TRADING TECHNIQUES

Using fractals in forex

BY LESLIE K. MCNEW

Traders are well aware of market-based fractal relationships — spatial similarities that can be captured across scales. All you need are the right tools.

Here, we will illustrate the profitability of trading a currency position using strategies based on the Fractal Market Hypothesis as discussed by Edgar Peters and Benoit Mandelbrot. We'll look at the fractal from the slant of a time series analysis provided by Mandelbrot in 1963.

Capturing trends across time frames
How market periods are self-similar
Forex profits on short-term time frames

Mandelbrot found that cotton prices (1900-1963) were not normally distributed and instead showed clusters around the mean with a greater frequency of extreme variations (the tails) than that found in a normal distribution. This type of distribution is known as leptokurtic: A distribution that displays a positive value of excess kurtosis or sharpness of the peak of the graph of distribution. In other words, it has a higher peak than a normal curve and "fat tails" or higher density of values at the extreme end of the probability curve. Fat tails imply greater risk and suggest a nonlinear stochastic process. Assets that exhibit price jumps also display fat tail distributions.

Mandelbrot's analysis led him to coin the term "fractal," although he did not provide a concise definition. Fractals are not limited to geometric patterns found in nature (some common fractals include seashells, snowflakes, ferns, coastlines and broccoli), but can also describe processes in time.

Fractals exhibit two quantifiable characteristics: Self-similarity and the fractal dimension. Self-similarity means that the parts are related to the whole. Peters puts it best: "The object or the process is similar at different scales, spatial or temporal, statistically. Each scale resembles other scales, but is not identical."

An object is said to be self-similar if it looks "roughly" the same on any scale. For this discussion, we assert that the trends found on a four-hour spot euro candlestick chart are fractal shapes: Each trend roughly resembles other trends, but they are never the same.

The fractal dimension measures how, in our particular case, a time series (a set of historical data) deviates. A line has dimension of 1, a plane has a dimension of 2, and a cube has a dimension of 3. A random line has a fractal dimension of 1.5. If a fractal dimension of a time series is greater than 1 but less than 1.5, then this particular time series exists between a straight line and a Gaussian random walk. Again, Peters proposes an excellent definition: "Regarding a time series, the fractal dimension measures how jagged the time series is."

For more from Leslie, go to **futuresmag.com/McNew**

We accept the fractal market hypothesis as stated by Peter and discussed below. Various empirical studies show that financial assets produced skewed and fat tail return distributions (Mandelbrot, 1963; Fama, 1965; Hols, et al., 1991). In fact, the frequency distribution of currency returns has a higher peak and fatter tails than U.S. stocks or bonds. We define a short-term investment horizon as a period of less than five years and a long-term investment horizon of greater than four years.

Portraying the market in five basic points:

1. The market is stable when it consists of investors covering a large number of investment horizons. This ensures that there is ample liquidity for traders.
2. The information set is more related to market sentiment and technical factors in the short-term than in the longer-term. As investment horizons increase, longer-term fundamental information dominates. Thus, price changes may reflect information important only to

that investment horizon.

3. If an event occurs that makes the validity of fundamental information questionable, long-term investors either stop participating in the market or begin trading based on short-term information. When the overall investment horizon of the market shrinks to a uniform level, the market becomes unstable. There are no long-term investors to stabilize the market by offering liquidity to short-term traders.
4. Prices reflect a combination of short-term technical trading and long-term fundamental valuation. Thus, short-term price changes are likely to be more volatile or "noisier" than long-term trades. The underlying trend in the market reflects changes in the fundamental (economic) environment. There is no reason to believe that the length of short-term trends is related to the long-term economic trend.
5. If a security has no tie to the economic cycle, then there will be no long-term trend. Trading, liquidity and short-term information will dominate.

Setting up the trade

Here are the basic facts of our trade scenario:

- Book balance $10,000
- Position size €10,000
- Instrument: Spot €. No transactions fees are paid when trading spot forex.
- Backtested data: Spot four-hour data (Jan. 1, 2007 to June 30, 2013). Data provided from www.fxcm.com
- Trading periods: Execute only on four-hour candlestick window. Testing on data with 1 a.m., 5 a.m., 9 a.m., 1 p.m., 5 p.m. and 9 p.m. candlestick windows. The importance of the timing of the four-hour candlestick is stressed. Monthly U.S. economic information is released at 8:30 a.m. and 10 a.m. It is important to note that the model transaction will occur after/before possible periods of market stress (like on the release of the monthly employment data).
- Book leverage: Approximately 1.30
- Trend indicator: Exponential moving averages (EMA). Fast EMA: 10-period (10 periods of four-hour blocks of data). Slow EMA: 20. This pair was backtested as optimal for this currency and this time interval.
- Transaction times: At the open of each four-hour candlestick. No other transactions are allowed
- Transaction limit: 40 pips per €10,000 position; each pip is worth $1
- Transaction stop: 20 pips

ORDER OF ANALYSIS

We do not take a position if the EMAs cross during a four-hour window. The position is only entered at the open of the next bar.

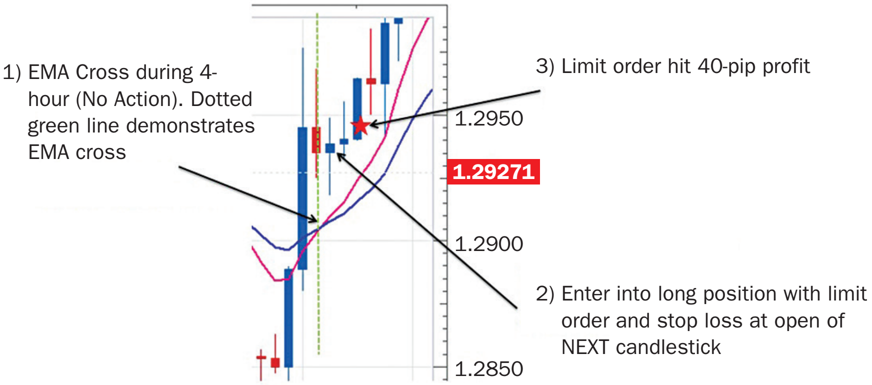

WHICH CAME FIRST?

On some candlesticks, we can't be sure if the profit limit or stop loss was hit first.

L	M	N	O	P	Q	R
Limit Order	Limit Order Hit	Limit ORDER HIT for TRADE		Stop Loss	Stop Loss Hit	Stop Loss for Trade Hit?
20				10		
0.002				0.001		
1.1453	NO	FALSE		1.14830		Default
1.1844	YES	TRUE		1.18140	YES	TRUE
1.1844	NO	TRUE		1.18140	NO	TRUE
1.1844	NO	TRUE		1.18140	NO	TRUE

We have examined a set of trades that are low risk, provide consistent low returns with a leverage of less than 1.5 and can be automated, which ensures low human capital fees. We consider this group of trades the "annuity" trade of the portfolio, or the first step of a return pyramid for a speculative portfolio. In terms of a baseball metaphor, this model is the first base of firm profit and not a home-run trade.

As the model operates in the short-term, we use technical indicators, particularly EMAs, to indicate the possibility of a trend. This is our only attempt to create some logic out of the noise that is produced at the short end of the market. It has been shown that short-term investors rely heavily on technical indicators.

Think of the greed and fear patterns (the positive and negative price movements) of the four-hour candlestick chart as a country's coastline. Determining the length of a country's coastline is not as simple as it appears, as first considered by L. F. Richardson (1881-1953) and sometimes known as the Richardson effect (Mandelbrot, 1983). In fact, the answer depends on the length of the ruler you use for the measurements. A shorter ruler measures more of the sinuosity of bays and inlets than a larger one, so the estimated length continues to increase as the ruler length decreases.

Traders do not know the optimal "ruler" to use to catch the maximum amount of profit for each "inlet" of price

METHOD COMPARISON

Here we can see a quick comparison of the three methods of determining whether the stop or limit was hit on the few instances the candlesticks were large enough to capture both.

	Candlestick method	Optimistic method	Conservative method
Year		Cumulative returns	
2007	3.04%	3.04%	3.04%
2008	1.75%	2.91%	-2.33%
2009	5.53%	6.60%	2.19%
2010	7.41%	7.25%	4.47%
2011	2.52%	3.46%	0.56%
2012	6.89%	7.65%	6.66%
2013	0.46%	0.89%	-0.52%
BT ITD	**30.840%**	**36.240%**	**14.640%**
6 months	**0.461%**	**0.889%**	**-0.521%**
Avg. year return	**3.94%**	**4.54%**	**2.01%**

CUMULATIVE RESULTS

Our basic model provides consistent returns over time in the euro.

Cumulative summary statistics fractal model A EUR/USD 01/01/2007 - 06/30/2013 Candlestick method returns	
Best month	1.86%
Worst month	-1.36%
Positive months	75.64% 59
Negative months	24.36% 19
Trades win	45.16%
Trades lose	54.84%

movement (trend). Our ruler is the limit order and our inlets are the trends of the four-hour market as depicted by the 10- and 20-period EMAs. Each inlet has two legs: The long trend (10 EMA > 20 EMA) and the short trend (20 EMA > 10 EMA). Because the optimal limit order is not known and the frequency and magnitude of each inlet are not consistent, we aim for profitability by taking a small bite out of each leg of each market wave.

The optimal EMA period lengths, limit orders and stop order amounts for the trading model were determined through backtesting using data from 2007 to June 2013 (four-hour data from www.fxcm.com). Regarding limit orders and stop orders, we looked for a combination that supplied consistent profits with low risk levels. Our backtesting return analysis for this model is found in "Cumulative results" (above). All returns are produced using CFA-recommended methodology: Geometrically linked returns.

Regarding the use of leverage in this model, the mean hedge fund industry leverage is approximately 2.13 with a standard deviation of 0.616. Hence, we sought to construct a model that targeted this industry average.

Trade mechanics

A transaction only will be considered at the open of each four-hour window and, if necessary, executed. This means that there are six four-hour candlestick windows in a daily 24-hour period and, thus, there are only six possible periods of transaction.

To open a trade, evaluation of the pair of EMAs occurs. If 10 EMA > 20 EMA, then a long position is taken. If 20 EMA > 10 EMA, then a short position is taken. At the time of trade entrance, both limit and stop orders are placed 40 and 20 pips away from the entry price, respectively. The transaction is automatically exited when the limit or stop order is hit. Our currency platform is FXCM and these orders are executed with little slippage except in the rare instances of complete market chaos.

There is no transaction on the anticipation of an EMA pair cross. The EMA signal must be firmly in place for trade entrance (see "Order of analysis," page 29). Because of this, the model is considered a "lagging" one. The trade entry only occurs firmly after the EMA signal and only at the time of the open of the four-hour candlestick window. The exit of the trade occurs on a pre-set limit or stop-order basis, or change in trend direction.

The model would be more profitable if the transaction took place as close to the actual EMA cross as possible, without the imposed time lag of execution only at the four-hour window. However, our available dataset for backtesting limited us to the use of the four-hour candlestick for trade entrance.

The ambiguity issue

This model was designed to trade on and was historically optimized using four-hour EUR/USD candlestick data. We have found that backtesting the model produced, on some occasions, an ambiguity issue: It is impossible to tell in a transaction which was hit first in the candlestick data — the limit or the stop. In other words, if the candlestick price depth is such that both the limit and the stock were met inside the pricing period, it may not be certain which order was executed first.

But the ambiguity issue does not significantly alter the return profile of the model. For example, in 2007-2013 there were approximately 10,000 model entries and the question of ambiguity arose approximately 50 times. The spreadsheet screenshot in "Which came first?" (page 29) illustrates the small possibility of the "double count" of both the stop order and the limit order being executed within the same observation period.

There were three methodologies used to test the bounds of historical model returns in light of the ambiguity issue: The most optimistic method is to assume that all limit orders hit before stop orders. The most conservative method is to assume that all stop orders hit before limit orders. Then there's the method that we settled on, the candlestick method.

In implementing the candlestick method, if the model is in a long position and the four-hour candlestick closes positively, then we assume that the limit order was met before the stop order executes. However, if the candlestick closes negatively, then the stop order hits before the limit order is met. If the model is in a short position and the four-hour candlestick closely negatively, then the limit order hits before the stop order is met. However, if the candlestick closes positively, then the stop order hits before the limit order is met. Finally, if the four-hour candlestick closes in a neutral position, the opening price equals the closing price, we assume that the limit order is met before the stop order is hit.

The table "Method comparison" (left) displays the results of the three tests to resolve the ambiguity issue. Again, we use the candlestick method when displaying our results. We feel this is the most logical use of the data. In addition, because our model is actually a "lagging" model (we do not enter a trade until after the EMA cross signal), the model itself has a built-in degree of profit conservatism.

This simple model is an attempt to catch the self-similar trends located within the four-hour candlestick price set of the spot EUR/USD. We know that these trends occur, but we are not sure of their magnitude or their temporal periodicity. Using exponential moving averages as our trend indicator, we capture a small bite of each trend. Our goals are low cost of execution, low risk and stable returns. This model is meant to provide consistent returns to add to the net return of the firm; think trading this spot four-hour model on five currencies at an average yearly return of more than 15% with low to no human capital cost.

Leslie K. McNew is the executive in residence, University of Scranton, Kania School of Management. Reach her at lesmcnew@gmail.com. Backtesting of this model was provided by Hussien O. Saleh and Peyton E. Veith of the University of Dayton School of Business Administration. Model validation and supervisory work were provided by Christine Muench, credit and risk analyst at E.ON Global Commodities North America. Additional model validation was provided by Zachary Hadaway.