OSCILLATING SERIES AND NEGATIVE AUTOCORRELATIONS

Updated: April 19, 2006

Jacob: Random walks, white noise processes, and mean reverting AR(1) models are often appropriate. What about oscillating series (that is, negative autoregressive parametrers) and negative autocorrelations? Negative autoregressive parameters and negative autocorrelations seem unusual.

Rachel: They are indeed unusual. As the textbook states about the Durbin-Watson statistic, positive serial correlation is common but negative serial correlation is rare.

Jacob: How might negative autocorrelations occur?

Rachel: True negative autocorrelations occur in multi-period cycles. But multi-period cycles are uncommon, and they can rarely be modeled. Property-casualty underwriting cycles average 6 to 8 years in length, but no one has yet modeled them successfully. If we could predict the turning points or the severity of an underwriting cycle, the cycle would probably not follow the predictions.

Oscillating series and negative autocorrelations often stem from errors in the modeling process. These errors have occurred in some student projects for interest rate modeling. We review them here for candidates working on student projects.

~ Random fluctuations that *look like* oscillations (but are not)
~ Measurement error with short time intervals, a stable series, and a steady trend
~ Taking second or third differences when first or second differences are white noise
~ Combining time periods with upward and downward drifts
~ Modeling trend as a random walk

Jacob: If we observe an oscillating pattern or negative autocorrelations, what do we do?

Rachel: Consider first if the pattern is just random fluctuation. A white noise process has random autocorrelations, half of which are negative. A line graph connecting a white noise process might appear to be a cycle, since we tend to see patterns even in random data. Some candidates assume this is an oscillating series (which it is not), and deduce that an AR(1) model with a negative ϕ_1 is needed.

If more than half the sample autocorrelations are within one standard deviation from zero, you may be seeing patterns in a white noise process. Use Bartlett's test and the Box-Pierce Q statistic to check for a white noise process before you assume an oscillating series.

Jacob: What if the process is not white noise? What if the time series truly oscillates?

Rachel: Check if you are taking frequent differences of a time series with a low drift. This occurs with daily differences of the Moody's investment yield rate.

DAILY VS MONTHLY RATES

Jacob: How does the time interval affect the observed patterns? What time intervals should we use for interest rates? Should we use daily changes or monthly changes in the interest rate? If we have daily figures, why not use them? Doesn't that improve the model?

Rachel: We consider interest rate volatility and the accuracy of measurement.

~ If rates are volatile and change frequently, we use shorter time intervals.
~ If interest rates are rounded and are not volatile, too frequent time intervals may cause spurious patterns.

The NEAS web site shows *daily* values for Moody's long-term corporate bond yield, since Moody's announces the rate each business day. You may combine rates into a monthly average, or you may use the rate on the first day of each month or week. Fitting an ARIMA model to the daily rates (or their first differences) may create spurious effects.

Suppose interest rates are a random walk with a drift of 0.5 basis points a day. Tomorrow's rate is forecasted as today's rate plus 0.005%. Interest rates increase by 2½ basis points a week, or 52 × 0.025% = 1.30% a year. This is a reasonable interest rate process.

A random walk is not stationary. The first differences are a stationary white noise process. The mean of the white noise process is the drift of the random walk, or 0.005%.

If the interest rate is 8.002% on day 1, the expected rates for days 2, 3, 4, … are

{8.007, 8.012, 8.017, 8.022, 8.027, 8.032, 8.037, 8.042, 8.047, 8.0452 …}

We measure interest rates in basis points, or 0.01%. With two decimal place accuracy, the expected rates are

{8.01%, 8.01%, 8.02%, 8.02%, 8.03%, 8.03%, 8.04%, 8.04%, 8.05%, 8.05% …}

Because of the rounding, the first differences are an oscillating series:

{0.00, 0.01, 0.00, 0.01, 0.00, 0.01, 0.00, 0.01, 0.00, … }

We model this series as AR(1) with $\delta = 0.01$ and $\phi_1 = -1$, not as a white noise process. This gives a mean of 0.01% / (1 – (–1)) = 0.005%, which is correct, but the oscillating pattern and the autoregressive parameter are not correct. This is a *spurious* relation; the true interest rate first differences are a white noise process with no oscillation.

Jacob: Would we see this in practice?

Rachel: The actual Moody's long-term corporate bond yield has declined by about 25 basis points a year over the past two decades. This may create spurious effects if we fit daily first differences to an ARIMA model.

Jacob: Can't we spot this effect without much trouble?

Rachel: The effect of measurement errors is not always easy to spot. Suppose interest rates are a random walk with a drift of 0.4 basis points per day. The expected rate increases by 2 basis points a week, or 1% a year. The first differences are a white noise process with a mean of 0.004%. If the interest rate is 8.00% on day 1, the forecasts for days 2, 3, 4, … are

{8.004, 8.008, 8.012, 8.016, 8.02, 8.024, 8.028, 8.032, 8.036, 8.040 …}

With two decimal place accuracy, the forecasts are

{8.00%, 8,01%, 8.01%, 8.02%, 8.02%, 8.02%, 8.03%, 8.03%, 8.04%, 8.04% …}

The first differences are an oscillating series:

{0.00, 0.01, 0.00, 0.01, 0.00, 0.00, 0.01, 0.00, 0.01, 0.00, … }

The sample autocorrelations are negative for lags 1 and 3, positive for lags 2 and 4, and +1 for lag 5. We might model this series as an AR(5) process instead of white noise. The spurious process is not always simple year to year oscillation.

Jacob: Can we spot this effect from the consistent pattern?

Rachel: When we combine measurement error and stochasticity, the patterns are hard to spot. As an exercise, add a random white noise process to the pattern above. As the stochasticity increases, it becomes harder to see the pattern.

Jacob: Don't the Durbin-Watson statistic and the correlogram smooth the stochasticity and uncover the true patterns?

Rachel: If the sample autocorrelations are caused by real factors, the Durbin-Watson statistic and the correlogram smooth the random fluctuations and uncover the true pattern. If the sample autocorrelations are caused by measurement error, the Durbin-Watson statistic and the correlogram have the same effect: they smooth the random fluctuations and uncover the measurement error, which now looks like a true pattern. The statistical tests do not differentiate between measurement error and true patterns.

One Difference Too Many

Jacob: You mention that taking too many differences can cause spurious effects. What does this mean? If first differences are a white noise process, aren't second differences also a white noise process?

Rachel: If interest rates are a random walk with a drift of zero, the *first* differences are a white noise process with a mean of zero. We examine the time series process for the *second* differences.

Suppose the *second* difference in period *j* is positive, such as +Z. We estimate the likely *second* difference in period *j+1*.

~ The expected *first* difference in periods *j* and *j*-1 is zero, since this is a white noise process with a mean of zero.
~ The *second* difference in period *j* is positive, so the *first* difference may be high in period *j* or low in period *j*-1.
~ The pattern in a white noise process is symmetric, so our estimate of the *first* difference is –½Z in period *j*-1 and +½Z in period *j*.
~ The *expected first* difference in period *j*+1 is zero, so the *expected second* difference is –½Z.
~ The autocorrelation for lag 1 is negative.
~ This is AR(1) with $\phi_1 < 0$, or an ARIMA(1,2,0) process with $\phi_1 < 0$

Jacob: If taking differences can lead to spurious patterns, why do the textbook authors take differences so often?

Rachel: Stochasticity makes it hard to see the pattern in the first differences. The authors *examine* the second differences. They use them only if the first differences are not stationary and the second differences are more stable. If the first differences are a stationary time series, we don't use second differences.

Linear Trend and Random Walks

Jacob: You mention trends vs random walks as another cause of spurious correlations. What does this mean?

Rachel: Stock prices, interest rates, commodity prices, business sales, and many other time series may be geometric or additive Brownian motion; that is, random walks with or without first taking logarithms. Some actuarial series, such as average claim severity, may be better modeled as linear or exponential trends.

Jacob: Average claim severity should be like stock prices and commodity prices. What causes this difference?

Rachel: The difference stems from sampling error. The prices for stocks and commodities are known with certainty, since the markets are liquid for these items. The expected average claim severity is estimated from a small sample of claims that occur.

~ If the claim severity trend is 10%, the expected claim severity in 20X8 is 10% greater than the expected claim severity in 20X7.

~ If we don't know the expected severity in 20X7 with certainty, because our sample size is too small, we may estimate the expected claim severity in 20X8 from a trend line fit to several past years.

If we model a trend as a random walk, the residuals have an oscillating pattern.

Jacob: Isn't a random walk just a trend process with stochasticity?

Rachel: A random walk is different from trend. Consider a linear *trend of zero* vs a random walk with a *drift of zero*. The linear trend of zero is white noise, not a random walk.

We relate this to the previous discussion. If we presume the process is a random walk, we take first differences to get a white noise process. But since the process is actually white noise, taking first differences gives an oscillating AR(1) process with $\phi_1 < 0$.

Jacob: What if the series has a non-zero trend? Does that change the autocorrelation?

Rachel: If the interest rates have a trend of 10 basis points a period and the random walk has a drift of 10 basis points a period, the logic remains the same.

Jacob: The previous discussion shows that first differences of a white noise process are a random walk with a negative ϕ_1. If we model a linear trend as a random walk, what happens?

Rachel: Consider first a linear trend of zero, which is a white noise process. Suppose the expected interest rate is 8% each month, but we model it as a random walk with a drift of zero. Let the actual interest rate be 8% in January 20X6.

The forecasted interest rate is 8% for February 20X6 as well. Interest rates are stochastic, so the actual interest rate may be higher or lower. Suppose the actual interest rate is 9%. The stochastic term for February 20X6 is +1%.

If we model the interest rate as an AR(1) process with no drift, the forecast for March 20X6 is 9%. The expected interest rate is actually 8%, so the expected error term in our model is –1%. The residuals have a negative serial correlation.

NEGATIVE SERIAL CORRELATION

Jacob: What should we do if we find negative serial correlation?

Rachel: Examine whether each of these spurious effects may be occurring. Long-term interest rates change slowly. If we model daily changes in twenty year rates with two decimal point precision, we expect spurious effects from measurement error.

Check whether you are taking too many differences. Some candidates presume that taking differences can't hurt. The opposite is true. If a time series is a white noise process, taking differences gives spurious results.

Jacob: Suppose the sample autocorrelations are high. Shouldn't we take differences?

Rachel: We *examine* the differences; we don't always use them. Even if the time series is a white noise process, 10% of the time its sample autocorrelation function gives a Box-Pierce Q statistic above the 90% critical value.

Jacob: What if the differences give an AR(1) process with a negative ϕ_1?

Rachel: If the *differences* give an AR(1) process with a negative ϕ_1, we go back and re-examine the sample autocorrelation function of the time series.

NEGATIVE SAMPLE AUTOCORRELATIONS

Jacob: What if the sample autocorrelations are negative for many lags in a row? This is not statistical fluctuation or measurement error, and it occurs often in the student projects on interest rates. What causes this?

Rachel: If the sample autocorrelation function is *positive* for many lags in a row, such as lags 1 through 30, we assume the interest rates follow a random walk. We sometimes find that the sample autocorrelations are *positive* for lags 1 through 20 and then *negative* for lags 21 through 40. The negative autocorrelation usually stems from a positive drift in the first half of the series and a negative drift in the second half of the series (or vice versa).

Jacob: Why does the drift change? Isn't our goal to model this change in the drift?

Rachel: Two types of causes lead to the change in drift; one is real and one is spurious.

Real: Some business cycles have automatic corrective effects. For example, some economists believe that prosperity builds up forces leading to recession, and recession builds up forces leading to prosperity. Similarly, rising interest rates build up forces leading to declining interest rates, and vice versa. If this is true, we should model the alternating interest rates eras as an interest rate cycle.

Spurious: Other economists do not believe interest rates have cycles. But changes in regulation caused by changes in regulators give the *appearance* of interest rate cycles.

Illustration: Suppose economist Y believes the Federal Reserve Board should raise interest rates to strengthen the economy, and economist Z believes the Federal Reserve Board should reduce interest rates to strengthen the economy. If the chairmanship of the FED changes from economist Y to economist Z, the rising pattern of interest rates may shift to a declining pattern.

Jacob: Don't we model this by an ARIMA process?

Rachel: We model endogenous effects by ARIMA processes, such as prosperous years creating forces that bring a recession. We do not model exogenous changes, such as a new FED chairman. Instead, we separate the time series into two eras and model each one by a simple ARIMA process.