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.