

Back-Adjusting Futures Contracts

Bob Fulks
May 11, 2000

Back-Adjusting Futures Contracts

Bob Fulks
May 11, 2000

Introduction

To backtest a trading system for trading futures contracts, we would like to have a long duration of price data on which to test our trading system. The problem is that futures contracts expire periodically and the data for each contract lasts only a few weeks or months. So we need some way to create a long series of price data from a sequence of contract prices. This paper discusses the various ways it can be done and explores the advantages and disadvantages of each method. My experience is with only the S&P futures contract so I will use it as an example although the same principles apply to any future contract.

The Problem

The issue arises because the pricing of a futures contract is slightly different than the price of the underlying commodity since it includes other factors. This difference is called the "premium". For the S&P Futures this difference includes the cost of interest and the dividends of the S&P stocks.

The theoretical value is called the "fair value". It is the price at which investing in the underlying commodity has the same return as investing in the futures contract. For agricultural futures, the difference can include such things as storage costs, etc. For the S&P futures it is calculated as follows:

$$\text{Futures_Price} = \text{Cash_Price} * (1 + d * (i - v) / 365)$$

where:

i = interest rate for fair value calculation = about 5% now

v = dividend rate of the S&P cash index = about 1% now

d = days_to_expiration

At rollover, the days_to_expiration number jumps, causing the jump in the premium. With today's values, the jump in price is about 1%, or about 12 points on the S&P futures contract. In the examples that follow, for simplicity, we will use the 12 points as the size of the jump at expiration, keeping in mind that the actual number does depend upon the level of the index, and interest and dividend rates.

The interest rate term arises because with the futures contract, we are using leverage for which we are not paying interest. Thus, the cost of the interest gets built into the price of the futures contract.

The dividend term arises because with futures, we do not get the dividends that we would have gotten if we had bought all the stocks in the S&P. (This is approximated as an average but since dividends occur at different times, an accurate simulation would include the exact expected dividends and when they were paid.)

The actual price difference between the futures contract and the underlying can and does deviate from this theoretical value on a minute-by-minute basis. But the difference is usually very short lived because arbitrage players step in to buy one and sell the other and this activity keeps the relative prices closely tracking fair value. There are often fairly big differences in the reported daily "closing prices" since the futures markets close at a different time than the cash market and a lot can happen between the close of the two markets.

Thus, we have a discontinuity in the price of the futures contract at the expiration that must be accounted for in some way for backtesting.

Possible Methods

The primary alternatives are:

1. Splice contracts together without price adjustment.

This causes large price jumps at splice points. The price jumps cause two problems.

- They distort the operation of most trading indicators and automatic trading systems. For example a 14-day simple moving average would mix some of the prices from the old contract and the higher prices from the new contract giving a distorted picture of what is happening.
- They can cause large trading profits and losses to be included in backtest results that a trader would not have experienced in actual trading. For example, if our system was long one contract before the expiration and we sold after expiration, then the system would include the 12-point artificial jump in price for an apparent profit of $12 * 250 = \$3000$. If we didn't notice this, we might think our trading system was very profitable. Even worse, if we were optimizing the parameters of our trading system, some of the parameter values might cause us to hold the position through the transition and some might not, causing big discontinuities in the results and causing us to get false optimum parameter values.

As a result of these problems, this method is unsuitable for most backtesting.

2. Close out trades at "roll over" to the new contract at expiration.

This is what we have to do in real life at the end of a contract. To do this, we exit the old position and re-enter a new position in the new contract. But this is trickier than it appears in backtesting since we would have to make sure that any indicators or averages we are using do not simultaneously look at some old and some new price values.

Some people prevent this by using "bridge data". For example, if your trading system uses 14 days of past data as part of its calculations, you would need to artificially create 14 days of bridge data from the old contract, increased in price by the 12 points, to get the trading system initialized at the new contract values. This can be difficult to do if we are trying to test over a long price series including several contracts.

3. Splice contracts together with forward price adjustment at contract boundaries.

Subtracting the 12-point difference from each new futures price before we use it does this. On the following contract, the offset would be 24 points, etc. This method eliminates the need to manipulate old price data but causes the futures contract prices to keep diverging further and further from actual prices as time goes on.

4. Splice contracts together with backward price adjustment at contract boundaries.

This is the method most often employed. Adjusting the price of the old contracts by adding the offset (12 points in our example) accomplishes this. With this method the most recent continuous futures contract prices are same as the current contract prices, but previous contract prices are offset. Backward adjustment is much more difficult to do, because all past prices have to be recalculated at each new contract boundary.

Some people add the difference to past contracts and some multiply all old data by a factor, 1.01 in our example. Adding the offset is most common since it keeps the old price values lined up with the tick values. For example, if a contract moved in 1/32-point increments, we would like the increments to remain on 1/32 boundaries far back in time. Using the multiplying factor would slowly cause them to deviate.

Either of the method 3 or 4 above gives us a series of prices with no discontinuities. They maintain the bar-to-bar increments as they were in the original contracts. Since most trading systems use the bar-to-bar differences these methods will not distort the trades. You can test if your trading system is sensitive to the absolute level of prices by adding some constant, such as 100 points, to all prices. If the trades and profits remain the same, your system is independent of the absolute level of prices and will not be affected by the back adjustment process.

On some commodities, the back adjustment is negative and this can make the price go negative over time. This is easy to handle by simply adding a constant value to all prices when backtesting to assure that all prices remain positive.

5. Continuously adjust the price series over time ("Perpetual")

This approach adds some of the prices from the current contract with some from the next contract. The continuous futures contract value initially would be composed of a large percentage of the current contract and a small percentage of the next contract. As a contract expiration date becomes closer, a progressively smaller percentage of the current contract and larger percentage of the new contract is used. This results in a smooth blend from one contract price level to the next.

The calculations are shown below:

Assume we are merging the prices of the 9/98 and the 12/98 contracts.

Assume the following can represent the price in each contract:

$$P1 = C + F * d1 + s1$$

$$P2 = C + F * d2 + s2$$

where:

P1 is the price of the 9/98 contract on any bar

P2 is the price of the 12/98 contract on any bar

d1 = days to expiration on the 9/98 contract

d2 = days to expiration on the 12/98 contract

s1 = a residual value indicating how the actual price differs from "fair value" on the 9/98 contract

s2 = a residual value indicating how the actual price differs from "fair value" on the 12/98 contract

F is the "Fair Value Factor" defined as follows:

$$F = C * (i - v) / 365 \text{ (about 13\% then)}$$

C = the value of the S&P cash index = about 1186 then

i = interest rate for fair value calculation = about 5% then

v = dividend rate of the S&P cash index = about 1% then

By manipulating the algebra, you can show that the merged contract price, P_x , equals:

$$P_x = (91 * F) + (s1 * d1 / 91) + (s2 * (1 - d1 / 91))$$

This consists of three terms:

- $(91 * F) = (91 * 0.13) =$ about 12 then. This is the premium at the start of each contract due to the fair value calculation.
- $(s1 * d1 / 91) =$ the residual value of the 9/98 contract weighted by the days remaining on the 9/98 contract.
- $(s2 * (1 - d1 / 91)) =$ the residual value of the 12/98 contract inversely weighted by the days remaining on the 9/98 contract.

So there is a fairly constant offset due to the fair value rates, plus the weighted average of the residual price values of the two contracts. It will be about 1% higher than the S&P cash index at all times with no discontinuities and noisier because the volatility of the futures contract is greater than that of the cash index. The residual value of the combined price series will have volatility less than either of its two components since uncorrelated random variations will partially cancel.

The characteristics of the resulting adjusted price will be quite different than the price of either contract. The prices will not fall on multiples of the minimum price movement such as 1/32 in the case of bonds. We could round to those increments but this would introduce a new source of noise. The volatility of the resulting price will be less than the real contracts, particularly near expirations. Finally the price will always be about 1% higher than the S&P cash index. This is unlike the futures price, which decreases at about a 4% per year rate over time. All of these factors can affect the performance of a trading system and could make the results of backtesting differ from trading a single contract.

For a long-term system, one that stayed in the market for many months, using such data would be very convenient since it would eliminate the need to ever worry about expirations in backtesting.

For a medium-term system, say one that stayed in the market for many days to weeks, such distortions would be insignificant.

For short-term trading, say under a few days, they become important.

For a day-trading system they could be serious sources of error. Systems that trade on the volatility of the price data will not work as well on such data since the volatility of the merged data is less than that of either contract.

6. Use continuous adjustment for both backtesting and trading

Actually, trading a mix of the two contracts can eliminate the distortions mentioned above. This is only practical if you are trading perhaps over 20 contracts so that you can adjust the number of each in the proper proportions over time. It also adds a level of complexity to the trading system.

Special Cases

The processes described above are suitable for most cases. But there are special cases that cannot be covered by any general technique.

For example:

- "Beginning with the March 2000 contract, the notional coupons of the CBOT® Treasury bond, 10-year, 5-year, and 2-year Treasury note futures contracts will change from 8% to 6%." Thus, all past data on the 8% contracts will not be suitable for backtesting the 8% contract.
- "The Chicago Mercantile Exchange (CME) announced that it has increased the size of the random length lumber futures contract beginning with the January 2000 contract expiration. The contract is currently sized at 80,000 board feet and will increase to 110,000 board feet due to the size and capacity of railcars used to ship lumber." Thus, the prices for the old contract will not be applicable to the new contract.

There is no general way to handle such special cases. Frequently, guidance for how to make adjustments is available directly from the web sites of the exchanges where these commodities are traded.

Summary

There is no "best" method in an absolute sense. All the methods have advantages and disadvantages. The best method in particular cases depends on the type of market analysis being done or the type of trading strategy being used.

For instance, trading systems that compare current prices to distant past prices probably will back-test more realistically with Method 5, because there is less long-term price distortion. However, Method 5 tends to produce very unrealistic results with trading strategies that depend upon price waveshape measurement, because there will be significant medium and short-term waveshape distortion. Neural networks, exotic waveshape filters, wavelet methods, Fourier methods, and many other advanced methods do not back test realistically with that method. Method 3 or 4 will be much better to use in those cases. Method 4, using backward price adjustment is the most popular.

My Conclusions

I use Method 4 - back adjusted - because my objective is to simulate the behavior of my trading system on actual contract data as accurately as possible. My trading systems use the bar-to-bar price differences so are unaffected by adding a constant to all prices. Since method 4 (and method 3) maintains the exact bar-to-bar price changes of each original contract, this is what I prefer.

Needless to say, other people might consider other factors more important. I have no interest in having my systems perform more poorly on backtesting than they would have on actual contracts. (Maybe you could call it some sort of "stress testing".)

If I were using a long-term period, I would definitely use method 5 - the continuously adjusted contract but that is not what I am trading.

-Bob Fulks