Monitoring for a Phase Transition in a Time Series of Wigner Matrices

Nina Dörnemann [1]

joint work with Piotr Kokoszka [2], Tim Kutta [1], Sunmin Lee [2]

[1]Department of Mathematics, Aarhus University, Denmark
[2]Department of Statistics, Colorado State University, Fort Collins, USA

Wigner matrices

Wigner matrices

Definition

A Wigner matrix $\mathbf{W}_n \in \mathbb{R}^{n\times n}$ is a random matrix whose entries $(w_{ij})_{1\le i,j\le n}$ are independent, up to the symmetry constraint $w_{ij} = w_{ji}$, and satisfy the conditions

$$\mathbb{E}[w_{ij}] = 0, \quad \mathbb{E}[w_{ij}^2] = \sigma^2 + \delta_{ij} c, \quad \max_{1\le i,j\le n} \mathbb{E}|w_{ij}|^p \le c_p,$$

Wigner matrices

Definition

A Wigner matrix $\mathbf{W}_n \in \mathbb{R}^{n\times n}$ is a random matrix whose entries $(w_{ij})_{1\le i,j\le n}$ are independent, up to the symmetry constraint $w_{ij} = w_{ji}$, and satisfy the conditions

$$\mathbb{E}[w_{ij}] = 0, \quad \mathbb{E}[w_{ij}^2] = \sigma^2 + \delta_{ij} c, \quad \max_{1\le i,j\le n} \mathbb{E}|w_{ij}|^p \le c_p,$$

Example: Gaussian Orthogonal Ensemble

Let $\mathbf{G} \in \mathbb{R}^{n\times n}$ be a random matrix with i.i.d. $\mathcal{N}(0,1)$ entries. Then,

$$\mathbf{W}_n = \frac{1}{\sqrt{2}}(\mathbf{G} + \mathbf{G}^\top)$$

is a Wigner matrix with $\sigma^2 = 1, c = 1$.

Signal detection in deformed Wigner matrices

When is it possible to detect a rank-one signal in an additively deformed Wigner matrix $\mathbf{M}_n$?

$$\mathbf{M}_n = s\mathbf{x}\mathbf{x}^\top + \frac{1}{\sqrt{n}}\mathbf{W}_n$$

$(s > 0,\ \|\mathbf{x}\| = 1)$

Signal detection in deformed Wigner matrices

When is it possible to detect a rank-one signal in an additively deformed Wigner matrix $\mathbf{M}_n$?

$$\mathbf{M}_n = s\mathbf{x}\mathbf{x}^\top + \frac{1}{\sqrt{n}}\mathbf{W}_n$$

$(s > 0,\ \|\mathbf{x}\| = 1)$

Answers from random matrix theory (as $n \to \infty$):

	$s < \sigma$	$s > \sigma$
Eigenvalue	$\lambda_1(\mathbf{M}_n) \xrightarrow{P} 2\sigma$	$\lambda_1(\mathbf{M}_n) \xrightarrow{P} s + \sigma^2 s^{-1}$
Eigenvector	$\langle \mathbf{u}_1(\mathbf{M}_n), \mathbf{x}\rangle^2 \xrightarrow{P} 0$	$\langle \mathbf{u}_1(\mathbf{M}_n), \mathbf{x}\rangle^2 \xrightarrow{P} 1 - \sigma^2 s^{-2}$

(Capitaine et. al., 2009; Lee and Schnelli, 2015; Pizzo et al., 2013)

Signal detection in deformed Wigner matrices (Summary)

When is it possible to detect a rank-one signal in an additively deformed Wigner matrix $\mathbf{M}_n$?

$$\mathbf{M}_n = s\mathbf{x}\mathbf{x}^\top + \frac{1}{\sqrt{n}}\mathbf{W}_n$$

($s > 0$, $\|\mathbf{x}\| = 1$ possibly random)

- **subcritical**: no information about the signal
- **supercritical**: both the maximum eigenvalue and its corresponding eigenvector contain information about the signal

Sequential change-point detection

Sequential change point detection

There exist mulitple paradigms
"Chu, C. S. J., Stinchcombe, M. and White, H. (1996).
Monitoring structural change. Econometrica 64 1045-1065."

Idea: Let's say we want to monitor for a change in the mean.

- Training sample $X_1, X_2, \ldots, X_m$ with constant mean

Sequential change point detection

There exist mulitple paradigms
"Chu, C. S. J., Stinchcombe, M. and White, H. (1996).
Monitoring structural change. Econometrica 64 1045-1065."

Idea: Let's say we want to monitor for a change in the mean.

- Training sample $X_1, X_2, \ldots, X_m$ with constant mean
- Then, data arrive sequentially $X_{m+1}, X_{m+2}, \ldots$

Sequential change point detection

There exist mulitple paradigms
"Chu, C. S. J., Stinchcombe, M. and White, H. (1996).
Monitoring structural change. Econometrica 64 1045-1065."

Idea: Let's say we want to monitor for a change in the mean.

- Training sample $X_1, X_2, \dots, X_m$ with constant mean
- Then, data arrive sequentially $X_{m+1}, X_{m+2}, \dots$
- At each point k, use $X_1, \dots., X_{m+k}$ to decide between

$$H_0 : \mathbb{E}X_1 = \mathbb{E}X_2 = \dots$$
$$H_1 : \exists k^\star : \ \mathbb{E}X_1 = \mathbb{E}X_2 = \dots = \mathbb{E}X_{m+k^*} \neq \mathbb{E}X_{m+k^*+1} = \dots,$$

where $k^\star$ is a unknown change point.

Sequential change point detection - Illustration

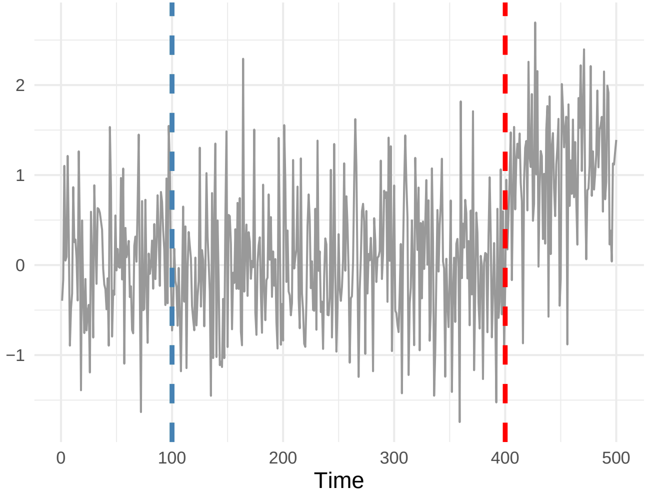

Sequential change point detection - Illustration

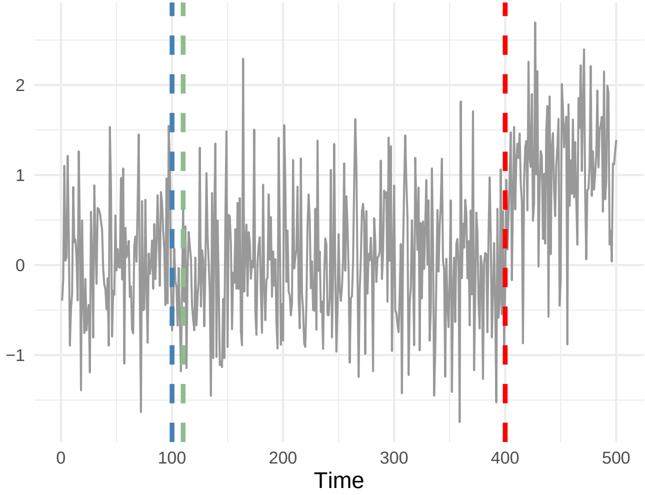

Sequential change point detection - Illustration

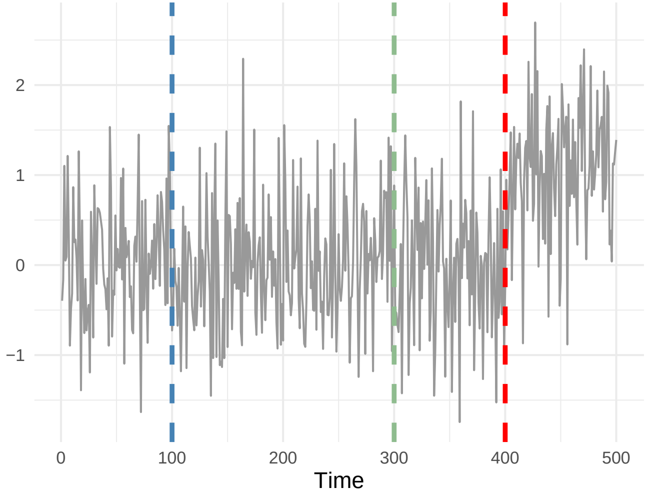

Sequential change point detection - Illustration

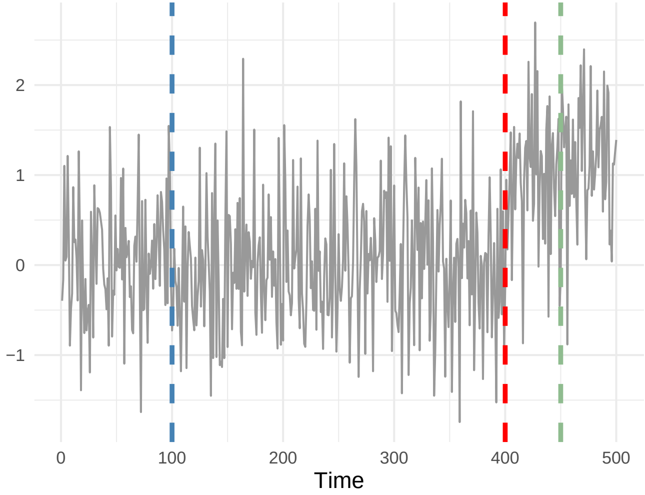

Statistical problem

Statistical Model for Regime Transition Detection

Model Setup:

- Data: $\mathbf{M}_t = s_t \mathbf{x}_{n,t} \mathbf{x}_{n,t}^\top + \frac{1}{\sqrt{n}} \mathbf{W}_{n,t}$
- $\mathbb{P}^{(1)}$: subcritical regime, initial state, support $[0, \sigma)$
- $\mathbb{P}^{(2)}$: supercritical regime, support (σ, ∞)
- $s_t \sim \mathbb{P}^{(1)}$ or $s_t \sim \mathbb{P}^{(2)}$

Statistical Model for Regime Transition Detection

Model Setup:

- Data: $\mathbf{M}_t = s_t \mathbf{x}_{n,t} \mathbf{x}_{n,t}^\top + \frac{1}{\sqrt{n}} \mathbf{W}_{n,t}$
- $\mathbb{P}^{(1)}$: subcritical regime, initial state, support $[0, \sigma)$
- $\mathbb{P}^{(2)}$: supercritical regime, support (σ, ∞)
- $s_t \sim \mathbb{P}^{(1)}$ or $s_t \sim \mathbb{P}^{(2)}$

Training Sample:

- $s_1, \ldots, s_m \sim \mathbb{P}^{(1)}$
- Signal in $\mathbf{M}_1, \ldots, \mathbf{M}_m$ cannot be detected

Statistical Model for Regime Transition Detection

Model Setup:

- Data: $\mathbf{M}_t = s_t \mathbf{x}_{n,t} \mathbf{x}_{n,t}^\top + \frac{1}{\sqrt{n}} \mathbf{W}_{n,t}$
- $\mathbb{P}^{(1)}$: subcritical regime, initial state, support $[0, \sigma)$
- $\mathbb{P}^{(2)}$: supercritical regime, support (σ, ∞)
- $s_t \sim \mathbb{P}^{(1)}$ or $s_t \sim \mathbb{P}^{(2)}$

Training Sample:

- $s_1, \dots, s_m \sim \mathbb{P}^{(1)}$
- Signal in $\mathbf{M}_1, \dots, \mathbf{M}_m$ cannot be detected

Monitoring Phase:

- At each time point k, use $\mathbf{M}_1, \dots, \mathbf{M}_{m+k}$ to decide whether

$$\mathrm{H}_0 : s_t \sim \mathbb{P}^{(1)} \ \forall t$$

$$\mathrm{H}_1 : s_1, \dots, s_{m+k^*} \sim \mathbb{P}^{(1)}, \ s_{m+k^*+1}, s_{m+k^*+2}, \dots \sim \mathbb{P}^{(2)},$$

where $k^\star$ is a unknown change point.

Methodological Criteria

Goal: A test for the emergence of a supercritical signal in a weakly dependent time series of high-dimensional matrices

Methodological Criteria

Goal: A test for the emergence of a supercritical signal in a weakly dependent time series of high-dimensional matrices

Some details:

- Asymptotics in $n, m \to \infty$ (dimension of the matrix and training sample size, respectively):

  $$m \asymp n^{\theta}$$

  for some constant $\theta \in (0, \infty)$
- For any n, the sequence $(\mathbf{W}_{n,t}, \mathbf{x}_{n,t}, s_t)_{t \in \mathbb{Z}}$ is strictly stationary.
- The triangular array $(\mathbf{W}_{n,t}, \mathbf{x}_{n,t}, s_t)_{t \in \mathbb{Z}}$ is ϕ-mixing.

Methodological Criteria

Goal: A test for the emergence of a supercritical signal in a weakly dependent time series of high-dimensional matrices

Some details:

- Asymptotics in $n, m \to \infty$ (dimension of the matrix and training sample size, respectively):

$$m \asymp n^{\theta}$$

  for some constant $\theta \in (0, \infty)$
- For any n, the sequence $(\mathbf{W}_{n,t}, \mathbf{x}_{n,t}, s_t)_{t \in \mathbb{Z}}$ is strictly stationary.
- The triangular array $(\mathbf{W}_{n,t}, \mathbf{x}_{n,t}, s_t)_{t \in \mathbb{Z}}$ is ϕ-mixing.

Two aims:

- Approximation of a nominal level $\alpha \in (0, 1)$ under H_0
- Power consistency under H_1

Test statistic

A first change-point detector I

- If λ_t is the largest eigenvalue of $\mathbf{M}_t$ under H_0, then

$$n^{2/3}\left(\lambda_t - 2\sigma\right) \to \begin{cases} TW_t & : \text{ if } s_t \sim \mathbb{P}^{(1)}, \\ \infty & : \text{ if } s_t \sim \mathbb{P}^{(2)}. \end{cases}$$

A first change-point detector I

- If λ_t is the largest eigenvalue of $\mathbf{M}_t$ under H_0, then

$$n^{2/3}\,(\lambda_t - 2\sigma) \to \begin{cases} TW_t & : \text{ if } s_t \sim \mathbb{P}^{(1)}, \\ \infty & : \text{ if } s_t \sim \mathbb{P}^{(2)}. \end{cases}$$

- CUSUM statistic (see, e.g., Horváth et al. (2003)):

$$\begin{aligned} \tilde{\Gamma}_{n,m}(k) &:= \frac{\sqrt{m}}{m+k}\left(\sum_{t=m+1}^{m+k} n^{2/3}(\lambda_t - 2\sigma) - \frac{k}{m}\sum_{t=1}^{m} n^{2/3}(\lambda_t - 2\sigma)\right) \\ &= \frac{\sqrt{m}}{m+k}\left(\sum_{t=m+1}^{m+k} n^{2/3}\lambda_t - \frac{k}{m}\sum_{t=1}^{m} n^{2/3}\lambda_t\right) \end{aligned}$$

A first change-point detector II

Can we use

$$\tilde{\Gamma}_{n,m}(k) = \frac{\sqrt{m}}{m+k}\left(\sum_{t=m+1}^{m+k} n^{2/3}\lambda_t - \frac{k}{m}\sum_{t=1}^{m} n^{2/3}\lambda_t\right)$$

for testing H_0 vs. H_1?

A first change-point detector II

Can we use

$$\tilde{\Gamma}_{n,m}(k) = \frac{\sqrt{m}}{m+k}\left(\sum_{t=m+1}^{m+k} n^{2/3}\lambda_t - \frac{k}{m}\sum_{t=1}^{m} n^{2/3}\lambda_t \right)$$

for testing H_0 vs. H_1? **Two problems:**

- **Practical:** Limiting distribution is still governed by the unknown time dependence of the sequence $(\mathbf{W}_t, \mathbf{x}_t, s_t)_t$
- **Theoretical:** Under H_0, we hope that

$$\tilde{\Gamma}_{n,m}(k) \approx \frac{\sqrt{m}}{m+k}\left(\sum_{t=m+1}^{m+k} TW_t - \frac{k}{m}\sum_{t=1}^{m} TW_t \right),$$

A first change-point detector II

Can we use

$$\tilde{\Gamma}_{n,m}(k) = \frac{\sqrt{m}}{m+k}\Bigg(\sum_{t=m+1}^{m+k} n^{2/3}\lambda_t - \frac{k}{m}\sum_{t=1}^{m} n^{2/3}\lambda_t\Bigg)$$

for testing H_0 vs. H_1? **Two problems:**

- **Practical:** Limiting distribution is still governed by the unknown time dependence of the sequence $(\mathbf{W}_t, \mathbf{x}_t, s_t)_t$
- **Theoretical:** Under H_0, we hope that

$$\tilde{\Gamma}_{n,m}(k) \approx \frac{\sqrt{m}}{m+k}\Bigg(\sum_{t=m+1}^{m+k} TW_t - \frac{k}{m}\sum_{t=1}^{m} TW_t\Bigg),$$

Limiting results for $\tilde{\Gamma}_{n,m}(k)$ are only know if $n \to \infty$ and m, k are fixed. It is not even known whether the moments of $n^{2/3}(\lambda_t - 2\sigma)$ converge to those of TW_t!

A first change-point detector II

Can we use

$$\tilde{\Gamma}_{n,m}(k) = \frac{\sqrt{m}}{m+k}\Bigg(\sum_{t=m+1}^{m+k} n^{2/3}\lambda_t - \frac{k}{m}\sum_{t=1}^{m} n^{2/3}\lambda_t\Bigg)$$

for testing H_0 vs. H_1? **Two problems:**

- **Practical:** Limiting distribution is still governed by the unknown time dependence of the sequence $(\mathbf{W}_t, \mathbf{x}_t, s_t)_t$
- **Theoretical:** Under H_0, we hope that

$$\tilde{\Gamma}_{n,m}(k) \approx \frac{\sqrt{m}}{m+k}\Bigg(\sum_{t=m+1}^{m+k} TW_t - \frac{k}{m}\sum_{t=1}^{m} TW_t\Bigg),$$

Limiting results for $\tilde{\Gamma}_{n,m}(k)$ are only know if $n \to \infty$ and m, k are fixed. It is not even known whether the moments of $n^{2/3}(\lambda_t - 2\sigma)$ converge to those of TW_t!

A self-normalized detector

- **Idea**: Division of $\tilde{\Gamma}_{n,m}(k)$ by a **"normalizer"** V_m that cancels out the long-run variance and hence makes the normalized statistic asymptotically pivotal.

A self-normalized detector

- **Idea**: Division of $\tilde{\Gamma}_{n,m}(k)$ by a **"normalizer"** V_m that cancels out the long-run variance and hence makes the normalized statistic asymptotically pivotal.
- Self-normalization $\neq$ estimation of long-run variance !

A self-normalized detector

- **Idea**: Division of $\tilde{\Gamma}_{n,m}(k)$ by a **"normalizer"** V_m that cancels out the long-run variance and hence makes the normalized statistic asymptotically pivotal.
- Self-normalization $\neq$ estimation of long-run variance !
- Normalizer drawing on the training sample:

$$V_m := \frac{n^{2/3}}{m^{3/2}} \sum_{t=1}^{m} \left| \sum_{s=1}^{t} \lambda_s - \frac{t}{m} \sum_{s=1}^{m} \lambda_s \right|,$$

A self-normalized detector

- **Idea**: Division of $\tilde{\Gamma}_{n,m}(k)$ by a **"normalizer"** V_m that cancels out the long-run variance and hence makes the normalized statistic asymptotically pivotal.
- Self-normalization $\neq$ estimation of long-run variance !
- Normalizer drawing on the training sample:

$$V_m := \frac{n^{2/3}}{m^{3/2}} \sum_{t=1}^{m} \Big| \sum_{s=1}^{t} \lambda_s - \frac{t}{m} \sum_{s=1}^{m} \lambda_s \Big|,$$

- Self-normalized detector:

$$\Gamma_{n,m}(k) = \frac{\tilde{\Gamma}_{n,m}(k)}{V_m}$$

Stylized results

Weak convergence under H_0

Under H_0 and adequate technical assumptions, as $m, n \to \infty$

$$\sup_{1 \le k < \infty} \Gamma_{n,m}(k) \xrightarrow{d} \sup_{0 \le x < \infty} \frac{[B(1+x) - B(1)] - xB(1)}{(1+x) \cdot \int_0^1 |B(s) - sB(1)| ds},$$

where B is the standard Brownian motion.

Methodology

Denoting by $q_{1-\alpha}$ the upper α-quantile of the limiting distribution, we reject the hypothesis H_0, if for some $k \geq 1$,

$$\Gamma_{n,m}(k) > q_{1-\alpha}.$$

Methodology

Denoting by $q_{1-\alpha}$ the upper α-quantile of the limiting distribution, we reject the hypothesis H_0, if for some $k \geq 1$,

$$\Gamma_{n,m}(k) > q_{1-\alpha}.$$

Theoretical guarantees

Under H_0,

$$\lim_{m\to\infty} \mathbb{P}\left(\sup_{1\leq k<\infty} \Gamma_{n,m}(k) > q_{1-\alpha}\right) = \alpha.$$

Under H_1 and $k^\star \asymp m^D$ for some $D > 0$, it holds that

$$\lim_{m\to\infty} \mathbb{P}\left(\sup_{1\leq k<\infty} \Gamma_{n,m}(k) > q_{1-\alpha}\right) = 1.$$

Some technical details

Main technical results under H_0

To show the convergence

$$\sup_{1 \le k < \infty} \Gamma_{n,m}(k) \xrightarrow{d} \sup_{0 \le x < \infty} \frac{[B(1+x) - B(1)] - xB(1)}{(1+x) \cdot \int_0^1 |B(s) - sB(1)| ds}$$

under H_0, we need two technical results:

Main technical results under H_0

To show the convergence

$$\sup_{1 \le k < \infty} \Gamma_{n,m}(k) \xrightarrow{d} \sup_{0 \le x < \infty} \frac{[B(1+x) - B(1)] - xB(1)}{(1+x) \cdot \int_0^1 |B(s) - sB(1)| ds}$$

under H_0, we need two technical results:

1. **Gaussian approximation** for partial sum process $P_m(x),\ x \in [0, T_m]$, of eigenvalues on a growing interval $[0, T_m]$.

Main technical results under H_0

To show the convergence

$$\sup_{1 \le k < \infty} \Gamma_{n,m}(k) \stackrel{d}{\to} \sup_{0 \le x < \infty} \frac{[B(1+x) - B(1)] - xB(1)}{(1+x) \cdot \int_0^1 |B(s) - sB(1)| ds}$$

under H_0, we need two technical results:

1. **Gaussian approximation** for partial sum process $P_m(x),\ x \in [0, T_m]$, of eigenvalues on a growing interval $[0, T_m]$.
   $\Rightarrow$ Approximate the denominator V_m in the self-normalized statistic by a transformation of a Brownian motion and similarly for the numerator $\tilde{\Gamma}_{n,m}(k)$ for values of $k \le mT_m$.

Main technical results under H_0

To show the convergence

$$\sup_{1 \leq k < \infty} \Gamma_{n,m}(k) \xrightarrow{d} \sup_{0 \leq x < \infty} \frac{[B(1+x) - B(1)] - xB(1)}{(1+x) \cdot \int_0^1 |B(s) - sB(1)| ds}$$

under H_0, we need two technical results:

1. **Gaussian approximation** for partial sum process $P_m(x),\ x \in [0, T_m]$, of eigenvalues on a growing interval $[0, T_m]$.
   $\Rightarrow$ Approximate the denominator V_m in the self-normalized statistic by a transformation of a Brownian motion and similarly for the numerator $\tilde{\Gamma}_{n,m}(k)$ for values of $k \leq mT_m$.
2. For large k, we need appropriate **tail bounds**.

Gaussian approximation of $P_m(x)$ for $x \in [0, T_m]$ (1/2)

As discussed previously, we cannot formulate a Gaussian approximation directly for $\tilde{\Gamma}_{n,m}(x)$

$$\tilde{\Gamma}_{n,m}(k) = \frac{\sqrt{m}}{m+k}\left(\sum_{t=m+1}^{m+k} n^{2/3}\lambda_t - \frac{k}{m}\sum_{t=1}^{m} n^{2/3}\lambda_t\right).$$

Gaussian approximation of $P_m(x)$ for $x \in [0, T_m]$ (1/2)

As discussed previously, we cannot formulate a Gaussian approximation directly for $\tilde{\Gamma}_{n,m}(x)$

$$\tilde{\Gamma}_{n,m}(k) = \frac{\sqrt{m}}{m+k}\Bigg(\sum_{t=m+1}^{m+k} n^{2/3}\lambda_t - \frac{k}{m}\sum_{t=1}^{m} n^{2/3}\lambda_t\Bigg).$$

Idea: Make use of our self-normalized approach and define a process $P_m(x)$ that we like to work with!

Gaussian approximation of $P_m(x)$ for $x \in [0, T_m]$ (2/2)

$$P_m(x) := \frac{1}{\sqrt{m\tau_n}} \sum_{t=1}^{\lfloor mx \rfloor} n^{2/3}[\lambda_t - b^{(n)}], \qquad x \in [0, T_m]$$

Gaussian approximation of $P_m(x)$ for $x \in [0, T_m]$ (2/2)

$$P_m(x) := \frac{1}{\sqrt{m\tau_n}} \sum_{t=1}^{\lfloor mx \rfloor} n^{2/3}[\lambda_t - b^{(n)}], \qquad x \in [0, T_m]$$

- **Truncation of eigenvalues**: For $\varepsilon > 0$, define the event

$$\Lambda_t := \{|\lambda_t - 2\sigma| < n^{\varepsilon - 2/3}\}.$$

  Λ_t occurs with high probability for any t in the subcritical regime!

Gaussian approximation of $P_m(x)$ for $x \in [0, T_m]$ (2/2)

$$P_m(x) := \frac{1}{\sqrt{m\tau_n}} \sum_{t=1}^{\lfloor mx \rfloor} n^{2/3}[\lambda_t - b^{(n)}], \qquad x \in [0, T_m]$$

- **Truncation of eigenvalues**: For $\varepsilon > 0$, define the event

$$\Lambda_t := \{|\lambda_t - 2\sigma| < n^{\varepsilon - 2/3}\}.$$

  Λ_t occurs with high probability for any t in the subcritical regime!

- **Centering:** $b^{(n)}$ is the conditional mean

$$b^{(n)} := \mathbb{E}[\lambda_1 | \Lambda_1].$$

Gaussian approximation of $P_m(x)$ for $x \in [0, T_m]$ (2/2)

$$P_m(x) := \frac{1}{\sqrt{m\tau_n}} \sum_{t=1}^{\lfloor mx \rfloor} n^{2/3}[\lambda_t - b^{(n)}], \qquad x \in [0, T_m]$$

- **Truncation of eigenvalues**: For $\varepsilon > 0$, define the event

$$\Lambda_t := \{|\lambda_t - 2\sigma| < n^{\varepsilon - 2/3}\}.$$

  Λ_t occurs with high probability for any t in the subcritical regime!
- **Centering:** $b^{(n)}$ is the conditional mean

$$b^{(n)} := \mathbb{E}[\lambda_1 | \Lambda_1].$$

- **Standardization:** Letting $Y_t := \mathbb{I}\{\Lambda_t\}(\lambda_t - b^{(n)})$, we define τ_n as the finite sample variance

$$\tau_n := \mathbb{E}\left(\frac{1}{\sqrt{\lfloor mT_m \rfloor}} \sum_{t=1}^{\lfloor mT_m \rfloor} Y_t\right)^2.$$

Tail bound

We show a tail bound of the type:

$$\sup_{m^{1+\zeta}<k<\infty} \Gamma_{n,m}(k) = \frac{-P_m(1)}{V_m/\sqrt{\tau_n}} + o_P(1).$$

Tail bound

We show a tail bound of the type:

$$\sup_{m^{1+\zeta}<k<\infty} \Gamma_{n,m}(k) = \frac{-P_m(1)}{V_m/\sqrt{\tau_n}} + o_P(1).$$

Idea for the proof:

- **Small $k \leq mT_m$:** Gaussian approximation applies.

Tail bound

We show a tail bound of the type:

$$\sup_{m^{1+\zeta}<k<\infty} \Gamma_{n,m}(k) = \frac{-P_m(1)}{V_m/\sqrt{\tau_n}} + o_P(1).$$

Idea for the proof:

- **Small** $k \leq mT_m$**:** Gaussian approximation applies.
- **Intermediate** $mT_m < k \leq m^D$**:** Use bounds for sums of bounded, mixing random variables (even though the bound depends on n).

Tail bound

We show a tail bound of the type:

$$\sup_{m^{1+\zeta}<k<\infty} \Gamma_{n,m}(k) = \frac{-P_m(1)}{V_m/\sqrt{\tau_n}} + o_P(1).$$

Idea for the proof:

- **Small** $k \leq mT_m$**:** Gaussian approximation applies.
- **Intermediate** $mT_m < k \leq m^D$**:** Use bounds for sums of bounded, mixing random variables (even though the bound depends on n).
- **Large** $k > m^D$**:** Temporal discounting $(1/(m+k))$ suppresses fluctuations.

Conclusion

Take-home messages

- **Phase transition detection** in high-dimensional random matrices is possible in **real time**.

Take-home messages

- **Phase transition detection** in high-dimensional random matrices is possible in **real time**.
- **Self-normalized detector** based on extremal eigenvalues: no tuning parameters needed.

Take-home messages

- **Phase transition detection** in high-dimensional random matrices is possible in **real time**.
- **Self-normalized detector** based on extremal eigenvalues: no tuning parameters needed.
- **Theory** based on random matrix theory and Gaussian approximation.

Take-home messages

- **Phase transition detection** in high-dimensional random matrices is possible in **real time**.
- **Self-normalized detector** based on extremal eigenvalues: no tuning parameters needed.
- **Theory** based on random matrix theory and Gaussian approximation.
- **Applications**: pollution monitoring, primate social interactions

Thank you for your attention!

Feel free to check out our preprint:
Dörnemann, N., Kokoszka P., Kutta T. and Lee, S. (2025). Monitoring for a Phase Transition in a Time Series of Wigner Matrices. arxiv:2507.04983.