

SECOND-ORDER MATRIX CONCENTRATION INEQUALITIES

JOEL A. TROPP

ABSTRACT. Matrix concentration inequalities give bounds for the spectral-norm deviation of a random matrix from its expected value. These results have a weak dimensional dependence that is sometimes, but not always, necessary. This paper identifies one of the sources of the dimensional term and exploits this insight to develop sharper matrix concentration inequalities. In particular, this analysis delivers two refinements of the matrix Khintchine inequality that use information beyond the matrix variance to reduce or eliminate the dimensional dependence.

1. MOTIVATION

Matrix concentration inequalities provide spectral information about a random matrix that depends smoothly on many independent random variables. In recent years, these results have become a dominant tool in applied random matrix theory. There are several reasons for the success of this approach.

- **Flexibility.** Matrix concentration applies to a wide range of random matrix models. In particular, we can obtain bounds for the spectral norm of a sum of independent random matrices in terms of the properties of the summands.
- **Ease of Use.** For many applications, matrix concentration tools require only a small amount of matrix analysis. No expertise in random matrix theory is required to invoke the results.
- **Power.** For a large class of examples, including independent sums, matrix concentration bounds are provably close to optimal.

See the monograph [Tro15] for an overview of this theory and a comprehensive bibliography.

The matrix concentration inequalities in the literature are suboptimal for certain examples because of a weak dependence on the dimension of the random matrix. Removing this dimensional term is difficult because there are many situations where it is necessary. The purpose of this paper is to identify one of the sources of the dimensional factor. Using this insight, we will develop some new matrix concentration inequalities that are qualitatively better than the current generation of results, although they sacrifice some of our desiderata. Ultimately, we hope that this line of research will lead to general tools for applied random matrix theory that are flexible, easy to use, and that give sharp results in most cases.

2. THE MATRIX KHINTCHINE INEQUALITY

To set the stage, we present and discuss the primordial matrix concentration result, the *matrix Khintchine inequality*, which describes the behavior of a special random matrix model, called a *matrix Gaussian series*. This result already exhibits the key features of more sophisticated matrix concentration inequalities, and it can be used to derive concentration bounds for more general models. As such, the matrix Khintchine inequality serves as a natural starting point for deeper investigations.

2.1. Matrix Gaussian Series. In this work, we focus on an important class of random matrices that has a lot of modeling power but still supports an interesting theory.

Definition 2.1 (Matrix Gaussian Series). Consider fixed Hermitian matrices $\mathbf{H}_1, \dots, \mathbf{H}_n$ with common dimension d , and let $\{\gamma_1, \dots, \gamma_n\}$ be an independent family of standard normal random variables. Construct the random matrix

$$\mathbf{X} := \sum_{i=1}^n \gamma_i \mathbf{H}_i. \tag{2.1}$$

Date: 13 March 2015. Revised 21 April 2015 and 3 August 2016.

2010 Mathematics Subject Classification. Primary: 60B20. Secondary: 60F10, 60G50, 60G42.

Key words and phrases. Concentration inequality, moment inequality, random matrix.

Email: jtropp@cms.caltech.edu. Tel: 626.395.5957.

We refer to a random matrix with this form as a matrix Gaussian series with Hermitian coefficients or, for brevity, an *Hermitian matrix Gaussian series*.

Matrix Gaussian series enjoy a surprising amount of modeling power. It is easy to see that we can express any random Hermitian matrix with jointly Gaussian entries in the form (2.1). More generally, we can use matrix Gaussian series to analyze a sum of independent, zero-mean, random, Hermitian matrices $\mathbf{Y}_1, \dots, \mathbf{Y}_n$. Indeed, for any norm $\|\cdot\|$ on matrices,

$$\mathbb{E} \left\| \sum_{i=1}^n \mathbf{Y}_i \right\| \leq \sqrt{2\pi} \cdot \mathbb{E} \left[\mathbb{E} \left[\left\| \sum_{i=1}^n \gamma_i \mathbf{Y}_i \right\| \mid \mathbf{Y}_1, \dots, \mathbf{Y}_n \right] \right]. \quad (2.2)$$

The process of passing from an independent sum to a conditional Gaussian series is called *symmetrization*. See [LT91, Lem. 6.3 and Eqn. (4.8)] for details about this calculation. Furthermore, some techniques for Gaussian series can be adapted to study independent sums directly without the artifice of symmetrization.

Note that our restriction to Hermitian matrices is not really a limitation. We can also analyze a rectangular matrix \mathbf{Z} with jointly Gaussian entries by working with the Hermitian dilation of \mathbf{Z} , sometimes known as the Jordan–Wielandt matrix. See [Tro15, Sec. 2.1.16] for more information on this approach.

2.2. The Matrix Variance. Many matrix concentration inequalities are expressed most naturally in terms of a matrix extension of the variance.

Definition 2.2 (Matrix Variance). Let \mathbf{X} be a random Hermitian matrix. The *matrix variance* is the deterministic matrix

$$\mathbf{Var}(\mathbf{X}) := \mathbb{E} \mathbf{X}^2 - (\mathbb{E} \mathbf{X})^2.$$

We use the convention that the power binds before the expectation.

In particular, consider a matrix Gaussian series $\mathbf{X} := \sum_{i=1}^n \gamma_i \mathbf{H}_i$. It is easy to verify that

$$\mathbf{Var}(\mathbf{X}) = \mathbb{E} \mathbf{X}^2 = \sum_{i,j=1}^n \mathbb{E}[\gamma_i \gamma_j] \cdot \mathbf{H}_i \mathbf{H}_j = \sum_{i=1}^n \mathbf{H}_i^2.$$

We see that the matrix variance has a clean expression in terms of the coefficients of the Gaussian series, so it is easy to compute in practice.

2.3. The Matrix Khintchine Inequality. The matrix Khintchine inequality is a fundamental fact about the behavior of matrix Gaussian series. The first version of this result was established by Lust-Piquard [LP86], and the constants were refined in the papers [Pis98, Buc01]. The version here is adapted from [MJC⁺14, Sec. 7.1].

Proposition 2.3 (Matrix Khintchine). *Consider an Hermitian matrix Gaussian series $\mathbf{X} := \sum_{i=1}^n \gamma_i \mathbf{H}_i$, as in (2.1). Introduce the matrix standard deviation parameter*

$$\sigma_q(\mathbf{X}) := \|\mathbf{Var}(\mathbf{X})^{1/2}\|_q = \left\| \left(\sum_{i=1}^n \mathbf{H}_i^2 \right)^{1/2} \right\|_q \quad \text{for } q \geq 1. \quad (2.3)$$

Then, for each integer $p \geq 1$,

$$\sigma_{2p}(\mathbf{X}) \leq \left(\mathbb{E} \|\mathbf{X}\|_{2p}^{2p} \right)^{1/(2p)} \leq \sqrt{2p-1} \cdot \sigma_{2p}(\mathbf{X}). \quad (2.4)$$

The symbol $\|\cdot\|_q$ denotes the Schatten q -norm.

The lower bound in (2.4) is simply Jensen's inequality. Section 7 contains a short proof of the upper bound.

The matrix Khintchine inequality also yields an estimate for the spectral norm of a matrix Gaussian series. This type of result is often more useful in practice.

Corollary 2.4 (Matrix Khintchine: Spectral Norm). *Consider an Hermitian matrix Gaussian series $\mathbf{X} := \sum_{i=1}^n \gamma_i \mathbf{H}_i$ with dimension d , as in (2.1). Introduce the matrix standard deviation parameter*

$$\sigma(\mathbf{X}) := \|\mathbf{Var}(\mathbf{X})\|^{1/2} = \left\| \sum_{i=1}^n \mathbf{H}_i^2 \right\|^{1/2}.$$

Then

$$\frac{1}{\sqrt{2}} \cdot \sigma(\mathbf{X}) \leq \mathbb{E} \|\mathbf{X}\| \leq \sqrt{e(1+2\log d)} \cdot \sigma(\mathbf{X}). \quad (2.5)$$

The symbol $\|\cdot\|$ denotes the spectral norm, also known as the ℓ_2 operator norm.

Proof Sketch. For the upper bound, observe that

$$\mathbb{E} \|\mathbf{X}\| \leq \left(\mathbb{E} \|\mathbf{X}\|_{2p}^{2p} \right)^{1/(2p)} \leq \sqrt{2p-1} \cdot \|\mathbf{Var}(\mathbf{X})^{1/2}\|_{2p} \leq d^{1/(2p)} \sqrt{2p-1} \cdot \|\mathbf{Var}(\mathbf{X})\|^{1/2}.$$

Indeed, the spectral norm is bounded above by the Schatten $2p$ -norm, and we can apply Lyapunov's inequality to increase the order of the moment from one to $2p$. Invoke Proposition 2.3, and bound the trace in terms of the spectral norm again. Finally, set $p = \lceil \log d \rceil$, and simplify the constants.

For the lower bound, note that

$$\mathbb{E} \|\mathbf{X}\| \geq \frac{1}{\sqrt{2}} \left(\mathbb{E} \left\| \sum_{i=1}^n \gamma_i \mathbf{H}_i \right\|^2 \right)^{1/2} = \frac{1}{\sqrt{2}} \left(\mathbb{E} \|\mathbf{X}^2\| \right)^{1/2} \geq \frac{1}{\sqrt{2}} \|\mathbf{Var}(\mathbf{X})\|^{1/2}.$$

The first relation follows from the optimal Khintchine–Kahane inequality [LO94]; the last is Jensen's. \square

2.4. Two Examples. The bound (2.5) shows that the matrix standard deviation controls the expected norm of a matrix Gaussian series up to a factor that is logarithmic in the dimension of the random matrix. One may wonder whether the lower branch or the upper branch of (2.5) gives the more accurate result. In fact, natural examples demonstrate that both extremes of behavior occur.

For an integer $d \geq 1$, define

$$\mathbf{X}_{\text{diag}} := \mathbf{X}_{\text{diag}}(d) := \begin{bmatrix} \gamma_1 & & & & \\ & \gamma_2 & & & \\ & & \gamma_3 & & \\ & & & \ddots & \\ & & & & \gamma_d \end{bmatrix}. \quad (2.6)$$

That is, \mathbf{X}_{diag} is a $d \times d$ diagonal matrix whose entries $\{\gamma_i : 1 \leq i \leq d\}$ are independent standard normal variables. Second, define

$$\mathbf{X}_{\text{goe}} := \mathbf{X}_{\text{goe}}(d) := \frac{1}{\sqrt{2d}} (\mathbf{G} + \mathbf{G}^*) \quad \text{where} \quad \mathbf{G} := \mathbf{G}(d) := \begin{bmatrix} \gamma_{11} & \gamma_{12} & \cdots & \gamma_{1d} \\ \gamma_{21} & \gamma_{22} & \cdots & \gamma_{2d} \\ \vdots & \vdots & \ddots & \vdots \\ \gamma_{d1} & \gamma_{d2} & \cdots & \gamma_{dd} \end{bmatrix} \quad (2.7)$$

The symbol $*$ denotes conjugate transposition. Up to scaling, the $d \times d$ random matrix \mathbf{X}_{goe} is the Hermitian part of a matrix \mathbf{G} whose entries $\{\gamma_{ij} : 1 \leq i, j \leq d\}$ are independent standard normal variables. The sequence $\{\mathbf{X}_{\text{goe}}(d) : d = 1, 2, 3, \dots\}$ is called the *Gaussian orthogonal ensemble* (GOE).

To apply the matrix Khintchine inequality, we represent each matrix as an Hermitian Gaussian series:

$$\mathbf{X}_{\text{diag}} = \sum_{i=1}^d \gamma_i \mathbf{E}_{ii} \quad \text{and} \quad \mathbf{X}_{\text{goe}} = \frac{1}{\sqrt{2d}} \sum_{1 \leq i, j \leq d} \gamma_{ij} (\mathbf{E}_{ij} + \mathbf{E}_{ji}).$$

We have written \mathbf{E}_{ij} for the $d \times d$ matrix with a one in the (i, j) position and zeros elsewhere. Respectively, the matrix variances satisfy

$$\mathbf{Var}(\mathbf{X}_{\text{diag}}) = \mathbf{I} \quad \text{and} \quad \mathbf{Var}(\mathbf{X}_{\text{goe}}) = (1 + d^{-1}) \cdot \mathbf{I}.$$

The bound (2.5) delivers

$$\frac{1}{\sqrt{2}} \leq \mathbb{E} \|\mathbf{X}\| \lesssim \sqrt{2e \log d} \quad \text{for } \mathbf{X} = \mathbf{X}_{\text{diag}} \text{ or } \mathbf{X} = \mathbf{X}_{\text{goe}}. \quad (2.8)$$

The relations \lesssim and \approx suppress lower-order terms. In each case, the ratio between the lower and upper bound has order $\sqrt{\log d}$. The matrix Khintchine inequality does not provide more precise information.

On the other hand, for these examples, detailed spectral information is available:

$$\mathbb{E} \|\mathbf{X}_{\text{goe}}\| \approx 2 \quad \text{and} \quad \mathbb{E} \|\mathbf{X}_{\text{diag}}\| \approx \sqrt{2 \log d}. \quad (2.9)$$

See [Tao12, Sec. 2.3] for a proof of the result on the GOE matrix; the bound for the diagonal matrix depends on the familiar calculation of the expected maximum of d independent standard normal random variables. We see that the norm of the GOE matrix is close to the lower bound provided by (2.5), while the norm of the diagonal matrix is close to the upper bound.

2.5. A Question. Corollary 2.4 shows that the matrix variance controls the expected norm of a matrix Gaussian series. On the other hand, the two examples in the previous section demonstrate that we need more information than the variance to determine the norm up to a constant factor. Therefore, we must ask...

Are there parameters that allow us to calculate the norm of a matrix Gaussian series more precisely than the matrix variance?

This paper provides the first affirmative answer to this question.

3. BEYOND THE MATRIX KHINTCHINE INEQUALITY

This section presents new results that improve on the matrix Khintchine inequality, Proposition 2.3. First, we motivate the type of parameters that arise when we try to refine this result. Then we define a quantity, called the *matrix alignment parameter*, that describes how the coefficients in the matrix Gaussian series interact with each other. In Section 3.3, we use the alignment parameter to state a new bound that provides a uniform improvement over the matrix Khintchine inequality. Further refinements are possible if we consider random matrices with highly symmetric distributions, so we introduce the class of strongly isotropic random matrices in Section 3.5. Section 3.6 contains a matrix Khintchine inequality for matrix Gaussian series that are strongly isotropic. This bound is good enough to compute the norm of a large GOE matrix exactly. Finally, in Sections 3.7 and 3.8, we discuss extensions and related work.

3.1. Prospects. What kind of parameters might allow us to refine Proposition 2.3? The result is already an identity for $p = 1$. For inspiration, let us work out what happens when $p = 2$:

$$\begin{aligned} \mathbb{E} \|\mathbf{X}\|_4^4 &= \mathbb{E} \operatorname{tr} \left(\sum_{i=1}^n \gamma_i \mathbf{H}_i \right)^4 = \sum_{i,j,k,\ell=1}^n \mathbb{E}[\gamma_i \gamma_j \gamma_k \gamma_\ell] \cdot \operatorname{tr} [\mathbf{H}_i \mathbf{H}_j \mathbf{H}_k \mathbf{H}_\ell] \\ &= 2 \operatorname{tr} \left(\sum_{i,j=1}^n \mathbf{H}_i^2 \right)^2 + \operatorname{tr} \left(\sum_{i,j=1}^n \mathbf{H}_i \mathbf{H}_j \mathbf{H}_i \mathbf{H}_j \right) =: 2 \operatorname{tr} \mathbf{Var}(\mathbf{X})^2 + \operatorname{tr} \mathbf{\Delta}. \end{aligned}$$

We use the convention that powers bind before the trace. The product of Gaussian variables has expectation zero unless the indices are paired. In the last expression, the first term comes from the cases where $i = j$ and $k = \ell$ or where $i = \ell$ and $j = k$; the second term comes from the case where $i = k$ and $j = \ell$. Once again, the matrix variance $\mathbf{Var}(\mathbf{X})$ emerges, but we have a new second-order term $\mathbf{\Delta}$ that arises from the summands where the indices alternate: (i, j, i, j) .¹

In a sense, the matrix $\mathbf{\Delta}$ reflects the extent to which the coefficient matrices are aligned. When the family $\{\mathbf{H}_i\}$ commutes, the matrix $\mathbf{\Delta} = \mathbf{Var}(\mathbf{X})^2$, so the second-order term provides no new information. More generally, whenever the coefficients commute, the quantity $(\mathbb{E} \|\mathbf{X}\|_{2p}^{2p})^{1/(2p)}$ can be expressed in terms of the matrix variance and the number p , and the matrix Khintchine inequality, Proposition 2.3, gives an estimate of the correct order. In other words, commuting coefficients are the worst possible circumstance. Most previous work on matrix concentration implicitly uses this worst-case model in the analysis.

To achieve better results, we need to account for how the coefficient matrices \mathbf{H}_i interact with each other. The calculation above suggests that the matrix $\mathbf{\Delta}$ might contain the information we need. Heuristically, when the coefficients fail to commute, the matrix $\mathbf{\Delta}$ should be small. As we will see, this idea is fruitful, but we need a parameter more discerning than $\mathbf{\Delta}$.

Let us summarize this discussion in the following observation:

To improve on the matrix Khintchine inequality, we must quantify the extent to which the coefficient matrices commute.

Our work builds on this intuition to establish new matrix concentration inequalities.

3.2. The Matrix Alignment Parameter. In this section, we introduce a new parameter for a matrix Gaussian series that describes how much the coefficients commute with each other. In later sections, we will present extensions of the matrix Khintchine inequality that rely on this parameter.

Definition 3.1 (Matrix Alignment Parameter). Let $\mathbf{H}_1, \dots, \mathbf{H}_n$ be Hermitian matrices with dimension d . For each $p \geq 1$, the *matrix alignment parameter* of this sequence is the quantity

$$w_p := \max_{\mathbf{Q}_\ell} \left\| \left| \sum_{i,j=1}^n \mathbf{H}_i \mathbf{Q}_1 \mathbf{H}_j \mathbf{Q}_2 \mathbf{H}_i \mathbf{Q}_3 \mathbf{H}_j \right|^{1/4} \right\|_p \quad \text{and} \quad w := w_\infty. \quad (3.1)$$

¹A related observation animates the theory of free probability, which gives a fine description of certain large random matrices. The key fact about centered, free random variables Y and Z is that crossing moments, such as $\varphi(YZY Z)$, must vanish [NS06].

The maximum takes place over a triple $(\mathbf{Q}_1, \mathbf{Q}_2, \mathbf{Q}_3)$ of unitary matrices with dimension d . The matrix absolute value is defined as $|\mathbf{B}| := (\mathbf{B}^* \mathbf{B})^{1/2}$.

Roughly, the matrix alignment parameter (3.1) describes how well the matrices $\mathbf{H}_1, \dots, \mathbf{H}_n$ can be aligned with each other under the worst choices of coordinates.

The quantity (3.1) appears mysterious, so it is worth a few paragraphs to clarify its meaning. First, let us compare the alignment parameter with the matrix standard deviation parameter (2.3) that appears in the matrix Khintchine inequality.

Proposition 3.2 (Standard Deviation versus Alignment). *Let $\mathbf{H}_1, \dots, \mathbf{H}_n$ be Hermitian matrices. Define the standard deviation and alignment parameters*

$$\sigma_p := \left\| \left(\sum_{i=1}^n \mathbf{H}_i \right)^{1/2} \right\|_p \quad \text{and} \quad w_p := \max_{\mathbf{Q}_\ell} \left\| \left| \sum_{i,j=1}^n \mathbf{H}_i \mathbf{Q}_1 \mathbf{H}_j \mathbf{Q}_2 \mathbf{H}_i \mathbf{Q}_3 \mathbf{H}_j \right|^{1/4} \right\|_p \quad \text{for } p \geq 1.$$

Then

$$w_p \leq \sigma_p \quad \text{for all } p \geq 4.$$

The proof of Proposition 3.2 appears in Section 8.9.

Next, let us return to the examples in the introduction. In Section 4, we provide detailed calculations of the standard deviation and alignment parameters. For the diagonal Gaussian series \mathbf{X}_{diag} defined in (2.6), we have

$$\sigma(\mathbf{X}_{\text{diag}}) = 1 \quad \text{and} \quad w(\mathbf{X}_{\text{diag}}) = 1.$$

For the GOE matrix \mathbf{X}_{goe} defined in (2.7),

$$\sigma(\mathbf{X}_{\text{goe}}) = 1 + d^{-1} \quad \text{and} \quad w(\mathbf{X}_{\text{goe}}) \leq (4d)^{-1/4}. \quad (3.2)$$

The matrix alignment parameter can tell the two examples apart, while the matrix standard deviation cannot!

Remark 3.3 (Notation for Alignment). Here and elsewhere, we abuse notation by writing $w_p(\mathbf{X})$ and $w(\mathbf{X})$ for the alignment parameter of a matrix Gaussian series $\mathbf{X} := \sum_{i=1}^n \gamma_i \mathbf{H}_i$, even though w is a function of the coefficient matrices \mathbf{H}_i in the representation of the series.

Remark 3.4 (Are the Unitaries Necessary?). At this stage, it may seem capricious to include the unitary matrices in the definition (3.1). In fact, the example in Section 4.3 demonstrates that the alignment parameter would lose its value if we were to remove the unitary matrices. On the other hand, there are situations where the unitary matrices are not completely arbitrary, as discussed in Section 8.1.

3.3. A Second-Order Matrix Khintchine Inequality. The first major result of this paper is an improvement on the matrix Khintchine inequality. This theorem uses the second-order information in the alignment parameter to obtain better bounds.

Theorem 3.5 (Second-Order Matrix Khintchine). *Consider an Hermitian matrix Gaussian series $\mathbf{X} := \sum_{i=1}^n \gamma_i \mathbf{H}_i$, as in (2.1). Define the matrix standard deviation and matrix alignment parameters*

$$\sigma_q(\mathbf{X}) := \left\| \left(\sum_{i=1}^n \mathbf{H}_i^2 \right)^{1/2} \right\|_q \quad \text{and} \quad w_q(\mathbf{X}) := \max_{\mathbf{Q}_\ell} \left\| \left| \sum_{i,j=1}^n \mathbf{H}_i \mathbf{Q}_1 \mathbf{H}_j \mathbf{Q}_2 \mathbf{H}_i \mathbf{Q}_3 \mathbf{H}_j \right|^{1/4} \right\|_q \quad \text{for } q \geq 1.$$

The maximum takes place over a triple $(\mathbf{Q}_1, \mathbf{Q}_2, \mathbf{Q}_3)$ of unitary matrices. Then, for each integer $p \geq 3$,

$$\left(\mathbb{E} \|\mathbf{X}\|_{2p}^{2p} \right)^{1/(2p)} \leq 3 \sqrt[4]{2p-5} \cdot \sigma_{2p}(\mathbf{X}) + \sqrt{2p-4} \cdot w_{2p}(\mathbf{X}). \quad (3.3)$$

The symbol $\|\cdot\|_q$ denotes the Schatten q -norm.

The proof of Theorem 3.5 appears in Section 8.

We can also derive bounds for the spectral norm of a matrix Gaussian series.

Corollary 3.6 (Second-Order Matrix Khintchine: Spectral Norm). *Consider an Hermitian matrix Gaussian series $\mathbf{X} := \sum_{i=1}^n \gamma_i \mathbf{H}_i$ with dimension $d \geq 8$, as in (2.1). Define the matrix standard deviation and matrix alignment parameters*

$$\sigma(\mathbf{X}) := \left\| \sum_{i=1}^n \mathbf{H}_i^2 \right\|^{1/2} \quad \text{and} \quad w(\mathbf{X}) := \max_{\mathbf{Q}_\ell} \left\| \sum_{i,j=1}^n \mathbf{H}_i \mathbf{Q}_1 \mathbf{H}_j \mathbf{Q}_2 \mathbf{H}_i \mathbf{Q}_3 \mathbf{H}_j \right\|^{1/4}.$$

The maximum ranges over a triple $(\mathbf{Q}_1, \mathbf{Q}_2, \mathbf{Q}_3)$ of unitary matrices. Then

$$\mathbb{E} \|\mathbf{X}\| \leq 3\sigma(\mathbf{X}) \sqrt[4]{2e \log d} + w(\mathbf{X}) \sqrt{2e \log d}. \quad (3.4)$$

The symbol $\|\cdot\|$ denotes the spectral norm.

The result follows from Theorem 3.5 by setting $p = \lceil \log d \rceil$. The potential gain in (3.4) over (2.5) comes from the reduction of the power on the first logarithm from one-half to one-quarter.

3.4. Matrix Khintchine versus Second-Order Matrix Khintchine. Let us make some comparisons between Proposition 2.3 and Theorem 3.5. First, recall that the alignment parameter is dominated by the standard deviation parameter: $w_q(\mathbf{X}) \leq \sigma_q(\mathbf{X})$ for $q \geq 4$ because of Proposition 3.2. Therefore, the bound (3.3) implies that

$$(\mathbb{E} \|\mathbf{X}\|_{2p}^{2p})^{1/(2p)} \leq (3\sqrt[4]{2p-1} + \sqrt{2p-1}) \cdot \sigma_{2p}(\mathbf{X}) \quad \text{for } p = 3, 4, 5, \dots$$

This is very close to the prediction from Proposition 2.3, so Theorem 3.5 is never significantly worse.

On the other hand, there are situations where Theorem 3.5 gives qualitatively better results. In particular, for the GOE matrix $\mathbf{X}_{\text{goe}}(d)$, the bound (3.3) and the calculation (3.2) yield

$$\mathbb{E} \|\mathbf{X}_{\text{goe}}(d)\| \leq 3\sqrt[4]{2e \log d} + \frac{\sqrt{2e \log d}}{\sqrt[4]{4d}} \quad \text{for } d \geq 8.$$

This estimate beats our first attempt in (2.8), but it still falls short of the correct estimate $\mathbb{E} \|\mathbf{X}_{\text{goe}}\| \approx 2$.

3.5. Strongly Isotropic Random Matrices. As we have seen, Theorem 3.5 offers a qualitative improvement over the matrix Khintchine inequality, Proposition 2.3. Nevertheless, the new result still lacks the power to determine the norm of the GOE matrix correctly. We can obtain more satisfactory results by specializing our attention to a class of random matrices with highly symmetric distributions.

Definition 3.7 (Strong Isotropy). Let \mathbf{X} be a random Hermitian matrix. We say that \mathbf{X} is *strongly isotropic* when

$$\mathbb{E} \mathbf{X}^p = (\mathbb{E} \bar{\text{tr}} \mathbf{X}^p) \cdot \mathbf{I} \quad \text{for } p = 0, 1, 2, \dots$$

The symbol $\bar{\text{tr}}$ denotes the normalized trace: $\bar{\text{tr}} \mathbf{A} := d^{-1} \text{tr} \mathbf{A}$ when \mathbf{A} has dimension d .

The easiest way to check that a random matrix is strongly isotropic is to exploit symmetry properties of the distribution. We offer one of many possible results in this direction [CT14, Lem. 7.1].

Proposition 3.8 (Strong Isotropy: Sufficient Condition). *Let \mathbf{X} be a random Hermitian matrix. Suppose that the distribution of \mathbf{X} is invariant under signed permutation:*

$$\mathbf{X} \sim \mathbf{\Pi}^* \mathbf{X} \mathbf{\Pi} \quad \text{for every signed permutation } \mathbf{\Pi}.$$

Then \mathbf{X} is strongly isotropic. The symbol \sim refers to equality of distribution. A signed permutation is a square matrix that has precisely one nonzero entry in each row or column, this entry taking the values ± 1 .

Proof. Suppose that $\mathbf{\Pi}$ is a signed permutation, drawn uniformly at random. For $p = 0, 1, 2, 3, \dots$,

$$\mathbb{E} \mathbf{X}^p = \mathbb{E} [(\mathbf{\Pi}^* \mathbf{X} \mathbf{\Pi})^p] = \mathbb{E} [\mathbb{E} [\mathbf{\Pi}^* \mathbf{X}^p \mathbf{\Pi} \mid \mathbf{X}]] = \mathbb{E} [(\bar{\text{tr}} \mathbf{X}^p) \cdot \mathbf{I}] = (\mathbb{E} \bar{\text{tr}} \mathbf{X}^p) \cdot \mathbf{I}$$

The first relation uses invariance under signed permutation, and the second relies on the fact that signed permutations are unitary. Averaging a fixed matrix over signed permutations yields the identity times the normalized trace of the matrix. \square

Proposition 3.8 applies to many types of random matrices. In particular, the diagonal Gaussian matrix \mathbf{X}_{diag} and the GOE matrix \mathbf{X}_{goe} are both strongly isotropic because of this result. Other types of distributional symmetry can also lead to strong isotropy.

Remark 3.9 (Group Orbits). Here is a more general class of matrix Gaussian series where we can verify strong isotropy using abstract arguments. Let \mathcal{G} be a unitary representation of a finite group, and let \mathbf{A} be a fixed Hermitian matrix with the same dimension. Consider the random Hermitian matrix

$$\mathbf{X} := \sum_{\mathbf{U} \in \mathcal{G}} \gamma_{\mathbf{U}} \cdot \mathbf{U} \mathbf{A} \mathbf{U}^*$$

where $\{\gamma_{\mathbf{U}} : \mathbf{U} \in \mathcal{G}\}$ is an independent family of standard normal variables. Since \mathcal{G} acts on itself by permutation,

$$\mathbf{U} \mathbf{X} \mathbf{U}^* \sim \mathbf{X} \quad \text{for each } \mathbf{U} \in \mathcal{G}.$$

This observation allows us to perform averaging arguments like the one in Proposition 3.8.

There are several ways to apply this property to argue that \mathbf{X} is strongly isotropic. For example, it suffices that

$$\mathcal{G}' := \{\mathbf{M} : \mathbf{M}\mathbf{U} = \mathbf{U}\mathbf{M} \text{ for all } \mathbf{U} \in \mathcal{G}\} = \{z\mathbf{I} : z \in \mathbb{C}\}.$$

It is also sufficient that $\{\mathbf{U}\mathbf{a} : \mathbf{U} \in \mathcal{G}\}$ forms a (complete) tight frame for every vector \mathbf{a} ; see the paper [VW08] for some situations where this condition holds.

Remark 3.10 (Spherical Designs). A *spherical t -design* is a collection $\{\mathbf{u}_i : i = 1, \dots, N\}$ of points on the unit sphere \mathbb{S}^{d-1} in \mathbb{R}^d with the property that

$$\int_{\mathbb{S}^{d-1}} \varphi(\mathbf{u}) d\mu(\mathbf{u}) = \frac{1}{N} \sum_{i=1}^N \varphi(\mathbf{u}_i)$$

where φ is an arbitrary algebraic polynomial in d variables with degree t and $d\mu$ is the Haar measure on the sphere \mathbb{S}^{d-1} . See the paper [BRV13] for existence results and background references.

Given a spherical t -design, consider the random matrix

$$\mathbf{X} := \sum_{i=1}^N \gamma_i \mathbf{u}_i \mathbf{u}_i^*.$$

where $\{\gamma_i : i = 1, \dots, N\}$ is an independent family of standard normal variables. By construction, this random matrix has the property that

$$\mathbb{E} \mathbf{X}^p = (\mathbb{E} \bar{\text{tr}} \mathbf{X}^p) \cdot \mathbf{I} \quad \text{for } p = 0, 1, 2, \dots, \lfloor t/2 \rfloor.$$

This variant of the strong isotropy property is sufficient for many purposes, provided that $t \approx \log d$.

3.6. A Second-Order Khintchine Inequality under Strong Isotropy. The second major result of this paper is a second-order matrix Khintchine inequality that is valid for matrix Gaussian series with the strong isotropy property. Like Theorem 3.5, this result uses the alignment parameter to control the norm of the random matrix.

Theorem 3.11 (Second-Order Khintchine under Strong Isotropy). *Consider an Hermitian matrix Gaussian series $\mathbf{X} := \sum_{i=1}^n \gamma_i \mathbf{H}_i$ with dimension d , as in (2.1), and assume that \mathbf{X} is strongly isotropic. Introduce the matrix standard deviation and matrix alignment parameters:*

$$\sigma(\mathbf{X}) := \left\| \sum_{i=1}^n \mathbf{H}_i^2 \right\|^{1/2} \quad \text{and} \quad w(\mathbf{X}) := \max_{\mathbf{Q}_\ell} \left\| \sum_{i,j=1}^n \mathbf{H}_i \mathbf{Q}_1 \mathbf{H}_j \mathbf{Q}_2 \mathbf{H}_i \mathbf{Q}_3 \mathbf{H}_j \right\|^{1/4}.$$

The maximum ranges over a triple $(\mathbf{Q}_1, \mathbf{Q}_2, \mathbf{Q}_3)$ of unitary matrices. Then, for each integer $p \geq 1$,

$$(\mathbb{E} \|\mathbf{X}\|_{2p}^{2p})^{1/(2p)} \leq [2\sigma(\mathbf{X}) + 2^{1/4} p^{5/4} w(\mathbf{X})] \cdot d^{1/(2p)}.$$

The symbol $\|\cdot\|$ refers to the spectral norm, while $\|\cdot\|_q$ is the Schatten q -norm.

The proof of this result appears in Section 9, where we also establish a lower bound.

Theorem 3.11 shows that the moments of the random matrix \mathbf{X} are controlled by the standard deviation $\sigma(\mathbf{X})$ whenever $p^{5/4} w(\mathbf{X}) \ll \sigma(\mathbf{X})$. If we take $p = \lceil \log d \rceil$, the Schatten $2p$ -norm is essentially the same as the spectral norm, and the dimensional factor on the right-hand side is negligible. Therefore,

$$w(\mathbf{X}) \log^{5/4} d \ll \sigma(\mathbf{X}) \quad \text{implies} \quad \mathbb{E} \|\mathbf{X}\| \lesssim 2\sigma(\mathbf{X}).$$

In the presence of strong isotropy, the spectral norm of a matrix Gaussian series is comparable with the standard deviation $\sigma(\mathbf{X})$ whenever the alignment parameter $w(\mathbf{X})$ is relatively small!

In particular, we can apply this result to the GOE matrix \mathbf{X}_{goe} because of Proposition 3.8. The calculation (3.2) of the standard deviation and alignment parameters ensures that $\mathbb{E} \|\mathbf{X}_{\text{goe}}\| \lesssim 2$. As we observed in (2.9), this bound is sharp. For this example, we can even take $p \approx d^{1/5}$, which leads to very good probability bounds via Markov's inequality. Furthermore, a more detailed version of Theorem 3.11, appearing in Section 9, is precise enough to show that the semicircle law is the limiting spectral distribution of the GOE.

On the other hand, the dependence on the exponent p in Theorem 3.11 is suboptimal. This point is evident when we consider the diagonal Gaussian matrix $\mathbf{X}_{\text{diag}}(d)$. Indeed, Theorem 3.11 only implies the bound

$$\mathbb{E} \|\mathbf{X}_{\text{diag}}\| \leq \text{const} \cdot \log^{5/4} d.$$

As we observed in (2.9), the power on the logarithm should be one-half.

3.7. Discussion. This paper opens a new chapter in the theory of matrix concentration and noncommutative moment inequalities. Our main technical contribution is to demonstrate that the matrix Khintchine inequality, Proposition 2.3, is not the last word on the behavior of a matrix Gaussian series. Indeed, we have shown that the matrix variance does not contain sufficient information to determine the expected norm of a matrix Gaussian series. We have also identified another quantity, the matrix alignment parameter, that allows us to obtain better bounds for every matrix Gaussian series. Furthermore, in the presence of more extensive distributional information, it is even possible to obtain numerically sharp bounds for the norm of certain matrix Gaussian series.

There are a number of ways to extend the ideas and results in this paper:

Higher-Order Alignment: If we consider alignment parameters involving $2k$ coefficient matrices, it is possible to improve the term $p^{1/4}\sigma_{2p}$ in Theorem 3.5 to $p^{1/(2k)}\sigma_{2p}$. See Section 8.1 for some additional details.

Other Matrix Series: We can use exchangeable pairs techniques [MJC⁺14] to study matrix series of the form $\mathbf{X} := \sum_{i=1}^n \xi_i \mathbf{H}_i$ where $\{\xi_i\}$ is an independent family of scalar random variables. This approach is potentially quite interesting when the ξ_i are Bernoulli (that is, 0–1) random variables.

Independent Sums: We can use conditioning and symmetrization, as in (2.2), to apply Theorem 3.5 to a sum of independent random matrices. See [CGT12, App.] for an example of this type of argument.

Rectangular Matrices: The techniques here also give results for rectangular random matrices by way of the Hermitian dilation [Tro15, Sec. 2.1.16]. In this setting, a different notion of strong isotropy becomes relevant; see Section 9.1.

We have not elaborated on these ideas because there is also evidence that alignment parameters will not lead to final results on matrix concentration.

3.8. Related Work. There are very few techniques in the literature on random matrices that satisfy all three of our three requirements: flexibility, ease of use, and power. In particular, for many practical applications, it is important to be able to work with an arbitrary sum of independent random matrices. We have chosen to study matrix Gaussian series because they are the simplest instance of this model, and they may lead to further insights about the general problem.

Most classical work in random matrix theory concerns very special classes of random matrices; the books [BS10, Tao12] provide an overview of some of the main lines of research in this field. There are some specific subareas of random matrix theory that address more general models. The monograph [NS06] gives an introduction to free probability. The book chapter [Ver12] describes a collection of methods from Banach space geometry. The monograph [Tro15] covers the theory of matrix concentration. The last three works have a wide scope of applicability, but none of them provides the ultimate description of the behavior of a sum of independent random matrices.

There is one specific strand of research that we would like to draw out because it is very close in spirit to this paper. Recently, Bandeira & van Handel [BV14] and van Handel [vH15] have studied the behavior of a real symmetric Gaussian matrix whose entries are independent and centered but have inhomogeneous variances (the *independent-entry model*). A $d \times d$ random matrix from this class can be written as

$$\mathbf{X}_{\text{indep}} := \sum_{i,j=1}^d a_{ij} \gamma_{ij} \cdot (\mathbf{E}_{ij} + \mathbf{E}_{ji}) \quad \text{for } a_{ij} \in \mathbb{R}.$$

As usual, $\{\gamma_{ij}\}$ is an independent family of standard normal random variables, and we assume that $a_{ij} = a_{ji}$ without loss of generality.

To situate this model in the context of our work, observe that matrix Gaussian series are significantly more general than the independent-entry model. The strongly isotropic model is incomparable with the independent-entry model. To see why, recall that strongly isotropic matrices can have dependent entries. At the same time, $\mathbb{E} \mathbf{X}_{\text{indep}}^p$ is diagonal for each integer $p \geq 0$, but it need not be a scalar matrix.

For the independent-entry model, Bandeira & van Handel [BV14] established the following (sharp) bound:

$$\mathbb{E} \|\mathbf{X}_{\text{indep}}\| \lesssim 2\sigma(\mathbf{X}_{\text{indep}}) + \text{const} \cdot \max_{ij} |a_{ij}| \cdot \sqrt{\log d} \quad (3.5)$$

The maximum entry $\max_{ij} |a_{ij}|$ plays the same role in this formula as the matrix alignment parameter plays in this paper. The paper [BV14] leans heavily on the independence assumption, so it is not clear whether the ideas extend to a more general setting.

To compare the result (3.5) with the work here, we can compute the matrix alignment parameter for the independent-entry model using a difficult extension of the calculation in Section 4.2. This effort yields

$$w(\mathbf{X}_{\text{indep}}) \approx \left(\max_i \sum_j |a_{ij}|^4 \right)^{1/4}.$$

We see that the matrix alignment parameter is somewhat larger than the maximum entry $\max_{ij} |a_{ij}|$. Thus, for the independent model, Theorem 3.5 gives us a better result than the classical Khintchine inequality, Proposition 2.3, but it is somewhat weaker than (3.5). Theorem 3.11 would give a result close to the bound (3.5), but it does not always apply because the independent-entry model need not be strongly isotropic.

The independent-entry model is not adequate to reach results with the same power and scope as the current generation of matrix concentration bounds [Tro15]. Nevertheless, the estimate (3.5) strongly suggests that there are better ways of summarizing the interactions of the coefficients in an Hermitian matrix Gaussian series $\mathbf{X} := \sum_{i=1}^n \gamma_i \mathbf{H}_i$ than the alignment parameter $w(\mathbf{X})$. One possibility is the weak variance parameter:

$$\sigma_\star(\mathbf{X}) := \sup_{\|\mathbf{u}\|=\|\mathbf{v}\|=1} \left(\sum_{i=1}^n |\langle \mathbf{u}, \mathbf{H}_i \mathbf{v} \rangle|^2 \right)^{1/2}.$$

For the independent-entry model, this quantity reduces to $\text{const} \cdot \max_{ij} |a_{ij}|$. The idea of considering $\sigma_\star(\mathbf{X})$ is motivated by the discussion in [Tro12, Sec. 4], as well as the work in [BV14, vH15]. Unfortunately, at this stage, it is not clear whether there are any parameters that allow us to obtain a simple description of the behavior of a Gaussian matrix in the absence of burdensome independence or isotropy assumptions. This is a frontier for future work.

4. COMPUTATION OF THE MATRIX ALIGNMENT PARAMETERS

In this section, we show how to compute the matrix alignment parameter for the two random matrices in the introduction, the diagonal Gaussian matrix and the GOE matrix. Afterward, we show by example that neither Theorem 3.5 nor Theorem 3.11 can hold if we remove the unitary factors from the matrix alignment parameter.

4.1. A Diagonal Gaussian Matrix. The diagonal Gaussian matrix takes the form

$$\mathbf{X}_{\text{diag}} := \sum_{i=1}^d \gamma_i \mathbf{E}_{ii}.$$

The matrix variance $\mathbf{Var}(\mathbf{X}_{\text{diag}}) = \mathbb{E} \mathbf{X}_{\text{diag}}^2 = \mathbf{I}$. It follows that the matrix standard deviation parameters, defined in (2.3), satisfy

$$\sigma_p(\mathbf{X}_{\text{diag}}) = \|\mathbf{Var}(\mathbf{X}_{\text{diag}})^{1/2}\|_p = d^{1/p} \quad \text{for } 1 \leq p \leq \infty.$$

We will show that the matrix alignment parameters, defined in (3.1), satisfy

$$w_p(\mathbf{X}_{\text{diag}}) = d^{1/p} \quad \text{for } 4 \leq p \leq \infty.$$

Thus, for this example, the second-order matrix Khintchine inequalities, Theorem 3.5 and Theorem 3.11, do not improve over the matrix Khintchine inequality, Proposition 2.3. This outcome is natural, given that the classical result is essentially optimal in this case.

Let us evaluate the matrix alignment parameter. For a triple $(\mathbf{Q}, \mathbf{S}, \mathbf{U})$ of unitary matrices, form the sum

$$\mathbf{W}(\mathbf{Q}, \mathbf{S}, \mathbf{U}) := \sum_{i,j=1}^d \mathbf{E}_{ii} \mathbf{Q} \mathbf{E}_{jj} \mathbf{S} \mathbf{E}_{ii} \mathbf{U} \mathbf{E}_{jj} = \sum_{i,j=1}^d q_{ij} s_{ji} u_{ij} \cdot \mathbf{E}_{ij} = \mathbf{Q} \odot \mathbf{S}^t \odot \mathbf{U}.$$

We have written \odot for the Schur (i.e., componentwise) product, and t is the transpose operation. When $\mathbf{Q} = \mathbf{S} = \mathbf{U} = \mathbf{I}$, the sum collapses: $\mathbf{W}(\mathbf{I}, \mathbf{I}, \mathbf{I}) = \mathbf{I}$. Therefore,

$$w_p(\mathbf{X}_{\text{diag}}) = \max_{\mathbf{Q}, \mathbf{S}, \mathbf{U}} \|\mathbf{W}(\mathbf{Q}, \mathbf{S}, \mathbf{U})\|_p^{1/4} \geq \|\mathbf{I}\|_p = d^{1/p} \quad \text{for } p \geq 1.$$

But Proposition 3.2 shows that

$$w_p(\mathbf{X}_{\text{diag}}) \leq \sigma_p(\mathbf{X}_{\text{diag}}) = d^{1/p} \quad \text{for each } p \geq 4.$$

Therefore, $w_p(\mathbf{X}_{\text{diag}}) = \sigma_p(\mathbf{X}_{\text{diag}}) = d^{1/p}$ for $p \geq 4$. The result for $p = \infty$ follows when we take limits.

Remark 4.1 (Commutativity). A similar calculation is valid whenever the family $\{\mathbf{H}_i\}$ of coefficient matrices in the matrix Gaussian series (2.1) commutes.

4.2. **A GOE Matrix.** The GOE matrix takes the form

$$\mathbf{X}_{\text{goe}} := \frac{1}{\sqrt{2d}} \sum_{i,j=1}^d \gamma_{ij} (\mathbf{E}_{ij} + \mathbf{E}_{ji}).$$

An easy calculation shows that the matrix variance satisfies

$$\mathbf{Var}(\mathbf{X}_{\text{goe}}) = \mathbb{E} \mathbf{X}_{\text{goe}}^2 = \frac{1}{2d} \sum_{i,j=1}^d (\mathbf{E}_{ij} + \mathbf{E}_{ji})^2 = (1 + d^{-1}) \cdot \mathbf{I}.$$

Therefore, the matrix standard deviation parameters, defined in (2.3), equal

$$\sigma_p(\mathbf{X}_{\text{goe}}) = \|\mathbf{Var}(\mathbf{X}_{\text{goe}})^{1/2}\|_p = \sqrt{1 + d^{-1}} \cdot d^{1/p} \quad \text{for } 1 \leq p \leq \infty.$$

We will demonstrate that the matrix alignment parameters, defined in (3.1), satisfy

$$w_p(\mathbf{X}_{\text{goe}}) \leq (d^{-1} + 3d^{-2})^{1/4} \cdot d^{1/p} \quad \text{for } 4 \leq p \leq \infty.$$

When d is large, the matrix alignment parameters are much smaller than the matrix standard deviation parameters. As a consequence, the second-order matrix Khintchine inequalities deliver a substantial gain over the classical matrix Khintchine inequality.

Let us compute the matrix alignment parameter. For a triple $(\mathbf{Q}, \mathbf{S}, \mathbf{U})$ of unitary matrices, introduce the (unnormalized) sum

$$\mathbf{W}(\mathbf{Q}, \mathbf{S}, \mathbf{U}) := \sum_{i_1, i_2, j_1, j_2=1}^d (\mathbf{E}_{i_1 i_2} + \mathbf{E}_{i_2 i_1}) \mathbf{Q} (\mathbf{E}_{j_1 j_2} + \mathbf{E}_{j_2 j_1}) \mathbf{S} (\mathbf{E}_{i_1 i_2} + \mathbf{E}_{i_2 i_1}) \mathbf{U} (\mathbf{E}_{j_1 j_2} + \mathbf{E}_{j_2 j_1}).$$

It is not hard to evaluate this sum if we take care. First, distribute terms:

$$\begin{aligned} \mathbf{W}(\mathbf{Q}, \mathbf{S}, \mathbf{U}) = \sum_{i_1, i_2, j_1, j_2=1}^d & \left[(q_{i_2 j_1} s_{j_2 i_1} u_{i_2 j_2} + q_{i_2 j_1} s_{j_2 i_2} u_{i_1 j_2} + q_{i_2 j_2} s_{j_1 i_1} u_{i_2 j_2} + q_{i_2 j_2} s_{j_1 i_2} u_{i_1 j_2}) \cdot \mathbf{E}_{i_1 j_1} \right. \\ & + (q_{i_2 j_1} s_{j_2 i_1} u_{i_2 j_1} + q_{i_2 j_1} s_{j_2 i_2} u_{i_1 j_1} + q_{i_2 j_2} s_{j_1 i_1} u_{i_2 j_1} + q_{i_2 j_2} s_{j_1 i_2} u_{i_1 j_1}) \cdot \mathbf{E}_{i_1 j_2} \\ & + (q_{i_1 j_1} s_{j_2 i_1} u_{i_2 j_2} + q_{i_1 j_1} s_{j_2 i_2} u_{i_1 j_2} + q_{i_1 j_2} s_{j_1 i_1} u_{i_2 j_2} + q_{i_1 j_2} s_{j_1 i_2} u_{i_1 j_2}) \cdot \mathbf{E}_{i_2 j_1} \\ & \left. + (q_{i_1 j_1} s_{j_2 i_1} u_{i_2 j_1} + q_{i_1 j_1} s_{j_2 i_2} u_{i_1 j_1} + q_{i_1 j_2} s_{j_1 i_1} u_{i_2 j_1} + q_{i_1 j_2} s_{j_1 i_2} u_{i_1 j_1}) \cdot \mathbf{E}_{i_2 j_2} \right]. \end{aligned}$$

In each line, we can sum through the two free indices to identify four matrix products. For example, in the first line, we can sum on i_2 and j_2 . This step yields

$$\begin{aligned} \mathbf{W}(\mathbf{Q}, \mathbf{S}, \mathbf{U}) = \sum_{i_1, j_1=1}^d & (\mathbf{s}^t \mathbf{U}^t \mathbf{Q} + \mathbf{U} \mathbf{S} \mathbf{Q} + \text{tr}(\mathbf{Q}^t \mathbf{U}) \cdot \mathbf{s}^t + \mathbf{U} \mathbf{Q}^t \mathbf{s}^t)_{i_1 j_1} \cdot \mathbf{E}_{i_1 j_1} \\ & + \sum_{i_1, j_2=1}^d (\text{tr}(\mathbf{Q}^t \mathbf{U}) \cdot \mathbf{s}^t + \mathbf{U} \mathbf{Q}^t \mathbf{s}^t + \mathbf{s}^t \mathbf{U}^t \mathbf{Q} + \mathbf{U} \mathbf{S} \mathbf{Q})_{i_1 j_2} \cdot \mathbf{E}_{i_1 j_2} \\ & + \sum_{i_2, j_1=1}^d (\mathbf{U} \mathbf{S} \mathbf{Q} + \mathbf{s}^t \mathbf{U}^t \mathbf{Q} + \mathbf{U} \mathbf{Q}^t \mathbf{s}^t + \text{tr}(\mathbf{Q} \mathbf{U}) \cdot \mathbf{s}^t)_{i_2 j_1} \cdot \mathbf{E}_{i_2 j_1} \\ & + \sum_{i_2, j_2=1}^d (\mathbf{U} \mathbf{Q}^t \mathbf{s}^t + \text{tr}(\mathbf{Q} \mathbf{U}) \cdot \mathbf{s}^t + \mathbf{U} \mathbf{S} \mathbf{Q} + \mathbf{s}^t \mathbf{U}^t \mathbf{Q})_{i_2 j_2} \cdot \mathbf{E}_{i_2 j_2}. \end{aligned}$$

Sum through the remaining indices to reach

$$\begin{aligned} \mathbf{W}(\mathbf{Q}, \mathbf{S}, \mathbf{U}) = & (\mathbf{s}^t \mathbf{U}^t \mathbf{Q} + \mathbf{U} \mathbf{S} \mathbf{Q} + \text{tr}(\mathbf{Q}^t \mathbf{U}) \cdot \mathbf{s}^t + \mathbf{U} \mathbf{Q}^t \mathbf{s}^t) \\ & + (\text{tr}(\mathbf{Q}^t \mathbf{U}) \cdot \mathbf{s}^t + \mathbf{U} \mathbf{Q}^t \mathbf{s}^t + \mathbf{s}^t \mathbf{U}^t \mathbf{Q} + \mathbf{U} \mathbf{S} \mathbf{Q}) \\ & + (\mathbf{U} \mathbf{S} \mathbf{Q} + \mathbf{s}^t \mathbf{U}^t \mathbf{Q} + \mathbf{U} \mathbf{Q}^t \mathbf{s}^t + \text{tr}(\mathbf{Q} \mathbf{U}) \cdot \mathbf{s}^t) \\ & + (\mathbf{U} \mathbf{Q}^t \mathbf{s}^t + \text{tr}(\mathbf{Q} \mathbf{U}) \cdot \mathbf{s}^t + \mathbf{U} \mathbf{S} \mathbf{Q} + \mathbf{s}^t \mathbf{U}^t \mathbf{Q}). \end{aligned}$$

Twelve of the sixteen terms are unitary matrices, and the remaining four are scaled unitary matrices. Furthermore, each trace is bounded in magnitude by d , the worst case being $\mathbf{Q} = \mathbf{U} = \mathbf{I}$. Applying the definition of the Schatten norm, the triangle inequality, and unitary invariance, we find that

$$\|\mathbf{W}(\mathbf{Q}, \mathbf{S}, \mathbf{U})\|_p^{1/4} = \|\mathbf{W}(\mathbf{Q}, \mathbf{S}, \mathbf{U})\|_{p/4}^{1/4} \leq ((4d + 12) \cdot \|\mathbf{I}\|_{p/4})^{1/4} = (4d + 12)^{1/4} \cdot d^{1/p} \quad \text{for } p \geq 4.$$

To compute $w_p(\mathbf{X}_{\text{goe}})$, we must reintroduce the scaling $(2d)^{-1/2}$, which gives the advertised result:

$$w_p(\mathbf{X}_{\text{goe}}) \leq (2d)^{-1/2} \cdot (4d + 12)^{1/4} \cdot d^{1/p} = (d^{-1} + 3d^{-2})^{1/4} \cdot d^{1/p}.$$

To obtain the bound for $p = \infty$, we simply take limits.

4.3. The Unitaries are Necessary. Suppose that $\mathbf{X} := \sum_{i=1}^n \gamma_i \mathbf{H}_i$ is an Hermitian matrix Gaussian series with dimension d , and let $\sigma(\mathbf{X})$ be the matrix standard deviation (2.3). Consider the alternative alignment parameter

$$\delta(\mathbf{X}) := \left\| \sum_{i,j=1}^n \mathbf{H}_i \mathbf{H}_j \mathbf{H}_i \mathbf{H}_j \right\|^{1/4}.$$

This quantity is suggested by the discussion in Section 3.1. Consider a general estimate of the form

$$\mathbb{E} \|\mathbf{X}\| \leq f(d) \cdot \sigma(\mathbf{X}) + g(d) \cdot \delta(\mathbf{X}). \quad (4.1)$$

We will demonstrate that, for every choice of the function g , there is a lower bound $f(d) \geq \text{const} \cdot \sqrt{\log d}$. From this claim, we deduce that it is impossible to improve over the classical Khintchine inequality by using the second-order quantity $\delta(\mathbf{X})$. Therefore, the unitary matrices in the alignment parameter $w(\mathbf{X})$ play a critical role. Most of this argument was developed by Afonso Bandeira; we are grateful to him for allowing us to include it.

Introduce the Pauli spin matrices

$$\mathbf{H}_1 := \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} \quad \mathbf{H}_2 := \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \quad \mathbf{H}_3 := \begin{bmatrix} 0 & i \\ -i & 0 \end{bmatrix}.$$

These matrices are Hermitian and unitary, so $\mathbf{H}_i^2 = \mathbf{I}$ for $i = 1, 2, 3$. Furthermore, they satisfy the relations $(\mathbf{H}_i \mathbf{H}_j)^2 = -\mathbf{I}$ when $i \neq j$. Next, define $\mathbf{H}_0 := \sqrt{\alpha} \mathbf{I}$, where $\alpha := 2\sqrt{3} - 3$. Calculate that

$$\begin{aligned} \sum_{i,j=0}^3 \mathbf{H}_i \mathbf{H}_j \mathbf{H}_i \mathbf{H}_j &= \sum_{i=0}^4 \mathbf{H}_i^4 + \sum_{j=1}^3 \mathbf{H}_0 \mathbf{H}_j \mathbf{H}_0 \mathbf{H}_j + \sum_{i=1}^3 \mathbf{H}_i \mathbf{H}_0 \mathbf{H}_i \mathbf{H}_0 + \sum_{\substack{i,j=1 \\ i \neq j}}^3 \mathbf{H}_i \mathbf{H}_j \mathbf{H}_i \mathbf{H}_j \\ &= (\alpha^2 + 3) \mathbf{I} + 6\alpha \mathbf{I} - 6\mathbf{I} = (\alpha^2 + 6\alpha - 3) \mathbf{I} = \mathbf{0}. \end{aligned}$$

Indeed, α is a positive root of the quadratic.

Consider the two-dimensional Gaussian series \mathbf{Y} generated by the matrices $\mathbf{H}_0, \dots, \mathbf{H}_3$:

$$\mathbf{Y} := \sum_{i=0}^3 \gamma_i \mathbf{H}_i.$$

As usual, $\{\gamma_i\}$ is an independent family of standard normal variables. For the series \mathbf{Y} , we have already shown that the alternative alignment parameter $\delta(\mathbf{Y}) = 0$. Let us compute the variance and standard deviation:

$$\text{Var}(\mathbf{Y}) = \sum_{i=0}^3 \mathbf{H}_i^2 = (\alpha + 3) \mathbf{I} = 2\sqrt{3} \mathbf{I} \quad \text{and} \quad \sigma(\mathbf{Y}) = \|\text{Var}(\mathbf{Y})\|^{1/2} = 12^{1/4}.$$

Expanding the random matrix \mathbf{Y} in coordinates, we also find that

$$\mathbf{Y} = \begin{bmatrix} \sqrt{\alpha} \gamma_0 + \gamma_1 & \gamma_2 + i\gamma_3 \\ \gamma_2 - i\gamma_3 & \sqrt{\alpha} \gamma_0 - \gamma_1 \end{bmatrix}.$$

Therefore, the top-left entry $(\mathbf{Y})_{11}$ is a centered normal random variable with variance $1 + \alpha = 2(\sqrt{3} - 1)$.

To obtain the counterexample to the bound (4.1), fix an integer $d \geq 1$. Let $\mathbf{Y}_1, \dots, \mathbf{Y}_d$ be independent copies of the two-dimensional Gaussian series \mathbf{Y} , and construct the $2d$ -dimensional matrix Gaussian series

$$\mathbf{X}_{\text{spin}} := \mathbf{Y}_1 \oplus \dots \oplus \mathbf{Y}_d = \sum_{j=1}^d \mathbf{E}_{jj} \otimes \mathbf{Y}_j \sim \sum_{j=1}^d \sum_{i=0}^3 \gamma_{ij} (\mathbf{E}_{jj} \otimes \mathbf{H}_i).$$

We have written \oplus for direct sum and \otimes for the Kronecker product; the matrices \mathbf{E}_{jj} are the diagonal units with dimension $d \times d$; and $\{\gamma_{ij}\}$ is an independent family of standard normal variables.

Extending the calculations above, we find that $\sigma(\mathbf{X}_{\text{spin}}) = 12^{1/4}$ and $\delta(\mathbf{X}_{\text{spin}}) = 0$. Meanwhile, the norm of \mathbf{X}_{spin} is bounded below by the absolute value of each of its diagonal entries. In particular,

$$\mathbb{E} \|\mathbf{X}_{\text{spin}}\| \geq \mathbb{E} \max_j |(Y_j)_{11}| \geq \text{const} \cdot (2(\sqrt{3} - 1))^{1/2} \cdot \sqrt{\log d}.$$

We have used the fact that the expected maximum of d independent standard normal variables is proportional to $\sqrt{\log d}$. Assuming that (4.1) is valid, we can sequence these estimates to obtain

$$\text{const} \cdot \sqrt{\log d} \leq f(d) \cdot \sigma(\mathbf{X}_{\text{spin}}) + g(d) \cdot \delta(\mathbf{X}_{\text{spin}}) = 12^{1/4} \cdot f(d).$$

Therefore, the function $f(d)$ must grow at least as fast as $\sqrt{\log d}$. We conclude that a bound of the form (4.1) can never improve over the classical matrix Khintchine inequality.

5. NOTATION & BACKGROUND

Before we enter into the body of the paper, let us set some additional notation and state a few background results. First, \mathbb{M}_d denotes the complex linear space of $d \times d$ matrices with complex entries. We write \mathbb{H}_d for the real-linear subspace of \mathbb{M}_d that consists of Hermitian matrices. The symbol $*$ represents conjugate transposition. We write $\mathbf{0}$ for the zero matrix and \mathbf{I} for the identity. The matrix \mathbf{E}_{ij} has a one in the (i, j) position and zeros elsewhere. The dimensions of these matrices are typically determined by context.

For an Hermitian matrix \mathbf{A} , we define the integer powers \mathbf{A}^p for $p = 0, 1, 2, 3, \dots$ in the usual way by iterated multiplication. For a positive-semidefinite matrix \mathbf{P} , we can also define complex powers \mathbf{P}^z by raising each eigenvalue of \mathbf{P} to the power z while maintaining the eigenvectors. In particular, $\mathbf{P}^{1/2}$ is the unique positive-semidefinite square root of \mathbf{P} . The matrix absolute value is defined for a general matrix \mathbf{B} by the rule $|\mathbf{B}| := (\mathbf{B}^* \mathbf{B})^{1/2}$. Note that $|\mathbf{P}| = \mathbf{P}$ when \mathbf{P} is positive semidefinite.

The *trace* and *normalized trace* of a matrix are given by

$$\mathrm{tr} \mathbf{B} := \sum_{i=1}^d b_{ii} \quad \text{and} \quad \bar{\mathrm{tr}} \mathbf{B} := \frac{1}{d} \sum_{i=1}^d b_{ii} \quad \text{for } \mathbf{B} \in \mathbb{M}_d.$$

We use the convention that a power binds before the trace to avoid unnecessary parentheses; powers also bind before expectation. The Schatten p -norm is defined for an arbitrary matrix \mathbf{B} via the rule

$$\|\mathbf{B}\|_p := (\mathrm{tr} |\mathbf{B}|^p)^{1/p} \quad \text{for } p \geq 1.$$

The Schatten ∞ -norm $\|\cdot\|_\infty$ coincides with the spectral norm $\|\cdot\|$. This work uses both trace powers and Schatten norms, depending on which one is conceptually clearer. We require some Hölder inequalities involving the trace and the Schatten norms. For matrices $\mathbf{A}, \mathbf{B} \in \mathbb{M}_d$ and $\varrho \geq 1$,

$$|\mathrm{tr}(\mathbf{A}\mathbf{B})| \leq (\mathrm{tr} |\mathbf{A}|^\varrho)^{1/\varrho} \cdot (\mathrm{tr} |\mathbf{B}|^{\varrho'})^{1/\varrho'} \quad \text{where } \varrho' := \varrho/(\varrho - 1). \quad (5.1)$$

Furthermore,

$$\|\mathbf{A}^* \mathbf{B}\|_\varrho^2 \leq \|\mathbf{A}^* \mathbf{A}\|_\varrho \cdot \|\mathbf{B}^* \mathbf{B}\|_\varrho. \quad (5.2)$$

These results are drawn from [Bha97, Chap. IV].

6. THE TRACE MOMENTS OF A MATRIX GAUSSIAN SERIES

For each major result in this paper, the starting point is a formula for the trace moments of a matrix Gaussian series.

Lemma 6.1 (Trace Moment Identity). *Let $\mathbf{X} := \sum_{i=1}^n \gamma_i \mathbf{H}_i$ be an Hermitian matrix Gaussian series, as in (2.1). For each integer $p \geq 1$, we have the identity*

$$\mathbb{E} \mathrm{tr} \mathbf{X}^{2p} = \sum_{q=0}^{2p-2} \sum_{i=1}^n \mathbb{E} \mathrm{tr} [\mathbf{H}_i \mathbf{X}^q \mathbf{H}_i \mathbf{X}^{2p-2-q}]. \quad (6.1)$$

The easy proof of Lemma 6.1 appears in the next two subsections.

Integration by parts is not foreign in the study of Gaussian random matrices; for example, see [AGZ10, Sec. 2.4.1] or [Kem13, Sec. 9]. The exchangeable pairs method for establishing matrix concentration is also based on an elementary, but conceptually challenging, analog of integration by parts [MJC⁺14, Lem. 2.4]. Aside from these works, we are not aware of any application of related techniques to prove results on matrix concentration.

6.1. Preliminaries. To obtain Lemma 6.1, the main auxiliary tool is the classical integration by parts formula for a function of a standard normal vector [NP12, Lem. 1.1.1]. In the form required here, the result can be derived with basic calculus.

Fact 6.2 (Gaussian Integration by Parts). *Let $\boldsymbol{\gamma} \in \mathbb{R}^n$ be a vector with independent standard normal entries, and let $f: \mathbb{R}^n \rightarrow \mathbb{R}$ be a function whose derivative is absolutely integrable with respect to the standard normal measure. Then*

$$\sum_{i=1}^n \mathbb{E} [\gamma_i \cdot f(\boldsymbol{\gamma})] = \sum_{i=1}^n \mathbb{E} [(\partial_i f)(\boldsymbol{\gamma})].$$

The symbol ∂_i denotes differentiation with respect to the i th coordinate.

We also use a well-known formula for the derivative of a matrix power [Bha97, Sec. X.4].

Fact 6.3 (Derivative of a Matrix Power). *Let $\mathbf{A} : \mathbb{R} \rightarrow \mathbb{M}_d$ be a differentiable function. For each integer $\rho \geq 1$,*

$$\frac{d}{du} (\mathbf{A}(u)^\rho) = \sum_{k=0}^{\rho-1} \mathbf{A}(u)^k \cdot \frac{d}{du} \mathbf{A}(u) \cdot \mathbf{A}(u)^{\rho-1-k}. \quad (6.2)$$

In particular,

$$\frac{d}{du} \operatorname{tr} \mathbf{A}(u)^\rho = \rho \cdot \operatorname{tr} \left[\mathbf{A}(u)^{\rho-1} \cdot \frac{d}{du} \mathbf{A}(u) \right].$$

The symbol \cdot refers to ordinary matrix multiplication.

6.2. Proof of Lemma 6.1. Let us treat the random matrix \mathbf{X} as a matrix-valued function of the standard normal vector $\boldsymbol{\gamma} := (\gamma_1, \dots, \gamma_n)$. That is,

$$\mathbf{X} = \mathbf{X}(\boldsymbol{\gamma}) = \sum_{i=1}^n \gamma_i \mathbf{H}_i.$$

Write $\mathbf{X} = \mathbf{X} \cdot \mathbf{X}^{2p-1}$ and distribute the sum in the first factor:

$$\mathbb{E} \operatorname{tr} \mathbf{X}^{2p} = \mathbb{E} \operatorname{tr} \left[\left(\sum_{i=1}^n \gamma_i \mathbf{H}_i \right) \mathbf{X}^{2p-1} \right] = \sum_{i=1}^n \mathbb{E} \left[\gamma_i \cdot \operatorname{tr} \left[\mathbf{H}_i \mathbf{X}^{2p-1} \right] \right]$$

The Gaussian integration by parts formula, Fact 6.2, implies that

$$\mathbb{E} \operatorname{tr} \mathbf{X}^{2p} = \sum_{i=1}^n \mathbb{E} \operatorname{tr} \left[\mathbf{H}_i \cdot \partial_i (\mathbf{X}^{2p-1}) \right].$$

Since $\partial_i \mathbf{X} = \mathbf{H}_i$, the derivative formula (6.2) yields

$$\mathbb{E} \operatorname{tr} \mathbf{X}^{2p} = \sum_{i=1}^n \operatorname{tr} \left[\mathbf{H}_i \cdot \sum_{q=0}^{2p-2} \mathbf{X}^q \mathbf{H}_i \mathbf{X}^{2p-2-q} \right] = \sum_{q=0}^{2p-2} \sum_{i=1}^n \operatorname{tr} \left[\mathbf{H}_i \mathbf{X}^q \mathbf{H}_i \mathbf{X}^{2p-2-q} \right].$$

This completes the proof of the formula (6.1).

7. A SHORT PROOF OF THE MATRIX KHINTCHINE INEQUALITY

Historically, proofs of the matrix Khintchine inequality have been rather complicated, but the result is actually an immediate consequence of Lemma 6.1. We will present this argument in detail because it has not appeared in the literature. Furthermore, the approach serves as a template for the more sophisticated theorems that are the main contributions of this paper. Let us restate Proposition 2.3 in the form that we will establish it.

Proposition 7.1 (Matrix Khintchine). *Let $\mathbf{X} := \sum_{i=1}^n \gamma_i \mathbf{H}_i$ be an Hermitian matrix Gaussian series, as in (2.1). Define the matrix variance and standard deviation parameters*

$$\mathbf{V} := \mathbf{Var}(\mathbf{X}) = \sum_{i=1}^n \mathbf{H}_i^2 \quad \text{and} \quad \sigma_{2q} := (\operatorname{tr} \mathbf{V}^q)^{1/(2q)} \quad \text{for each } q \geq 1. \quad (7.1)$$

Then, for each integer $p \geq 1$,

$$(\mathbb{E} \operatorname{tr} \mathbf{X}^{2p})^{1/(2p)} \leq \sqrt{2p-1} \cdot \sigma_{2p}.$$

The short proof of Proposition 7.1 appears in the next two sections. The approach parallels the exchangeable pairs method that has been used to establish the matrix Khintchine inequality for Rademacher series [MJC⁺14, Cor. 7.3]. Here, we replace exchangeable pairs with the conceptually simpler argument based on Gaussian integration by parts. To reach the statement of Proposition 2.3, we simply rewrite the trace in terms of a Schatten norm.

Remark 7.2 (Noninteger Moments). Our proof of Proposition 7.1 can be adapted to obtain moment bounds for all $p \geq 2$. See [MJC⁺14, Cor. 7.3] for a closely related argument.

7.1. Preliminaries. The main idea in the proof is to simplify the trace moment identity (6.1) with an elementary matrix inequality. Anticipating subsequent arguments, we state the inequality in greater generality than we need right now.

Proposition 7.3. *Suppose that \mathbf{H} and \mathbf{A} are Hermitian matrices of the same size. Let q and r be integers that satisfy $0 \leq q \leq r$. For each real number s in the range $0 \leq s \leq \min\{q, r-q\}$,*

$$\operatorname{tr} [\mathbf{H} \mathbf{A}^q \mathbf{H} \mathbf{A}^{r-q}] \leq \operatorname{tr} [\mathbf{H} |\mathbf{A}|^s \mathbf{H} |\mathbf{A}|^{r-s}].$$

The proof of Proposition 7.3 depends on a numerical fact. For nonnegative numbers α and β , the function $\theta \mapsto \alpha^\theta \beta^{1-\theta} + \alpha^{1-\theta} \beta^\theta$ is convex on the interval $[0, 1]$, and it achieves its minimum at $\theta = \frac{1}{2}$. Therefore,

$$\alpha^\theta \beta^{1-\theta} + \alpha^{1-\theta} \beta^\theta \leq \alpha^{\theta'} \beta^{1-\theta'} + \alpha^{1-\theta'} \beta^{\theta'} \quad \text{when } 0 \leq \theta' \leq \min\{\theta, 1-\theta\}. \quad (7.2)$$

We need to lift this scalar inequality to matrices.

Proof. Without loss of generality, we may change coordinates so that \mathbf{A} is diagonal: $\mathbf{A} = \sum_i a_i \mathbf{E}_{ii}$. Expanding both copies of \mathbf{A} ,

$$\mathrm{tr}[\mathbf{H}\mathbf{A}^q\mathbf{H}\mathbf{A}^{r-q}] = \sum_{i,j} a_i^q a_j^{r-q} \cdot \mathrm{tr}[\mathbf{H}\mathbf{E}_{ii}\mathbf{H}\mathbf{E}_{jj}] = \sum_{i,j} \frac{1}{2} (a_i^q a_j^{r-q} + a_i^{r-q} a_j^q) \cdot \mathrm{tr}[\mathbf{H}\mathbf{E}_{ii}\mathbf{H}\mathbf{E}_{jj}].$$

After we take absolute values, the inequality (7.2) implies that

$$a_i^q a_j^{r-q} + a_i^{r-q} a_j^q \leq |a_i|^s |a_j|^{r-s} + |a_i|^{r-s} |a_j|^s.$$

The remaining trace is nonnegative: $\mathrm{tr}[\mathbf{H}\mathbf{E}_{ii}\mathbf{H}\mathbf{E}_{jj}] = |h_{ij}|^2$, where h_{ij} are the components of the matrix \mathbf{H} . As a consequence,

$$\mathrm{tr}[\mathbf{H}\mathbf{A}^q\mathbf{H}\mathbf{A}^{r-q}] \leq \sum_{i,j} \frac{1}{2} (|a_i|^s |a_j|^{r-s} + |a_i|^{r-s} |a_j|^s) \cdot \mathrm{tr}[\mathbf{H}\mathbf{E}_{ii}\mathbf{H}\mathbf{E}_{jj}] = \mathrm{tr}[\mathbf{H}|\mathbf{A}|^s \mathbf{H}|\mathbf{A}|^{r-s}].$$

To reach the last identity, we reversed our steps to reassemble the sum into a trace. \square

7.2. Proof of the Matrix Khintchine Inequality. We may now establish Proposition 7.1. Let us introduce notation for the quantity of interest:

$$E^{2p} := \mathbb{E} \mathrm{tr} \mathbf{X}^{2p}.$$

Use the integration by parts result, Lemma 6.1, to rewrite the trace moment:

$$E^{2p} = \sum_{q=0}^{2p-2} \sum_{i=1}^n \mathbb{E} \mathrm{tr}[\mathbf{H}_i \mathbf{X}^q \mathbf{H}_i \mathbf{X}^{2p-2-q}].$$

For each choice of q , apply the matrix inequality from Proposition 7.3 with $r = 2p - 2$ and $s = 0$ to reach

$$E^{2p} \leq (2p-1) \sum_{i=1}^n \mathbb{E} \mathrm{tr}[\mathbf{H}_i^2 \mathbf{X}^{2(p-1)}] = (2p-1) \cdot \mathbb{E} \mathrm{tr}[\mathbf{V} \mathbf{X}^{2(p-1)}]$$

We have identified the matrix variance \mathbf{V} defined in (7.1).

Next, let us identify a copy of E on the right-hand side and solve the resulting algebraic inequality. To that end, invoke Hölder's inequality (5.1) for the trace with $\rho = p$ and $\rho' = p/(p-1)$:

$$\begin{aligned} E^{2p} &\leq (2p-1) \cdot (\mathrm{tr} \mathbf{V}^p)^{1/p} \cdot \mathbb{E} (\mathrm{tr} \mathbf{X}^{2p})^{(p-1)/p} \\ &\leq (2p-1) \cdot \sigma_{2p}^2 \cdot (\mathbb{E} \mathrm{tr} \mathbf{X}^{2p})^{(p-1)/p} = (2p-1) \cdot \sigma_{2p}^2 \cdot E^{2(p-1)}. \end{aligned}$$

We have identified the quantity σ_{2p} from (7.1). The second inequality is Lyapunov's. Since the unknown E is nonnegative, we can solve the polynomial inequality to reach

$$E \leq \sqrt{2p-1} \cdot \sigma_{2p}.$$

This is the required result.

8. A SECOND-ORDER MATRIX KHINTCHINE INEQUALITY

In this Section, we prove Theorem 3.5, the second-order matrix Khintchine inequality. Let us restate the result in the form that we will establish it.

Theorem 8.1 (Second-Order Matrix Khintchine). *Let $\mathbf{X} = \sum_{i=1}^n \gamma_i \mathbf{H}_i$ be an Hermitian matrix Gaussian series, as in (2.1). Define the matrix variance and standard deviation parameter*

$$\mathbf{V} := \mathbf{Var}(\mathbf{X}) = \sum_{i=1}^n \mathbf{H}_i^2 \quad \text{and} \quad \sigma_{2p} := (\mathrm{tr} \mathbf{V}^p)^{1/(2p)} \quad \text{for } p \geq 1. \quad (8.1)$$

Define the matrix alignment parameter

$$w_{2p} := \max_{\mathbf{Q}_\ell} \left(\mathrm{tr} \left| \sum_{i,j=1}^n \mathbf{H}_i \mathbf{Q}_1 \mathbf{H}_j \mathbf{Q}_2 \mathbf{H}_i \mathbf{Q}_3 \mathbf{H}_j \right|^{p/2} \right)^{1/(2p)} \quad \text{for } p \geq 1 \quad (8.2)$$

where the maximum ranges over a triple $(\mathbf{Q}_1, \mathbf{Q}_2, \mathbf{Q}_3)$ of unitary matrices. Then, for each integer $p \geq 3$,

$$(\mathbb{E} \mathrm{tr} \mathbf{X}^{2p})^{1/(2p)} \leq 3^4 \sqrt{2p-5} \cdot \sigma_{2p} + \sqrt{2p-4} \cdot w_{2p}. \quad (8.3)$$

The proof of Theorem 8.1 will occupy us for the rest of the section. To reach the statement in the introduction, we rewrite traces in terms of Schatten norms. We also provide the proof of Proposition 3.2 in Section 8.9.

8.1. Discussion. Before we establish Theorem 8.1, let us spend a moment to discuss the proof of this result. Theorem 8.1 is based on the same pattern of argument as the matrix Khintchine inequality, Proposition 7.1. This time, we apply Proposition 7.3 more surgically to control the terms in the trace moment identity from Lemma 6.1. The most significant new observation is that we can use complex interpolation to reorganize the products of matrices that arise during the calculation.

We can refine this argument in several ways. First, if we apply complex interpolation with more care, it is possible to define the matrix alignment parameter (8.2) as a maximum over the set

$$\{\mathbf{Q}_1, \mathbf{Q}_2, \mathbf{Q}_3 \text{ are commuting unitaries and } \mathbf{Q}_\ell = \mathbf{I} \text{ for some } \ell\}.$$

Given that commuting matrices are simultaneously diagonalizable, this improvement might make it easier to bound the matrix alignment parameters.

Second, it is quite clear from the proof that we can proceed beyond the second-order terms. For example, for an integer $p \geq 3$, we can obtain results in terms of the third-order quantities

$$w_{2p,1} := \max_{\mathbf{Q}_\ell} \left(\text{tr} \left| \sum_{i,j,k=1}^n \mathbf{H}_i \mathbf{Q}_1 \mathbf{H}_j \mathbf{Q}_2 \mathbf{H}_k \mathbf{Q}_3 \mathbf{H}_i \mathbf{Q}_4 \mathbf{H}_j \mathbf{Q}_5 \mathbf{H}_k \right|^{p/3} \right)^{1/(2p)}$$

$$w_{2p,2} := \max_{\mathbf{Q}_\ell} \left(\text{tr} \left| \sum_{i,j,k=1}^n \mathbf{H}_i \mathbf{Q}_1 \mathbf{H}_j \mathbf{Q}_2 \mathbf{H}_k \mathbf{Q}_3 \mathbf{H}_i \mathbf{Q}_4 \mathbf{H}_k \mathbf{Q}_5 \mathbf{H}_j \right|^{p/3} \right)^{1/(2p)}.$$

The ordering of indices is (i, j, k, i, j, k) and (i, j, k, i, k, j) , respectively. This refinement allows us reduce the order of coefficient on the standard deviation term σ_{2p} in (8.3) to $p^{1/6}$. Unfortunately, we must also compute both alignment parameters $w_{2p,1}$ and $w_{2p,2}$, instead of just w_{2p} . This observation shows why it is unproductive to press forward with this approach. Indeed, the number of orderings of indices grows super-exponentially as we consider longer products, which is an awful prospect for applications.

8.2. Preliminaries. In the proof of Theorem 8.1, we will use two interpolation results to reorganize products of matrices. The first one is a type of matrix Hölder inequality [LP86, Cor. 1]. Here is a version of the result specialized to our setting.

Fact 8.2 (Lust-Piquard). *Consider a finite sequence $(\mathbf{A}_1, \dots, \mathbf{A}_n)$ of Hermitian matrices with the same dimension, and let \mathbf{B} be a positive-semidefinite matrix of the same dimension. For each number $\rho \geq 2$,*

$$\left(\text{tr} \left(\sum_{i=1}^n \mathbf{A}_i \mathbf{B} \mathbf{A}_i \right)^{\rho/2} \right)^{2/\rho} \leq \left(\text{tr} \left(\sum_{i=1}^n \mathbf{A}_i^2 \right)^\rho \right)^{1/\rho} \cdot \left(\text{tr} \mathbf{B}^\rho \right)^{1/\rho}.$$

See [PX97, Lem. 1.1] for a proof based on the Hadamard Three-Lines Theorem [Gar07, Prop. 9.1.1].

The second result is a more complicated interpolation for a multilinear function whose arguments are powers of random matrices.

Proposition 8.3 (Multilinear Interpolation). *Suppose that $F: (\mathbb{M}_d)^k \rightarrow \mathbb{C}$ is a multilinear function. Fix nonnegative integers $\alpha_1, \dots, \alpha_k$ with $\sum_{i=1}^k \alpha_i = \alpha$. Let $\mathbf{Y}_i \in \mathbb{H}_d$ be random matrices, not necessarily independent, for which $\mathbb{E} \|\mathbf{Y}_i\|^\alpha < \infty$. Then*

$$\left| \mathbb{E} F(\mathbf{Y}_1^{\alpha_1}, \dots, \mathbf{Y}_k^{\alpha_k}) \right| \leq \max_{i=1, \dots, k} \mathbb{E} \max_{\mathbf{Q}_\ell} |F(\mathbf{Q}_1, \dots, \mathbf{Q}_{i-1}, \mathbf{Q}_i \mathbf{Y}_i^\alpha, \mathbf{Q}_{i+1}, \dots, \mathbf{Q}_k)|.$$

In this expression, each \mathbf{Q}_ℓ is a (random) unitary matrix that commutes with \mathbf{Y}_ℓ .

As with Fact 8.2, the proof of Proposition 8.3 depends on the Hadamard Three-Lines Theorem [Gar07, Prop. 9.1.1]. The argument is standard but somewhat involved, so we postpone the details to Appendix A.

8.3. The Overture. Let us commence with the proof of Theorem 8.1. The initial steps are similar with the argument that leads to the matrix Khintchine inequality, Proposition 7.1. Introduce notation for the quantity of interest:

$$E^{2p} := \mathbb{E} \text{tr} \mathbf{X}^{2p} = \sum_{q=0}^{2p-2} \sum_{i=1}^n \mathbb{E} \text{tr} [\mathbf{H}_i \mathbf{X}^q \mathbf{H}_i \mathbf{X}^{2p-2-q}]. \quad (8.4)$$

The identity follows from the integration by parts result, Lemma 6.1.

This time, we make finer estimates for the summands in (8.4). Apply Proposition 7.3 with $s = 0$ to the terms where $q \in \{0, 1, 2p-3, 2p-2\}$. For the remaining $2p-5$ values of the exponent q , apply Proposition 7.3 with $s = 1$. We reach the bound

$$E^{2p} \leq 4 \sum_{i=1}^n \mathbb{E} \text{tr} [\mathbf{H}_i^2 \mathbf{X}^{2p-2}] + (2p-5) \sum_{i=1}^n \mathbb{E} \text{tr} [\mathbf{H}_i \mathbf{X}^2 \mathbf{H}_i \mathbf{X}^{2p-4}]. \quad (8.5)$$

We can take advantage of the fact that the \mathbf{H}_i are interleaved with the powers \mathbf{X}^r of the random matrix in the second term.

8.4. The First Term. To treat the first term on the right-hand side of (8.5), simply repeat the arguments from Section 7.2 to obtain a bound in terms of the quantity E . We have

$$\sum_{i=1}^n \mathbb{E} \operatorname{tr} [\mathbf{H}_i^2 \mathbf{X}^{2p-2}] = \mathbb{E} \operatorname{tr} [\mathbf{V} \mathbf{X}^{2p-2}] \leq (\operatorname{tr} \mathbf{V}^p)^{1/p} \cdot (\mathbb{E} \operatorname{tr} \mathbf{X}^{2p})^{2(p-1)} = \sigma_{2p}^2 \cdot E^{2(p-1)}. \quad (8.6)$$

The quantities \mathbf{V} and σ_{2p} are defined in (8.1), and we have identified a copy of E .

8.5. Integration by Parts, Again. To continue, we want to break down the matrix \mathbf{X}^2 that appears in the second term on the right-hand side of (8.5). To do so, we perform another Gaussian integration by parts. Write $\mathbf{X}^2 = \sum_{j=1}^n \gamma_j \mathbf{H}_j \mathbf{X}$, and invoke Fact 6.2 to obtain

$$\begin{aligned} \sum_{i=1}^n \mathbb{E} \operatorname{tr} [\mathbf{H}_i \mathbf{X}^2 \mathbf{H}_i \mathbf{X}^{2p-4}] &= \sum_{i,j=1}^n \mathbb{E} [\gamma_j \cdot \operatorname{tr} [\mathbf{H}_i \mathbf{H}_j \mathbf{X} \mathbf{H}_i \mathbf{X}^{2p-4}]] \\ &= \sum_{i,j=1}^n \mathbb{E} \operatorname{tr} [\mathbf{H}_i \mathbf{H}_j^2 \mathbf{H}_i \mathbf{X}^{2p-4}] + \sum_{i,j=1}^n \mathbb{E} \operatorname{tr} \left[\mathbf{H}_i \mathbf{H}_j \mathbf{X} \mathbf{H}_i \left(\sum_{r=0}^{2p-5} \mathbf{X}^r \mathbf{H}_j \mathbf{X}^{2p-5-r} \right) \right]. \end{aligned} \quad (8.7)$$

This result follows from the product rule and the formula (6.2) for the derivative of a power. We will bound the first term on the right-hand side of (8.7) in terms of the standard deviation parameter σ_{2p} , and the second term will lead to the matrix alignment parameter w_{2p} .

8.6. Finding the Standard Deviation Parameter. Let us address the first term on the right-hand side of (8.7). First, draw the sum back into the trace and identify the matrix variance \mathbf{V} , defined in (8.1):

$$\sum_{i,j=1}^n \mathbb{E} \operatorname{tr} [\mathbf{H}_i \mathbf{H}_j^2 \mathbf{H}_i \mathbf{X}^{2p-4}] = \mathbb{E} \operatorname{tr} \left[\left(\sum_{i=1}^n \mathbf{H}_i \mathbf{V} \mathbf{H}_i \right) \mathbf{X}^{2(p-2)} \right].$$

To isolate the random matrix \mathbf{X} , apply Hölder's inequality (5.1) with exponents $\varrho = p/2$ and $\varrho' = p/(p-2)$, and follow up with Lyapunov's inequality. Thus,

$$\mathbb{E} \operatorname{tr} \left[\left(\sum_{i=1}^n \mathbf{H}_i \mathbf{V} \mathbf{H}_i \right) \mathbf{X}^{2(p-2)} \right] \leq \left(\operatorname{tr} \left(\sum_{i=1}^n \mathbf{H}_i \mathbf{V} \mathbf{H}_i \right)^{p/2} \right)^{2/p} \cdot (\mathbb{E} \operatorname{tr} \mathbf{X}^{2p})^{(p-2)/p}.$$

The Lust-Piquard inequality, Fact 8.2, with $\varrho = p$ implies that

$$\left(\operatorname{tr} \left(\sum_{i=1}^n \mathbf{H}_i \mathbf{V} \mathbf{H}_i \right)^{p/2} \right)^{2/p} \leq \left(\operatorname{tr} \left(\sum_{i=1}^n \mathbf{H}_i^2 \right)^p \right)^{1/p} \cdot (\operatorname{tr} \mathbf{V}^p)^{1/p} = (\operatorname{tr} \mathbf{V}^p)^{2/p} = \sigma_{2p}^4.$$

Once again, we identified \mathbf{V} and σ_{2p} from (8.1). Combine the last three displays to arrive at

$$\sum_{i,j=1}^n \mathbb{E} \operatorname{tr} [\mathbf{H}_i \mathbf{H}_j^2 \mathbf{H}_i \mathbf{X}^{2p-4}] \leq \sigma_{2p}^4 \cdot (\mathbb{E} \operatorname{tr} \mathbf{X}^{2p})^{(p-2)/p} = \sigma_{2p}^4 \cdot E^{2(p-2)}. \quad (8.8)$$

We have identified another copy of E .

8.7. Finding the Matrix Alignment Parameter. It remains to study the second term on the right-hand side of (8.7). Rearranging the sums, we write this object as

$$\sum_{r=0}^{2p-5} \sum_{i,j=1}^n \mathbb{E} \operatorname{tr} [\mathbf{H}_i \mathbf{H}_j \mathbf{X} \mathbf{H}_i \mathbf{X}^r \mathbf{H}_j \mathbf{X}^{2p-5-r}].$$

We can apply the interpolation result, Proposition 8.3, to consolidate the powers of the random matrix \mathbf{X} . Consider the multilinear function

$$F(\mathbf{A}_1, \mathbf{A}_2, \mathbf{A}_3) := \sum_{i,j=1}^n \operatorname{tr} [\mathbf{H}_i \mathbf{H}_j \mathbf{A}_1 \mathbf{H}_i \mathbf{A}_2 \mathbf{H}_j \mathbf{A}_3].$$

Since \mathbf{X} is a matrix Gaussian series, it has moments of all orders. Therefore, for each index r ,

$$\begin{aligned} &|\mathbb{E} \operatorname{tr} [\mathbf{H}_i \mathbf{H}_j \mathbf{X} \mathbf{H}_i \mathbf{X}^r \mathbf{H}_j \mathbf{X}^{2p-5-r}]| \\ &\leq \max \left\{ \mathbb{E} \max_{\mathbf{Q}_\ell} |F(\mathbf{Q}_1 \mathbf{X}^{2p-4}, \mathbf{Q}_2, \mathbf{Q}_3)|, \mathbb{E} \max_{\mathbf{Q}_\ell} |F(\mathbf{Q}_1, \mathbf{Q}_2 \mathbf{X}^{2p-4}, \mathbf{Q}_3)|, \mathbb{E} \max_{\mathbf{Q}_\ell} |F(\mathbf{Q}_1, \mathbf{Q}_2, \mathbf{Q}_3 \mathbf{X}^{2p-4})| \right\}. \end{aligned}$$

All three terms in the maximum admit the same bound, so we may as well consider the third one:

$$\begin{aligned}
\mathbb{E} \max_{\mathbf{Q}_\ell} |F(\mathbf{Q}_1, \mathbf{Q}_2, \mathbf{Q}_3 \mathbf{X}^{2p-4})| &= \mathbb{E} \max_{\mathbf{Q}_\ell} \left| \sum_{i,j=1}^n \text{tr} [\mathbf{H}_i \mathbf{H}_j \mathbf{Q}_1 \mathbf{H}_i \mathbf{Q}_2 \mathbf{H}_j \mathbf{Q}_3 \mathbf{X}^{2(p-2)}] \right| \\
&= \mathbb{E} \max_{\mathbf{Q}_\ell} \left| \text{tr} \left[\left(\sum_{i,j=1}^n \mathbf{Q}_3 \mathbf{H}_i \mathbf{H}_j \mathbf{Q}_1 \mathbf{H}_i \mathbf{Q}_2 \mathbf{H}_j \right) \mathbf{X}^{2(p-2)} \right] \right| \\
&\leq \mathbb{E} \max_{\mathbf{Q}_\ell} \left[\left(\text{tr} \left| \sum_{i,j=1}^n \mathbf{H}_i \mathbf{H}_j \mathbf{Q}_1 \mathbf{H}_i \mathbf{Q}_2 \mathbf{H}_j \right|^{p/2} \right)^{2/p} \cdot (\text{tr} \mathbf{X}^{2p})^{(p-2)/p} \right] \\
&\leq \max_{\mathbf{Q}_\ell} \left(\text{tr} \left| \sum_{i,j=1}^n \mathbf{H}_i \mathbf{H}_j \mathbf{Q}_1 \mathbf{H}_i \mathbf{Q}_2 \mathbf{H}_j \right|^{p/2} \right)^{2/p} \cdot (\mathbb{E} \text{tr} \mathbf{X}^{2p})^{(p-2)/p} \\
&\leq w_{2p}^4 \cdot E^{2(p-2)}.
\end{aligned}$$

The first step is the definition of F . To reach the second line, we use the fact that \mathbf{Q}_3 commutes with \mathbf{X} , then we cycle the trace. The third line is Hölder's inequality (5.1) with $\varrho = p/2$ and $\varrho' = p/(p-2)$, and we have used the left unitary invariance of the matrix absolute value to delete \mathbf{Q}_3 . Next, take the maximum over all unitary matrices, and apply Lyapunov's inequality to draw the expectation into the term involving \mathbf{X} . Finally, identify the quantity E and note that the maximum is bounded by the alignment parameter w_{2p}^4 , defined in (8.2). Similar calculations are valid for the other two terms, whence

$$|\mathbb{E} \text{tr} [\mathbf{H}_i \mathbf{H}_j \mathbf{X} \mathbf{H}_i \mathbf{X}^r \mathbf{H}_j \mathbf{X}^{2p-5-r}]| \leq w_{2p}^4 \cdot E^{2(p-2)}.$$

Since there are $2p-4$ possible choices of r , we determine that

$$\sum_{r=0}^{2p-5} \sum_{i,j=1}^n \mathbb{E} \text{tr} [\mathbf{H}_i \mathbf{H}_j \mathbf{X} \mathbf{H}_i \mathbf{X}^r \mathbf{H}_j \mathbf{X}^{2p-5-r}] \leq (2p-4) \cdot w_{2p}^4 \cdot E^{2(p-2)}. \quad (8.9)$$

The main part of the argument is finished.

8.8. Putting the Pieces Together. To conclude, we merge the bounds we have obtained and solve the resulting inequality for the quantity E . Combine (8.5), (8.6), (8.7), (8.8), and (8.9) to reach

$$E^{2p} \leq 4\sigma_{2p}^2 \cdot E^{2(p-1)} + (2p-5) \cdot \left[\sigma_{2p}^4 + (2p-4) \cdot w_{2p}^4 \right] \cdot E^{2(p-2)}.$$

Clearing factors of E , we reach the inequality

$$E^4 \leq 4\sigma_{2p}^2 \cdot E^2 + (2p-5) \cdot \left[\sigma_{2p}^4 + (2p-4) \cdot w_{2p}^4 \right].$$

If α and β are nonnegative numbers, each nonnegative solution to the quadratic inequality $t^2 \leq \alpha t + \beta$ must satisfy $t \leq \alpha + \sqrt{\beta}$. It follows that

$$E^2 \leq 4\sigma_{2p}^2 + \sqrt{2p-5} \cdot \left[\sigma_{2p}^4 + (2p-4) \cdot w_{2p}^4 \right]^{1/2}.$$

Take the square root, and invoke subadditivity of the square root (twice) to reach

$$E \leq \left(2 + \sqrt[4]{2p-5} \right) \cdot \sigma_{2p} + \sqrt[4]{(2p-5)(2p-4)} \cdot w_{2p}.$$

Finally, we simplify the numerical constants to arrive at (8.3).

8.9. Comparison of Standard Deviation and Alignment Parameters. Our last task in this section is to establish Proposition 3.2, which states that the alignment parameter w_{2p} never exceeds the standard deviation σ_{2p} . The easiest way to obtain this result is to use block matrices and inequalities for the Schatten norm.

Fix an integer $p \geq 2$, and fix a triple $(\mathbf{Q}_1, \mathbf{Q}_2, \mathbf{Q}_3)$ of unitary matrices. Consider the quantity

$$S := \left(\text{tr} \left| \sum_{i,j=1}^n \mathbf{H}_i \mathbf{Q}_1 \mathbf{H}_j \mathbf{Q}_2 \mathbf{H}_i \mathbf{Q}_3 \mathbf{H}_j \right|^{p/2} \right)^{2/p}.$$

To establish Proposition 3.2, it suffices to show that $S \leq \sigma_{2p}^4$. Using block matrices and converting the trace into a Schatten norm, we can write

$$S = \left\| \left[\begin{array}{c} \vdots \\ \mathbf{Q}_2^* \mathbf{H}_j \mathbf{Q}_1^* \mathbf{H}_i \\ \vdots \end{array} \right]^* \left[\begin{array}{c} \vdots \\ \mathbf{H}_i \mathbf{Q}_3 \mathbf{H}_j \\ \vdots \end{array} \right] \right\|_{p/2}.$$

The entries of the block column matrices are indexed by pairs (i, j) , arranged in lexicographic order. Invoke the Cauchy–Schwarz inequality (5.2) for Schatten norms with $\rho = p/2$:

$$S^2 \leq \left\| \begin{bmatrix} \vdots \\ \mathbf{Q}_2^* \mathbf{H}_j \mathbf{Q}_1^* \mathbf{H}_i \\ \vdots \end{bmatrix}^* \begin{bmatrix} \vdots \\ \mathbf{Q}_2^* \mathbf{H}_j \mathbf{Q}_1^* \mathbf{H}_i \\ \vdots \end{bmatrix} \right\|_{p/2} \times \left\| \begin{bmatrix} \vdots \\ \mathbf{H}_i \mathbf{Q}_3 \mathbf{H}_j \\ \vdots \end{bmatrix}^* \begin{bmatrix} \vdots \\ \mathbf{H}_i \mathbf{Q}_3 \mathbf{H}_j \\ \vdots \end{bmatrix} \right\|_{p/2}.$$

Write each product of two block matrices as a sum:

$$S^2 \leq \left\| \sum_{i,j=1}^n \mathbf{H}_i \mathbf{Q}_1^* \mathbf{H}_j^2 \mathbf{Q}_1 \mathbf{H}_i \right\|_{p/2} \times \left\| \sum_{i,j=1}^n \mathbf{H}_j \mathbf{Q}_3^* \mathbf{H}_i^2 \mathbf{Q}_3 \mathbf{H}_j \right\|_{p/2}.$$

The two factors have the same form, so it suffices to bound the first one. Indeed,

$$\left\| \sum_{i,j=1}^n \mathbf{H}_i \mathbf{Q}_1^* \mathbf{H}_j^2 \mathbf{Q}_1 \mathbf{H}_i \right\|_{p/2} = \left\| \sum_{i=1}^n \mathbf{H}_i \mathbf{Q}_1^* \mathbf{V} \mathbf{Q}_1 \mathbf{H}_i \right\|_{p/2} \leq \left\| \sum_{i=1}^n \mathbf{H}_i^2 \right\|_p \cdot \left\| \mathbf{Q}_1^* \mathbf{V} \mathbf{Q}_1 \right\|_p = \|\mathbf{V}\|_p^2 = \sigma_{2p}^4.$$

We have identified the matrix variance \mathbf{V} , defined in (8.1), and then we applied the Lust–Piquard inequality, Fact 8.2, with $\rho = p$. We identified \mathbf{V} again, invoked unitary invariance of the Schatten norm, and then we recognized the quantity σ_{2p} from (8.1). In summary, we have established that $S^2 \leq \sigma_{2p}^8$. This is what we needed to show.

9. SECOND-ORDER MATRIX KHINTCHINE UNDER STRONG ISOTROPY

In this section, we prove an extension of Theorem 3.11 that gives both lower and upper bounds for the trace moments of a strongly isotropic matrix Gaussian series.

Theorem 9.1 (Second-Order Matrix Khintchine under Strong Isotropy). *Let $\mathbf{X} := \sum_{i=1}^n \gamma_i \mathbf{H}_i$ be an Hermitian matrix Gaussian series, as in (2.1). Assume that \mathbf{X} has the strong isotropy property*

$$\mathbb{E} \mathbf{X}^p = (\mathbb{E} \bar{\text{tr}} \mathbf{X}^p) \cdot \mathbf{I} \quad \text{for each } p = 0, 1, 2, \dots \quad (9.1)$$

Define the matrix standard deviation parameter and matrix alignment parameter

$$\sigma := \left\| \sum_{i=1}^n \mathbf{H}_i^2 \right\|^{1/2} \quad \text{and} \quad w := \max_{\mathbf{Q}_e} \left\| \sum_{i,j=1}^n \mathbf{H}_i \mathbf{Q}_1 \mathbf{H}_j \mathbf{Q}_2 \mathbf{H}_i \mathbf{Q}_3 \mathbf{H}_j \right\|^{1/4}. \quad (9.2)$$

The maximum ranges over a triple $(\mathbf{Q}_1, \mathbf{Q}_2, \mathbf{Q}_3)$ of unitary matrices. Then, for each integer $p \geq 1$,

$$\text{Cat}_p^{1/(2p)} \cdot \sigma \cdot [1 - (pw/\sigma)^4]_+^{1/(2p)} \leq (\mathbb{E} \bar{\text{tr}} \mathbf{X}^{2p})^{1/(2p)} \leq \text{Cat}_p^{1/(2p)} \cdot \sigma + 2^{1/4} p^{5/4} \cdot w.$$

The lower bound also requires that $p^{7/4} w \leq 0.7\sigma$. We have written Cat_p for the p th Catalan number, the function $[a]_+ := \max\{a, 0\}$, and $\bar{\text{tr}}$ is the normalized trace.

The proof of this result appears below, starting in Section 9.3. To reach the statement of Theorem 3.11 in the introduction, we rewrite normalized traces in terms of Schatten norms. Fact 9.2 (below) states that the Catalan numbers satisfy the bound $\text{Cat}_p \leq 4^p$ for each $p = 1, 2, 3, \dots$, which gives an explicit numerical form for the upper bound.

9.1. Discussion. Before we establish Theorem 9.1, let us comment on the proof and the meaning of this result. The most important observation is that the estimate is extremely accurate, at least for some examples. In particular, for the GOE matrix \mathbf{X}_{goe} defined in (2.7), we showed in Section 4.2 that the standard deviation parameter $\sigma \approx 1$ while the alignment parameter $w \approx d^{-1/4}$. Therefore, Theorem 9.1 implies that

$$\mathbb{E} \bar{\text{tr}} \mathbf{X}_{\text{goe}}^{2p} \approx \text{Cat}_p \quad \text{when } p \ll d^{1/7}.$$

This estimate is sufficient to prove that the limiting spectral distribution of the GOE is the semicircle law. See [Tao12, Sec. 2.3] for details about how to derive the law from the trace moments. Furthermore, Markov’s inequality implies that the norm $\|\mathbf{X}_{\text{goe}}\| \approx 2$ with high probability.

The proof of Theorem 9.1 has a lot in common with the arguments leading up to Proposition 7.1 and Theorem 8.1. The main innovation is that we can use the strong isotropy to imitate a moment identity that would hold in free probability. This idea allows us to remove the dependence on p from the standard deviation term.

Although it may seem that the proof requires the matrix \mathbf{X} to be a Gaussian series, there are analogous techniques, based on the theory of exchangeable pairs [MJC⁺14], that allow us to deal with other types of random matrix series. This observation has the potential to lead to universality laws. It is also clear from the argument that we could prove related results with an approximate form of strong isotropy.

In addition, it is possible to extend these ideas to a rectangular matrix Gaussian series $\mathbf{Z} := \sum_{i=1}^n \gamma_i \mathbf{S}_i \in \mathbb{C}^{d_1 \times d_2}$. In this case, we consider the Hermitian dilation:

$$\mathcal{H}(\mathbf{Z}) = \begin{bmatrix} \mathbf{0} & \mathbf{Z} \\ \mathbf{Z}^* & \mathbf{0} \end{bmatrix}.$$

The correct analog of strong isotropy is that

$$\mathbb{E} \mathcal{H}(\mathbf{Z})^{2p} = \begin{bmatrix} (\mathbb{E} \operatorname{tr}(\mathbf{Z}\mathbf{Z}^*))^p \cdot \mathbf{I} & \mathbf{0} \\ \mathbf{0} & (\mathbb{E} \operatorname{tr}(\mathbf{Z}^*\mathbf{Z}))^p \cdot \mathbf{I} \end{bmatrix} \quad \text{for } p = 0, 1, 2, \dots$$

This observation allows us to obtain sharp bounds for the trace moments of rectangular Gaussian matrices. In this fashion, we can even show that the limiting spectral density of a sequence of rectangular Gaussian matrices is the Marcenko–Pastur distribution, provided that the aspect ratio of the sequence is held constant.

Finally, we remark that similar arguments can be applied to obtain algebraic relations for the Stieltjes transform of the matrix \mathbf{X} . This approach may lead more directly to limit laws for sequences of random matrices with increasing dimension. See [AGZ10, Sec. 2.4.1] or [Kem13, Sec. 9] for an argument of this species.

9.2. Preliminaries. Aside from the results we have collected so far, the proof of Theorem 9.1 requires a few additional ingredients. First, we state some of the basic and well-known properties of the Catalan numbers.

Fact 9.2 (Catalan Numbers). *The p th Catalan number is defined by the formula*

$$\operatorname{Cat}_p := \frac{1}{p+1} \binom{2p}{p} \quad \text{for } p = 0, 1, 2, \dots \tag{9.3}$$

In particular, $p \mapsto \operatorname{Cat}_p$ is nondecreasing, and $\operatorname{Cat}_p \leq 4^p$ for each p . The Catalan numbers satisfy the recursion

$$\operatorname{Cat}_0 = 1 \quad \text{and} \quad \operatorname{Cat}_{p+1} = \sum_{q=0}^p \operatorname{Cat}_q \operatorname{Cat}_{p-q}. \tag{9.4}$$

The next result is a covariance identity for a product of centered functions of a Gaussian vector [NP12, Thm. 2.9.1]. It can be regarded as a refinement of the Poincaré inequality, which provides a bound for the variance of a centered function of a Gaussian vector.

Fact 9.3 (Gaussian Covariance Identity). *Let $\boldsymbol{\gamma}, \boldsymbol{\gamma}' \in \mathbb{R}^n$ be independent standard normal vectors. Let $f, g : \mathbb{R}^n \rightarrow \mathbb{C}$ be functions whose derivatives are square integrable with respect to the standard normal measure, and assume that $\mathbb{E} f(\boldsymbol{\gamma}) = \mathbb{E} g(\boldsymbol{\gamma}) = 0$. Then*

$$\mathbb{E} [f(\boldsymbol{\gamma}) \cdot g(\boldsymbol{\gamma})] = \int_0^\infty \frac{dt}{e^t} \sum_{j=1}^n \mathbb{E} [(\partial_j f)(\boldsymbol{\gamma}) \cdot (\partial_j g)(\boldsymbol{\gamma}_t)] \quad \text{where } \boldsymbol{\gamma}_t := e^{-t} \boldsymbol{\gamma} + \sqrt{1 - e^{-2t}} \boldsymbol{\gamma}'.$$

The symbol ∂_j refers to differentiation with respect to the j th coordinate.

The usual statement of this result involves the Ornstein–Uhlenbeck semigroup, but we have given a more elementary formulation.

Finally, we need a bound for the solution to a certain type of polynomial inequality. This estimate is related to Fujiwara’s inequality [Mar66, Sec. 27]. We include a proof sketch since we could not locate the precise statement in the literature.

Proposition 9.4 (Polynomial Inequalities). *Consider an integer $k \geq 3$, and fix positive numbers α and β . Then*

$$u^k \leq \alpha + \beta u^{k-2} \quad \text{implies} \quad u \leq \alpha^{1/k} + \beta^{1/2}$$

Proof Sketch. Consider the polynomial $\varphi : u \mapsto u^k - \beta u^{k-2} - \alpha$. The Descartes Rule of Signs implies that φ has exactly one positive root, say u_+ . Furthermore, $\varphi(u) \leq 0$ for a positive number u if and only if $u \leq u_+$ because $\varphi(0) < 0$. By direct calculation, one may verify that $u_\star = \alpha^{1/k} + \beta^{1/2}$ satisfies $\varphi(u_\star) > 0$, which means that $u_+ < u_\star$. We conclude that $\varphi(u) \leq 0$ implies $u \leq u_+ < u_\star$. \square

9.3. The Normalized Trace Moments. Let us commence the proof of Theorem 9.1. First, we introduce notation for the normalized trace moments of the matrix Gaussian series:

$$\mu_p := \mathbb{E} \bar{\text{tr}} \mathbf{X}^p \quad \text{for } p = 0, 1, 2, \dots$$

It is clear that $\mu_0 = 1$. Since \mathbf{X} is a symmetric random variable, the odd trace moments are zero:

$$\mu_{2p+1} = 0 \quad \text{for } p = 0, 1, 2, \dots$$

It remains to calculate the even trace moments.

We can obtain the second moment from a simple argument:

$$\sum_{i=1}^n \mathbf{H}_i^2 = \mathbb{E} \mathbf{X}^2 = (\mathbb{E} \bar{\text{tr}} \mathbf{X}^2) \cdot \mathbf{I} = \mu_2 \cdot \mathbf{I}. \quad (9.5)$$

The first identity follows from a direct calculation using the definition (2.1) of the matrix Gaussian series. The second identity is the strong isotropy hypothesis (9.1), and the last relation is the definition of μ_2 . Take the spectral norm of (9.5) to see that

$$\mu_2 = \sigma^2. \quad (9.6)$$

We have identified the standard deviation parameter, defined in (9.2).

9.4. Representation of Higher-Order Moments. The major challenge is to compute the rest of the even moments. As usual, the first step is to invoke Gaussian integration by parts. For each integer $p \geq 1$, Lemma 6.1 implies that

$$\mu_{2(p+1)} = \sum_{q=0}^{2p} \sum_{i=1}^n \mathbb{E} \bar{\text{tr}} [\mathbf{H}_i \mathbf{X}^q \mathbf{H}_i \mathbf{X}^{2p-q}].$$

We are considering $\mu_{2(p+1)}$ instead of μ_{2p} because it makes the argument cleaner. To analyze this expression, we will examine each index q separately and subject each one to the same treatment.

Fix an index $0 \leq q \leq 2p$. First, we center both \mathbf{X}^q and \mathbf{X}^{2p-q} by adding and subtracting their expectations:

$$\begin{aligned} \sum_{i=1}^n \mathbb{E} \bar{\text{tr}} [\mathbf{H}_i \mathbf{X}^q \mathbf{H}_i \mathbf{X}^{2p-q}] &= \sum_{i=1}^n \bar{\text{tr}} [\mathbf{H}_i (\mathbb{E} \mathbf{X}^q) \mathbf{H}_i (\mathbb{E} \mathbf{X}^{2p-q})] \\ &\quad + \sum_{i=1}^n \mathbb{E} \bar{\text{tr}} [\mathbf{H}_i (\mathbf{X}^q - \mathbb{E} \mathbf{X}^q) \mathbf{H}_i (\mathbf{X}^{2p-q} - \mathbb{E} \mathbf{X}^{2p-q})]. \end{aligned}$$

The cross-terms vanish because each one has zero mean. It is productive to think of the first sum on the right-hand side as an approximation to the left-hand side, while the second sum is a perturbation.

Let us focus on the first sum on the right-hand side of the last display. We can use the strong isotropy hypothesis (9.1) to simplify this expression:

$$\begin{aligned} \sum_{i=1}^n \bar{\text{tr}} [\mathbf{H}_i (\mathbb{E} \mathbf{X}^q) \mathbf{H}_i (\mathbb{E} \mathbf{X}^{2p-q})] &= \sum_{i=1}^n \bar{\text{tr}} [\mathbf{H}_i ((\mathbb{E} \bar{\text{tr}} \mathbf{X}^q) \cdot \mathbf{I}) \mathbf{H}_i ((\mathbb{E} \bar{\text{tr}} \mathbf{X}^{2p-q}) \cdot \mathbf{I})] \\ &= (\bar{\text{tr}} \sum_{i=1}^n \mathbf{H}_i^2) (\mathbb{E} \bar{\text{tr}} \mathbf{X}^q) (\mathbb{E} \bar{\text{tr}} \mathbf{X}^{2p-q}) \\ &= \sigma^2 \cdot \mu_q \mu_{2p-q}. \end{aligned}$$

The last identity follows from (9.5) and (9.6). As a side note, our motivation here is to imitate the moment identity that would hold if \mathbf{X} and \mathbf{H}_i were free from each other, in the sense of free probability.

Finally, we combine the last three displays to reach

$$\mu_{2(p+1)} = \sigma^2 \cdot \sum_{q=0}^p \mu_{2q} \mu_{2(p-q)} + \sum_{q=1}^{2p-1} \sum_{i=1}^n \mathbb{E} \bar{\text{tr}} [\mathbf{H}_i (\mathbf{X}^q - \mathbb{E} \mathbf{X}^q) \mathbf{H}_i (\mathbf{X}^{2p-q} - \mathbb{E} \mathbf{X}^{2p-q})]. \quad (9.7)$$

Observe that we have modified the indexing of both sums. This step depends on the facts that $\mu_q = 0$ for odd q and that $\mathbf{X}^0 = \mathbf{I}$.

9.5. The Perturbation Term. The next step in the argument is to bound the perturbation term in (9.7) in terms of the alignment parameter w , defined in (9.2). We will use the Gaussian covariance identity, Fact 9.3.

To that end, let us explain how to write each summand in the perturbation term as a covariance. Let \mathbf{H} be a real, diagonal matrix: $\mathbf{H} = \text{diag}(h_1, \dots, h_d)$. Expanding the normalized trace, using α, β for coordinate indices, we find that

$$\mathbb{E} \bar{\text{tr}} [\mathbf{H} (\mathbf{X}^q - \mathbb{E} \mathbf{X}^q) \mathbf{H} (\mathbf{X}^{2p-q} - \mathbb{E} \mathbf{X}^{2p-q})] = \frac{1}{d} \sum_{\alpha, \beta=1}^d h_\alpha h_\beta \cdot \mathbb{E} [(\mathbf{X}^q - \mathbb{E} \mathbf{X}^q)_{\alpha\beta} (\mathbf{X}^{2p-q} - \mathbb{E} \mathbf{X}^{2p-q})_{\beta\alpha}].$$

To apply the Gaussian covariance identity to the expectation, we introduce a parameterized family $\{\mathbf{X}_t : t \geq 0\}$ of random matrices where

$$\mathbf{X}_t := \sum_{k=1}^n (e^{-t} \gamma_k + \sqrt{1 - e^{-2t}} \gamma'_k) \cdot \mathbf{H}_k \quad \text{where } \boldsymbol{\gamma}' \text{ is an independent copy of } \boldsymbol{\gamma}.$$

Observe that \mathbf{X} and \mathbf{X}_t have the same distribution, although they are dependent. Fact 9.3 and Fact 6.3 deliver

$$\begin{aligned} & \mathbb{E}[(\mathbf{X}^q - \mathbb{E}\mathbf{X}^q)_{\alpha\beta}(\mathbf{X}^{2p-q} - \mathbb{E}\mathbf{X}^{2p-q})_{\beta\alpha}] \\ &= \int_0^\infty \frac{dt}{e^t} \sum_{j=1}^n \mathbb{E} \left[\left(\sum_{r=0}^{q-1} \mathbf{X}^r \mathbf{H}_j \mathbf{X}^{q-1-r} \right)_{\alpha\beta} \left(\sum_{s=0}^{2p-q-1} \mathbf{X}_t^s \mathbf{H}_j \mathbf{X}_t^{2p-q-1-s} \right)_{\beta\alpha} \right] \\ &= \sum_{r=0}^{q-1} \sum_{s=0}^{2p-q-1} \int_0^\infty \frac{dt}{e^t} \sum_{j=1}^n \mathbb{E} [(\mathbf{X}^r \mathbf{H}_j \mathbf{X}^{q-1-r})_{\alpha\beta} (\mathbf{X}_t^s \mathbf{H}_j \mathbf{X}_t^{2p-q-1-s})_{\beta\alpha}]. \end{aligned}$$

Combining these formulas and expressing the result in terms of the normalized trace again, we find that

$$\begin{aligned} & \mathbb{E} \bar{\text{tr}} [\mathbf{H}(\mathbf{X}^q - \mathbb{E}\mathbf{X}^q) \mathbf{H}(\mathbf{X}^{2p-q} - \mathbb{E}\mathbf{X}^{2p-q})] \\ &= \sum_{r=0}^{q-1} \sum_{s=0}^{2p-q-1} \int_0^\infty \frac{dt}{e^t} \sum_{j=1}^n \mathbb{E} \bar{\text{tr}} [\mathbf{H}(\mathbf{X}^r \mathbf{H}_j \mathbf{X}^{q-1-r}) \mathbf{H}(\mathbf{X}_t^s \mathbf{H}_j \mathbf{X}_t^{2p-q-1-s})]. \end{aligned}$$

In fact, this expression is valid for any Hermitian matrix \mathbf{H} because of the unitary invariance of the trace. Summing the last identity over $\mathbf{H} = \mathbf{H}_i$, we reach

$$\begin{aligned} & \sum_{i=1}^n \mathbb{E} \bar{\text{tr}} [\mathbf{H}_i(\mathbf{X}^q - \mathbb{E}\mathbf{X}^q) \mathbf{H}_i(\mathbf{X}^{2p-q} - \mathbb{E}\mathbf{X}^{2p-q})] \\ &= \sum_{r=0}^{q-1} \sum_{s=0}^{2p-q-1} \int_0^\infty \frac{dt}{e^t} \sum_{i,j=1}^n \mathbb{E} \bar{\text{tr}} [\mathbf{H}_i \mathbf{X}^r \mathbf{H}_j \mathbf{X}^{q-1-r} \mathbf{H}_i \mathbf{X}_t^s \mathbf{H}_j \mathbf{X}_t^{2p-q-1-s}]. \quad (9.8) \end{aligned}$$

At this point, the alignment parameter w starts to become visible.

9.6. Finding the Matrix Alignment Parameter. Our next goal is to control the expression (9.8) in terms of the alignment parameter w . To do so, we use the interpolation result, Proposition 8.3, to bound the sum over (i, j) . For each choice of indices (q, r, s) , we obtain the estimate

$$\begin{aligned} \left| \sum_{i,j=1}^n \mathbb{E} \bar{\text{tr}} [\mathbf{H}_i \mathbf{X}^r \mathbf{H}_j \mathbf{X}^{q-1-r} \mathbf{H}_i \mathbf{X}_t^s \mathbf{H}_j \mathbf{X}_t^{2p-q-1-s}] \right| &\leq \max \left\{ \mathbb{E} \max_{\mathbf{Q}_\ell} \left| \sum_{i,j=1}^n \bar{\text{tr}} [\mathbf{H}_i \mathbf{Q}_1 \mathbf{X}^{2p-2} \mathbf{H}_j \mathbf{Q}_2 \mathbf{H}_i \mathbf{Q}_3 \mathbf{H}_j \mathbf{Q}_4] \right|, \right. \\ &\quad \mathbb{E} \max_{\mathbf{Q}_\ell} \left| \sum_{i,j=1}^n \bar{\text{tr}} [\mathbf{H}_i \mathbf{Q}_1 \mathbf{H}_j \mathbf{Q}_2 \mathbf{X}^{2p-2} \mathbf{H}_i \mathbf{Q}_3 \mathbf{H}_j \mathbf{Q}_4] \right|, \\ &\quad \mathbb{E} \max_{\mathbf{Q}_\ell} \left| \sum_{i,j=1}^n \bar{\text{tr}} [\mathbf{H}_i \mathbf{Q}_1 \mathbf{H}_j \mathbf{Q}_2 \mathbf{H}_i \mathbf{Q}_3 \mathbf{X}_t^{2p-2} \mathbf{H}_j \mathbf{Q}_4] \right|, \\ &\quad \left. \mathbb{E} \max_{\mathbf{Q}_\ell} \left| \sum_{i,j=1}^n \bar{\text{tr}} [\mathbf{H}_i \mathbf{Q}_1 \mathbf{H}_j \mathbf{Q}_2 \mathbf{H}_i \mathbf{Q}_3 \mathbf{H}_j \mathbf{Q}_4 \mathbf{X}_t^{2p-2}] \right| \right\}. \end{aligned}$$

Each random unitary matrix \mathbf{Q}_ℓ commutes with the corresponding random matrix \mathbf{X} or \mathbf{X}_t .

As in Section 8.7, we can bound each term in the maximum in the same fashion. For example, consider the fourth term:

$$\begin{aligned} \mathbb{E} \max_{\mathbf{Q}_\ell} \left| \sum_{i,j=1}^n \bar{\text{tr}} [\mathbf{H}_i \mathbf{Q}_1 \mathbf{H}_j \mathbf{Q}_2 \mathbf{H}_i \mathbf{Q}_3 \mathbf{H}_j \mathbf{Q}_4 \mathbf{X}_t^{2p-2}] \right| &\leq \mathbb{E} \max_{\mathbf{Q}_\ell} \left[\left\| \sum_{i,j=1}^n \mathbf{H}_i \mathbf{Q}_1 \mathbf{H}_j \mathbf{Q}_2 \mathbf{H}_i \mathbf{Q}_3 \mathbf{H}_j \right\| \cdot (\bar{\text{tr}} \mathbf{X}_t^{2p-2}) \right] \\ &\leq \max_{\mathbf{Q}_\ell} \left\| \sum_{i,j=1}^n \mathbf{H}_i \mathbf{Q}_1 \mathbf{H}_j \mathbf{Q}_2 \mathbf{H}_i \mathbf{Q}_3 \mathbf{H}_j \right\| \cdot (\mathbb{E} \bar{\text{tr}} \mathbf{X}_t^{2p-2}) \\ &= w^4 \cdot \mu_{2(p-1)}. \end{aligned}$$

The first step is Hölder's inequality for the trace, and the second step is Hölder's inequality for the expectation. In the last line, we recall that \mathbf{X}_t has the same distribution as \mathbf{X} to identify $\mu_{2(p-1)}$. Finally, we recognize the matrix alignment parameter w , defined in (9.2).

In summary, we have shown that

$$\left| \sum_{i,j=1}^n \mathbb{E} \bar{\text{tr}} [\mathbf{H}_i \mathbf{X}^r \mathbf{H}_j \mathbf{X}^{q-1-r} \mathbf{H}_i \mathbf{X}_t^s \mathbf{H}_j \mathbf{X}_t^{2p-q-1-s}] \right| \leq w^4 \cdot \mu_{2(p-1)}.$$

Introduce this bound into (9.8) to arrive at

$$\left| \sum_{i=1}^n \mathbb{E} \bar{\text{tr}} [\mathbf{H}_i(\mathbf{X}^q - \mathbb{E}\mathbf{X}^q) \mathbf{H}_i(\mathbf{X}^{2p-q} - \mathbb{E}\mathbf{X}^{2p-q})] \right| \leq q(2p-q) \cdot w^4 \cdot \mu_{2(p-1)} \leq p^2 \cdot w^4 \cdot \mu_{2(p-1)}.$$

We have used the numerical inequality $u(a-u) \leq a^2/4$, valid for $u \in \mathbb{R}$. Finally, we sum this expression over the index q to conclude that

$$\left| \sum_{q=0}^{2p-1} \sum_{i=1}^n \mathbb{E} \bar{\text{tr}} [\mathbf{H}_i(\mathbf{X}^q - \mathbb{E}\mathbf{X}^q) \mathbf{H}_i(\mathbf{X}^{2p-q} - \mathbb{E}\mathbf{X}^{2p-q})] \right| \leq 2p^3 \cdot w^4 \cdot \mu_{2(p-1)}. \quad (9.9)$$

This is the required bound for the perturbation term in (9.7).

9.7. A Recursion for the Trace Moments. In view of (9.7) and (9.9), we have shown that

$$\mu_{2(p+1)} = \sigma^2 \cdot \sum_{q=0}^p \mu_{2q} \mu_{2(p-q)} \pm 2p^3 \cdot w^4 \cdot \mu_{2(p-1)} \quad \text{for } p = 1, 2, 3, \dots \quad (9.10)$$

We have written \pm to indicate that the expression contains both a lower bound and an upper bound for the normalized trace moment $\mu_{2(p+1)}$.

In the next two sections, we will solve this recursion to obtain explicit bounds on the trace moments. First, we obtain the upper bound

$$\mu_{2p}^{1/(2p)} \leq \text{Cat}_p^{1/(2p)} \cdot \sigma + 2^{1/4} p^{5/4} \cdot w \quad \text{for } p = 1, 2, 3, \dots \quad (9.11)$$

This result gives us a Khintchine-type inequality. Afterward, assuming $p^{7/4} w \leq 0.7\sigma$, we establish the lower bound

$$\mu_{2p}^{1/(2p)} \geq \text{Cat}_p^{1/(2p)} \cdot \sigma \cdot [1 - (pw/\sigma)^4]_+^{1/(2p)} \quad \text{for } p = 1, 2, 3, \dots \quad (9.12)$$

Together these estimates yield the statement of Theorem 9.1.

9.8. Solving the Recursion: Upper Bound. We begin with the proof of the upper bound (9.11). The first step in the argument is to remove the lag term $\mu_{2(p-1)}$ from the recursion (9.10) using moment comparison. Fix an integer $p \geq 1$. Observe that

$$\mu_{2(p-1)} = \mathbb{E} \bar{\text{tr}} X^{2(p-1)} \leq \mathbb{E} (\bar{\text{tr}} X^{2(p+1)})^{(p-1)/(p+1)} \leq (\mathbb{E} \bar{\text{tr}} X^{2(p+1)})^{(p-1)/(p+1)} = \mu_{2(p+1)}^{(p-1)/(p+1)}.$$

The first inequality holds because $q \mapsto (\bar{\text{tr}} A^q)^{1/q}$ is increasing for any positive-semidefinite matrix A , while the second inequality is Lyapunov's. Introduce this estimate into the recursion (9.10) to obtain

$$\mu_{2(p+1)} \leq \sigma^2 \cdot \sum_{q=0}^p \mu_{2q} \mu_{2(p-q)} + 2p^3 w^4 \cdot \mu_{2(p+1)}^{(p-1)/(p+1)}.$$

This is a polynomial inequality of the form $u^{p+1} \leq \alpha + \beta u^{p-1}$, so Proposition 9.4 ensures that $u \leq \alpha^{1/(p+1)} + \beta^{1/2}$. In other words,

$$\mu_{2(p+1)}^{1/(p+1)} \leq \left(\sigma^2 \cdot \sum_{q=0}^p \mu_{2q} \mu_{2(p-q)} \right)^{1/(p+1)} + \sqrt{2} p^{3/2} \cdot w^2 \quad \text{for } p = 1, 2, 3, \dots \quad (9.13)$$

Using this formula, we will apply induction to prove that

$$\mu_{2p}^{1/p} \leq \text{Cat}_p^{1/p} \cdot \sigma^2 + \sqrt{2} p^{5/2} \cdot w^2 \quad \text{when } p = 1, 2, 3, \dots \quad (9.14)$$

The stated result (9.11) follows from (9.14) once we take the square root and invoke subadditivity.

Let us commence the induction. The formula (9.14) holds for $p = 1$ because $\mu_2 = \sigma^2$, as noted in (9.6). Assuming that the bound (9.14) holds for each integer p in the range $1, 2, 3, \dots, r$, we will verify that the same bound is also valid for $p = r + 1$. For any integer q in the range $1 \leq q \leq r$, the bound (9.14) implies that

$$\mu_{2q} \leq \text{Cat}_q \cdot \sigma^{2q} \cdot \left(1 + \frac{\sqrt{2} q^{5/2} w^2}{\text{Cat}_q^{1/q} \sigma^2} \right)^q \leq \text{Cat}_q \cdot \sigma^{2q} \cdot \exp \left(\frac{\sqrt{2} q^{7/2} w^2}{\sigma^2} \right). \quad (9.15)$$

Using the definition (9.3) of the Catalan numbers, one may verify that $q \mapsto q^{5/2} \text{Cat}_q^{-1/q}$ is increasing, so

$$\mu_{2q} \leq \text{Cat}_q \cdot \sigma^{2q} \cdot \left(1 + \frac{\sqrt{2} (r+1)^{5/2} w^2}{\text{Cat}_{r+1}^{1/(r+1)} \sigma^2} \right)^q \quad \text{for } q = 0, 1, 2, \dots, r.$$

The case $q = 0$ follows by inspection. Now, the latter bound and the recursion (9.4) for Catalan numbers together imply that

$$\sigma^2 \cdot \sum_{q=0}^r \mu_{2q} \mu_{2(r-q)} \leq \text{Cat}_{r+1} \cdot \sigma^{2(r+1)} \cdot \left(1 + \frac{\sqrt{2} (r+1)^{5/2} w^2}{\text{Cat}_{r+1}^{1/(r+1)} \sigma^2} \right)^r.$$

Take the $r + 1$ root to determine that

$$\begin{aligned} \left(\sigma^2 \cdot \sum_{q=0}^r \mu_{2q} \mu_{2(r-q)} \right)^{1/(r+1)} &\leq \text{Cat}_{r+1}^{1/(r+1)} \cdot \sigma^2 \cdot \left(1 + \frac{r}{r+1} \cdot \frac{\sqrt{2} (r+1)^{5/2} w^2}{\text{Cat}_{r+1}^{1/(r+1)} \sigma^2} \right), \\ &= \text{Cat}_{r+1}^{1/(r+1)} \cdot \sigma^2 + \sqrt{2} (r+1)^{5/2} w^2 - \sqrt{2} (r+1)^{3/2} w^2. \end{aligned}$$

We have used the numerical inequality $(1+x)^\alpha \leq 1 + \alpha x$, valid for $0 \leq \alpha \leq 1$ and $x \geq 0$. Combine this estimate with the recursive bound (9.13) for $p = r$ to obtain

$$\mu_{2(r+1)}^{1/(r+1)} \leq \text{Cat}_{r+1}^{1/(r+1)} \cdot \sigma^2 + \sqrt{2} (r+1)^{5/2} w^2.$$

We see that (9.14) holds for $p = r + 1$, and the induction may proceed.

9.9. Solving the Recursion: Lower Bound. We turn to the proof of the lower bound (9.12). Assuming that $p^{7/4}w \leq 0.7\sigma$, we will use induction to show that

$$\mu_{2p} \geq \sigma^{2p} \cdot \text{Cat}_p \cdot [1 - (pw/\sigma)^4]_+ \quad \text{for } p = 0, 1, 2, 3, \dots \quad (9.16)$$

The result (9.12) follows once we take the $(2p)$ th root.

To begin the induction, recall that $\mu_0 = 1$, so the formula (9.16) is valid for $p = 0$. Suppose now that (9.16) is valid for each integer p in the range $0, 1, 2, \dots, r$. We will verify the formula for $p = r + 1$. The lower branch of the recursion (9.10) states that

$$\mu_{2(r+1)} \geq \sigma^2 \cdot \sum_{q=0}^r \mu_{2q} \mu_{2(r-q)} - 2r^3 w^4 \mu_{2(r-1)}.$$

The induction hypothesis (9.16) yields

$$\begin{aligned} \sigma^2 \cdot \sum_{q=0}^r \mu_{2q} \mu_{2(r-q)} &\geq \sigma^{2(p+1)} \cdot \sum_{q=0}^r \text{Cat}_q \text{Cat}_{r-q} \cdot [1 - (qw/\sigma)^4 - ((r-q)w/\sigma)^4] \\ &\geq \text{Cat}_{r+1} \cdot \sigma^{2(p+1)} \cdot (rw/\sigma)^4. \end{aligned}$$

We have used the fact that $q \mapsto q^4 + (r-q)^4$ achieves its maximum value on $[0, r]$ at one of the endpoints because of convexity. We also applied the recursive formula (9.4) for the Catalan numbers. The bound (9.15) implies that

$$\mu_{2(r-1)} \leq 2 \text{Cat}_{r-1} \sigma^{2(r-1)} \quad \text{when} \quad \frac{r^{7/4}w}{\sigma} \leq \sqrt{\frac{\log 2}{\sqrt{2}}} \approx 0.7.$$

Therefore,

$$2r^3 w^4 \mu_{2(r-1)} \leq 4 \text{Cat}_{r-1} \cdot \sigma^{2(r-1)} \cdot r^3 w^4 \leq 4 \text{Cat}_{r+1} \cdot \sigma^{2(r+1)} \cdot r^3 (w/\sigma)^4.$$

The second inequality holds because the Catalan numbers are nondecreasing. Combine the last three displays to arrive at

$$\mu_{2(r+1)} \geq \text{Cat}_{r+1} \cdot \sigma^{2(r+1)} \cdot [1 - (r^4 + 4r^3)(w/\sigma)^4] \geq \text{Cat}_{r+1} \cdot \sigma^{2(r+1)} \cdot [1 - (r+1)^4 (w/\sigma)^4].$$

We have verified the formula (9.16) for $p = r + 1$, which completes the proof.

APPENDIX A. INTERPOLATION RESULTS

In this appendix, we establish Proposition 8.3, the interpolation inequality for a multilinear function of a random matrix, whose proof appears below in Appendix A.4.

A.1. Multivariate Complex Interpolation. The interpolation result we use in the body of the paper is a consequence of a more general theorem on interpolation for a function of several complex variables.

Proposition A.1 (Multivariate Complex Interpolation). *Let k be a natural number. For a positive number α , define the simplicial prism*

$$\Delta_k(\alpha) := \left\{ (c_1, \dots, c_k) \in \mathbb{C}^k : \text{Re } c_i \geq 0 \text{ for each } i \text{ and } \sum_{i=1}^k \text{Re } c_i \leq \alpha \right\}.$$

Consider a bounded, continuous function $G : \Delta_k(\alpha) \rightarrow \mathbb{C}$. For each pair (i, j) of distinct indices and each $\mathbf{c} \in \Delta_k(\alpha)$, assume that G has the analytic section property:

$$z \mapsto G(c_1, \dots, c_i + z, \dots, c_j - z, \dots, c_k) \quad \text{is analytic on } -\text{Re } c_i < \text{Re } z < \text{Re } c_j. \quad (\text{A.1})$$

Then, for each $\mathbf{z} \in \Delta_k(\alpha)$ with $\beta := \sum_{i=1}^k \text{Re } z_i > 0$,

$$|G(z_1, \dots, z_k)| \leq \left(\prod_{i=1}^k \sup_{t_\ell \in \mathbb{R}} |G(it_1, \dots, it_{i-1}, \beta + it_i, it_{i+1}, \dots, it_k)|^{\text{Re } z_i} \right)^{1/\beta}. \quad (\text{A.2})$$

We establish Proposition A.1 in the next two sections. The argument relies on the same principles that support standard univariate complex interpolation. Although it seems likely that a result of this form already appears in the literature, we were not able to locate a reference.

A.2. Preliminaries. Proposition A.1 depends on Hadamard's Three-Lines Theorem [Gar07, Prop. 9.1.1].

Proposition A.2 (Three-Lines Theorem). *Consider the vertical strip $\Delta_1(1)$ in the complex plane:*

$$\Delta_1(1) := \{z \in \mathbb{C} : 0 \leq \operatorname{Re} z \leq 1\}.$$

Consider a bounded, continuous function $f : \Delta_1(1) \rightarrow \mathbb{C}$ that is analytic in the interior of $\Delta_1(1)$. For each $\theta \in [0, 1]$,

$$\sup_{t \in \mathbb{R}} |f(\theta + it)| \leq \sup_{t \in \mathbb{R}} |f(1 + it)|^\theta \cdot \sup_{t \in \mathbb{R}} |f(it)|^{1-\theta}.$$

As we will see, this result delivers the $k = 2$ case of Proposition A.1.

A.3. Proof of Proposition A.1. The proof of the multivariate interpolation result, Proposition A.1, follows by induction on the number k of arguments.

Let us begin with the base cases. When the function G has one argument only, the inequality (A.2) is obviously true. Next, consider a bivariate function $G_2 : \Delta_2(\alpha) \rightarrow \mathbb{C}$ that is bounded and continuous and has the analytic section property (A.1). Fix a point $\mathbf{z} \in \Delta_2(\alpha)$ with $\beta := \operatorname{Re} z_1 + \operatorname{Re} z_2 > 0$. Define the bounded, continuous function

$$f(y) := G_2(\beta y + i \operatorname{Im} z_1, \beta(1 - y) + i \operatorname{Im} z_2) \quad \text{for } y \in \Delta_1(1).$$

The assumption (A.1) implies that f is analytic on $0 < \operatorname{Re} y < 1$. Select $\theta = \operatorname{Re} z_1 / \beta$, which gives $1 - \theta = \operatorname{Re} z_2 / \beta$. An application of the Three-Lines Theorem, Proposition A.2, implies that

$$|G_2(z_1, z_2)| = |f(\theta)| \leq \sup_{t \in \mathbb{R}} |f(1 + it)|^\theta \cdot \sup_{t \in \mathbb{R}} |f(it)|^{1-\theta}.$$

Introducing the definition of f and simplifying,

$$\begin{aligned} |G_2(z_1, z_2)| &\leq \sup_{t \in \mathbb{R}} |G_2(\beta(1 + it) + i \operatorname{Im} z_1, -\beta it + i \operatorname{Im} z_2)|^{\operatorname{Re} z_1 / \beta} \\ &\quad \times \sup_{t \in \mathbb{R}} |G_2(\beta it + i \operatorname{Im} z_1, \beta(1 - it) + i \operatorname{Im} z_2)|^{\operatorname{Re} z_2 / \beta} \\ &\leq \left(\sup_{s_1, s_2 \in \mathbb{R}} |G_2(\beta + is_1, is_2)|^{\operatorname{Re} z_1} \cdot \sup_{s_1, s_2 \in \mathbb{R}} |G_2(is_1, \beta + is_2)|^{\operatorname{Re} z_2} \right)^{1/\beta}. \end{aligned} \quad (\text{A.3})$$

This is the $k = 2$ case of Proposition A.1.

Fix a positive integer k , and suppose that we have established the inequality (A.2) for functions with $k - 1$ arguments. In other words, assume that $G_{k-1} : \Delta_{k-1}(\alpha') \rightarrow \mathbb{C}$ is bounded and continuous, and it has the analytic section property (A.1). Then, for $\mathbf{z} \in \Delta_{k-1}(\alpha')$ with $\beta' := \sum_{i=1}^{k-1} \operatorname{Re} z_i > 0$,

$$|G_{k-1}(z_1, \dots, z_{k-1})| \leq \left(\prod_{i=1}^{k-1} \sup_{t_i \in \mathbb{R}} |G_{k-1}(it_1, \dots, \beta' + it_i, \dots, it_{k-1})|^{\operatorname{Re} z_i} \right)^{1/\beta'}. \quad (\text{A.4})$$

We need to extend this result to functions with k variables.

Consider a bounded, continuous function $G_k : \Delta_k(\alpha) \rightarrow \mathbb{C}$ with the analytic section property (A.1). Fix a complex vector $\mathbf{z} \in \Delta_k(\alpha)$ with $\beta := \sum_{i=1}^k \operatorname{Re} z_i > 0$. Define the number $\beta' := \sum_{i=1}^{k-1} \operatorname{Re} z_i$. When $\beta' = 0$, the formula (A.2) is trivial for $G = G_k$ because $\operatorname{Re} z_1 = \dots = \operatorname{Re} z_{k-1} = 0$. Therefore, we may assume that $\beta' > 0$. Introduce the function

$$G_{k-1}(y_1, \dots, y_{k-1}) := G_k(y_1, \dots, y_{k-1}, z_k) \quad \text{for } \mathbf{y} \in \Delta_{k-1}(\beta').$$

One may verify that G_{k-1} inherits boundedness, continuity, and analytic sections from G_k . Therefore, the induction hypothesis (A.4) gives

$$|G_k(z_1, \dots, z_{k-1}, z_k)| \leq \left(\prod_{i=1}^{k-1} \sup_{t_i \in \mathbb{R}} |G_k(it_1, \dots, \beta' + it_i, \dots, it_{k-1}, z_k)|^{\operatorname{Re} z_i} \right)^{1/\beta'}. \quad (\text{A.5})$$

For each fixed choice of the index i and of the numbers $t_1, \dots, t_{k-1} \in \mathbb{R}$, consider the function

$$G_2(y_i, y_k) := G_k(it_1, \dots, y_i, \dots, it_{k-1}, y_k) \quad \text{for } (y_i, y_k) \in \Delta_2(\alpha).$$

Since $\beta' + \operatorname{Re} z_k = \beta$, the bivariate case (A.3) provides that

$$\begin{aligned} |G_k(it_1, \dots, \beta + it_i, \dots, it_{k-1}, z_k)| &\leq \sup_{s_i, s_k \in \mathbb{R}} |G_k(it_1, \dots, \beta + is_i, \dots, it_{k-1}, is_k)|^{\beta' / \beta} \\ &\quad \times \sup_{s_i, s_k \in \mathbb{R}} |G_k(it_1, \dots, is_i, \dots, it_{k-1}, \beta + is_k)|^{1 - \beta' / \beta}. \end{aligned} \quad (\text{A.6})$$

Combine the bounds (A.5) and (A.6) to reach

$$\begin{aligned} |G_k(z_1, \dots, z_{k-1}, z_k)| &\leq \left(\prod_{i=1}^{k-1} \sup_{t_\ell \in \mathbb{R}} |G_k(it_1, \dots, \beta + it_i, \dots, it_k)|^{\operatorname{Re} z_i} \right)^{1/\beta} \\ &\quad \times \left(\prod_{i=1}^{k-1} \sup_{t_\ell \in \mathbb{R}} |G_k(it_1, \dots, it_i, \dots, \beta + it_k)|^{\operatorname{Re} z_i} \right)^{(\beta - \beta')/(\beta' \beta)} \end{aligned}$$

Since $\beta' = \sum_{i=1}^{k-1} \operatorname{Re} z_i$ and $\beta - \beta' = \operatorname{Re} z_k$, we see that the second product has the same form as the $i = k$ term of the first product. Thus,

$$|G_k(z_1, \dots, z_k)| \leq \left(\prod_{i=1}^k \sup_{t_\ell \in \mathbb{R}} |G_k(it_1, \dots, \beta + it_i, \dots, it_k)|^{\operatorname{Re} z_i} \right)^{1/\beta}.$$

This step completes the induction. We have established Proposition A.1.

A.4. Interpolation for a Multilinear Function of Random Matrices. We are now prepared to establish the interpolation result, Proposition 8.3, for a multilinear function of random matrices. We will actually establish a somewhat more precise version, which we state here.

Proposition A.3 (Refined Multilinear Interpolation). *Suppose that $F : (\mathbb{M}_d)^k \rightarrow \mathbb{C}$ is a multilinear function. Fix nonnegative integers $\alpha_1, \dots, \alpha_k$ with $\sum_{i=1}^k \alpha_i = \alpha$. Let $\mathbf{Y}_i \in \mathbb{H}_d$ be random Hermitian matrices, not necessarily independent, for which $\mathbb{E} \|\mathbf{Y}_i\|^\alpha < \infty$. Then*

$$|\mathbb{E} F(\mathbf{Y}_1^{\alpha_1}, \dots, \mathbf{Y}_k^{\alpha_k})| \leq \left(\prod_{i=1}^k \left[\mathbb{E} \max_{\mathbf{Q}_\ell} |F(\mathbf{Q}_1, \dots, \mathbf{Q}_{i-1}, \mathbf{Q}_i \mathbf{Y}_i^{\alpha_i}, \mathbf{Q}_{i+1}, \dots, \mathbf{Q}_k)| \right]^{\alpha_i} \right)^{1/\alpha}.$$

Each \mathbf{Q}_ℓ is a unitary matrix that commutes with \mathbf{Y}_ℓ .

We establish this result in the next section.

Observe that Proposition A.3 immediately implies Proposition 8.3, the interpolation result that we use in the body of the paper. Indeed, we recognize the large parenthesis on the right-hand side as a geometric mean, and we bound the geometric mean by the maximum of its components.

A.5. Proof of Proposition A.3. By a perturbative argument, we may assume that each matrix \mathbf{Y}_i is almost surely nonsingular. Indeed, for a parameter $\varepsilon > 0$, we can replace each \mathbf{Y}_i by the modified matrix $\tilde{\mathbf{Y}}_i := \mathbf{Y}_i + \varepsilon \gamma_i \mathbf{I}$, where $\{\gamma_i\}$ is an independent family of standard normal variables. After completing the argument, we can draw ε down to zero to obtain the inequality for the original random matrices.

The first step in the argument is to perform a polar factorization of each random Hermitian matrix: $\mathbf{Y}_i = \mathbf{U}_i \mathbf{P}_i$ where \mathbf{U}_i is unitary, \mathbf{P}_i is almost surely positive definite, and the two factors commute with \mathbf{Y}_i for each index i . For clarity of argument, we introduce the unitary matrices $\mathbf{S}_i = \mathbf{U}_i^{\alpha_i}$. With this notation,

$$|\mathbb{E} F(\mathbf{Y}_1^{\alpha_1}, \dots, \mathbf{Y}_k^{\alpha_k})| = |\mathbb{E} F(\mathbf{S}_1 \mathbf{P}_1^{\alpha_1}, \dots, \mathbf{S}_k \mathbf{P}_k^{\alpha_k})| \tag{A.7}$$

We will perform interpolation only on the positive-definite matrices.

Next, we introduce a complex-valued function by replacing the powers α_i with complex variables:

$$G : (z_1, \dots, z_k) \mapsto \mathbb{E} F(\mathbf{S}_1 \mathbf{P}_1^{z_1}, \dots, \mathbf{S}_k \mathbf{P}_k^{z_k}) \quad \text{for } \mathbf{z} \in \Delta_k(\alpha).$$

The set $\Delta_k(\alpha)$ is the simplicial prism defined in the statement of Proposition A.1.

Claim A.4. *The function G is bounded, continuous, and has analytic sections.*

These are the properties required to apply the interpolation result, Proposition A.1.

Let us assume that Claim A.4 holds so that we can complete the proof. The relation (A.7) and Proposition A.1 imply that

$$\begin{aligned} |\mathbb{E} F(\mathbf{Y}_1^{\alpha_1}, \dots, \mathbf{Y}_k^{\alpha_k})| &= |\mathbb{E} F(\mathbf{S}_1 \mathbf{P}_1^{\alpha_1}, \dots, \mathbf{S}_k \mathbf{P}_k^{\alpha_k})| \\ &\leq \left(\prod_{i=1}^k \sup_{t_\ell \in \mathbb{R}} \left| \mathbb{E} F(\mathbf{S}_1 \mathbf{P}_1^{it_1}, \dots, \mathbf{S}_i \mathbf{P}_i^{\alpha + it_i}, \dots, \mathbf{S}_k \mathbf{P}_k^{it_k}) \right|^{\alpha_i} \right)^{1/\alpha}. \end{aligned}$$

Fix an index i in the product. Introduce the unitary matrix $\mathbf{Q}_{ii}(t_i) = \mathbf{S}_i \mathbf{P}_i^{it_i} \mathbf{U}_i^{-\alpha}$, where \mathbf{U}_i is the polar factor of \mathbf{Y}_i . It follows that

$$\mathbf{S}_i \mathbf{P}_i^{\alpha + it_i} = (\mathbf{S}_i \mathbf{P}_i^{it_i} \mathbf{U}_i^{-\alpha}) (\mathbf{U}_i^\alpha \mathbf{P}_i^\alpha) = \mathbf{Q}_{ii}(t_i) \mathbf{Y}_i^\alpha$$

Similarly, for each $\ell \neq i$, we can define $\mathbf{Q}_{i\ell}(t_\ell) = \mathbf{S}_\ell \mathbf{P}_\ell^{it_\ell}$. Therefore,

$$\begin{aligned} |\mathbb{E} F(\mathbf{Y}_1^{\alpha_1}, \dots, \mathbf{Y}_k^{\alpha_k})| &\leq \left(\prod_{i=1}^k \sup_{t_\ell \in \mathbb{R}} |\mathbb{E} F(\mathbf{Q}_{i1}(t_1), \dots, \mathbf{Q}_{ii}(t_i) \mathbf{Y}_i^\alpha, \dots, \mathbf{Q}_{ik}(t_k))|^{\alpha_i} \right)^{1/\alpha} \\ &\leq \left(\prod_{i=1}^k \left[\mathbb{E} \max_{\mathbf{Q}_\ell} |F(\mathbf{Q}_1, \dots, \mathbf{Q}_i \mathbf{Y}_i^\alpha, \dots, \mathbf{Q}_k)| \right]^{\alpha_i} \right)^{1/\alpha}. \end{aligned}$$

By construction $\mathbf{Q}_{i\ell}(t_\ell)$ commutes with \mathbf{Y}_ℓ for each index ℓ . In the second line, we apply Jensen's inequality. Then we relax the supremum to include all unitary matrices \mathbf{Q}_ℓ that commute with the corresponding \mathbf{Y}_ℓ . We can replace the supremum with a maximum since the unitary group is compact and the function F is continuous. This is what we needed to show.

Finally, we must verify Claim A.4. Each multilinear function $F: (\mathbb{M}_d)^k \rightarrow \mathbb{C}$ is bounded and continuous:

$$|F(\mathbf{A}_1, \dots, \mathbf{A}_k)| \leq \text{Const} \cdot \prod_{i=1}^k \|\mathbf{A}_i\|.$$

Fix a point $\mathbf{z} \in \Delta_k(\alpha)$, and let $\beta := \sum_{i=1}^k \text{Re } z_i$. Applying this observation to the function G ,

$$\begin{aligned} |G(z_1, \dots, z_k)| &= |\mathbb{E} F(\mathbf{S}_1 \mathbf{P}_1^{z_1}, \dots, \mathbf{S}_k \mathbf{P}_k^{z_k})| \leq \text{Const} \cdot \mathbb{E} \left[\prod_{i=1}^k \|\mathbf{S}_i \mathbf{P}_i^{z_i}\| \right] \\ &= \text{Const} \cdot \mathbb{E} \left[\prod_{i=1}^k \|\mathbf{Y}_i\|^{\text{Re } z_i} \right] \leq \text{Const} \cdot \mathbb{E} \left[\frac{1}{\beta} \sum_{i=1}^k \|\mathbf{Y}_i\| \right]^\beta. \end{aligned}$$

The first estimate follows from Jensen's inequality and the bound on the multilinear function F . The second inequality depends on the unitary invariance of the spectral norm, the identity $\|\mathbf{P}^z\| = \|\mathbf{P}\|^{\text{Re } z}$, and the polar decomposition $\mathbf{Y}_i = \mathbf{U}_i \mathbf{P}_i$. The last bound is the inequality between the geometric and arithmetic mean. Since $\mathbb{E} \|\mathbf{Y}_i\|^\alpha < \infty$, and $\beta \leq \alpha$, we conclude that G is bounded. Since F is continuous, an application of the Dominated Convergence Theorem shows that G is a continuous function as well.

The proof that G has analytic sections is similar. Fix a vector $\mathbf{c} \in \Delta_k(\alpha)$. Since F is multilinear, it is easy to check that the map

$$z \mapsto F(\mathbf{S}_1 \mathbf{P}_1^{c_1}, \dots, \mathbf{S}_i \mathbf{P}_i^{c_i+z}, \dots, \mathbf{S}_j \mathbf{P}_j^{c_j-z}, \dots, \mathbf{S}_k \mathbf{P}_k^{c_k}) \text{ is analytic on } -\text{Re } c_i < z < \text{Re } c_j$$

for any fixed choice of \mathbf{S}_ℓ and \mathbf{P}_ℓ and any pair (i, j) of distinct indices. Together, the Morera Theorem and the Fubini–Tonelli Theorem allow us to conclude that

$$z \mapsto \mathbb{E} F(\mathbf{S}_1 \mathbf{P}_1^{c_1}, \dots, \mathbf{S}_i \mathbf{P}_i^{c_i+z}, \dots, \mathbf{S}_j \mathbf{P}_j^{c_j-z}, \dots, \mathbf{S}_k \mathbf{P}_k^{c_k})$$

also is analytic. Therefore, the analytic section property (A.1) is in force. Claim A.4 is established.

ACKNOWLEDGMENTS

Afonso Bandeira is responsible for the argument in Section 4.3, and Ramon van Handel has offered critical comments. Parts of this research were completed at Mathematisches Forschungsinstitut Oberwolfach (MFO) and at Instituto Nacional de Matemática Pura e Aplicada (IMPA) in Rio de Janeiro. The author gratefully acknowledges support from ONR award N00014-11-1002, a Sloan Research Fellowship, and the Gordon & Betty Moore Foundation.

REFERENCES

- [AGZ10] G. W. Anderson, A. Guionnet, and O. Zeitouni. *An introduction to random matrices*, volume 118 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 2010.
- [BV14] A. Bandeira and R. van Handel. Sharp nonasymptotic bounds on the norm of random matrices with independent entries. Available at <http://www.arXiv.org/abs/1408.6185>, Aug. 2014.
- [Bha97] R. Bhatia. *Matrix analysis*, volume 169 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1997.
- [BRV13] A. Bondarenko, D. Radchenko, and M. Viazovska. Optimal asymptotic bounds for spherical designs. *Ann. of Math. (2)*, 178(2):443–452, 2013.
- [BS10] Z. Bai and J. W. Silverstein. *Spectral analysis of large dimensional random matrices*. Springer Series in Statistics. Springer, New York, second edition, 2010.
- [Buc01] A. Buchholz. Operator Khintchine inequality in non-commutative probability. *Math. Ann.*, 319:1–16, 2001.
- [CGT12] R. Y. Chen, A. Gittens, and J. A. Tropp. The masked sample covariance estimator: an analysis using matrix concentration inequalities. *Inf. Inference*, 1(1):2–20, 2012.
- [CT14] R. Y. Chen and J. A. Tropp. Subadditivity of matrix φ -entropy and concentration of random matrices. *Electron. J. Probab.*, 19:no. 27, 30, 2014.
- [Gar07] D. J. H. Garling. *Inequalities: a journey into linear analysis*. Cambridge University Press, Cambridge, 2007.

- [Kem13] T. Kemp. Math 247a: Introduction to random matrix theory. Available at <http://www.math.ucsd.edu/~tkemp/247A/247A.Notes.pdf>, Nov. 2013.
- [LO94] R. Latała and K. Oleszkiewicz. On the best constant in the Khinchin-Kahane inequality. *Studia Math.*, 109(1):101–104, 1994.
- [LP86] F. Lust-Piquard. Inégalités de Khintchine dans C_p ($1 < p < \infty$). *C. R. Math. Acad. Sci. Paris*, 303(7):289–292, 1986.
- [LT91] M. Ledoux and M. Talagrand. *Probability in Banach Spaces: Isoperimetry and Processes*. Springer, Berlin, 1991.
- [Mar66] M. Marden. *Geometry of polynomials*. Second edition. Mathematical Surveys, No. 3. American Mathematical Society, Providence, R.I., 1966.
- [MJC⁺14] L. Mackey, M. I. Jordan, R. Y. Chen, B. Farrell, and J. A. Tropp. Matrix concentration inequalities via the method of exchangeable pairs. *Ann. Probab.*, 42(3):906–945, 2014.
- [NP12] I. Nourdin and G. Peccati. *Normal approximations with Malliavin calculus*, volume 192 of *Cambridge Tracts in Mathematics*. Cambridge University Press, Cambridge, 2012. From Stein’s method to universality.
- [NS06] A. Nica and R. Speicher. *Lectures on the combinatorics of free probability*, volume 335 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, Cambridge, 2006.
- [Pis98] G. Pisier. Non-commutative vector valued L_p -spaces and completely p -summing maps. *Astérisque*, (247):vi+131, 1998.
- [PX97] G. Pisier and Q. Xu. Non-commutative martingale inequalities. *Comm. Math. Phys.*, 189(3):667–698, 1997.
- [Tao12] T. Tao. *Topics in random matrix theory*, volume 132 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2012.
- [Tro12] J. A. Tropp. User-friendly tail bounds for sums of random matrices. *Found. Comput. Math.*, 12(4):389–434, 2012.
- [Tro15] J. A. Tropp. *An Introduction to Matrix Concentration Inequalities*. Foundations and Trends in Machine Learning. Now, 2015. To appear. Available at <http://www.arXiv.org/abs/1501.01571>.
- [vH15] R. van Handel. On the spectral norm of inhomogeneous random matrices. Available at <http://www.arXiv.org/abs/1502.05003>, Feb. 2015.
- [Ver12] R. Vershynin. Introduction to the non-asymptotic analysis of random matrices. In *Compressed sensing*, pages 210–268. Cambridge Univ. Press, Cambridge, 2012.
- [VW08] R. Vale and S. Waldron. Tight frames generated by finite nonabelian groups. *Numer. Algorithms*, 48(1-3):11–27, 2008.