EIGENVECTOR OVERLAPS
(AN RMT TALK)

J.-Ph. Bouchaud (CFM/Imperial/ENS)

(Joint work with Romain Allez, Joel Bun & Marc Potters)

$$\mathbf{M} = \mathbf{C} + \mathbf{O}\mathbf{B}\mathbf{O}^{\dagger}$$

Randomly Perturbed Matrices

Questions in this talk:

- How similar are
  - the eigenvectors of a « pure » matrix **C** and those of a noisy observation of **C**?
  - the eigenvectors of two independent noisy observations of **C**?
- So what?

Models of Randomly Perturbed Matrices

(Free) Additive noise

$$\mathbf{M} = \mathbf{C} + \mathbf{O}\mathbf{B}\mathbf{O}^{\dagger}$$

« Pure system »
« Signal »

« Noise »
B diag,
O random rotation

(Free) Multiplicative noise

$$\mathbf{M} = \sqrt{\mathbf{C}}\mathbf{O}\mathbf{B}\mathbf{O}^{\dagger}\sqrt{\mathbf{C}}$$

« Pure system »
« Signal »

« Noise »
B diag,
O random rotation

3

Models of Randomly Perturbed Matrices

Additive noise	Multiplicative noise
$\mathbf{M} = \mathbf{C} + \mathbf{O}\mathbf{B}\mathbf{O}^\dagger$	$\mathbf{M} = \sqrt{\mathbf{C}}\mathbf{O}\mathbf{B}\mathbf{O}^\dagger\sqrt{\mathbf{C}}$

➢ Additive examples:

- Inference of **C** given **M** + an observation noise model, eg **B** = **W**(igner)
- Quantum mechanics with a time dependent perturbation; localisation
- Dyson Brownian motion: $\mathbf{O}\mathbf{B}\mathbf{O}^\dagger$ = **W**(t) Brownian noise
  → stochastic evolution of eigenvalues & eigenvectors

➢ Multiplicative example:

- Empirical **M** vs. « True » covariance matrix **C**;
  $\mathbf{OBO}^t = \mathbf{XX}^t = \mathbf{W}$(ishart), where **X** is a N x T white noise matrix

Objects of interest: Definitions

$$\boxed{\Phi(\lambda_i, c_j) := N\mathbb{E}\left[\langle \mathbf{u}_i | \mathbf{v}_j \rangle^2\right]}$$

« Overlap »

Eigenvector of **M**

Eigenvector of **C**

Notes:

- N = size of the matrices, N >> 1 in the sequel
- **E**[..]: average over small intervals of λ >> 1/N
- The overlaps are quickly of order 1/N:

In the Dyson picture, some finite hybridisation takes place at each « collision » between eigenvalues → t_{eq} ~ 1/N

$$d|\psi_i^t\rangle = -\frac{1}{2N}\sum_{j\neq i}\frac{dt}{(\lambda_i(t)-\lambda_j(t))^2}|\psi_i^t\rangle + \frac{1}{\sqrt{N}}\sum_{j\neq i}\frac{dw_{ij}(t)}{\lambda_i(t)-\lambda_j(t)}|\psi_j^t\rangle$$

Objects of interest: Definitions

Resolvent: a central tool in RMT

$$\boxed{\mathbf{G}_{\mathbf{M}}(z) := (z\mathbf{I}_N - \mathbf{M})^{-1}}$$

Stieltjes transform and spectral density (or eigenvalue distribution)

$$\operatorname{Im} \mathfrak{g}_{\mathbf{M}}(\lambda - i\eta) \equiv \operatorname{Im} \frac{1}{N} \operatorname{Tr}\left[\mathbf{G}_{\mathbf{M}}(\lambda - i\eta)\right] = \pi\, \rho_{\mathbf{M}}(\lambda)$$

Overlaps:

$$\langle \mathbf{v}_i | \operatorname{Im} \mathbf{G}_{\mathbf{M}}(\lambda - i\eta) | \mathbf{v}_i \rangle \approx \pi \rho_{\mathbf{M}}(\lambda) \Phi(\lambda, c_i)$$

Note: everywhere the « resolution » $\eta \to 0$ but $\gg 1/N$

Objects of interest: Definitions

R-Transform

$$\mathcal{B}_{\mathbf{M}}(\mathfrak{g}_{\mathbf{M}}(z)) = z. \qquad \mathcal{R}_{\mathbf{M}}(z) := \mathcal{B}_{\mathbf{M}}(z) - \frac{1}{z}$$

e.g. the ***R***-transform of a Wigner matrix is ***R***(z)=σ^2 z

S-Transform

$$\mathcal{T}_{\mathbf{M}}(z) = z\mathfrak{g}_{\mathbf{M}}(z) - 1, \qquad \mathcal{S}_{\mathbf{M}}(z) := \frac{z+1}{z\mathcal{T}_{\mathbf{M}}^{-1}(z)}$$

e.g. the ***S***-transform of a Wishart matrix is ***S***(z)=1/(1+qz) with: **q=N/T**

Main Theoretical Result (J. Bun, R. Allez JPB, M. Potters, IEEE 2016)

Additive noise

$$\boxed{\langle \mathbf{G}_{\mathbf{M}}(z)\rangle = \mathbf{G}_{\mathbf{C}}(Z(z))}$$

$$Z(z) = z - \mathcal{R}_{\mathbf{B}}(\mathfrak{g}_{\mathbf{M}}(z))$$

Multiplicative noise

$$\boxed{z\langle \mathbf{G}_{\mathbf{M}}(z)\rangle = Z(z)\mathbf{G}_{\mathbf{C}}(Z(z))}$$

$$Z(z) = z\mathcal{S}_{\mathbf{B}}(z\mathfrak{g}_{\mathbf{M}}(z) - 1)$$

8

Notes:

- Results obtained using a replica representation of the resolvent + low rank HCIZ
- Taking the trace of these matrix equalities recovers the « free » convolution rules:

$$\boxed{\mathcal{R}_{\mathbf{M}}(z) = \mathcal{R}_{\mathbf{C}}(z) + \mathcal{R}_{\mathbf{B}}(z)}$$

$$\boxed{\mathcal{S}_{\mathbf{M}}(u) = \mathcal{S}_{\mathbf{C}}(u)\mathcal{S}_{\mathbf{B}}(u)}$$

Overlaps: simplified results

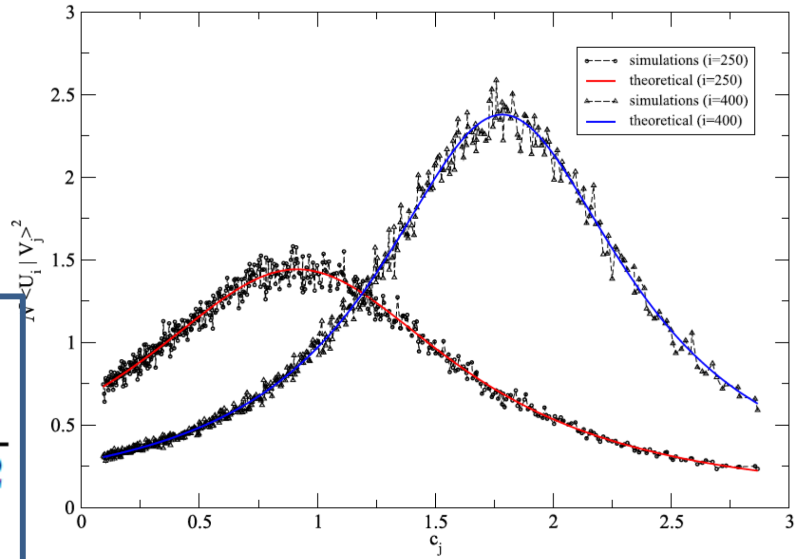

Additive noise when **B=W** (not necessarily Gaussian)

$$\boxed{\Phi(\lambda, c) = \frac{\sigma^2}{(c - \lambda + \sigma^2 \mathfrak{h}_{\mathbf{M}}(\lambda))^2 + \sigma^4 \pi^2 \rho_{\mathbf{M}}(\lambda)^2}}$$

Notes:

- Tends to a delta function when σ=0 (no noise)
- Cauchy-like formula with power-law tail decrease for large $|c - \lambda|$
- Note: True for all « Wigner-like » matrices (not necessarily Gaussian)

Empirical covariance matrices (multiplicative noise)

$$\boxed{\Phi(\lambda, c) = \frac{qc\lambda}{(c(1-q) - \lambda + qc\lambda \mathfrak{h}_{\mathbf{M}}(\lambda))^2 + q^2 \lambda^2 c^2 \pi^2 \rho_{\mathbf{M}}(\lambda)^2}}$$

Notes:

- Result first obtained by Ledoit & Péché
- Tends to a delta function when q=0 (infinite T for a fixed N)

Overlaps: the case of an outlier (again!)

Suppose **C** is of rank one, with its single non zero eigenvalue γ and **B** = **W**(t) a Brownian matrix noise

- Applying the above formalism (to order 1/N) in the additive case leads to a spectrum of **M** composed of

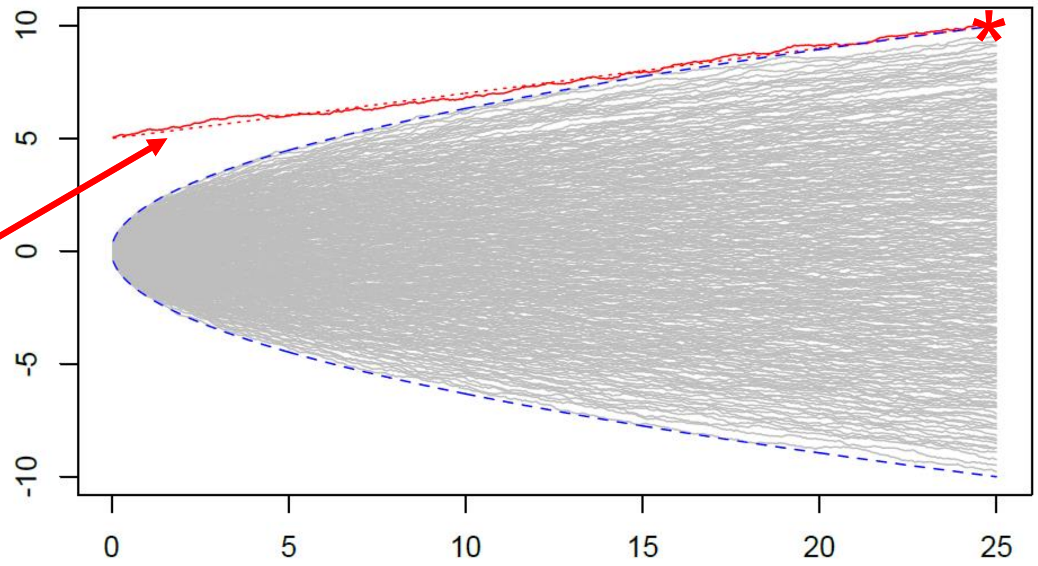

- A Wigner semi-circle of radius $2\sigma t^{1/2}$
- An isolated eigenvalue $\lambda^* = \gamma + \sigma^2 t/\gamma$

as long as $t < t^* = (\gamma/\sigma)^2$

- For $t > t^*$ the isolated eigenvalue disappears in the Wigner sea (BBP transition)

- As for the overlaps, the above results hold for the bulk; the isolated eigenvector keeps an overlap = $1 - (t/t^*)$ with its initial direction (conj: $\sim N^{-1/3}$ at t*??)

From Overlaps to Rotationally Invariant Estimators

- Assume one has no prior about **C**
- What is the best L_2 estimator $\Xi(\mathbf{M})$ of **C** knowing **M**?
- Without any indication about the directions of the eigenvectors of **C**, one is stuck with those of **M**:

$$\boxed{\Xi(\mathbf{M}) = \sum_{i=1}^{N} \xi_i \left|\mathbf{u}_i\right\rangle \left\langle\mathbf{u}_i\right|}$$

  where the ξ must be determined

- From L_2 optimality, the ξ are in principle given by $$\widehat{\xi}_i = \sum_{j=1}^{N} \langle\mathbf{u}_i|\mathbf{v}_j\rangle^2 c_j$$
- But the c's and v's are unknown…(« Oracle » estimator)

From Overlaps to Rotationally Invariant Estimators

$$\boxed{\widehat{\xi}_i = \sum_{j=1}^{N} \langle \mathbf{u}_i | \mathbf{v}_j \rangle^2 c_j}$$

- The high dimensional « miracle »

$$\begin{aligned}\widehat{\xi}_i &\underset{N\to\infty}{=} \int c\,\rho_{\mathbf{C}}(c)\,\Phi(\lambda_i, c)\,\mathrm{d}c. \\ &= \frac{1}{N\pi\rho_{\mathbf{M}}(\lambda_i)} \lim_{z\to\lambda_i - i0^+} \operatorname{Im}\operatorname{Tr}\left[\mathbf{G}_{\mathbf{M}}(z)\mathbf{C}\right]\end{aligned}$$

- In the <u>additive case</u>: $\widehat{\xi}_i = F_1(\lambda_i); \quad F_1(\lambda) = \lambda - \alpha_1(\lambda) - \beta_1(\lambda)\mathfrak{h}_{\mathbf{M}}(\lambda)$

Note 1: everything only depends on **M** !
Note 2: the formula is F(x)=Sx/(S+N) for Gaussian **C** and **B**

$$\begin{cases}\alpha_1(\lambda) := \operatorname{Re}\left[\mathcal{R}_{\mathbf{B}}\left(\mathfrak{h}_{\mathbf{M}}(\lambda) + \mathrm{i}\pi\rho_{\mathbf{M}}(\lambda)\right)\right] \\ \beta_1(\lambda) := \dfrac{\operatorname{Im}\left[\mathcal{R}_{\mathbf{B}}\left(\mathfrak{h}_{\mathbf{M}}(\lambda) + \mathrm{i}\pi\rho_{\mathbf{M}}(\lambda)\right)\right]}{\pi\rho_{\mathbf{M}}(\lambda)}\end{cases}$$

From Overlaps to Rotationally Invariant Estimators

- The multiplicative case

$$\widehat{\xi}_i = F_2(\lambda_i); \quad F_2(\lambda) = \lambda\gamma_{\mathbf{B}}(\lambda) + (\lambda\mathfrak{h}_{\mathbf{M}}(\lambda) - 1)\omega_{\mathbf{B}}(\lambda)$$

- The empirical covariance matrix case (Ledoit-Péché)

$$F_2(\lambda) = \frac{\lambda}{(1 - q + q\lambda\mathfrak{h}_{\mathbf{M}}(\lambda))^2 + q^2\lambda^2\pi^2\rho_{\mathbf{M}}^2(\lambda)}$$

- Non-linear « shrinkage », only requires **M**

(Inverse-Wishart **C** with time dependent volatilities)

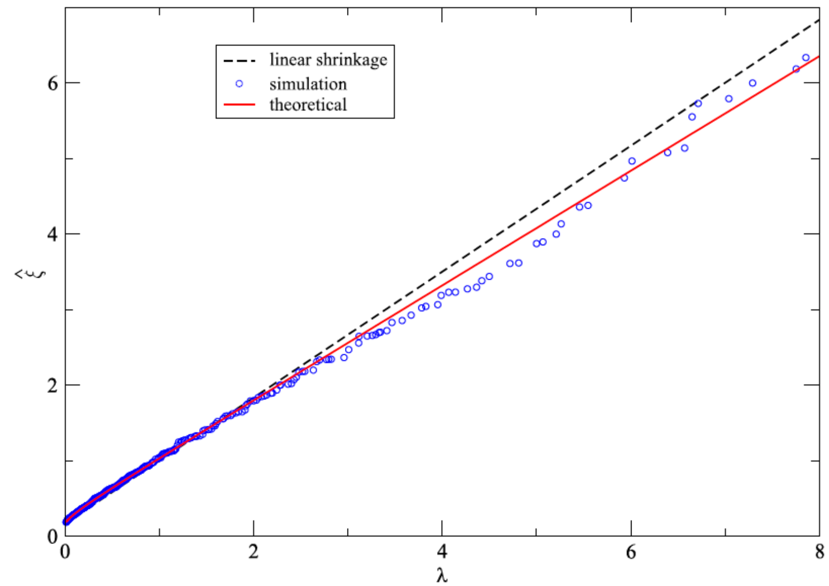

Overlaps between independent realisations

- Extending the above techniques allows us to compute the overlap

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

for *two independent* realisations, e.g. **M** = **C** + **W** and **M'** = **C** + **W'**

- The result is cumbersome but explicit, *both for the multiplicative & additive cases, e.g.*

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

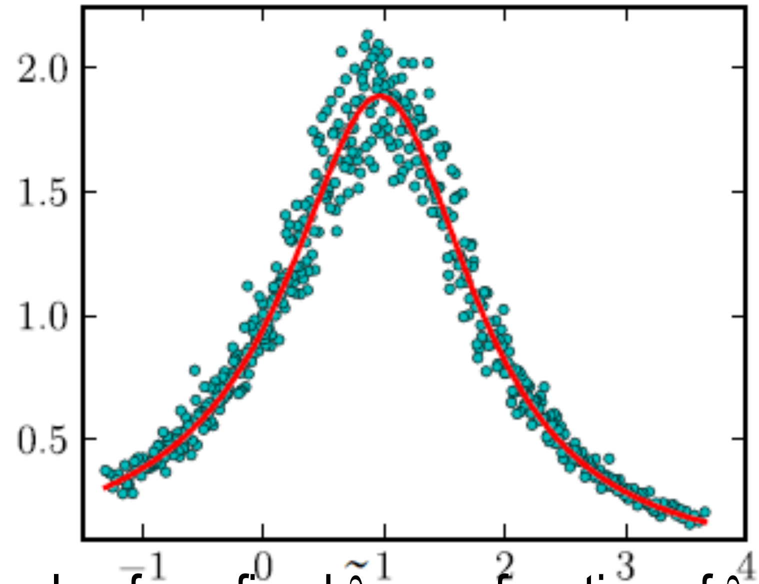

Overlap for a fixed $\tilde{\lambda}$ as a function of λ

- Again, the formula does not depend explicitly on the (possibly unknown) **C**
- It can be used to test whether **M** and **M'** originate from the same (unknown) **C**
- Again, universal within the whole class of Wigner/Wishart like matrices

Overlaps between independent realisations

- The case of financial covariance matrices: is the « true » underlying correlation structure stable in time?

(Different time periods + Bootstrap)

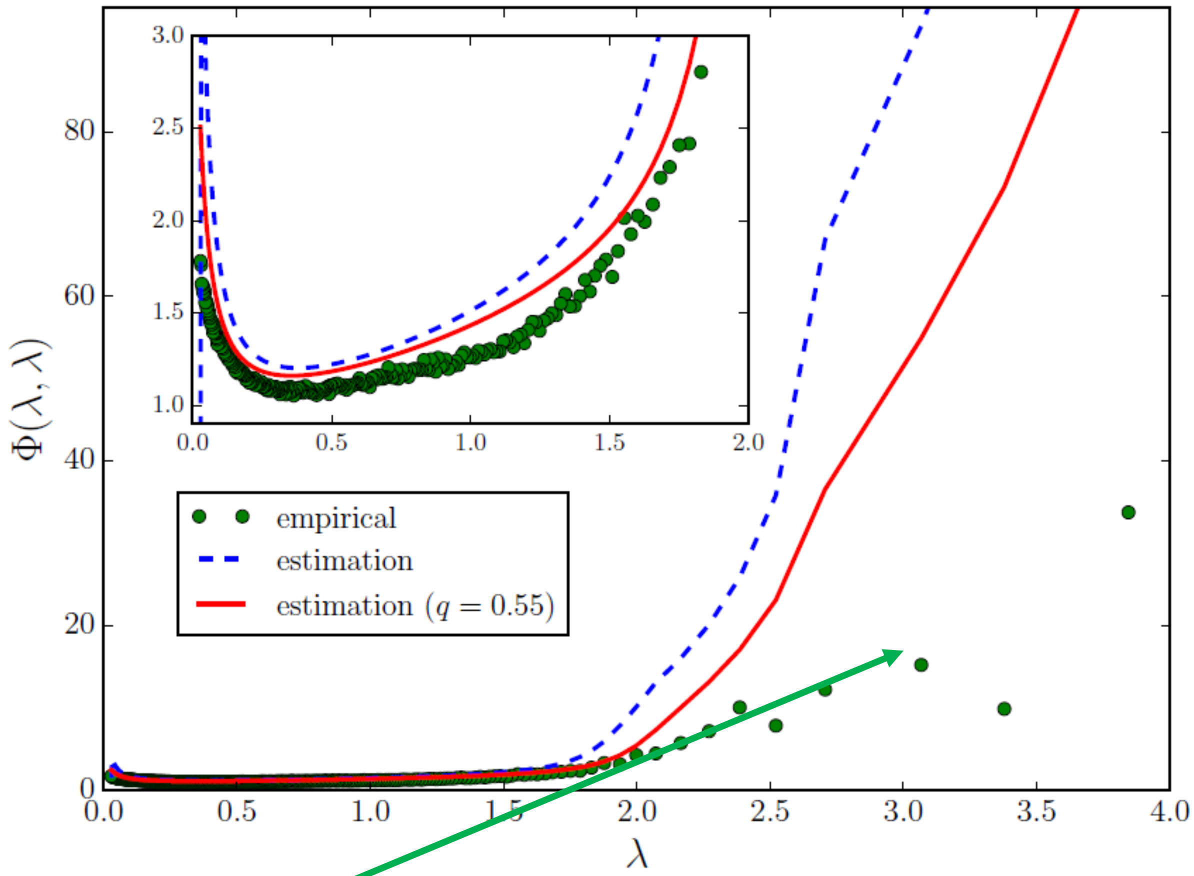

- Large eigenvectors are **unstable** (cf. R. Allez, JPB)
- Important for portfolio optimisation (uncontrolled risk exposure to large modes)
- « Eyeballing » test: can it be turned into a true statistical test?

Overlaps between independent realisations

- An ugly formula for $\Phi(\lambda, \tilde{\lambda}) := N\mathbb{E}\big[\langle \mathbf{u}_\lambda , \tilde{\mathbf{u}}_{\tilde{\lambda}} \rangle^2\big]$ but a simple interpretation
- From the previous overlaps of **M**'s with **C** one gets:

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

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

- « Ergodic hypothesis »: all $\varepsilon(\mu, \lambda)$ for different μ, λ are *independent*

$$\boxed{\Phi(\lambda, \lambda') = \int \mathrm{d}\mu \varrho_C(\mu) \Phi_0(\mu, \lambda)\, \Phi_0(\mu, \lambda')}$$

→ A simple « triangle » formula (that appears to depend on **C**)

From the convolution formula back to the Oracle estimator

- Consider $\nu_i(q) := \langle \mathbf{u}_i , \tilde{\mathbf{S}}\mathbf{u}_i \rangle$ where $\tilde{\mathbf{S}}$ is an *independent* realisation of the covariance matrix

- Then using the convolution formula, one can easily show that $\nu_i(q)$ coincides with $\widehat{\xi}_i$ In other words, $\tilde{\mathbf{S}}$ can be used as a proxy to **C** in the Oracle formula

- This cross-validation, or « out of sample » estimator simplifies considerably the numerical estimation of $\widehat{\xi}_i$

$$\widehat{\xi}_i = \sum_{j=1}^{N} \langle \mathbf{u}_i | \mathbf{v}_j \rangle^2 c_j$$

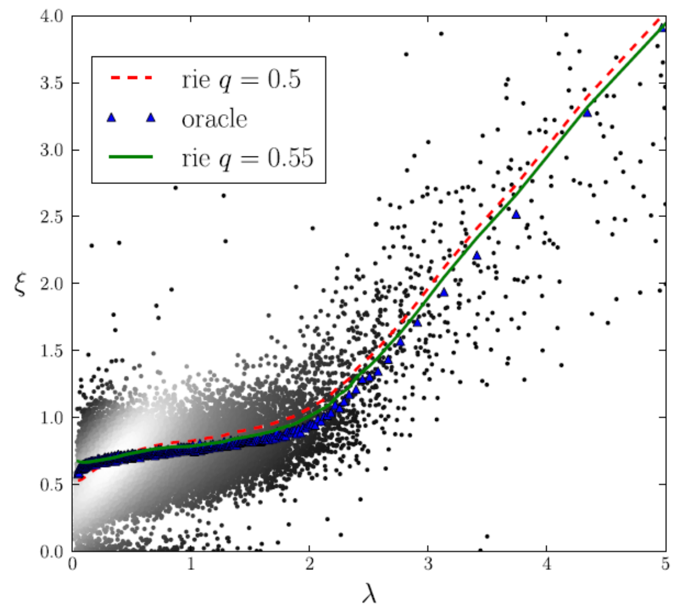

- Free Random Matrices results for Stieltjes transforms can be extended to the full resolvant matrix → access to overlaps
- Large dimension « miracles »:
  - The Oracle estimator can be estimated
  - The hypothesis that large matrices are generated from the same underlying matrix **C** can be tested without knowing **C**

Conclusions/Open problems

- True statistical test at large N ?
- RIE for cross-correlation SVDs (*en route* with F Benaych & M Potters)
- Overlaps for covariances matrices computed on overlapping periods?
- Dyson motion description for correlation matrices?
- Generalisation of Freeness, interpolating between commuting and free?
- Beyond RIE? Prior on eigenvectors?