



Time Series Analysis and Forecasting

Seminar 12: Spectral Analysis and Frequency-Domain Methods



Daniel Traian PELE

Bucharest University of Economic Studies

IDA Institute Digital Assets

Blockchain Research Center

AI4EFin Artificial Intelligence for Energy Finance

Romanian Academy, Institute for Economic Forecasting

MSCA Digital Finance

Seminar Outline

- ▣ **Part I: Year 3 Level** – Log-returns, periodogram, inconsistency, AR(1) spectrum
- ▣ **Part II: Master Level** – Welch vs. multitaper, AR estimation, HP filter
- ▣ **Part III: PhD Extension** – GPH estimator, Local Whittle, cross-spectral analysis
- ▣ **Discussion Topics** – Choosing estimators, interpreting spectra, HP filter debate
- ▣ **AI-Assisted Exercise** – Critical thinking with AI tools
- ▣ **Take-Home Exercises** – Extended practice
- ▣ **Summary** – Key concepts by level
- ▣ **Bibliography** – References and code

Exercise 1: Log>Returns

Question

Given a price series P_t , define log-returns $r_t = \log(P_t/P_{t-1})$.

Exercise

- **(A)** Plot the price series P_t and the return series r_t . What visual features do you notice?
- **(B)** Compute the sample ACF of r_t and of r_t^2 for lags $h = 1, \dots, 20$.
- **(C)** What does a near-zero ACF of r_t suggest about predictability? What do the significant ACF values of r_t^2 indicate?

Answer on next slide...

Exercise 1: Solution

Solution

- ▣ **(A)** P_t typically shows a non-stationary trending pattern; r_t fluctuates around zero with periods of high and low volatility (volatility clustering).
- ▣ **(B)** Sample ACF of r_t : most values fall within $\pm 1.96/\sqrt{n}$ confidence bands $\Rightarrow$ returns are approximately uncorrelated. Sample ACF of r_t^2 : significant positive values at many lags $\Rightarrow$ squared returns are correlated.
- ▣ **(C)** Near-zero $\text{ACF}(r_t) \Rightarrow$ returns are close to **white noise**: past returns carry little information to forecast future returns (weak-form market efficiency). Significant $\text{ACF}(r_t^2) \Rightarrow$ **ARCH/GARCH effects**: volatility is predictable even when returns are not.

Exercise 2: Raw Periodogram

Question

Given $n = 256$ observations, the periodogram at Fourier frequency $\omega_j = 2\pi j/n$ is:

$$I(\omega_j) = \frac{1}{2\pi n} |d(\omega_j)|^2, \quad d(\omega_j) = \sum_{t=1}^n Y_t e^{-i\omega_j t}$$

Exercise

- ▣ **(A)** Plot $I(\omega_j)$ against $\omega_j \in [0, \pi]$. How do you identify the dominant frequency?
- ▣ **(B)** If the series has a *quarterly* pattern (period = 4), at which angular frequency ω^* do you expect the largest periodogram peak?
- ▣ **(C)** What does a completely flat periodogram indicate?

Answer on next slide...

Exercise 2: Solution

Solution

- ▣ **(A)** The dominant frequency is the ω_j at which $I(\omega_j)$ reaches its global maximum. The corresponding period is $T = 2\pi/\omega_j$ time units.
- ▣ **(B)** A quarterly pattern has period $T = 4$, so $\omega^* = 2\pi/4 = \pi/2 \approx 1.571$ rad. This is the 64th Fourier frequency for $n = 256$ (since $j^* = n/T = 64$).
- ▣ **(C)** A flat periodogram $I(\omega_j) \approx c$ for all j indicates **white noise**: all frequencies contribute equally to the series variance ($\int_0^\pi f(\omega) d\omega = \gamma(0)$, where $f(\omega) = \sigma^2/(2\pi) = \text{const}$).

 TSA_ch12_periodogram

Exercise 3: Periodogram Inconsistency

Question

Show that the periodogram is an *inconsistent* estimator of the spectral density.

Exercise

- ▣ **(A)** State the asymptotic distribution of $I(\omega_j)$ as $n \rightarrow \infty$.
- ▣ **(B)** Derive $\text{Var}[I(\omega_j)]$ as a function of n . Does it go to zero?
- ▣ **(C)** Propose a solution that produces a consistent estimator, and explain the bias-variance tradeoff involved.

Answer on next slide...

Exercise 3: Solution

Solution

- ▣ **(A)** For $\omega_j \neq 0, \pi$: $\frac{2 I(\omega_j)}{f(\omega_j)} \xrightarrow{d} \chi_2^2 \Rightarrow I(\omega_j) \stackrel{d}{\approx} f(\omega_j) \cdot \chi_2^2/2$
- ▣ **(B)** Since χ_2^2 has variance 4: $\text{Var}[I(\omega_j)] \approx f(\omega_j)^2 \cdot 4/4 = f(\omega_j)^2$. This does **not** decrease with n $\Rightarrow$ periodogram is **inconsistent**.
- ▣ **(C) Smoothing** over m adjacent ordinates: $\hat{f}(\omega_j) = \frac{1}{2m+1} \sum_{k=-m}^m I(\omega_{j+k})$. Variance $\propto 1/m$, bandwidth $\propto m/n$. More smoothing $\Rightarrow$ lower variance but higher bias.

 TSA_ch12_periodogram

Exercise 4: AR(1) Spectrum

Question

Consider the AR(1) process $Y_t = 0.8 Y_{t-1} + \varepsilon_t$ with $\sigma_\varepsilon^2 = 1$.

Exercise

- (A) Derive the spectral density $f(\omega)$ for $\omega \in [0, \pi]$.
- (B) Compute the ratio $f(0)/f(\pi)$. Interpret the result.
- (C) What is a *red spectrum*? Give an economic example.

Answer on next slide...

Exercise 4: Solution

Solution

- ▣ (A) For AR(1) with $\phi = 0.8$: $f(\omega) = \frac{1}{2\pi} \cdot \frac{1}{1.64 - 1.6 \cos \omega}$
- ▣ (B) $f(0) = \frac{25}{2\pi}$, $f(\pi) = \frac{1}{2\pi \cdot 3.24}$, so $\frac{f(0)}{f(\pi)} = \left(\frac{1+\phi}{1-\phi}\right)^2 = \left(\frac{1.8}{0.2}\right)^2 = 81$. Most variance is concentrated at low frequencies.
- ▣ (C) A **red spectrum** is dominated by low frequencies: most variance is at slow oscillations (long cycles). GDP, consumption, and stock prices all exhibit red spectra — they move slowly and trend persistently, with little energy at high frequencies.

 TSA_ch12_spectral_density

Exercise 5: Welch vs. Multitaper Estimation

Question

Compare two spectral estimators on the same series of $n = 512$ observations: **Welch**: $L = 4$ overlapping segments; **Multitaper**: $NW = 4$, $K = 7$ DPSS tapers.

Exercise

- ▣ (A) Compute the frequency resolution (bandwidth) for each estimator.
- ▣ (B) Which estimator has lower variance? Why?
- ▣ (C) Describe the bias-variance tradeoff for each method.

Answer on next slide...

Exercise 5: Solution

Solution

- ▣ **(A)** Welch ($L = 4$, $n_{\text{seg}} = 128$): $\Delta\omega_{\text{Welch}} \approx 2/128 \approx 0.0156$ cyc/obs. Multitaper ($NW = 4$, $n = 512$): $\Delta\omega_{\text{MT}} \approx 2NW/n = 8/512 = 0.0156$ cyc/obs. Both have the same frequency resolution for these parameters.
- ▣ **(B)** Multitaper averages $K = 7$ independent eigen-spectra $\Rightarrow$ variance $\propto 1/K$. For equal bandwidth, multitaper has lower variance: equivalent degrees of freedom $\approx 2K = 14$ vs $\approx 2L = 8$.
- ▣ **(C)** Both methods trade frequency resolution for variance reduction. Multitaper achieves near-optimal sidelobe suppression via DPSS tapers, minimising spectral leakage — preferred for most applications.

Exercise 6: AR Spectral Estimation

Question

Consider the AR(2) process:

$$Y_t = 1.5 Y_{t-1} - 0.7 Y_{t-2} + \varepsilon_t, \quad \sigma_\varepsilon^2 = 1$$

Exercise

- ▣ **(A)** Derive the parametric spectral density $\hat{f}_{AR}(\omega)$.
- ▣ **(B)** Find the peak frequency ω^* of the spectrum analytically.
- ▣ **(C)** Convert ω^* to a cycle period and interpret economically.

Answer on next slide...

Exercise 6: Solution

Solution

- ▣ **(A)** With $\phi_1 = 1.5$, $\phi_2 = -0.7$: $f(\omega) = \frac{1}{2\pi |1 - 1.5e^{-i\omega} + 0.7e^{-2i\omega}|^2}$
- ▣ **(B)** Differentiating $|A(e^{-i\omega})|^2$ numerically: $\omega^* \approx 0.47$ rad. (The formula $\arccos(\phi_1/2|\phi_2|)$ yields no real solution here, so numerical optimisation is used.)
- ▣ **(C)** Period $= 2\pi/\omega^* \approx 2\pi/0.47 \approx 13.4$ observations. For quarterly data, this corresponds to a **3.3-year business cycle** — consistent with typical inventory and investment cycles.

 TSA_ch12_spectral_density

Exercise 7: HP Filter Transfer Function

Question

The HP filter extracts trend τ_t by minimising $\sum(Y_t - \tau_t)^2 + \lambda \sum(\Delta^2 \tau_t)^2$. The squared gain of the *cycle* component is: $|A_{HP}(\omega)|^2 = \frac{\lambda [2(1 - \cos \omega)]^2}{1 + \lambda [2(1 - \cos \omega)]^2}$

Exercise

- **(A)** Evaluate $|A_{HP}(\omega)|^2$ at $\omega = 0$ and $\omega = \pi$ for any λ .
- **(B)** Find the half-power cutoff frequency for $\lambda = 1600$.
- **(C)** Summarise Hamilton (2018)'s critique of the HP filter.

Answer on next slide...

Exercise 7: Solution

Solution

- ▣ **(A)** At $\omega = 0$: $2(1 - \cos 0) = 0 \Rightarrow |A_{HP}(0)|^2 = 0$ (trend removed). At $\omega = \pi$: $|A_{HP}(\pi)|^2 = 16\lambda/(1 + 16\lambda) \approx 1$ (high freq. passed).
- ▣ **(B)** Setting $|A_{HP}(\omega)|^2 = 0.5$: $\lambda[2(1 - \cos \omega)]^2 = 1$, so $2(1 - \cos \omega) = 1600^{-1/2} = 0.025$, giving $\omega^* \approx 0.158$ rad $\approx 2\pi/40$ (period ≈ 40 quarters = 10 years).
- ▣ **(C) Hamilton (2018)**: HP filtering is equivalent to a regression on lagged values. The resulting cycle is spuriously correlated at all lags and imposes ARMA(4,4) structure on white noise — generating *artificial* business cycles. Hamilton recommends a simple 2-step regression alternative.

Exercise 8: GPH Estimator

Question

For ARFIMA(0, d ,0), the GPH estimator fits OLS to:

$$\log I(\omega_j) = c - d \log[4 \sin^2(\omega_j/2)] + u_j, \quad j = 1, \dots, m$$

Exercise

- (A) What is the OLS estimator $\hat{d}$? How is it computed?
- (B) What bandwidth m should be used? Theoretical recommendations?
- (C) What is the asymptotic standard error of $\hat{d}$?

Answer on next slide...

Exercise 8: Solution

Solution

- ▣ **(A)** Let $x_j = -\log[4 \sin^2(\omega_j/2)]$ and $y_j = \log I(\omega_j)$: $\hat{d} = \frac{\sum_{j=1}^m (x_j - \bar{x})(y_j - \bar{y})}{\sum_{j=1}^m (x_j - \bar{x})^2}$ (OLS slope)
- ▣ **(B)** Theory: $m = O(n^\alpha)$, $\alpha \in (0, 1)$. Common practice: $m \approx n^{0.5}$ to $n^{0.6}$. Too small $m \Rightarrow$ high variance; too large $m \Rightarrow$ bias from short-memory contamination.
- ▣ **(C)** $SE(\hat{d}) = \pi/\sqrt{6m}$. For $n = 1000$, $m = 32$: $SE \approx 0.226$. Hurst: $H = d + 0.5$; persistent series have $d > 0$ ($H > 0.5$).

Exercise 9: Local Whittle Estimator

Question

Compare the GPH and Local Whittle estimators for a series of $n = 1000$ observations with true $d = 0.3$.

Exercise

- **(A)** Write the Local Whittle objective function and state what it minimises.
- **(B)** What bandwidth m is used, and what is the asymptotic variance of $\hat{d}_{LW}$?
- **(C)** Why is Local Whittle more efficient than GPH?

Answer on next slide...

Exercise 9: Solution

Solution

- ▣ **(A)** Local Whittle minimises the restricted Whittle log-likelihood:

$$Q(d) = \log G(d) - 2d \bar{\lambda}_m, \quad G(d) = \frac{1}{m} \sum_{j=1}^m \frac{I(\omega_j)}{\omega_j^{-2d}}, \quad \hat{d}_{LW} = \arg \min_d Q(d)$$

- ▣ **(B)** Bandwidth $m = O(n^{0.6-0.8})$; for $n = 1000$ typically $m \approx 100-250$. Asymptotic variance: $\text{Var}(\hat{d}_{LW}) = 1/(4m)$. For $m = 100$: SD = 0.05.
- ▣ **(C)** GPH uses OLS on log-periodogram (heavy-tailed χ_2^2 errors); LW uses a proper likelihood approximation. LW is asymptotically efficient: $\text{Var} = 1/(4m)$ vs GPH's $\approx 1.645/m$ for the same m .

Exercise 10: Cross-Spectral Analysis

Question

Suppose $Y_t = X_{t-4} + \varepsilon_t$ where $\varepsilon_t \perp X_s$ for all s, t .

Exercise

- ▣ **(A)** Derive the cross-spectrum $S_{XY}(\omega) = e^{-4i\omega} S_{XX}(\omega)$. What is the phase function $\phi(\omega)$?
- ▣ **(B)** Compute the squared coherence $C_{XY}^2(\omega)$: (i) without noise, (ii) with noise $\varepsilon_t \sim (0, \sigma_\varepsilon^2)$.
- ▣ **(C)** What does the phase slope $-d\phi/d\omega$ tell you, and what is its value here?

Answer on next slide...

Exercise 10: Solution

Solution

- ▣ **(A)** Transfer function $H(\omega) = e^{-4i\omega}$, so $S_{XY}(\omega) = e^{-4i\omega} S_{XX}(\omega)$. The phase is $\phi(\omega) = -4\omega$ (linear in ω).
- ▣ **(B)** $C_{XY}^2(\omega) = |S_{XY}|^2 / (S_{XX} \cdot S_{YY})$: (i) No noise: $S_{YY} = S_{XX} \Rightarrow C_{XY}^2(\omega) = 1$ (perfect linear relationship). (ii) With noise: $C_{XY}^2(\omega) = S_{XX}(\omega) / [S_{XX}(\omega) + \sigma_\varepsilon^2 / (2\pi)] < 1$.
- ▣ **(C)** The **group delay** $= -d\phi/d\omega = 4$ periods at ALL frequencies. A constant group delay confirms the system introduces a **pure shift of 4 periods** with no phase distortion.

Discussion: Choosing the Right Spectral Estimator

Key Question

You have $n = 1000$ observations. Which spectral estimator do you use and why?

Decision Framework

1. **Need a quick look at dominant frequencies?**
 - ▶ Raw periodogram — fast but noisy
2. **Need low variance with good sidelobe suppression?**
 - ▶ Multitaper (DPSS) — the modern standard
3. **Short series or nonstationary segments?**
 - ▶ Welch (overlapping segments) — reduces sensitivity to nonstationarity
4. **Series known to follow an AR/ARMA model?**
 - ▶ Parametric AR spectrum — smooth and model-consistent

Discussion: Interpreting Economic Spectra

Key Question

You compute the spectral density of quarterly GDP growth and find most power at frequencies below $\pi/16$. What does this imply?

Interpretation Guide

- Frequency $\omega < \pi/16 \Rightarrow$ period > 32 quarters = 8 years: **long cycles**
- The spectrum is **red**: persistent, slowly-moving component dominates
- Business cycle band: $\omega \in [\pi/16, \pi/3]$ (periods 6–32 quarters)

Practical Rules

- Trend**: power near $\omega \approx 0$ (very low frequency)
- Business cycle**: power at 6–32 quarter periods (Baxter-King band)
- Seasonal**: sharp peaks at $\omega = \pi/2$ (quarterly), $2\pi/12$ (monthly)
- White noise**: flat spectrum across all frequencies

Discussion: HP Filter in Practice

Key Question

Should economists stop using the HP filter? What are the alternatives?

Arguments Against (Hamilton 2018)

- ▣ Imposes spurious ARMA(4,4) structure on white noise
- ▣ Generates artificial cycles not present in the data
- ▣ End-of-sample estimates are unreliable
- ▣ Results are not invariant to the choice of λ

Alternatives

- ▣ **Hamilton filter**: regress Y_{t+h} on $Y_t, Y_{t-1}, Y_{t-2}, Y_{t-3}$; use residual as cycle
- ▣ **Bandpass filter** (Baxter-King, Christiano-Fitzgerald): isolate business-cycle frequencies directly
- ▣ **Unobserved components model**: explicitly model trend + cycle with state-space methods

AI Exercise: Critical Thinking

Prompt to test in ChatGPT / Claude / Copilot

“Estimate the spectral density of quarterly US real GDP growth (1947–present) using the periodogram, Welch method, and multitaper. Identify the dominant business-cycle frequency. Test for long memory using the GPH estimator and compare with Local Whittle.”

Exercise:

1. Did the AI use the correct spectral normalisation ($1/(2\pi)$ convention)?
2. Are the Welch segments and multitaper bandwidth parameters justified?
3. Does the identified business-cycle peak correspond to the expected 6–32 quarter range?
4. Is the GPH bandwidth m chosen optimally (e.g. $m \approx n^{0.5}$)?
5. Do the GPH and Local Whittle estimates agree? If not, why might they differ?

Warning: AI-generated code may run without errors and look professional. *That does not mean it is correct.*

 TSA_ch12_long_memory

Take-Home Exercises

1. **Theoretical:** Show that for white noise $\varepsilon_t \sim \text{WN}(0, \sigma^2)$, the spectral density is $f(\omega) = \sigma^2/(2\pi)$ for all $\omega \in [0, \pi]$.
2. **Computation:** For the AR(2) process $Y_t = 0.6Y_{t-1} - 0.3Y_{t-2} + \varepsilon_t$:
 - ▶ Derive the spectral density $f(\omega)$ analytically
 - ▶ Simulate $n = 512$ observations and compare the periodogram with the theoretical spectrum
 - ▶ Apply Welch and multitaper estimation; compare all three
3. **Applied:** Download monthly sunspot numbers (available at sidc.be):
 - ▶ Estimate the spectral density and identify the ≈ 11 -year cycle
 - ▶ Test for long memory; report $\hat{d}$ via GPH and Local Whittle
4. **Critical Thinking:** Why does applying the HP filter with $\lambda = 1600$ to monthly data (instead of quarterly) produce misleading results?

 TSA_ch12_sunspot

Exercise Solution Hints

Hints

1. By the Wiener-Khinchin theorem, $f(\omega) = \frac{1}{2\pi} \sum_{h=-\infty}^{\infty} \gamma(h) e^{-i\omega h}$. For WN: $\gamma(0) = \sigma^2$, $\gamma(h) = 0$ for $h \neq 0$, so $f(\omega) = \sigma^2/(2\pi)$ — flat.
2. AR(2) spectrum: $f(\omega) = \frac{\sigma^2}{2\pi |1 - 0.6e^{-i\omega} + 0.3e^{-2i\omega}|^2}$. The process has a spectral peak (complex roots). Verify numerically: the peak should lie between 0 and $\pi/2$.
3. Sunspot cycle: expect a sharp peak near $\omega = 2\pi/(11 \times 12) \approx 0.048$ rad/month. Long-memory tests typically give $\hat{d} \approx 0.3$ – 0.5 for sunspot numbers.
4. HP filter calibration: $\lambda = 1600$ is designed for quarterly data. For monthly data, the correct value is $\lambda \approx 1600 \times 3^4 = 129,600$. Using $\lambda = 1600$ on monthly data cuts off at the wrong frequency, distorting the extracted cycle.

Summary: Chapter 12

Part	Exercises
Year 3	Ex. 1–4
Master	Ex. 5–7
PhD	Ex. 8–10

Key tradeoff: Frequency resolution $\longleftrightarrow$ Statistical variance — more smoothing $\Rightarrow$ lower variance but cannot resolve nearby spectral peaks

Questions?

Bibliography I

Spectral Analysis Textbooks

- ▣ Hyndman, R.J., & Athanasopoulos, G. (2021). *Forecasting: Principles and Practice*, 3rd ed., OTexts.
- ▣ Shumway, R.H., & Stoffer, D.S. (2017). *Time Series Analysis and Its Applications*, 4th ed., Springer.
- ▣ Percival, D.B., & Walden, A.T. (1993). *Spectral Analysis for Physical Applications*. Cambridge University Press.

Long Memory and Filters

- ▣ Geweke & Porter-Hudak (1983). Estimation of long memory time series models. *J. Time Ser. Anal.*, 4, 221–238.
- ▣ Robinson (1995). Gaussian semiparametric estimation of long range dependence. *Ann. Statist.*, 23, 1630–1661.
- ▣ Hamilton (2018). Why you should never use the HP filter. *Rev. Econ. Statist.*, 100(5), 831–843.

Bibliography II

Wavelets and Advanced Methods

- ▣ Mallat, S. (2009). *A Wavelet Tour of Signal Processing*, 3rd ed., Academic Press.
- ▣ Gençay, R., Selçuk, F., & Whitcher, B. (2002). *An Introduction to Wavelets and Other Filtering Methods in Finance and Economics*. Academic Press.
- ▣ Baxter, M., & King, R.G. (1999). Measuring business cycles: Approximate band-pass filters. *Rev. Econ. Statist.*, 81(4), 575–593.

Online Resources and Code

- ▣ **Quantlet**: <https://quantlet.com> — Code repository for quantitative methods
- ▣ **Quantinar**: <https://quantinar.com> — Quantitative methods learning platform
- ▣ **GitHub TSA**: https://github.com/danpele/TSA/tree/main/TSA_ch12 — Python code for this seminar

Thank You!

Questions?

Seminar materials are available at: <https://danpele.github.io/Time-Series-Analysis/>



Quantlet



Quantinar