Overlaps between eigenvectors of correlated random matrices

Joël Bun,[1] Jean-Philippe Bouchaud,[2] and Marc Potters[2]
[1]*Independent Researcher, France*
[2]*Capital Fund Management, 23–25 rue de l'Université, 75 007 Paris, France*

(Received 31 March 2016; published 29 November 2018)

We obtain general, exact formulas for the overlaps between the eigenvectors of large correlated random matrices, with additive or multiplicative noise. These results have potential applications in many different contexts, from quantum thermalization to high-dimensional statistics. We find that the overlaps only depend on measurable quantities, and do *not* require the knowledge of the underlying "true" (noiseless) matrices. We apply our results to the case of empirical correlation matrices, that allow us to estimate reliably the width of the spectrum of the true correlation matrix, even when the latter is very close to the identity. We illustrate our results on the example of stock returns correlations, which clearly reveal a nontrivial structure for the bulk eigenvalues. We also apply our results to the problem of matrix denoising in high dimensions.

DOI: 10.1103/PhysRevE.98.052145

I. INTRODUCTION

The structure of the eigenvalues and eigenvectors of large random matrices is of primary importance in many different contexts, from quantum mechanics to high-dimensional data analysis. Correspondingly, random matrix theory (RMT) has established itself as a major discipline, at the frontier between theoretical physics, mathematics, probability theory, and applied statistics, with a somewhat intimidating corpus of knowledge [1]. One of the most striking applications of RMT concerns quantum chaos and quantum transport [2], with renewed interest coming from problems of quantum ergodicity ("eigenstate thermalization") [3,4], entanglement, and dissipation (for recent reviews, see Refs. [5,6]). In the context of signal processing, RMT is of primary importance in the analysis of high-dimensional statistics [7–9], wireless communication channels [10,11], etc. Other examples cover chemical physics [12] or the dynamics of complex systems—from random ecologies [13] to glasses and spin glasses [14].

Whereas the spectral properties of random matrices have been investigated at length, the interest has recently shifted to the statistical properties of their eigenvectors—see, e.g., Refs. [3,15–22] for more recent papers. In particular, suppose that we investigate a random system with N constituents whose interactions are governed by a fixed matrix $\mathbf{C}$, which is often call *population* (or *pure*). However, we rarely observed $\mathbf{C}$ directly in practice, but we rather collect realizations (or *samples*) of the system from which we try to infer $\mathbf{C}$. By definition, the inferred matrix is subject to measurement noise, which is why we shall call it the sample (or noisy) matrix $\mathbf{S}$. To make this setting more concrete, one can think of $\mathbf{C}$ as a covariance matrix and $\mathbf{S}$ its empirical estimate. The main question is how does the sample matrix resemble those of the population one? As alluded to above, we have been able answer this question for the eigenvalues since the seminal work of Ref. [23], but for the eigenvectors the answer is more recent. More precisely, we obtained in Ref. [24] explicit formulas for the overlaps between these pure and noisy eigenvectors for a wide class of random matrices, generalizing results obtained for sample covariance and correlation matrices of Ref. [16]—obtained as $\mathbf{S} = \sqrt{\mathbf{C}}\mathcal{W}\sqrt{\mathbf{C}}$, where $\mathcal{W}$ is a Wishart matrix (see Supplemental Material [25] for a precise definition)—and for matrices of the form $\mathbf{S} = \mathbf{C} + \mathbf{W}$, where $\mathbf{W}$ is a symmetric Gaussian random matrix [19,26,27]. Even though these models are very specific, they captures true physical systems such as dressed Hamiltonian [4], random ecologies [13], large sensor networks [11], or financial correlation matrices [28], for instance.

In the present paper, we want to generalize these results to the overlaps between the eigenvectors of *two* different realizations of such random matrices that remain correlated through their common part $\mathbf{C}$. For example, imagine one measures the sample correlation matrix of the same process, but on two nonoverlapping time intervals, characterized by two independent realizations of the Wishart noises $\mathcal{W}$ and $\tilde{\mathcal{W}}$. How close are the corresponding eigenvectors expected to be? We provide exact, explicit formulas for these overlaps in the high-dimensional regime. Precisely, we give a transparent interpretation to our formulas and generalize them to various cases, in particular, when the noises are correlated. Perhaps surprisingly, these overlaps can be evaluated from the empirical spectrum of $\mathbf{S}$ only, i.e., *without any prior knowledge* of the pure matrix $\mathbf{C}$ itself. We emphasize here that the noisy matrices $\mathcal{W}$ and $\tilde{\mathcal{W}}$ are drawn from the same distribution but do not necessarily share the same parameters. So, for instance, the two independent Wishart noises can come from independent sets of different sizes. These results lead us to propose a statistical test based on these overlaps that allows one to determine whether two realizations of the random matrix $\mathbf{S}$ and $\tilde{\mathbf{S}}$ indeed correspond to the very same underlying "true" matrix $\mathbf{C}$ either in the multiplicative and additive cases defined above. We shall also revisit the theory of rotational invariant estimators (RIEs) [29] that encountered much attention recently (see Refs. [16,28,30,31] to cite a few).

II. THEORETICAL RESULTS

A. Inversion formula

Throughout the following, we consider $N \times N$ symmetric random matrices and denote by $\lambda_1 \geqslant \lambda_2 \geqslant \cdots \geqslant \lambda_N$ the eigenvalues of $\mathbf{S}$ and by $\mathbf{u}_1, \mathbf{u}_2, \ldots, \mathbf{u}_N$ the corresponding eigenvectors. Similarly, we denote by $\tilde{\lambda}_1 \geqslant \tilde{\lambda}_2 \geqslant \cdots \geqslant \tilde{\lambda}_N$ the eigenvalues of $\tilde{\mathbf{S}}$ and by $\tilde{\mathbf{u}}_1, \tilde{\mathbf{u}}_2, \ldots, \tilde{\mathbf{u}}_N$ the associated eigenvectors. Note that we will sometimes index the eigenvectors by their corresponding eigenvalues for convenience. We emphasize that we allow the parameters that describe the randomness of $\mathbf{S}$ and $\tilde{\mathbf{S}}$ to be different. The central objects that we focus on in this study are the asymptotic ($N \to \infty$) *scaled, mean-squared overlaps* [32],

$$\Phi(\lambda, \tilde{\lambda}) := N\mathbb{E}[\langle \mathbf{u}_\lambda, \tilde{\mathbf{u}}_{\tilde{\lambda}} \rangle^2], \tag{1}$$

that remain $\mathcal{O}(1)$ in the limit $N \to \infty$. In the above equation, the expectation $\mathbb{E}$ can be interpreted either as an average over different realizations of the randomness or, for fixed randomness, as an average over small "slices" of eigenvalues of width $\eta = d\lambda \gg N^{-1}$, such that the result becomes self-averaging in the large N limit. We will study the asymptotic behavior of (1) using the complex function

$$\psi_N(z, \tilde{z}) := \mathbb{E}\left[\frac{1}{N} \operatorname{Tr}[(z - \mathbf{S})^{-1}(\tilde{z} - \tilde{\mathbf{S}})^{-1}]\right], \tag{2}$$

where $z, \tilde{z} \in \mathbf{C}$. For large random matrices, we expect the eigenvalues $[\lambda_i]_{i\in[[1,N]]}$ and $[\tilde{\lambda}_i]_{i\in[[1,N]]}$ to stick to their *classical* locations, i.e., smoothly allocated with respect to the quantile of the spectral density. Differently said, the sample eigenvalues become deterministic in the large N limit. Hence, we obtain after taking the continuous limit that $\psi_N(z, \tilde{z}) \sim \psi(z, \tilde{z})$ where the limiting value is given by

$$\psi(z, \tilde{z}) := \int\int \frac{\varrho(\lambda)}{z - \lambda} \frac{\tilde{\varrho}(\tilde{\lambda})}{\tilde{z} - \tilde{\lambda}} \Phi(\lambda, \tilde{\lambda}) d\lambda d\tilde{\lambda}, \tag{3}$$

with ϱ and $\tilde{\varrho}$ are the spectral densities of $\mathbf{S}$ and $\tilde{\mathbf{S}}$. Then, it suffices to compute

$$\psi(x - i\eta, y \pm i\eta) = \int\int \frac{(x - \lambda + i\eta)}{(x - \lambda)^2 + \eta^2} \frac{(y - \tilde{\lambda} \mp i\eta)}{(y - \tilde{\lambda})^2 + \eta^2} \times \varrho(\lambda)\tilde{\varrho}(\tilde{\lambda})\Phi(\lambda, \tilde{\lambda}) d\lambda d\tilde{\lambda},$$

to deduce that

$$\operatorname{Re}[\psi(x - i\eta, y + i\eta) - \psi(x - i\eta, y - i\eta)] = 2\iint \frac{\eta\varrho(\lambda)}{(x - \lambda)^2 + \eta^2} \frac{\eta\tilde{\varrho}(\tilde{\lambda})}{(y - \tilde{\lambda})^2 + \eta^2} \Phi(\lambda, \tilde{\lambda}) d\lambda d\tilde{\lambda}. \tag{4}$$

We may now invoke the Sokhotski-Plemelj identity to obtain the *inversion* formula

$$\Phi(\lambda, \tilde{\lambda}) = \frac{\operatorname{Re}[\psi_0(z, \bar{\tilde{z}}) - \psi_0(z, \tilde{z})]}{2\pi^2 \varrho(\lambda)\tilde{\varrho}(\tilde{\lambda})}, \tag{5}$$

with $z = \lambda - i\eta$, $\bar{z}$ its complex conjugate, and $\psi_0 \equiv \lim_{\eta\downarrow 0^+} \psi$. The orthonormality of eigenbases implies the following normalization for $\Phi(\lambda, \tilde{\lambda})$,

$$\int \Phi(\lambda, \tilde{\lambda})\rho(\lambda) d\lambda = 1. \tag{6}$$

Equations (2) and (5) can be computed explicitly in two trivial cases. First, if $\mathbf{S}$ and $\tilde{\mathbf{S}}$ are *free*, then $\psi(z, \tilde{z}) = \mathfrak{g}(z)\tilde{\mathfrak{g}}(\tilde{z})$ with $\mathfrak{g}(z)$ is the Stieltjes transform of $\mathbf{S}$ defined as the limiting value of $N^{-1}\operatorname{Tr}[(z - \mathbf{S})^{-1}]$. Equation (5) then gives $\Phi(\lambda, \tilde{\lambda}) = 1$ as expected since *free* matrices can be understood as matrices with random eigenvectors relative to one another. The other trivial case is when $\tilde{\mathbf{S}} = \mathbf{S}$. In that case, we have $\psi(z, \tilde{z}) = [\mathfrak{g}(z) - \mathfrak{g}(\tilde{z})]/(\tilde{z} - z)$ and $\Phi(\lambda, \tilde{\lambda}) = \delta(\lambda - \tilde{\lambda})/\rho(\lambda)$, i.e., the only overlap is the self-overlap.

Equation (5) tells us that in the high-dimensional regime, we can study the mean-squared overlap (1) through the bivariate function $\psi(z, \tilde{z})$ which is easier to handle using tools from RMT (see below).

B. Multiplicative noise

The study of the asymptotic behavior of the function ψ requires one to control the resolvent of $\mathbf{S}$ and $\tilde{\mathbf{S}}$ entrywise. It was shown recently that one can approximate these (random) resolvents entrywise by deterministic equivalent quantities [24,33,34]. We begin first with the Wishart multiplicative noise $\boldsymbol{\mathcal{W}}$ that we introduced in the Introduction. More precisely, let $\mathbf{S} := \sqrt{\mathbf{C}}\boldsymbol{\mathcal{W}}\sqrt{\mathbf{C}}$ and $\tilde{\mathbf{S}} := \sqrt{\mathbf{C}}\tilde{\boldsymbol{\mathcal{W}}}\sqrt{\mathbf{C}}$, where $\boldsymbol{\mathcal{W}}$ and $\tilde{\boldsymbol{\mathcal{W}}}$ are two independent Wishart matrices with possibly two different observation ratios $q := N/T$ and $\tilde{q} := N/\tilde{T}$. By independence, we have

$$\psi_N(z, \tilde{z}) = \frac{1}{N}\sum_{k,l}^{N} \mathbb{E}_{\mathcal{P}}\big[(z - \mathbf{S})^{-1}_{kl}\big]\mathbb{E}_{\tilde{\mathcal{P}}}\big[(\tilde{z} - \tilde{\mathbf{S}})^{-1}_{kl}\big], \tag{7}$$

where $\mathbb{E}_{\mathcal{P}}[\cdot]$ ($\mathbb{E}_{\tilde{\mathcal{P}}}[\cdot]$) denotes the expectation value over the probability measure $\mathcal{P}$ ($\tilde{\mathcal{P}}$) associated with $\mathbf{S}$ ($\tilde{\mathbf{S}}$). Then, we use the *deterministic* estimate of the resolvent of $\mathbf{S}$ which yields for $N \to \infty$ [24,33,34]

$$\mathbb{E}_{\mathcal{P}}\big[(z - \mathbf{S})^{-1}_{kl}\big] \sim \zeta(z)[z\zeta(z) - \mathbf{C}]^{-1}_{kl} + O(N^{-1/2}), \tag{8}$$

where we defined

$$\zeta(z) := \frac{1}{1 - q + qz\mathfrak{g}(z)}. \tag{9}$$

The estimate (8) holds as well for $\tilde{\mathbf{S}}$ by replacing ζ, q, and $\mathfrak{g}$ with $\tilde{\zeta}$, $\tilde{q}$, and $\tilde{\mathfrak{g}}$. Note that we can deduce the fixed-point equation associated with $\mathfrak{g}(z)$ by taking the normalized trace in (8), and this yields for $N \to \infty$,

$$\mathfrak{g}(z) \sim \zeta(z)\mathfrak{g}_{\mathbf{C}}[z\zeta(z)], \tag{10}$$

where $\mathfrak{g}_{\mathbf{C}}$ is the Stieltjes transform associated with the pure matrix $\mathbf{C}$. Again, the value of $\tilde{\mathfrak{g}}$ is obtained from (10) by replacing ζ with $\tilde{\zeta}$.

By plugging Eq. (9) into Eq. (7), we get

$$\psi_N(z, \tilde{z}) \sim \frac{1}{N}\operatorname{Tr}\{\zeta(z)[z\zeta(z) - \mathbf{C}]^{-1}\tilde{\zeta}(\tilde{z})[\tilde{z}\tilde{\zeta}(\tilde{z}) - \mathbf{C}]^{-1}\},$$

and then, using the identity

$$[z\zeta(z) - \mathbf{C}]^{-1}[\tilde{z}\tilde{\zeta}(\tilde{z}) - \mathbf{C}]^{-1} = \frac{1}{\tilde{z}\tilde{\zeta}(\tilde{z}) - z\zeta(z)}\{[z\zeta(z) - \mathbf{C}]^{-1} - [\tilde{z}\tilde{\zeta}(\tilde{z}) - \mathbf{C}]^{-1}\}, \tag{11}$$

we obtain

$$\psi_N(z,\tilde{z}) \sim \frac{\zeta(z)\tilde{\zeta}(\tilde{z})}{\tilde{z}\tilde{\zeta}(\tilde{z}) - z\zeta(z)} \frac{1}{N} \mathrm{Tr}\{[z\zeta(z) - \mathbf{C}]^{-1} - [\tilde{z}\tilde{\zeta}(\tilde{z}) - \mathbf{C}]^{-1}\}.$$

From this last equation, we deduce

$$\psi_N(z,\tilde{z}) \sim \frac{1}{\tilde{z}\tilde{\zeta}(\tilde{z}) - z\zeta(z)} \left(\frac{\tilde{\zeta}(\tilde{z})}{N} \mathrm{Tr}\{\zeta(z)[z\zeta(z) - \mathbf{C}]^{-1}\} - \frac{\zeta(z)}{N} \mathrm{Tr}\{\tilde{\zeta}(\tilde{z})[\tilde{\zeta}(\tilde{z}) - \mathbf{C}]^{-1}\} \right).$$

One notices from (10) that the two normalized trace terms in the latter equation are exactly given by $\mathfrak{g}(z)$ and $\tilde{\mathfrak{g}}(\tilde{z})$ in the large N limit. We therefore conclude that in the case of a multiplicative Wishart perturbation, the asymptotic value of (2) reads

$$\psi(z,\tilde{z}) \sim \frac{\tilde{\zeta}(\tilde{z})\mathfrak{g}(z) - \zeta(z)\tilde{\mathfrak{g}}(\tilde{z})}{\tilde{z}\tilde{\zeta}(\tilde{z}) - z\zeta(z)}, \tag{12}$$

which holds for any $q = \mathcal{O}(1)$ and $\tilde{q} = \mathcal{O}(1)$. The striking observation in Eq. (12) is that the result does not depend *explicitly* on the population matrix $\mathbf{C}$ that we wish to estimate. This feature is crucial since it indicates that we shall be able to characterize the mean-squared overlap (1) in terms of observable variables only.

Now that we have determined the asymptotic value $\psi(z,\tilde{z})$, let us now compute the main quantity of interest, i.e., Eq. (1). To that end, it is convenient to work with the complex function $m(z) := 1/[z\zeta(z)]$. Indeed, by expressing (12) in terms of the function m, we end up with

$$\psi(z,\tilde{z}) = \frac{1}{q\tilde{q}z\tilde{z}} \left[\frac{(\tilde{q}z - q\tilde{z})\tilde{m}^2}{m - \tilde{m}} + \frac{(q-\tilde{q})\tilde{m}}{m-\tilde{m}} \right] + \frac{m+\tilde{m}}{q\tilde{z}} - \frac{1-q}{qz\tilde{z}}.$$

Defining $m_0(\lambda) = \lim_{\eta\downarrow 0} m(\lambda - i\eta) \equiv m_R(\lambda) + i m_I(\lambda)$, one obtains after some elementary computations (see Supplemental Material [25] for details) the following general result,

$$\Phi_{q,\tilde{q}}(\lambda,\tilde{\lambda}) = \frac{2(\tilde{q}\lambda - q\tilde{\lambda})\alpha(\lambda,\tilde{\lambda}) + (\tilde{q}-q)\beta(\lambda,\tilde{\lambda})}{\lambda\tilde{\lambda}\gamma(\lambda,\tilde{\lambda})}, \tag{13}$$

where we defined

$$\begin{aligned}
\alpha(\lambda,\tilde{\lambda}) &:= m_R(\lambda)|\tilde{m}_0(\tilde{\lambda})|^2 - \tilde{m}_R(\tilde{\lambda})|m_0(\lambda)|^2,\\
\beta(\lambda,\tilde{\lambda}) &:= |\tilde{m}_0(\tilde{\lambda})|^2 - |m_0(\lambda)|^2,\\
\gamma(\lambda,\tilde{\lambda}') &:= \{[m_R(\lambda) - m_R(\tilde{\lambda})]^2 + [m_I(\lambda) + \tilde{m}_I(\tilde{\lambda})]^2\}\\
&\quad \times \{[m_R(\lambda) - \tilde{m}_R(\tilde{\lambda})]^2 + [m_I(\lambda) - \tilde{m}_I(\tilde{\lambda})]^2\}.
\end{aligned} \tag{14}$$

The final result (13) is invariant under the exchange of (q,λ) with $(\tilde{q},\tilde{\lambda})$ and does not depend on $\mathbf{C}$ explicitly, as expected. We will see in the next section that it is a crucial feature in order to establish an observable stability test. This result holds for independent samples and as long as the parameters q and $\tilde{q}$ are of order $\mathcal{O}(1)$ [see Ref. [34] for a rigorous statement on the applicability of Eq. (8)].

It is easy to show that this formula reproduces the mean-squared overlap between a given sample eigenvector and its true value in the limit $\tilde{q} \to 0$ (see, e.g., Refs. [16,24]). To prove our claim, let us consider $\tilde{q} \to 0$, for which $\tilde{\lambda} \to \mu$, where μ denotes the corresponding population eigenvalue [35]. In that specific framework, we have $\tilde{m}_R = 1/\mu$ and $\tilde{m}_I = 0$. Hence, we deduce from (13) that

$$\Phi_{q,\tilde{q}\to 0}(\lambda,\mu) = \frac{q\mu}{\lambda|1-\mu m_0(\lambda)|^2}, \tag{15}$$

which is exactly the result of Ref. [16].

Next, we look at the case where we split our data sets in two windows of the same size ($q = \tilde{q}$) which is relevant when ones wishes to measure the stability of the eigenvectors associated to the same eigenvalue. For $q = \tilde{q}$, the eigenvalues λ and $\tilde{\lambda}$ are now distributed according to the same density function so that $\tilde{m}(\tilde{\lambda}) = m(\tilde{\lambda})$. Moreover, we infer from (14) that the contribution of $\beta(\lambda,\tilde{\lambda})$ in (13) vanishes. The self-overlap limit $\tilde{\lambda} \to \lambda$ needs to be handled with care as the formula (13) seems ill defined when $q = \tilde{q}$. Nevertheless, if we write $\tilde{\lambda} = \lambda + \varepsilon$ with $\varepsilon > 0$, one has

$$\begin{aligned}
\alpha(\lambda,\lambda+\varepsilon) &= \varepsilon^2(|m_0|^2\partial_\lambda m_R - m_R\partial_\lambda|m_0|^2) + \mathcal{O}(\varepsilon^3),\\
\gamma(\lambda,\lambda+\varepsilon) &= 4\varepsilon^2 m_I^2(\lambda)|\partial_\lambda m_0(\lambda)|^2 + \mathcal{O}(\varepsilon^3).
\end{aligned}$$

As a consequence, we conclude by plugging these two expressions into (13) and then setting $\varepsilon = 0$ that the *self-overlap* is given by

$$\Phi(\lambda,\lambda) = \frac{q}{2\lambda^2} \frac{|m_0(\lambda)|^4 \partial_\lambda[m_R(\lambda)/|m_0(\lambda)|^2]}{m_I^2(\lambda)|\partial_\lambda m_0(\lambda)|^2}. \tag{16}$$

We now explain how we can extend these results to more general multiplicative noise. More specifically, let us consider matrices of the form $\mathbf{S} = \sqrt{\mathbf{C}}\mathbf{O}\mathbf{B}\mathbf{O}^*\sqrt{\mathbf{C}}$, where $\mathbf{O}$ is a random matrix chosen in the orthogonal group $O(N)$ according to the Haar measure and $\mathbf{B}$ is a given random matrix independent from $\mathbf{C}$ and $\mathbf{O}$ (see, e.g., Refs. [21,24,36] for similar models). The framework investigated above corresponds to the case where $\mathbf{O}\mathbf{B}\mathbf{O}^*$ is a Wishart matrix. Using the results of Ref. [24], we find for this general model that (12) still holds with $\zeta(z) = \mathcal{S}_{\mathbf{B}}(z\mathfrak{g}(z) - 1)$, where $\mathcal{S}_{\mathbf{B}}$ is the so-called Voiculescu's $\mathcal{S}$-transform of the $\mathbf{B}$ matrix [37]. If $\mathbf{B} = \mathcal{W}$, then $\mathcal{S}_{\mathbf{B}}(\omega) = 1/(1+q\omega)$. However, it seems difficult at this stage to obtain an explicit formula for the mean-squared overlap (1) in this general case since the analytic structure of the $\mathcal{S}$-transform depends on the choice of $\mathbf{B}$.

C. Additive noise

All the arguments that we use for the multiplicative noise model can be repeated nearly verbatim for the additive noise. In the case of additive real symmetric Gaussian noise, referred to as the Gaussian orthogonal ensemble (GOE) in the literature, we have $\mathbf{S} = \mathbf{C} + \mathbf{W}$ and $\tilde{\mathbf{S}} = \mathbf{C} + \tilde{\mathbf{W}}$ with $\mathbf{W}$ and $\tilde{\mathbf{W}}$ two independent GOE matrices with variance σ^2 and $\tilde{\sigma}^2$. We assume throughout the following that both variances are finite and hence do not scale with the dimension N. Each entry of the resolvent $(z - \mathbf{S})^{-1}$ may also be approximated by a deterministic value in the high-dimensional regime [19,24,34],

$$\mathbb{E}_{\mathcal{P}}\big[(z-\mathbf{S})^{-1}_{kl}\big] \sim \big[\zeta^a(z) - \mathbf{C}^{-1}_{kl}\big], \tag{17}$$

where we defined [38]

$$\zeta^a(z) := z - \sigma^2 \mathfrak{g}(z). \tag{18}$$

Once again, the asymptotic limit (17) holds for $\tilde{\mathbf{S}}$ by replacing ζ^a, $\mathfrak{g}$, and σ by $\tilde{\zeta}^a$, $\tilde{\mathfrak{g}}$, and $\tilde{\sigma}$. By performing the same computations as above, we obtain for the limiting value of ψ_N,

$$\psi^a(z,\tilde{z}) = \frac{\mathfrak{g}(z) - \tilde{\mathfrak{g}}(\tilde{z})}{\tilde{\zeta}^a(\tilde{z}) - \zeta^a(z)}. \tag{19}$$

As for Eq. (12), Eq. (19) depends only on *a priori* observable quantities since it does not involve explicitly the unknown population matrix $\mathbf{C}$. Consequently, we will obtain an observable expression for (1) using the inversion formula (5).

We can now turn on the computation of the mean-squared overlap (1) in the additive Gaussian model. From (19), we find that

$$\lim_{\eta\to 0}[\psi^a(\lambda - i\eta, \tilde{\lambda} + i\eta) - \psi^a(\lambda - i\eta, \tilde{\lambda} - i\eta)]$$
$$= \frac{\mathfrak{g}_0(\tilde{\zeta}^a - \overline{\tilde{\zeta}^a}) + \zeta^a(\overline{\tilde{\mathfrak{g}}_0} - \tilde{\mathfrak{g}}_0) + \mathfrak{g}_0\overline{\tilde{\zeta}^a} - \overline{\tilde{\mathfrak{g}}_0}\tilde{\zeta}^a}{(\overline{\tilde{\zeta}^a} - \zeta^a)(\tilde{\zeta}^a - \zeta^a)},$$

where we used the notation $\mathfrak{g}_0 \equiv \lim_{\eta\downarrow 0^+} \mathfrak{g}$. Defining $\zeta^a_0 = \lim_{\eta\downarrow 0^+} \zeta^a(\lambda - i\eta) \equiv \zeta^a_R + i\zeta^a_I$ and performing similar algebraic manipulations as above (see Supplemental Material [25] for details), we eventually get

$$\Phi^a_{\sigma,\tilde{\sigma}}(\lambda,\tilde{\lambda}) = \frac{2(\tilde{\sigma}^2\lambda - \sigma^2\tilde{\lambda})\big(\zeta^a_R - \tilde{\zeta}^a_R\big) + (\sigma^2 - \tilde{\sigma}^2)\beta^a(\lambda,\tilde{\lambda})}{\gamma^a(\lambda,\tilde{\lambda})}, \tag{20}$$

with $\gamma^a(\lambda,\tilde{\lambda})$ given by the same expression as $\gamma(\lambda,\tilde{\lambda})$ in Eq. (14) with the substitutions $m_R \to \zeta^a_R$ and $m_I \to \zeta^a_I$ and

$$\beta^a(\lambda,\tilde{\lambda}) := \big[\zeta^a_R(\lambda) - \zeta^a_I(\lambda)\big]\big[\zeta^a_R(\lambda) + \zeta^a_I(\lambda)\big] - \big[\tilde{\zeta}^a_R(\tilde{\lambda}) - \tilde{\zeta}^a_I(\tilde{\lambda})\big]\big[\tilde{\zeta}^a_R(\tilde{\lambda}) - \tilde{\zeta}^a_I(\tilde{\lambda})\big]. \tag{21}$$

We notice that the contribution of the term β^a again vanishes when we suppose that both Gaussian noises have the same variance. As in the multiplicative case, the self-overlap $\Phi^a(\lambda,\lambda)$ is reached by expanding (20) in powers of ε with $\tilde{\lambda} = \lambda + \varepsilon$. This eventually yields, by taking $\varepsilon = 0$,

$$\Phi^a(\lambda,\lambda) = \frac{\sigma^2}{2}\frac{\partial_\lambda \zeta^a_R(\lambda)}{\big[\zeta^a_I(\lambda)\big]^2 + \big|\partial_\lambda \zeta^a_0(\lambda)\big|^2}. \tag{22}$$

Similarly to the multiplicative case, the additive model can be generalized to $\mathbf{S} = \mathbf{C} + \mathbf{OBO}^*$ with the same definitions for $\mathbf{B}$ and $\mathbf{O}$. In that case, the above result (19) holds but now with $\xi(z) = z - \mathcal{R}_{\mathbf{B}}[\mathfrak{g}(z)]$, where $\mathcal{R}_{\mathbf{B}}(z)$ is the $\mathcal{R}$-transform of the $\mathbf{B}$ matrix [37]—which is simply equal to $\mathcal{R}_{\mathbf{B}}(z) = \sigma^2 z$ when $\mathbf{B} = \mathbf{W}$ is a Gaussian random matrix, as considered above.

Another interesting and important extension of the result (20) is when the noises $\mathbf{W}$, $\tilde{\mathbf{W}}$ are correlated—while the above calculations referred to independent noises. In the additive case, the trick is to realize that one can always write (in law) $\mathbf{W} = \sqrt{\rho}\mathbf{W}_0 + \sqrt{1-\rho}\mathbf{W}_1$ and $\tilde{\mathbf{W}} = \sqrt{\rho}\mathbf{W}_0 + \sqrt{1-\rho}\mathbf{W}_2$, where $\mathbf{W}_1, \mathbf{W}_2$ are now independent, as above. Since our formulas do not rely on the common matrix $\mathbf{C}$, it can therefore be replaced by $\mathbf{C} + \sqrt{\rho}\mathbf{W}_0$. Then, Eq. (22) trivially holds with σ^2 replaced by $\sigma_1^2 - \sigma_1\sigma_2\rho$ with σ_1 and σ_2 the width of the noisy matrices $\mathbf{W}_1$ and $\mathbf{W}_2$ (see Supplemental Material [25] for more details). Similarly, $\tilde{\sigma}^2$ is replaced by $\sigma_2^2 - \sigma_1\sigma_2\rho$. Note that in the case where the noise parameters are identical, σ^2 is simply multiplied by $1-\rho$. The corresponding shape of $\Phi^a(\lambda,\lambda)$ for different values of ρ and $\sigma = \tilde{\sigma}$ is shown in Fig. 1. We also provide in the inset a comparison with synthetic data for a fixed $\rho = 0.54$, $\sigma = 1$. The empirical average is taken over 200 realizations of the noise and, again, the agreement is excellent.

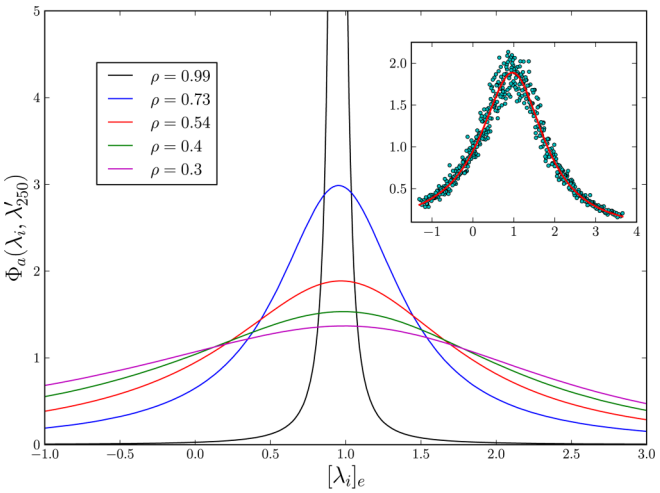

FIG. 1. Main figure: Evaluation of the self-overlap $\Phi_a(\lambda,\tilde{\lambda})$ for a fixed $\tilde{\lambda} \approx 0.95$ as a function of λ for $N = 500$, $\sigma = 1$, and for different values of ρ. The population matrix $\mathbf{C}$ is given by a (white) Wishart matrix with parameter $T = 2N$. Inset: We compare the theoretical prediction $\Phi_a(\lambda,\tilde{\lambda} \approx 0.95)$ for a fixed $\rho = 0.54$ with synthetic data. The empirical averages (blue points) are obtained from 100 independent realizations of $\mathbf{S}$.

Considering correlated noises in the multiplicative model is of crucial importance since it describes the case of correlation matrices measured on overlapping periods, such that $\mathbf{S} = \sqrt{\mathbf{C}}[\mathcal{W}_0 + \mathcal{W}]\sqrt{\mathbf{C}}$ and $\mathbf{S}' = \sqrt{\mathbf{C}}[\mathcal{W}_0 + \tilde{\mathcal{W}}]\sqrt{\mathbf{C}}$ (see Sec. III B below). This case turns out to be more subtle and is the subject of ongoing investigations.

D. A convolution formula

Before investigating some concrete applications of these formulas, let us end this theoretical section with an interesting interpretation of the above formalism. We first introduce the set of eigenvectors $\mathbf{v}_\mu$ of the pure matrix $\mathbf{C}$, labeled by the eigenvalue μ. We define the overlaps between the $\mathbf{u}$'s and the $\mathbf{v}$'s as $\sqrt{\Phi_0(\mu,\lambda)/N} \times \varepsilon(\mu,\lambda)$, where $\Phi_0(\mu,\lambda)$ were explicitly computed in Ref. [24] for a wide class of problems, and $\varepsilon(\mu,\lambda)$ are random variables of unit variance. Now, one can always decompose the $\mathbf{u}$'s as

$$\mathbf{u}_\lambda = \frac{1}{\sqrt{N}}\int d\mu \varrho_C(\mu)\sqrt{\Phi_0(\mu,\lambda)}\varepsilon(\mu,\lambda)\mathbf{v}_\mu, \tag{23}$$

where ϱ_C is the spectral density of $\mathbf{C}$. Using the orthonormality of the $\mathbf{v}$'s, one then finds

$$\langle \mathbf{u}_\lambda, \mathbf{u}'_{\lambda'} \rangle = \frac{1}{N} \int d\mu \varrho_C(\mu) \sqrt{\Phi_0(\mu,\lambda)\Phi_0(\mu,\lambda')}\varepsilon(\mu,\lambda)\varepsilon(\mu,\lambda').$$

If we square this last expression and average over the noise, and make an "ergodic hypothesis" [3] according to which all signs $\varepsilon(\mu,\lambda)$ are in fact independent from one another, one finds the following, rather intuitive, convolution result for squared overlaps,

$$\Phi(\lambda,\lambda') = \int d\mu \varrho_C(\mu)\Phi_0(\mu,\lambda)\Phi_0(\mu,\lambda'). \tag{24}$$

It turns out that this expression is completely general and exactly equivalent to Eqs. (13) and (20) in the corresponding cases. However, whereas this expression still contains some explicit dependence on the structure of the pure matrix $\mathbf{C}$, it has disappeared in Eqs. (13) and (20). Nevertheless, this second interpretation will be useful in order to obtain an efficient way to estimate $\mathbf{C}$ from large noisy matrices.

III. APPLICATION

Equations (13) and (20) are exact in the high-dimensional limit (HDL). Note that from Ref. [34], we expect these results to hold with fluctuations of order $N^{-1/2}$ but a more rigorous analysis of the subleading terms can be useful for practical purposes. We leave this question for future work. Throughout this section, we shall focus on the case of sample covariance and correlation matrices but most of the results that will follow can be easily transposed to the additive noise. We emphasize that in the special case of sample correlation matrices, the HDL is defined as

$$N, T, \tilde{T} \to \infty \quad \text{with} \quad q = \mathcal{O}(1), \quad \tilde{q} = \mathcal{O}(1), \tag{25}$$

where T is the sample size. We will assume throughout this section that the variance of each variable can be estimated independently with great accuracy in the HDL so that we will not distinguish further covariances and correlations henceforth.

The first application concerns a stability test for the eigenvectors of large correlation matrices. More precisely, we investigate whether or not the mean-squared overlap between the eigenvectors of two correlation matrices measured with nearby nonoverlapping samples is entirely explain by measurement noise. Differently said, we test the hypothesis that the dynamics of the eigenvectors is captured by the sample correlation matrix model. The second application involves the convolution formula. In particular, we link our results with the theory of RIEs that provides significant improvement over classical sample estimates in the HDL (see Ref. [28] for a recent review).

A. Eigenvector stability

The first application deals with the stability of the eigenvectors in the case of two nonoverlapping adjacent samples. In order to give more insight, we begin with a theoretical example where the true correlation matrix $\mathbf{C}$ is an inverse Wishart matrix of parameter $\kappa \in (0,\infty)$, that corresponds to $1/q$ for Wishart matrices (see Ref. [24] for details). In that case, the function $m(z)$ can be explicitly computed. This finally leads to

$$\Phi(\lambda,\lambda) = \frac{\upsilon(\lambda+2q\kappa)^2}{2q\kappa(2\lambda(\upsilon+\kappa)-\lambda^2\kappa+\kappa(2q\upsilon-1))}, \tag{26}$$

with $\upsilon := 1+q\kappa$ and λ is within the interval $[\lambda^-,\lambda^+]$, where the edges are given by $\lambda^\pm = \kappa^{-1}[\upsilon+\kappa \pm \sqrt{(2\kappa+1)(2q\kappa+1)}]$. An interesting limit corresponds to $\kappa \to \infty$, where $\mathbf{C}$ tends to the identity matrix, and the overlaps are expected to become all equal to $1/N$. Indeed, one finds, for a fixed q,

$$\Phi(\lambda,\lambda') \underset{\kappa\to\infty}{\sim} \left[1 + \frac{(\lambda-1)(\lambda'-1)}{2q^2\kappa} + \mathcal{O}\left(\frac{1}{\kappa^2}\right)\right], \tag{27}$$

which is in fact *universal* in this limit, provided the eigenvalue spectrum of $\mathbf{C}$ has a variance given by $(2\kappa)^{-1} \to 0$ [39]. This formula is interesting insofar as it allows one to estimate the width of the eigenvalue distribution of $\mathbf{C}$, even when it is close to the identity matrix, i.e., $\kappa \gg 1$. One could think of directly using information on the empirical spectrum, for example, the Marčenko-Pastur prediction $\mathrm{Tr}\,\mathbf{C}^{-1} = (\mathbf{1}-\mathbf{q})\mathrm{Tr}\,\mathbf{S}^{-1}$, that in principle allows one extract the parameter κ through $1+(2\kappa)^{-1} = (1-q)\mathrm{Tr}\,\mathbf{S}^{-1}/N$. However, this second method is numerically unstable and very imprecise when $\kappa \gg 1$ and finite N (for one thing, the right-hand side can be negative, which would lead to a negative variance). Our formula based on overlaps avoids these difficulties. Note that we can generalize Eq. (27) to any two noise matrices. Indeed, provided that $\mathcal{S}_\mathbf{C}(\omega) = 1-\sigma^2\omega$ with $\sigma^2 := N^{-1}\,\mathrm{Tr}\,\mathbf{C^2} - \mathbf{1}$, one has for $\sigma^2 \to 0$,

$$\Phi(\lambda,\lambda') = 1 + \sigma^2[2\mathfrak{g}_R(\lambda)-1][2\tilde{\mathfrak{g}}_R(\tilde{\lambda})-1] + \mathcal{O}(\sigma^4), \tag{28}$$

where the subscript R has the same meaning as above.

As an illustration, we check the validity of Eq. (26) in Fig. 2 with $\kappa = 10$, $N = 500$, and $q = 0.5$. More precisely, we determine the empirical average overlap as follows: We consider 50 independent realization of the Wishart noise $\mathcal{W}$. For each pair of samples we compute a smoothed overlap as

$$[\langle \mathbf{u}_i, \tilde{\mathbf{u}}_i \rangle^2] = \frac{1}{Z_i}\sum_{j=1}^{N} \frac{\langle \mathbf{u}_i, \tilde{\mathbf{u}}_j \rangle^2}{(\lambda_i - \lambda'_j)^2 + \eta^2}, \tag{29}$$

with $Z_i = \sum_{k=1}^N [(\lambda_i - \lambda'_k)^2 + \eta^2]^{-1}$ the normalization constant and η the width of the Cauchy kernel, that we choose to be $N^{-1/2}$ in such a way that $N^{-1} \ll \eta \ll 1$. We then average this quantity over all pairs for a given value of i to obtain $[\langle \mathbf{u}_i, \tilde{\mathbf{u}}_i \rangle^2]_e$, and plot the resulting quantity as a function of the average eigenvalue position $[\lambda_i]_e$. We observe that the agreement with Eq. (26) is excellent, even when the true underlying matrix $\mathbf{C}$ is close to the identity matrix. Note that the empirical estimate Eq. (29) is universal, i.e., independent of the underlying structure of $\mathbf{C}$.

For a general and arbitrary population matrix $\mathbf{C}$, evaluating Eq. (13)—or Eq. (20)—is rather difficult because of finite-size effects, especially for the multiplicative case. Indeed, when we consider multiplicative noises, the eigenvalues of $\mathbf{S}$ are confined to stay positive meaning the presence of a *hard* wall

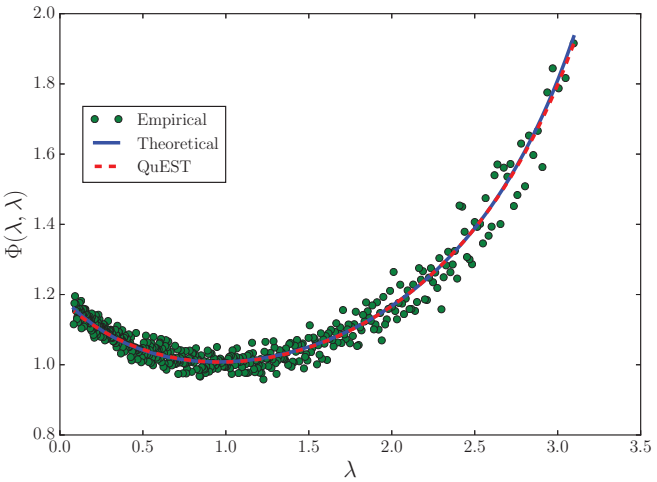

FIG. 2. Evaluation of the self-overlap for an inverse-Wishart population matrix with $N = 500$ and $\kappa = 10$ using formula (27). The blue plain line corresponds to the case where we know exactly the true eigenvalues while the red dotted line is obtained from the estimated population eigenvalues using the QuEST algorithm (see Ref. [40]). The empirical average is plotted by the green points and is obtained from (29) over 50 realizations of $\mathcal{W}$.

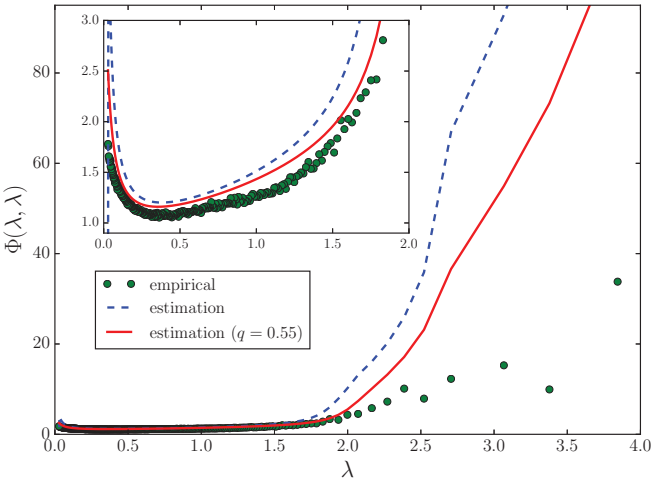

FIG. 3. Evaluation of the self-overlap $\Phi(\lambda, \lambda)$ as a function of the sample eigenvalues λ using the $N = 300$ most liquid U.S. equities from 2004 to 2013. We split the data into two nonoverlapping periods with the same sample size of 1200 business days. For each period, we randomly select $T = 600$ days and we repeat $B = 100$ bootstraps of the original data. The empirical self-overlap is computed using Eq. (29) over these 100 bootstraps (green points) and the limiting formula (16) is estimated using the QuEST algorithm with $q = 0.5$ (blue dashed line). We also provide the estimation we get using the same effective observation ratio $q_{\text{eff}} = 0.55$ that accounts for correlation and heavy tail effects [28]. Inset: Focus in the bulk of eigenvalues.

at the origin [41]. Consequently, the use of local laws to estimate the Stieltjes transform $\mathfrak{g}(z)$, as in Ref. [34], often leads to noisy results for very small eigenvalues. Hence, the determination of Eq. (13) from real data is rather difficult and one has to resort to numerical regularization schemes to do so. In the specific case of the sample correlation matrix, one possible solution is to invert the celebrated Marčenko-Pastur equation [31,40] to infer the eigenvalues of the population matrix $\mathbf{C}$. Once this is done, one can evaluate the Stieltjes transform $\mathfrak{g}(z)$ with high precision even near the origin. In the following, we shall use the so-called Quantized Eigenvalues Sampling Transform (QuEST) numerical scheme of Ref. [40] to obtain these *pure* eigenvalues. We plot in Fig. 2 the results obtained when using the estimated population eigenvalues from the QuEST algorithm (red dotted line) and note that the agreement is quite remarkable.

Now that we have an estimate of the population eigenvalues, we can investigate an application to real data. Here, we study the case of the U.S. stock market but the results below can be extended to other regions [28]. The difficulty when dealing with real data is to measure the empirical mean-squared overlaps (29) between two nonoverlapping correlation matrices $\mathbf{S}$ and $\tilde{\mathbf{S}}$ as in Eq. (29) because we may not have enough data points to evaluate accurately an average over the noise as required in Eq. (1). To circumvent this problem, we use a bootstrap procedure to increase the size of the data [42]: We take a total period of 2400 business days from 2004 to 2013 for the $N = 300$ most liquid assets of the S&P 500 index that we split into two nonoverlapping subsets of the same size of 1200 days, corresponding to 2004 to 2008 and 2008 to 2013. We restrict to $N = 300$ stocks such that all of them are present throughout the whole period from 2004 to 2013. Then, for each subset and each bootstrap sample $b \in \{1, \ldots, B\}$, we select randomly $T = 600$ distinct days to construct two "independent" sample correlation matrices $\mathbf{S}_b$ and $\tilde{\mathbf{S}}_b$, with $q = \tilde{q} = N/T = 0.5$. We then compute the empirical mean-squared overlap (1) and also the theoretical limit (13)—using the QuEST algorithm—from these B bootstrap data sets.

For our simulations, we set $B = 100$ and plot in Fig. 3 the resulting estimation of Eq. (1) we get from the QuEST algorithm (blue dashed line) and the empirical bootstrap estimate (29) (green points) using U.S. stocks. We also perform the estimation with an effective observation ratio $q_{\text{eff}} = 0.55$ (red plain line) as advocated in Ref. [28] for the S&P 500 index, to account for correlation or heavy tail effects.

It is clear from Fig. 3 that the eigenvectors associated with large eigenvalues are not well described by the theory: We notice a discrepancy between the (estimated) theoretical curve and the empirical one even after accounting for an effective ratio q_{eff}. The difference is even worse for the market mode (not shown). This is presumably related to the fact that the largest eigenvectors are expected to genuinely evolve with time, as already argued in Ref. [17]. Note also the gap at the left edge between the theoretical and empirical prediction in the inset of Fig. 3 that is partly corrected with q_{eff}. This suggests that one can still improve the Marčenko-Pastur framework by adding, e.g., autocorrelation or heavy tailed entries which allows one to widen the spectral density of $\mathbf{S}$ (see, e.g., Refs. [36,43] for autocorrelation and [44–47] for heavy tailed entries). Finally, all these remarks hold for other markets as well [28].

B. Rotational invariant estimator

Aside from the statistics of eigenvectors, the theoretical framework presented is actually very useful for the

estimation of $\mathbf{C}$ from large noisy matrices in the specific class of rotational invariant estimators (see Ref. [28] and references therein for a recent review on this topic). For this class of estimators, we keep the observed eigenvectors $\mathbf{u}_i$ of $\mathbf{S} := \sqrt{\mathbf{C}}\mathcal{W}\sqrt{\mathbf{C}}$ and replace the eigenvalues with the *oracle* eigenvalues,

$$\xi_i \equiv \xi_i(q) := \langle \mathbf{u}_i, \mathbf{C}\mathbf{u}_i \rangle. \tag{30}$$

There exist several ways to approximate this estimator in the HDL: using the anisotropic local law [28], numerical inversion of the Marčenko-Pastur equation [40], and the so-called cross-validation (CV) estimator of Ref. [30]. The remarkable feature of all these methods is that the resulting estimator only depends on observable quantities, while it is clearly not the case for (30). Even if the first two techniques are interesting on their own, we will rather focus on the last one as it turns to be related to the convolution formula derived in Sec. II D.

From now on, we consider the multiplicative but the following arguments can easily be generalized for the additive case. If we have a second sample $\tilde{\mathbf{S}}$, we can estimate the oracle with the quantity

$$\nu_i \equiv \nu_i(q) := \langle \mathbf{u}_i, \tilde{\mathbf{S}}\mathbf{u}_i \rangle, \tag{31}$$

where we recall that $\mathbf{u}_i$ is independent from $\tilde{\mathbf{S}} := \sqrt{\mathbf{C}}\tilde{\mathcal{W}}\sqrt{\mathbf{C}}$. We again assume that we are in the regime (25) and we thus infer from (24) that

$$\begin{aligned}
\nu_i &\sim \int \tilde{\varrho}(\tilde{\lambda})\Phi(\lambda,\tilde{\lambda})\tilde{\lambda}d\tilde{\lambda} \\
&= \int \tilde{\varrho}(\tilde{\lambda})\left[\int \varrho_{\mathbf{C}}(\mu)\Phi_0(\lambda,\mu)\Phi_0(\tilde{\lambda},\mu)d\mu\right]\tilde{\lambda}d\tilde{\lambda} \\
&= \int \varrho_{\mathbf{C}}(\mu)\Phi_0(\lambda,\mu)\left[\int \tilde{\varrho}(\tilde{\lambda})\Phi_0(\tilde{\lambda},\mu)\tilde{\lambda}d\tilde{\lambda}\right]d\mu.
\end{aligned} \tag{32}$$

The term in the brackets can be simplified by using the very definition of $\tilde{\mathbf{S}}$,

$$\begin{aligned}
\int \tilde{\varrho}(\tilde{\lambda})\Phi_0(\tilde{\lambda},\mu_j)\tilde{\lambda}d\tilde{\lambda} &\approx \langle \mathbf{v}_j, \tilde{\mathbf{S}}\mathbf{v}_j \rangle \\
&= \langle \mathbf{C}^{1/2}\mathbf{v_j}, \tilde{\mathcal{W}}\mathbf{C}^{1/2}\mathbf{v_j} \rangle,
\end{aligned}$$

and we rewrite this last line owing to the eigenvalue equation as

$$\begin{aligned}
\int \tilde{\varrho}(\tilde{\lambda})\Phi_0(\tilde{\lambda},\mu_j)\tilde{\lambda}d\tilde{\lambda} &\approx \mu_j \langle \mathbf{v}_j, \tilde{\mathcal{W}}\mathbf{v}_j \rangle \\
&= \mu_j \sum_{k=1}^{N} \omega_k \mathbb{E}[\langle \mathbf{w}_k, \mathbf{v}_j \rangle^2],
\end{aligned}$$

with ω_k the kth eigenvalue of the white Wishart matrix $\tilde{\mathcal{W}}$ and $\mathbf{w}_k$ the corresponding eigenvector. Finally, we invoke that $\mathbb{E}[\langle \mathbf{w}_k, \mathbf{v}_j \rangle^2] = N^{-1}$ for all $j \in [[1, N]]$ to conclude that in the HDL,

$$\sum_{k=1}^{N} \omega_k \langle \mathbf{w}_k, \mathbf{v}_j \rangle^2 = 1, \tag{33}$$

and we therefore have

$$\int \tilde{\varrho}(\tilde{\lambda})\Phi_0(\tilde{\lambda},\mu_j)\tilde{\lambda}d\tilde{\lambda} \approx \mu_j. \tag{34}$$

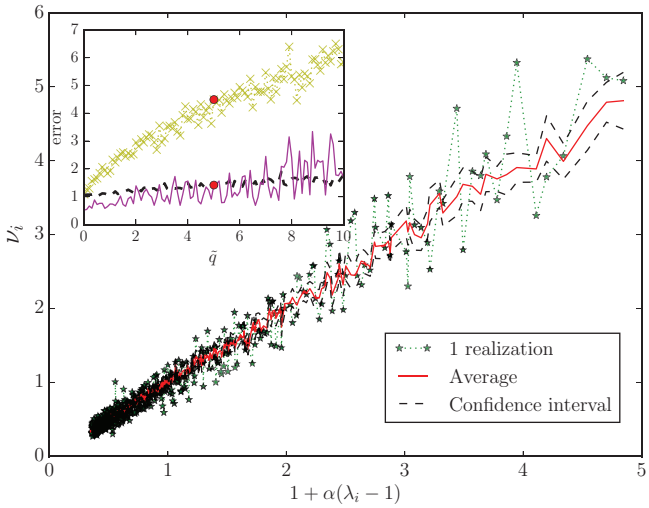

FIG. 4. Main figure: Evaluation of Eq. (35) when $\mathbf{C}$ is a 500×500 inverse-Wishart matrix with parameter $\kappa = 0.5$. The first noisy matrix $\mathbf{S}$, which is the one we wish to denoise, is drawn from a Wishart distribution with a parameter $q = 0.5$. The second noisy matrix $\tilde{\mathbf{S}}$ is also a Wishart matrix but with a parameter $\tilde{q} = 5$. The x axis is given by the true asymptotic value, namely, the linear shrinkage estimator with intensity $\alpha = 1/(1+2q\kappa)$. The y axis are the eigenvalues obtained from (35). The dotted star line is the result obtained from one sample, the plain red line is the average results from 20 independent realizations of $\tilde{\mathbf{S}}$, and the dark dashed lines give the confidence interval. Inset: Squared error as a function of $\tilde{q}$ where we kept $\mathbf{S}$ fixed with the same N and T. The yellow crossed line corresponds to the error over a single realization of $\tilde{\mathbf{S}}$ for Eq. (35) and the purple plain line to its sorted version. The dashed black line corresponds to the average value of Eq. (35) over 20 realizations of $\tilde{\mathbf{S}}$. The red point gives the error for $\tilde{q} = 5$, i.e., the sample shown in the main figure.

Plugging this last equation into (32), we obtain for any $\tilde{q} = \mathcal{O}(1)$ the following result,

$$\nu_i(q) \sim \int \varrho_{\mathbf{C}}(\mu)\Phi_0(\lambda,\mu)\mu d\mu \equiv \xi_i(q), \tag{35}$$

where the last equivalence comes from the definition (30) of the oracle estimator in the continuous limit.

This result is very interesting and indicates that one can approximate the oracle estimator (30) by considering the quadratic form between the eigenvectors of a given realization of $\mathbf{C}$—say, $\mathbf{S}$—and another realization of $\mathbf{C}$—say, $\tilde{\mathbf{S}}$—even if the latter is characterized by a different value of the quality ratio $\tilde{q} \neq q$.

To illustrate this last point, let us consider an inverse-Wishart matrix with parameter $\kappa = 0.5$ as the population correlation matrix of size $N = 500$. Both noisy matrices are drawn from a multivariate Gaussian distribution but with different parameters: The first noisy matrix $\mathbf{S}$ is computed using $T = 1000$ while the second one $\tilde{\mathbf{S}}$ corresponds to $\tilde{T} = 100$. With this prior, the oracle estimator is given in the HDL by the linear shrinkage with intensity $\alpha = 1/(1+2q\kappa)$ and $q = N/T$. In Fig. 4, we plot the prediction obtained from (35) with $\mathbf{S}$ fixed and a single realization of $\tilde{\mathbf{S}}$ (dotted star line) and we see that although noisy, the prediction is already

fairly accurate. We also plot in the same figure the average value of (35) over 20 independent realizations of $\tilde{\mathbf{S}}$ (plain red line) and we observe that the agreement with limiting value (given by the line $y = x$ in Fig. 4) is excellent, with only very small fluctuations (see the confidence interval given by the blue shaded area).

Quite surprisingly, we can significantly improve the accuracy of the estimation by applying an *ad hoc* regularization even for a single realization of $\tilde{\mathbf{S}}$. More specifically, we see from Fig. 4 that the prediction (35) does not necessarily preserve the order of the eigenvalues due to the finite size of the sample. However, observing a nonmonotonic cleaning scheme may be an unwanted feature within a rotational invariant assumption. Indeed, there is no reason *a priori* to expect that it is optimal to modify the order of the eigenvalues, that is to say, the variance associated with the principal components.

There are several ways to regularize the estimation obtained from (4). We can either sort the cleaned eigenvalues [28] or perform an *isotonic* regression [30]. We provide an illustration of the sorting regularization in the inset of Fig. 4 over a single realization of $\tilde{\mathbf{S}}$ (purple plain line). The improvement over (35) (yellow crossed line) in terms of squared error is significant. Moreover, even if the estimation becomes quite noisy for large values of $\tilde{q}$, we notice that the quality of estimation is still on par with the average value over 20 realizations of $\tilde{\mathbf{S}}$ (red plain line), which is quite remarkable. We also want to emphasize that in all cases, the error we obtain is always less than 9, which is the error we get when keeping the sample eigenvalues of $\mathbf{S}$ (see the inset of Fig. 4).

A takeaway from Fig. 4 is that we can indeed use the result (35) even when $\tilde{q} \gg q$. Hence, this gives a simple way to check the quality of the estimation by comparing the in-sample result (i.e., the estimator we obtain using the information of $\mathbf{S}$) and the out-of-sample result (obtained with the information of $\tilde{\mathbf{S}}$). It therefore demonstrates the validity of the out-of-sample test of Ref. [28] for assessing the quality of *empirical* optimal RIE using financial data. The second finding is that it is possible to estimate quite accurately the optimal oracle for a fixed value of q by using a relatively small amount of independent data from a *single* realization.

C. Cross-validation estimator

We can now apply the analysis of the previous section to the cross-validation (CV) estimator proposed in Ref. [30]. The idea behind this estimator is to split a single data set into two disjoint sets (potentially of unequal sizes) and build the matrices $\mathbf{S}$ and $\tilde{\mathbf{S}}$ and compute $\nu_i(q)$ as above. We then average over all possible choices of split. More precisely, suppose that we want to estimate (30) out of T independent samples that we split into K nonoverlapping sets whose indices are denoted by $\{\mathcal{I}_\varsigma\}_{\varsigma=1}^K$. The CV estimator then reads

$$\nu_i^{\mathrm{cv}}(q_\varsigma) := \frac{1}{K}\sum_{\varsigma=1}^{K}\sum_{t\in\mathcal{I}_\varsigma}\left\langle \mathbf{u}_i^{(\varsigma)}, \frac{\mathbf{x}_t\mathbf{x}_t^*}{|\mathcal{I}_\varsigma|}\mathbf{u}_i^{(\varsigma)}\right\rangle, \tag{36}$$

where each set $\mathcal{I}_\varsigma$ has an equal size such that $K|\mathcal{I}_\varsigma| = T$, $\mathbf{u}_i^{(\varsigma)}$ is the eigenvector associated with the ith eigenvalue obtained from the sample correlation matrix in which we removed all the observations belonging to the set $\mathcal{I}_\varsigma$ for a fixed

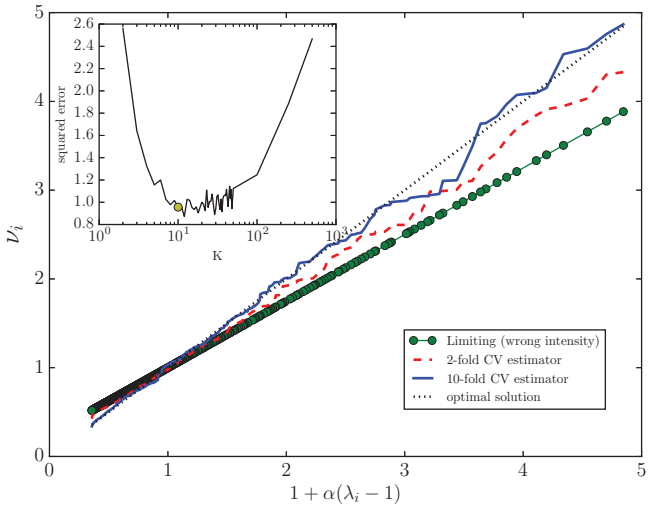

FIG. 5. Main figure: Evaluation of Eq. (36) using the same experience than in Fig. 4. The x axis represents to the optimal cleaning [with $\alpha = 1/(1+2q\kappa)$] while the y axis are the eigenvalues obtained from (36) for a single realization of $\mathbf{S}$. The black dotted line is the option solution, the blue plain line is the sorted tenfold CV estimator ($|\mathcal{I}_\varsigma| = 100$), and the red dashed line corresponds to the sorted twofold case ($|\mathcal{I}_\varsigma| = 500$). We also provide the linear shrinkage with a wrong intensity $\alpha = 1/(1+2q_\varsigma\kappa)$, where $q_\varsigma = 1$ corresponds to the case $|\mathcal{I}_\varsigma| = 500$. Inset: We plot the squared error of (36) for the same realization of $\mathbf{S}$ as a function of K. We see that the minimum is attained at $K = 10$ (yellow point).

$\varsigma \in [[1, K]]$, and $q_\varsigma := N/(T - |\mathcal{I}_\varsigma|)$ is the corresponding observation ratio. If we denote by $\tilde{\mathbf{S}}_\varsigma$ the sample correlation matrix associated with the observations of the set $\mathcal{I}_\varsigma$, then we can rewrite (36) as

$$\nu_i^{\mathrm{cv}} = \frac{1}{K}\sum_{\varsigma=1}^{K}\left\langle \mathbf{u}_i^{(\varsigma)}, \tilde{\mathbf{S}}_\varsigma\mathbf{u}_i^{(\varsigma)}\right\rangle, \tag{37}$$

which can be thought as an average version of the estimator (35).

The crucial difference resides in the observation ratios q_ς associated with the eigenvector $\mathbf{u}_i^{(\varsigma)}$ and $\tilde{q}_\varsigma := N/|\mathcal{I}_\varsigma|$ associated with the matrix $\tilde{\mathbf{S}}_\varsigma$. Hence, we clearly have a trade-off in the choice of the cardinality of the $\mathcal{I}_\varsigma$: Choosing $|\mathcal{I}_\varsigma|$ too large implies that q_ς deviates strongly from $q := N/T$ while a too small cardinality leads to the noisy limit $\tilde{q}_\varsigma \sim N$ in which the convergence (35) becomes dubious. We therefore understand from this rewriting why the *leave-one-out* case, i.e., $|\mathcal{I}_\varsigma| = 1$, will not return a reliable estimate of the optimal value (30) even after averaging since $\tilde{q}_\varsigma = N$. For a suitable value of the cardinality $|\mathcal{I}_\varsigma|$, one can expect that $\tilde{q}_\varsigma$ is small enough in order to be in the regime of Eq. (35) and, more importantly, that $\mathbf{u}_i^{(\varsigma)}$ is not too far from $\mathbf{u}_i(q)$. In that case, provided that the samples are independent of each other, we shall have

$$\nu_i^{\mathrm{cv}}(q_\varsigma) \sim \nu_i(q). \tag{38}$$

We provide a numerical check of this result in Fig. 5 in the case of a sorted tenfold CV estimate (blue plain line) where we reconsider the same configuration of Fig. 4. Note that a

tenfold CV leads to a "test" set of size $|\mathcal{I}_\varsigma| = 100$ meaning that $q_\varsigma \approx 0.55$, which is not that far from the true value of q. We see that the agreement with the optimal value (black dotted line) is excellent, showing that for this specific choice of $|\mathcal{I}_\varsigma|$, the convergence (38) holds. In order to illustrate the trade-off discussed above, we also plot the twofold CV estimator for which $q_\varsigma = 1$. In that case, the result we obtain from (36) rather coincides with the linear shrinkage with an intensity of $\alpha = 1/(1 + 2q_\varsigma\kappa)$. This demonstrates that the choice of K in (36) is crucial in order to obtain a consistent estimation of (30) (see the inset of Fig. 5).

From a practical viewpoint, the estimator (37) is quite appealing since it can be generalized to more general classes of random processes. It provides a simple tool to approximate the oracle estimator (30) with great accuracy in the high-dimensional regime beyond the sample correlation matrix model. From a theoretical perspective, we see that there is still some work to be done in order to understand the convergence expressed by (38). Indeed, the relationship between $\mathbf{u}_i^{(\varsigma)}$ and $\mathbf{u}_i$, which corresponds to the overlapping and correlated case discussed above in the additive case, should provide insights to establish the convergence result as a function of the cardinality of $|\mathcal{I}_\varsigma|$. Moreover, the study of subleading terms in Eq. (35) can be of particular interest to understand the phase transition that occurs depending on the value of q_ς and $\tilde{q}_\varsigma$.

IV. CONCLUSION

In summary, we have provided general, exact formulas for the overlaps between the eigenvectors of large correlated random matrices, with additive or multiplicative noise. Remarkably, these results do not require the knowledge of the underlying "pure" matrix and have a broad range of applications in different contexts. We showed that the cross-sample eigenvector overlaps provide unprecedented information about the structure of the true eigenvalue spectrum, much beyond that contained in the empirical spectrum itself. For example, the width of the bulk of the spectrum of the true underlying correlation matrix can be reliably estimated, even when the latter is very close to the identity matrix. We have illustrated our results on the example of stock returns correlations, that clearly reveal a nontrivial structure for the bulk eigenvalues. We have also discussed the application to matrix denoising and we saw in particular that it is possible to make use of these overlaps between independent samples to approximate the so-called oracle estimator with great accuracy.

ACKNOWLEDGMENTS

We thank R. Allez, D. Bartz, R. Bénichou, R. Chicheportiche, B. Collins, A. Rej, E. Sérié, and D. Ullmo for very useful discussions.

APPENDIX A: WISHART MATRICES

While GOE (or Wigner) matrices are now well known in physics with several applications, it is not the case for Wishart matrices. We provide in this Appendix exact definitions of the model and discuss briefly its asymptotic behavior. It is believed that the first random matrix model comes from the statistician Wishart in Ref. [49], which considers the sample covariance matrix of N random variables. More precisely, let $\mathbf{X}$ be an $N \times T$ random matrix where each column is a realization of a multivariate Gaussian with zero mean and population covariance matrix $\mathbf{C}$. The Wishart matrix $\boldsymbol{\mathcal{W}}$ is then defined by

$$\boldsymbol{\mathcal{W}} := \frac{1}{T}\mathbf{X}\mathbf{X}^*. \tag{A1}$$

We see that $\mathbb{E}[\boldsymbol{\mathcal{W}}] = \mathbf{C}$ and this is why this matrix is very useful in practice. The asymptotic behavior of $\boldsymbol{\mathcal{W}}$ is universal and depends only on the parameter $q := N/T$. The universality means that the statistics of the eigenvalues and the mean-squared overlaps of $\boldsymbol{\mathcal{W}}$ do not depend on the exact realization of $\mathbf{X}$. More precisely, the entries of the resolvent $(z\mathbf{I}_N - \boldsymbol{\mathcal{W}})^{-1}$ are self-averaging and statistics about the spectrum of $\boldsymbol{\mathcal{W}}$ are derived from it. Furthermore, we stress that the Gaussian assumption is not important but simplifies the computation. We refer the reader to Ref. [34] for a rigorous presentation of the universality of Wishart matrices.

A Wishart matrix with covariance matrix $\mathbf{C}$ can be written as $\boldsymbol{\mathcal{W}} = \sqrt{\mathbf{C}}\boldsymbol{\mathcal{W}}_\mathbf{I}\sqrt{\mathbf{C}}$, where $\boldsymbol{\mathcal{W}}_\mathbf{I}$ is a *white Wishart*, i.e., a Wishart with the identity as its covariance matrix. In this article, all the Wishart matrices are white Wishart so we drop the term *white* and denote them $\boldsymbol{\mathcal{W}}$.

APPENDIX B: MEAN-SQUARED OVERLAP FOR INDEPENDENT SAMPLE COVARIANCE MATRICES

We keep the notations of Sec. II B and shall often omit the arguments z and $\tilde{z}$ in m and $\tilde{m}$ in the following when there is no confusion. Using the definition $m(z) = 1/[z\zeta(z)]$, we rewrite (12) as

$$\psi(z,\tilde{z}) \sim \frac{1}{1/\tilde{m} - 1/m}\left[\frac{\mathfrak{g}}{\tilde{z}\tilde{m}} - \frac{\tilde{\mathfrak{g}}}{zm}\right], \tag{B1}$$

which is equivalent to

$$\psi(z,\tilde{z}) \sim \frac{1}{z\tilde{z}}\frac{1}{m - \tilde{m}}[z\mathfrak{g}m - \tilde{z}\tilde{\mathfrak{g}}\tilde{m}]. \tag{B2}$$

We can then express the function ψ as a function of m and $\tilde{m}$ only,

$$\psi(z,\tilde{z}) \sim \frac{1}{q\tilde{q}z\tilde{z}}\left[\frac{(\tilde{q}z - q\tilde{z})\tilde{m}^2}{m - \tilde{m}} + \frac{(q - \tilde{q})\tilde{m}}{m - \tilde{m}}\right] + \frac{m + \tilde{m}}{q\tilde{z}} - \frac{1 - q}{qz\tilde{z}}. \tag{B3}$$

To use the inversion formula Eq. (5), we need to evaluate $\psi(\lambda - i\eta, \tilde{\lambda} \pm i\eta)$. Using the shorthanded notation $m_0(\lambda) = \lim_{\eta\to 0} m(\lambda - i\eta)$, we find

$$\begin{aligned}\lim_{\eta\to 0}&[\psi(\lambda - i\eta, \tilde{\lambda} + i\eta) - \psi(\lambda - i\eta, \tilde{\lambda} - i\eta)] \\ &\sim \frac{\tilde{q}\lambda - q\tilde{\lambda}}{q\tilde{q}\lambda\tilde{\lambda}}\frac{m_0[\overline{\tilde{m}_0^2} - \tilde{m}_0^2] + \tilde{m}_0\overline{\tilde{m}_0}[\tilde{m}_0 - \overline{\tilde{m}_0}]}{(m_0 - \overline{\tilde{m}_0})(m_0 - \tilde{m}_0)} \\ &\quad + \frac{\tilde{q} - q}{q\tilde{q}\lambda\tilde{\lambda}}\frac{m_0[\overline{\tilde{m}_0} - \tilde{m}_0]}{(m_0 - \overline{\tilde{m}_0})(m_0 - \tilde{m}_0)} + \mathrm{Im},\end{aligned} \tag{B4}$$

where we omitted the explicit expressions of the imaginary part since this is not important for Eq. (5). Then, using the

representation $m_0 = m_R + i m_I$ and $\tilde{m}_0 = \tilde{m}_R + i\tilde{m}_I$, one finds

$$m_0\big[\overline{\tilde{m}_0^2} - \tilde{m}_0^2\big] + \tilde{m}_0\overline{\tilde{m}_0}[\overline{\tilde{m}_0} - \tilde{m}_0] = 2\tilde{m}_I\big[2m_I\tilde{m}_R + i\big(\tilde{m}_R^2 + \tilde{m}_R^2 - 2m_R\tilde{m}_R\big)\big],$$
$$\overline{m_0}[\tilde{m}_0 - \overline{\tilde{m}_0}] = 2\tilde{m}_I[m_I - i m_R], \tag{B5}$$

and

$$(m_0 - \overline{\tilde{m}_0})(m_0 - \tilde{m}_0) = (m_R - \tilde{m}_R)^2 - \big(m_I^2 - \tilde{m}_I^2\big) + 2i m_I(m_R - \tilde{m}_R). \tag{B6}$$

Straightforward computations yields

$$|\big(m_0 - \overline{\tilde{m}_0}\big)(m_0 - \tilde{m}_0)|^2 = \big[(m_R - \tilde{m}_R)^2 - \big(m_I^2 - \tilde{m}_I^2\big)\big]^2 + 4m_I^2(m_R - \tilde{m}_R)^2,$$
$$= \big[(m_R - \tilde{m}_R)^2 + \big(m_I^2 + \tilde{m}_I^2\big)^2\big]\big[(m_R - \tilde{m}_R)^2 + \big(m_I^2 - \tilde{m}_I^2\big)^2\big], \tag{B7}$$

which is exactly the denominator in (13). For the numerator, the elementary complex analysis in Eq. (B4) yields

$$\big(m_0\big[\overline{\tilde{m}_0^2} - \tilde{m}_0^2\big] + \tilde{m}_0\overline{\tilde{m}_0}[\overline{\tilde{m}_0} - \tilde{m}_0]\big) \times \overline{(m_0 - \overline{\tilde{m}_0})(m_0 - \tilde{m}_0)} = 4m_I\tilde{m}_I[m_R|\tilde{m}_0|^2 - \tilde{m}_R|m_0|^2], \tag{B8}$$

and

$$m_0[\overline{\tilde{m}_0} - \tilde{m}_0] \times \overline{(m_0 - \overline{\tilde{m}_0})(m_0 - \tilde{m}_0)} = 2m_I\tilde{m}_I[|\tilde{m}_0|^2 - |m_0|^2]. \tag{B9}$$

By regrouping these last three equations with the prefactors in (B4), and recalling that $m_I(\lambda) = \pi q\varrho(\lambda)$ and $\tilde{m}_I(\tilde{\lambda}) = \pi\tilde{q}\tilde{\varrho}(\tilde{\lambda})$, we obtain by using the inversion formula (5) the following result,

$$\Phi_{q,\tilde{q}}(\lambda,\tilde{\lambda}) = \frac{2(\tilde{q}\lambda - q\tilde{\lambda})[m_R|\tilde{m}_0|^2 - \tilde{m}_R|m_0|^2] + (\tilde{q} - q)[|\tilde{m}_0|^2 - |m_0|^2]}{\lambda\tilde{\lambda}[(m_R - \tilde{m}_R)^2 + (m_I + \tilde{m}_I)^2][(m_R - \tilde{m}_R)^2 + (m_I - \tilde{m}_I)^2]}, \tag{B10}$$

which is exactly Eq. (13).

APPENDIX C: MEAN-SQUARED OVERLAP FOR DEFORMED GOE MATRICES

The derivation of the overlaps (20) for two independent deformed GOEs is very similar to sample covariance matrices. Hence, we shall omit most details that can be obtained by following the arguments of the above Appendix. Again, we shall skip the arguments λ and $\tilde{\lambda}$ where there is no confusion.

In Sec. II C, we obtained

$$\lim_{\eta\to 0}[\psi(\lambda - i\eta, \tilde{\lambda} + i\eta) - \psi(\lambda - i\eta, \tilde{\lambda} - i\eta)] = \frac{\mathfrak{g}_0(\tilde{\zeta}^a - \overline{\tilde{\zeta}^a}) + \zeta^a(\overline{\tilde{\mathfrak{g}}_0} - \tilde{\mathfrak{g}}_0) + \mathfrak{g}_0\overline{\tilde{\zeta}^a} - \overline{\tilde{\mathfrak{g}}_0}\tilde{\zeta}^a}{(\overline{\tilde{\zeta}^a} - \zeta^a)(\tilde{\zeta}^a - \zeta^a)}. \tag{C1}$$

By proceeding as above [see Eq. (B4) and thereafter], we find

$$\mathfrak{g}_0(\tilde{\zeta}^a - \overline{\tilde{\zeta}^a}) + \zeta^a(\overline{\tilde{\mathfrak{g}}_0} - \tilde{\mathfrak{g}}_0) + \mathfrak{g}_0\overline{\tilde{\zeta}^a} - \overline{\tilde{\mathfrak{g}}_0}\tilde{\zeta}^a = 2\big[\big(\zeta_I^a\tilde{\mathfrak{g}}_I - \mathfrak{g}_I\tilde{\zeta}_I^a\big) + i\big(\tilde{\zeta}_I^a\big(\mathfrak{g}_R - \tilde{\mathfrak{g}}_R\big) - \tilde{\mathfrak{g}}_I(\zeta_R^a - \tilde{\zeta}_R^a)\big)\big] \tag{C2}$$

and

$$(\overline{\tilde{\zeta}^a} - \zeta^a)(\tilde{\zeta}^a - \zeta^a) = \big(\zeta_R^a - \tilde{\zeta}_R^a\big)^2 + \big[\big(\tilde{\zeta}_I^a\big)^2 - \big(\zeta_I^a\big)^2\big] + 2i\zeta_I^a\big(\tilde{\zeta}_R^a - \zeta_R^a\big). \tag{C3}$$

Hence, by putting these last two equations into (C1) and then using (5), we get after some straightforward computations

$$\Phi_a(\lambda,\tilde{\lambda}) = \frac{(\sigma^2 + \tilde{\sigma}^2)\big(\zeta_R^a - \tilde{\zeta}_R^a\big)^2 + 2\sigma^2\tilde{\sigma}^2(\mathfrak{g}_R - \tilde{\mathfrak{g}}_R)\big(\zeta_R^a - \tilde{\zeta}_R^a\big) - (\sigma^2 - \tilde{\sigma}^2)\big[\big(\zeta_I^a\big)^2 - \big(\tilde{\zeta}_I^a\big)^2\big]}{\big[\big(\zeta_R^a - \tilde{\zeta}_R^a\big)^2 + \big(\zeta_I^a + \tilde{\zeta}_I^a\big)^2\big]\big[\big(\zeta_R^a - \tilde{\zeta}_R^a\big)^2 + \big(\zeta_I^a - \tilde{\zeta}_I^a\big)^2\big]}, \tag{C4}$$

which is exactly (20) after some manipulations. When the common matrix $\mathbf{C}$ is itself GOE, then all matrices are GOEs and more explicit results can be given. We specialize here in the case where the variance of the two noises are the same $\sigma^2 = \tilde{\sigma}^2$, and we get

$$\Phi_a(\lambda,\tilde{\lambda}) = \frac{(\sigma_{\mathbf{C}^2} + \sigma^2)(1 - f^2)}{f^2(\lambda^2 + \tilde{\lambda}^2) - f(1 + f^2)\lambda\tilde{\lambda} + (1 - f^2)^2(\sigma_{\mathbf{C}^2} + \sigma^2)}, \tag{C5}$$

where $\sigma_{\mathbf{C}^2}$ the variance of $\mathbf{C}$ and $f = \sigma_{\mathbf{C}}^2/(\sigma_{\mathbf{C}}^2 + \sigma^2)$.

APPENDIX D: THE CASE OF CORRELATED GAUSSIAN ADDITIVE NOISES

In this Appendix, we give the exact derivation of Eq. (20) with correlated noises. Let

$$\mathbf{S} = \mathbf{C} + \mathbf{W}_1, \quad \tilde{\mathbf{S}} = \mathbf{C} + \mathbf{W}_2, \tag{D1}$$

where $\mathbf{W}_1, \mathbf{W}_2$ are two correlated GOE matrices (independent from $\mathbf{C}$) satisfying

$$\langle \mathbf{W}_1 \rangle_{\mathcal{N}} = 0, \quad \langle \mathbf{W}_2 \rangle_{\mathcal{N}} = 0, \quad \mathrm{Cov}(\mathbf{W}_1, \mathbf{W}_2) = \begin{pmatrix} \sigma_1^2 & \rho_{12} \\ \rho_{12} & \sigma_2^2 \end{pmatrix}, \tag{D2}$$

where we denoted by $\mathcal{N}$ the Gaussian measure and used the abbreviation $\rho_{12} = \rho\sigma_1\sigma_2$. Using the stability of GOE under addition, let us rewrite the noise terms as

$$\mathbf{W}_1 = \mathbf{A} + \mathbf{B}_1, \quad \mathbf{W}_2 = \mathbf{A} + \mathbf{B}_2, \tag{D3}$$

where $\mathbf{A}$ that satisfies

$$\langle \mathbf{A} \rangle_{\mathcal{N}} = 0, \quad \langle \mathbf{A}^2 \rangle_{\mathcal{N}} = \rho_{12}, \tag{D4}$$

and $\mathbf{B}_1$ and $\mathbf{B}_2$ are two GOEs matrices independent from $\mathbf{A}$ with

$$\begin{cases} \langle \mathbf{B}_1 \rangle_{\mathcal{N}} = 0, \quad \langle \mathbf{B}_2 \rangle_{\mathcal{N}} = 0, \\ \langle \mathbf{B}_1^2 \rangle_{\mathcal{N}} = \sigma_1^2 - \rho_{12}, \quad \langle \mathbf{B}_2^2 \rangle_{\mathcal{N}} = \sigma_2^2 - \rho_{12}, \\ \langle \mathbf{B}_1 \mathbf{B}_2 \rangle_{\mathcal{N}} = 0. \end{cases} \tag{D5}$$

One can check that this parametrization yields exactly the correlation structure of Eq. (D2). Therefore, using (D3) into (D1), we have the equivalence (in law)

$$\mathbf{S}_1 = \mathbf{D} + \mathbf{B}_1, \quad \tilde{\mathbf{S}} = \mathbf{D} + \mathbf{B}_2, \tag{D6}$$

where we defined

$$\mathbf{D} := \mathbf{C} + \mathbf{A} \overset{\text{law}}{=} \mathbf{C} + \sqrt{\rho}\mathbf{W}_0, \tag{D7}$$

with $\mathbf{W}_0$ a GOE matrix with variance $\sigma_1\sigma_2$. Since the noises are now independent and that the mean-squared overlap Φ_a, given in Eq. (20), is "independent" from the exact structure of $\mathbf{C}$, we can therefore replace $\mathbf{C}$ by $\mathbf{D}$. Hence, we deduce that the overlaps for this model will again be given by Eq. (20) with $\sigma^2 = \sigma_1^2 - \rho_{12}$, and $\tilde{\sigma}^2 = \sigma_2^2 - \rho_{12}$, as announced.

[1] *The Oxford Handbook of Random Matrix Theory*, edited by G. Akemann, J. Baik, and P. Di Francesco (Oxford University Press, Oxford, UK, 2011).
[2] C. W. Beenakker, Rev. Mod. Phys. **69**, 731 (1997).
[3] J. M. Deutsch, Phys. Rev. A **43**, 2046 (1991).
[4] G. Ithier and F. Benaych-Georges, arXiv:1510.04352.
[5] R. Nandkishore and D. A. Huse, Annu. Rev. Condens. Matter Phys. **6**, 15 (2015).
[6] J. Eisert, M. Friesdorf, and C. Gogolin, Nat. Phys. **11**, 124 (2015).
[7] I. M. Johnstone, in *Proceedings of the International Congress of Mathematicians I* (European Mathematical Society, 2007), pp. 307–333.
[8] J.-P. Bouchaud and M. Potters, in *The Oxford Handbook of Random Matrix Theory* (Ref. [1]).
[9] D. Paul and A. Aue, J. Stat. Plan. Inference **150**, 1 (2014).
[10] A. M. Tulino and S. Verdú, Commun. Inf. Theory **1**, 1 (2004).
[11] R. Couillet and M. Debbah, *Random Matrix Methods for Wireless Communications* (Cambridge University Press, Cambridge, MA, 2011).
[12] A. A. Lee, M. P. Brenner, and L. J. Colwell, Phys. Rev. Lett. **119**, 208101 (2017).
[13] R. M. May, Nature (London) **238**, 413 (1972).
[14] Y. V. Fyodorov, Phys. Rev. Lett. **92**, 240601 (2004).
[15] M. Wilkinson and P. N. Walker, J. Phys. A: Math. Gen. **28**, 6143 (1995).
[16] O. Ledoit and S. Péché, Probab. Theory Relat. Fields **151**, 233 (2011).
[17] R. Allez and J.-P. Bouchaud, Phys. Rev. E **86**, 046202 (2012).
[18] P. Bourgade, L. Erdős, H.-T. Yau, and J. Yin, Commun. Pure Appl. Math. **69**, 1815 (2015).
[19] R. Allez and J.-P. Bouchaud, Random Matrices: Theory Appl. **3**, 1450010 (2014).
[20] A. Bloemendal, A. Knowles, H.-T. Yau, and J. Yin, Probab. Theory Relat. Fields **164**, 459 (2016).
[21] R. Couillet and F. Benaych-Georges, Elect. J. Stat. **10**, 1393 (2016).
[22] R. Monasson and D. Villamaina, Europhys. Lett. **112**, 50001 (2015).
[23] V. A. Marčhenko and L. A. Pastur, Math. USSR Sb. **1**, 457 (1967).
[24] J. Bun, R. Allez, J.-P. Bouchaud, and M. Potters, IEEE Trans. Inf. Theory **62**, 7475 (2016).
[25] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevE.98.052145 for further details on Wishart matrices and the derivation of Eqs. (13) and (20) for both uncorrelated and correlated noises.
[26] P. Biane, Quantum Probab. Commun. **11**, 55 (2003).
[27] R. Allez, J. Bun, and J.-P. Bouchaud, arXiv:1412.7108.
[28] J. Bun, J.-P. Bouchaud, and M. Potters, Phys. Rep. **666**, 1 (2017).
[29] W. James and C. Stein, in *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability* (University of California Press, Oakland, 1961), Vol. 1, pp. 361–379.
[30] D. Bartz, arXiv:1611.00798.
[31] N. E. Karoui, Ann Stat. **36**, 2757 (2008).
[32] As the dimension goes to infinity, one has to be cautious about the indexing of eigenvectors. A standard convention in the large dimensional limit is to use the eigenvalues as labels for the eigenvectors. This means that $\mathbf{u}_i$ is the eigenvector associated with the eigenvalue λ_i and that becomes $\mathbf{v}_\lambda$ as $N \to \infty$.
[33] Z. Burda, A. Görlich, A. Jarosz, and J. Jurkiewicz, Physica A **343**, 295 (2004).
[34] A. Knowles and J. Yin, Probab. Theory Relat. Fields **169**, 257 (2017).
[35] T. W. Anderson, Ann. Math. Stat. **34**, 122 (1963).
[36] Z. Burda, J. Jurkiewicz, and B. Wacław, Phys. Rev. E **71**, 026111 (2005).
[37] D. Voiculescu, Invent. Math. **104**, 201 (1991).
[38] Here and henceforth, the superscript a denotes the additive noise.

[39] The analysis for the additive case leads to a very similar result. More precisely, taking $\mathbf{C} = I_N + \mathbf{W}_0$ with $\mathbf{W}_0$ a GOE (independent from $\mathbf{W}$ and $\mathbf{W}'$) of variance $\sigma_0^2 \to 0$, one finds that $\Phi_a(\lambda, \lambda')$ is given by exactly the same formula (27) with the substitution $2q^2\kappa \to \sigma^4/\sigma_0^2$.

[40] O. Ledoit and M. Wolf, Computat Stat. Data Anal. **115**, 199 (2017).

[41] Note that this problem does not appear in the additive case, so one can safely use concentration results from Ref. [34].

[42] This technique is especially useful in machine learning and we refer the reader to, e.g., Sec. 7.11 of Ref. [48] for a more detailed explanation.

[43] D. Bartz and K.-R. Müller, in *Advances in Neural Information Processing Systems* (MIT Press, Cambridge, MA, 2014), pp. 1592–1600.

[44] G. Biroli, J. P. Bouchaud, and M. Potters, Acta Phys. Polon. Ser. B **39**, 4009 (2008).

[45] Z. Burda, J. Jurkiewicz, M. A. Nowak, G. Papp, and I. Zahed, Physica A **343**, 694 (2004).

[46] N. El Karoui, Ann. Appl. Probab. **19**, 2362 (2009).

[47] R. Couillet, F. Pascal, and J. W. Silverstein, J. Multivariate Anal. **139**, 56 (2015).

[48] J. Friedman, T. Hastie, and R. Tibshirani, *The Elements of Statistical Learning* (Springer, New York, 2001), Vol. 1.

[49] J. Wishart, Biometrika **20A**, 32 (1928).