Algorithmic Trading
Session 5
Trade Signal Generation III
Mean Reversion Strategies

Oliver Steinki, CFA, FRM

Outline

Introduction

Where Do We Stand in the Algo Prop Trading Framework?

- As we have seen, algorithmic proprietary trading strategies can be broken down into three subsequent steps: Signal Generation, Trade Implementation and Performance Analysis

- The first step, **Signal Generation**, defines when and how to trade. For example, in a moving average strategy, the crossing of the shorter running moving average over the longer running moving average triggers when to trade. Next to long and short, the signal can also be neutral (do nothing). Using moving averages to generate long/short trading signals is an example choice of how to trade

- Sessions 3 – 6 deal with the question of deciding when and how to trade
  - **Session 3**: Finding Suitable Trading Strategies and Avoiding Common Pitfalls
  - **Session 4:** Backtesting
  - **Today's Session 5:** Mean Reversion Strategies
  - **Session 6:** Momentum Strategies

Introduction

Mean Reversion vs. Momentum

- Trading strategies can be profitable **only if securities prices are either mean-reverting or trending**. Otherwise, they are random walking, and trading will be futile. If you believe that prices are mean reverting and that they are currently low relative to some reference price, you should buy now and plan to sell higher later. However, if you believe the prices are trending and that they are currently low, you should (short) sell now and plan to buy at an even lower price later. The opposite is true if you believe prices are high

- Academic research has indicated that **stock prices are on average very close to random walking**. However, this does not mean that under certain special conditions, they cannot exhibit some degree of mean reversion or trending behaviour. Furthermore, at any given time, stock prices can be both mean reverting and trending depending on the time horizon you are interested in. Constructing a trading strategy is essentially a matter of determining if the prices under certain conditions and for a certain time horizon will be mean reverting or trending, and what the initial reference price should be at any given time

Mean Reversion

What is Mean Reversion?

- Most price series are not mean reverting, but are geometric random walks. The returns, not the prices, are the ones that usually randomly distribute around a mean of zero. Unfortunately, one cannot trade directly on the mean reversion of returns (One should not confuse mean reversion of returns with anti-serial-correlation of returns, which we can definitely trade on. But anti-serial-correlation of returns is the same as the mean reversion of prices). Those few price series that are found to be mean reverting are called **stationary**, and we will see the statistical tests (ADF test, the Hurst exponent and Variance Ratio test) for stationarity

- Fortunately, we can manufacture many more mean-reverting price series than there are traded assets because we can often combine two or more individual price series that are not mean reverting into a portfolio whose net market value (i.e., price) is mean reverting. Those price series that can be combined this way are called **cointegrating,** and we will describe the statistical tests (CADF test and Johansen test) for cointegration

- Also, as a by-product of the Johansen test, we can determine the exact weightings of each asset in order to create a mean reverting portfolio. Because of this possibility of artificially creating stationary portfolios, there are numerous opportunities available for mean reversion traders

Stationarity

Mean Reversion vs. Stationarity and Tests to Discover It

- Mean reversion and stationarity are **two equivalent ways** of looking at the same type of price series, but these two ways give rise to two different statistical tests for such series
- The **mathematical description of a mean-reverting** price series is that the change of the price series in the next period is proportional to the difference between the mean price and the current price. This gives rise to the **ADF test**, which tests whether we can reject the null hypothesis that the proportionality constant is zero
- However, the **mathematical description of a stationary** price series is that the variance of the log of the prices increases slower than that of a geometric random walk. That is, their variance is a sublinear function of time, rather than a linear function, as in the case of a geometric random walk. This sublinear function is usually approximated by τ^{2H}, where τ is the time separating two price measurements, and H is the so-called **Hurst exponent**, which is less than 0.5 if the price series is indeed stationary (and equal to 0.5 if the price series is a geometric random walk). The **Variance Ratio test** can be used to see whether we can reject the null hypothesis that the Hurst exponent is actually 0.5.
- Note that **stationarity is somewhat of a misnomer**: It doesn't mean that the prices are necessarily range bound, with a variance that is independent of time and thus a Hurst exponent of zero. It merely means that the variance increases slower than normal diffusion.

Mean Reversion
Augmented Dickey Fuller Test

- If a price series is mean reverting, then the current price level will tell us something about what the price's next move will be: If the price level is higher than the mean, the next move will be a downward move; if the price level is lower than the mean, the next move will be an upward move. The ADF test is based on just this observation. We can describe the price changes using a linear model:

$$\Delta y(t) = \lambda y(t-1) + \mu + \beta t + \alpha_1 \Delta y(t-1) + \cdots + \alpha_k \Delta y(t-k) + \varepsilon$$

- The **ADF test will find out if $\boldsymbol{\lambda = 0}$**. If the hypothesis $\lambda = 0$ can be rejected, it means that the next move of the asset is dependent on the current level and therefore not random

- The statisticians Dickey and Fuller described the distribution of this test statistic and tabulated the critical values for us, so we can look up for any value of $\lambda/\mathrm{SE}(\lambda)$ whether the hypothesis can be rejected at, say, the 95 percent probability level

- Since we expect mean regression, $\lambda/\mathrm{SE}(\lambda)$ has to be negative, and it has to be more negative than the critical value for the hypothesis to be rejected. The critical values themselves depend on the sample size and whether we assume that the price series has a non-zero mean $-\mu/\lambda$ or a steady drift $-\beta t/\lambda$. Most practitioners assume the drift term to be zero

Stationarity
Hurst Exponent

- Intuitively speaking, a "stationary" price series means that the prices diffuse from its initial value more slowly than a geometric random walk would. Mathematically, we can determine the nature of the price series by measuring this speed of diffusion. The speed of diffusion can be characterized by the variance

$$Var(\tau) = \langle |z(t+\tau) - z(t)|^2 \rangle$$

  where z is the log prices (z = log(y)), τ is an arbitrary time lag, and $\langle | \ldots | \rangle$ an average over all t. For a geometric random walk, we know that

$$\langle |z(t+\tau) - z(t)|^2 \rangle \sim \tau$$

  The $\sim$ means that this relationship turns into an equality with some proportionality constant for large τ, but it may deviate from a straight line for small τ. But if the (log) price series is mean reverting or trending (i.e., has positive correlations between sequential price moves), the last equation won't hold. Instead, we can write:

$$\langle |z(t+\tau) - z(t)|^2 \rangle \sim \tau^{2H}$$

- This is the definition of the Hurst exponent *H*. For a price series exhibiting geometric random walk, $H = 0.5$. But for a mean-reverting series, $H < 0.5$, and for a trending series, $H > 0.5$. As *H* decreases toward zero, the price series is more mean reverting, and as *H* increases toward 1, the price series is increasingly trending; thus, *H* serves also as an indicator for the degree of mean reversion or trendiness

Stationarity
Variance Ratio Test

- Because of finite sample size, we need to know the statistical significance of an estimated value of H to be sure whether we can reject the null hypothesis that H is really 0.5. This hypothesis test is provided by the Variance Ratio test. It simply tests whether

$$\frac{Var(z(t)-z(t-\tau))}{\tau\, Var(z(t)-z(t-1))}$$

is equal to 1. The outputs of this test are **h** and **pValue:** h = 1 means rejection of the random walk hypothesis at the 90 percent confidence level, h = 0 means it may be a random walk. pValue gives the probability that the null hypothesis (random walk) is true

Mean Reversion

Half Life of Mean Reversion I

- The statistical tests we described for mean reversion or stationarity are very demanding, with their requirements of at least 90 percent certainty. But in practical trading, we can often be profitable with much less certainty. In this section, we shall find another way to interpret the λ coefficient of the ADF Equation so that we know whether it is negative enough to make a trading strategy practical, even if we cannot reject the null hypothesis that its actual value is zero with 90 percent certainty in an ADF test. We shall find that **λ is a measure of how long it takes for a price to mean revert**

- To reveal this new interpretation, it is only necessary to transform the discrete time series Equation of ADF to a differential form so that the changes in prices become infinitesimal quantities. Furthermore, we ignore the drift and the lagged differences and end up with the Ornstein Uhlenbeck formula for a mean reveriting process:

$$dy(t) = (\lambda y(t-1) + \mu)dt + d\varepsilon$$

In the discrete form of this equation, the linear regression of $\Delta y(t)$ against $\Delta y(t-1)$ gave us λ. This value carries over to the differential form, but it also allows for an analytical solution for the expected value of $y(t)$:

$$E\big(y(t)\big) = y_0 \exp(\lambda t) - \frac{\mu}{\lambda}(1 - \exp(\lambda t))$$

Mean Reversion

Half Life of Mean Reversion II

- Remembering that λ is negative for a mean-reverting process, this tells us that the expected value of the price decays exponentially to the value $-\mu/\lambda$ with the half-life of decay equals to $-log(2)/\lambda$. This connection between a regression coefficient λ and the half-life of mean reversion is very useful for algorithmic traders.
  - First, if we find that λ is positive, this means the price series is not at all mean reverting, and we shouldn't even attempt to write a mean-reverting strategy to trade it.
  - Second, if λ is very close to zero, this means the half-life will be very long, and a mean-reverting trading strategy will not be very profitable because we won't be able to complete many round-trip trades in a given time period.
  - Third, λ also determines a natural time scale for many parameters in our strategy. For example, if the half life is 20 days, we shouldn't use a look-back of 5 days to compute a moving average or standard deviation for a mean-reversion strategy. Often, setting the lookback to equal a small multiple of the half-life is close to optimal, and doing so will allow us to avoid brute-force optimization of a free parameter based on the performance of a trading strategy

Mean Reversion

Trading Mean Reversion

- Once we determine that a price series is mean reverting, and that the half life of mean reversion for a price series is short enough for our trading horizon, we can easily trade this price series profitably using a simple linear strategy: determine the normalized deviation of the price (moving average divided by the moving standard deviation of the price) from its moving average, and maintain the number of units in this asset negatively proportional to this normalized deviation

- The look-back for the moving average and standard deviation can be set to equal the half-life

- You might wonder why it is necessary to use a moving average or standard deviation for a mean-reverting strategy at all. If a price series is stationary, shouldn't its mean and standard deviation be fixed forever? Though we usually assume the mean of a price series to be fixed, in practice it may change slowly, e.g. due to changes in the economy or corporate management. As for the standard deviation, recall that the Hurst exponent equation implies even a "stationary" price series with $0 < H < 0.5$ has a variance that increases with time, though not as rapidly as a geometric random walk. So it is appropriate to use moving average and standard deviation to allow ourselves to adapt to an ever-evolving mean and standard deviation, and also to capture profit more quickly

Cointegration

Cointegrated Augmented Dickey Fuller Test

- Unfortunately, most financial price series are not stationary or mean reverting. But, fortunately, we are not confined to trading those "prefabricated" financial price series: We can proactively create a portfolio of individual price series so that the market value (or price) series of this portfolio is stationary. This is the notion of cointegration: If we can find a stationary linear combination of several non-stationary price series, then these price series are called cointegrated. The most common combination is that of two price series: We are long one asset and simultaneously short another asset, with an appropriate allocation of capital to each asset.

- Why do we need any new tests for the stationarity of the portfolio price series, when we already have the trusty ADF and Variance Ratio tests for stationarity? The answer is that given a number of price series, we do not know *a priori* what hedge ratios we should use to combine them to form a stationary portfolio

- Just because a set of price series is cointegrating does not mean that *any* random linear combination of them will form a stationary portfolio. But pursuing this line of thought further, what if we first determine the optimal hedge ratio by running a linear regression fit between two price series, use this hedge ratio to form a portfolio, and then finally run a stationarity test on this portfolio price series? This is essentially what the CADF test does

Cointegration

Johansen Test for Cointegration

- In order to test for cointegration of more than two variables, we need to use the Johansen test. To understand this test, let's generalize the discrete version of the ADF equation to the case where the price variable $y(t)$ are actually vectors representing multiple price series, and the coefficients λ and α are actually matrices. We will assume $\beta t = 0$ for simplicity. Using English and Greek capital letters to represent vectors and matrices respectively, we can rewrite the ADF equation as:

$$\Delta Y(t) = \Delta Y(t-1) + M + A_1 \Delta Y(t-1) + \cdots + A_k \Delta Y(t-k) + \varepsilon_t$$

  Just as in the univariate case, we do not have cointegration if $\Delta = 0$. Let's denote the rank of Λ as r, and the number of price series n.

- The number of independent portfolios that can be formed by various linear combinations of the cointegrating price series is equal to r. The Johansen test will calculate r for us in two different ways, both based on the eigenvector decomposition of Λ. One test produces the so-called trace statistic, and the other one produces the eigen statistic. We do not need to worry what they are exactly, since many programming packages will provide critical values for each statistic to allow us to test whether we can reject the null hypotheses that $r = 0$ (no cointegrating relationship).

- As a useful by-product, the eigenvectors found can be used as our hedge ratios for the individual price series to form a stationary portfolio.

Summary and Questions

- Mean reversion means that the change in price is proportional to the difference between the mean price and the current price. Stationarity means that prices diffuse slower than a geometric random walk
- The ADF test is designed to test for mean reversion. The Hurst exponent and Variance Ratio tests are designed to test for stationarity
- Half-life of mean reversion measures how quickly a price series reverts to its mean, and is a good predictor of the profitability or Sharpe ratio of a mean reverting trading strategy when applied to this price series
- One can combine two or more non-stationary price series to form a stationary portfolio, these price series are called cointegrating. Cointegration can be measured by either CADF test or Johansen test
- The eigenvectors generated from the Johansen test can be used as hedge ratios to form a stationary portfolio out of the input price series, and the one with the largest eigenvalue is the one with the shortest half-life

- Questions?

Sources

- Quantitative Trading: How to Build Your Own Algorithmic Trading Business by Ernest Chan
- Algorithmic Trading: Winning Strategies and Their Rationale by Ernest Chan
- The Mathematics of Money Management: Risk Analysis Techniques for Traders by Ralph Vince