TS Module 16: ARIMA forecasting

(The attached PDF file has better formatting.)

Intuition: Values, Residuals, and Forecasts

This posting deals with non-stationary ARIMA processes whose first or second differences are stationary ARMA processes. See also the homework assignment 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 Y_t are for the ARIMA process.
- The μ, ϕ, and θ parameters are for the underlying ARMA process ∇Y_t.
  - The ARIMA process is not stationary and does not have a mean.
  - The variance of the ARIMA process is not finite.

Exam problems relate ARIMA parameters, expected values, forecasts, and actual values. Some exam problems compute the ARIMA parameters from forecasts and actual values. To do so, compute parameters, forecasts, actual values, and residuals for the underlying ARMA process. In practice, we compute parameters, expected values, and forecasts from observed (actual) values. Those computations use the method of moments, non-linear regression, and maximum likelihood. Final exam problems may also give expected values or forecasts, for which the solution is simpler.

- 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.

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

- The forecasts are expected values, so the residual = actual value minus the forecast.
- The values of the ARIMA and ARMA processes differ, but the residuals are the same.

The rationale for this relation is:

- The residual is the actual value minus the forecast.
  - The ARIMA forecast for Period t+1 is the ARIMA value for Period *t* + ARMA forecast for Period *t+1*.
  - The actual ARIMA value for Period t+1 is the ARIMA value for Period *t* + 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 $Y'_t = \nabla Y_t = Y_t - Y_{t-1}$, 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).

Most final exam problems use first differences, not second differences, when computing the parameters of the underlying ARMA process.

Derive the μ, ϕ, and θ 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.

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.

An exam problem may

- 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.

Cryer and Chan refer to this as updating forecasts. The revised forecasts can be solved by first principles or by formula. The final exam problems use first principles.

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 and are not tested on the final exam.

- 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 them.

- Some ARMA processes have a constant parameter; Cryer and Chan call this θ_0. The constant parameter is not the mean.

- 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 16.1: The values for an ARIMA process are

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

A. What are the forecast (expected) 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 forecast and actual values for the ARMA process of first differences are the forecast and actual values for Period T+1 in the ARIMA process minus subtract the actual Period T ARIMA value.

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 $ARIMA_t = ARIMA_{t-1} + ARMA_t$.
- The ARIMA process is the cumulative sum of the ARMA values.

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

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

What is the estimated mean of the underlying ARMA process?

Solution 16.2: 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.

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.

The AR(1) process of first differences has parameters: μ, θ_0, and ϕ. From any two of these parameters, we derive the third. Some exam problems say 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. Use the difference of these two values to forecast future first differences. Add the 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 or neither μ nor ϕ, 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. 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.