Introduction to Time Series Analysis. Lecture 19.

1. Review: Spectral density, rational spectra.
2. Linear filters.
3. Frequency response of linear filters.
4. Spectral estimation
5. Sample autocovariance
6. Discrete Fourier transform and the periodogram

Review: Spectral density

If a time series $\{X_t\}$ has autocovariance γ satisfying $\sum_{h=-\infty}^{\infty} |\gamma(h)| < \infty$, then we define its **spectral density** as

$$f(\nu) = \sum_{h=-\infty}^{\infty} \gamma(h) e^{-2\pi i \nu h}$$

for $-\infty < \nu < \infty$. We have

$$\gamma(h) = \int_{-1/2}^{1/2} e^{2\pi i \nu h} f(\nu)\, d\nu.$$

Review: Rational spectra

For a linear time series with $MA(\infty)$ polynomial ψ,

$$f(\nu) = \sigma_w^2 \left|\psi\left(e^{2\pi i\nu}\right)\right|^2.$$

If it is an ARMA(p,q), we have

$$\begin{aligned} f(\nu) &= \sigma_w^2 \left|\frac{\theta(e^{-2\pi i\nu})}{\phi\left(e^{-2\pi i\nu}\right)}\right|^2 \\ &= \sigma_w^2 \frac{\theta_q^2 \prod_{j=1}^q \left|e^{-2\pi i\nu} - z_j\right|^2}{\phi_p^2 \prod_{j=1}^p \left|e^{-2\pi i\nu} - p_j\right|^2}, \end{aligned}$$

where $z_1, \ldots, z_q$ are the zeros (roots of $\theta(z)$)
and $p_1, \ldots, p_p$ are the poles (roots of $\phi(z)$).

Time-invariant linear filters

A filter is an operator; given a time series $\{X_t\}$, it maps to a time series $\{Y_t\}$. We can think of a linear process $X_t = \sum_{j=0}^{\infty} \psi_j W_{t-j}$ as the output of a *causal linear filter* with a white noise input.

A time series $\{Y_t\}$ is the output of a linear filter $A = \{a_{t,j} : t, j \in \mathbb{Z}\}$ with input $\{X_t\}$ if

$$Y_t = \sum_{j=-\infty}^{\infty} a_{t,j} X_j.$$

If $a_{t,t-j}$ is independent of t ($a_{t,t-j} = \psi_j$), then we say that the filter is *time-invariant*.
If $\psi_j = 0$ for $j < 0$, we say the filter ψ is *causal*.

We'll see that the name 'filter' arises from the frequency domain viewpoint.

Time-invariant linear filters: Examples

1. $Y_t = X_{-t}$ is linear, but not time-invariant.
2. $Y_t = \frac{1}{3}(X_{t-1} + X_t + X_{t+1})$ is linear, time-invariant, but not causal:

$$\psi_j = \begin{cases} \frac{1}{3} & \text{if } |j| \leq 1, \\ 0 & \text{otherwise.} \end{cases}$$

3. For polynomials $\phi(B), \theta(B)$ with roots outside the unit circle, $\psi(B) = \theta(B)/\phi(B)$ is a linear, time-invariant, causal filter.

Time-invariant linear filters

The operation

$$\sum_{j=-\infty}^{\infty} \psi_j X_{t-j}$$

is called the *convolution* of X with ψ.

Time-invariant linear filters

The sequence ψ is also called the *impulse response*, since the output $\{Y_t\}$ of the linear filter in response to a *unit impulse,*

$$X_t = \begin{cases} 1 & \text{if } t = 0, \\ 0 & \text{otherwise,} \end{cases}$$

is

$$Y_t = \psi(B)X_t = \sum_{j=-\infty}^{\infty} \psi_j X_{t-j} = \psi_t.$$

Frequency response of a time-invariant linear filter

Suppose that $\{X_t\}$ has spectral density $f_x(\nu)$ and ψ is *stable*, that is, $\sum_{j=-\infty}^{\infty} |\psi_j| < \infty$. Then $Y_t = \psi(B)X_t$ has spectral density

$$f_y(\nu) = \left|\psi\left(e^{2\pi i\nu}\right)\right|^2 f_x(\nu).$$

The function $\nu \mapsto \psi(e^{2\pi i\nu})$ (the polynomial $\psi(z)$ evaluated on the unit circle) is known as the *frequency response* or *transfer function* of the linear filter.

The squared modulus, $\nu \mapsto |\psi(e^{2\pi i\nu})|^2$ is known as the *power transfer function* of the filter.

Frequency response of a time-invariant linear filter

For stable ψ, $Y_t = \psi(B)X_t$ has spectral density

$$f_y(\nu) = \left|\psi\left(e^{2\pi i\nu}\right)\right|^2 f_x(\nu).$$

We have seen that a linear process, $Y_t = \psi(B)W_t$, is a special case, since $f_y(\nu) = |\psi(e^{2\pi i\nu})|^2\sigma_w^2 = |\psi(e^{2\pi i\nu})|^2 f_w(\nu)$.

When we pass a time series $\{X_t\}$ through a linear filter, the spectral density is multiplied, frequency-by-frequency, by the squared modulus of the frequency response $\nu \mapsto |\psi(e^{2\pi i\nu})|^2$.

This is a version of the equality $\mathrm{Var}(aX) = a^2\mathrm{Var}(X)$, but the equality is true for the component of the variance at every frequency.

This is also the origin of the name 'filter.'

Frequency response of a filter: Details

Why is $f_y(\nu) = \left|\psi\left(e^{2\pi i\nu}\right)\right|^2 f_x(\nu)$? First,

$$\begin{aligned}
\gamma_y(h) &= \mathrm{E}\left[\sum_{j=-\infty}^{\infty} \psi_j X_{t-j} \sum_{k=-\infty}^{\infty} \psi_k X_{t+h-k}\right] \\
&= \sum_{j=-\infty}^{\infty} \psi_j \sum_{k=-\infty}^{\infty} \psi_k \mathrm{E}\left[X_{t+h-k} X_{t-j}\right] \\
&= \sum_{j=-\infty}^{\infty} \psi_j \sum_{k=-\infty}^{\infty} \psi_k \gamma_x(h+j-k) = \sum_{j=-\infty}^{\infty} \psi_j \sum_{l=-\infty}^{\infty} \psi_{h+j-l} \gamma_x(l).
\end{aligned}$$

It is easy to check that $\sum_{j=-\infty}^{\infty} |\psi_j| < \infty$ and $\sum_{h=-\infty}^{\infty} |\gamma_x(h)| < \infty$ imply that $\sum_{h=-\infty}^{\infty} |\gamma_y(h)| < \infty$. Thus, the spectral density of y is defined.

Frequency response of a filter: Details

$$
\begin{aligned}
f_y(\nu) &= \sum_{h=-\infty}^{\infty} \gamma(h) e^{-2\pi i \nu h} \\
&= \sum_{h=-\infty}^{\infty} \sum_{j=-\infty}^{\infty} \psi_j \sum_{l=-\infty}^{\infty} \psi_{h+j-l} \gamma_x(l) e^{-2\pi i \nu h} \\
&= \sum_{j=-\infty}^{\infty} \psi_j e^{2\pi i \nu j} \sum_{l=-\infty}^{\infty} \gamma_x(l) e^{-2\pi i \nu l} \sum_{h=-\infty}^{\infty} \psi_{h+j-l} e^{-2\pi i \nu (h+j-l)} \\
&= \psi(e^{2\pi i \nu j}) f_x(\nu) \sum_{h=-\infty}^{\infty} \psi_h e^{-2\pi i \nu h} \\
&= \left| \psi(e^{2\pi i \nu j}) \right|^2 f_x(\nu).
\end{aligned}
$$

Frequency response: Examples

For a linear process $Y_t = \psi(B)W_t$, $f_y(\nu) = \left|\psi\left(e^{2\pi i\nu}\right)\right|^2 \sigma_w^2$.

For an ARMA model, $\psi(B) = \theta(B)/\phi(B)$, so $\{Y_t\}$ has the rational spectrum

$$\begin{aligned} f_y(\nu) &= \sigma_w^2 \left| \frac{\theta(e^{-2\pi i\nu})}{\phi\left(e^{-2\pi i\nu}\right)} \right|^2 \\ &= \sigma_w^2 \frac{\theta_q^2 \prod_{j=1}^{q} \left|e^{-2\pi i\nu} - z_j\right|^2}{\phi_p^2 \prod_{j=1}^{p} \left|e^{-2\pi i\nu} - p_j\right|^2}, \end{aligned}$$

where p_j and z_j are the poles and zeros of the rational function $z \mapsto \theta(z)/\phi(z)$.

Frequency response: Examples

Consider the moving average

$$Y_t = \frac{1}{2k+1} \sum_{j=-k}^{k} X_{t-j}.$$

This is a time invariant linear filter (but it is not causal). Its transfer function is the Dirichlet kernel

$$\begin{aligned}
\psi(e^{-2\pi i\nu}) = D_k(2\pi\nu) &= \frac{1}{2k+1} \sum_{j=-k}^{k} e^{-2\pi i j\nu} \\
&= \begin{cases} 1 & \text{if } \nu = 0, \\ \frac{\sin(2\pi(k+1/2)\nu)}{(2k+1)\sin(\pi\nu)} & \text{otherwise.} \end{cases}
\end{aligned}$$

Example: Moving average

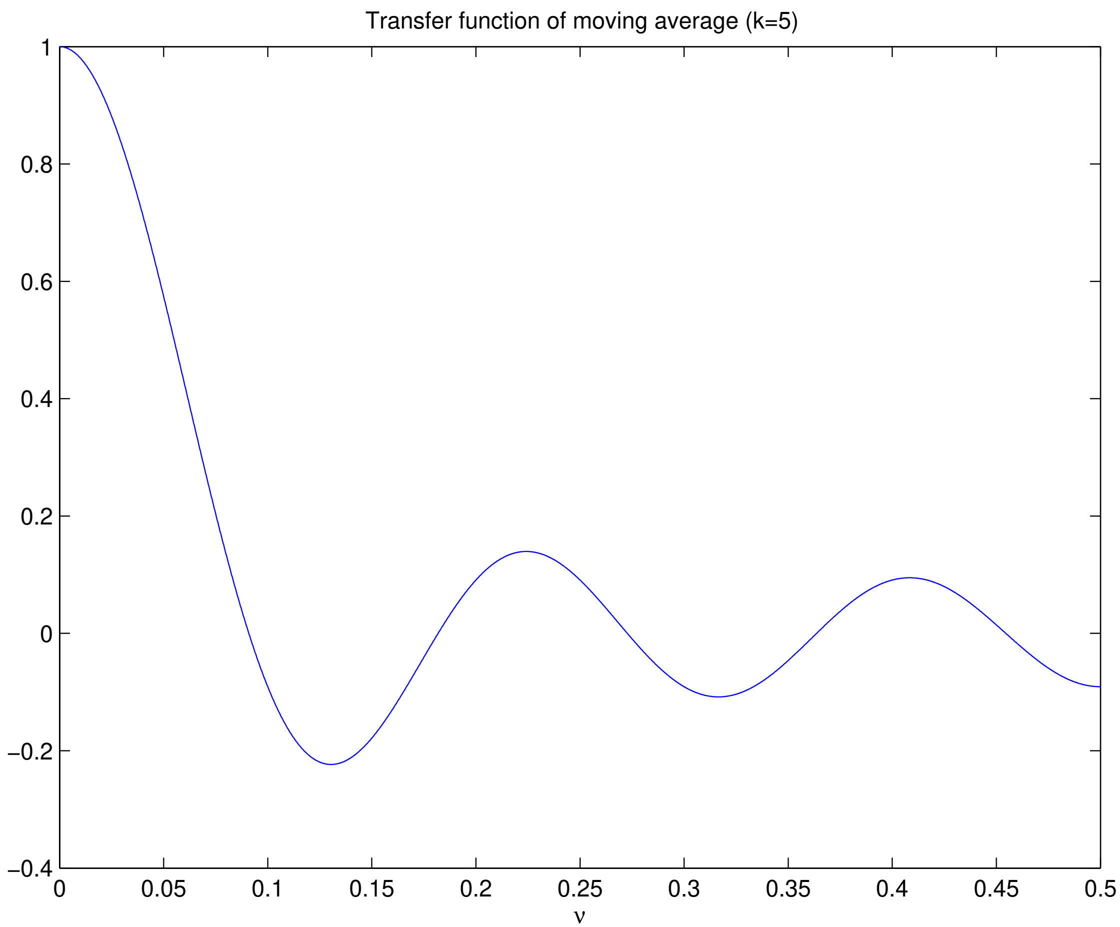

Example: Moving average

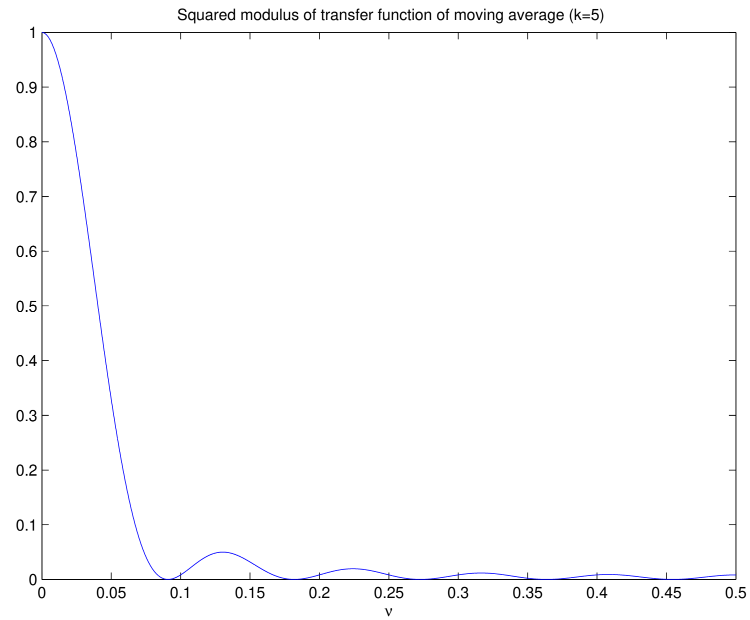

This is a *low-pass filter*: It preserves low frequencies and diminishes high frequencies. It is often used to estimate a monotonic trend component of a series.

Example: Differencing

Consider the first difference

$$Y_t = (1 - B)X_t.$$

This is a time invariant, causal, linear filter.

Its transfer function is

$$\psi(e^{-2\pi i\nu}) = 1 - e^{-2\pi i\nu},$$

$$\text{so} \qquad |\psi(e^{-2\pi i\nu})|^2 = 2(1 - \cos(2\pi\nu)).$$

Example: Differencing

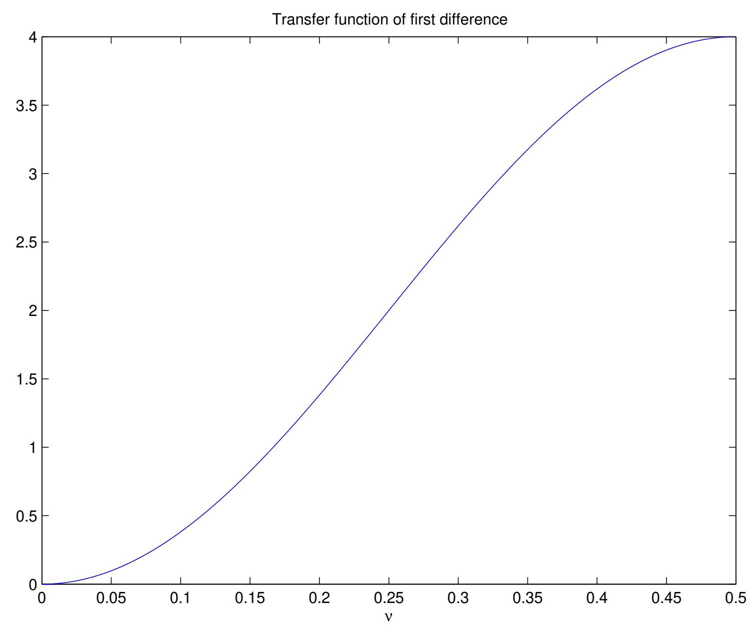

This is a *high-pass filter*: It preserves high frequencies and diminishes low frequencies. It is often used to eliminate a trend component of a series.

Estimating the Spectrum: Outline

- We have seen that the spectral density gives an alternative view of stationary time series.
- Given a realization $x_1, \ldots, x_n$ of a time series, how can we estimate the spectral density?
- One approach: replace $\gamma(\cdot)$ in the definition

$$f(\nu) = \sum_{h=-\infty}^{\infty} \gamma(h) e^{-2\pi i \nu h},$$

  with the sample autocovariance $\hat{\gamma}(\cdot)$.
- Another approach, called the *periodogram*: compute $I(\nu)$, the squared modulus of the discrete Fourier transform (at frequencies $\nu = k/n$).

Estimating the spectrum: Outline

- These two approaches are *identical* at the Fourier frequencies $\nu = k/n$.
- The asymptotic expectation of the periodogram $I(\nu)$ is $f(\nu)$. We can derive some asymptotic properties, and hence do hypothesis testing.
- Unfortunately, the asymptotic variance of $I(\nu)$ is constant. It is not a consistent estimator of $f(\nu)$.
- We can reduce the variance by smoothing the periodogram—averaging over adjacent frequencies. If we average over a narrower range as $n \to \infty$, we can obtain a consistent estimator of the spectral density.

Estimating the spectrum: Sample autocovariance

Idea: use the sample autocovariance $\hat{\gamma}(\cdot)$, defined by

$$\hat{\gamma}(h) = \frac{1}{n} \sum_{t=1}^{n-|h|} (x_{t+|h|} - \bar{x})(x_t - \bar{x}), \qquad \text{for } -n < h < n,$$

as an estimate of the autocovariance $\gamma(\cdot)$, and then use a sample version of

$$f(\nu) = \sum_{h=-\infty}^{\infty} \gamma(h) e^{-2\pi i \nu h},$$

That is, for $-1/2 \leq \nu \leq 1/2$, estimate $f(\nu)$ with

$$\hat{f}(\nu) = \sum_{h=-n+1}^{n-1} \hat{\gamma}(h) e^{-2\pi i \nu h}.$$

Estimating the spectrum: Periodogram

Another approach to estimating the spectrum is called the periodogram. It was proposed in 1897 by Arthur Schuster (at Owens College, which later became part of the University of Manchester), who used it to investigate periodicity in the occurrence of earthquakes, and in sunspot activity.

Arthur Schuster, "On Lunar and Solar Periodicities of Earthquakes," *Proceedings of the Royal Society of London*, Vol. 61 (1897), pp. 455–465.

To define the periodogram, we need to introduce the *discrete Fourier transform* of a finite sequence $x_1, \ldots, x_n$.