CTAs: Which Trend is Your Friend?

Fabian Dori
Head of Portfolio Management at 1741 Asset Management Ltd.

Manuel Krieger
Director at 1741 Asset Management Ltd.

Urs Schubiger
Chief Investment Officer of 1741 Asset Management Ltd.

Daniel Torgler, CAIA
Associate at 1741 Asset Management Ltd.

1. Introduction

The occurrence of predictable trends within financial markets is inconsistent with the assumptions of classical financial theory and efficient markets hypothesis. Nevertheless, it can be empirically validated that market prices can be subject to trends. But, which trends should you measure? Which trend is your friend?

2. Measuring trends

In recent decades, the literature has recorded the development of a host of different trend measurement methods. They can essentially be divided into what are known as linear and non-linear approaches. Rather than focusing on the specific methodology of these techniques, the intention of this article is to point out that no single approach systematically delivers better results across all dimensions. In fact, the suitability of a given method depends on certain context-specific questions – which trend do you want to measure, and which characteristics of the measurement do you view as especially important? Which trend should be your friend? To simplify matters, we use moving averages in the following discussion. These can be understood intuitively, and they are easy to implement. They may therefore be regarded as representative of the various methods.

The first control variable that is used to specify the nature of the measured trend is the number of observations, which directly determines whether the short-term, medium-term, or long-term trend is to be measured. The more observations we process in order to measure the trend, the longer the measured trend will be. This point brings up one primary characteristic of the measurement, which seems to be unavoidable. The measurement of the trend is generally subject to a delay that, in turn, depends on the number of observations. In the case of moving averages, the trend over the defined time interval is subject by design to a delay corresponding to about one half of the observations. The measurement of a trend over 36 days is therefore delayed by about 18 days. It follows that the measured trend lags behind the observed price by half of the period that is used.

Second, in addition to the inherent delay in trend measurement, the precision of the measurement plays a key part. This factor reflects how precisely the trend component is extracted from the time series, without including too much information that is not relevant to the trend measurement, while ensuring that sufficient relevant information is included. The following related result applies as well: our ability to extract the trend improves as we include more observations – but again, this entails a longer delay (and vice-versa). We have to decide whether to focus on precision and to accept the disadvantage of the delay, or whether it is necessary to avoid a measurement delay at the expense of precision. When measuring shorter trends, it is advisable to select a method that focuses on precision, because the inertia of the trend means that the delay is insignificant in the shorter term. When dealing with long-term momentum cycles, however, precision is less significant due to the large number of observations; on the other hand, it is important to have the shortest possible measurement delay.

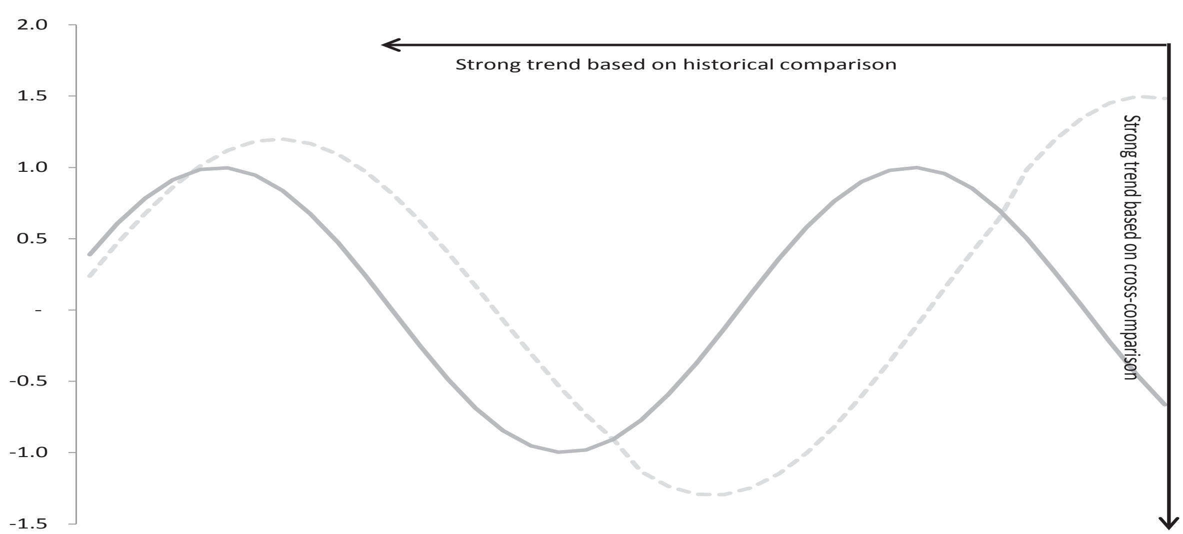

Exhibit 1: Historic comparison and cross-comparison of trend strength

Once the trend has been measured, the question arises as to how we can process this measurement in a context. A single measured value in isolation is of only limited use. So how can we find out whether the trend in question is strong or weak? Is that trend a strong friend? To answer this question, the current trend can be compared with its own history in order to determine its strength. Alternatively, we can compare it with current trend measurements for a group of different time series that we regard as comparable.

Exhibit 1 illustrates this principle and shows that the current trend strength for the dashed-line time series is similar not only in historical terms, but also as compared to the trend for the solid-line time series. This provides us with a dual view that will enable us to determine the true strength of the trend. These methods of calibrating the information obtained from a trend measurement constitute an essential factor for achieving intrinsic stability and for a sound and broad-based determination of the position size (derived from the trend measurement).

3. Selecting trends

Once the trend has been measured and calibrated, the next question that arises is: which trend is your friend? Based on the trend measurement, how should the position be sized appropriately on the market? Here, we draw a basic distinction between two signal categories – binary signals and steady signals.

Exhibit 2 summarises their main characteristics. Binary signals only make use of the prefixed (plus or minus) sign from the trend measurement. This application is based on the assumption that the trend strength does not contain any information that can be evaluated for financial purposes. In periods when market trends are unambiguous, the signal gives rise to very little turnover. But in the absence of a clearly predominant trend in the markets, turnover increases substantially. When executing the trades on the market, we should therefore bear in mind that there could be a resultant impact on market prices. Sensitivity to trendless markets is very high, because the fixed position size leads to an over-allocation in weak trends. In technical terms, the minimum of trend information is processed in this situation.

The second category (the steady signal) closely follows the trend measurement. The stronger the measurement, the larger the position will be. The assumption is that the persistence of a trend can also be applied to its strength, so a stronger trend implies better returns. Markets characterised by extreme trends (as in 2008, for example) will result in profitable returns. At the same time, the loss will be less than if a binary signal is used because positions in trendless markets tend to be smaller. This is compounded by the fact that turnover in trendless markets is lower than in markets characterised by trends, where positions are frequently adjusted according to the increase in strength. The impact on market prices is smaller because the adjustments made frequently are minor. With a steady signal, the information content is utilised in full.

Response function signals can be identified as a subcategory within the category of steady signals. These signals are influenced by a number of considerations. First and foremost is the concept of qualified selection, according to which the maximum trend is not necessarily the most profitable trend in every case. Consequently, for instance, the position can only be increased when the trend strength increases. But if the measured

	Binary signal	Steady signal	Response function
Assumption	Persistence and trend strength are not correlated	Persistence and trend strength are positively correlated	Non-linear correlation requires qualified selection
Information content	Lowest	Highest	Can be modelled
Signal frequency	Two signals, very frequent	Steady signals	Can be modelled
Position changes	Usually very major	Usually minor	Fairly minor
Market impact	Potentially substantial	Potentially low	Potentially quite low
Turnover	Very high for volatile markets, very low for quiet markets	Low for trendless markets, high for trending markets	Can be modelled

Exhibit 2: Characteristics of trend signals

trend continues to rise, the position is reduced again – in order to take the accumulated profits. Of course, a trend cannot continue to strengthen indefinitely and trend reversals often entail severe losses, especially if the trends in question are strong. In the extreme range of the signal, therefore, it is even possible in principle to take a position contrary to the signal. This is based on the assumption that exaggerated trends tend to produce mean reverting (random walk) behavior. However, one drawback of this signal function is that in markets with extreme trends (e.g., in 2008), the returns earned are not as high as those achieved with steady or even binary signals. Numerous variants are conceivable within this third category, based not only on statistical analyses, but also on fundamental convictions.

The frequency of the respective daily signals is an issue that arises in connection with all strategies. Signals that depend on the prefixed binary sign have two frequencies, but they are very numerous: there will frequently be positive as well as negative signals. On the other hand, scaled signals are determined by the distribution of the trend measurement and we may assume that the stronger a signal is, the less often it will occur. This fact is of fundamental importance when deciding how to deploy a trend-following strategy, and which trend to follow for this purpose.

4. Empirical results

In order to analyze signals of the three different types on an empirical basis, they were applied to a universe of 96 instruments (including 30 currency pairs, 19 equity indices, 11 government bonds, 8 money market, and 28 commodity instruments) over the period from 1993 until 2013. The data were recorded exclusive of transaction costs.

The various approaches are compared on the basis of two different risk-adjusted returns. First, they are compared using the information ratio (defined as annualised return divided by annualised volatility). The second comparison uses the ratio of annualised return to maximum drawdown.

Information ratio	**Currencies**	**Equities**	**Bonds**	**STIR**	**Commodities**
	0.35	0.43	0.54	0.94	0.65
Return / maximum DD	**Currencies**	**Equities**	**Bonds**	**STIR**	**Commodities**
	0.17	0.16	0.28	0.38	0.23
Return attribution	**Currencies**	**Equities**	**Bonds**	**STIR**	**Commodities**
Long side	+	+	+	+	+
Short side	-	-	-	-	+

Exhibit 3: Binary Signal

Information ratio	**Currencies**	**Equities**	**Bonds**	**STIR**	**Commodities**
	0.23	0.12	0.42	0.83	0.67
Return / maximum DD	**Currencies**	**Equities**	**Bonds**	**STIR**	**Commodities**
	0.09	0.04	0.22	0.26	0.36
Return attribution	**Currencies**	**Equities**	**Bonds**	**STIR**	**Commodities**
Long side	+	+	+	+	+
Short side	+	-	-	-	+

Exhibit 4: Steady Signal

Information ratio	**Currencies**	**Equities**	**Bonds**	**STIR**	**Commodities**
	0.33	0.40	0.61	0.99	0.82
Return / maximum DD	**Currencies**	**Equities**	**Bonds**	**STIR**	**Commodities**
	0.18	0.20	0.35	0.44	0.41
Return attribution	**Currencies**	**Equities**	**Bonds**	**STIR**	**Commodities**
Long side	+	+	+	+	+
Short side	+	+	-	-	+

Exhibit 5: Response Function Signal

As regards the information ratio, the empirical comparison shows that the binary method and the response function signal produce comparable results, although the latter signal may produce slight outperformance. For virtually all asset classes, on the other hand, the steady signal produces the lowest return per unit of risk entered into. The steady signal only performs marginally better than the binary method in the commodities asset class. These relationships can be attributed to the characteristics of the individual approaches. The essential difference between the three methods is explained by the conversion of trend strength into a trading position. With extreme signals, therefore, we can deduce from the results that the steady approach is likely to tend towards large positions that impact the risk/return ratio negatively in the event of trend reversals. Accordingly, very pronounced trends seem to go hand-in-hand with disproportionately high risk.

The ratio of annualized return to maximum drawdown may be regarded as a measurement of the signal's stability. Based on this yardstick, the response function approach prevails over the other two methods across all asset classes. This leads one to conclude that a slow increase on the inception of a trend together with a continuous reduction as excessive trends emerge can evidently make a substantial contribution towards the stability of the trend-following process.

Finally, it is interesting to note that the long positions for all signal types deliver a positive contribution to return regardless of the asset class, whereas clear differences are evident in the case of short positions. The binary signal performs worst in this regard. Commodities are the only asset class where the binary signal can generate positive returns from short positions. By contrast, the steady signal even manages to produce a positive return on the currency side as well. In this case, the short side is understood as a position against the interest-rate advantage that is normally present in the currency pair. The best performance is achieved by the response function signal, which is able to produce a positive contribution to return on commodities, currencies, and equities. Interest-bearing bonds and STIRS are the only classes where it delivers a negative return.

5. Conclusions

In summary, we may state that the time-frame used for the measurement represents the first key parameter. It was also demonstrated that a conflict between the precision and the delay of a trend measurement is an inherent feature of momentum-oriented models. Once a trend has been measured, the next step is to assess its strength. As well as using the trend's own history for this purpose, we can also correlate it with the trend of comparable time series. The information obtained from the trend measurement can be converted into a trading signal with the help of three functions, each of which displays different characteristics.

All of the foregoing boils down to the conclusion that our – "Which trend is your friend?" – cannot be answered solely on the basis of the data. Rather, we must first consider the various aspects in order to decide which method we regard as appropriate (in which environment). The approach is not "method follows trend," but "trend is defined by method." The length of the measured trend, the way the signal is compared and, ultimately, the conversion of the signal into a position are key factors that determine the character of the trend sequence, and they can be used for various CTA classifications.

The first classification relates to the length of the measured trends. It defines whether the trends to be measured are short-term, medium-term, or long-term. Although almost every manager uses a slightly different definition here, the rule of thumb for guidance purposes is that short-term trends range from intraday to one week, medium-term trends refer to a timeframe of between one and about six months, while long-term trends continue for more than one year. The second classification differentiates the method of trend comparison. The objective here is to ascertain whether the strength of the trend is calibrated in absolute terms, in relation to the history or on the basis of a cross-comparison. Finally, a distinction can be drawn according to whether a binary signal, a steady signal, or even a response function signal is used for the conversion into a trading position. These categories can be used to classify the majority of CTAs so that we can compare their performance more adequately.

Author Bios

Fabian Dori, M.A. HSG, CFA, FRM is Head of Portfolio Management at 1741 Asset Management Ltd. and has been a member of the executive board since 2013. Before joining 1741 in 2012, he was a quantitative analyst and portfolio manager at Wegelin & Co. Private Bankers, developing and managing quantitative investment strategies in the field of currencies, volatility, and multi-asset class. Fabian Dori holds a master's degree in quantitative economics & finance from the University of St. Gallen. A Swiss citizen, he speaks German, English, French, and Spanish.

Manuel Krieger, M.A. HSG is a director at 1741 Asset Management Ltd., member of the Global Diversification portfolio management team, as well as the fund manager of the 1741 Diversified Trends fund. Before joining 1741 in 2012, he was portfolio manager at Wegelin & Co. Private Bankers. Manuel Krieger has a Master's degree in Information Management from the University of Fribourg and was awarded top honors for his master thesis in the area of artificial intelligence. Manuel Krieger also holds a Master's degree in Finance from the University of St. Gallen. He is Swiss and speaks German, English, and French.

Urs Schubiger, M.Sc. ETH, lic.iur. is Chief Investment Officer of 1741 Asset Management Ltd. and member of the executive board since 2012. From 2009 onward he was responsible for Research, Modelling, and Technology in the Asset Management division of Wegelin & Co. Private Bankers. He was named a partner with limited liability of Wegelin & Co. in 2011. Before joining Wegelin & Co., he was a quantitative analyst in research and portfolio management for the U.S. Equity Statistical Arbitrage Group of UBS O'Connor in Chicago. Urs Schubiger holds a master's degree in mathematics from the Federal Institute of Technology (ETH) in Zurich and a master's degree in law from the University of Basel. A Swiss citizen, he speaks German and English.

Daniel Torgler, M.A. HSG, CAIA, is an associate at 1741 Asset Management Ltd. where he is responsible for the 1741 Global Diversification strategy. Before joining 1741 Asset Management Ltd. in 2012, he was a research analyst at Man Investments in Switzerland and the UK, a credit analyst in leveraged finance at Pemba Credit Advisers, and a portfolio manager at Wegelin & Co. Private Bankers. Daniel Torgler holds a Master's degree from the University of St. Gallen. Daniel Torgler is a Swiss citizen and speaks German and English.