TS Module 17 forecast variance: practice problems

(The attached PDF file has better formatting.)

[Final exam problems may ask about variances, standard deviations, or confidence interval of forecasts. This time series course does not explain how to estimate confidence intervals if the ARIMA parameters are not known with certainty. The final exam problems may assume the parameters are known with certainty or they may give t-values.]

** Exercise Question 17.1: Forecast Variance

An AR(1) model for a time series of T observations has parameters ϕ and θ_0 which are known with certainty; they are not estimates. The variance of the error term is σ^2_ε.

A. What is the variance of the one period ahead forecast?
B. What is the variance of the two periods ahead forecast?
C. What is the variance of the three periods ahead forecast?
D. What is the variance of the n periods ahead forecast as $n \rightarrow \infty$?

Part A: The one period ahead forecast is $\hat{y}_{t+1} = \theta_0 + \phi \times y_t + \epsilon_t$. All items in this expression are known except for ϵ_t, so the variance of the forecast is the variance of $\epsilon_t = \sigma^2_\varepsilon$.

Part B: The two periods ahead forecast is $\hat{y}_{t+1} = \theta_0 + \phi \times y_{t+1} + \epsilon_{t+1}$.

Rewrite this as $\hat{y}_{t+1} = \theta_0 + \phi \times (\theta_0 + \phi \times y_t + \epsilon_t) + \epsilon_{t+1} = \theta_0 + \phi \times (\theta_0 + \phi \times y_t) + (\phi \times \epsilon_t) + \epsilon_{t+1}$

The random variables in this expression are ϵ_t and ϵ_{t+1}, both of which have a variance of σ^2_ε, so the variance of this expression is $(1 + \phi^2) \times \sigma^2_\varepsilon$.

Part C: The same reasoning shows the variance of the three periods ahead forecast is $(1 + \phi^2 + \phi^4) \times \sigma^2_\varepsilon$

Part D: The variance of the n periods ahead forecast as $n \rightarrow \infty$ is

$$(1 + \phi^2 + \phi^4 + \dots) \times \sigma^2_\varepsilon = \sigma^2_\varepsilon / (1 - \phi^2)$$

Jacob: If we know all the parameters, the forecast is certain; why does it have a variance?

Rachel: The forecast itself is certain; it has no variance. The actual value next period has a variance. The problem means:

> *If we ran this experiment 100 times and collected 100 actual values of next year's values, what would be their variance?*

The forecast variance means the variance of the forecasted value.

** Question 17.2: Optimum Forecast

The optimum forecast with a time series model is defined as the forecast which has

A. The minimum mean-square forecast error.
B. The minimum R^2.
C. The maximum R^2.
D. The minimum Q-statistic.
E. The maximum Q-statistic.

Answer 17.2: A

Jacob: If the optimal forecast is the forecast with the minimum mean squared error, why do we focus on the Box-Pierce Q statistic?

Rachel: If we posit an ARIMA process, we observe the Box-Pierce Q statistic. We never observe the mean squared error of the forecast until we observe the forecasted values.

Jacob: Why don't we examine the mean squared error if we model all values except the last one with ARIMA processes, and pick the process that gives the smallest error?

Rachel: The time series is stochastic. By random fluctuations, a poor model may forecast exactly in any one value.

Jacob: Why don't we forecast the last 20 or 30 values?

Rachel: Suppose the entire time series has 50 values. The ARIMA process found from the first 20 or 30 values may differ from the ARIMA process found for the entire 50 values.

Jacob: Suppose the time series has 1,000 values.

Rachel: One ARIMA process may do well if the time series doesn't turn but remains in a steady trend or drift. Another ARIMA process may do well at predicting turning points. Suppose the time series turns in the last 20 points. The ARIMA process that best predicts the last 20 points may not be the ARIMA process that best predicts the next 20 points.