TS Module 1 Time series overview

(The attached PDF file has better formatting.)

- Model building
- Time series plots

Read Section 1.1, "Examples of time series," on pages 1-8. These example introduce the book; you are not tested on them.

Read Section 1.2, "Model building strategy," on page 8. Know the three steps in the Box and Jenkins ARIMA modeling approach: specification, filling, and diagnostics (top of page 8), and the principle of parsimony (middle of page 8). The final exam tests these concepts.

The authors provide an R packages that draws the graphics in the textbook and produces the figures. Their appendix explains how to run the R functions.

You don't need to know R for this course, but it is well worth learning R for your general actuarial career. An hour or two spent downloading R and learning its basics help you in many ways, not just in the time series on-line course.

This module and the last four modules (discussing the student project) have no homework assignments. The course has 19 modules with homework assignments. You must complete 80% of these, or 15 homework assignments.

The on-line course has a final exam date, but no due dates for homework assignments or the student project. Actuarial candidates have many responsibilities, and you may not have the time to complete a homework assignment or the student project.

To avoid falling behind, do the homework assignments as you review the textbook. Send in the homework assignments (by regular mail) to the NEAS office in batches, such as modules 2-7, 8-13, and 14-20. Put each homework assignment on a separate page.

TS Module 2 Time series concepts

(The attached PDF file has better formatting.)

- Stochastic processes
- Means, variances, and covariances
- Stationarity

Read Section 2.2.1, “Means, variances, and covariances,” on pages 11-12.

Know equations 2.2.1 through 2.2.7 on pages 11 and 12. You know these equations from other work; they are definitions and formulas for variances and covariances.

Read the “Random walk” on pages 12-14. Know equations 2.2.8 through 2.2.13 on pages 12 and 13. Random walks occur in financial and actuarial work. These relations are tested on the final exam.

Read “A moving average” on pages 14-16, and know equations 2.2.14 through 2.2.16. Focus on the derivation of equation 2.2.15 on page 16. You use the same reasoning for most of the time series in this textbook.

Read Section 2.3, “Stationarity,” on pages 17-18, and know equations 2.3.1 and 2.3.2. The equations are apply to all stationary series.

Read “White noise” on page 17, and know equation 2.2.3.

Module 2 seems easy, but these are the building blocks of time series analysis. As you work through the modules, verify the derivations of the equations. The final exam problems are easy if you understand the principles. Some problems ask you to back into parameters of a time series. The intuition for the relations in this module are essential.

Module 2: Time series concepts HW

(The attached PDF file has better formatting.)

Homework assignment: equally weighted moving average

This homework assignment uses the material on pages 14-15 (“A moving average”).

Let $Y_t = 1/5 \times (\epsilon_t + \epsilon_{t-1} + \epsilon_{t-2} + \epsilon_{t-3} + \epsilon_{t-4})$ and $\sigma^2_e = 100$.

A. What is $\gamma_{t,t}$, the variance of Y_t?
B. What is $\gamma_{t,t-3}$, the covariance of Y_t and Y_{t-3}?

Write out the derivations in the format shown on page 15.

TS Module 3 Trends

(The attached PDF file has better formatting.)

- Deterministic vs stochastic trends
- Estimation of a constant mean

Read Section 3.1, “Deterministic vs stochastic trends,” on pages 27-28. The authors say that “Many authors use the word trend only for a slowly changing mean function, such as a linear time trend, and use the term seasonal component for a mean function that varies cyclically. We do not find it useful to make such distinctions here.” In practice, statisticians distinguish seasonal effects from long-term trends. Many student projects examine cycles in time series and separate them from long-term trends.

Read Section 3.2, “Estimation of a constant mean,” on pages 28-30. Know equations 3.2.1 through 3.2.5 on pages 28 and 29; you will not be tested on equation 3.2.6. The final exam tests the example on the bottom of page 28 and the formulas on page 29. It does not give complex time series or test equation 3.2.3.

The authors give many examples of time series graphs. All graphs show the R code at the bottom, and the data is in the TSA package. No knowledge of R is required for this course. But reproducing the graphs helps you understand how the time series parameters affect the sample autocorrelations and other output.

Module 3: Trends HW

(The attached PDF file has better formatting.)

Homework assignment: MA(1) Process: Variance of mean

Five MA(1) processes with 50 observations are listed below. The variance of ϵ_t is 1.

A. For each process, what is the variance of $\overline{y}$, the average of the Y observations?
B. How does the pattern of the first time series differ from that of the last time series?
C. Explain intuitively why oscillating patterns have lower variances of their means.

1. $Y_t = \mu + e_t + e_{t-1}$
2. $Y_t = \mu + e_t + \frac{1}{2} e_{t-1}$
3. $Y_t = \mu + e_t$
4. $Y_t = \mu + e_t - \frac{1}{2} e_{t-1}$
5. $Y_t = \mu + e_t - e_{t-1}$

(See page 50 of the Cryer and Chan text, Exercise 3.2)

TS Module 4 Regression methods

(The attached PDF file has better formatting.)

- Regression methods
- Interpreting regression output
- Residual analysis

Read Section 3.3, "Regression methods," on pages 30-33. If you have taken a regression course, this material is easy. We use linear regression to fit autoregressive processes. You are not responsible for cosine trends on pages 34-36.

Read Section 3.5, "Interpreting regression output," on pages 40-42. The material is covered in the regression analysis course, and it should be familiar to you.

Read Section 3.6, "Residual analysis," on pages 42-50. Focus on q-q plots and the sample autocorrelation function. Know equation 3.6.2 on page 46 and Exhibit 3.17 on page 50.

The final exam problems for the time series course do not test regression analysis. The homework assignment for this module asks you to form the confidence interval for the regression coefficient; this is basic regression.

The student project requires you to run regressions. If you use Excel, you need fit only autoregressive processes, which you can do with linear regression. If you use more sophisticated statistical software, you can fit moving average and mixed models.

If you have not taken a regression course, the time series course is difficult. You can take regression analysis and time series at the same time, since the regression needed for the time series course is taught in the first few modules of the regression analysis course.

TS Module 4: Regression methods HW

(The attached PDF file has better formatting.)

Homework assignment: Residuals

Cryer and Chan show the following Exhibit 3.10.

Exhibit 3.10 Standardized Residuals versus Fitted Values for the Temperature Seasonal Means Model

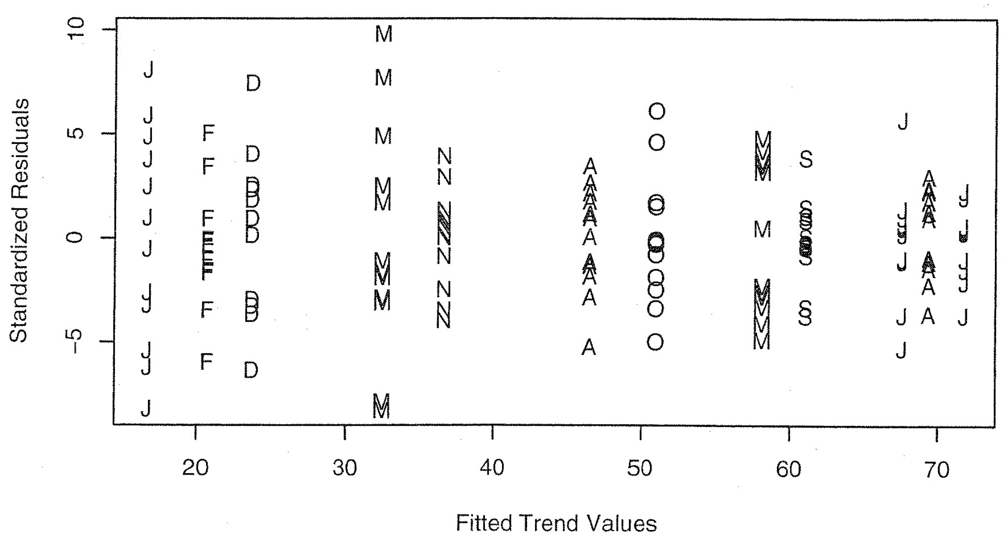

A. Temperatures vary by month. Why aren't the residuals high in July and August and low in January and February?
B. An actuary looking at this exhibit concludes that temperatures vary more in December than in November. Why is this not a reasonable conclusion?
C. A toy store has high sales in December and low sales in February. Would you expect the residuals to have a higher variance in December or February?
D. Why does the reasoning for toy sales not apply to daily temperature?

For Part C, suppose expected sales are $500,000 in December and $50,000 in February. Think of December as 10 February's packed into one month. What is the ratio of the variances if the 10 February's are independent? What is the ratio if the 10 February's are perfectly correlated?

For Part D, the daily temperature is in arbitrary units. A day in August is not like two days or ten days of February packed together.

TS Module 5 Stationary moving average processes

(The attached PDF file has better formatting.)

- General linear processes
- Moving average processes

Read Section 4.1, “General linear processes,” on pages 55-56. Know equation 4.1.6 on page 56 and its derivation on the top half of the page.

Read Section 4.2, “Moving average processes,” on page 57-58.

Note the negative sign for θ in equation 4.2.1. Don’t err by using “$+\theta$” on an exam problem. Know equations 4.2.2 on page 57; they are tested on the final exam. Know the table on the top of page 58. An exam problem may give you θ_1 and ask for ρ_1 (or *vice versa*).

As you work through the modules, keep the parameters distinct. The true parameters are unknown; we must estimate them. These are the ϕ and θ parameters. These parameters imply the autocorrelation function, the ρ parameters. We observe sample autocorrelations, or the *r* parameters, from which we back into estimates of the ϕ and θ parameters.

Pages 58 through 62 are mostly graphs. Understand what the graphs show; you need not memorize their shapes, but you must know the principles of a high or low autocorrelation.

Don’t just flip pages. The authors often show two or more graphs, with different values of a time series parameter. Understand how the parameter affects the shape of the graph.

Read pages 62-65 on moving average processes.

Know equation 4.2.3 on page 63. Work through the derivation on pages 62-63. After a few exercises, the procedure is not hard.

Know equation 4.2.4 on page 65. The final exam tests 4.2.4, but not 4.2.5.

TS Module 5: Stationary processes HW

(The attached PDF file has better formatting.)

Homework assignment: general linear process

A time series has the form $Y_t = \epsilon_t + \phi \times \epsilon_{t-1} - \phi^2 \times \epsilon_{t-2} + \phi^3 \times \epsilon_{t-3} - \ldots$

The plus and minus signs alternate. $\phi = 0.2$ and $\sigma^2_e = 9$.

A. What is γ_0, the variance of Y_t? Show the derivation.
B. What is γ_1, the covariance of Y_t and Y_{t-1}? Show the derivation.
C. What is ρ_2, the correlation of Y_t and Y_{t-2}? Show the derivation.

(Show the algebra for the derivations. One or two lines is sufficient for each part.)

TS Module 6 Stationary autoregressive processes

(The attached PDF file has better formatting.)

- Autoregressive processes
- Autocorrelation functions

Read Section 4.2, “Autoregressive processes,” on pages 66-70. Know equations 4.3.1 through 4.3.6 on pages 66 and 67. These equations are simple, but you need them for the mixed autoregressive moving average processes. The time series textbook builds rapidly. The first modules are easy, and if you understand the relations, the later modules are clear. If the skip the concepts in the early modules, the later modules are difficult.

Read the two short sections:

- “The general linear process for the AR(1) model” on pages 70-71.
- “Stationarity of an AR(1) process” on page 71.

Know equations 4.3.7 and 4.3.8 on page 70. These equations restate the results from the previous sub-section.

Read from “Second-order autoregressive process” on pages 71 through the middle of page 73, stopping after equation 4.3.15. You are not responsible for the material from “Although Equation (4.3.13) … through the end of the page.”

Review Exhibit 4.18 on page 74. The final exam gives autocorrelations at various lags and asks what type of ARMA or ARIMA process might cause them.

Read “Variance of the AR(2) model” on page 75. You need not memorize equations 4.3.19 and 4.3.20. An exam problem asking for the variance will give the equation.

Read “The ψ-coefficients of the AR(2) Model” on page 75, stopping after the explanation of equation 4.3.21. You are not responsible for the last three equations on the page, starting from “One can also show that …” until the end of the page.

Read “General autoregressive process” on pages 76-77. You use Yule-Walker equations for the student project. The final exam has simple problems using Yule-Walker equations.

The rest of the textbook builds on the concepts in the early modules. We combine moving average and autoregressive processes, with seasonality and differences (integration).

TS Module 6: Stationary autoregressive processes HW

(The attached PDF file has better formatting.)

Homework assignment: AR(2) process

An AR(2) process has ϕ_1 = 0.2 or –0.2 and ϕ_2 = ranging

- from 0.2 to 0.7 in steps of 0.1.
- from –0.2 to –0.9 in steps of –0.1.

Complete the table below, showing ρ_1 and ρ_2, the autocorrelations of lags 1 and 2. Use an Excel spreadsheet (or other software) and form the table by coding the cell formulas. Print the Excel spreadsheet and send it in as your homework assignment.

ϕ_1	ϕ_2	ρ_1	ρ_2	ϕ_1	ϕ_2	ρ_1	ρ_2
0.2	0.2			-0.2	0.2		
0.2	0.3			-0.2	0.3		
0.2	0.4			-0.2	0.4		
0.2	0.5			-0.2	0.5		
0.2	0.6			-0.2	0.6		
0.2	0.7			-0.2	0.7		
0.2	-0.2			-0.2	-0.2		
0.2	-0.3			-0.2	-0.3		
0.2	-0.4			-0.2	-0.4		
0.2	-0.5			-0.2	-0.5		
0.2	-0.6			-0.2	-0.6		
0.2	-0.7			-0.2	-0.7		
0.2	-0.8			-0.2	-0.8		
0.2	-0.9			-0.2	-0.9		

TS Module 7 Stationary mixed processes

(The attached PDF file has better formatting.)

- Mixed autoregressive moving average processes
- Invertibility

Read Section 4.4, “Mixed autoregressive moving average processes,” on pages 77-79.

Know equations 4.4.3, 4.4.4, and 4.4.5 on page 78 for the ARMA(1,1) process.

Read Section 4.5, “Invertibility,” on pages 79-81. Know the statement on page 78:

“If $|\theta| < 1$, the MA(1) model can be inverted into an infinite order autoregressive model. We say that the MA(1) model is invertible if and only if $|\theta| < 1$.”

The authors emphasize parsimony and simplicity. The previous textbook for the time series course modeled some time series with complex processes, with many moving average and autoregressive parameters. Cryer and Chan concentrate on simple models. If you model a time series with more than four or five parameters, you don’t have a good model. Most student projects conclude that an AR(1), AR(2), ARMA(1,1), or MA(1) model works best, or that first or second differences of the series can be modeled by one of these processes.

TS Module 7: stationary mixed processes HW

(The attached PDF file has better formatting.)

Homework assignment: mixed autoregressive moving average process

An ARMA(1,1) process has $\sigma^2 = 1$, $\theta_1 = 0.4$, and $\phi_1 = 0.6$.

A. What is the value of γ_0?
B. What is the value of γ_1?
C. What is the value of ρ_1?
D. What is the value of ρ_2?

TS Module 7 Stationary mixed processes

(The attached PDF file has better formatting.)

Filter Representation (ψ parameters)

We use ϕ parameters for autoregressive models and θ parameters for moving average models. These parameters have different definitions:

- The ϕ parameters relate future time series values to past time series values.
- The θ parameters relate future time series values to past residuals.

Moving average parameters have a finite memory, and autoregressive parameters have an infinite memory.

- For an MA(1) process, a random fluctuation in period T affects the time series value in period T+1 only.
- For an AR(1) process, a random fluctuation in period T affects the time series value in all future periods.

We can convert a ϕ parameter to an infinite series of θ parameters.

Illustration: A $\varphi_1 = 0.500$ is equivalent to an infinite series of θ parameters

$$\theta_1 = -0.500,\ \theta_2 = -0.250,\ \theta_3 = -0.125,\ \ldots \text{ where } \theta_j = -(0.500^j).$$

One might wonder: *Why convert a single parameter to an infinite series?*

Answer: Each θ parameter affects one future value. To estimate variances of forecasts, we convert autoregressive parameters into sets of moving average parameters. We call the new model a filter representation and we represent the new parameters by ψ variables.

Take heed: The ψ parameters have the opposite sign of the θ parameters: $\theta = 0.450$ is $\psi = -0.450$. The model is the same, but the signs of the coefficients are reversed.

$$y_t = \delta + \varepsilon_t - \theta_1\ \varepsilon_{t-1} \text{ is the same as } y_t = \delta + \varepsilon_t + \psi_1\ \varepsilon_{t-1}$$

The exercise below emphasizes the intuition.

- Once you master the intuition, the formulas are easy.
- If you memorize the formulas by rote, you forget them.

General Format

The general form of a filter representation is $y_t - \mu = \psi_0\ \varepsilon_t + \psi_1\ \varepsilon_{t-1} + \psi_2\ \varepsilon_{t-2} + \ldots$

- For a moving average model, $\mu = \delta = \theta_0$ (parameter name depends on the textbook).
- The filter representation converts the time series to a mean of zero.
  - For the values of the original time series, add the original mean.

Both moving average and autoregressive processes have filter representations.

- If the time series has only moving average parameters, $\psi_j = -\theta_j$.
- If the time series has autoregressive parameters, each ϕ_j is a series of ψ_j's.

We examine the filter representation for autoregressive models and mixed models.

Question 1.1: AR(1) Filter representation

An AR(1) model with $\varphi_1 = 0.6$ is converted to a filter representation

$$y_t - \mu = \psi_0\,\varepsilon_t + \psi_1\,\varepsilon_{t-1} + \psi_2\,\varepsilon_{t-2} + \ldots$$

A. What is ψ_0?
B. What is ψ_1?
C. What is ψ_2?
D. What is ψ_j?

Part A: ψ_0 is one for all ARIMA models. It is generally not shown in the filter representation.

Part B: If the *current error term* increases by 1 unit, the *current value* increases by one unit. The one period ahead forecast changes by $1 \times \varphi_1 = 1 \times 0.6 = 0.6$, so $\psi_1 = \varphi_1$.

Part C: If the one period ahead forecast changes by $1 \times \varphi_1 = 1 \times 0.6 = 0.6$, the two periods ahead forecast changes by $0.6 \times \varphi_1 = 0.6^2$, so $\psi_2 = \varphi_1^{\,2}$.

Part D: The same reasoning shows that $\psi_j = (\varphi_1)^j$.

Question 1.2: ARMA(1,1) filter representation

An ARMA(1,1) model with $\varphi_1 = 0.6$, $\theta_1 = 0.4$ is converted to a filter representation $y_t - \mu = \psi_0\ \varepsilon_t + \psi_1\ \varepsilon_{t-1} + \psi_2\ \varepsilon_{t-2} + \ldots$

A. What is ψ_0?
B. What is ψ_1?
C. What is ψ_2?
D. What is ψ_j?

Part A: ψ_0 is one for all ARIMA models.

Part B: Suppose the *current error term* increases by 1 unit.

- The moving average part of the ARMA process changes the forecast by $1 \times -\theta_1 = 1 \times -0.4 = -0.4$.
- If the *current error term* increases by one unit, the *current value* increases by one unit.
- The autoregressive part of the ARMA process changes the forecast by $1 \times \varphi_1 = 1 \times 0.6 = 0.6$.

The combined change in the forecast is $-0.4 + 0.6 = 0.2$. The change in the one period ahead forecast is $\varphi_1 - \theta_1$.

Take heed: The negative sign reflects the convention that moving average parameters are the negative of the moving average coefficients.

Part C: The one period ahead forecast increases 0.2 units (the result in Part B), so the two periods ahead forecast increases $0.2 \times \varphi_1 = 0.2 \times 0.6 = 0.12$ units.

Part D: Repeating the reasoning above gives $\psi_j = 0.2^{j-1} \times 0.6$.

Question 1.3: ARMA(2,1) filter representation

An ARMA(2,1) model with $\varphi_1 = 0.6$, $\varphi_2 = -0.3$, $\theta_1 = 0.4$ is converted to a filter representation $y_t - \mu = \psi_0\ \varepsilon_t + \psi_1\ \varepsilon_{t-1} + \psi_2\ \varepsilon_{t-2} + \dots$

A. What is ψ_0?
B. What is ψ_1?
C. What is ψ_2?
D. What is ψ_3?

Part A: ψ_0 is one for all ARIMA models. It is generally not shown in the filter representation.

Part B: Suppose the *current error term* increases by 1 unit.

- The moving average part of the ARMA process changes the forecast by $1 \times -\theta_1 = 1 \times -0.4 = -0.4$.
- If the *current error term* increases by one unit, the *current value* increases by one unit.
- The autoregressive part of the ARMA process changes the forecast by $1 \times \varphi_1 = 1 \times 0.6 = 0.6$.

The combined change in the forecast is $-0.4 + 0.6 = 0.2$. The change in the one period ahead forecast is $\varphi_1 - \theta_1$.

Part C: A 1 unit increase in the current error term increases the two periods ahead forecast two ways in this exercise:

- The one period ahead forecast increases 0.2 units (the result in Part A), so the two periods ahead forecast increases $0.2 \times \varphi_1 = 0.2 \times 0.6 = 0.12$ units.
- The *current value* increases 1 unit, so the φ_2 parameter causes the two periods ahead forecast to increase -0.3 units.

The change in the two periods ahead forecast is $0.12 - 0.3 = -0.18$ units, so $\psi_2 = -0.18$.

Take heed: The θ_1 parameter does not affect forecasts two or more periods ahead: an MA(1) process has a memory of one period. In contrast, an AR(1) process has an infinite memory. The φ_1 parameter affects all future forecasts.

Part D: If the number of periods ahead is greater than the maximum of *p* and *q* (2 and 1 in this exercise), the direct effects of the parameters is zero. We compute the combined effects: $\psi_3 = \varphi_1 \times \psi_2 + \varphi_2 \times \psi_1 = 0.6 \times -0.18 - 0.3 \times 0.2 = -0.168$.

Exercise 1.4: AR(2) Model

An AR(2) model $y_t - \mu = \varphi_1 (y_{t-1} - \mu) + \varphi_2 (y_{t-2} - \mu) + \varepsilon_t$ has $\varphi_1 = 0.4$ and $\varphi_2 = -0.5$. We convert this model to an infinite moving average model, or the filter representation

$$y_t - \mu = \varepsilon_t + \psi_1 \varepsilon_{t-1} + \psi_2 \varepsilon_{t-2} + \dots$$

A. What is ψ_1?
B. What is ψ_2?
C. What is ψ_3?

Part A: Suppose the *residual* in Period *T* increases one unit. We examine the effect on the *value* in Period T+1.

- The *current value* increases 1 unit.
- The φ_1 coefficient causes next period's value to increase 0.4 units.

Part B: Suppose the *residual* in Period *T* increases one unit. We examine the effect on the *value* in Period T+2.

- The current value increases 1 unit.
- The φ_2 coefficient causes the two periods ahead value to increase –0.5 units.
- The φ_1 coefficient has a two step effect. It causes next period's value to increase 0.4 units and the value in the following period to increase 0.4 × 0.4 = 0.16 units.

The net change in the two periods ahead value is –0.5 + 0.16 = –0.34.

- The AR(2) formula is: $\psi_2 = \varphi_1^2 + \varphi_2 = 0.4^2 - 0.5 = -0.340$.
- The explanation above is the intuition for this formula.

Part C: We use all permutations: $\varphi_1 \times \varphi_1 \times \varphi_1$, $\varphi_1 \times \varphi_2$, and $\varphi_2 \times \varphi_1$ =

$$0.4^3 + 2 \times 0.4 \times -0.5 = -0.336$$

For this part of the exercise, the subscript of ψ is greater than the order of the ARMA process. Instead of working out all the permutations, we multiply each ϕ_j coefficient by the ψ_{k-j} coefficient. We multiply ϕ_1 by ψ_2 and ϕ_2 by ψ_1 = $0.4 \times -0.34 + -0.5 \times 0.4 = -0.336$

Take heed: The formulas are simple permutations.

- Focus on the intuition, not on memorizing formulas.
- The final exam problems can all be solved with first principles.

TS Module 8 Non-stationary time series basics

(The attached PDF file has better formatting.)

- Variable transformations
- Stationarity through differencing

Read Section 5.1, “Stationarity through differencing,” on pages 88-92. Know equation 5.1.10 on page 90 and its derivation. Distinguish between σ^2_ε and σ^2_e in this equation.

Read again the last paragraph on page 90 and review Exhibit 5.4 on page 91. Most actuarial time series are not stationary. For your student project, you take first and second differences, and you might also take logarithms. The homework assignment shows how a loss cost trend is made stationary by logarithms and first differences.

Cryer and Chan do not stress changes in the time series over time. The authors know how to judge if the parameters are stable, but they keep the statistics at a first year level.

For the student project, ask yourself whether the time series itself has changed. The module on the interest rate time series on the NEAS web site stresses the three interest rate eras affecting the time series.

TS Module 8: Non-stationary time series basics HW

(The attached PDF file has better formatting.)

Homework assignment: Stationarity through differencing and logarithms

Automobile liability claim severities have a geometric trend of +8% per annum.
The average claim severity in year t is the average claim severity in year t-1, plus or minus a random error term.

A. Is the time series of average claim severities stationary?
B. Is the first difference of this time series stationary?
C. Is the second difference of this time series stationary?
D. Is the logarithm of this time series stationary?
E. What transformation makes the time series stationary?

TS Module 9 Non-stationary ARIMA time series

(The attached PDF file has better formatting.)

- ARIMA process
- Constant terms in ARIMA models

Read Section 5.2, “ARIMA models,” on pages 92-97. Read the material for the concepts; the final exam does not test the equations. Know how taking first or second differences makes the process stationary.

For actuarial time series, such as loss cost trends, inflation indices, stock prices, and dollar values, first take logarithms and then take first differences. The authors mention this, but it is easy to forget.

Read Section 5.3, “Constant terms in ARIMA models, on pages 97-98.

Know equations 5.3.16 and 5.3.17 on the bottom on page 97; they are tested frequently on the final exam.

- Only the ϕ_j terms are in the denominator of the expression for μ.
- The constant θ_0 term is in the numerator of the expression for μ.

The previous textbook used for this on-line course used δ instead of θ_0. Some practice problems on the discussion forum still have δ. Cryer and Chan use θ instead of θ_1 for an MA(1) process and ϕ instead of ϕ_1 for an AR(1) process. The final exam problems use the notation if the Cryer and Chan textbook, but some practice problems have other notation.

Read Section 5.4, “Other transformations,” on page 98-100. Know equation 5.4.3 in the middle of page 99. Many actuarial time series are percentage changes.

Power transformations on pages 101-102 are covered in the regression analysis course. They are not tested in the time series course. But they are needed for proper modeling of actuarial time series. If you have not taken the regression analysis course with the John Fox textbook, read these two pages.

TS Module 9: Non-stationary ARIMA time series HW

(The attached PDF file has better formatting.)

Homework assignment: Non-stationary autoregressive process

A time series $Y_t = \beta \times Y_{t-1} + \epsilon_t$ has σ_ε^2 = 3, where *k* is a constant. (The textbook has β = 3.)

A. What is the variance of Y_t as a function of β and *t*?
B. What is $\rho(y_t, y_{t-k})$ as a function of β, k, and *t*?

See equations 5.1.4 and 5.1.5 on page 89. Show the derivations for an arbitrary β.

TS Module 10 Autocorrelation functions

(The attached PDF file has better formatting.)

- Sample autocorrelation function
- Partial autocorrelation function

Read the introduction to Chapter 6, “Model specification,” on page 109. Know the three bullet points at the top of the page; they are tested on the final exam and you may structure your student project in three steps.

Read Section 6.1, “Sample autocorrelation function,” on pages 109-112. Know equation 6.1.1 on the bottom of page 109.

The denominator of the sample autocorrelation function has n terms and the numerator has n-k terms. If we did not adjust in this fashion, the sample autocorrelation function for a white noise process would increase (in absolute value) as the lag increases.

The discussion forum for the time series student project has an Excel worksheet that shows why we need to adjust the number of terms in the numerator and denominator.

The final exam problems may give a set of values and ask for the sample autocorrelations of lag 1, 2, or 3, as the homework assignment does. Make sure you use the proper number of terms in the numerator and denominator.

Know equation 6.1.3 on the bottom of page 110. You will not be tested on equations 6.1.2 or 6.1.4.

Know equations 6.1.5 and 6.1.6 on the top of page 111. You are not responsible for equations 6.1.7 and 6.1.8 in the middle of page 111.

Know the last paragraph of this section on page 112.

The discussion forum for the time series student project has an Excel worksheet with a VBA macro that forms correlograms. See the project template for daily temperature, which forms a correlogram from 100 years of daily temperature readings. The large number of computations may slow down your computer if you have an old model. If you use statistical software with functions for sample autocorrelations, the built-in code is more efficient.

Read Section 6.1, “Partial autocorrelation function,” on pages 112-114.

Know equation 6.2.3 on page 113 and equations 6.2.4, 6.2.5, and 6.2.6 on page 114.

You are not responsible for pages 115 through the end of this section on page 117.

TS Module 10: autocorrelation functions HW

(The attached PDF file has better formatting.)

Homework assignment: Sample autocorrelations

A time series has ten elements: {10, 8, 9, 11, 13, 12, 10, 8, 7, 12}.

A. What is the sample autocorrelation of lag 1?
B. What is the sample autocorrelation of lag 2?
C. What is the sample autocorrelation of lag 3?

Show the derivations with a table like the one below. Remember to use the proper number of terms in the numerator, depending on the lag.

Entry	*Entry*	*Deviation*	*Deviation Squared*	*Cross Product Lag1*	*Cross Product Lag2*	*Cross Product Lag3*
1	10					
2	8					
3	9					
4	11					
5	13					
6	12					
7	10					
8	8					
9	7					
10	12					
Avg/tot						
Autocorr						

TS Module 10 Sample autocorrelation functions practice problems

(The attached PDF file has better formatting.)

Question 1.1: Sample Autocorrelation Function

The sample autocorrelation of lag $k \approx (-0.366)^k$ for *all* $k > 1$, and the sample autocorrelation of lag 1 is –0.900. The time series is most likely which of the following choices?

A. AR(1)
B. MA(1)
C. ARMA(1,1)
D. ARIMA(1,1,1)
E. A random walk

Answer 1.1: C

A stationary autoregressive model has geometrically declining autocorrelations for lags more than its order. If the order is p, the lags for p+1 are higher are geometrically declining. This is true here, so we presume an AR(1) process.

If the time series is AR(1), the sample autocorrelation for lag 1 should be about –0.366. It is –0.900, so we assume the series also has a moving average component of order 1.

Question 1.2: Sample Autocorrelation Function

For a time series of 1,600 observations, the sample autocorrelation function of lag *k* is $\approx 0.366 \times 1.2^{-k}$ for $k < 4$. For $k \geq 4$, the sample autocorrelations are normally distributed with a mean of zero and a standard deviation of 2.5%. The time series is probably

A. Stationary and Autoregressive of order 3
B. Stationary and Moving Average of order 3
C. Non-stationary
D. A random walk with a drift for three periods
E. A combination of stationary autoregressive of order 3 and a white noise process

Answer 1.2: B

Statement B: For $k \geq 4$ and 1,600 observations, the sample autocorrelations are normally distributed with a mean of zero and a standard deviation of 2.5%; these are the sample autocorrelations of a white noise process. A moving average time series has sample autocorrelations that drop off to a white noise process after its order (3 in this problem).

Statement A: An autoregressive process has geometrically declining autocorrelations for lags greater than its order.

Statements C and D: A non-stationary time series would not have autocorrelations that drop off to a random walk after 3 periods. A random walk is not stationary.

Statement E: A stochastic time series has white noise built in; adding white noise doesn't change anything.

Question 1.3: Sample Autocorrelation Function

If the sample autocorrelations for a time series of 1,600 observations for the first five lags are 0.461, 0.021, –0.017, 0.025, and –0.009, the time series is most likely which of the following choices?

A. AR(1) with $\varphi_1 \approx 0.45$
B. MA(1)
C. ARMA(1,1) with $\varphi_1 \approx 0.45$
D. ARIMA(1,1,1) with $\varphi_1 \approx 0.45$
E. A random walk

Answer 1.3: B

The sample autocorrelations decline to zero after the first lag, with random fluctuations within the white noise limits. The process is presumably moving average of order 1.

The process could also have an AR(1) parameter with $\varphi_1 < 0.15$, but we have no reason to assume an autoregressive parameter. If $\varphi_1 \approx 0.45$, the sample autocorrelation of lag 2 should be significantly more than zero.

Question 1.4: Covariances

We examine γ_0, γ_1, and γ_2, the covariances from a stationary time series for lags of 0, 1, and 2. Which of the following is true?

γ_0 is the variance, which is constant for a stationary time series, so the autocorrelations are the covariances divided by the variance. The autocorrelations have a maximum absolute value of one, and the variance is positive.

A. $\gamma_0 \geq \gamma_1$
B. $\gamma_1 \geq \gamma_2$
C. $\gamma_2 \geq \gamma_1$
D. $\gamma_1 + \gamma_2 \geq \gamma_0$
E. If $\gamma_1 \geq 0$, $\gamma_2 \geq 0$

Answer 1.4: A

The covariances of the time series can increase or decrease with the lag.

Illustration: For an MA(q) process with $\theta_j = 0$ for $1 \leq j \leq q-1$, the covariances are 0 for lags of 1 through q-1 but non-zero for a lag of q.

The variances of all elements of a stationary time series are the same, so none of the covariances can exceed the variance.

All five choices can be true. Only choice A is always true.

Question 1.5: Sample Autocorrelation Function

A time series is {7, 9, 10, 13, 9, 11, 8, 11, 12, 10, 8, 12}. What is $\hat{\rho}_2$,
the *sample* autocorrelation function at lag 2? Use the data in the table below.

T	*Y_t*	*(3)*	*(4)*	*(5)*
1	8	-4	4	16
2	12	0	0	0
3	11	-1	2	1
4	16	4	8	16
5	10	-2	6	4
6	14	2	4	4
7	9	-3	-3	9
8	14	2	2	4
9	13	1	-3	1
10	13	1	3	1
11	9	-3		9
12	15	3		9
Total	144	0.00	20.0	74.0

- Column 3 is $y_t - \bar{y}$
- Column 4 is $(y_t - \bar{y}) \times (y_{t+2}\ \bar{y})$
- Column 5 is $(y_t - \bar{y})^2$

A. –0.2727
B. –0.1200
C. +0.1212
D. +0.2700
E. +1.1200

Answer 1.5: D

$\bar{y} = \sum y_t / 12 = 144/12 = 12.000$

$\hat{\rho}_2 = [\sum (y_t - \bar{y}) \times (y_{t+2} - \bar{y})] / \sum (y_t - \bar{y})^2 = 20 / 74 = 0.270$

TS Module 11 Simulated and actual time series

(The attached PDF file has better formatting.)

- Specification of simulated time series
- Specification of actual time series

Read Section 6.3, “Specification of simulated time series,” on pages 117-124. The text shows how to use correlograms to identify the time series. You use these tools for your student project.

Read Section 6.4, “Non-stationarity,” on pages 124-128.

Know the problems of over-differencing on pages 126-128. Some student projects make this error at first. A candidate may feel that the correlogram does not approach zero fast enough and takes a second difference. Sometimes this is correct; more often it is wrong.

Be sure that differencing is warranted in your project. If you take a second difference, say why it is justified. The time series may be a combination of two ARIMA(1,1,0) processes with different values for μ or ϕ. Takings second differences obscures the true parameters.

For your student project, consider taking logarithms before first differences. If you have a long enough time series, such as average claim severities in nominal dollars for forty years, you see the exponential curve. For a short time series, such as twelve months of daily stock prices, you won’t see the exponential pattern in the sample points.

The final exam does not test the Dickey-Fuller Unit-Root test on pages 128-130. You may want to use this tool in your student project, though. It provides a quantitative test for non-stationarity that you may use in addition to graphic analsis.

Read Section 6.6, “Specification of actual time series,” on pages 133-140. The final exam does not test these time series, but this section helps you in your student project.

TS Module 11: simulated and actual time series HW

(The attached PDF file has better formatting)

Homework assignment: Partial autocorrelations

[Partial autocorrelations are covered in Module 10, along with sample autocorrelations.]

- A stationary ARMA process has ρ_2 = 0.20.
- ρ_1 ranges from 0.2 to 0.7 in units of 0.1.

A. Graph the partial autocorrelation of lag 2 (ϕ_{22}) as a function of ρ_1.
B. Explain why the partial autocorrelation is positive for low ρ_1 and negative for high ρ_1.

TS Module 12 Parameter estimation method of moments

(The attached PDF file has better formatting.)

- Method of moments
- Autoregressive, moving average, and mixed models

Read Section 7.1, “Method of moments,” on pages 149-154. Know equation 7.1.1 on the bottom of page 149 and equations 7.1.2 and 7.1.3 on the top of page 150. An exam problem may give the sample autocorrelations for the first two lags of an AR(2) process and ask for ϕ_1 and ϕ_2, which you solve using equation 7.1.2.

The final exam does not ask Yule-Walker equations for processes not illustrated in the text. But know how to use the method of moments for your student project.

- Use linear regression for autoregressive processes with Excel’s regression add-in.
- If you do not have other statistical software, you must use Yule-Walker equations for moving average and mixed processes.

Moving average models: Know equations 7.1.4 on the bottom of page 150. The final exam gives the sample autocorrelation for an MA(1) process and asks for θ_1.

Know equations 7.1.5 and 7.1.6 in the middle of page 151. The final exam will give ρ_1 and ρ_2 for an ARMA(1,1) process and ask for the estimates of ϕ and θ.

You are not responsible for “Estimates of the noise variance” on pages 151-152: equations 7.1.8 and 7.1.9 on the bottom of page 151 and equation 7.1.10 on the top of page 152.

Read “Numerical examples” on pages 152-154. These are illustrations; you are not tested on the equations in this section.

TS Module 12: Parameter estimation Yule-Walker equations

(The attached PDF file has better formatting.)

Use the Yule-Walker equations to derive initial estimates of the ARMA coefficients. Know how to solve the Yule-Walker equations for AR(1), AR(2), and MA(1) processes.

- A student project might also use Yule-Walker equations for MA(2) and ARMA models.
- For the final exam, focus on the equations for AR(1), AR(2), and MA(1) models.

Exercise 1.1: MA(1) model and Yule-Walker equations

An MA(1) model has an estimated ρ_1 of –0.35. What is the Yule-Walker initial estimate for θ_1 if it lies between –1 and +1?

Solution 1.1: An MA(1) model has $\rho_1 = \dfrac{-\theta_1}{\left(1+\theta_1^2\right)}$.

We invert the equation to get $\theta_1 = \dfrac{-1 \pm \sqrt{1-4\rho_1^2}}{2\rho_1}$

We compute $(-1 + (1 - 4 \times 035^2)^{0.5}) / (2 \times -0.35) = 0.408$

The final exam uses multiple choice questions. To avoid arithmetic errors, after solving the problem, check that it gives the correct autocorrelation.

The table below shows selected MA(1) values for ρ_1 and θ_1. Note several items:

For a given value of ρ_1, two values of θ_1 may solve the Yule-Walker equation. The exam problem may specify bounds for θ_1, such as an absolute value less than one. The textbook expresses this as the MA(1) model is invertible.

For an invertible MA(1) model, ρ_1 and θ_1 have opposite signs, reflecting the sign convention for the moving average parameter.

Know several limiting cases.

- As $\theta_1 \rightarrow$ zero, $\rho_1 \rightarrow$ zero
- As $\theta_1 \rightarrow$ one, $\rho_1 \rightarrow$ negative one half (–0.5)
- As $\theta_1 \rightarrow$ infinity, $\rho_1 \rightarrow$ negative one (–1)

θ_1	ρ_1	θ_1	ρ_1
0.1	-0.0990	-0.1000	0.0990
0.2	-0.1923	-0.2000	0.1923
0.3	-0.2752	-0.3000	0.2752
0.4	-0.3448	-0.4000	0.3448
0.5	-0.4000	-0.5000	0.4000
0.6	-0.4412	-0.6000	0.4412
0.7	-0.4698	-0.7000	0.4698
0.8	-0.4878	-0.8000	0.4878
0.9	-0.4972	-0.9000	0.4972

TS Module 12: Parameter estimation method of moments HW

(The attached PDF file has better formatting.)

Homework assignment: Method of moments

An ARMA(1,1) process has $r_1 = -0.25$ and $r_2 = -0.125$.

A. What is ϕ, the autoregressive parameter?
B. What is θ, the moving average parameter?

TS Module 13 Parameter estimation least squares

(The attached PDF file has better formatting.)

- Least squares estimation
- Autoregressive models

Read Section 7.2, “Least squares estimation,” on pages 154-156. Know how to solve for the parameters of an autoregressive process using least squares estimation. You are not responsible for nonlinear regression (with numerical methods) for processes having a moving average component.

Know equations 7.2.5 on page 155 and equations 7.2.9 and 7.2.10 in the middle of page 156. These are the same results as given by the Yule-Walker equations. You need not memorize the exact formulas for small samples. If an exam problem uses these formulas, it gives them to you on the exam.

TS Module 13: Parameter estimation least squares HW

(The attached PDF file has better formatting.)

Homework assignment: Estimating parameters by regression

An AR(1) process has the following values:

0.44 1.05 0.62 0.72 1.08 1.24 1.42 1.35 1.50

A. Estimate the parameter ϕ by regression analysis.
B. What are 95% confidence intervals for the value of ϕ?
C. You initially believed that ϕ is 50%. Should you reject this assumption?

The time series course does not teach regression analysis. You are assumed to know how to run a regression analysis, and you must run regressions for the student project.

Use the Excel *REGRESSION* add-in. The 95% confidence interval is the estimated β ± the *t*-value × the standard error of β. The *t*-value depends on the number of observations. Excel has a built-in function giving the *t*-value for a sample of N observations.

TS Module 14 Model diagnostics

(The attached PDF file has better formatting.)

- Residual analysis
- q-q plots

Read Section 8.1, “Residual analysis,” on pages 174-179. Know equation 8.1.3 on page 175.

We test if residuals lie within a 95% confidence interval by plotting standardized residuals (residuals divided by their standard deviation). See Exhibit 8.1 on page 176, Exhibit 8.2 on page 177, and Exhibit 8.3 on page 178.

We test if the residuals are normally distributed with q-q plots. See Exhibits 8.4 and 8.5 on page 179 and Exhibit 8.6 on page 180.

Standardized residuals and q-q plots are covered in the regression analysis course. You use these techniques in the student project.

The final exam does not test standardized residuals. But it may give a q-q plot (quantile comparison plot) and ask what it means. The homework assignment shows the types of q-q plots that may be tested.

TS Module 14: Model diagnostics HW

(The **two** attached PDF files have better formatting.)

Homework assignment: quantile comparison (q-q) plots

Quantile comparison plots are explained in the regression analysis on-line course, and they are used also in the time series course. If this homework assignment is difficult, review the module of quantile comparison plots in the regression analysis course. (Module 3, "Quantile comparison plots," on pages 34-37, especially Figures 3.8 and 3.9; you can search for *quantile comparison plots* or *q-q plots* on the internet to see several examples.)

The four figures below show quantile comparison plots for four distributions. For each one

A. Is the distribution symmetric, right skewed, or left skewed? Explain how the quantile comparison plot shows this.
B. If the distribution is symmetric, is it heavy tailed or thin tailed? Explain how the quantile comparison plot shows this.

Quantile comparison plots are a useful tool for actuarial work, so it is worth knowing how to use them. For your student project, you may test if the residuals of an ARIMA process are normally distributed by forming a quantile comparison plot.

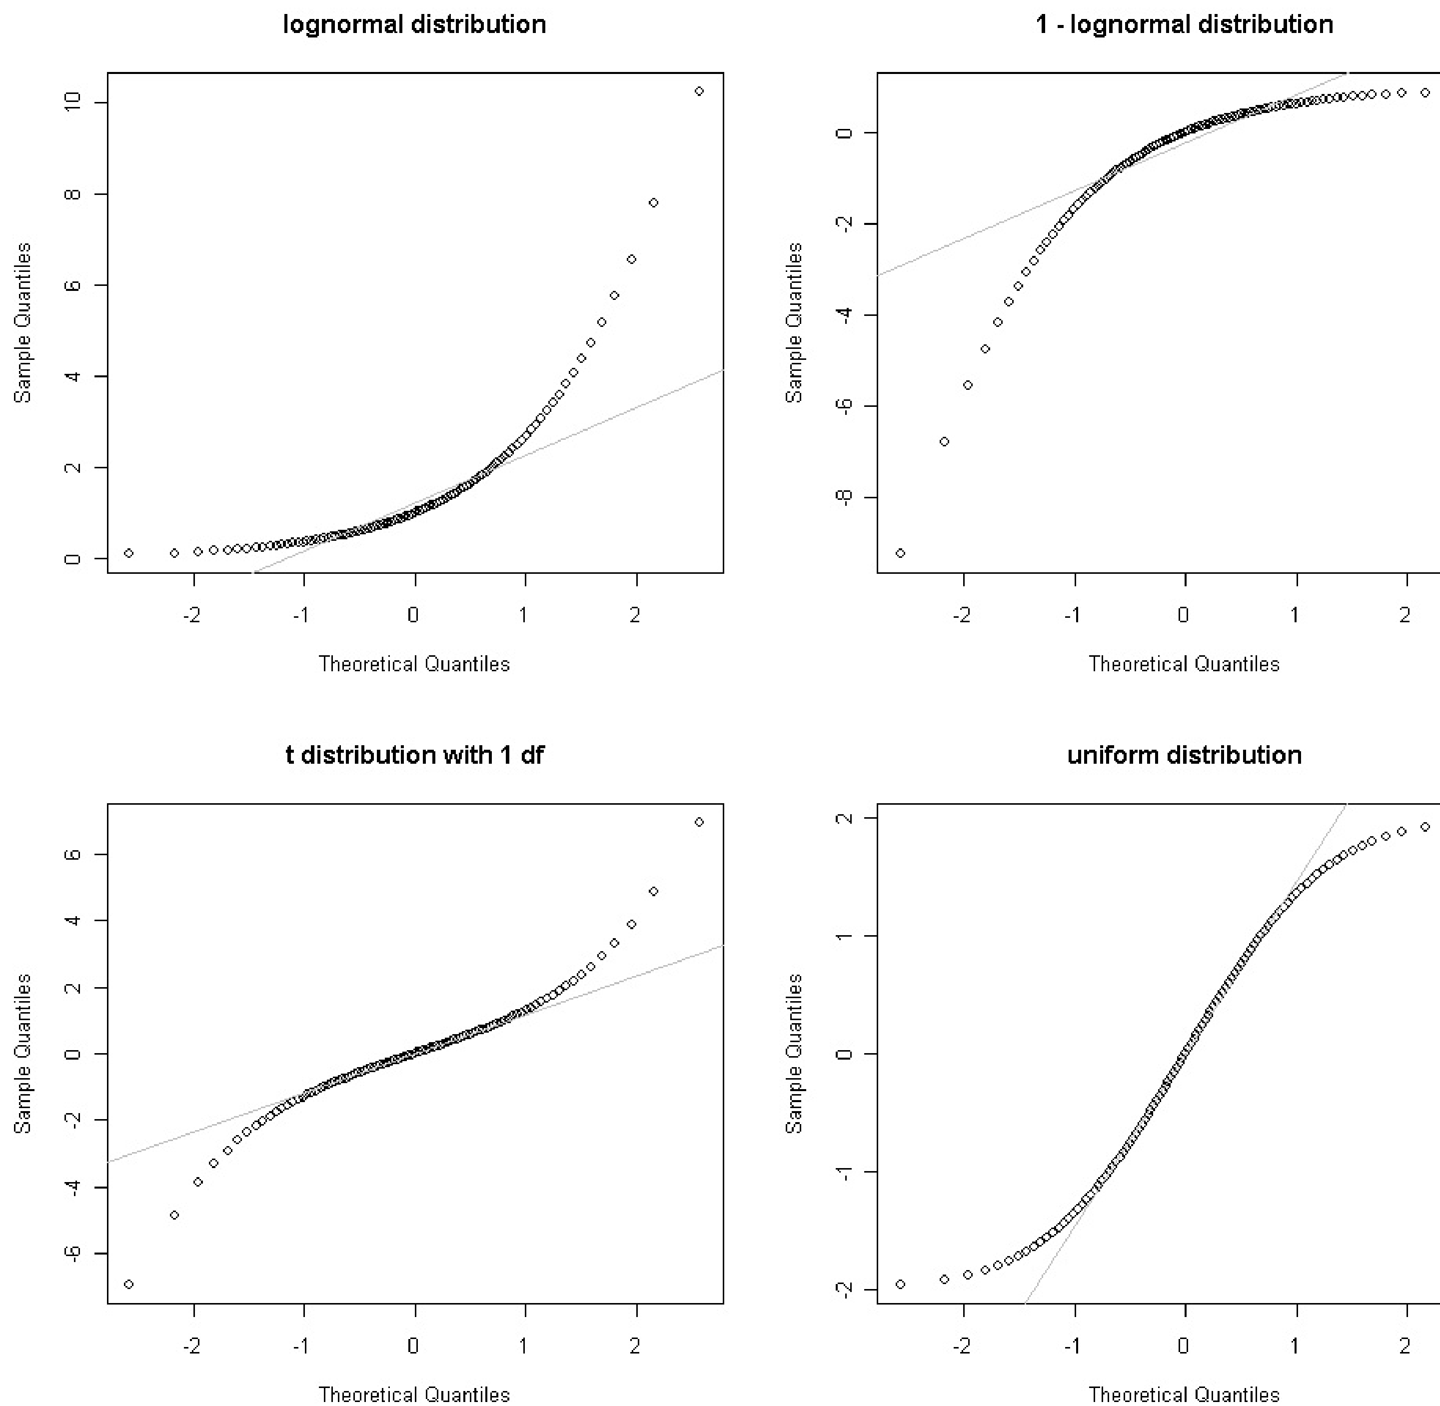

lognormal distribution
1 - lognormal distribution
t distribution with 1 df
uniform distribution
Sample Quantiles
Theoretical Quantiles

TS Module 14: Model diagnostics HW

(The attached PDF file has better formatting.)

Homework assignment: quantile comparison (q-q) plots

Quantile comparison plots are explained in the regression analysis on-line course, and they are used also in the time series course. If this homework assignment is difficult, review the module of quantile comparison plots in the regression analysis course. (Module 3, "Quantile comparison plots," on pages 34-37, especially Figures 3.8 and 3.9; you can search for *quantile comparison plots* or *q-q plots* on the internet to see several examples.)

The four figures below show quantile comparison plots for four distributions. For each one

A. Is the distribution symmetric, right skewed, or left skewed? Explain how the quantile comparison plot shows this.
B. If the distribution is symmetric, is it heavy tailed or thin tailed? Explain how the quantile comparison plot shows this.

Quantile comparison plots are a useful tool for actuarial work, so it is worth knowing how to use them. For your student project, you may test if the residuals of an ARIMA process are normally distributed by forming a quantile comparison plot.

TS Module 14 Model Diagnostics intuition

(The attached PDF file has better formatting.)

Diagnostic checking is especially important for the *student project* , for which you estimate one or more ARIMA processes and check which one is best. You use several methods:

In-sample tests examine the Box-Pierce Q statistic or the Ljung-Box Q statistic to see if the residuals of the ARIMA model are a white noise process.

Out-of-sample tests examine the mean squared error of the ARIMA models to see which one is the best predictor.

Diagnostic testing is both art and science. Random fluctuations and changes in the model parameters over time force us to rely on judgment in many cases.

The final exam tests objective items.

Numerical problems test the variance or standard deviation of a white noise process, the value of Bartlett's test, or the computation of the Box-Pierce Q statistic and Ljung-Box Q statistic.

Multiple choice true-false questions test the principles of diagnostic checking.

We review several topics that are often tested on the final exam.

We fit an ARMA(p,q) model to a time series and check if the model is specified correctly.

A. We compare the autocorrelation function for the simulated series (the time series generated by the model) with the *sample* autocorrelation function of the original series. If the two series differ materially, the ARMA process may not be correctly specified.

B. If the autocorrelation function of the ARMA process and the sample autocorrelation function of the original time series are similar, we compute the residuals of the model. We often assume the error terms before the first observed value are zero and the values before the first observed value are the mean.

C. If the model is correctly specified, the residuals should resemble a white noise process.

D. If the model is correctly specified, the residual autocorrelations are uncorrelated, normally distributed random variables with mean 0 and standard deviation 1/T, where T is the number of observations in the time series.

E. The Q statistic, where Q =

, is approximately distributed as chi-square with (K – p – q) degrees of freedom. Cryer and Chan use a more exact statistic. The final exam tests both the (unadjusted) Box-Pierce Q statistic and the Ljung-Box Q statistic.

Statement A: Suppose the sample autocorrelations are 0.800, 0.650, 0.500, 0.400, 0.350, and 0.250 for the first six lags and we try to fit an MA(2) model.

Use the Yule-Walker equations or nonlinear regression to estimate è$_0$, è$_1$, and è$_2$.

Compare the autocorrelation function for the model with the sample autocorrelation function of the original time series.

The autocorrelation function for the MA(2) model drops to zero after the second lag, but the sample autocorrelation function of the original time series does not drop to zero. We infer that the time series is not an MA(2) model.

This example is simple. Given the sample autocorrelations, we should not even have tried an MA(2) model. Other example are more complex.

This comparison does not have strict rules. No ARIMA process fits perfectly, and selecting the best model is both art and science. In a statistical project, we overlay the correlogram with the autocorrelation function of the model being tested, and we judge if the differences are random fluctuations.

Distinguish the two sides of this comparison:

The sample autocorrelations are empirical data. They do not depend on the model.

The autocorrelations reflect the fitted process. You select a model, fit parameters, and derive the autocorrelations.

There are different functions; be sure to differentiate them.

The sample autocorrelations are distorted by random fluctuations. They are estimated from empirical data, with adjustments for the degrees of freedom at the later lags. This adjustment is built into the sample autocorrelation function.

The autocorrelations are derived algebraically. If we know the exact parameters of the ARIMA process, we know the exact autocorrelations.

The time series is stochastic. The model may be correct, but random fluctuations cause unusual sample autocorrelations. Know how to form confidence intervals.

Statement B: Residuals, time series, and fitted processes.

Residuals are discussed so often it seems that time series have inherent residuals.

The residuals of the time series are not known until we specify a model. A time series with no model has no residuals.

The ARIMA process by itself has an error term, not residuals. The realization of the ARIMA process has residuals.

The assumptions in *Statement B* are a simple method of computing the residuals. In theory, we can estimate slightly better residuals, but the extra effort is not worth the slight gain in accuracy. The simple assumptions cause the residuals for the first few terms to be *slightly* over-stated, but the over-statement is not material.

Statement C: White noise process

The residuals are slightly over-stated and autocorrelated for the first few terms, but this discrepancy is not material. The residuals *resemble* a white noise process; they are not exactly a white noise process. The exam problems do not harp on this distinction.

Take heed: We test the residuals to validate the fitted model. If we fit an AR(1) process, the residuals resemble a white noise process, not a random walk or an AR(1) process.

Checking the residuals is an in-sample test.
Out-of-sample tests are also important.

We use both in-sample and out-of-sample tests.

In-sample tests compare the past estimates with the observed values.
Out-of-sample tests compare the forecasts with future values.

Your student project should leave out several values for out-of-sample tests.

Illustration: For a time series of monthly interest rates or sales or gas consumption, we may use years 20X0 through 20X8 to fit the model and year 20X9 to check the model.

For final exam problems, distinguish between in-sample and out-of-sample tests. Know the tests used for each, and how we compare different models.

Statement D: The variance is 1/T; the standard deviation is the square root of 1/T.

Take heed: The exam problems ask about

The distribution, which is normal, not ÷-squared, lognormal, or other.

The variance or standard deviation: we use the number of observations, not the degrees of freedom. We don't use T – p – q.

Keep several principles in mind:

As T increases, the sum of squared errors of the time series increases. It is proportional to T – p – q. In most scenarios, p and q are small and T is large, so the sum of squared errors increases roughly in proportion to T.

As T increases, the expected variance of the time series doesn't change. It may increase or decrease, but it is *unbiased*, so *we don't expect it* to increase or decrease.

As T increases, the variance of the sample autocorrelations decreases in proportion to 1/T if the residuals are a white noise process.

Statement E: The term *approximately* is used because the residuals are not exactly a white noise process.

Take heed: Know the formula and use of the Box-Pierce Q statistic and the Ljung-Box Q statistic. We don't use all the residuals. If we have 200 observations in the time series, we might use sample autocorrelations from lag 5 to 35.The first few sample autocorrelations have slight serial correlation even for a white noise process and correlations of higher lags and less stable

TS Module 15 Forecasting basics

(The attached PDF file has better formatting.)

- Minimum mean squared error forecasting
- Deterministic Trends
- ARIMA forecasting: autoregressive processes

Read Section 9.1, “Minimum mean squared error,“ on page 191.

Read Section 9.2, “Deterministic Trends,” on pages 191-193. Note that the forecasts are unbiased and the forecast error variance is constant (equation 9.2.4 on page 192).

Read Section 9.3, “ARIMA forecasting,” on page 194-197, stopping before the MA(1) heading. Know equation 9.3.8 on page 194.

Know the variance for AR(1) forecasts: equation 9.3.16 on the top of page 196 and equation 9.3.17 on the bottom of page 197.

Forecasting is the objective of time series analysis. A final exam problem may combine the pieces of time series analysis. It may give you data to estimate the ARIMA process and ask for the variance of the one period or two periods ahead forecast. The discussion forum has practice problems to help you prepare.

TS Module 15: Forecasting basics HW

(The attached PDF file has better formatting.)

Homework assignment: ARIMA(1,1,0) forecasts

An ARIMA(1,1,0) process has 40 observations y_t, t = 1, 2, …, 40, with y_{40} = 60 and y_{39} = 50.

This time series is *not* stationary, but its *first differences* are a stationary AR(1) process.

The parameter θ_0 of the stationary AR(1) time series of first differences is 5.

The 1 period ahead forecast $\hat{y}_{40}(1)$ is 60.

We determine the 2 period ahead forecast $\hat{y}_{40}(2)$.

A. What is the most recent value of the autoregressive model of first differences? Derive this value from the most recent two values of the ARIMA(1,1,0) process.
B. What is the one period ahead forecast of the first differences? Derive this value from the the one period ahead forecast of the ARIMA(1,1,0) process.
C. What is the parameter ϕ_1 of the AR(1) process of first differences? Derive this parameter from the 1 period ahead forecast.
D. What is the two periods ahead forecast of the AR(1) process of first differences? Use the parameter of the AR(1) process.
E. What is the two periods ahead forecast of the ARIMA(1,1,0) process? Derive this from the two periods ahead forecast of the AR(1) process.

TS Module 15 Mean reversion intuition

(The attached PDF file has better formatting.)

- The time series courses emphasizes the intuition of ARIMA processes.
- Exam problems use scenarios and probabilities to test the intuition.

Autoregressive processes have mean reversion.

Question 1.1: Mean reversion

An AR(1) model of 90 day Treasury bill yields has a mean of 5.00%, φ_1 of 0.60, and σ = 0.4% per month. The values are for the first day of each month.

The date now is January 15, 20X8. Let

- A_T = the *absolute value* of the difference between January's Treasury bill yield and 5%.
- A_{T+1} = *absolute value* of the difference between February's Treasury bill yield and 5%.

For which of the following values of the January 20X8 90 day Treasury bill yield is the probability *greatest* that $A_{T+1} > A_T$?

A. 4.3%
B. 4.6%
C. 4.9%
D. 5.2%
E. 5.5%

Answer 1.1: C

If the problem asks: "For which value of the January 20X8 90 day Treasury bill yield is the probability *greatest* that $y_{T+1} > y_T$," we choose the *lowest* value of y_T (Choice A).

For the *absolute* values, use the following reasoning. If $y_t = \mu$, the difference now is zero. The time series is stochastic with continuous values, so the probability that the value will again be the mean next period is infinitesimally small (i.e., close to zero). If the time series has discrete values, the probability is low. The absolute value of the difference from the mean increases, whether the actual interest rate increases or decreases..

The mean reversion is proportional, so the farther y_t now is from the mean, the more it moves toward the mean. The stochastic term may still move it away from the mean. But the stochastic term is constant for all values of the time series (since it is stationary). As the deterministic movement toward the mean increases, the less likely it will be offset by the stochastic movement.
Stationary autoregressive processes show mean reversion.

If the current value is not the mean, the forecasts eventually move toward the mean. Suppose the mean is μ and the most recent value of the time series is y_T.

- For an AR(1) model, the mean reversion is immediate. Next period's forecast is between μ and y_T.
- For an AR(p) model, the mean reversion may be delayed p-1 periods. The p periods ahead forecast is between μ and y_T.

Question 1.2: Mean reversion

An AR(1) model of 90 day Treasury bill yields has a mean of 5.00%, φ_1 of 0.60, and σ = 0.4% per month. The values are for the first day of each month.

The date now is January 15, 20X8. Let

- A_T = the absolute value of the difference between January's Treasury bill yield and 5%.
- A_{T+1} = absolute value of the difference between February's Treasury bill yield and 5%.

For which of the following values of the January 20X8 90 day Treasury bill yield is the probability *smallest* that $A_{T+1} > A_T$?

A. 4.3%
B. 4.6%
C. 4.9%
D. 5.2%
E. 5.5%

Answer 1.2: A

To grasp the intuition, reason through the following sequence.

If the January 20X8 Treasury bill yield is 5%, A_T is zero. In this exercise, the Treasury bill yield is an ARMA process with a error term. The probability that the February 20X8 Treasury bill yield is also 5% is zero (with no rounding) and very small even with rounding. It is likely that $A_{T+1} > A_T$.

If the January 20X8 Treasury bill yield is 5.01%, A_T is 0.01%. The forecast for February is 5.006%. The actual February yield has a normal distribution with a mean of 5.006% and a standard deviation of 0.4%. The 95% confidence interval for the February 20X8 Treasury bill yield is 4.206% to 5.806%. The probability that the February yield is between 4.99% and 5.01% is very small. We can compute the exact probability with a cumulative density function for the normal distribution. The exam problems do not ask for precise probabilities.

If the January 20X8 Treasury bill yield is 10%, A_T is 5%. The forecast for February is 8%. The actual February yield has a normal distribution with a mean of 8% and a standard deviation of 0.4%. The 95% confidence interval for the February 20X8 Treasury bill yield is 7.2% to 8.8%. The probability that the February yield is between 0% and 10% is almost 100%. This probability is the cdf for ±5 standard deviations from the mean.

TS Module 15 Moving average forecasts practice problems

(The attached PDF file has better formatting.)

For a moving average model, the forecast depends on the mean, the past residuals, and the moving average parameters.

- The past values don't affect the forecasts of a moving average model, but they affect the forecasts of all other ARMA models and all ARIMA models.
- The moving average parameters are the negative of the moving average coefficients.

Question 1.1: MA(3) Model

An MA(3) model of 100 values with $\theta_1 = -2$, $\theta_2 = -3$, and $\theta_3 = -4$ has a mean of 50 and a standard error of 2 ($\sigma^2 = 4$).

The residuals in Periods 98, 99, and 100 are zero.

The value in Period 101 is 51.

A. What is the residual in Period 101?
B. What is the forecast for Period 102?
C. What is the forecast for Period 103?
D. What is the forecast for Period 104?
E. What is the forecast for Period 105?

Part A: Let Z be the residual in Period 101. We have

$$51 = 50 - (-2) \times 0 - (-3) \times 0 - (-4) \times 0 + Z$$
$$\Rightarrow Z = 1$$

Part B: The forecast for Period 102 is $50 - (-2) \times 1 - (-3) \times 0 - (-4) \times 0 = 52$

Part C: The forecast for Period 103 is $50 - (-2) \times 0 - (-3) \times 1 - (-4) \times 0 = 53$

Part D: The forecast for Period 104 is $50 - (-2) \times 0 - (-3) \times 0 - (-4) \times 1 = 54$

Part E: The forecast for Period 105 is $50 - (-2) \times 0 - (-3) \times 0 - (-4) \times 0 = 50$

TS Module 16 ARIMA Forecasting

(The attached PDF file has better formatting.)

- ARIMA forecasting: moving average processes
- Random walk with drift
- ARMA(1,1) process

Read from the MA(1) heading on page 197 through page 198, stopping at the "Random walk with drift" heading. Know equation 9.3.21 at the middle of page 197 and equation 9.3.22 at the top of page 198. The forecasts are simple for the MA(1) model.

Read from "Random walk with drift" on page 198 through ARMA(p,q) on page 199.

Know equation 9.3.26 in the middle of page 198 and equation 9.3.27 on page 199. As the authors say: "In contrast to the stationary case, here Var($\epsilon_t(l)$) grows without limit as the forecast lead time l increases."

Read from ARMA(p,q) on page 199 until "To argue the validity …" at the bottom of page 200. Know the formulas for the ARMA(1,1) process: Equation 9.3.30 on page 199 and equation 9.3.32 at the top of page 200.

Read "Non-stationary models" from page 200 through end of page 201. You don't have to memorize the equations in this sub-section.

TS Module 16: ARIMA Forecasting HW

(The attached PDF file has better formatting.)

Homework assignment: ARIMA(0,1,1) process

The estimated (forecast) and actual values for Periods 48, 49, and 50 of an ARIMA(0,1,1) process are shown below.

Forecasts are *one period ahead* forecasts: $\hat{y}_{48}(1)$ for Period 49 and $\hat{y}_{49}(1)$ for Period 50.

Period	*Forecast*	*Actual*
48	70.5	71.5
49	72.0	74.0
50	73.0	74.8

- The estimated and actual values are for the ARIMA(0,1,1) time series.
- The values of μ and θ_1 are for the ARMA model of first differences. (Cryer and Chan use θ for an ARMA process, not θ_1.)

To solve the homework assignment, use the following steps:

- Determine the residual for each period for the ARIMA(0,1,1) model.
- These are also the residuals for the ARMA process of first differences.
- Determine the forecasts and actual values for the MA(1) process of first differences for the last two periods.
- Write equations for these forecasts in terms of the mean, the residual in the previous period, and θ_1. Remember that θ_1 is the *negative* of the moving average parameter.
- You have a pair of linear equations with two unknowns, μ and θ_1.
- Solve for μ and θ_1 and verify that these values give the forecasts in the table.
- Use the residual for Period 50 and the values of μ and θ_1 to forecast Period 51.
- Derive the forecast for the original ARIMA time series for Period 51.

A. What is the mean μ of the ARMA model of first differences?
B. What is θ_1 of the ARMA model of first differences?
C. What is the forecasted value of the ARIMA(0,1,1) process for Period 51?

Show the derivations of the parameters and the forecast.

TS Module 16: ARIMA forecasting

(The attached PDF file has better formatting.)

Intuition: Values, Residuals, and Forecasts

This posting explains non-stationary ARIMA processes whose first or second differences are stationary ARMA processes. See also the practice problems for this module.

- The homework assignment for this module relies on the explanations here.
- The final exam problems on ARIMA processes are modeled on the exercises here.

Know what each item in the ARIMA process refers to.

- The elements and forecasts of the time series are for the ARIMA process.
- The μ, ϕ, and θ parameters are for the underlying ARMA process.
  - The ARIMA process is not stationary and does not have a mean.

Exam problems may derive the parameters from forecasts and actual values. To do so, compute the forecasts, actual values, and residuals for the underlying ARMA process.

- The step-by-step guide uses first differences (d=1) and one-period ahead forecasts.
- The procedures for (i) second differences and (ii) other forecasts are similar.

The step-by-step guide below summaries all the forecasting modules (15-18).

Step #1: Compute residuals from forecasts and actual values.

The values of the ARIMA and ARMA processes differ, but the residuals are the same.

- The residual is the actual value minus the forecast.
  - The ARIMA forecast for Period t+1 is the ARIMA value for Period t plus the ARMA forecast for Period t+1.
  - The actual ARIMA value for Period t+1 is the ARIMA value for Period t plus the actual ARMA value for Period t+1.
- The residual of the underlying ARMA process is the residual of the ARIMA process.

If Y is an ARIMA process and $Y'_t = Y_t - Y_{t-1}$, the residual for Y'_t = the residual for Y_t.

Step #2: Compute the ARMA forecasts and actual values

- The ARMA value for Period t+1 is the difference of ARIMA values for Periods t and t+1.
- The ARMA forecast for Period t+1 is the ARIMA *forecast* for Period t+1 minus the *actual ARIMA value* for Period t.

Take heed: The ARMA forecast is *not* the difference of ARIMA forecasts for two periods.

Step #3: Derive the ARMA parameters

An exam problem may specify an ARIMA process, which implies a certain ARMA process.

- The first differences of ARIMA(p,1,q) is ARMA(p,q).
- The second differences of ARIMA(p,2,q) is ARMA(p,q).

Derive the μ, ϕ_1, and θ_1 parameters.

- From the values, forecasts, and residuals, write linear equations of the parameters.
- An exam problem may have N equations in N unknowns, where N = 1, 2, or 3.

An exam problem may ask for ARIMA forecasts, variances, and confidence intervals. From the parameters of the underlying ARMA process, compute forecasts, variances, and confidence intervals for the ARIMA time series. The exam problem may

- Give the parameters of the ARMA process.
- Derive the parameters from ARIMA values and forecasts in past periods.

Step #4: Derive forecasts for the ARMA and ARIMA processes

- Derive the forecasts for the underlying ARMA process.
- The most recent value of the ARIMA time series plus the cumulative sum of the next *k* ARMA forecasts is the *k* periods ahead ARIMA forecast.

Step #5: Variances

The exam problem may give the value of σ^2 or data from which to derive the value. This value is for variances and covariances, not for correlations. Understand the derivation of the values; don't just memorize formulas.

- Write the forecasts as linear combinations of scalars and error terms.
- The error terms are independent random variables with a variance of σ^2.
- The variance of the linear combination is the linear combination of the variances using the squares of the coefficients.

An exam problem may test

- Forecasts for future periods for the ARIMA process
- Variances, standard deviations, and confidence intervals for the ARIMA forecasts

The exam problem may also

- Give forecasts at time T for one or more periods.

- Give the observed value at time T+1.
- Derive new forecasts for one or more periods.

Exam problems use simple ARIMA processes [ARIMA(1,1,0), ARIMA(0,1,1), ARIMA(2,1,0), ARIMA(0,1,2), and ARIMA(1,1,1)], for which the solutions can be derived easily.

The following principles are used in many problems.

- Take first differences to convert an ARIMA process to an ARMA process. If the ARIMA process has *d* = 2, take second differences. ARIMA process with d > 2 are rare.

- The values of μ, ϕ, and θ are for the ARMA model of first differences, not for the original ARIMA process. An ARIMA process is not stationary, and it does not have a mean or variance. An exam problem may give these parameters or derive these parameters.

- Some authors have an δ parameter; Cryer and Chan call this θ_0.

- Integrate the ARMA process for forecasts of the ARIMA process. For *d* = 1, integrate once; for *d* = 2, integrate twice.

- The ARIMA residuals are the residuals of the underlying ARMA process.

- The *mean* of the ARMA process of first differences is the *drift* of the ARIMA process.

ARMA PROCESSES FROM ARIMA VALUES

An exam problem may give values of the ARIMA process to derive a residual for the ARMA process (or vice versa).

Exercise 1.1: The values for an ARIMA process are

Period	*Forecast*	*Actual*
T	13	15
T + 1	17	18

A. What are the forecast and actual values for Period T+1 for the underlying ARMA process?
B. What are the residuals for Period T and T+1 for the underlying ARMA process?

Part A: The values for the ARMA process of first differences subtract the actual Period T value from the forecast and actual values for Period T+1 in the ARIMA process.

Period	*Forecast*	*Actual*
T	–	–
T + 1	2	3

Take heed: The ARMA forecast for Period T+1 is 17 – 15, not 17 – 13.

Part B: The residuals are the same for the ARMA and ARIMA processes.

- The residual for Period T+1 is 3 – 2 = 1.
- The residual for Period T+1 is 18 – 17 = 1.

Take heed: Don't make the error that $ARIMA_t = ARMA_{t-1} + ARMA_t$. The correct relation is $ARIMA_t = ARIMA_{t-1} + ARMA_t$.

MEANS AND DRIFTS

Some exam problems give the parameters of the ARMA process of first differences to derive the drift of the ARIMA process. The *mean* of the ARMA process of first differences is the *drift* of the ARIMA process.

Exercise 1.2: An ARIMA process has the following values for Periods 10 to 15:

Period	[redacted]	*ARIMA*
10		100.0
11		102.0
12		103.6
13		105.4
14		108.3
15		110.0

What is the estimated mean of the underlying ARMA process?

Solution 1.2: The values of the ARMA process of first differences are in the ARMA column.

Period	*ARMA*	*ARIMA*
10		100.0
11	2.0	102.0
12	1.6	103.6
13	1.8	105.4
14	2.9	108.3
15	1.7	110.0
Average	2.0	

We derive the ARIMA process by integrating the ARMA process. The ARIMA process is not stationary: it has a drift, not a mean.

- The mean of the ARMA process of first differences is 2.
- The drift of the ARIMA process is 2.

Forecasts and confidence intervals

Exam problems may derive forecasts or confidence intervals for ARIMA processes.

- Some problems give the parameters of the ARMA process of first differences and the current value for the ARIMA process.
  - You derive the forecast for the ARIMA process.
- Some problems give estimated and actual values for the ARIMA process. You derive
  - The parameters of the ARMA process of first differences.
  - The forecasts for the ARIMA process.

The textbook gives equations for many of the relations. The intuition is straight-forward, but the equations are hard to memorize. For optimal exam preparation, focus on the intuition to convert ARIMA to ARMA and vice versa.

Derive the answers by first principles. To check your work: take first differences of the final ARIMA process to verify they form the ARMA process with the given parameters.

ARIMA(1,1,0) and ARIMA(0,1,1) models are AR(1) models and MA(1) models with a final step that integrates (sums) the terms.

The ARIMA(1,1,0) and ARIMA(0,1,1) time series are not stationary.

The deterministic part of the time series is

- linearly increasing if the mean of the first differences is positive
- linearly decreasing if the mean of the first differences is negative.

The ARIMA process has no mean, so it is not stationary.

ARIMA(1,1,0) problems

The exam problem may give

- actual and forecasted values of the ARIMA(1,1,0) time series
- the parameters of the AR(1) model of first differences.

Pindyck and Rubinfeld text: The AR(1) process of first differences has parameters μ, δ, and φ_1. From any two of these parameters, we derive the third. We use δ and φ_1 to solve the exam problem. The problem may give μ and φ_1, from which we derive δ.

Cryer and Chan text: The AR(1) process of first differences has parameters: μ, θ_0, and φ. From any two of these parameters, we derive the third. We use θ_0, and φ to solve the exam problem. The problem may give μ and φ, from which we derive θ_0.

An exam problem may say the mean of the AR(1) process (the first differences) is zero. This simplifies the work.

We need the most recent first difference to forecast future values. The exam problem may give the two most recent values of the original time series. We use the difference of these two values to forecast future first differences. We add these first differences to the most recent value of the original time series to get the forecasts for the original time series.

The exam problem may give only μ or φ_1 or neither μ nor φ_1, and it may give actual and forecast values of the original time series for several periods. We back into the values of μ and φ_1. Exam problems use simple algebra to derive the parameters, not the regression analysis used for the student projects.

If the problem gives only values, not forecasts, we derive the autoregressive parameters with linear regression and the moving average parameters with non-linear regression. For an ARMA process, we need non-linear regression. Exam problems use Yule-Walker equations to derive parameters, not regression analysis.

If the problem gives forecasts as well as values, we use algebra to derive parameters.

ARIMA(0,1,1) Problems

The exam problem may give the parameters of the MA(1) process of first differences. We need the most recent residual of the first differences to forecast future values. The residual of the first differences is the residual of the time series values. The exam problem may give the forecasted and actual values for the current period.

The exam problem may give only one moving average parameter, and it may give actual and forecast values of the original time series for several periods. We back into the values of μ and θ_1. Exam problems use simple algebra to derive parameters, not the regression analysis used for the student projects.

Variances

Some exam problems derive the variance of the forecasts. We explain the solution method for an ARIMA(0,1,1) process.

Extending the solution from ARIMA(0,1,1) to ARIMA(0,1,2) is like extending from MA(1) to MA(2). Most exam problems use low order ARIMA processes.

For exam problems with ARIMA(1,1,0) and ARIMA(1,1,1) processes,

- Convert to the AR(1) or ARMA(1,1) process of the first differences.
- Compute the variances of the forecasts of the corresponding moving average process.
- Calculate the variances of the forecasts of the ARIMA process.

Illustration: An ARIMA(0,1,1) model has $\mu = 1$, $\theta = 0.8$, and $\sigma_\varepsilon^2 = 4$. the most recent observation is $y_T = 0$ and its residual $\epsilon_T = 0$. These are parameters of the moving average time series that is the first differences of the ARIMA(0,1,1) time series.

The residual for the original time series is also the residual of the first differences, since a 1 unit increase in the original time series causes a 1 unit increase in the first difference.

- The forecasts of the future first differences are the mean for all lags, since the most recent residual is zero. The forecasts are 1, 1, 1, …
- For forecasts of the original time series, use the most recent value plus the sum of the forecasts of the first differences up through the specified lag. These forecasts are
  - Lag 1: 0 + 1 = 1
  - Lag 2: 0 + 1 + 1 = 2
  - Lag 3: 0 + 1 + 1 + 1 = 3

Common Errors

The variances of the forecasts of the ARMA process are σ^2 for the first lag and $(1 + \theta_1^2) \times \sigma^2$ for subsequent lags. Do **not** just add the variances of the forecasts

- Lag 1: $1 \times \sigma^2$
- Lag 2: $1 \times \sigma^2 + (1 + \theta_1^2) \times \sigma^2 = (2 + \theta_1^2) \times \sigma^2$
- Lag 3: $(2 + \theta_1^2) \times \sigma^2 + (1 + \theta_1^2) \times \sigma^2 = (3 + 2\,\theta_1^2) \times \sigma^2$

We can not add the variances because the forecasts are not independent.

- For a white noise process, the forecasts are independent.
- For an ARMA process, the serial correlation is not zero.

Exercise 1.3: Binary Residuals: ARIMA vs ARMA

Standard ARMA processes assume a normal distribution for the residuals. This exercise uses a binary residual to clarify the variance of ARMA vs ARIMA processes.

An MA(1) process has θ_1 = 1. The residual in each period is 1 or –1 with 50% chance of each. An ARIMA(0,1,1) process is the cumulative sum of the MA(1) process.

A. What are the variances of the 1, 2, & 3 period ahead forecasts for the ARMA process?
B. What are the variances of the 1, 2, & 3 period ahead forecasts for the ARIMA process?

Part A: The variance of the error term is $[1^2 + (-1)^2] / 2 = 1$. This is a population variance, so we divide by N, not by N-1.

- In a sample variance, the error terms are not necessarily +1 and –1.
- The population variance uses the distribution of error terms.

The one period ahead forecast for the ARMA process is $\mu + \epsilon_t - \theta_1 \times \epsilon_{t-1}$.

- The parameters μ and θ_1 are scalars with no variance.
- ϵ_{t-1} has already occurred, so it is also a scalar.
- ϵ_t has not yet occurred, so it is a random variable with a variance of 1.

⇛ The variance of $\mu + \epsilon_t - \theta_1 \times \epsilon_{t-1}$ is 1.

Part B: The one period ahead forecast for the ARIMA process is the one period ahead forecast for the ARMA process plus the most recent value of the ARIMA process.

- The most recent value of the ARIMA process has already occurred, so it is a scalar.
- The variance of the one period ahead forecast of the ARIMA process is the same as the variance of the one period ahead forecast of the ARMA process.

Part C: The two periods ahead forecast for the ARMA process is $\mu + \epsilon_{t+1} - \theta_1 \times \epsilon_t$.

- The parameters μ and θ_1 are scalars with no variance.
- ϵ_t and ϵ_{t+1} have not yet occurred.
- They are independent random variables with variances of 2.
- The variance of $\theta_1 \times \epsilon_{t-1}$ is $\theta_1^2 \times \text{var}(\epsilon_{t-1})$.
- The variance of the difference of two independent random variables is the sum of the variances of each random variable.

⇛ The variance of $\mu + \epsilon_t - \theta_1 \times \epsilon_{t-1}$ is $2 + \theta_1^2 \times 2 = 2 \times (1 + \theta_1^2) = 4$, since $\theta_1 = 1$.

Part D: The two periods ahead forecast for the ARMA process is the sum of

- The most recent value of the ARIMA process: Y_{t-1}
- The one period ahead forecast of the ARMA process.
- The two periods ahead forecast of the ARMA process.

$= Y_{t-1} + \mu + \epsilon_t - \theta_1 \times \epsilon_{t-1} + \mu + \epsilon_{t+1} - \theta_1 \times \epsilon_t$
$= [Y_{t-1} + 2 \times \mu - \theta_1 \times \epsilon_{t-1}] + [\epsilon_{t+1} + (1 - \theta_1) \times \epsilon_t]$

We have grouped the two periods ahead forecast into

- a sum of scalars (with no variance) and
- a sum of random variables.

$\theta_1 = 1$ so $(1 - \theta_1) = 0$, and the variance of the sum of random variables is $\text{var}(\epsilon_{t+1})$.

Intuition: Let $\mu = 0$ and $Y_{t-1} = 0$. The table below shows the possible values of the two periods ahead value of the ARMA process.

	Residuals		*Values*		
Scenario	ϵ_t	ϵ_{t+1}	*ARMA(t)*	*ARMA(t+1)*	*ARIMA(t+1)*
1	-1	-1	-1	0	-1
2	-1	1	-1	2	1
3	1	-1	1	-2	-1
4	1	1	1	0	1

For the ARMA process in Period t+1:

- The mean value is $(2 + 0 + 0 + -2) / 4 = 0$.
- The variance is $(2^2 + 0^2 + 0^2 + (-2)^2) / 4 = 8 / 4 = 2$.

For the ARIMA process in Period t+1:

- The mean value is $(-1 + 1 + -1 + 1) / 4 = 0$.
- The variance is $((-1)^2 + 1^2 + (-1)^2 + 1^2) / 4 = 4 / 4 = 1$.

Illustration: ARMA and ARIMA residuals

- Suppose the one period ahead forecast for the ARMA process is 1 and the actual value is 2, giving a residual of 1.
- The one period ahead forecast for the ARIMA process is Y_t + 1 and the actual value is Y_t + 2, giving a residual of 1.

The residual of 1 in the one period ahead forecast has two effects:

- The value of the first differences *increases* by 1 in the first forecast period.
- The value of the first differences *decreases* by 0.8 × 1 in the second forecast period.

The two effects on the original ARIMA time series are

- The value increases by 1 in the first forecast period.
- The value increases by 1 – 0.8 × 1 = 0.2 in the second forecast period.

For the ARIMA(0,1,1) process, each error term has a variance of σ^2. The effect on the forecast of lag *k* is

- σ^2 × the error term in future period *k*.
- $(1 - \theta_1)^2 \times \sigma^2$ × the error term in future period *k-1*.
- $(1 - \theta_1)^2 \times \sigma^2$ × the error term in future period *k-2*.
- **...**

The error terms are independent, so these variances are independent. The integrated ARIMA time series uses the sum of these variances. The variance of the *k* periods ahead forecast is $[1 + (k-1) \times (1 - \theta_1)^2] \times \sigma^2$.

Study recommendation: Review the intuition and practice problems. The exam problems are similar to the practice problems.

Take heed: Remember the sign convention for the moving average parameters. A moving average parameter of θ_j adds $-\theta_j \times \epsilon_{t-j}$ in Period *t*.

- A *positive* θ_1 in an ARIMA(0,1,1) model offsets the changes in adjacent periods.
  - For the *effect on the two periods ahead forecast*, a residual of 1 with a θ_1 of 0.8 is like a residual of 0.2 and a θ_1 of zero.
- A negative θ_1 in an ARIMA(0,1,1) model causes parallel changes in adjacent periods.
  - For the *effect on the two periods ahead forecast*, a residual of 1 with a θ_1 of –0.8 is like a residual of 1.8 and a θ_1 of zero.

Illustration: ARIMA(0,1,1) with $\theta_1 = 1$

If $\theta_1 = 1$, a residual in Period *t* causes an equal and opposite change in Period *t*+1. The only uncertainty in the ARIMA forecast is the residual in the forecast period.

Exercise 1.4: ARIMA(0,1,1) model

An ARIMA(0,1,1) process has $\theta = 0.8$, $\mu = 1$, and $\sigma_\varepsilon^2 = 4$ as parameters of the MA(1) time series that is the first differences of the ARIMA process. The forecast for the most recent period (Period T) was $\hat{Y}_T = 22.2$ and the actual value was $Y_T = 22.7$.

A. What is the residual of the ARIMA(0,1,1) process for the most recent period?
B. What is the residual of the underlying MA(1) process for the most recent period?
C. What is the one period ahead forecast for the underlying MA(1) process?
D. What is the one period ahead forecast for the ARIMA(0,1,1) process?
E. What is the two periods ahead forecast for the underlying MA(1) process?
F. What is the two periods ahead forecast for the ARIMA(0,1,1) process?
G. What is the variance of the one period ahead forecast for the MA(1) process?
H. What is the variance of the one period ahead forecast for the ARIMA(0,1,1) process?
I. What is the variance of the two periods ahead forecast for the MA(1) process?
J. What is the variance of the two periods ahead forecast for the ARIMA(0,1,1) process?

Part A: The residual of the ARIMA(0,1,1) process for Period T is 22.7 – 22.2 = 0.5.

Part B: The residual of the underlying MA(1) process for Period T is also 0.5.

The residual for an ARIMA process is also the residual for the underlying ARMA process.

- ARIMA forecast = actual ARMA value in the previous period + ARMA forecast.
- Actual ARIMA value = actual ARMA value in the previous period + actual ARMA value.

- The ARMA residual = the actual ARMA value – the ARMA forecast.
- The ARIMA residual = the actual ARIMA value – the ARIMA forecast
    - = the actual ARMA value – the ARMA forecast = the ARMA residual

Part C: The one period ahead forecast for the underlying MA(1) process is

$$\epsilon_{t+1} + \mu - \theta_1 \times \epsilon_t = 1 - 0.8 \times 0.5 = \epsilon_{t+1} + 0.6.$$

The mean of an error term ϵ is zero, so the forecast is 0.6.

Part D: The one period ahead forecast for the ARIMA(0,1,1) process is the most recent ARIMA value + the underlying ARMA forecast = 22.7 + 0.6 = 23.3. *Take heed:* The one period ahead forecast begins with the most recent value, not the most recent forecast.

Part E: The two period ahead forecast for the underlying MA(1) process is $\mu = 1$. The expected value of each future residual is zero, so the forecast is the mean.

Part F: The two period ahead forecast for the ARIMA(0,1,1) process is the one period ahead ARIMA forecast + the two periods ahead ARMA forecast = 23.3 + 1 = 24.3.

Part E: The *k* period ahead forecast (k > 1) for the underlying MA(1) process is μ = 1.

Part F: The *k* period ahead forecast (k > 1) for the ARIMA(0,1,1) process is

Y_T + ∑MA(1) forecasts = 22.7 + 0.6 + (k – 1) × 1 = 32.3 + k.

Part G: The variance of the one period ahead forecast for the MA(1) process is

$$\begin{aligned}
&\text{var}(\epsilon_{t+1} + \mu - \theta_1 \times \epsilon_t) = \text{var}(\epsilon_{t+1}) + \text{var}(\mu - \theta_1 \times \epsilon_t) \\
= \; &\text{var}(\epsilon_{t+1} + 0.6) = \text{var}(\epsilon_{t+1}) + \text{var}(0.6) = \text{var}(\epsilon_{t+1}) = \sigma^2 \quad = 4
\end{aligned}$$

Part H: The variance of the one period ahead forecast for the ARIMA(0,1,1) process is

$$\begin{aligned}
&\text{var}(\epsilon_{t+1} + Y_T + \mu - \theta_1 \times \epsilon_t) = \text{var}(\epsilon_{t+1}) + \text{var}(Y_T + \mu - \theta_1 \times \epsilon_t) \\
= \; &\text{var}(\epsilon_{t+1} + 23.3) = \text{var}(\epsilon_{t+1}) + \text{var}(23.3) = \text{var}(\epsilon_{t+1}) = \sigma^2 \quad = 4
\end{aligned}$$

Part I: The variance of the two periods ahead forecast for the MA(1) process is

$$\begin{aligned}
&\text{var}(\epsilon_{t+2} + \mu - \theta_1 \times \epsilon_{t+1}) = \text{var}(\epsilon_{t+2} - \theta_1 \times \epsilon_{t+1}) + \text{var}(\mu) \\
= \; &\text{var}(\epsilon_{t+1}) + \text{var}(-\theta_1 \times \epsilon_{t+1}) + \text{var}(1.0) \\
= \; &\text{var}(\epsilon_{t+1}) + \theta_1^2 \times \text{var}(\epsilon_{t+1}) + \text{var}(1.0) = (1 + \theta_1^2) \times \text{var}(\epsilon_{t+1}) = 1.64 \times \sigma^2 \quad = 6.560
\end{aligned}$$

Part J: The variance of the two periods ahead forecast for the ARIMA(0,1,1) process is

$$\begin{aligned}
&\text{var}(\epsilon_{t+2} + \mu - \theta_1 \times \epsilon_{t+1} + Y_T + \mu - \theta_1 \times \epsilon_t) \\
= \; &\text{var}(\epsilon_{t+2} + (1 - \theta_1) \times \epsilon_{t+1}) + \text{var}(Y_T + 2 \times \mu - \theta_1 \times \epsilon_t) \\
= \; &\text{var}(\epsilon_{t+2}) + \text{var}(1 - \theta_1 \times \epsilon_{t+1}) + \text{var}(24.3) \\
= \; &\text{var}(\epsilon_{t+2}) + (1 - \theta_1)^2 \times \text{var}(\epsilon_{t+1}) = \{ 1 + (1 - \theta_1)^2 \} \times \text{var}(\epsilon_t) = 1.04 \times \sigma^2 \quad = 4.160
\end{aligned}$$

WHITE NOISE PROCESS AND RANDOM WALK

Exercise 1.5: White Noise Process and Random Walk

- A white noise process has parameters μ = 1 and σ^2 = 4.
  - A white noise process is an MA(1) process with θ_1 = 0.
- A random walk with drift = 1 and σ^2 = 4 is the cumulative sum of the white noise terms.
- The most recent value of the random walk is 10.

A. What are the one- and two period ahead forecasts of the white noise process?
B. What are the one- and two period ahead forecasts of the random walk?
C. What are the variances of the one- and two period ahead forecasts of the white noise process?
D. What are the variances of the one- and two period ahead forecasts of the random walk?

Part A: A white noise process has no serial correlation. Each value is independent of the other values. The expected value (forecast) is the mean μ (= 1 in this exercise).

Part B: The random walk is the cumulative sum of the white noise process. The most recent value is 10, so the one period ahead forecast is 11, the two periods ahead forecast is 12, and so forth.

Part C: Each value of the white noise process is $\mu + \epsilon_t$. The mean μ is a scalar with no variance, and the residual ϵ_t is a random variable with a variance of σ^2 = 4, so each forecast has a variance of 4.

Part D: Each value of the random walk is $\sum\mu + \sum\epsilon_t$. The mean μ is a scalar with no variance, so the cumulative sum of the scalars also has no variance.

- The error term in a future value is a random variable with a variance of σ^2.
- The residual in an observed value is a realization of the random variable, which is a scalar with no variance.
- The error terms are independent with a variance of σ^2 = 4 for each one.
- The variance of a sum of independent random variables is the sum of the variances.
- The variance of the *k* period ahead forecast k × σ^2.

We are tempted to justify this reasoning by a comparison with white noise and random walk processes.

- For a white noise process, θ_1 = 0 so θ_1^2 = 0. The variance of each forecast is σ_ε^2.
- The integrated series is a random walk, for which the variance of the *k* periods ahead forecast is k × σ_ε^2.

But this reasoning is not correct.

- For a random walk, the variances of each period are uncorrelated.
- For an ARIMA(0,1,1) model, the variances of adjoining periods are correlated.

We use first principles to solve this exercise. Suppose the residual for the one period ahead value of the first differences is +1.

- The residual for the one period ahead value of the original time series is also +1.
- The residual for the two period ahead value of the original time series is +1 – 0.8 × +1 = +0.2.

Let y be the time series of first differences. The residual of +1 in period t+1 means that y_{t+1} increases by +1. Since {y} is an MA(1) model, y_{t+2} decreases by +1 × $-\theta_1$. The two periods ahead value of the integrated time series increases by 1 + 1 × $-\theta_1$.

Intuition: Suppose θ_1 = 1. Every random error is exactly offset the next period. If the residual in period *t* is +5, so the value in period *t* increases by 5, and the value in period *t+1* decreases by 5. For the integrated (original) time series, the residual of +5 in period *t* causes the value in period *t* to increase by +5 and no change in the value in period *t+1*.

We consider the variance of the two periods ahead forecast. For simplicity, suppose the MA(1) model of first differences has a mean of zero (so δ = 0).

- The one period ahead value increases by ϵ_{t+1}.
- The two periods ahead value increases by ϵ_{t+2} and decreases by 0.8 × ϵ_{t+1}.

For the integrated (original) time series

- The one period ahead value increases by ϵ_{t+1}.
- The two periods ahead value increases by ϵ_{t+2} + (1 – 0.8) × ϵ_{t+1}.

Part C: We extend this reasoning to all future forecasts:

For the integrated (original) time series

- The one period ahead value increases by ϵ_{t+1}.
- The two periods ahead value increases by ϵ_{t+2} + (1 – 0.8) × ϵ_{t+1}.
- The three periods ahead value increases by ϵ_{t+3} + (1 – 0.8) × ϵ_{t+2} + (1 – 0.8) × ϵ_{t+1}.

The MA(1) model is stationary, so all the error terms has the same variance.

- The variance of the two periods ahead forecast is σ_ε^2 + (1 – 0.8 σ_ε^2 .

- The variance of the three periods ahead forecast is σ_ε^2 + 2 × (1 – 0.8 σ_ε^2 .
- The variance of the *k* periods ahead forecast is σ_ε^2 + (*k* – 1) × (1 – 0.8 σ_ε^2 .

Question 1.6: ARIMA(0,1,1) Process

A time series follows an ARIMA(0,1,1) process. The forecast and actual values for Periods 48, 49, and 50 are shown below. The forecasts are *one period ahead* forecasts: $\hat{y}_{48}(1)$ for Period 49 and $\hat{y}_{49}(1)$ for Period 50.

Period	*Forecast Values*	*Actual Values*
48	80.5	81.5
49	82.0	84.0
50	83.0	84.8

A. What are the forecasts and values of the underlying MA(1) process (first differences)?
B. What are the residuals of the underlying MA(1) process (first differences)?
C. What are the forecasts for Periods 49 and 50 as linear combinations of μ and θ_1?
D. What is the mean μ of the MA(1) process of first differences?
E. What is the θ_1 of the MA(1) process of first differences?
F. What is the forecast for Period 51?

Part A: The table below show the MA(1) forecasts and actual values.

- The MA(1) actual values are the first differences of the ARIMA(0,1,1) values.
- The MA(1) forecasts are the ARIMA forecasts minus the previous ARIMA actual value.

Period	*Forecast Values*	*Actual Values*
48	$80.5 - Y_{47}$	$81.5 - Y_{47}$
49	0.5	2.5
50	-1.0	0.8

We don't know the ARIMA(0,1,1) value for Period 47. We show the expression in the table above to compute the residual for Period 47.

Part B: The residual is the actual minus the forecast value.

Period	*Forecast Values*	*Actual Values*	*Residual*
48	$80.5 - Y_{47}$	$81.5 - Y_{47}$	1.0
49	0.5	2.5	2.0
50	-1.0	0.8	1.8

The residual for the first differences equals the residual for the original time series. We can compute the same residuals from the ARIMA process.

Period	*Forecast*	*Actual*	*Residual*
48	80.5	81.5	1.0
49	82.0	84.0	2.0
50	83.0	84.8	1.8

Part C: Each forecast equals $\mu - \theta_1 \times \epsilon_{t-1}$

For Period 49, the forecast of the first difference is 82.0 – 81.5 = 0.5. This forecast equals

$$\mu - \theta_1 \times \epsilon_{48}$$
$$0.5 = \mu - \theta_1 \times 1.0.$$

For Period 50, the forecast of the first difference is 83.0 – 84.0 = –1.0. This forecast equals

the mean – θ_1 × the residual of Period 49

$$-1.0 = \mu - \theta_1 \times 2.0.$$

Parts D and E: We solve the simultaneous linear equations for $\mu = 2$ and $\theta_1 = 1.5$.

Part F: The residual in Period 50 is 84.8 – 83.0 = 1.8. The first difference for Period 51 is 2.0 – 1.5 × 1.8 = -0.700 . The forecast for Period 51 is 84.8 – 0.70 = 84.100.

Exercise 1.7: ARIMA(0,1,1) Model

An ARIMA(0,1,1) model for a time series of 60 observations, y_t, t = 1, 2, …, 60, has $\mu = 0$ and $\theta_1 = 0.4$. The forecast of the next observation, y_{61}, is 38. The actual value of y_{61} is 39. We continue to use the same ARIMA model.

What is the new forecast of y_{62}?

Solution 1.7: The residual in period 61 is 39 – 38 = 1. The forecasted first difference is 0 – 0.4 × 1 = – 0.4. The forecasted value for Period 62 is 39 + –0.4 = 38.6.

TS Module 17 Forecasting bounds

(The attached PDF file has better formatting.)

- Prediction limits
- Forecasting illustrations

Read Section 9.4, “Prediction limits,” on pages 203-204. Know equation 9.4.2 on page 203 and the formula for an AR(1) process in the middle of page 204.

Read Section 9.5, “Forecasting illustrations,” on pages 204-206. This section is graphics; there are no formulas to learn. Your student project will probably forecast the time series, and this section shows illustrations.

TS Module 17: Forecasting bounds HW

(The attached PDF file has better formatting.)

Homework assignment: Random walk with drift

An insurer's capital follows a random walk with a drift of $10 million a month and a volatility of $40 million a month. The initial capital is $200 million.

A random walk is an ARIMA(0,1,0) process. The capital *changes* are a white noise process, with a mean μ of $10 million a month and a σ of $40 million a month.

A. What is the distribution of capital after one month? (What is the type of distribution, such as normal, lognormal, uniform, or something else? Use the characteristics of a white noise process. What is the mean of the distribution after one month? Use the starting capital and the drift. What is the standard deviation after one month? The volatility is the standard deviation per unit of time, not the variance per unit of time. It is σ, not σ^2.)
B. What is the distribution of capital after six months? (The serial correlation is zero, so the capital changes are independent and the variances are additive. Derive σ^2 for one month from σ, add the σ^2's for six months, and derive the σ after six months.)
C. What is the distribution of capital after one year?
D. What are the probabilities of insolvency at the end of six months and one year? (You have a distribution with a mean μ and a standard deviation σ. Find the probability that a random draw from this distribution is less than zero. Use the cumulative distribution function of a standard normal distribution. Excel has a built-in function for this value.)
E. At what time in the future is the probability of insolvency greatest? (Write an equation for the probability as a function of (i) the mean of the distribution at time *t* and (ii) the standard deviation of the distribution at time *t*. To maximize this probability, set its first derivative to zero, and solve for *t*.)

TS Module 17 Forecast confidence intervals intuition

(The attached PDF file has better formatting.)

If the parameters of the time series are known with certainty, a 95% confidence interval is

The mean forecast ± 1.96 standard deviations of the forecast error.

We often use ±2 standard deviations for a 95% confidence interval, since the parameters are not known with certainty. The 95% is arbitrary, and using 1.96 gives a false aura of precision. Unless an exam problem specifies the number of standard deviations, you may use either 1.96 or 2.00.

The confidence intervals get wider as the lag increases and

- Either approach a maximum or reach a maximum for ARMA processes.
- Become infinitely wide for ARIMA processes.

The middle of the confidence interval *eventually* moves toward the mean of the time series.

Take heed: An ARIMA process that is not stationary has no mean.

Take heed: For an oscillating ARMA process, confidence intervals for two successive lags may not even overlap. Even for non-oscillating ARMA processes, confidence intervals for successive periods may not overlap if σ is small and the significance level is strict (small).

The exercises below highlights items to focus on. The first exercise is more difficult than the final exam problems. Most final exam problems use numerical examples and simple scenarios. The first exercise uses algebra to show the general relations.

The term *Prediction Interval* is used to mean confidence interval. We are P% confident that the predicted value of the time series falls within a certain interval.

Exercise 1.1: Prediction Intervals

We use an AR(1) process to model a time series of N observations, y_t, t = 1, 2, …, N.

The mean of the AR(1) process is μ and the most recent observed value is $y_T = \mu + k$.

- Let S be the 95% confidence interval for the one period ahead forecast.
- Let S′ be the 95% confidence interval for the two periods ahead forecast.

Which of the following is true?

A. If $\mu - k$ lies within S, then it lies within S′.
B. If $\mu - k$ lies within S′, then it lies within S.
C. If $\mu + 2k$ lies within S, then it lies within S′
D. If $\mu + 2k$ lies within S′, then it lies within S.
E. $\mu + k$ lies within both S and S′.

Solution: A. The confidence interval has two changes from the first to the second period:

- The confidence interval becomes wider. The one period, the variance of the forecast is the variance of the time series. For the two periods ahead forecast, the variance of the forecast adds the variance stemming from the ARIMA parameters (the previous period's value or residual).
- The mean (center) of the confidence interval moves. For an AR(1) model, the center moves closer to the mean of the time series (mean reversion). For all ARMA processes, the center of the confidence interval *eventually* moves toward the mean.

Take heed: For an ARMA process where $\phi_1 = \theta_1$, the width of the confidence interval is the same for the one and two periods ahead forecasts. Some exam problems tests this fact.

For most ARIMA models, the center of the confidence interval moves toward the mean of the time series but does not cross it.

Take heed: Exceptions are of two types:

- some autoregressive models oscillate
- some moving average models with large residuals in the most recent periods may move away from the mean for one or two forecasts.

Oscillation: An AR(1) model or an ARMA(1,1) model with a negative ϕ_1 oscillates. Even a high order ARIMA model with a large and negative ϕ_1 parameter oscillates. (*Large* means larger than the positive φ_2 and higher order autoregressive parameters.)

Statement A is true: The confidence interval for the second forecast period moves closer to the mean and it becomes wider. If $\mu - k$ lies within S it lies within S′.

Statement B is false: Suppose $\phi_1 \approx -1$ and $\sigma < k$.

- The center of the one period ahead forecast confidence interval is $\approx \mu - k$, so this point falls in the 95% confidence interval.
- The center of the two periods ahead forecast confidence interval is $\approx \mu + k$. The 95% confidence interval has a width of about $4 \times \sigma$, or $2 \times \sigma$ on each side. If σ is less than k, the point $\mu + k$ does not fall in the 95% confidence interval.

Statement C is false: If ϕ_1 is close to zero, the confidence interval for the second period is about the same size as the confidence interval for the first period. The difference between the two periods is that the center of the confidence interval moves toward the mean. The point $\mu + 2k$ may lie within S but not within S′.

Statement D is false: If $\phi_1 \approx 1$, the confidence interval doubles in width from S to S′. The center stays $\mu + k$. If $\frac{1}{4} k < \sigma < \frac{1}{2} k$, the point $\mu + 2k$ is in S′ but not S.

Statement E is false: If the ARIMA process is a white noise process with mean μ and standard deviation σ, the 95% confidence interval for all periods is $\mu \pm 1.96\ \sigma$. Suppose $k = 3\ \sigma$, so the most recent value was an outlier. The value $\mu + k$ does not fall in either confidence interval.

{The following two questions go together.}

Question 1.2: Confidence intervals of forecasts

- An AR(1) model has a mean (μ) of 0, ϕ_1 = 0.95, and σ = 1.
- The *z* value for a 95% confidence interval is 1.96.
- The Period T value (the most recent value) is 2.00

The point 4.2 falls in the 95% confidence interval for which of the funds withheld forecasts: one period ahead, two periods ahead, and three periods ahead?

A. One period ahead forecast only
B. Three periods ahead forecast only.
C. All but the one period ahead forecast.
D. All but the three periods ahead forecast.
E. All three forecasts.

Answer 1.2: C

Question 1.3: Confidence intervals of forecasts

- An AR(1) model has a mean (μ) of 0, ϕ_1 = 0.95, and σ = 1.
- The *z* value for a 95% confidence interval is 1.96.
- The Period T value (the most recent value) is 2.00

The point –0.5 falls in the 95% confidence interval for which of the funds withheld forecasts: one period ahead, two periods ahead, and three periods ahead?

A. One period ahead forecast only
B. Three periods ahead forecast only.
C. All but the one period ahead forecast.
D. All but the three periods ahead forecast.
E. All three forecasts.

Answer 1.3: C

The means, variances, standard deviations, and 95% confidence intervals are

Forecast (Periods Ahead)	*Mean*	*Variance*	*Standard Deviation*	*95% Confidence Interval* *Lower Bound*	*Upper Bound*
One Period	1.900	1.000	1.000	-0.060	3.860
Two Periods	1.805	1.903	1.379	-0.898	4.508
Three Periods	1.715	2.717	1.648	-1.516	4.945

The principles of AR(1) forecasts are

- For a high φ_1 (in absolute value):
  - The confidence interval becomes wider as the forecast lag increases.
  - The center of the confidence interval moves slowly toward the mean μ.
- For a low φ_1 (in absolute value):
  - The confidence interval stays about the same width as the forecast lag increases.
  - The center of the confidence interval move rapidly to the mean μ.

Take heed: If the autoregressive parameter is low (near zero), the width of the confidence intervals does not change much, but the center of the confidence interval changes. Exam problems have a variety of correct answers.

The exercise below explains the effects of the ARIMA parameters on confidence intervals.

Exercise 1.4: Prediction Intervals

We use an AR(1) process to model a time series of N observations, y_t, t = 1, 2, …, N. The variance of the error term is σ_ε^2, the autoregressive parameter is φ_1, the mean of the AR(1) process is μ, and the most recent observed value is $y_T = \mu + k$.

- Let S be the 95% confidence interval for the one period ahead forecast.
- Let S′ be the 95% confidence interval for the two periods ahead forecast.

Which of the following increases the probability that μ – k falls within S′?

A. μ increases / decreases (in absolute value)
B. k increases / decreases (in absolute value)
C. N increases / decreases (assume the parameters are estimated from observed values)
D. σ_ε^2 increases / decreases
E. φ_1 increases / decreases (in absolute value)

This exercise asks about S′, the confidence interval for the two periods ahead forecast. The solution discusses S as well, the confidence interval for the one period ahead forecast.

Part A: The mean μ does not affect the probability, since the points μ + k and μ – k are displacements from μ. The textbook often assumes μ = 0 to simplify the intuition.

Part B: As *k* increases:

- The centers of the confidence intervals move rapidly toward the mean μ.
- The width of the confidence intervals is not affected by *k*.

Illustration: Suppose the variance of the error term σ_ε^2 is 1 and the autoregressive parameter φ_1 is 50%.

If k = 1:

- The one period ahead forecast is μ + 0.5 and two periods ahead forecast is μ + 0.25.
- The point μ – 1 is 1.25 away from the center of the confidence interval.
- The variance of the forecast is $(1 + 0.5^2) \times 1 = 1.25$.
- The standard deviation of the forecast is $1.25^{0.5} = 1.118$.
- The point μ – 1 is within two standard deviations of the forecast mean.

If k = 4:

- The one period ahead forecast is μ + 2 and the two periods ahead forecast is μ + 1.
- The point μ – 4 is 5 units away from the center of the confidence interval.

- The variance of the forecast is $(1 + 0.5^2) \times 1 = 1.25$.
- The standard deviation of the forecast is $1.25^{0.5} = 1.118$.
- The point $\mu - 4$ is *not* within two standard deviations of the forecast.

Part C: In N is small, the forecast has more uncertainty. The 95% confidence interval decreases as N increases, and the probability that $\mu - k$ falls within the confidence interval decreases. If N is large, these relations are still true but their effect is very small.

The effect of N on the confidence interval is discussed in the regression analysis course, not the time series course. The variances for the time series course assume the ARIMA parameters are known. If the parameters are known (not estimated from observed values), the number of observations N does not affect the confidence intervals.

Know the relation of N to the width of the confidence interval if the parameters are not known with certainty. You are not responsible for the mathematics of this relation.

Part D: As σ_ε^2 increases, the 95% confidence interval becomes wider, so the probability that $\mu - k$ falls within the confidence interval increases.

Part E:

- As ϕ_1 increases, the mean reversion is weaker and the variance of the forecast is larger.
- As ϕ_1 decreases, the mean reversion is stronger and the variance of the forecast is smaller.

As ϕ_1 increases, the center of the confidence interval is farther from the mean μ and its width is larger.

- The center of the two periods ahead confidence interval is $\mu + \phi_1^2 \times k$.
- The point $\mu - k$ is $(1 + \phi_1^2) \times k$ below the mean.
- The forecast variance is $(1 + \phi_1^2) \times \sigma_\varepsilon^2$ and its standard deviation is $(1 + \phi_1^2)^{0.5} \times \sigma_\varepsilon$.
- The width of a confidence interval depends on the standard deviation of the forecast.

The ratio $(1 + \phi_1^2) \times k$ to $(1 + \phi_1^2)^{0.5} \times \sigma_\varepsilon = (1 + \phi_1^2)^{0.5} \times (k / \sigma_\varepsilon)$.
As ϕ_1 increases, this ratio becomes larger.

ARMA(1,1) PROCESS

The variance of the one period ahead forecast is the variance of the time series, as is true for all ARIMA processes.

For a *two periods ahead forecast*, the first order lag coefficients affect the variance of the forecast. The estimated variance is

- $(1 + \varphi_1^2) \times \sigma^2$ for the autoregressive model
- $(1 + \theta_1^2) \times \sigma^2$ for the moving average model.

For an ARMA(1,1) model, the variance of the one period ahead forecast is

$$(1 + [\varphi_1 - \theta_1]^2) \times \sigma^2.$$

- The moving average coefficient is the negative of the moving average parameter.
- We add the coefficients = we take the difference of the parameters.

Intuition: Suppose the current value is the mean and the current residual is zero.

- The one period ahead and two periods ahead forecasts are the mean for both AR(1) and MA(1) models.
- The variance of the one period ahead forecast is the variance of the error term.
- The additional variance of the two periods ahead forecast depends on the variance of the error term × ϕ_1^2 for the AR(1) model and the variance of the error term × θ_1^2 for the MA(1) model.

For subsequent forecasts, autoregressive and moving average models differ.

For a given error term, the moving average parameter affects a single forecast. The error term in the next period (T+1) affects the means of the following forecasts:

- θ_1 affects the two periods ahead forecast
- θ_2 affects the three periods ahead forecast
- θ_3 affects the four periods ahead forecast

For the variances of the forecasts:

- θ_1 affects forecasts two or more periods ahead
- θ_2 affects forecasts three or more periods ahead
- θ_3 affects forecasts four or more periods ahead

An autoregressive process has interactions among the parameters. The mean of each forecast depends on all the ϕ parameters up to that period.

- The three periods ahead forecast is affected by φ_1^2 and φ_2 .
- The four periods ahead forecast is affected by φ_1^3, $2 \times \varphi_1 \times \varphi_2$, and φ_3 .

PRINCIPLES OF INTERACTION

- Moving average parameters do not interact with each other.
  - If θ_j and θ_k affect the forecast, the variances are θ_j^2 and θ_k^2.
  - The errors are independent, so the combined variance is $\theta_j^2 + \theta_k^2$.
- Autoregressive parameters interact with each other.
  - A forecast of L periods uses all permutations of autoregressive parameters whose subscripts sum to L-1.

Take heed: The combinations are *multiplicative*, but the subscripts *sum* to L-1.

Take heed: Autoregressive parameters may be used more than once.

Illustration: If L = 5 periods, one permutation is $\varphi_1 \times \varphi_2 \times \varphi_1$

Take heed: We consider all permutations, not just distinct combinations.

Illustration: For L = 5 periods, the permutations $\varphi_2 \times \varphi_1 \times \varphi_1$; $\varphi_1 \times \varphi_2 \times \varphi_1$; $\varphi_1 \times \varphi_1 \times \varphi_2$ are three separate permutations.

1. Moving average parameters and autoregressive parameters for the same period are added.

Illustration: The two periods ahead forecast is affected by φ_1 and $-\theta_1$. They joint effect of an ARMA(1,1) model is $\varphi_1 - \theta_1$.

2. Moving average parameters and autoregressive parameters interact in one direction only. The moving average parameter is used only once and must be the first term.

Illustration: For L = 6, the subscripts add to 5. The first term in the permutation is a moving average parameter or an autoregressive parameter. Subsequent terms are autoregressive parameters.

Exercise 1.5: Prediction Intervals

We use an AR(1) process to model a time series of 200 interest rate observations, y_t, t = 1, 2, …, 200. We estimate $\hat{\phi}_1$ = 0.663 a $\hat{\sigma}$ = 8 basis points (0.08%).

We forecast interest rates for the next four periods: 201, 202, 203, and 204. Assume the 95% confidence interval for a sample of 200 observations uses a *t* value of 2.00.

A. What is the variance of the one period ahead forecast?
B. What is the width of the 95% confidence interval for the one period ahead forecast?
C. What is the variance of the two periods ahead forecast?
D. What is the width of the 95% confidence interval for the two periods ahead forecast?
E. What is the variance of the three periods ahead forecast?
F. What is the width of the 95% confidence interval for the three periods ahead forecast?
G. Which increases the width of the 95% confidence interval more for the *two* periods ahead forecast: a 20% increase in the estimate of φ_1 or a 20% increase in the estimate of σ_ε^2?
H. Which increases the width of the 95% confidence interval more for the *three* periods ahead forecast: a 20% increase in the estimate of φ_1 or a 20% increase in the estimate of σ_ε^2?

Part A: The variance of the one period ahead forecast is the square of the standard deviation of the error term, or 8^2 = 64 basis points squared. (For simplicity, we use basis points as the unit of measurement.)

Take heed: The variance is the square of the standard deviation, so the units are squared.

Part B: The standard deviation of the one period ahead forecast is 8 basis points. The 95% confidence interval has a width of 2 × *t* value × 8 basis points = 32 basis points.

Part C: The variance of the two periods ahead forecast is $(1 + 0.663^2)$ × 64 basis points = 1.440 × 64 basis points = 92.160 basis points.

- φ_1 does not have units of measurement; that is, it is unit-less.
- It is the same whether we use integers, percentage points, or basis points.

Part D: The standard deviation of the two periods ahead forecast is $92.160^{0.5}$ = 9.600 basis points. The 95% confidence interval has a width of 2 × *t* value × 9.6 basis points = 38.400 basis points.

Part E: The variance of the three periods ahead forecast is $(1 + 0.663^2 + 0.663^4)$ × 64 basis points = 1.633 × 64 basis points = 106.432 basis points.

Part F: The standard deviation of the two periods ahead forecast is $106.432^{0.5}$ = 10.317 basis points. The 95% confidence interval has a width of 2 × *t* value × 10.317 basis points = 41.268 basis points.

Part G: We compare the two changes:

- Increasing the estimate of φ_1 by 20% increases the variance of the two periods ahead forecast to $(1 + [1.2 \times 0.663]^2) \times 64$ basis points = 1.633 × 64 basis points = 106.432 basis points.
- Increasing the estimate of σ_ε^2 by 20% increases the variance of the two periods ahead forecast to $(1 + 0.663^2) \times 1.2 \times 64$ basis points = 1.440 × 76.8 basis points = 110.592 basis points.

Increasing the estimate of σ_ε^2 by 20% has the larger effect on the variance of the forecast and on the width of the 95% confidence interval.

Part H: We compare the two changes:

- Increasing the estimate of φ_1 by 20% increases the variance of the two periods ahead forecast to $(1 + [1.2 \times 0.663]^2 + [1.2 \times 0.663]^4) \times 64$ basis points = 2.034 × 64 basis points = 130.176 basis points.
- Increasing the estimate of σ_ε^2 by 20% increases the variance of the two periods ahead forecast to $(1 + 0.663^2 + 0.663^4) \times 1.2 \times 64$ basis points = 1.633 × 76.8 basis points = 127.718 basis points.

Increasing the estimate of φ_1 by 20% has the larger effect on the variance of the forecast and on the width of the 95% confidence interval.

TS Module 18 Forecast updates and weights

(The attached PDF file has better formatting.)

- Updating ARIMA forecasts
- Forecast weights

Read Section 9.6, “Updating ARIMA forecasts,” on page 207. Know equation 9.6.1 in the middle of page 207.

Read Section 9.7, “Forecast weights,” on pages 207-209.

Know equation 9.7.4 on page 208. Depending on the data in an exam problem, you may also use equation 9.7.5 on the bottom of page 208 or equation 9.7.6 on the top of page 209. The three equations are equivalent.

TS Module 18: Forecast updates and weights HW

(The attached PDF file has better formatting.)

Homework assignment: ARIMA(0,1,1) forecasts

An ARIMA(0,1,1) model for a time series of 100 observations, y_t, t = 1, 2, …, 100, has θ_1 = 0.4.

- The forecast of the next observation, y_{101}, is 25.
- The actual value of y_{101} is 26.
- The forecast of the next observation, y_{102}, is 26.
- The actual value of y_{102} is 26.

We continue to use the same ARIMA model. That is, we don't re-estimate the parameters with the additional data. We forecast y_{103}, the ARIMA value in the next period.

A. From the actual and forecasted values of y_{101}, derive the residual for the ARMA model of the first differences.
B. From the actual value of y_{101} and the forecasted value of y_{102}, derive the forecasted value for Period 102 for the ARMA model of the first differences.
C. This forecasted value for Period 102 is a function of μ, θ_1, and the residual for Period 101. Derive the μ (mean) of the ARMA model of first differences.
D. From the actual and forecasted values of y_{102}, derive the residual for the ARMA model of the first differences for Period 102.
E. Using this residual, determine the forecasted first difference for the next period.
F. From the forecasted first difference, derive the forecasted value of the original time series.

The values of μ and θ_1 are the coefficients of the ARMA process for the first differences. (Cryer and Chan use θ for an MA(1) process, not θ_1.)

TS Module 18 Forecast updates intuition

(The attached PDF file has better formatting.)

Exam problems relating ARIMA coefficients to forecasts may

- Give the ARIMA parameters and derive forecasts.
- Give forecasts and derive the ARIMA parameters.

Some exam problems do both. They give

- the forecasts for the next one or more periods
- the actual value in the next period
- ask for the new forecasts for the next one or more periods

To solve these problems

- Derive the ARIMA parameters from the old forecasts.
- Derive the new forecasts from the new values.

The exam problems generally give the type of process, such as AR(1) or ARMA(1,1).

- Some problems derive all the ARIMA parameters from the forecasts.
- Some problems give the mean and derive ϕ_1 or θ_1 (or other parameters).

Take heed: Suppose the exam problem gives the forecast for Period 101 and the actual observed value for Period 101. We use the observed value for the autoregressive part of the model and the residual for the moving average part of the model.

This posting uses the parameter θ_0 (δ):

- The Cryer and Chan textbook uses θ_0.
- The previous textbook for this course uses δ.

Exercise 1.1: Revised Forecasts AR(1) Model

An AR(1) model of 300 observations has $\hat{\mu}$ = $\hat{\phi}$ = 0.5, y_{300} = 2.6

- We *forecast* the next three periods: $\hat{y}_{300}(1)$, $\hat{y}_{300}(2)$, and $\hat{y}_{300}(3)$.
- In period 301, the *actual value* is 3.2.

A. What is estimated θ_0 (δ) for this AR(1) model?
B. In Period 300, what are the forecasts for periods 301, 302, and 303, or $\hat{y}_{300}(1)$, $\hat{y}_{300}(2)$, and $\hat{y}_{300}(3)$?
C. If the actual value in Period 301 is 3.2, what are the revised forecasts for periods 302 and 303, or $\hat{y}_{301}(1)$ and $\hat{y}_{301}(2)$?

Part A: We use the estimated mean μ and autoregressive parameter ϕ:

$$\theta_0\ (\delta) = \mu \times (1 - \varphi_1) = 3 \times (1 - 0.5) = 1.500$$

We use the same formula for all ARIMA models to derive θ_0 (δ) from μ or μ from θ_0 (δ).

Part B: We solve for the forecasts step by step:

- The forecast for Period 301 is 1.500 + 0.5 × 2.600 = 2.800.
- The forecast for Period 302 is 1.500 + 0.5 × 2.800 = 2.900.
- The forecast for Period 303 is 1.500 + 0.5 × 2.900 = 2.950.

We can also write the AR(1) process as a mean reversion, and use differences from the mean. An AR(1) model is $y_t = \theta_0\ (\delta) + \varphi_1 y_{t-1} + \varepsilon_t$. We write θ_0 (δ) in terms of the mean μ as $\theta_0\ (\delta) = \mu \times (1 - \varphi_1)$. We rewrite the AR(1) formula as

$$y_t = \mu \times (1 - \varphi_1) + \varphi_1 y_{t-1} + \varepsilon_t \Rightarrow$$
$$y_t - \mu = \varphi_1 (y_{t-1} - \mu) + \varepsilon_t$$

- The difference from the mean shrinks each period by a constant proportion.
- The actual value is then distorted by a random fluctuation ε_t.

The shrinkage from the mean is called *mean reversion*.

- Mean reversion does *not* imply that the *actual values* get closer to the mean.
- Mean reversion implies that the *forecasts* get closer to the mean as the forecast interval increases.

The *forecasts* are scalars, not random variables; they have no error term. If $y_t - \mu = k$, then

- $\hat{y}_t(1) - \mu = k \times \varphi_1$

- $\hat{y}_t(2) - \mu = k \times \varphi_1^2$
- $\hat{y}_t(3) - \mu = k \times \varphi_1^3$
- $\hat{y}_t(n) - \mu = k \times \varphi_1^n$

We determine the forecasts in each future period:

- For Period 300, $y_t - \mu$, (the difference from mean) is 2.6 – 3.0 = –0.4
- For Period 301, $y_t - \mu$, (the difference from mean) is –0.4 × 0.5 = –0.2
- For Period 302, $y_t - \mu$, (the difference from mean) is –0.2 × 0.5 = –0.1
- For Period 303, $y_t - \mu$, (the difference from mean) is –0.1 × 0.5 = –0.05

Part C: In Period 301, the actual value is 3.2. We revise the forecasts for Periods 302 and 303.

- The forecast for Period 302 is 1.500 + 0.5 × 3.200 = 3.100.
- The forecast for Period 303 is 1.500 + 0.5 × 3.100 = 3.050.

Alternatively, we use differences from the mean. We start one period later, with the actual value in Period 301.

- For Period 301, $y_t - \mu$, (the difference from mean) is 3.2 – 3.0 = +0.2
- For Period 302, $y_t - \mu$, (the difference from mean) is +0.2 × 0.5 = +0.1
- For Period 303, $y_t - \mu$, (the difference from mean) is +0.1 × 0.5 = +0.05

The difference from the mean technique is often easier. On exam problems,

- subtract the mean from the given values
- derive the forecasts (and variances, standard deviations, and confidence intervals)
- add back the mean.

This avoids errors in adding θ_0 (δ) to each term. If any autoregressive parameters are negative, using differences from the mean clarifies the pattern, since oscillations are around zero.

Exercise 1.2: AR(1) Forecasts

An AR(1) model of 90 day Treasury bill yields has a mean of 5.00%. The values are for the first day of each month.

- In December 20X7, we estimate the 90 day Treasury bill yield as 5.40% for January 20X8 and 5.05% for April 20X8.
- In January 20X7, the actual 90 day Treasury bill yield is 5.80%. We do not change the estimates for θ_0 (δ) or φ_1.

A. What does it mean that an AR(1) model is mean reverting at a constant proportional rate?
B. What is φ_1 for this AR(1) model?
C. What are the original estimates (in December 20X7) for February 20X8 and March 20X8?
D. What are the revised estimates (in January 20X8) for February, March, and April 20X8?

Part A: An AR(1) model is $y_t = \theta_0 (\delta) + \varphi_1 y_{t-1} + \varepsilon_t$.

We write θ_0 (δ) in terms of the mean μ as $\theta_0 (\delta) = \mu \times (1 - \varphi_1)$. We rewrite the AR(1) formula as

$$y_t = \mu \times (1 - \varphi_1) + \varphi_1 y_{t-1} + \varepsilon_t \Rightarrow$$
$$y_t - \mu = \varphi_1 (y_{t-1} - \mu) + \varepsilon_t$$

The difference between the value and the mean shrinks each period by a constant proportion. The actual value is then distorted by a random fluctuation ε_t. The forecasts are scalars, not random variables; they have no error term. If $y_t - \mu = k$, then

- $\hat{y}_t(1) - \mu = k \times \varphi_1$
- $\hat{y}_t(2) - \mu = k \times \varphi_1^2$
- $\hat{y}_t(3) - \mu = k \times \varphi_1^3$
- $\hat{y}_t(n) - \mu = k \times \varphi_1^n$

Given the *current value and any forecast*, we derive φ_1 as $(\hat{y}_t(n) - \mu) / (y_t - \mu) = \varphi_1^n$

Given *any two forecasts*, we derive φ_1 as $(\hat{y}_t(n) - \mu) / (\hat{y}_t(m) - \mu) = \varphi_1^{n-m}$

In this exercise,

- The one period ahead forecast is 5.4%, so the difference from the mean is 0.4%.
- The four periods ahead forecast is 5.05%, so the difference from the mean is 0.05%.

For the formula, n = 4 and m = 1, so we have

$0.05\% / 0.4\% = \varphi_1{}^3 \Rightarrow \varphi_1 = [\ 0.05\% / 0.4\%\ ]^{1/3} = 0.500$

Take heed: A third degree polynomial equation has three roots. This equation has one real root and two imaginary roots. We *discard the imaginary roots and choose the real root*. In general, if (n – m) is an odd number (1, 3, 5, …), we have one real root.

If (n – m) is an even number (2, 4, 6, …), we have two real roots. If φ_1 solves the equation, so does $-\varphi_1$.

The exercise may give the pattern of the forecasts or the pattern of the autocorrelations.

- If the forecasts approach the mean asymptotically, φ_1 is positive.
- If the forecasts oscillate about the mean, φ_1 is negative.

Illustration: In this exercise, had the problem given the forecasts for January and May, the difference 5 – 1 = 4 is even. We would not know the sign of φ_1. We would know the forecasts for March and July, but not the forecasts for February, April, and June.

Part C: The forecasted difference from the mean in January 20X8 is 5.40% – 5.00% = 0.40%. The forecasted differences from the mean for February and March 20X8 are

- February 20X8: 0.40% × 0.500 = 0.20%
- March 20X8: 0.20% × 0.500 = 0.10%

The forecasts add back the mean:

- February 20X8: 0.20% + 5.00% = 5.20%
- March 20X8: 0.10% + 5.00% = 5.10%

Part D: The actual difference from the mean in January 20X8 is 5.80% – 5.00% = 0.80%. The forecasted differences from the mean for February, March, and April 20X8 are

- February 20X8: 0.80% × 0.500 = 0.40%
- March 20X8: 0.40% × 0.500 = 0.20%
- April 20X8: 0.20% × 0.500 = 0.10%

The forecasts add back the mean:

- February 20X8: 0.40% + 5.00% = 5.40%
- March 20X8: 0.20% + 5.00% = 5.20%
- April 20X8: 0.10% + 5.00% = 5.10%

Exercise 1.3: AR(1) Forecasts

An AR(1) model of 90 day Treasury bill yields has a mean of 5.00%. The values are for the first day of each month.

- In December 20X7, we estimate the 90 day Treasury bill yield as 5.40% for January 20X8 and 5.05% for April 20X8.
- In January 20X7, the actual 90 day Treasury bill yield is 5.80%. We do not change the estimates for θ_0 (δ) or φ_1.

A. What is φ_1 for this AR(1) model?
B. What are the original estimates (in December 20X7) for February 20X8 and March 20X8?
C. What are the revised estimates (in January 20X8) for February, March, and April 20X8?

Part A: An AR(1) model is $y_t = \theta_0 (\delta) + \varphi_1 y_{t-1} + \varepsilon_t$.

We write θ_0 (δ) in terms of the mean μ as $\theta_0 (\delta) = \mu \times (1 - \varphi_1)$. We rewrite the AR(1) formula as

$$y_t = \mu \times (1 - \varphi_1) + \varphi_1 y_{t-1} + \varepsilon_t \Rrightarrow$$
$$y_t - \mu = \varphi_1 (y_{t-1} - \mu) + \varepsilon_t$$

The difference between the value and the mean shrinks each period by a constant proportion. The actual value is then distorted by a random fluctuation ε_t. The forecasts are scalars, not random variables; they have no error term. If $y_t - \mu = k$, then

- $\hat{y}_t(1) - \mu = k \times \varphi_1$
- $\hat{y}_t(2) - \mu = k \times \varphi_1^2$
- $\hat{y}_t(3) - \mu = k \times \varphi_1^3$
- $\hat{y}_t(n) - \mu = k \times \varphi_1^n$

Given the *current value and any forecast*, we derive φ_1 as $(\hat{y}_t(n) - \mu) / (y_t - \mu) = \varphi_1^n$

Given *any two forecasts*, we derive φ_1 as $(\hat{y}_t(m) - \mu) / (\hat{y}_t(n) - \mu) = \varphi_1^{n-m}$

In this exercise,

- The one period ahead forecast is 5.4%, so the difference from the mean is 0.4%.
- The four periods ahead forecast is 5.05%, so the difference from the mean is 0.05%.

For the formula, n = 4 and m = 1, so we have

$0.05\% / 0.4\% = \varphi_1^{\,3} \Rrightarrow \varphi_1 = [\, 0.05\% / 0.4\% \,]^{1/3} = 0.500$

Take heed: A third degree polynomial equation has three roots. This equation has one real root and two imaginary roots. We *discard the imaginary roots and choose the real root*. In general, if (n – m) is an odd number (1, 3, 5, …), we have one real root.

If (n – m) is an even number (2, 4, 6, …), we have two real roots. If φ_1 solves the equation, so does $-\varphi_1$.

The exercise may give the pattern of the forecasts or the pattern of the autocorrelations.

- If the forecasts approach the mean asymptotically, φ_1 is positive.
- If the forecasts oscillate about the mean, φ_1 is negative.

Illustration: In this exercise, had the problem given the forecasts for January and May, the difference 5 – 1 = 4 is even. We would not know the sign of φ_1. We would know the forecasts for March and July, but not the forecasts for February, April, and June.

Part C: The forecasted difference from the mean in January 20X8 is 5.40% – 5.00% = 0.40%. The forecasted differences from the mean for February and March 20X8 are

- February 20X8: 0.40% × 0.500 = 0.20%
- March 20X8: 0.20% × 0.500 = 0.10%

The forecasts add back the mean:

- February 20X8: 0.20% + 5.00% = 5.20%
- March 20X8: 0.10% + 5.00% = 5.10%

Part D: The actual difference from the mean in January 20X8 is 5.80% – 5.00% = 0.80%. The forecasted differences from the mean for February, March, and April 20X8 are

- February 20X8: 0.80% × 0.500 = 0.40%
- March 20X8: 0.40% × 0.500 = 0.20%
- April 20X8: 0.20% × 0.500 = 0.10%

The forecasts add back the mean:

- February 20X8: 0.40% + 5.00% = 5.40%
- March 20X8: 0.20% + 5.00% = 5.20%
- April 20X8: 0.10% + 5.00% = 5.10%

Question 1.4: ARIMA(0,1,1) process

An ARIMA(0,1,1) model for a time series of 25 observations, y_t, t = 1, 2, …, 25 has $\theta_1 = 0.6$, $\hat{y}_{24}(1) = 11$, $y_{25} = 12$, and $\hat{y}_{25}(1) = 13$.

A. What is $\hat{y}_{25}(2)$?
B. If the actual observed value in Period 26 is 14, what is $\hat{y}_{26}(1)$?

Part A: The residual in Period 25 is 12 – 11 = 1. This is also the residual of the MA(1) process of the first differences. The forecast for the next period is μ (the mean of the first differences) – 0.6 × 1. The forecasted first difference is 13 – 12 = 1, so μ = 1 + 0.6 = 1.6.

Part B: In Period 25, the two periods ahead forecast of the first differences for Period 27 is the mean of 1.6. The observed value of the original time series in Period 26 is 13, so the forecast for Period 27 of the original time series is 13 + 1.6 = 14.6.

Part C: The observed value in Period 26 is 14, so the residual is 14 – 14.6 = –0.4. This is also the residual for the first differences. The forecasted first difference for Period 27 is 1.6 – 0.6 × –0.4 = 1.84. The observed value of the original time series in Period 26 is 14, so the forecasted value for Period 27 for the original time series is 14 + 1.84 = 15.84.

TS Module 19 Seasonal models basics

(The attached PDF file has better formatting.)

- Seasonal ARIMA models
- Multiplicative seasonal ARIMA models

Read Section 10.1, “Seasonal ARIMA models,” on pages 228-229. The seasonal moving average process has the same form as the basic moving average process. Distinguish between the non-seasonal θ and the seasonal Θ. In the next section, both parameters are used; don’t confuse them.

Read Section 10.2, “Multiplicative seasonal ARIMA models,” on pages 229-232. Know equations 10.2.2, 10.2.3, and 10.2.4 at the bottom on page 230 and equation 10.2.5 at the top of page 231. Spend a few minutes to derive each equation; they follow the same pattern as the equations for other moving average models.

Know equation 10.2.8 at the bottom of page 231 and equations 10.2.9, 10.2.10, and 10.2.11 at the top of page 232. These are similar to the equations for the AR(1) process.

TS Module 19: Seasonal models basics HW

(The attached PDF file has better formatting.)

Homework assignment: auto insurance seasonality

This homework assignment applies ARIMA modeling to auto insurance policy counts. The same logic applies to other lines of business as well. The homework assignment has three parts, each adding a piece to the insurance scenario.

Part A: Seasonality

The population in State W, the number of drivers, and the number of cars are stable.

- An auto insurer sells twelve month policies in State W.
    - Its renewal rate is 90% on average.
    - The actual renewal rate varies with its prices and those of its competitors.

For Part A of the homework assignment, all insurers charge the same state-made rates.

- Random fluctuation determines how much new business each insurer writes.
- You model policies written each month (both new and renewal) with an ARIMA process.

A. How would you model renewal policies? Do you use a ϕ (autoregressive) parameter or a θ (moving average) parameter?
B. Is the process stationary?
C. If the renewal rate were 100%, would the process be stationary?

Intuition: If the insurer writes more policies in January 20X1, it writes more policies in January 20X2, January 20X3, and so forth. A random fluctuation dies our slowly. What is the autocorrelation function? Is this a stationary time series?

Part B: We add a free market to this exercise.

- The insurer competes in a free market.
- The insurer revises its rates each year, and its competitors revise their rates at other dates during the year.

New policies sold (new business production) depends on the insurer's relative rate level compared with its competitors. If its rates are lower (higher) than its competitors', it writes more (fewer) new policies.

- The insurer does not expect higher or lower rates than its competitors charge.
- At any time, its rates may be higher or lower than its competitors charge, so its market share may grow or shrink.

D. How would you change the model? Do you add an autoregressive or a moving average term? Note that rate changes occur once a year, so if the insurer has high (low) rates now, it will probably have high (low) rates next month.
E. Does the free market increase or decrease the variance of the process? (When firms charge different prices and revise their prices periodically, is market share more or less variable?)

Part C: The insurer revises rates if its policy count is higher or lower than expected.

- If the insurer writes more policies than expected, it is afraid that its rates are too low, and it files for a rate increase.
- If the insurer writes fewer policies than expected, it is afraid that its rates are too high, and it files for a rate decrease.

F. If this rate change occurs immediately, how should you change the model? Do you add an autoregressive or a moving average term? (In practice, rate filings take several months to be approved. Assume this rate change takes effect immediately.)
G. If this rate change has a one month lag, how should you change the model? (For this part of the homework assignment, assume the insurer compares its actual vs expected policy count at the end of each month and changes the rate beginning either the next day or one month later. Use whichever assumption you prefer.)

TS Module 20 Seasonal models advanced

(The attached PDF file has better formatting.)

- Non-stationary seasonal ARIMA models
- Forecasting seasonal ARIMA models

Read Section 10.3, “Non-stationary seasonal ARIMA models,” on pages 233-234.

Read Section 10.4, “Model specification, fitting, and checking,” on pages 234-241. These are illustrations; they help you with your student project, but they no material to know for the final exam.

Read Section 10.5, “Forecasting seasonal ARIMA models,” on pages 241-245. These are examples; you need not memorize the equations. Any exam problems are structured to be answered intuitively, not by equation.

TS Module 20 Seasonal models advanced HW

(The attached PDF file has better formatting.)

Homework assignment: auto sales

New auto sales vary by month. New models are released in October, in time for holiday sales of November and December. Sales in January and February are low, since cold weather and the high sales in the previous months leave low demand for new cars.

Part A: Compare the seasonality of auto insurance in the homework assignment for the previous module with the seasonality of auto sales in this homework assignment. We can model seasonality by a seasonal autoregressive parameter or by seasonally adjusted data.

- Monthly temperature is seasonal. But if 20X8 has a particularly cold January, we don't expect January 20X9 to be colder than usual.
- Insurance renewals depend on the number of policies written one year ago (or six months ago). If January 20X8 has a higher than usual policy count, we expect a higher than usual policy count in January 20X9 as well.

1. How would you model the seasonality of auto insurance?
2. How would you model the seasonality of auto sales?

Part B: Cars are durable and expensive goods.

Cars last several years. If a consumer buys a new car in year X, his or her demand for new cars is low in year X+1 and rises steadily lover subsequent years.

Cars are expensive. In tough economic times (recessions, high unemployment), demand for new cars is weak. In prosperous years, built-up demand for new cars causes high sales.

GDP shows weak cycles. Durable goods (cars, housing) show strong cycles.

1. How would you model the long-term cycles of auto sales? Use an annual model, not a monthly model, so you can avoid the monthly seasonality.
2. Assume GDP follows a weak cycle:
   a. Consumers who have not bought a new car recently are more likely to buy one if the economy is strong.
   b. Consumers who have bought a new are recently are unlikely to buy a second one for several years.
   c. Explain how you might model this process.

This homework assignment does not have a single correct answer. Economists model the sales of durable goods many ways. This homework assignment asks you to demonstrate that you understand how autoregressive and moving average processes work.

TS Module 21 Building an ARIMA process

(The attached PDF file has better formatting.)

- Model specification, fitting, and checking
- Tools for model building: Excel VBA macros and R functions

Modules 21-24 help you design and complete the student project. They have no homework assignments; study for the final exam and begin your student project. The time series on-line course has 19 homework assignments, of which you must complete 15.

Read the step-by-step guide to ARIMA modeling. Your student project builds an ARIMA process to model an actual time series. You specify an ARIMA process, fit parameters, and verify that the actual time series conforms to the process. Getting started is often hard, so we provide a step-by-step guide.

Excel does not have all the tools needed for ARIMA modeling.

- If you use Excel, review the illustrative worksheets on the NEAS web site.
- If you use R, review Appendix A in the Cryer and Chan textbook, which describes the time series modeling tools in their TSA package.

Excel has a *REGRESSION* add-in that can fit autoregressive processes. Moving average and mixed models need nonlinear regression, for which Excel does not have built-in functions.

- If you use Excel for your student project, fit AR(1), AR(2), and (perhaps) AR(3) models.
- If the sample autocorrelation function indicates an MA(1) or ARMA(1,1) model, use the Yule-Walker equations.

As the textbook explains, the Yule-Walker equations are not efficient for moving average and mixed models. The fitted parameters may be far from the true coefficients, and your results may not predict well.

ARIMA modeling is both art and science. A time series that follows an ARIMA process exactly is rare, since changes in the environment make the parameters unstable.

Illustration: Interest rates might follow an AR(1) process in 20X1 with μ = 5% and ϕ = 40%. In 20X2, they might follow an AR(1) process in 20X1 with μ = 6% and ϕ = 30%.

Statisticians speak of interventions: changes in the environment that cause the model to change. Part of time series modeling is judging whether an apparent change is a random fluctuation or a true change in the process.

Some candidates conclude that ARIMA modeling lacks the economic sophistication of real models. They say we should study the fiscal and monetary policies that affect interest rates to form realistic forecasts.

This perspective is correct, but it misses the point. We use structural (economic) factors when we have the information, but we don't always have these data.

Illustration: An actuary is forecasting claim severity trends. In theory, the actuary should examine inflation indices and other economic data that affect severities, such as gas prices, unemployment, legal trends, and business growth. But the actuary is pressed for time, and simply fits an exponential curve to average claim severities of the past 12 quarters.

Time series modeling says we can do better. The actuary should take logarithms and first differences and then fit an ARMA process. The ARMA process will probably fit better than the simple exponential curve. A sophisticated econometric model might do even better, but it takes too long to build and the added precision may be small.

For your student project, you may find that ARIMA models sometimes do quite well.

TS Module 22 Student projects interest rates

(The attached PDF file has better formatting.)

- Internet data for time series processes
- Templates for other projects

Review the Excel worksheets and the discussion forum postings on interest rates.

- Post World War II Interest rates fall into three regimes, based on drift and volatility.
- The three regimes reflect macroeconomic policy, both fiscal and monetary.

Fitting an ARIMA model may not work well if the time series pattern changes.

Some time series are stable, such as daily temperature. The same pattern occurs for the past hundred years, with random fluctuations that are modeled by the ARIMA process. Sunspots and other physical or meteorological time series are similar.

Other time series change. Birth rates, divorce rates, crime rates, longevity, travel patterns, interest rates, inflation, GDP, and hundreds of other statistics change with the environment. Examine the time series and see whether you should use a single ARIMA process or two processes for different periods.

Review the templates for other time series projects on the NEAS web sites.

- Time series are shown for several economic indices.
- Thousands of internet sites has extensive time series data for student projects.

You can use the data on the NEAS web site. You will produce a better student project if you spend an hour surfing the internet and finding up-to-date and more complete data.

TS Module 23 Student projects sport statistics

(The attached PDF file has better formatting.)

- Sports statistics: time series of won-loss records
- Time series differences by team or league

Read the discussion forum postings on sports statistics.

- The NEAS web site has selected statistics for four sports: baseball, hockey, basketball, and football.
- Professional sports teams have web sites with extensive data on every player.

Several past student projects on sports statistics are posted on the NEAS web site. Review those projects to get ideas for your own.

If you do a student project on a favorite sports team, don't just use the data on the NEAS web site. Go to the web site of the sports team, where you will find a wealth of statistics. Design a project that differs from the standard template on the NEAS web site. You spend an extra hour finding data and designing your project, but you enjoy the project more. You form hypotheses that interest you, and you learn how to test them.

TS Module 24 Student projects weather and demographics

(The attached PDF file has better formatting.)

- Weather time series: daily temperature
- Demographic time series: birth rates and crime statistics

Read the discussion forum posting on daily temperature. The NEAS web site has statistics from 1,000+ U.S. weather stations for the past hundred years.

Be sure to examine the data for missing figures and adjust your time series analysis.

Review past student projects posted on the NEAS discussion forum for daily temperature, rainfall, snow-fall, and similar weather-related items. The internet has a wealth of data on global warming. The controversies surrounding these data are fascinating. Time series analysis show whether the historical data show warming of the earth over the past century.

Review the project templates and past student projects on birth rates, crime statistics, and other demographic time series. These give you ideas for your own project. The FBI has a wealth of data available on the internet. If you enjoy TV dramas on police investigations, you can analyze actual data showing crimes by location and year.

Births, marriages, divorces, and deaths are recorded faithfully in many countries. These items are perennially in the news: Are teen-age births increasing? Is the divorce rate increasing?