Graham Capital Management
Research Note, March 2018

The Speed of Trend-Following

Nastja Bethke[1], Daniel DeWoskin[1], Gary Dean[1], Erik Forseth[1]

Abstract
Trend-following strategies aim to profit from sustained directional moves in markets. A key decision a trend follower needs to make is what 'speed' they wish to be – in other words, do they want to capture short, intermediate or long-term trends – and to parameterize their models accordingly. In this paper we present typical formulations of trend-following strategies, and investigate how their speed is set by their parameterization. We also show that despite formulation differences, trend-followers can be very similar at their core, and therefore can be highly correlated. Using two well-publicized CTA indices, we consider what value might be added to a portfolio by pairing a typical trend-following strategy with a source of alpha that aims to capitalize on short-term market behavior.

[1] *Quantitative Research Analyst*

1. Trend-Following Strategies

The 'reactivity' of a trend following signal to changes in the overall direction of prices can be thought of as the *speed of trend-following*. By placing more weight on the recent past this reactivity is increased; conversely, greater weight given to prices further in the past reduces the relative impact recent observations have on the generated signal, making the strategy less reactive. Reactive signals are better able to capture short-term trends, but can be prone to whipsawing. Less reactive models are more steady, but might be late to recognize a significant price reversal.

One prominent trend-following strategy - the *moving average crossover* model - uses moving averages of past prices to generate a trading signal. The difference between two moving averages, a 'fast' and a 'slow' one, calculated over lookbacks of n_f and n_s days (with $1 < n_f < n_s$), respectively, determines whether to enter a long or a short position. Small values of these parameters corresponds to more reactive systems. If prices are trending up, this difference tends to be positive, as the average over the most recent n_f prices will tend to be larger than the average over the most recent n_s prices, and vice versa. Here we restrict ourselves to positions of unit magnitude, solely the sign of the trading signal determines the position. Figure 1 illustrates the behavior of the moving average crossover model: superimposed on the WTI crude oil price over the two last years are slow and fast moving averages of the price for different parameterizations $n_f \times n_s$. The times when the fast and slow moving averages coincide are marked with black dots; these are also the times when the trading signal, and thus the position, switches sign. Positive positions - arising when the difference between the fast and slow moving average is positive - are shaded in green (red shading indicates a negative position). These examples show both the overall similarity between the trading signals generated by the different parameterizations as well as their particular differences: the 10×30 moving average crossover switches positions more frequently and is seen to react to small price reversals that the other moving average crossover does not react to.

Breakout models, which compare the current price with a threshold to decide on whether to go long/short or neither, represent another popular type of strategy that can profit from price trends. The addition of a neutral state makes breakout models structurally different from moving average crossover models. The breakout threshold is determined by a past maximum or minimum price; this is typically termed a "price breakout". Alternatively, one can use a past price at a fixed lookback and a channel width (e.g. a fixed percentage or a rolling volatility estimate) to set an upper and lower price bound; this is called a "channel breakout". Figure 2 shows the behavior of a channel breakout for two different lookbacks and a fixed 5% channel half-width. While a moving average crossover's speed is determined by the pair $n_f \times n_s$, for channel breakout strategies the lookback is the defining factor. Again both breakout examples lead to comparable positions, with the shorter lookback resulting in faster position changes. Note also the similarity to the signals in the moving average crossover case: even though these models appear different superficially, they exhibit significant overlap, especially when suitable parameters are chosen.[1]

2. Correlation between Trend-Followers

When comparing trading strategies, the correlation between the resulting PnL is one important aspect that is taken into consideration, for example, when trying to achieve diversification or, on the contrary, to mimic the dynamics of some reference system. We therefore investigate the PnL correlation between trend-following strategies of different speeds. We backtest the respective signals using rolled futures prices for over 50 liquid markets - covering equity indices, fixed income, foreign exchange and commodities, with data starting in 2000. The PnL calculation incorporates estimated slippage costs.

First we focus on the moving average crossover formulation. A theoretical estimate for the correlation between two different moving average crossover signals and therefore strategies can be derived by assuming independent identically distributed (IID) market returns. Different pairs $n_f \times n_s$ effectively apply dif-

[1] In fact, the trading signal for both moving average crossover and breakout strategies can be expressed in terms of a weighted sum of past *returns*, with the parameters setting the length of history to be considered, as well as the weights applied to each historical return.

Figure 1. WTI crude price, 10×30 and 40×120 moving averages for the years 2016-2018. Position changes are marked with black dots, green-shaded areas and red-shaded areas correspond to positive and negative trading signals, respectively.

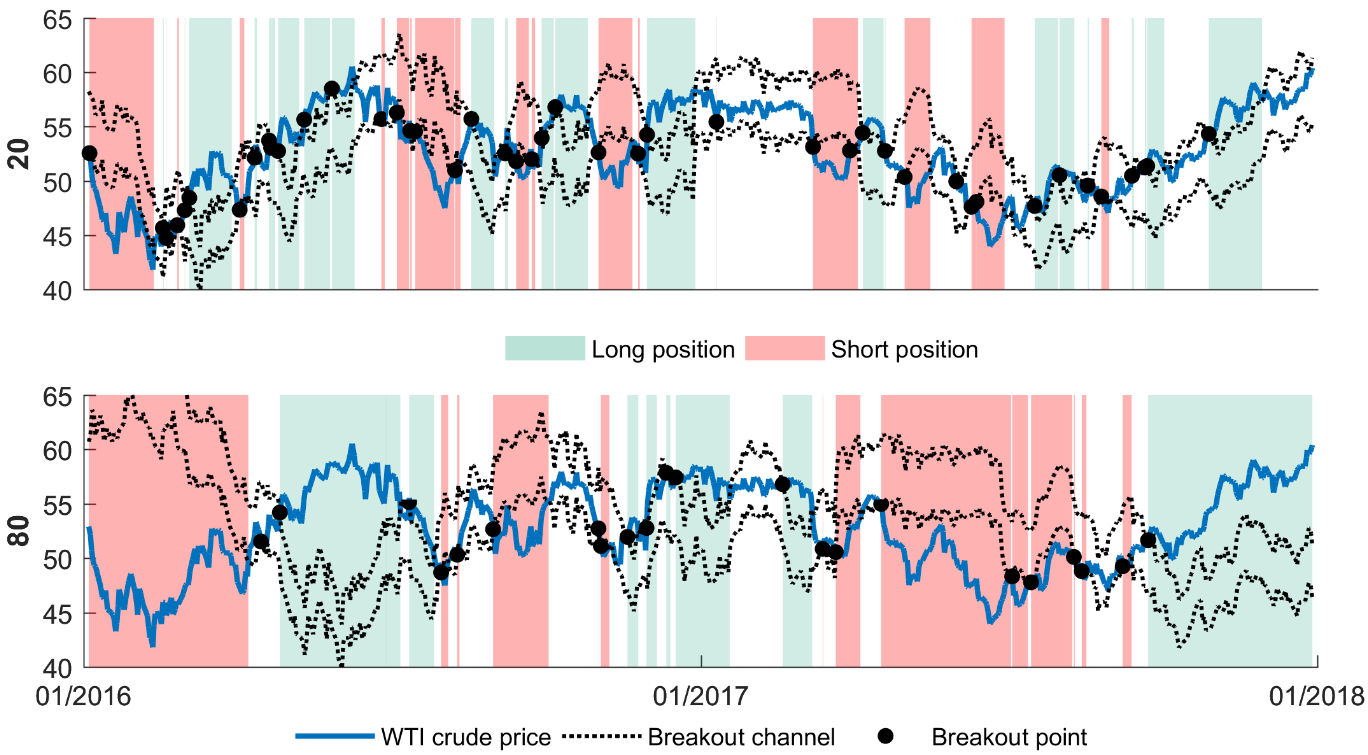

Figure 2. WTI crude price, 20- and 80-day lookback breakout channels for the years 2016-2018. Position changes are marked with black dots, red-shaded areas and green-shaded areas correspond to positive and negative trading signals, respectively. Note that the breakout model does not always take a position.

$n_f \times n_s$	10×30	20×60	30×90	40×120
10×30	100%	86%	69%	55%
20×60	-	100%	95%	86%
30×90	-	-	100%	97%
40×120	-	-	-	100%

Table 1a. Theoretical correlation between moving average crossover signals of different speeds $n_f \times n_s$.

$n_f \times n_s$	10×30	20×60	30×90	40×120
10×30	100%	82%	63%	49%
20×60	-	100%	90%	79%
30×90	-	-	100%	95%
40×120	-	-	-	100%

Table 1b. Actual correlation between moving average crossover strategies of different speeds $n_f \times n_s$.

lookback	20	40	60	80
20	100%	71%	58%	50%
40	-	100%	82%	71%
60	-	-	100%	87%
80	-	-	-	100%

Table 2a. Theoretical correlation between breakout signals with different lookbacks.

lookback	20	40	60	80
20	100%	74%	60%	54%
40	-	100%	83%	74%
60	-	-	100%	87%
80	-	-	-	100%

Table 2b. Actual correlation between breakout strategies with different lookbacks.

ferent weights to the same returns; for any two pairs we can calculate the theoretical correlation analytically. Table 1a shows the theoretical correlations between a range of pairs $n_f \times n_s$. For comparison, Table 1b shows the actual PnL correlations for moving average strategies of different speeds. These are in line with the theoretical estimates but tend to be lower. Differences between the theoretical and simulated correlations are due to the market returns not being IID and the PnL correlations reflecting an aggregate of different markets. Clearly correlations between moving average crossover strategies remain very high unless very different speeds are compared.

We show the analogous results for the channel breakout in Tables 2a and 2b, drawing a similar conclusion as for the moving average crossover: the signals/strategies for a wide range of lookbacks are highly correlated. Again, the agreement between theoretical and simulated correlations is good.[2]

Having considered the correlation structure for moving average crossovers and breakout models across different speeds, we briefly turn to the correlation between moving average crossover and breakout signals of a given speed. We find that each moving average crossover signal, specified by its pair $n_f \times n_s$, has a "best match" breakout signal, which will be about 95% correlated.[3] For the 10×30, 20×50 and 30×90 moving average crossovers, the breakout models that match best have lookbacks of around 30, 60 and 100 days, respectively. This high level of correlation for matching parameterizations is due to past prices being weighted in a similar fashion in the signal calculation of both models: both signals effectively take into account the same amount of information, and weight this information similarly.

3. How Important is the Speed of Trend-Following?

Despite different implementations and speeds, trend-following strategies tend to be highly correlated.[4] When aggregating different signals into a portfolio, we usually expect diversification benefits, that is improved performance when these additional signals capture a *distinct source of alpha*. Starting with a typical trend-follower, can one expect distinct alpha to come from much faster signals? Does merely modulating the speed of a raw signal result in diversification benefits? "Short-term trading" is a broad term for a diverse collection of potentially powerful strategies. Care must be taken when evaluating any model to see that it adds value to a portfolio; this holds true for short-term models as well.

For illustration, we consider two CTA indices: the SG Trend Index (TR) - an equally-weighted index tracking the ten largest trend-following CTAs, and the SG Short Term Traders Index (ST) - a volatility-weighted index tracking the performance of short-term CTAs with average holding period of less than ten days. We consider the ST Index as representative of a number of short-term funds. Figure 3 shows cumulative returns for each index since 2008 (re-scaled to 10% annualized volatility for comparison). We notice the qualitative similarities in the shape of the two curves: the correlation between the returns of these two series is 55% over this period. This matches our theoretical and backtested correlations between the fastest and slowest signals from Table 1 and Table 2 very well. The ST Index' correlation is thus consistent with that of a fast trend-follower. A priori, we expect this level of correlation to be too high for any substantial diversification benefit, especially given that one of the constituents has exhibited negative long-run performance. Indeed, the inclusion of the ST Index in a portfolio with the TR Index at any level hurts rather than helps. Considering the performance of the ST Index conditional on the performance of the TR Index, we find that its worst down days effectively coincide with those of the TR Index. Combined with the TR Index, the ST Index therefore offers little diversification benefit or protection.

[2]The theoretical estimate disregards the neutral state, which is equivalent to applying a zero channel width. The neutral state masks the underlying signal, and the correlation of an intermittent signal with a continuous one will be lowered. Considering aggregate PnL, however, the absence of a signal for any one market will not affect the overall PnL and correlation value much.

[3]This is true for a range of channel widths, starting from zero.

[4]High correlation does not necessarily imply similar performance, however.

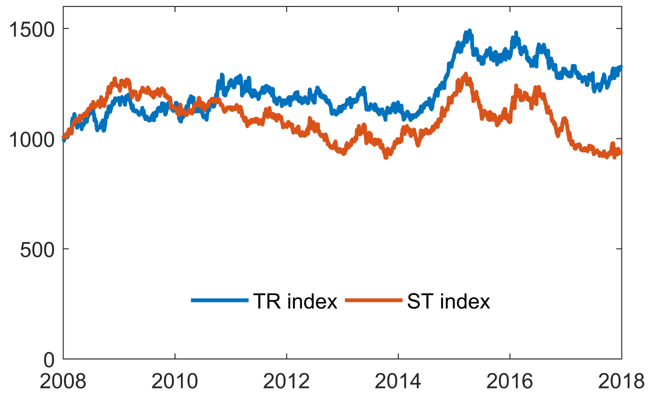

Figure 3. Cumulative returns of the TR and ST indices. For comparison, the returns of each index have been normalized to 10% annualized volatility.

4. Conclusion

Trend-followers of different speed and specification are highly correlated because of their common dependence on past prices. While we can design trend-followers of different speed, the exact speed of a raw signal - within a reasonable range - is not all that important in the performance of a strategy.

Short-term trading encompasses a collection of potentially effective strategies. The ability to capture market moves on various timescales can add robustness to a portfolio. However, when allocating to complementary strategies, investors must look at long-run and conditional correlations to assess possible diversification benefits. The ideal short-term strategy to pair with a trend-follower will be one that has zero, or preferably negative, conditional correlation on trend-following down days, while still retaining some non-negative, preferably positive, correlation on the up days.