Stochastic processes

Susan Thomas

August 12, 2009

Topics

- Defining stochastic processes
- Stationary stochastic processes
- Sample autocorrelations, partial autocorrelations,

The core problem of modelling a time series

1. To model and forecast a univariate time series.
2. To establish the relation between a set of univariate time series.

The first involves

1. finding out how much of the past information matters and how much comes from the latest innovation.
2. finding out the distribution and behaviour of the innovation.
3. finding out how much of the behaviour is “permanent” and how much is “temporary”.

Once this is established for the univariate time series, the multivariate modelling can begin.
In this course, we focus on the univariate time series modelling and forecasting problem.

The core problem of modelling a time series

- The biggest worry: whether the innovations are "stable" through time or not.
  If the innovation series (ϵ_t) are not stable – ie, the variance/covariances are a function of time – then it becomes difficult to understand what drives the DGP.
- This question of *stationarity* or *non–stationarity* of the time series is the first and important aspect of understanding a time series, x_t.
- The OLS perspective of estimation: need to explicitly deal with multicollinearity in time series.

Part I

Stochastic time series processes

What is a deterministic time series process?

- *Deterministic time series processes*: when the effect of time on the next observation in the time series is deterministic.

$$\begin{aligned} y_t &= a + bt + \eta_t \\ y_{t-1} &= a + b(t-1) + \eta_{t-1} \\ \Delta y_t = (1-L)y_t = y_t - y_{t-1} &= b + (\eta_t - \eta_{t-1}) \\ \mathrm{E}\Delta y_t &= b \quad \forall t \end{aligned}$$

Here, the expected change in y_t is independent of what the value of "t" really is.

What is a stochastic time series process?

- *Stochastic time series processes*: when the effect on the next observation depends upon what happened before.

$$
\begin{aligned}
y_t &= a + by_{t-1} + \epsilon_t \\
y_{t-1} &= a + by_{t-2} + \epsilon_{t-1} \\
\Delta y_t = (1-L)y_t = y_t - y_{t-1} &= b(1-L)y_{t-1} + (\epsilon_t - \epsilon_{t-1}) \\
\mathrm{E}\Delta y_t &= b\Delta y_{t-1}
\end{aligned}
$$

The expected change in y_t is "dynamic" – changes in value every time period.

Classification of models used for univariate time series

1. Linear models for the conditional expectation of the time series. For example,

$$\begin{aligned} x_t &= a + bx_{t-1} + \epsilon_t \\ \epsilon_t &\sim D(0, \sigma^2) \end{aligned}$$

   - Stationary : $-1 < b < 1$
   - Nonstationary : $b \geq 1; b \leq -1$

2. Linear models for the conditional variance. For example,

$$\begin{aligned} x_t &= a + bx_{t-1} + \epsilon_t \\ \epsilon_t &\sim D(0, \sigma_t^2) \\ \sigma_t^2 &= g_0 + g_1 \sigma_{t-1}^2 \end{aligned}$$

Univariate time series models, contd.

1. Nonlinear models for the conditional mean – smooth transition models which can be thought of as a linear model with time-varying parameters. Example: regime switching models.
2. More general non-inear models: where parameters of both the conditional expectation and variance of the time series variable is time-varying.

Part II

Time series in `samples.rda`

Stock market prices (daily): Nifty

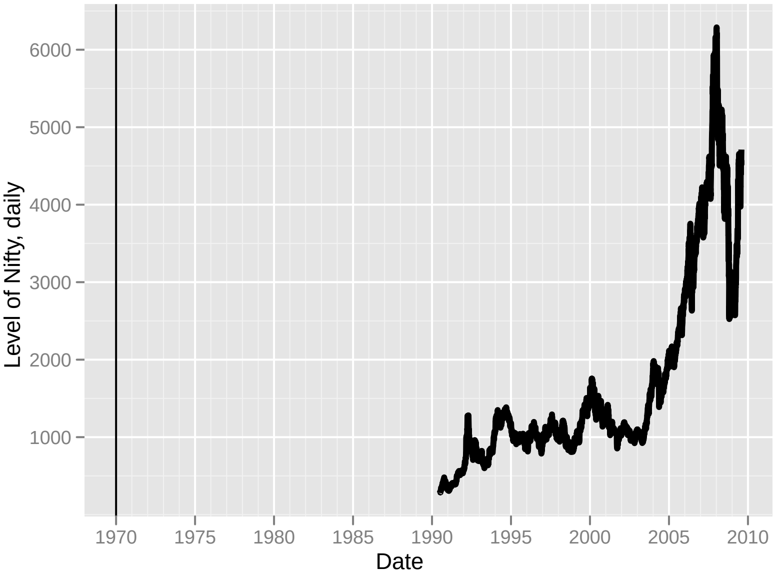

Stock market prices (daily): SNP500

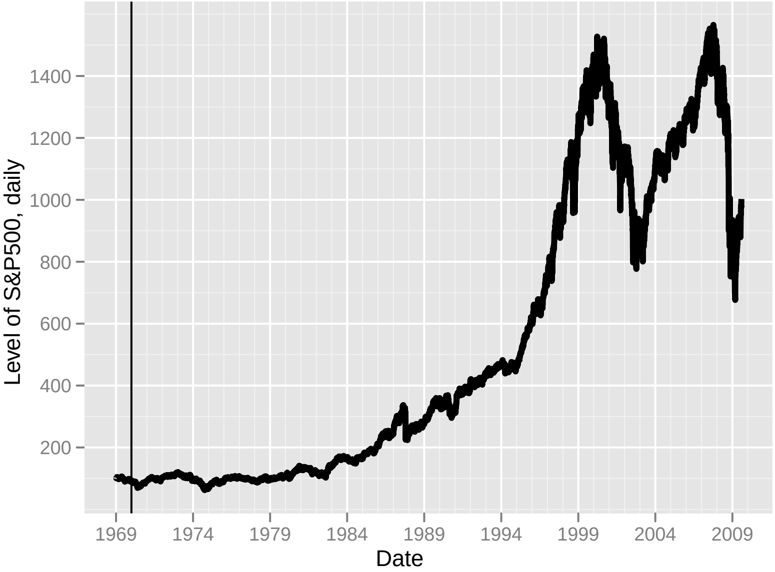

Stock market prices (daily): Infosys Technologies

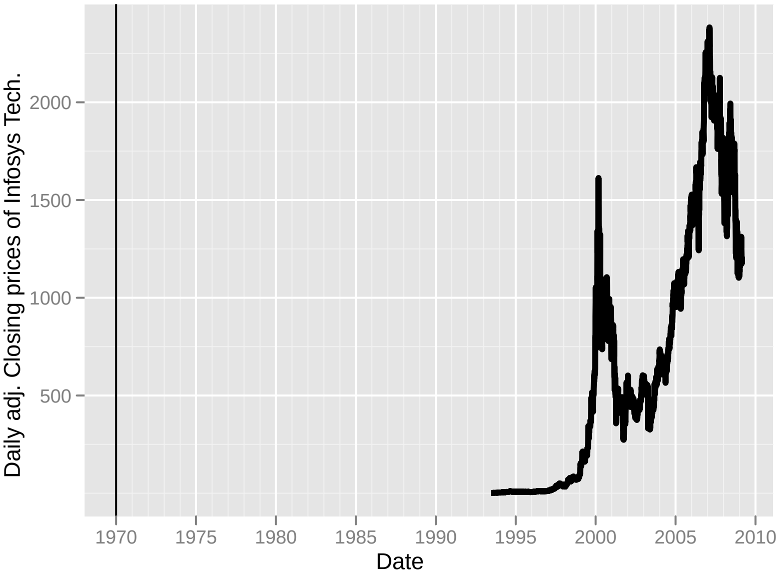

Stock market prices (daily): Reliance Industries Ltd.

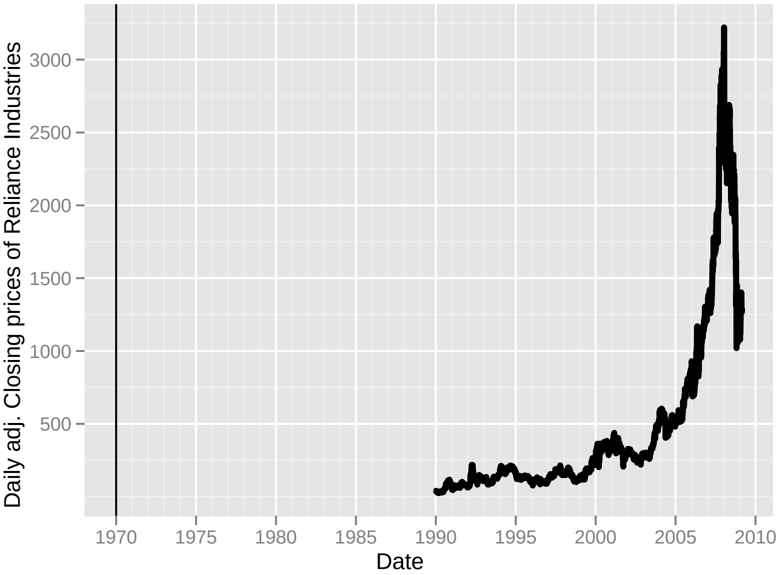

Stock market prices (daily): Nifty

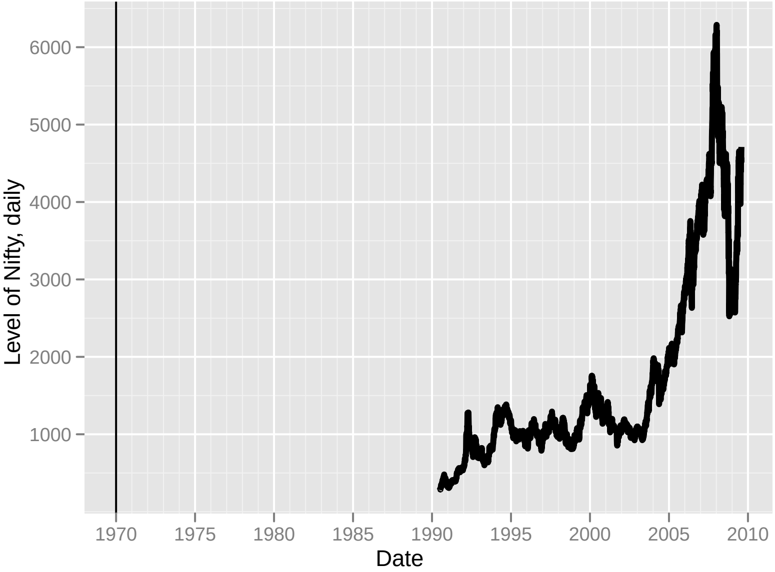

Stock market capitalisation (monthly): COSPI

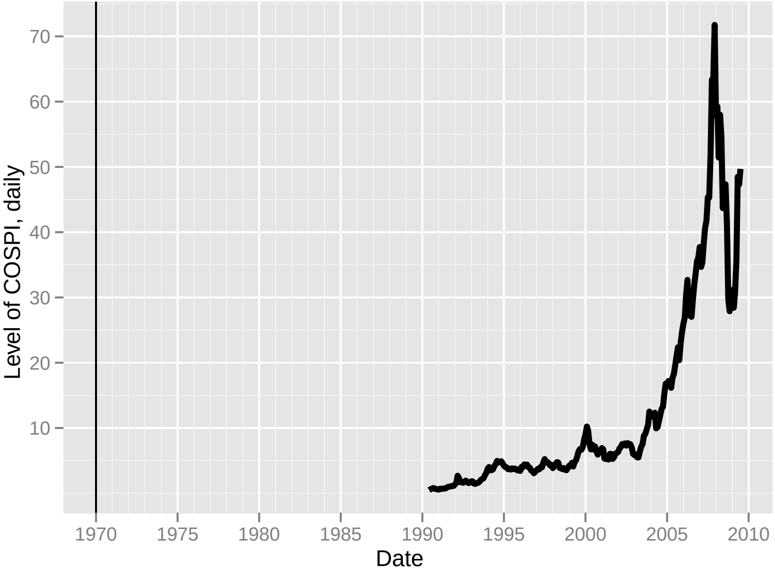

Bond market rates (daily): US 3-month interest rates

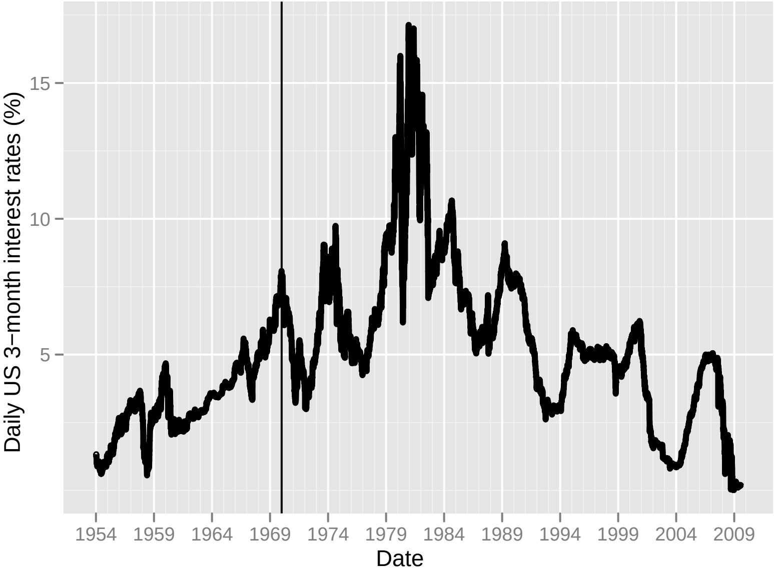

Bond market rates (daily): US 10-year interest rates

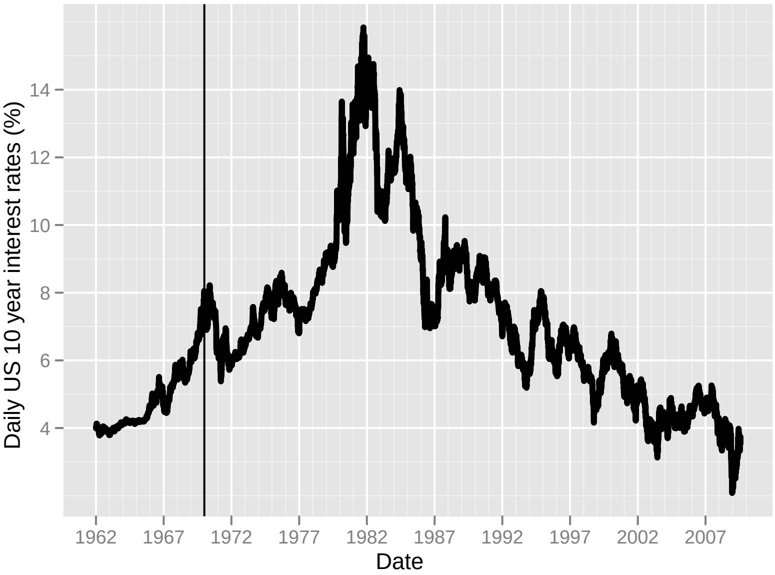

Bond market rates (daily): US BAA-rated interest rates

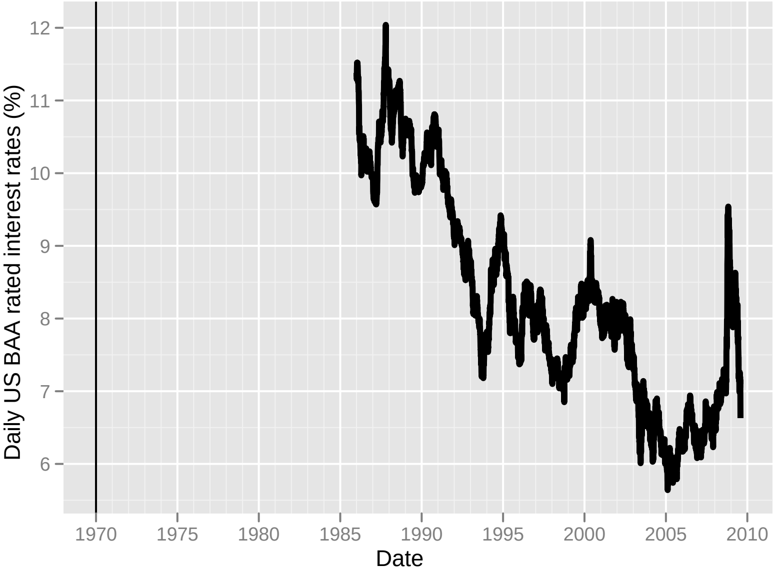

Monthly Indian 3-month interest rates (%)

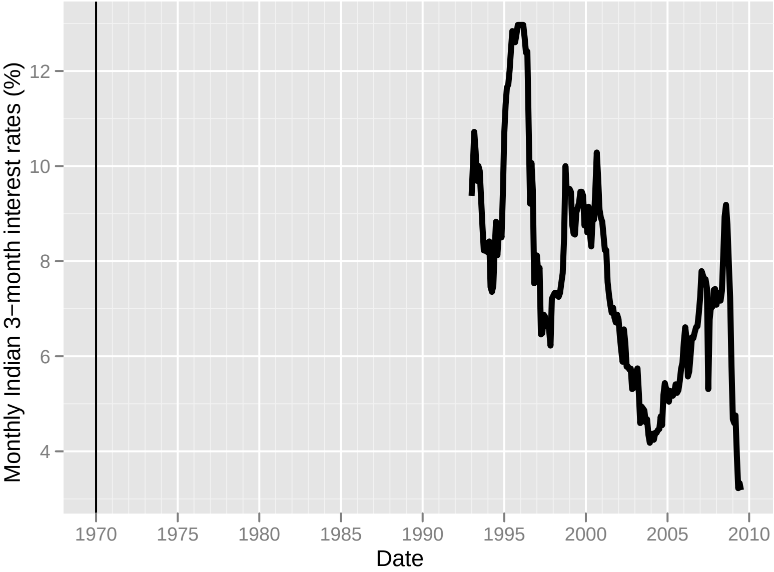

Monthly US 3-month interest rates (%)

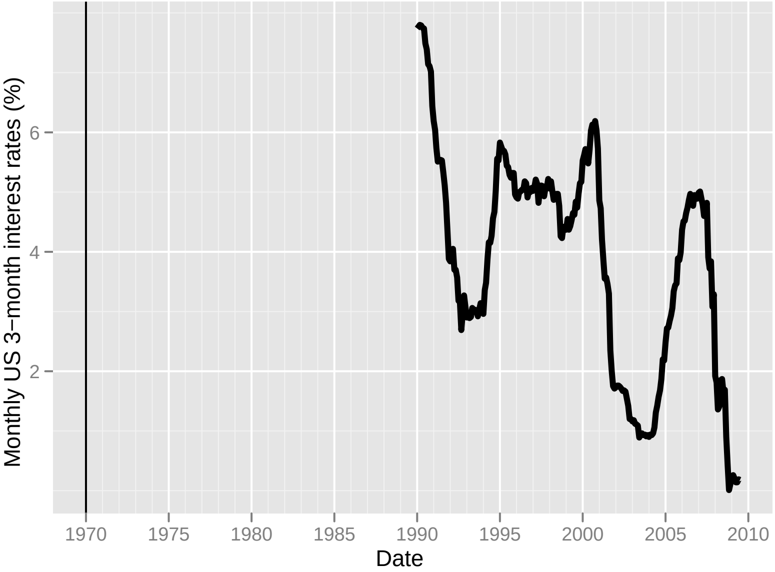

Indian IIP (monthly): Raw levels

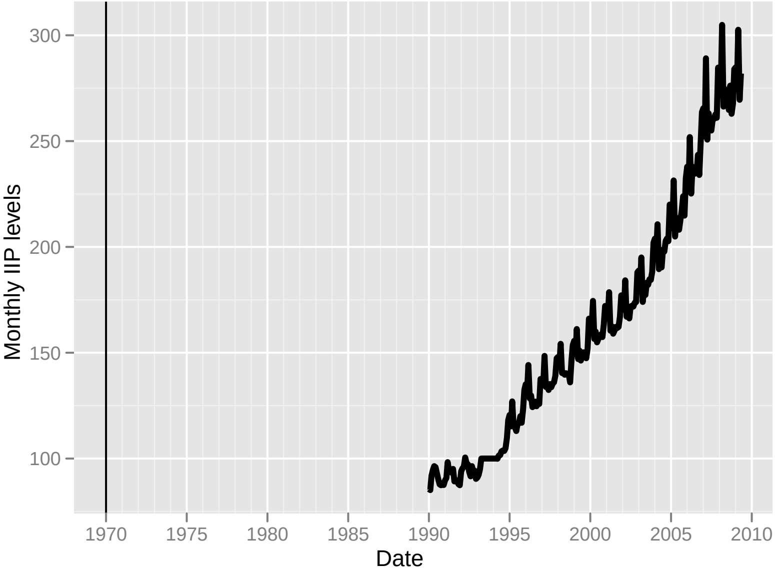

Indian IIP (monthly): Seasonally adjusted levels

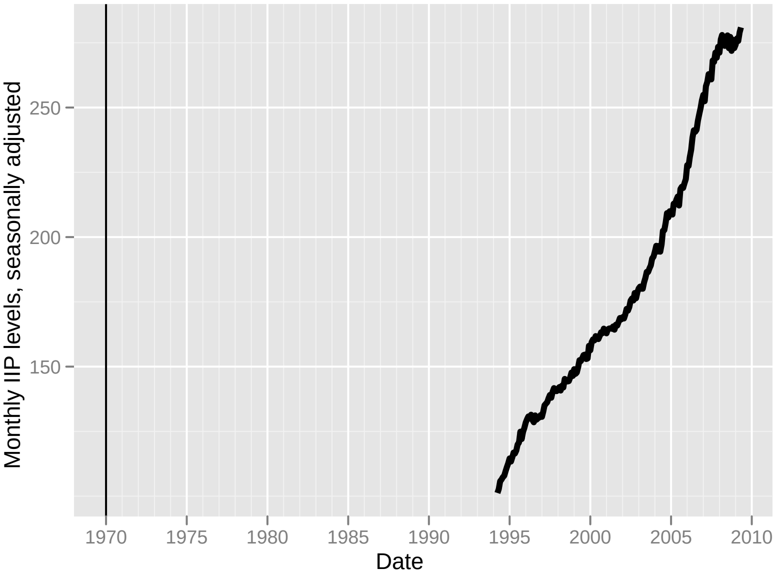

Part III

Linear models for the conditional expectation of stochastic time series processes

Stationary stochastic processes

- Any stochastic process is a collection of random variables

$$y_t, y_{t-1}, y_{t-2}, y_{t-3}, y_{t-4}, \cdots$$

  where each observation in the series is assumed to be generated by the previous observation.
  Ie, y_t is assumed to be generated/linked to y_{t-1}.
- The data itself may be only a sample – ie, there are observations before the first one and after the last one.
- Each t can be any frequency. However, the frequency is fixed for a given time series.

Stationary stochastic processes – Definition

y_t is stationary if it has first and second moments that are time-invariant.

- First moment: $\mathrm{E}(y_t) = \mu_y \quad \forall t \subset T$
- Second moments: $\mathrm{E}[(y_t - \mu_y)(y_{t-h} - \mu_y)] = \gamma_h \quad \forall t \subset T, \forall h$ such that $(t-h) \subset T$.
  where γ_h is called the autocovariance at lag of h with respect to the data.

Intuition

- The first condition: values of the series must fluctuate around a mean that is constant.
  For example, if we take subsets of data, the mean should be similar across all sets.
  Observation: None of the data in `samples.rda` appear to have this feature.
- The second condition implies that the variances and the covariances of the series remain constant with time as well.
  Note: Each covariance is called an *autocovariance*.
  A stationary process has variance/covariances that do not change with time.
  In addition, each covariance is a function of the number of lags h.

Nomenclature and norms

- Sometimes, a process is called *covariance stationary* if the first and the second moments are constant.
  A process is also called covariance stationary if $E(y_t - \mu_y)(y_{t-h} - \mu_y) = f(h) \forall h$
- A process is called "trend stationary" when it becomes stationary when a deterministic trend (like a term "$a + bt$") is removed from it.
- At the start of a DGP, it is possible that a series appears to not be covariance stationary until some "start-up" period. Then, this series is called "asymptotically stationary". If a process can be made stationary by modifying some initial values, that is called asymptotically stationary.

Tools to detect second moment autocovariances: autocorrelations and partial autocorrelations

- Unlike in the case of the consistency of first moments in the DGP, it is difficult to visually inspect data for stationarity in the second moments.
- Statistics used to test for second-moment stationarity: sample autocovariances or autocorrelations and partial autocorrelations.
  1. Autocorrelations (ACs):

$$\tilde{\rho}_h = \tilde{\gamma}_h / \tilde{\gamma}_0$$

     where $\gamma_h = E(y_t - \mu_y)(y_{t-h} - \mu_y)$
  2. Partial Autocorrelations (PACs): correlation between y_t, y_{t-h} conditional on $y_{t-1}, y_{t-2}, \ldots, y_{t-h+1}$.

Calculating sample autocorrelations

- ACs are calculated as

  autocovariance / variance.

  where the autocovariances are calculated at a fixed lag, *h*.
  ACs are denoted as **AC(h)**.
- Sample AC(h) is calculated as $\tilde{\gamma}_h$:

$$\tilde{\gamma}_h = \frac{1}{(T-h)} \sum_{t=h+1}^{T} (y_t - \bar{y})(y_{t-h} - \bar{y})$$

$$\bar{y} = \sum_{t=1}^{T} y_t$$

- **Observation**: $H_0 : \gamma_h = 0$ for stationary series.
  At worst, under the null of stationarity, γ_h should grow as *h*.
  For a stationary series, the sample ACs die out quickly.

Calculating sample partial autocorrelations

- PACs are calculated from the regression:

$$y_t = \nu + \alpha_1 y_{t-1} + \alpha_2 y_{t-2} + \alpha_3 y_{t-3} + ldots + \alpha_h y_{t-h} + u_t$$

  where sample PAC(h) = the OLS estimate $\hat{\alpha}_h$ in the above regression.
- **Observation**: $H_0 : \alpha_h \to 0$ as $h \to \infty$ for stationary series. For a stationary series, the sample PACs die out quickly.

ACs for daily levels of Nifty

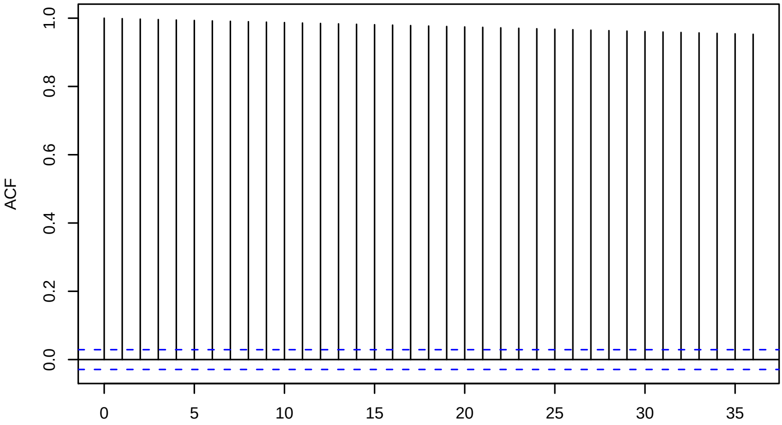

PACs for daily levels of Nifty

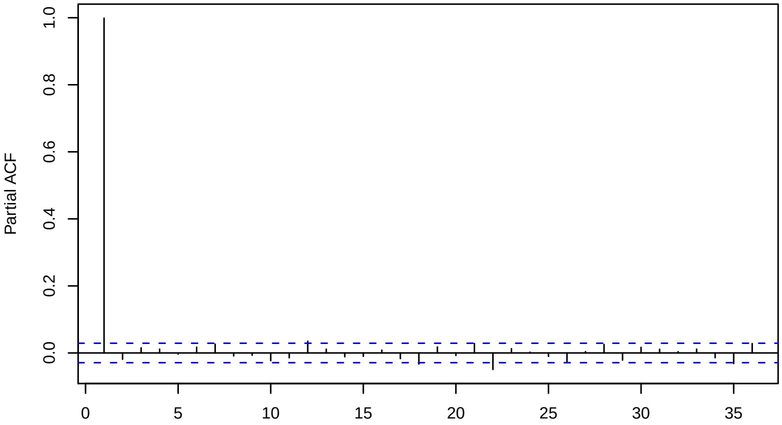

ACs for daily levels of S&P500

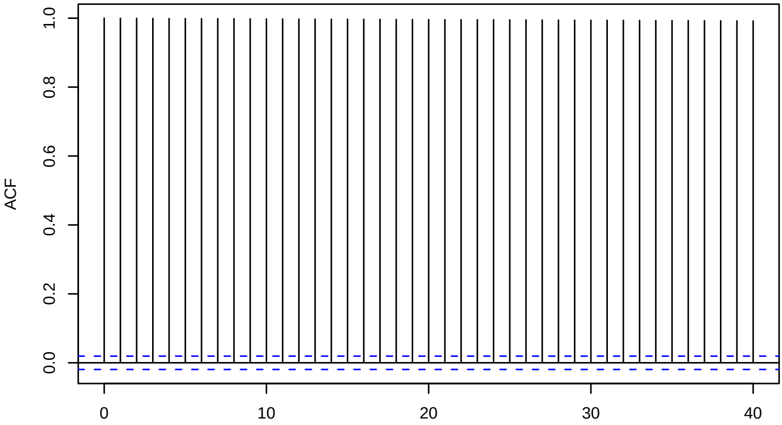

PACs for daily levels of S&P500

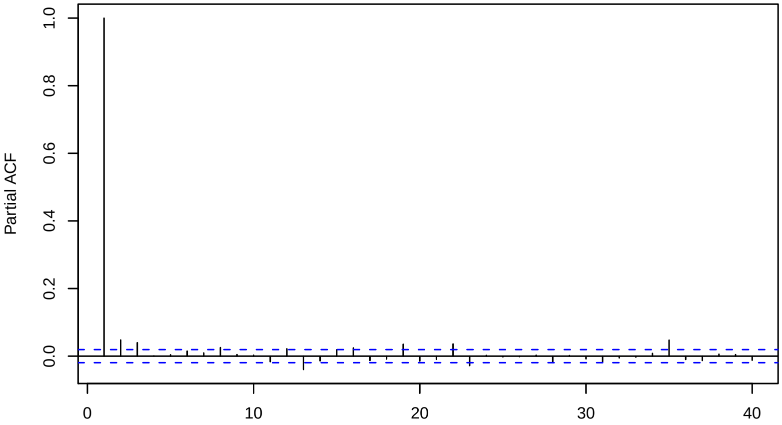

ACs for monthly levels of Indian IIP

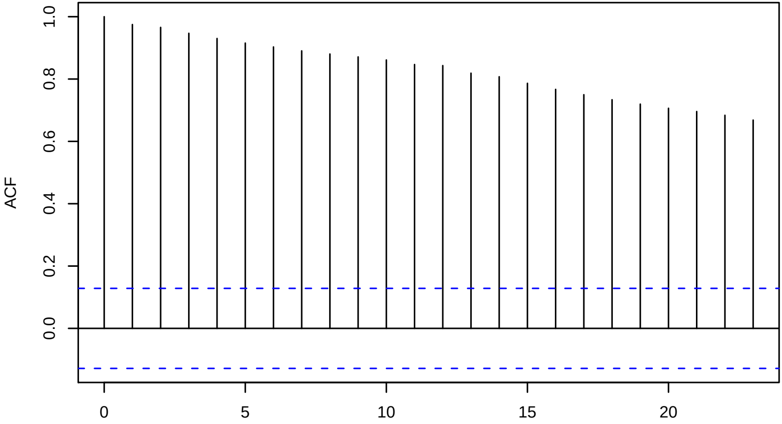

PACs for monthly levels of Indian IIP

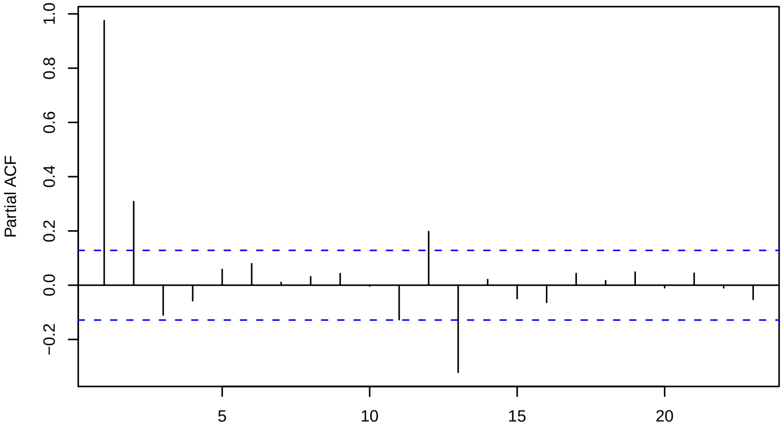

Part IV

Data transformations and filters

Log transformations and rates of change

- Sometimes simple transformations can move a series closer to stationarity.
  1. Sometimes it is observation about the data – does the series show larger fluctuations for larger values of the series?
     Suggested transform: a log transform may help to identify a trend in the series.
  2. Sometimes the data is explicitly seasonal.
     For instance, there is seasonality in the IIP data. This might be a deterministic or a stochastic seasonality.
  3. Sometimes, theory suggests the transformation.
     For instance, prices are modelled as log-normal. Then the difference of the log(prices) – $\Delta \log P_t = r_t$ or returns – become stationary.

ACs for daily levels of Nifty returns

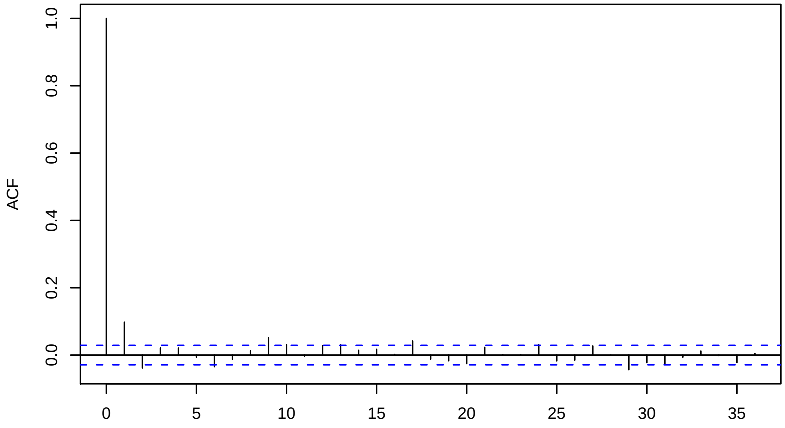

PACs for daily levels of Nifty returns

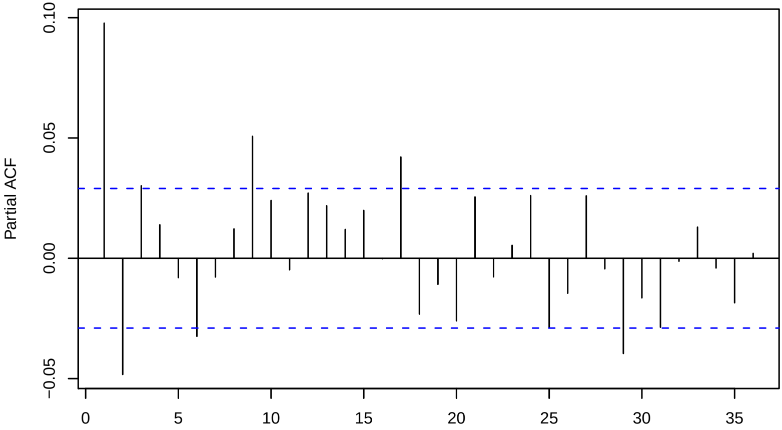

Data filtering

- Sometimes time series are filtered to remove a specific feature, transforming it from one series into another.
- Typically, a filter is a linear funciton.
- Example: $y_t, y_{t-1}, y_{t-2}, \ldots$ might be filtered to $x_t, x_{t-1}, \ldots$ using:

$$x_t = \sum_{j=-k}^{l} \omega_j y_{t-j}, t = k+1, \ldots, T-l$$

  where ω_j is a weight on lag j.
- Usually, the weights are designed to add upto 1.
- Filtering is often used to seasonally adjust data like quarterly data.
- A more generic filter is the Hodrick-Prescott filter, which is used to adjust cyclical data.
  Often used in the context of business cycle analysis.

Data filtering using moving averages

- Example: $\omega = (\frac{1}{8}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{8})$

$$x_t = \frac{1}{8}y_{t-2} + \frac{1}{4}y_{t-1} + \frac{1}{4}y_t + \frac{1}{4}y_{t+1} + \frac{1}{8}y_{t+2}$$

  gives x_t as a weighted moving average transformation of y_t
- Such type of filters are often used to remove "excessive noise" from the underlying data, which in turn helps identify patterns more readily.
- This can be more efficiently re-written as:

$$x_t = (\frac{1}{8}L^{-2} + \frac{1}{4}L^{-1} + \frac{1}{4}L^0 + \frac{1}{4}L^{+1} + \frac{1}{8}L^{+2})y_t$$

Data filtering using differencing

- $$x_t = \Delta y_t = y_t - y_{t-1}$$

  If y_t is a non-stationary process, this filter gives x_t as a stationary process.
- y_t is said to be an *integrated time series*.
- If by differencing once, the resultant series x_t becomes stationary, then y_t is said to be integrated to the order one. Typically written as y_t is an $I(1)$ series.
- Sometimes y_t can have *seasonal integration*.
  Example, if IIP is a seasonally stochastic monthly time series. Then,

  $$x_t = \Delta y_t = y_t - y_{t-12}$$

  will be the resulting stationary series.

Textbooks with detailed treatment of time series models

- Fuller 1976: *Introduction to statistical time series*
- Priestly 1981: *Spectral analysis and time series*
- Brockwell and Davis 1987: *Time series: Theory and methods*
- Hamilton 1994: *Time series analysis*