Walk-forward analysis of a volatility breakout system

Progressively updating a trading system to recent market data provides insight into how robust the technique is and allows it to adjust to changing market conditions.

BY DANIEL FERNANDEZ

A frequent question that arises when developing a trading strategy is how to properly determine specific parameters. Because market conditions are constantly changing, it's unrealistic to expect a system with fixed parameters to perform at a consistently high level over long time periods. But how do you decide if a strategy should use a certain stop-loss, profit target, or indicator period length rather than another? What's the best way to adjust parameters over time?

FIGURE 1: WALK-FORWARD TESTING

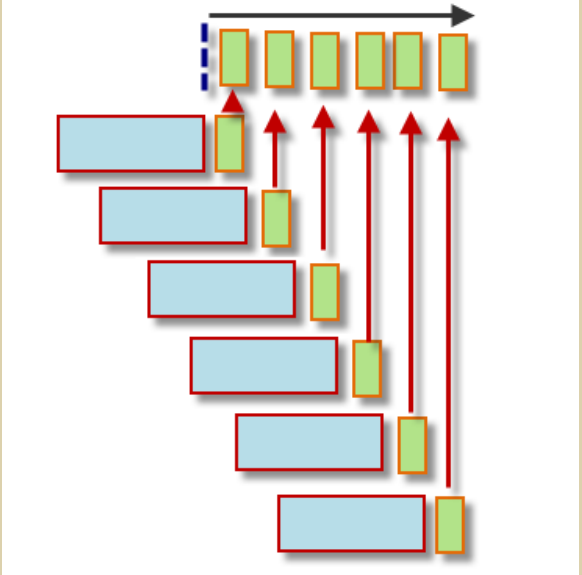

Walk forward analysis applies optimized parameters from the most recent period (blue rectangles to the subsequent forward-test periods (green rectangles), re-optimizing the parameters at regular intervals.

The answer to these questions usually comes through optimization — the process of testing multiple parameter values on historical data to see which set gives the best results — "best" being something that can be measured different ways. However, optimization presents another problem: The potential for excessive "fitting" of a system to past market conditions, which is likely to result in far worse results in future trading. For example, say a trading system uses a moving average as one of its inputs and a historical test shows that a length of 24 days produced the highest net profit, but that all other lengths between 20 and 30 days performed poorly. Using a 24-day moving average length would be a poor choice for future trading because it is likely a fluke; this parameter selection is merely fitted to the past data conditions, which are unlikely to be repeated in the future.

Walk-forward analysis is one of the most powerful techniques for evaluating whether a strategy can survive the parameter selection process — that is, whether the idea of optimizing and using a given set of "best" parameters is justified. In this procedure a strategy is optimized for a period of x bars and the results of the optimization are used for y bars in the future, at which point the procedure is repeated to obtain the parameter settings for the subsequent trading period. This process of "walking" the strategy forward through successive sets of new data is designed to keep the system as up to date as possible with changing market conditions.

Figure 1 illustrates how this is carried out. The blue rectangles represent the optimization periods and the green rectangles represent the subsequent forward-testing

periods to which the optimized parameter settings are applied. By running a 10-year test with constant re-optimization and the use of the "best" parameters through every new forward-testing period, we can get a very good idea of how our strategy handles changing parameter values, and whether doing this makes sense.

TABLE 1: STRATEGY STATISTICS

Annualized return	19.65%
Max. drawdown	32.87%
Win percent	58%
Reward/risk ratio	1.41
Ulcer Index	10.79
Profit factor	1.94
No. of trades	171

The strategy's performance statistics were solid — an average compound yearly profit of 19.65 percent and a maximum drawdown of 32.87 percent.

The system

To illustrate the walk-forward analysis process, we will use a daily time frame trend-following strategy (based on volatility breakout) that includes two parameters. Basically, the system takes a position in the direction of the most recent day's momentum (as determined by whether the close is above or below the open) as long as that momentum exceeds a certain threshold, the logic being that such a move will be followed by more price movement in the same direction:

1. Go long if the most recent day's close is above the open and the open-to-close range is greater than *x* pips.
2. Go short if the most recent day's close is below the open and the open-to-close range is greater than *x* pips.
3. Close a position after it has been open for *y* days.
4. Only one trade can be open at a time.
5. Trade 0.00001 times the current account balance per signal.

This type of strategy is perfect for illustrating walk-forward analysis because it has a limited number of parameters that can be optimized quickly and conveniently. Other than the position-sizing rule (step 5), the strategy has no inherent dynamic capabilities (such as volatility adapted indicators). As a result, the only way the system can possibly adapt to evolving market conditions is by changing parameters *x* (the intraday move threshold that determines trade entry) and *y* (the holding period).

The walk-forward process

The strategy was applied to the Euro/U.S. dollar pair (EUR/USD) over a 10-year period starting in January 2001 and ending in January 2011, with optimization of parameters x and y conducted every 50 calendar days on the most recent 300 calendar days. Parameter x was optimized from 10 and 500 pips in 10-pip increments (10, 20, 30, 40…500), while y was optimized between one and 50 days in one-day increments (1, 2, 3, 4…50). The best result for each optimization period was defined as the largest net profit; the parameter set that produced this result was used for the subsequent segment of the forward test (the next 50-day period). Trading costs of 2 pips per trade were applied throughout the test.

The primary advantage of this analysis is that it eliminates hindsight: Although we don't know which parameters will be used for each upcoming 50-day period, we have a very clear process for establishing how to derive these values from the previous data. In the end we have a result that tells us how our strategy will behave if we had applied a constant set of rules for the selection of parameters through the whole test.

Results

Figure 2 shows the system's equity curve using the parameters derived from the walk-forward process. The system was profitable overall, indicating the technique has an edge that is maintained in the long-term based on the constant re-optimization of the system's parameters. Table 1 shows that the strategy's basic performance metrics were

FIGURE 2: EQUITY CURVE

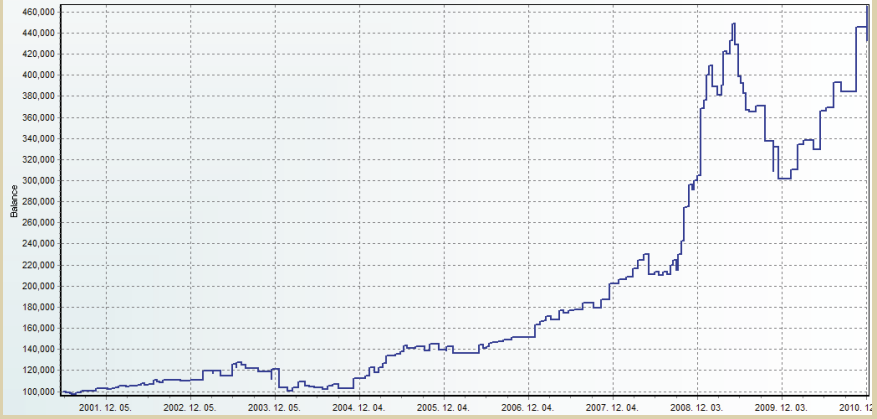

The system was profitable overall, with one major drawdown followed by a new equity high at the end of the test window.

FIGURE 3: MONTHLY RETURNS

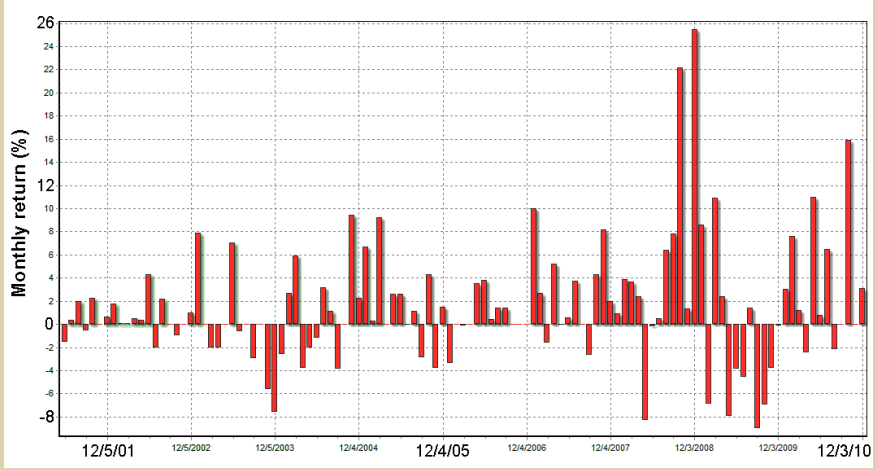

Re-optimization helped the strategy continue to perform during the financial crisis.

FIGURE 4: EVOLVING PARAMETER VALUES

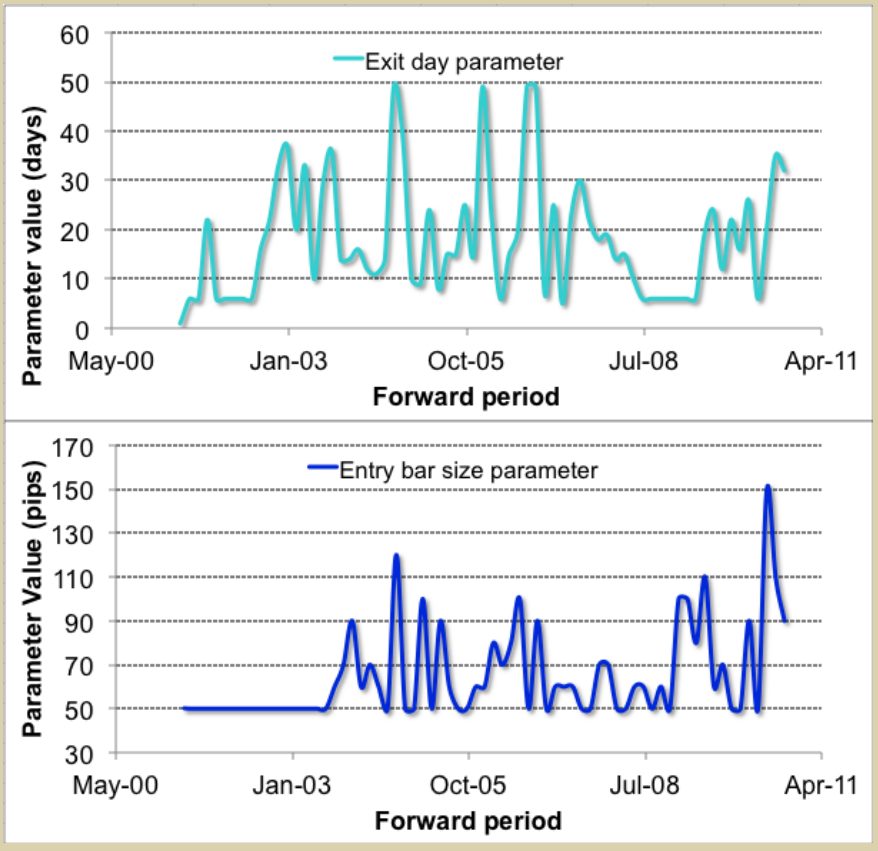

The exit (top) and entry (bottom) and parameters oscillated significantly over the course of the test.

quite good, with an average compound yearly profit of 19.65 percent and a maximum drawdown of 32.87 percent. The maximum drawdown length was 644 days, in line with what we would expect for a trend-following system on the daily time frame.

The monthly returns in Figure 3 show the best-performing conditions for this system have actually developed within the past three years (despite the most significant drawdown also occurring during this period). This highlights how re-optimization allowed the strategy to function during the financial crisis even though such extreme conditions had not previously appeared in the test period. When a strategy fails to produce profitable results in walk-forward analysis, it implies the strategy's best parameter choice always lagged the changes in market conditions. Such systems have a high probability of always giving curve-fitted results and there's a high chance their profitability is entirely based on hindsight.

It's also interesting how the optimum parameters evolved over time. These changes show us how market conditions have changed and if the system has actually adapted. Figure 4 shows the entry (x, bottom) and exit (y, top) parameters oscillated substantially throughout the test period, implying that the strategy's character has varied significantly in response

to changing market conditions. Table 2 shows the entry parameter ranged from 50 to 150 pips, had a median value of 60, and a mode (most frequently occurring value) of 50. (The estimated mode, which derives a hypothetical mode value based on a data set's median and average values, was also 50.) The exit parameter (number of days until exit) ranged from 1 to 49, with a median of 15 (average 18) and a mode of 6 (estimated mode of 9).

TABLE 2: PARAMETER RANGES

	x (entry)	y (exit)
Max.	150	49
Min.	50	1
Avg.	65	18
Med.	60	15
Mode	50	6
Est. mode	50	9

Over time, the optimum value for the entry parameter (x) was a little more stable than the optimum exit parameter (y).

When extreme drastic market changes occurred (e.g., during the financial crisis) the system adapted by drastically reducing the time to exit trades (from 30 to six days between July 2007 and March 2009) and increasing the threshold for opening a trade (from 50 to 100 pips over the same period) — in essence adapting the system to the fact that larger, shorter-term moves were the best bet.

Walk-forward analysis offers a way to build a trading strategy without any parameter-selection bias. You can evaluate a strategy over a long-term period using a defined set of rules for the choosing parameters, removing much of the potential hindsight and curve-fitting inherent in regular trading system creation. Using simple trading techniques, such as the one described here, and evaluating them through walk-forward analysis is an excellent way to determine their level of robustness.

For information on the author, see p. 4.