightarrow u_{2p} \rightarrow u_1$ in G_n of length $2p$. If we denote by $n_i(\mathbf{u})$ the number of distinct edges that are visited precisely i times by the cycle \mathbf{u} , then we can write [here $g \sim N(0, 1)$]

$$\mathbf{E}\text{Tr}[X^{2p}] = \sum_{\mathbf{u} \in [n]^{2p}} b_{u_1 u_2} b_{u_2 u_3} \cdots b_{u_{2p} u_1} \prod_{i \geq 1} \mathbf{E}[g^i]^{n_i(\mathbf{u})}.$$

A cycle \mathbf{u} is called *even* if it visits each distinct edge an even number of times, that is, if $n_i(\mathbf{u}) = 0$ whenever i is odd. As $\mathbf{E}[g^i] = 0$ when i is odd, it follows immediately that the sum in the above expression can be restricted to even cycles.

The *shape* $\mathbf{s}(\mathbf{u})$ of a cycle \mathbf{u} is obtained by relabeling the vertices in order of appearance. For example, the cycle $7 \rightarrow 3 \rightarrow 5 \rightarrow 4 \rightarrow 3 \rightarrow 5 \rightarrow 4 \rightarrow 3 \rightarrow 7$ has shape $1 \rightarrow 2 \rightarrow 3 \rightarrow 4 \rightarrow 2 \rightarrow 3 \rightarrow 4 \rightarrow 2 \rightarrow 1$. We denote by

$$\mathcal{S}_{2p} := \{\mathbf{s}(\mathbf{u}) : \mathbf{u} \text{ is an even cycle of length } 2p\}$$

the collection of shapes of even cycles, and we define the collection of even cycles with given shape \mathbf{s} starting (and ending) at a given point u as

$$\Gamma_{\mathbf{s}, u} := \{\mathbf{u} \in [n]^{2p} : \mathbf{s}(\mathbf{u}) = \mathbf{s}, u_1 = u\}$$

for any $u \in [n]$ and $\mathbf{s} \in \mathcal{S}_{2p}$. Clearly the edge counts $n_i(\mathbf{u})$ depend only on the shape $\mathbf{s}(\mathbf{u})$ of \mathbf{u} , and we can therefore unambiguously write $n_i(\mathbf{s})$ for the number of distinct edges visited i times by any cycle with shape \mathbf{s} . We then obtain

$$\mathbf{E}\text{Tr}[X^{2p}] = \sum_{u \in [n]} \sum_{\mathbf{s} \in \mathcal{S}_{2p}} \prod_{i \geq 1} \mathbf{E}[g^i]^{n_i(\mathbf{s})} \sum_{\mathbf{u} \in \Gamma_{\mathbf{s}, u}} b_{u_1 u_2} b_{u_2 u_3} \cdots b_{u_{2p} u_1}.$$

Finally, given any shape $\mathbf{s} = (s_1, \dots, s_{2p})$, we denote by $m(\mathbf{s}) = \max_i s_i$ the number of distinct vertices visited by any cycle with shape \mathbf{s} .

Now that we have set up a convenient bookkeeping device, the proof of Proposition 2.1 is surprisingly straightforward. It relies on two basic observations.

LEMMA 2.5. *Suppose that $\sigma_* \leq 1$. Then we have for any $u \in [n]$ and $\mathbf{s} \in \mathcal{S}_{2p}$*

$$\sum_{\mathbf{u} \in \Gamma_{\mathbf{s}, u}} b_{u_1 u_2} b_{u_2 u_3} \cdots b_{u_{2p} u_1} \leq \sigma^{2(m(\mathbf{s})-1)}.$$

In particular, it follows that

$$\mathbf{E} \operatorname{Tr}[X^{2p}] \leq n \sum_{\mathbf{s} \in \mathcal{S}_{2p}} \sigma^{2(m(\mathbf{s})-1)} \prod_{i \geq 1} \mathbf{E}[g^i]^{n_i(\mathbf{s})}.$$

PROOF. Fix an initial point u , and shape $\mathbf{s} = (s_1, \dots, s_{2p})$. Let

$$i(k) = \inf\{j : s_j = k\}$$

for $1 \leq k \leq m(\mathbf{s})$. That is, $i(k)$ is the first time in any cycle of shape \mathbf{s} at which its k th distinct vertex is visited [of course, $i(1) = 1$ by definition].

Now consider any cycle $\mathbf{u} \in \Gamma_{\mathbf{s}, u}$. As the cycle is even, the edge $\{u_{i(k)-1}, u_{i(k)}\}$ must be visited at least twice for every $2 \leq k \leq m(\mathbf{s})$. On the other hand, as the vertex $u_{i(k)}$ is visited for the first time at time $i(k)$, the edge $\{u_{i(k)-1}, u_{i(k)}\}$ must be distinct from the edges $\{u_{i(\ell)-1}, u_{i(\ell)}\}$ for all $\ell < k$. We can therefore estimate

$$\begin{aligned} \sum_{\mathbf{u} \in \Gamma_{\mathbf{s}, u}} b_{u_1 u_2} b_{u_2 u_3} \cdots b_{u_{2p} u_1} &\leq \sum_{\mathbf{u} \in \Gamma_{\mathbf{s}, u}} b_{u u_{i(2)}}^2 b_{u_{i(3)-1} u_{i(3)}}^2 \cdots b_{u_{i(m(\mathbf{s})-1)} u_{i(m(\mathbf{s}))}}^2 \\ &= \sum_{v_2 \neq \cdots \neq v_{m(\mathbf{s})}} b_{u v_2}^2 b_{v_{s_i(3)-1} v_3}^2 \cdots b_{v_{s_i(m(\mathbf{s})-1)} v_{m(\mathbf{s})}}^2, \end{aligned}$$

where we use that $\max_{i,j} |b_{ij}| = \sigma_* \leq 1$. As $s_{i(k)-1} < k$ by construction, it is readily seen that the quantity on the right-hand side is bounded by $\sigma^{2(m(\mathbf{s})-1)}$. \square

LEMMA 2.6. *Let Y_r be defined as in Proposition 2.1. Then for any $r > p$*

$$\mathbf{E} \operatorname{Tr}[Y_r^{2p}] = r \sum_{\mathbf{s} \in \mathcal{S}_{2p}} (r-1)(r-2) \cdots (r-m(\mathbf{s})+1) \prod_{i \geq 1} \mathbf{E}[g^i]^{n_i(\mathbf{s})}.$$

PROOF. In complete analogy with the identity for $\mathbf{E} \operatorname{Tr}[X^{2p}]$, we can write

$$\mathbf{E} \operatorname{Tr}[Y_r^{2p}] = \sum_{\mathbf{s} \in \mathcal{S}_{2p}} |\{\mathbf{u} \in [r]^{2p} : \mathbf{s}(\mathbf{u}) = \mathbf{s}\}| \prod_{i \geq 1} \mathbf{E}[g^i]^{n_i(\mathbf{s})}.$$

Each cycle $\mathbf{u} \in [r]^{2p}$ with given shape $\mathbf{s}(\mathbf{u}) = \mathbf{s}$ is uniquely defined by specifying its $m(\mathbf{s})$ distinct vertices. Thus as long as $m(\mathbf{s}) \leq r$, there are precisely

$$r(r-1)\cdots(r-m(\mathbf{s})+1)$$

such cycles. However, note that any even cycle of length $2p$ can visit at most $m(\mathbf{s}) \leq p+1$ distinct vertices, so the assumption $p < r$ implies the result. \square

We can now complete the proof.

PROOF OF PROPOSITION 2.1. Fix $p \in \mathbb{N}$, and let $r = \lceil \sigma^2 \rceil + p$. Then

$$(r-1)(r-2)\cdots(r-m(\mathbf{s})+1) \geq (\sigma^2 + p - m(\mathbf{s}) + 1)^{m(\mathbf{s})-1} \geq \sigma^{2(m(\mathbf{s})-1)}$$

for any $\mathbf{s} \in \mathcal{S}_{2p}$, where we have used that any even cycle of length $2p$ can visit at most $m(\mathbf{s}) \leq p+1$ distinct vertices. It remains to apply Lemmas 2.5 and 2.6. \square

3. Extensions and complements.

3.1. *Nonsymmetric matrices.* Let X be the $n \times m$ random rectangular matrix with $X_{ij} = g_{ij}b_{ij}$, where $\{g_{ij} : 1 \leq i \leq n, 1 \leq j \leq m\}$ are independent $N(0, 1)$ random variables and $\{b_{ij} : 1 \leq i \leq n, 1 \leq j \leq m\}$ are given scalars. While this matrix is not symmetric, one can immediately obtain a bound on $\mathbf{E}\|X\|$ from Theorem 1.1 by applying the latter to the symmetric matrix

$$\tilde{X} = \begin{bmatrix} 0 & X \\ X^* & 0 \end{bmatrix}.$$

Indeed, it is readily seen that $\|\tilde{X}\| = \|X\|$, so we obtain

$$\mathbf{E}\|X\| \leq (1 + \varepsilon) \left\{ 2(\sigma_1 \vee \sigma_2) + \frac{6}{\sqrt{\log(1 + \varepsilon)}} \sigma_* \sqrt{\log(n + m)} \right\}$$

for any $0 < \varepsilon \leq 1/2$ with

$$\sigma_1 := \max_i \sqrt{\sum_j b_{ij}^2}, \quad \sigma_2 := \max_j \sqrt{\sum_i b_{ij}^2}, \quad \sigma_* := \max_{ij} |b_{ij}|.$$

While this result is largely satisfactory, it does not lead to a sharp constant in the first term: it is known from asymptotic theory [10] that when $b_{ij} = 1$ for all i, j , we have $\mathbf{E}\|X\| \sim \sqrt{n} + \sqrt{m}$ as $n, m \rightarrow \infty$ with $n/m \rightarrow \gamma \in]0, \infty[$, while the above bound can only give the weaker inequality $\mathbf{E}\|X\| \leq 2(1 + o(1))(\sqrt{n} \vee \sqrt{m})$. The latter bound can therefore be off by as much as a factor 2.

We can regain the lost factor and also improve the logarithmic term by exploiting explicitly the bipartite structure of \tilde{X} in the proof of Theorem 1.1. This leads to the following sharp analogue of Theorem 1.1 for rectangular random matrices.

THEOREM 3.1. *Let X be the $n \times m$ matrix with $X_{ij} = g_{ij}b_{ij}$. Then*

$$\mathbf{E}\|X\| \leq (1 + \varepsilon) \left\{ \sigma_1 + \sigma_2 + \frac{5}{\sqrt{\log(1 + \varepsilon)}} \sigma_* \sqrt{\log(n \wedge m)} \right\}$$

for any $0 < \varepsilon \leq 1/2$.

As the proof of this result closely follows the proof of Theorem 3.1, we will only sketch the necessary modifications to the proof in the rectangular setting.

SKETCH OF PROOF. Let $G_{n,m} = ([n] \sqcup [m], E_{n,m})$ be the complete bipartite graph whose left and right vertices are indexed by $[n]$ and $[m]$, respectively (i.e., with edges $E_{n,m} = \{(u, v) : u \in [n], v \in [m]\}$). We begin by noting that

$$\begin{aligned} & \mathbf{E} \operatorname{Tr}[(XX^*)^p] \\ &= \sum_{\mathbf{u} \in [n]^p} \sum_{\mathbf{v} \in [m]^p} b_{u_1 v_1} b_{u_2 v_1} b_{u_2 v_2} b_{u_3 v_2} \cdots b_{u_p v_p} b_{u_1 v_p} \prod_{i \geq 1} \mathbf{E}[g^i]^{n_i(\mathbf{u}, \mathbf{v})}, \end{aligned}$$

where we denote by $n_i(\mathbf{u}, \mathbf{v})$ the number of distinct edges in $G_{n,m}$ that are visited precisely i times by the cycle $u_1 \rightarrow v_1 \rightarrow u_2 \rightarrow v_2 \rightarrow \cdots \rightarrow u_p \rightarrow v_p \rightarrow u_1$. In direct analogy with the symmetric case, we can define the collection \mathcal{S}_{2p} of shapes of even cycles of length $2p$, and by $\Gamma_{\mathbf{s}, u}$ the collection of cycles with given shape $\mathbf{s} \in \mathcal{S}_{2p}$ starting at a given point $u \in [n]$. We denote by $n_i(\mathbf{s})$ the number of distinct edges that are visited precisely i times by \mathbf{s} , and by $m_1(\mathbf{s})$ and $m_2(\mathbf{s})$ the number of distinct right and left vertices, respectively, that are visited by \mathbf{s} (i.e., the number of distinct vertices that appear in even and odd positions in the cycle).

It is now straightforward to adapt the proofs of Lemmas 2.5 and 2.6 to the present setting. Assuming that $\sigma_* \leq 1$, the analogue of Lemma 2.5 yields

$$\mathbf{E} \operatorname{Tr}[(XX^*)^p] \leq n \sum_{\mathbf{s} \in \mathcal{S}_{2p}} \sigma_1^{2m_1(\mathbf{s})} \sigma_2^{2(m_2(\mathbf{s})-1)} \prod_{i \geq 1} \mathbf{E}[g^i]^{n_i(\mathbf{s})}.$$

On the other hand, let $Y_{r,r'}$ be the $r \times r'$ matrix whose entries are independent $N(0, 1)$ random variables. Then the analogue of Lemma 2.6 yields

$$\mathbf{E} \operatorname{Tr}[(Y_{r,r'} Y_{r,r'}^*)^p]$$

$$\begin{aligned}
&= r \sum_{\mathbf{s} \in \mathcal{S}_{2p}} (r-1) \cdots (r - m_2(\mathbf{s}) + 1) r' (r' - 1) \cdots (r' - m_1(\mathbf{s}) + 1) \\
&\quad \times \prod_{i \geq 1} \mathbf{E}[g^i]^{n_i(\mathbf{s})}
\end{aligned}$$

when $r > p/2$ and $r' > p/2$. Choosing $r = \lceil \sigma_2^2 + p/2 \rceil$ and $r' = \lceil \sigma_1^2 + p/2 \rceil$ yields

$$\mathbf{E} \operatorname{Tr}[(XX^*)^p] \leq \frac{n}{r} \mathbf{E} \operatorname{Tr}[(Y_{r,r'} Y_{r,r'}^*)^p]$$

in analogy with Proposition 2.1. To complete the proof, we note that adapting the argument of Lemma 2.2 to the rectangular case (cf. [8]) yields

$$\mathbf{E}[\|Y_{r,r'}\|^{2p}]^{1/2p} \leq \sqrt{r} + \sqrt{r'} + 2\sqrt{p}.$$

We can therefore estimate (assuming without loss of generality that $\sigma_* = 1$)

$$\mathbf{E}\|X\| \leq \mathbf{E} \operatorname{Tr}[(XX^*)^p]^{1/2p} \leq n^{1/2p} \left\{ \sqrt{\lceil \sigma_1^2 + p/2 \rceil} + \sqrt{\lceil \sigma_2^2 + p/2 \rceil} + 2\sqrt{p} \right\}.$$

Choosing $p = \lceil \alpha \log n \rceil$ and proceeding as in the proof of Theorem 1.1 yields the result with a dimensional factor of $\sqrt{\log n}$ rather than $\sqrt{\log(n \wedge m)}$. However, as $\|X\| = \|X^*\|$, the latter bound follows by exchanging the roles of n and m . \square

3.2. Non-Gaussian variables. We have phrased our main results in terms of Gaussian random matrices for concreteness. However, note that the core argument of the proof of Theorem 1.1, the comparison principle of Proposition 2.1, did not depend at all on the Gaussian nature of the entries: it is only subsequently in Lemma 2.2 that we exploited the theory of Gaussian processes. The same observation applies to the proof of Theorem 3.1. As a consequence, we can develop various extensions of our main results to more general distributions of the entries.

Let us begin by considering the case of sub-Gaussian random variables.

COROLLARY 3.2. *Theorems 1.1 and 3.1 remain valid if the independent Gaussian random variables g_{ij} are replaced by independent symmetric random variables ξ_{ij} such that $\mathbf{E}[\xi_{ij}^{2p}] \leq \mathbf{E}[g^{2p}]$ for every $p \in \mathbb{N}$ and i, j [$g \sim N(0, 1)$].*

PROOF. As ξ_{ij} are assumed to be symmetric, $\mathbf{E}[\xi_{ij}^p] = 0$ when p is odd. It therefore follows readily by inspection of the proof that Proposition 2.1 (and its rectangular counterpart) remains valid under the present assumptions. \square

Corollary 3.2 implies, for example, that the conclusions of Theorems 1.1 and 3.1 hold verbatim when g_{ij} are replaced by independent Rademacher variables ε_{ij} , that is, $\mathbf{P}[\varepsilon_{ij} = \pm 1] = 1/2$; see Section 4.2 below for more on such matrices. The moment assumption $\mathbf{E}[\xi_{ij}^{2p}] \leq \mathbf{E}[g^{2p}]$ is somewhat unwieldy, however. We can obtain a similar result under standard sub-Gaussian tail assumptions.

COROLLARY 3.3. *If the independent Gaussian variables g_{ij} are replaced by independent random variables ξ_{ij} that are centered and sub-Gaussian in the sense*

$$\mathbf{E}[\xi_{ij}] = 0, \quad \mathbf{P}[|\xi_{ij}| \geq t] \leq Ce^{-t^2/2c} \quad \text{for all } t \geq 0 \text{ and } i, j,$$

then Theorems 1.1 and 3.1 remain valid up to a universal constant that depends on C and c only. That is, we have $\mathbf{E}\|X\| \lesssim \sigma + \sigma_\sqrt{\log n}$ in the case of Theorem 1.1, and $\mathbf{E}\|X\| \lesssim \sigma_1 + \sigma_2 + \sigma_*\sqrt{\log(n \wedge m)}$ in the case of Theorem 3.1.*

PROOF. Let X' be an independent copy of X . As $\mathbf{E}X' = 0$, we obtain by Jensen's inequality $\mathbf{E}\|X\| = \mathbf{E}\|X - \mathbf{E}X'\| \leq \mathbf{E}\|X - X'\|$. The entries of the matrix $X - X'$ are still sub-Gaussian (with the constants C, c increasing by at most a constant factor), but are now symmetric as well. We can therefore assume without loss of generality that ξ_{ij} are symmetric sub-Gaussian random variables. Using the integration formula $\mathbf{E}[\xi^{2p}] = \int_0^\infty \mathbf{P}[|\xi| \geq t^{1/2p}] dt$, it is readily shown that the random variables ξ_{ij}/K satisfy $\mathbf{E}[\xi_{ij}^{2p}] \leq \mathbf{E}[g^{2p}]$ for all $p \in \mathbb{N}$, where K is a constant that depends on C, c only. The result follows from Corollary 3.2. \square

REMARK 3.4. The main difference between Corollaries 3.2 and 3.3 is that the bound of Corollary 3.3 is multiplied by an additional constant factor as compared to Corollary 3.2. Thus the constant in front of the leading term in Corollary 3.3 is no longer sharp. This is of little consequence in many applications (particularly in nonasymptotic problems), but implies that we no longer capture the exact asymptotics of Wigner matrices. The latter can sometimes be recovered at the expense of increasing the logarithmic term in the estimate; see Corollary 3.6 below.

The sub-Gaussian assumption of Corollary 3.2 requires that the random variables ξ_{ij} have at worst Gaussian tails. For random variables with heavier tails, the conclusion of Theorem 1.1 cannot hold as stated. Consider, for example, the diagonal case where $b_{ij} = \mathbf{1}_{i=j}$, so that $\sigma = \sigma_* = 1$; then $\|X\| = \max_{i \leq n} |\xi_{ii}|$, which must grow faster in the heavy-tailed setting than the $\sim \sqrt{\log n}$ bound that would be obtained if the conclusion of Theorem 1.1

were valid. It seems reasonable to expect that in the case of heavy-tailed entries, the $\sqrt{\log n}$ rate must be changed to a quantity that controls the maximum of the heavy-tailed random variables under consideration. This is, roughly speaking, the content of the following result. (We will work in the setting of Theorem 1.1 for simplicity, though an entirely analogous result can be proved in the setting of Theorem 3.1.)

COROLLARY 3.5. *Let X be the $n \times n$ symmetric matrix with $X_{ij} = \xi_{ij} b_{ij}$, where $\{\xi_{ij} : i \geq j\}$ are independent centered random variables and $\{b_{ij} : i \geq j\}$ are given scalars. If $\mathbf{E}[\xi_{ij}^{2p}]^{1/2p} \leq K p^{\beta/2}$ for some $K, \beta > 0$ and all p, i, j , then*

$$\mathbf{E}\|X\| \lesssim \sigma + \sigma_* \log^{(\beta \vee 1)/2} n.$$

The universal constant in the above inequality depends on K, β only.

Let us note that as $\mathbf{E}[\xi^{2p}]^{1/2p} \lesssim \sqrt{p}$ for sub-Gaussian random variables ξ , Corollary 3.5 reduces to Corollary 3.3 in the sub-Gaussian setting. If we consider subexponential random variables, for example, then $\mathbf{E}[\xi^{2p}]^{1/2p} \lesssim p$, and thus $\sqrt{\log n}$ must be replaced by $\log n$ in the conclusion of Theorem 1.1. These scalings are precisely as expected, as the maximum of n independent random variables ξ_i with $\mathbf{E}[\xi_i^{2p}]^{1/2p} \sim p^{\beta/2}$ is of order $\log^{\beta/2} n$. Note, however, that the logarithmic factor only changes when the tails of the entries are heavier than Gaussian. The $\sqrt{\log n}$ factor cannot be reduced, in general, when the entries have lighter tails than Gaussian, as the universality property of many random matrix models leads to essentially Gaussian behavior; see Remark 4.8 for further discussion and examples.

PROOF OF COROLLARY 3.5. Symmetrizing as in the proof of Corollary 3.3, we can assume without loss of generality that ξ_{ij} are symmetric random variables. We will also assume without loss of generality that $\beta \geq 1$, as the case $\beta < 1$ is covered by Corollary 3.3.

Let g_{ij} and \tilde{g}_{ij} be i.i.d. $N(0, 1)$ random variables, and define $\eta_{ij} = g_{ij} |\tilde{g}_{ij}|^{\beta-1}$. Then η_{ij} are symmetric random variables, and by Stirling's formula

$$\mathbf{E}[\eta_{ij}^{2p}]^{1/2p} = \left[\frac{2^{p\beta}}{\pi} \Gamma\left(p + \frac{1}{2}\right) \Gamma\left(p(\beta - 1) + \frac{1}{2}\right) \right]^{1/2p} \gtrsim p^{\beta/2}.$$

If we denote by \tilde{X} the matrix with entries $\tilde{X}_{ij} = \eta_{ij} b_{ij}$, then it follows readily from the trace identities in the proof of Theorem 1.1 that

$$\mathbf{E} \operatorname{Tr}[X^{2p}] \leq C^{2p} \mathbf{E} \operatorname{Tr}[\tilde{X}^{2p}] \leq C^{2p} n \mathbf{E}\|\tilde{X}\|^{2p}$$

for all p and a universal constant C (depending on K, β). We therefore have

$$\mathbf{E}\|X\| \lesssim \mathbf{E}[\|\tilde{X}\|^{2\lceil \log n \rceil}]^{1/2\lceil \log n \rceil}.$$

Applying Theorem 1.1 conditionally on the variables $\tilde{g} = \{\tilde{g}_{ij}\}$ yields

$$\mathbf{E}[\|\tilde{X}\|^{2\lceil\log n\rceil}|\tilde{g}] \leq \left[C \max_i \sqrt{\sum_j b_{ij}^2 |\tilde{g}_{ij}|^{2\beta-2}} + C\sigma_* \max_{ij} |\tilde{g}_{ij}|^{\beta-1} \sqrt{\log n} \right]^{2\lceil\log n\rceil}$$

for another universal constant C . (While the statement of Theorem 1.1 only gives a bound on $\mathbf{E}\|X\|$, an inspection of the proof shows that what is in fact being bounded is the quantity $\mathbf{E}[\|X\|^{2\lceil\alpha\log n\rceil}]^{1/2\lceil\alpha\log n\rceil}$ with $\alpha = 1/2\log(1+\varepsilon)$.)

We must now estimate the expectation of the right-hand side of this equation. Note that $\|\max_{i\leq n} |Z_i|\|_p \leq \mathbf{E}[\sum_{i=1}^n |Z_i|^p]^{1/p} \leq n^{1/p} \max_{i\leq n} \|Z_i\|_p$ for any random variables Z_1, \dots, Z_n . Using $\|\tilde{g}_{ij}\|_p \lesssim \sqrt{p}$, a simple computation yields

$$\left\| \max_{ij} |\tilde{g}_{ij}|^{\beta-1} \right\|_{2\lceil\log n\rceil} \lesssim \log^{(\beta-1)/2} n.$$

Similarly, using the Rosenthal-type inequality of [6], Theorem 8, we obtain

$$\left\| \max_i \sum_j b_{ij}^2 |\tilde{g}_{ij}|^{2\beta-2} \right\|_{\lceil\log n\rceil} \lesssim \sigma^2 + \sigma_*^2 \left\| \max_{ij} |\tilde{g}_{ij}|^{\beta-1} \right\|_{2\lceil\log n\rceil}^2 \log n.$$

Substituting these estimates into the above expression completes the proof. \square

As was discussed above, the drawback of Corollary 3.3 and 3.5 is that an additional universal constant is introduced as compared to Theorem 1.1. In the following result, we have retained the sharp constant at the expense of a suboptimal scaling of the logarithmic term: for example, when applied to Gaussian entries, Corollary 3.6 yields $\log n$ instead of $\sqrt{\log n}$ in Theorem 1.1. Nonetheless, Corollary 3.6 can be useful in that it captures the sharp asymptotics of the edge of the spectrum of Wigner-type matrices as long as σ dominates the logarithmic term. Moreover, when the random variables ξ_{ij} are uniformly bounded, Corollary 3.6 is sharper than Corollary 3.2 in that the leading term depends on the variance rather than the uniform size of the entries; this will be exploited in Section 3.4 below.

COROLLARY 3.6. *Let X be the $n \times n$ symmetric random matrix with $X_{ij} = \xi_{ij} b_{ij}$, where $\{\xi_{ij} : i \geq j\}$ are independent symmetric random variables with unit variance and $\{b_{ij} : i \geq j\}$ are given scalars. Then we have for any $\alpha \geq 3$*

$$\mathbf{E}\|X\| \leq e^{2/\alpha} \left\{ 2\sigma + 14\alpha \max_{ij} \|\xi_{ij} b_{ij}\|_{2\lceil\alpha\log n\rceil} \sqrt{\log n} \right\}.$$

PROOF. Let ε_{ij} be i.i.d. Rademacher random variables independent of X , and denote by \tilde{X} the matrix with entries $\tilde{X}_{ij} = \varepsilon_{ij} X_{ij}$. As we assumed that ξ_{ij} are symmetric random variables, evidently X and \tilde{X} have the same distribution. We now apply Corollary 3.2 to \tilde{X} conditionally on the matrix X . This yields

$$\mathbf{E}\|X\| \leq (1 + \delta)\mathbf{E}\left[2\sqrt{\max_i \sum_j X_{ij}^2} + \frac{6}{\sqrt{\log(1 + \delta)}} \max_{ij} |X_{ij}| \sqrt{\log n}\right]$$

for any $0 < \delta \leq 1/2$. We can estimate

$$\mathbf{E} \max_{ij} |X_{ij}| \leq \left[\sum_{ij} \mathbf{E} |X_{ij}|^{2\lceil \alpha \log n \rceil} \right]^{1/2\lceil \alpha \log n \rceil} \leq e^{1/\alpha} \max_{ij} \|\xi_{ij} b_{ij}\|_{2\lceil \alpha \log n \rceil}.$$

On the other hand, by the Rosenthal-type inequality of [6], Theorem 8, we have

$$\left\| \sum_j X_{ij}^2 \right\|_{\lceil \alpha \log n \rceil} \leq (1 + \delta) \sum_j b_{ij}^2 + \frac{2}{\delta} \left\| \max_j X_{ij}^2 \right\|_{\lceil \alpha \log n \rceil}^{\lceil \alpha \log n \rceil},$$

so that

$$\mathbf{E} \left[\max_i \sum_j X_{ij}^2 \right] \leq e^{1/\alpha} \left\{ (1 + \delta)\sigma^2 + \frac{2e^{1/\alpha}}{\delta} \max_{ij} \|\xi_{ij} b_{ij}\|_{2\lceil \alpha \log n \rceil}^2 \lceil \alpha \log n \rceil \right\}.$$

Choosing α such that $e^{1/\alpha} = 1 + \delta$ and using $\delta \geq \log(1 + \delta)$, the result follows by combining the above estimates and straightforward manipulations. \square

3.3. Dimension-free bounds. A drawback of the results obtained so far is that they depend explicitly on the dimension n of the random matrix. This dependence is sharp in many natural situations; see Section 3.5 below. On the other hand, the results of Latała [12] and of Riemer and Schütt [17] have shown that it is possible to obtain dimension-free estimates, where n is replaced by an “effective dimension” that is defined in terms of a norm of the matrix of coefficients of the form

$$|(b_{ij})|_p := \left[\sum_{ij} |b_{ij}|^p \right]^{1/p}.$$

While the bounds of [12, 17] yield suboptimal results in many cases, a dimension-free formulation has at least two advantages. First, a low-dimensional matrix can be embedded in a high-dimensional space without changing its norm: for example, if all $b_{ij} = 0$ except $b_{11} = 1$, then $\mathbf{E}\|X\| \sim 1$, but Theorem 1.1 yields a bound of order $\sqrt{\log n}$. The advantage of a dimension-free bound is that it automatically adapts to high-dimensional matrices that

possess approximate low-dimensional structure. Second, dimension-free results can be used to study infinite-dimensional matrices, while Theorem 1.1 is not directly applicable in this setting.

The following result provides a dimension-free analogue of Theorems 1.1 and 3.1. To prove it, we apply the stratification technique developed in [17].

COROLLARY 3.7. *Let X be the $n \times n$ symmetric matrix with $X_{ij} = g_{ij}b_{ij}$, where $\{g_{ij} : i \geq j\}$ are i.i.d. $\sim N(0, 1)$ and $\{b_{ij} : i \geq j\}$ are given scalars. Then*

$$\mathbf{E}\|X\| \lesssim \sigma + \sigma_* \sqrt{\log \frac{|(b_{ij})|_p}{\sigma_*}}$$

for any $1 \leq p < 2$, where the universal constant depends on p only. Similarly, if X is the $n \times m$ random rectangular matrix $X_{ij} = g_{ij}b_{ij}$, then for any $1 \leq p < 2$

$$\mathbf{E}\|X\| \lesssim \sigma_1 + \sigma_2 + \sigma_* \sqrt{\log \frac{|(b_{ij})|_p}{\sigma_*}}.$$

PROOF. We will prove the result in the symmetric case; the proof in the rectangular case is identical. We also assume without loss of generality that $\sigma_* = 1$.

Define the matrices $X^{(k)}$ for $k \geq 1$ as $X_{ij}^{(k)} = g_{ij}b_{ij}\mathbf{1}_{2^{-k} < |b_{ij}| \leq 2^{-k+1}}$, so that

$$\mathbf{E}\|X\| = \mathbf{E}\left\| \sum_{k \geq 1} X^{(k)} \right\| \leq \mathbf{E}\left\| \sum_{k < k_0} X^{(k)} \right\| + \sum_{k \geq k_0} \mathbf{E}\|X^{(k)}\|$$

for a constant k_0 to be chosen appropriately in the sequel.

Denote by $c(k)$ the number of nonzero entries of $X^{(k)}$. Then

$$c(k) := |\{ij : 2^{-k} < |b_{ij}| \leq 2^{-k+1}\}| \leq 2^{kp} |(b_{ij})|_p^p.$$

Therefore, the nonzero entries of $X^{(k)}$ must be contained in a submatrix of size $c(k) \times c(k)$. Applying Theorem 1.1 to this submatrix yields

$$\mathbf{E}\|X^{(k)}\| \lesssim 2^{-k+1} \{\sqrt{c(k)} + \sqrt{\log c(k)}\} \lesssim 2^{-k(1-p/2)} |(b_{ij})|_p^{p/2}.$$

As $p < 2$, the right-hand side decays geometrically. Let k_0 be the smallest integer k such that $2^{-k(1-p/2)} |(b_{ij})|_p^{p/2} \leq 1$. Then we can estimate

$$\sum_{k \geq k_0} \mathbf{E}\|X^{(k)}\| \lesssim \sum_{k \geq k_0} 2^{-k(1-p/2)} |(b_{ij})|_p^{p/2} \leq \sum_{k \geq 0} 2^{-k(1-p/2)} \lesssim \sigma,$$

where we use $\sigma \geq \sigma_* = 1$. On the other hand, the matrix $\sum_{k < k_0} X^{(k)}$ has at most

$$\sum_{k < k_0} c(k) \leq \sum_{k < k_0} 2^{kp} |(b_{ij})|_p^p \lesssim 2^{k_0 p} |(b_{ij})|_p^p \lesssim |(b_{ij})|_p^{p+p^2/(2-p)}$$

entries by the definition of k_0 . Applying Theorem 1.1 completes the proof. \square

Note that the scaling in Corollary 3.7 improves as we increase p . Unfortunately, the constant blows up as $p \rightarrow 2$, so we need $p < 2$ to obtain a nontrivial result.

REMARK 3.8. Up to universal constants, the result of Corollary 3.7 is strictly better than that of Theorems 1.1 and 3.1 as $|(b_{ij})|_p \leq n^{2/p}\sigma_*$. It greatly improves the bounds of [17]. The bound of [12] is of a somewhat different nature: Latała proves the inequality $\mathbf{E}\|X\| \lesssim \sigma_1 + \sigma_2 + |(b_{ij})|_4$. The latter bound is significantly worse than Corollary 3.7 in most cases, but they are not strictly comparable.

Let us emphasize, however, that all the notions of effective dimension used here or in [12, 17] are essentially ad-hoc constructions. As will be shown in Section 3.5 below, the bounds of Theorems 1.1 and 3.1 are tight in situations where the coefficients b_{ij} exhibit a sufficient degree of homogeneity. The improvement provided by the dimension-free bounds is therefore of interest only in those cases where there is significant inhomogeneity in the magnitude of the coefficients, that is, in the presence of many scales. It is, however, unreasonable to expect that such inhomogeneity can be captured in a sharp manner by a norm of the form $|(b_{ij})|_p$. This is already illustrated by the simplest Gaussian examples: for example, if $b_{ii} = 1/\sqrt{\log i}$ and $b_{ij} = 0$ for $i \neq j$, then a standard Gaussian computation shows that $\mathbf{E}\|X\| \lesssim 1$, while all dimension-free bounds we have discussed grow at least as $\sqrt{\log n}$.

3.4. *Tail bounds.* Given explicit bounds on the expectation $\mathbf{E}\|X\|$, we can readily obtain nonasymptotic tail inequalities for $\|X\|$ by applying standard concentration techniques. In view of the significant utility of such tail inequalities in applications, we record some useful results along these lines here.

COROLLARY 3.9. *Under the assumptions of Theorem 1.1, we have*

$$\mathbf{P}\left[\|X\| \geq (1 + \varepsilon)\left\{2\sigma + \frac{6}{\sqrt{\log(1 + \varepsilon)}}\sigma_*\sqrt{\log n}\right\} + t\right] \leq e^{-t^2/4\sigma_*^2}$$

for any $0 < \varepsilon \leq 1/2$ and $t \geq 0$. In particular, for every $0 < \varepsilon \leq 1/2$ there exists a universal constant c_ε such that for every $t \geq 0$

$$\mathbf{P}[\|X\| \geq (1 + \varepsilon)2\sigma + t] \leq ne^{-t^2/c_\varepsilon\sigma_*^2}.$$

PROOF. As $\|X\| = \sup_v |\langle v, Xv \rangle|$ (the supremum is over the unit ball) and

$$\mathbf{E}[\langle v, Xv \rangle^2] = \sum_i b_{ii}^2 v_i^4 + 2 \sum_{i \neq j} b_{ij}^2 v_i^2 v_j^2 \leq 2\sigma_*^2,$$

the first inequality follows from Gaussian concentration [7], Theorem 5.8, and Theorem 1.1. For the second inequality, note that we can estimate

$$\begin{aligned} \mathbf{P}[\|X\| \geq (1 + \varepsilon)2\sigma + c_\varepsilon \sigma_* t] &\leq \mathbf{P}[\|X\| \geq (1 + \varepsilon)2\sigma + c'_\varepsilon \sigma_* \sqrt{\log n} + \sigma_* t] \\ &\leq e^{-t^2/4} \end{aligned}$$

for $t \geq 2\sqrt{\log n}$ (with $c_\varepsilon, c'_\varepsilon$ chosen in the obvious manner), while

$$\mathbf{P}[\|X\| \geq (1 + \varepsilon)2\sigma + c_\varepsilon \sigma_* t] \leq 1 \leq ne^{-t^2/4}$$

for $t \leq 2\sqrt{\log n}$. Combining these bounds completes the proof. \square

Tail bounds on $\|X\|$ have appeared widely in the recent literature under the name ‘‘matrix concentration inequalities’’; see [15, 24]. In the present setting, the corresponding result of this kind implies that for all $t \geq 0$

$$\mathbf{P}[\|X\| \geq t] \leq ne^{-t^2/8\sigma^2}.$$

The second inequality of Corollary 3.9 was stated for comparison with this matrix concentration bound. Unlike the matrix concentration bound, Corollary 3.9 is essentially optimal in that it captures not only the correct mean, but also the correct tail behavior of $\|X\|$ [13], Corollary 3.2 (Note that due to the factor $1 + \varepsilon$ in the leading term, we do not expect to see Tracy–Widom fluctuations at this scale.)

REMARK 3.10. Integrating the tail bound obtained by the matrix concentration method yields the estimate $\mathbf{E}\|X\| \lesssim \sigma\sqrt{\log n}$. This method therefore yields an alternative proof of the noncommutative Khintchine bound that was discussed in the Introduction. Combining this bound with concentration as in the proof of Corollary 3.9 already yields a better tail bound than the one obtained directly from the matrix concentration method. Nonetheless, it should be emphasized that the suboptimality of the above bound on the expected norm stems from the suboptimal tail behavior obtained by the matrix concentration method. Our sharp tail bounds help clarify the source of this inefficiency: the parameter σ should only control the mean of $\|X\|$, while the fluctuations are controlled entirely by σ_* .

An entirely analogous result can be obtained in the rectangular setting of Theorem 3.1. As the proof is identical, we simply state the result.

COROLLARY 3.11. *Under the assumptions of Theorem 3.1, we have*

$$\mathbf{P}\left[\|X\| \geq (1 + \varepsilon)\left\{\sigma_1 + \sigma_2 + \frac{5}{\sqrt{\log(1 + \varepsilon)}}\sigma_*\sqrt{\log(n \wedge m)}\right\} + t\right] \leq e^{-t^2/2\sigma_*^2}$$

for any $0 < \varepsilon \leq 1/2$ and $t \geq 0$. In particular, for every $0 < \varepsilon \leq 1/2$ there exists a universal constant c'_ε such that for every $t \geq 0$

$$\mathbf{P}[\|X\| \geq (1 + \varepsilon)(\sigma_1 + \sigma_2) + t] \leq (n \wedge m)e^{-t^2/c'_\varepsilon\sigma_*^2}.$$

The Gaussian concentration property used above is specific to Gaussian variables. However, there are many other situations where strong concentration results are available [7], and where similar results can be obtained. For example, if the Gaussian variables g_{ij} are replaced by symmetric random variables ξ_{ij} with $\|\xi_{ij}\|_\infty \leq 1$ (this captures in particular the case of Rademacher variables), Corollaries 3.9 and 3.11 remain valid with slightly larger universal constants $c_\varepsilon, c'_\varepsilon$. This follows from the identical proof, up to the replacement of Gaussian concentration by a form of Talagrand's concentration inequality [7], Theorem 6.10.

In the case of bounded entries, however, a more interesting question is whether it is possible to obtain tail bounds that capture the variance of the entries rather than their uniform norm (which is often much bigger than the variance), akin to the classical Bernstein inequality for sums of independent random variables. We presently develop a very useful result along these lines.

COROLLARY 3.12. *Let X be an $n \times n$ symmetric matrix whose entries X_{ij} are independent symmetric random variables. Then there exists for any $0 < \varepsilon \leq 1/2$ a universal constant \tilde{c}_ε such that for every $t \geq 0$*

$$\mathbf{P}[\|X\| \geq (1 + \varepsilon)2\tilde{\sigma} + t] \leq ne^{-t^2/\tilde{c}_\varepsilon\tilde{\sigma}_*^2},$$

where we have defined

$$\tilde{\sigma} := \max_i \sqrt{\sum_j \mathbf{E}[X_{ij}^2]}, \quad \tilde{\sigma}_* := \max_{ij} \|X_{ij}\|_\infty.$$

PROOF. Let $X_{ij} = \tilde{X}_{ij}\mathbf{E}[X_{ij}^2]^{1/2}$, so that \tilde{X}_{ij} have unit variance. Then

$$\mathbf{E}\|X\| \leq (1 + \varepsilon)2\tilde{\sigma} + C_\varepsilon\tilde{\sigma}_*\sqrt{\log n}$$

for a suitable constant C_ε by Corollary 3.6. On the other hand, a form of Talagrand's concentration inequality [7], Theorem 6.10, yields

$$\mathbf{P}[\|X\| \geq \mathbf{E}\|X\| + t] \leq e^{-t^2/c\tilde{\sigma}_*^2}$$

for all $t \geq 0$, where c is a universal constant. The proof is completed by combining these bounds as in the proof of Corollary 3.9. \square

Corollary 3.12 should be compared with the matrix Bernstein inequality in [24], which reads as follows in our setting (we omit the explicit constants):

$$\mathbf{P}[\|X\| \geq t] \leq ne^{-t^2/c(\tilde{\sigma}^2 + \tilde{\sigma}_*t)}.$$

While this result looks quite different at first sight than Corollary 3.12, the latter yields strictly better tail behavior up to universal constants: indeed, note that

$$e^{-t^2/c^2\tilde{\sigma}_*^2} \leq e^{1-2t/c\tilde{\sigma}_*} \leq 3e^{-2t^2/c(\tilde{\sigma}^2 + \tilde{\sigma}_*t)}$$

using $2x - 1 \leq x^2$. The discrepancy between these results is readily explained. In our sharp bounds, the variance term $\tilde{\sigma}$ only appears in the mean of $\|X\|$ and not in the fluctuations: the latter only depend on the uniform parameter $\tilde{\sigma}_*$ and do not capture the variance. A tail bound in terms of $\tilde{\sigma}$ and $\tilde{\sigma}_*$ should therefore indeed be of Hoeffding type, as in Corollary 3.12, rather than of Bernstein type as might be expected from the “matrix concentration” approach. Using a Bernstein form of Talagrand’s concentration inequality [14], Theorem 3, in the proof of Corollary 3.12 does not lead to any further improvement in the present setting.

REMARK 3.13. If X_{ij} in Corollary 3.12 are only assumed to be centered (rather than symmetric), we can symmetrize as in the proof of Corollary 3.3 to obtain

$$\mathbf{P}[\|X\| \geq (1 + \varepsilon)2\sqrt{2}\tilde{\sigma} + t] \leq ne^{-t^2/\tilde{c}_\varepsilon\tilde{\sigma}_*^2}.$$

Unfortunately, this results in an additional factor $\sqrt{2}$ in the leading term, which is suboptimal for Wigner matrices. We do not know whether it is possible, in general, to improve the constant when the entries are not symmetrically distributed.

Corollary 3.12 (and Corollary 3.6 which is used in its proof) also admit direct analogues in the setting of rectangular matrices. As the proofs are essentially identical to the ones given above, we leave such extensions to the reader.

3.5. *Lower bounds.* The main results of this paper provide upper bounds on $\mathbf{E}\|X\|$. However, a trivial lower bound already suffices to establish the sharpness of our upper bounds in many cases of interest, at least for Gaussian variables.

LEMMA 3.14. *In the setting of Theorem 1.1, we have*

$$\mathbf{E}\|X\| \gtrsim \sigma + \mathbf{E} \max_{ij} |b_{ij}g_{ij}|.$$

Similarly, in the setting of Theorem 3.1

$$\mathbf{E}\|X\| \gtrsim \sigma_1 + \sigma_2 + \mathbf{E} \max_{ij} |b_{ij}g_{ij}|.$$

PROOF. Let us prove the second inequality; the first inequality follows in a completely analogous manner. As $\|X\| \geq \max_{ij} |X_{ij}|$, it is trivial that

$$\mathbf{E}\|X\| \geq \mathbf{E} \max_{ij} |X_{ij}| = \mathbf{E} \max_{ij} |b_{ij}g_{ij}|.$$

On the other hand, as $\|X\| \geq \max_i \|Xe_i\|$ ($\{e_i\}$ is the canonical basis in \mathbb{R}^n),

$$\mathbf{E}\|X\| \geq \max_i \mathbf{E}\|Xe_i\| \gtrsim \max_i \mathbf{E}[\|Xe_i\|^2]^{1/2} = \sigma_2.$$

Here we use the estimate

$$\mathbf{E}[\|Xe_i\|^2] = (\mathbf{E}\|Xe_i\|)^2 + \text{Var}\|Xe_i\| \lesssim (\mathbf{E}\|Xe_i\|)^2,$$

where $\text{Var}\|Xe_i\| \leq \max_j b_{ji}^2 \lesssim \max_j \mathbf{E}[b_{ji}g_{ji}]^2 \leq (\mathbf{E}\|Xe_i\|)^2$ by the Gaussian Poincaré inequality [7], Theorem 3.20. Analogously, we obtain

$$\mathbf{E}\|X\| \geq \max_i \mathbf{E}\|X^*e_i\| \gtrsim \sigma_1.$$

Averaging these three lower bounds yields the conclusion. \square

This simple bound shows that our main results are sharp as long there are enough large coefficients b_{ij} . This is the content of the following easy bound.

COROLLARY 3.15. *In the setting of Theorem 1.1, suppose that*

$$|\{ij : |b_{ij}| \geq c\sigma_*\}| \geq n^\alpha$$

for some constants $c, \alpha > 0$. Then

$$\mathbf{E}\|X\| \asymp \sigma + \sigma_* \sqrt{\log n},$$

where the universal constant in the lower bound depends on c, α only. The analogous result holds in the setting of Theorem 3.1.

PROOF. Denote by I the set of indices in the statement of the corollary. Then

$$\mathbf{E} \max_{ij} |b_{ij}g_{ij}| \geq \mathbf{E} \max_{ij \in I} |b_{ij}g_{ij}| \geq c\sigma_* \mathbf{E} \max_{ij \in I} |g_{ij}| \gtrsim \sigma_* \sqrt{\log |I|} \gtrsim \sigma_* \sqrt{\log n},$$

where we use a standard lower bound on the maximum of independent $N(0,1)$ random variables. The proof is completed by applying Lemma 3.14. \square

For example, it follows that our main results are sharp as soon as every row of the matrix contains at least one large coefficient, that is, with magnitude of the same order as σ_* . This is the case in many natural examples of interest, and in these cases our results are optimal (up to the values of universal

constants). Of course, it quite possible that our bound is sharp even when the assumption of Corollary 3.15 fails: for example, in view of Lemma 3.14, our bound is sharp whenever $\sigma_*\sqrt{\log n} \lesssim \sigma$ regardless of any other feature of the problem.

Corollary 3.15 suggests that our main results can fail to be sharp when the sizes of the coefficients b_{ij} are very heterogeneous. If the matrix contains a few large entries and many small entries, one could still obtain good bounds by splitting the matrix into two parts and applying Theorem 1.1 to each part; this is the idea behind the dimension-free bounds of Corollary 3.7. However, when there are many different scales with few coefficients at each scale, such an approach cannot be expected to yield sharp results in general; see Remark 3.8 for a simple example.

REMARK 3.16. An intriguing observation that was made in [17] is that the trivial lower bound $\mathbf{E}\|X\| \geq \mathbf{E}\max_i \|Xe_i\|$ appears to be surprisingly sharp: we do not know of any example where the corresponding upper bound

$$\mathbf{E}\|X\| \stackrel{?}{\lesssim} \mathbf{E}\max_i \|Xe_i\|$$

fails. If such an inequality were to hold, the conclusion of Theorem 1.1 would follow easily using Gaussian concentration and a simple union bound. In fact, if this were the case, we could obtain an improvement of Theorem 1.1 in the following manner; cf. [22], Proposition 2.4.16. Note that, by Gaussian concentration,

$$\mathbf{P}\left[\max_i \{\|Xe_i\| - \mathbf{E}\|Xe_i\|\} > t\right] \leq \sum_k e^{-t^2/2\max_j b_{kj}^2} = \sum_k k^{-t^2/2\max_{ij} b_{ij}^2 \log i}.$$

Integrating this bound therefore gives

$$\begin{aligned} \mathbf{E}\|X\| &\stackrel{?}{\lesssim} \mathbf{E}\max_i \|Xe_i\| \\ &\leq \max_i \mathbf{E}\|Xe_i\| + \mathbf{E}\max_i \{\|Xe_i\| - \mathbf{E}\|Xe_i\|\} \\ &\lesssim \sigma + \max_{ij} |b_{ij}| \sqrt{\log i} \end{aligned}$$

which would yield a strict improvement on Theorem 1.1. [Note that there is no loss of generality in sorting the rows of (b_{ij}) to minimize the term $\max_{ij} |b_{ij}| \sqrt{\log i}$.] We do not know of any mechanism, however, that would give rise to such inequalities, and it is possible that the apparent sharpness of the quantity $\mathbf{E}\max_i \|Xe_i\|$ is simply due to the fact that it is of the same order as the bound of Theorem 1.1 in most natural examples. Regardless, it does not appear that our method of proof could be adapted to give rise to inequalities of this form.

REMARK 3.17. The conclusion of Corollary 3.15 relies heavily on the Gaussian nature of the entries. When the distributions of the entries are bounded, for example, it is possible that our bounds are no longer sharp. This issue will be discussed further in Section 4.2 below in the context of Rademacher matrices.

4. Examples.

4.1. *Sparse random matrices.* In the section, we consider the special case of Theorem 1.1 where the coefficients b_{ij} can take the values zero or one only. This is in essence a sparse counterpart of Wigner matrices in which a subset of the entries has been set to zero. This rather general model covers many interesting random matrix ensembles, including the case of random band matrices where $b_{ij} = \mathbf{1}_{|i-j| \leq k}$ that has been of significant recent interest [4, 11, 21].

Let us fix a matrix (b_{ij}) of $\{0, 1\}$ -valued coefficients. We immediately compute

$$\sigma^2 = k, \quad \sigma_* = 1,$$

where k is the maximal number of nonzero entries in any row of the matrix of coefficients (b_{ij}) . If we interpret (b_{ij}) as the adjacency matrix of a graph on n points, then k is simply the maximal degree of this graph. The following conclusion follows effortlessly from our main results.

COROLLARY 4.1. *Let X be the $n \times n$ symmetric random matrix with $X_{ij} = g_{ij}b_{ij}$, where $\{g_{ij}\}$ are independent $N(0, 1)$ variables and $b_{ij} \in \{0, 1\}$. Let k be the maximal number of nonzero entries in a row of (b_{ij}) . Then*

$$\mathbf{E}\|X\| \asymp \sqrt{k} + \sqrt{\log n},$$

provided that every row of (b_{ij}) has at least one nonzero entry.

PROOF. This is immediate from Theorem 1.1 and Lemma 3.14. \square

REMARK 4.2. If a row of (b_{ij}) is zero, then the corresponding column is zero as well by symmetry. We therefore lose nothing by removing this row and column, and we can apply Corollary 4.1 to the resulting lower-dimensional matrix. The assumption that every row of (b_{ij}) is nonzero is therefore completely innocuous.

Our bound evidently captures precisely the correct order of magnitude of the spectral norm of sparse random matrices. It is possible to obtain much sharper conclusions, however, from our main results. To motivate this, let us first quote a result that is stated in [4], Theorem 2.3, under weaker assumptions. (For simplicity, we adopt in the remainder of this section the setting and notation of Corollary 4.1.)

THEOREM 4.3. *Suppose each row of (b_{ij}) has exactly k nonzero entries. Then the empirical spectral distribution of X/\sqrt{k} converges to the semicircle law*

$$\frac{1}{n} \sum_{i=1}^n \delta_{\lambda_i(X/\sqrt{k})} \xrightarrow{k \rightarrow \infty} \frac{1}{2\pi} \sqrt{4-x^2} \mathbf{1}_{x \in [-2,2]} dx,$$

provided that $k = o(n)$. [Here $\lambda_1(X) \geq \dots \geq \lambda_n(X)$ are the eigenvalues of X .]

Theorem 4.3 shows that the bulk of the spectrum of X behaves precisely like that of a Wigner matrix under minimal assumptions. As the semicircle distribution has support $[-2, 2]$, one might assume that edge of the spectrum will converge to 2. That this is the case for Wigner matrices is a textbook result [1]. In the present case, however, we obtain a phase transition phenomenon.

COROLLARY 4.4. *Suppose that each row of (b_{ij}) has exactly k nonzero entries. Then the following phase transition occurs as $n \rightarrow \infty$:*

- If $k/\log n \rightarrow \infty$, then $\|X\|/\sqrt{k} \rightarrow 2$ in probability.
- If $k/\log n \rightarrow 0$, then $\|X\|/\sqrt{k} \rightarrow \infty$ in probability.
- If $k \sim \log n$, then $\{\|X\|/\sqrt{k}\}$ is bounded but may not converge to 2.

PROOF. If $k/\log n \rightarrow 0$, then $\|X\|/\sqrt{k} \geq \max_{ij} |g_{ij}b_{ij}|/\sqrt{k}$. As each row has a nonzero entry, the maximum is taken over at least $n/2$ independent $N(0, 1)$ random variables which is of order $\sqrt{2 \log(n/2)}$ as $n \rightarrow \infty$. Thus $\|X\|/\sqrt{k}$ diverges.

For $k/\log n \rightarrow \infty$, we note that Corollary 3.9 yields

$$\mathbf{P}[\|X\|/\sqrt{k} \geq 2 + \varepsilon] \leq ne^{-C_\varepsilon k} = n^{1-C_\varepsilon k/\log n}$$

for a suitable constant C_ε . Thus $\|X\|/\sqrt{k} \leq 2 + \varepsilon + o(1)$ for any $\varepsilon > 0$. On the other hand, Theorem 4.3 implies that $\|X\|/\sqrt{k} \geq 2 - \varepsilon - o(1)$ for any $\varepsilon > 0$.

If $k = a \log n$, Corollary 3.9 similarly yields that $\mathbf{P}[\|X\|/\sqrt{k} > C] \rightarrow 0$ for a sufficiently large constant C , so $\{\|X\|/\sqrt{k}\}$ is bounded. However, Lemma 3.14 shows that $\mathbf{E}\|X\|/\sqrt{k} \gtrsim a^{-1/2} > 3$ when a is sufficiently small. \square

REMARK 4.5. We have not investigated the precise behavior of $\|X\|/\sqrt{k}$ for the boundary case $k = a \log n$. The proof of Corollary 4.4 shows that if a is chosen sufficiently small, the rescaled norm remains bounded but strictly separated from the bulk as $n \rightarrow \infty$. We do not know whether this is the

case for all a , or whether the norm does in fact converge to 2 when a is sufficiently large. A precise investigation of this question is beyond the scope of this paper.

In the special case of band matrices, Corollary 4.4 was proved by Sodin [21] following an earlier suboptimal result of Khorunzhiy [11]. However, his combinatorial proof relies on the specific positions of the nonzero entries of (b_{ij}) . In a recent paper, Benaych-Georges and P ech e [4] showed in the general setting (i.e., without assuming specific positions of the entries) that $\|X\|/\sqrt{k} \rightarrow 2$ when $k/\log^9 n \rightarrow \infty$. To the best of our knowledge, however, the result of Corollary 4.4 is new. While this result is of independent interest, we particularly emphasize how effortlessly a sharp conclusion could be derived from the main results of this paper.

Beyond the Gaussian case, Corollaries 3.5 and 3.6 can be used to obtain similar results in the presence of heavy-tailed entries. Using Corollary 3.5, it can be shown that $\|X\|/\sqrt{k}$ remains bounded if and only if $\log^{\beta/2} n = O(k)$ when the entries ξ_{ij} have moments of order $\mathbf{E}[\xi_{ij}^{2p}]^{1/2p} \sim p^{\beta/2}$ with $\beta \geq 1$. This establishes the appropriate phase transition point in the heavy-tailed setting; however, we cannot conclude convergence to the edge of the semicircle due to the additional universal constant in Corollary 3.5. On the other hand, using Corollary 3.6 we can establish convergence to the edge of the semicircle under an assumption on the rate of growth of k that is suboptimal by a logarithmic factor; this is comparable to the results in [4], though we obtain a somewhat better scaling. The details are omitted.

4.2. Rademacher matrices. We have seen that our main results provide sharp bounds in many cases on the norm of matrices with independent Gaussian entries. While our upper bounds continue to hold for sub-Gaussian variables, this is not the case for the lower bounds in Section 3.5, and in this case we cannot expect our results to be sharp at the same level of generality. As a simple example, consider the case where X is the diagonal matrix with i.i.d. entries on the diagonal. If the entries are Gaussian, then $\|X\| \gtrsim \sqrt{\log n}$, so that Theorem 1.1 is sharp. If the entries are bounded, however, then $\|X\| \lesssim 1$. On the other hand, the universality property of Wigner matrices shows that Theorem 1.1 is sharp in this case even when adapted to bounded random variables (Corollary 3.3).

In view of these observations, it is natural to ask whether it is possible to obtain systematic improvement of our main results that captures the size of the norm of random matrices with bounded entries. For concreteness, let us consider the case of Rademacher matrices $X_{ij} = \varepsilon_{ij} b_{ij}$, where $\{\varepsilon_{ij}\}$ are independent Rademacher (symmetric Bernoulli) random variables. In this setting, we can immediately obtain a trivial but useful improvement on

Corollary 3.3. (In the rest of this section, we will consider symmetric matrices and universal constants for simplicity; analogous results for rectangular matrices or with explicit constants are easily obtained.)

COROLLARY 4.6. *Let X be the $n \times n$ symmetric random matrix with $X_{ij} = \varepsilon_{ij} b_{ij}$, where $\{\varepsilon_{ij}\}$ are independent Rademacher variables. Then*

$$\mathbf{E}\|X\| \lesssim (\sigma + \sigma_* \sqrt{\log n}) \wedge \|B\|,$$

where $B := (|b_{ij}|)$ is the matrix of absolute values of the coefficients.

PROOF. In view of Corollary 3.2, it suffices to show that $\mathbf{E}\|X\| \leq \|B\|$. Note, however, that this inequality even holds pointwise: indeed,

$$\|X\| = \sup_v \sum_{ij} \varepsilon_{ij} b_{ij} v_i v_j \leq \sup_v \sum_{ij} |b_{ij} v_i v_j| = \|B\|,$$

where the supremum is taken over the unit ball in \mathbb{R}^n . \square

Corollary 4.6 captures two reasons why a Rademacher matrix can have small norm: either it behaves like a Gaussian matrix with small norm, or its norm is uniformly bounded due to the boundedness of the matrix entries. This idea mirrors the basic ingredients in the general theory of Bernoulli processes [22], Chapter 5. While simple, Corollary 4.6 captures at least the Wigner and diagonal examples considered above, albeit in a somewhat ad-hoc manner. We will presently show that a less trivial result can be easily derived from Corollary 4.6 as well.

The norm of Rademacher matrices was first investigated in a general setting by Seginer [19]. Using a delicate combinatorial method, he proves in this case that $\mathbf{E}\|X\| \lesssim \sigma \log^{1/4} n$. The assumption of Rademacher entries is essential: that such a bound cannot hold in the Gaussian case is immediate from the diagonal matrix example. Let us show that this result is an easy consequence of Corollary 4.6.

COROLLARY 4.7. *Let X be the $n \times n$ symmetric random matrix with $X_{ij} = \varepsilon_{ij} b_{ij}$, where $\{\varepsilon_{ij}\}$ are independent Rademacher variables. Then*

$$\mathbf{E}\|X\| \lesssim \sigma \log^{1/4} n.$$

PROOF. Fix $u > 0$. Let us split the matrix into two parts $X = X^+ + X^-$, where $X_{ij}^+ = \varepsilon_{ij} b_{ij} \mathbf{1}_{|b_{ij}| > u}$ and $X_{ij}^- = \varepsilon_{ij} b_{ij} \mathbf{1}_{|b_{ij}| \leq u}$. For X^- , we can estimate

$$\mathbf{E}\|X^-\| \lesssim \sigma + u \sqrt{\log n}.$$

On the other hand, we estimate for X^+ by the Gershgorin circle theorem

$$\mathbf{E}\|X^+\| \leq \|(|b_{ij}|\mathbf{1}_{|b_{ij}|>u})\| \leq \max_i \sum_j |b_{ij}|\mathbf{1}_{|b_{ij}|>u} \leq \frac{\sigma^2}{u}.$$

We therefore obtain for any $u > 0$

$$\mathbf{E}\|X\| \lesssim \sigma + u\sqrt{\log n} + \frac{\sigma^2}{u}.$$

The proof is completed by optimizing over $u > 0$. \square

Corollary 4.7 not only recovers Seginer's result with a much simpler proof, but also effectively explains why the mysterious term $\log^{1/4} n$ arises. More generally, the method of proof suggests how Corollary 4.6 can be used efficiently: we should attempt to split the matrix X into two parts, such that one part is small by Theorem 1.1, and the other part is small uniformly. This idea also arises in a fundamental manner in the general theory of Bernoulli processes [22]. Unfortunately, it is generally not clear for a given matrix how to choose the best decomposition.

REMARK 4.8. In view of Corollary 4.1, one might hope that Corollary 4.6 (or a suitable adaptation of this bound) could yield sharp results in the general setting of sparse random matrices. The situation for Rademacher matrices turns out to be more delicate, however. To see this, let us consider two illuminating examples. In the following, let $k = \lceil \sqrt{\log n} \rceil$, and assume for simplicity that n/k is integer.

First, consider the block-diagonal matrix X of the form

$$X = \begin{bmatrix} X_1 & & & & \\ & X_2 & & 0 & \\ & & \cdot & & \\ & 0 & & \cdot & \\ & & & & X_{n/k} \end{bmatrix},$$

where each X_i is a $k \times k$ symmetric matrix with independent Rademacher entries. Such matrices are considered by Seginer in [19], who shows by an elementary argument that $\mathbf{E}\|X\| \sim \sqrt{\log n}$. Thus Theorem 1.1 already yields a sharp result (and, in particular, the logarithmic term in Theorem 1.1 cannot be eliminated).

On the other hand, Sodin [20] shows that if X is the Rademacher matrix where the coefficient matrix B is chosen to be a realization of the adjacency matrix of a random k -regular graph, then $\mathbf{E}\|X\| \sim \sqrt{k} \leq \log^{1/4} n$ with high probability. Thus in this case $\mathbf{E}\|X\| \sim \sigma$, and it appears that the logarithmic term in Theorem 1.1 is missing (evidently none of our bounds are sharp in this case).

Note, however, that in both these examples the parameters $\sigma, \sigma_*, \|B\|$ are identical: we have $\sigma = \sqrt{k}$, $\sigma_* = 1$, and $\|B\| = k$ (by the Perron–Frobenius theorem). In particular, there is no hope that the norm of sparse Rademacher matrices can be controlled using only the degree of the graph: the *structure* of the graph must come into play. It is an interesting open problem to understand precisely what aspect of this structure controls the norm of sparse Rademacher matrices. This question is closely connected to the study of random 2-lifts of graphs in combinatorics [5].

Acknowledgments. In April 2014, one of us (ASB) attended the workshop “Mathematical Physics meets Sparse Recovery” at Oberwolfach where the question of removing the logarithmic term in the noncommutative Khintchine inequality was raised by Joel Tropp in a more general context. While we ultimately took a very different approach, we thank Joel for motivating our interest in such problems which led us to develop the ideas reported here. We are grateful to Sasha Sodin, Nikhil Srivastava and Shahar Mendelson for several interesting discussions and for providing us with a number of helpful references. Afonso S. Bandeira would also like to thank Oberwolfach for its hospitality during the above-mentioned workshop.

REFERENCES

- [1] ANDERSON, G. W., GUIONNET, A. and ZEITOUNI, O. (2010). *An Introduction to Random Matrices. Cambridge Studies in Advanced Mathematics* **118**. Cambridge Univ. Press, Cambridge. [MR2760897](#)
- [2] AUBRUN, G. (2007). Sampling convex bodies: A random matrix approach. *Proc. Amer. Math. Soc.* **135** 1293–1303 (electronic). [MR2276637](#)
- [3] BAI, Z. D. and YIN, Y. Q. (1988). Necessary and sufficient conditions for almost sure convergence of the largest eigenvalue of a Wigner matrix. *Ann. Probab.* **16** 1729–1741. [MR0958213](#)
- [4] BENAYCH-GEORGES, F. and PÉCHÉ, S. (2014). Largest eigenvalues and eigenvectors of band or sparse random matrices. *Electron. Commun. Probab.* **19** no. 4, 9. [MR3164751](#)
- [5] BILU, Y. and LINIAL, N. (2006). Lifts, discrepancy and nearly optimal spectral gap. *Combinatorica* **26** 495–519. [MR2279667](#)
- [6] BOUCHERON, S., BOUSQUET, O., LUGOSI, G. and MASSART, P. (2005). Moment inequalities for functions of independent random variables. *Ann. Probab.* **33** 514–560. [MR2123200](#)
- [7] BOUCHERON, S., LUGOSI, G. and MASSART, P. (2013). *Concentration Inequalities: A Nonasymptotic Theory of Independence, With a Foreword by Michel Ledoux*. Oxford Univ. Press, Oxford. [MR3185193](#)
- [8] DAVIDSON, K. R. and SZAREK, S. J. (2001). Local operator theory, random matrices and Banach spaces. In *Handbook of the Geometry of Banach Spaces, Vol. I* 317–366. North-Holland, Amsterdam. [MR1863696](#)
- [9] FÜREDI, Z. and KOMLÓS, J. (1981). The eigenvalues of random symmetric matrices. *Combinatorica* **1** 233–241. [MR0637828](#)

- [10] GEMAN, S. (1980). A limit theorem for the norm of random matrices. *Ann. Probab.* **8** 252–261. [MR0566592](#)
- [11] KHORUNZHIY, O. (2008). Estimates for moments of random matrices with Gaussian elements. In *Séminaire de Probabilités XLI. Lecture Notes in Math.* **1934** 51–92. Springer, Berlin. [MR2483726](#)
- [12] LATAŁA, R. (2005). Some estimates of norms of random matrices. *Proc. Amer. Math. Soc.* **133** 1273–1282 (electronic). [MR2111932](#)
- [13] LEDOUX, M. and TALAGRAND, M. (1991). *Probability in Banach Spaces: Isoperimetry and Processes. Ergebnisse der Mathematik und Ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)]* **23**. Springer, Berlin. [MR1102015](#)
- [14] MASSART, P. (2000). About the constants in Talagrand’s concentration inequalities for empirical processes. *Ann. Probab.* **28** 863–884. [MR1782276](#)
- [15] OLIVEIRA, R. I. (2010). Sums of random Hermitian matrices and an inequality by Rudelson. *Electron. Commun. Probab.* **15** 203–212. [MR2653725](#)
- [16] PISIER, G. (2003). *Introduction to Operator Space Theory. London Mathematical Society Lecture Note Series* **294**. Cambridge Univ. Press, Cambridge. [MR2006539](#)
- [17] RIEMER, S. and SCHÜTT, C. (2013). On the expectation of the norm of random matrices with non-identically distributed entries. *Electron. J. Probab.* **18** no. 29, 13. [MR3035757](#)
- [18] RUDELSON, M. and VERSHYNIN, R. (2010). Non-asymptotic theory of random matrices: Extreme singular values. In *Proceedings of the International Congress of Mathematicians. Volume III* 1576–1602. Hindustan Book Agency, New Delhi. [MR2827856](#)
- [19] SEGNER, Y. (2000). The expected norm of random matrices. *Combin. Probab. Comput.* **9** 149–166. [MR1762786](#)
- [20] SODIN, S. (2009). The Tracy–Widom law for some sparse random matrices. *J. Stat. Phys.* **136** 834–841. [MR2545551](#)
- [21] SODIN, S. (2010). The spectral edge of some random band matrices. *Ann. of Math.* (2) **172** 2223–2251. [MR2726110](#)
- [22] TALAGRAND, M. (2014). *Upper and Lower Bounds for Stochastic Processes: Modern Methods and Classical Problems. Ergebnisse der Mathematik und Ihrer Grenzgebiete. 3. Folge. A Series of Modern Surveys in Mathematics [Results in Mathematics and Related Areas. 3rd Series. A Series of Modern Surveys in Mathematics]* **60**. Springer, Heidelberg. [MR3184689](#)
- [23] TAO, T. (2012). *Topics in Random Matrix Theory. Graduate Studies in Mathematics* **132**. Amer. Math. Soc., Providence, RI. [MR2906465](#)
- [24] TROPP, J. A. (2012). User-friendly tail bounds for sums of random matrices. *Found. Comput. Math.* **12** 389–434. [MR2946459](#)
- [25] VERSHYNIN, R. (2012). Introduction to the non-asymptotic analysis of random matrices. In *Compressed Sensing* 210–268. Cambridge Univ. Press, Cambridge. [MR2963170](#)

PROGRAM IN APPLIED
AND COMPUTATIONAL MATHEMATICS
PRINCETON UNIVERSITY
PRINCETON, NEW JERSEY 08544
USA
E-MAIL: ajsb@math.princeton.edu

DEPARTMENT OF OPERATIONS RESEARCH
AND FINANCIAL ENGINEERING
SHERRERD HALL 227
PRINCETON UNIVERSITY
PRINCETON, NEW JERSEY 08544
USA
E-MAIL: rvan@princeton.edu