Lecture 1: Fundamentals of Time Series

Ganesh Viswanath Natraj

IB9Y60: Empirical Finance
Warwick Business School

Monday 16th January, 2023

Organisation- Blended learning

THE UNIVERSITY OF WARWICK

- 3 hours of learning a week, the $1 + 1 + 1$ system of **blended learning**
  1. 1 hour of live lectures
  2. 1 hour of pre-recorded material
  3. 1 hour of seminars
- To do well, it is expected students continually revise topics and keep up with lectures and pre-recorded material.
- Pre-recorded material for the week is to be reviewed after the live lecture.
- Theory solutions for the seminar is covered in pre-recorded material for that week. Please have a go at solving the theory questions before reviewing this video!
- In addition, I will use Vevox questions during lecture to engage student understanding.

Organisation

- Lecture: Monday 2pm-3pm M1
- Seminars: (Starting week 3)
  1. Thursday 9am-10am 1.007
  2. Thursday 12pm-1pm 1.007
  3. Thursday 2pm-3pm 1.007
  4. Thursday 3pm-4pm 1.007
  5. Thursday 4pm-5pm 1.007
- Lectures will go through the theory, seminars will go through empirical methods using Matlab, the recommended language of the course
- You are welcome to use Python/R, however I do not expect you to use it during the course

Organisation- the Team

THE UNIVERSITY OF WARWICK

Lecturer

- Ganesh Viswanath-Natraj
- Email: ganesh.viswanath-natraj@wbs.ac.uk
- Office Hours: Monday 4pm-6pm 2.209

Seminar TA

- Junxuan Wang
- Email: junxuan.wang.19@mail.wbs.ac.uk
- Office Hours: Thursday 1pm-3pm PhD office 2.008

Organisation- Assessment

The structure of the course is 70% final exam, 20% group project and 10% class test.

Group Project (20%)

- 4 Questions in topics of econometric forecasting and cointegration, volatility modeling, PCA and factor analysis, and a 2 page research proposal on a topic in empirical finance.
- Submit names (5-6 people) to FinancePG@wbs.ac.uk by Monday 30th January.
- If you do not sort into groups by then, you will be randomized by the Masters office. Final group listings will be released early February.
- Expect release of project early February, and due date: **online submission 30th March, 12pm**.

Organisation- Assessment

Class Test (10%):

- Multiple Choice questions, theory and empirical
- Will cover topics 1 through to topics 4
- Date: **Friday 3rd March, 9:15am-10am**, more details provided closer to date.

Final Exam (70%):

- Details of Exam time/venue TBD
- Exam will cover all topics 1-9, mix of theory and empirical questions

Course Objectives

- Learn a series of econometric methods (VECM, volatility modeling)
- Learn how to apply these methods to financial data
- Learn a range of empirical stylised facts drawn from the analysis of financial markets; the rates, models of equity returns, the yield curve and exchange rates
- Learn how to test asset pricing models, (CAPM, Fama French, Consumption asset pricing models)

Lecture Topics

1. Time Series Fundamentals : ACF, PACF, ARMA, Lag-Operator, Stationarity, Information Criteria, Unit-Root Testing
2. Time Series Forecasting: Cointegration, VECM, model evaluation, empirical application: exchange rate forecasting
3. Volatility Modeling: Historical Volatility and Bloomberg Risk Metrics, ARCH, GARCH processes. Forecasting Volatility, Asymmetric Volatility Modeling
4. Value at Risk and Non-Normality: VaR models, non-normality (QQ Plot)

Lecture Topics

5. Capital Asset Pricing Model: Fama French Factors, Portfolio Analysis, Fama MacBeth
6. Factor analysis: Empirical application: factor analysis of currency excess returns. PCA
7. Generalized Method of Moments (GMM): Theory, Asymptotic properties, empirical application: estimating parameters of consumption asset pricing model.
8. Monte Carlo Simulations: Empirical applications: Black Scholes options pricing under normality and non-normality
9. Panel Data, Binary Dependent Variable Models: Empirical applications: Banking Competition, Predicting Default

Text Book

Recommended text

- Brooks, C. Introductory Econometrics for Finance (2019). Cambridge University Press.

Supplementary texts

- Ruppert,D. and Matteson,D. Statistics and Data Analysis for Financial Engineering with R Examples, 2nd edition, Springer, 2015.
- Campbell, J.Y., Lo, A.W. and MacKinley, A.C. (1997), The Econometrics of Financial Markets, Princeton.
- Enders, W. (2009), Applied Econometric Time Series, Wiley.
- Cochrane, J, Time Series Notes, https://faculty.chicagobooth.edu/john.cochrane/research /Papers/time_series_book.pdf.

Topic 8: Monte Carlo Simulations

- Can we numerically estimate the Black Scholes formula using the Monte Carlo method?
- Simulate the path of S_t from today until time to expiry, using the following formula for the stock price

$$S_t = S_{t-1} \cdot e^{((r-\frac{1}{2}\cdot\sigma^2)\cdot\delta_t+\sigma\cdot\sqrt{\delta_t}\cdot Z_t)}$$

- For N simulations, we calculate the payoff at time to expiry $T-t$, $\Pi = max(S_T - X, 0)$. The average discounted payoff over N is the value of the Call–this will be numerically equivalent to using Black Scholes formula.

$$\hat{C}_t = e^{-r(T-t)}\frac{\sum_{i=1}^{N}\Pi_i}{N}$$

Topic 8: Monte Carlo Simulations

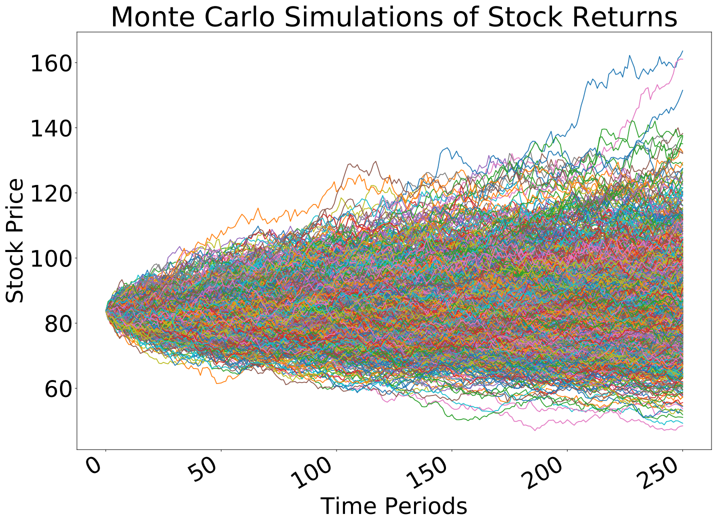

Roadmap of Lecture

THE UNIVERSITY OF WARWICK

- Definitions: Stationarity, Autocorrelation Function, Autoregressive (AR) and Moving Average (MA) processes
- Box-Jenkins approach for ARMA model selection
- Application: Modeling of SP 500 returns as an ARMA process
- Reference: Brooks Chapter 6

Definition: Types of Data

- Time series are data on one variable collected over time.
  $T > 1, N = 1$
- Examples: inflation or unemployment (monthly frequency), government budget deficit (annual) or stock price indices (intra-day frequency).
- Cross-sectional data are data on one or more variables collected at a single point in time. $T = 1, N > 1$
- Examples: consumption of households in the UK. Cost of all items in a grocery store.
- Panel data has dimensions of both time series and cross-sections.
  $T > 1, N > 1$
- Examples: a census survey of income of multiple households ($i = 1, 2, ..., N$) over multiple years ($t = 1, 2, ..., T$)

Definition: White Noise process

THE UNIVERSITY OF WARWICK

- A process ϵ_t is called a white noise (WN) process if has zero mean, constant variance, and the shocks are independent and identically distributed over time, also known as i.i.d.
    1. $\mathbb{E}[\epsilon_t] = 0$
    2. $\mathbb{E}[\epsilon_t \epsilon_{t-1}] = cov(\epsilon_t, \epsilon_{t-1}) = 0$
    3. $var(\epsilon_t) = var(\epsilon_t | \epsilon_{t-1}, \epsilon_{t-2}, ...) = \sigma^2_\epsilon$
- If $\epsilon_t \sim N(0, \sigma^2_\epsilon)$, then ϵ_t is called **Normal White Noise (NWN)**
- White noise processes are the building blocks of all time series!

Definition: White Noise process

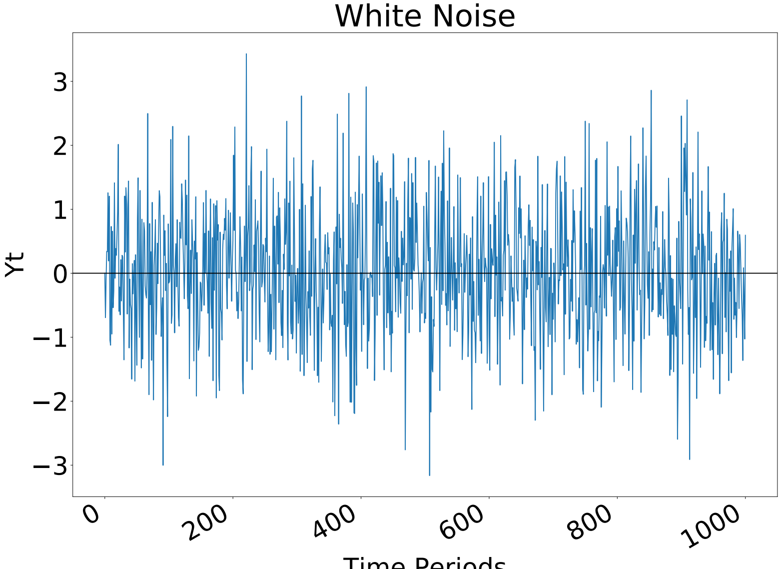

Definition: Strict Stationarity

- A process is strictly stationary if, for any values of $j_1, j_2, ..., j_n$ the joint distribution of $y_t, y_{t+j_1}, ..., y_{t+j_n}$ depends only on the intervals separating the dates $(j_1, j_2, ...j_n)$ and not on the dates themselves (t)
- Strict stationarity requires that the joint distribution of a stochastic process does not depend on time
- The only factor affecting the relationship between two observations is the gap between them.
- Strict stationarity is weaker than i.i.d since the process maybe serially correlated over time.

Definition: Weak Stationarity

- A process is covariance stationary if

$$\mathbb{E}[y_t] = \mu < \infty \qquad \forall t$$

$$var[y_t] = \sigma^2 < \infty \qquad \forall t$$

$$cov(y_t, y_{t+s}) = \gamma_s \qquad \forall t, s$$

- Covariance stationarity requires that both the unconditional mean and unconditional variance are finite and do not change over time
- Covariance stationarity does not necessarily imply strict stationarity

Definition: Lag Operator

THE UNIVERSITY OF WARWICK

- Lag operator is a convention used to denote lags of a time series variable

$$Ly_t = y_{t-1}$$
$$L^2 y_t = y_{t-2}$$
$$L^j y_t = y_{t-j}$$

- Lag polynomial

$$\phi(L) = 1 + \phi L + \phi^2 L^2 + ...$$

- Difference operator (we will come back to this next week!)

$$\Delta y_t = y_t - y_{t-1} = (1 - L) y_t$$

Example: Daily Stock Returns

Taking log of S&P 500, we find that SP data is non-stationary. It has an upward trend, suggesting that the mean (level) of the index is time variant

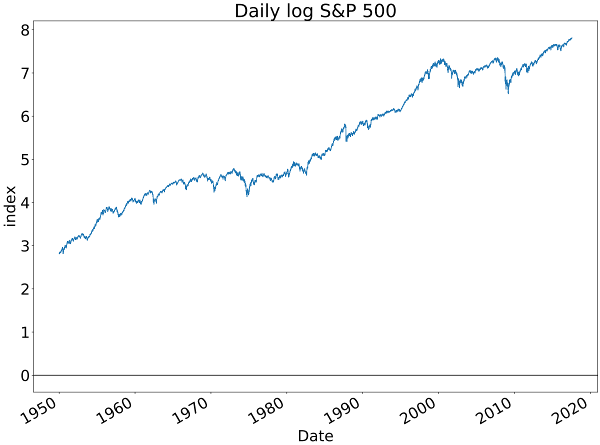

Example: Daily Stock Returns

We can construct stock returns, i.e. taking the first difference in logs.
$R_t = log(S_t) - log(S_{t-1})$

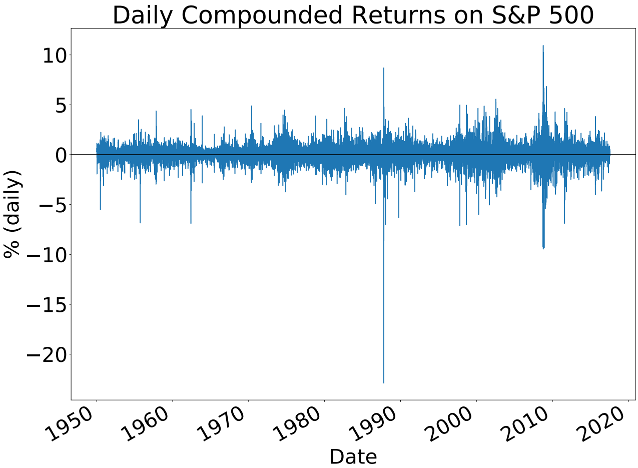

Example: Daily Stock Returns

Histograms suggest that returns can be approximated by a normal distribution, with fat tails

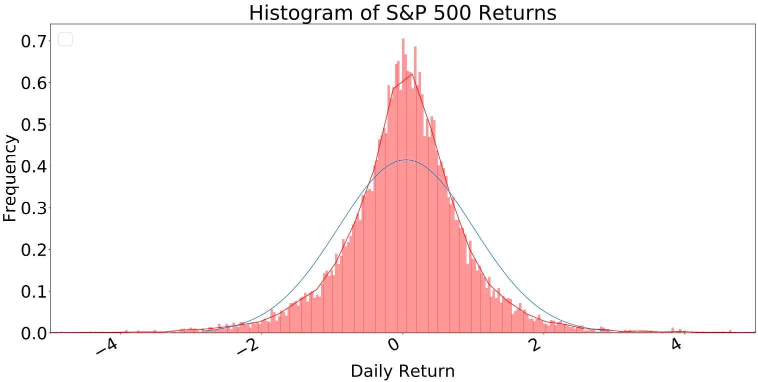

Autocorrelation Function

THE UNIVERSITY OF WARWICK

- Financial time series typically exhibit serial correlation. Knowing today's stock price helps to forecast tomorrow's stock price
- Autocovariance function

$$\gamma_j = cov(y_t, y_{t-j}) = \mathbb{E}[(y_t - \mu)(y_{t-j} - \mu)]$$

- Autocorrelation function:

$$\tau_j = \frac{\gamma_j}{\gamma_0}$$

- We can trace the autocorrelation function of a series as τ_j, and plot for $j = 0, 1, 2, ...$
- The autocorrelation plot gives us a sense of how backward looking a series is, i.e. how much of today's value is dependent on past information?

Autocorrelation Function: intuition

THE UNIVERSITY OF WARWICK

- Consider a regression of y_t on y_{t-j}

$$y_t = \phi_j y_{t-j} + \epsilon_t$$

- Assuming $\mathbb{E}[y_t] = 0$ and $var(y_t) = \sigma^2$

$$\mathbb{E}[\hat{\phi}_j] = \frac{cov(y_t, y_{t-j})}{var(y_t)} = \tau_j$$

- Autocorrelation coefficients are regression coefficients
- τ_j measures the unconditional correlation between y_t and y_{t-j}

Autocorrelation Function: testing significance

THE UNIVERSITY OF WARWICK

- Testing whether single autocorrelation coefficient is zero
- Under H_0, $\tau_j = 0$: $\sqrt{T}\hat{\tau}_j \rightarrow N(0,1)$
- A 95% confidence interval for $\hat{\tau}$ is given by $+- \frac{1.96}{\sqrt{T}}$
- Test whether the first h autocorrelation coefficients are jointly zero: **Ljung/Box test statistic**

$$Q = T(T+2)\sum_{k=1}^{h} \frac{\hat{\tau_k}^2}{T-k} \rightarrow \chi_h^2$$

Partial Autocorrelation Function

- Now consider the regressions

$$y_t = \phi_{1,1} y_{t-1} + \epsilon_t$$

$$y_t = \phi_{1,2} y_{t-1} + \phi_{2,2} y_{t-2} + \epsilon_t$$

- $\phi_{1,1}$ measures the unconditional correlation between y_t and $y_{t=1}$
- $\phi_{2,2}$ measures the correlation between y_t and y_{t-2} net of the correlation between y_t and y_{t-1}
- $\alpha_2 = \phi_{2,2}$ is the partial autocorrelation coefficient

Partial Autocorrelation Function

- The partial autocorrelation coefficient α_j measures the correlation of y_t and y_{t-j} net of the autocorrelations 1 to j-1
- Consider the recursive system

$$y_t = \phi_{1,1} y_{t-1} + \epsilon_t$$

$$y_t = \phi_{1,2} y_{t-1} + \phi_{2,2} y_{t-2} + \epsilon_t$$

$$y_t = \phi_{1,j} y_{t-1} + \phi_{2,j} y_{t-2} + ... + \phi_{j,j} y_{t-j} + \epsilon_t$$

- $\alpha_j = \phi_{j,j}$ is the measure of the marginal correlation of y_t and y_{t-j} net of the correlation between y_t and $y_{t-1}, y_{t-2}, ... y_{t-j+1}$

ACF and PACF Function of White Noise process

$$y_t = \epsilon_t \qquad \epsilon_t \sim N(0,1)$$

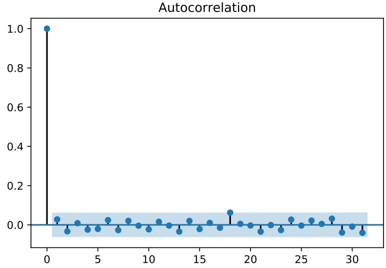

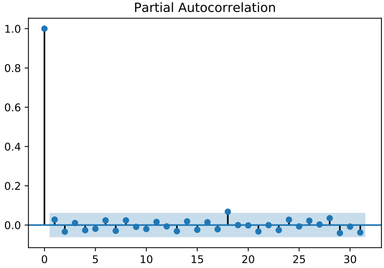

Definition: Autoregressive Process

- An Autoregressive process of order 1, or AR(1)

$$y_t = \mu + \phi y_{t-1} + \epsilon_t, \qquad \epsilon_t \sim WN(0, \sigma_\epsilon^2)$$

$$\mathbb{E}[y_t] = \frac{\mu}{1 - \phi_1}$$

- **Question**: What happens when $\phi > 1$?
- An Autoregressive process of order p, or AR(p)

$$y_t = \mu + \phi_1 y_{t-1} + \phi_2 y_{t-2} + ... + \phi_p y_{t-p} + \epsilon_t, \qquad \epsilon_t \sim WN(0, \sigma_\epsilon^2)$$

$$\mathbb{E}[y_t] = \frac{\mu}{1 - \phi_1 - \phi_2 - ... - \phi_p}$$

Conditions for Stationarity: AR

- For an AR(1) process,

$$y_t = \phi y_{t-1} + \epsilon_t \implies |\phi| < 1$$

- For an AR(p) process, the p roots of the following polynomial must lie outside the unit circle

$$(1 - \phi_1 L - ... - \phi_p L^p) y_t = \epsilon_t$$

$$(1 - \lambda_1 L)....(1 - \lambda_p L) y_t = \epsilon_t$$

$$\implies |\lambda_i| < 1 \forall i$$

Simulation of AR(1) process

We plot simulated series for $\phi = 0.5$ and $\phi = 0.95$. Note that $\phi = 0.95$ is much more persistent, whereas $\phi = 0.5$ is closer to a white noise process

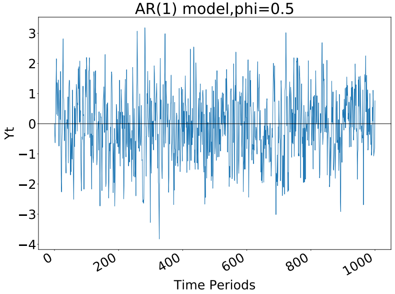

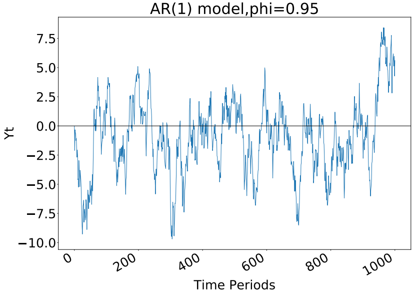

ACF and PACF Function of AR(1)

$$y_t = 0.5y_{t-1} + \epsilon_t \qquad \epsilon_t \sim N(0,1)$$

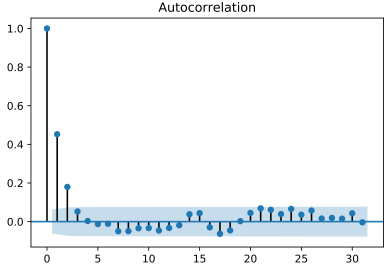

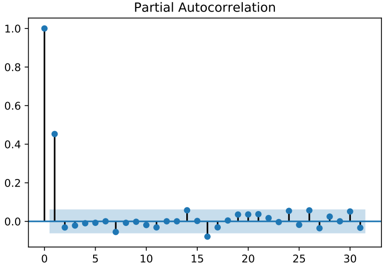

ACF and PACF Function of AR(1)

$$y_t = 0.95y_{t-1} + \epsilon_t \qquad \epsilon_t \sim N(0,1)$$

Higher persistence of $\phi = 0.95$ is seen through more gradual decay of ACF.

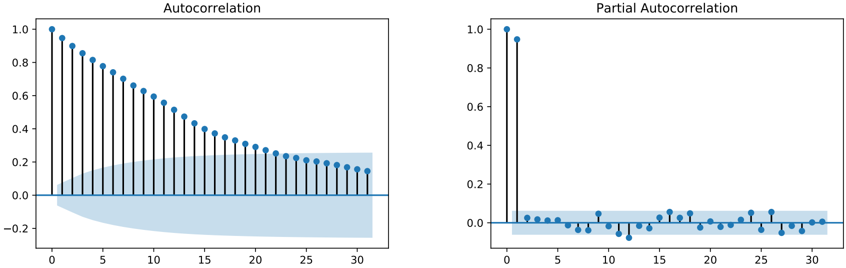

Moving Average Process

THE UNIVERSITY OF WARWICK

- A Moving Average process of order 1, or MA(1)

$$y_t = \mu + \epsilon_t, + \theta_1 \epsilon_{t-1}, \qquad \epsilon_t \sim WN(0, \sigma_\epsilon^2)$$

- MA(q) process

$$y_t = \mu + \epsilon_t + \theta_1 \epsilon_{t-1} + \theta_2 \epsilon_{t-2} + \ldots + \theta_q \epsilon_{t-q}$$

- The moving average process is covariance stationary if and only if

$$\sum_{j=0}^{q} \theta_j^2 < \infty$$

ACF and PACF Function of MA(1)

THE UNIVERSITY OF WARWICK

$$y_t = 0.5\epsilon_{t-1} + \epsilon_t \qquad \epsilon_t \sim N(0, 1)$$

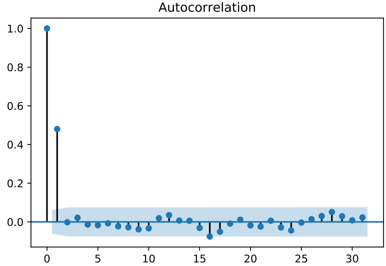

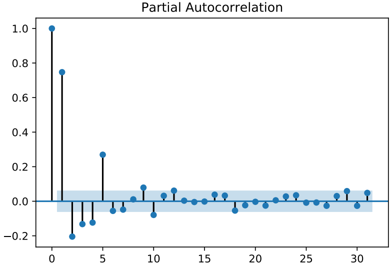

Heuristics with Autocorrelation and Partial Autocorrelation Functions

We can deduce the length of an AR and MA process by examining the ACF and PACF of a time series.

- An autoregressive process has
  1. A geometrically decaying ACF
  2. A number of significant coefficients of PACF = AR order
- A moving average process has
  1. A number of significant coefficients of ACF= MA order
  2. A geometrically decaying PACF

ACF and PACF for an AR(3)

$$y_t = 0.5y_{t-1} + 0.3y_{t-2} + 0.1y_{t-3} + \epsilon_t \qquad \epsilon_t \sim N(0,1)$$

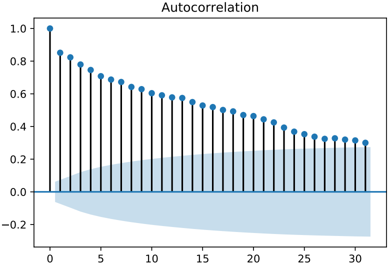

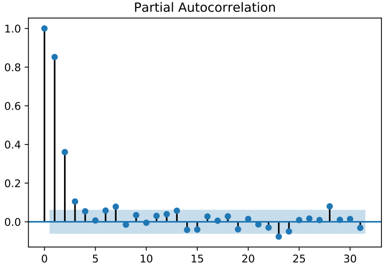

ACF and PACF for an MA(3)

$$y_t = 0.9\epsilon_{t-1} + 0.7\epsilon_{t-2} + 0.5\epsilon_{t-3} + \epsilon_t \qquad \epsilon_t \sim N(0,1)$$

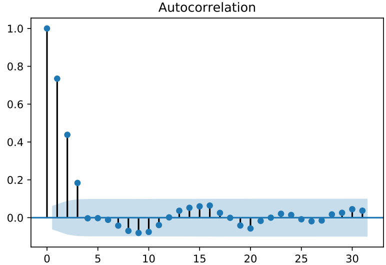

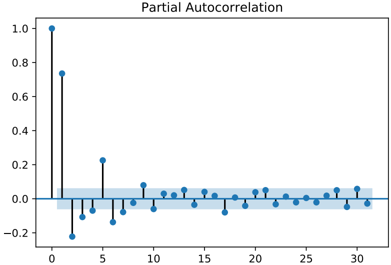

Recap and Looking Forward

- We have covered the following fundamentals of time series
  1. AR, MA processes
  2. Using ACF and PACF to deduct time series properties
  3. We also covered an important property of time series, **stationarity**.
- Please follow pre-recorded material in week 2:
  1. ARMA and the Box-Jenkins approach.
  2. Maximum Likelihood Estimation (MLE).
- ARMA is a combination of AR and MA processes.
- Box Jenkins is a procedure used to identify, estimate and check an AR/MA process.
- MLE is an alternative method to OLS–and is particularly useful to estimate parameters of a MA(q) process.
- **Next week**: Vector error correction models (VECM) use a method of **cointegration**: which takes a linear combination of two series to generate a stationary series.