Historical Performance of Trend Following

Jonathan Duke, Head of Investment Systems
David Harding, Founder and Chairman
Kate Land, PhD, Senior Scientist

Winton Capital Management

December 2013

Trend following is a well known trading strategy in which market forecasts are based solely on recent price movements. We document the performance of such strategies for a portfolio of large futures markets since the mid-1980s, and consider if there is any dependence on trading speed. In general we find trend following systems to be effective in forecasting future price movements, but we observe a significant and persistent decline in the forecasting ability of those with the fastest turnover.

1 Introduction

There is a long history of practitioners within the financial markets using 'technical analysis' to take investment decisions based on patterns in historical price data. These ideas became popular within the futures markets where professional investment advisors operate under the rubric of 'Commodity Trading Advisor' (CTA). The acronym CTA has now become almost synonymous with 'trend following' or 'momentum' strategies, where the pattern in question is simply whether a market price has risen or fallen in recent history. This tendency for prices to trend is described in statistical terms as 'auto-correlation' of returns.

It is only in recent history that trend following on major asset classes, via futures markets, has gained the attention of the academic community [1]. The focus of academic work has been on equity markets, where momentum is most commonly embodied as a cross-sectional effect used to predict relative, rather than absolute, returns of stocks [2, 3]. In this guise it still however exploits the existence of auto-correlation in returns, relative to the market. The evidence for momentum in equity prices is strong enough that it is now accepted as one of the standard equity 'risk factors' that can be used to explain the returns of any stock or portfolio, along with the market, value, and size factors [4, 5]. As it becomes better documented, returns from momentum signals are increasingly deemed to be 'beta' rather than 'alpha' in the world of equities.

In this study we document the performance of trend following strategies for different time horizons on data from the mid-1980s until today. Within the financial markets great efforts are ongoing to achieve the fastest possible communication speeds with exchanges, contributing to a perception that fastest is always best. However we find evidence that fast has not always been best in the world of trend following and futures markets, where we see a significant decline in the performance of the fastest strategies even before transaction costs are accounted for. In Section 2 we outline the markets that we use to document this effect and the rules of our trend following trading system. In Section 3 we document our results, and we conclude with a discussion in Section 4.

2 Data and Methodology

Futures

We use the front contract for 20 futures markets which are representative of their sector, highly liquid, global in their reach, and have a long history. The markets are listed in Table 1. For the subsequent analysis, we use the date range 1 Jan 1984 to 31 Oct 2013, and therefore we have nearly 30 years of results. 15 of our 20 assets have started trading by Jan 1984, and 18 by Sept 1986 and all 20 by Sept 1988 (Bunds is the last to start).

It is tempting to synthesise futures prices prior to these start dates in order to perform a longer historical

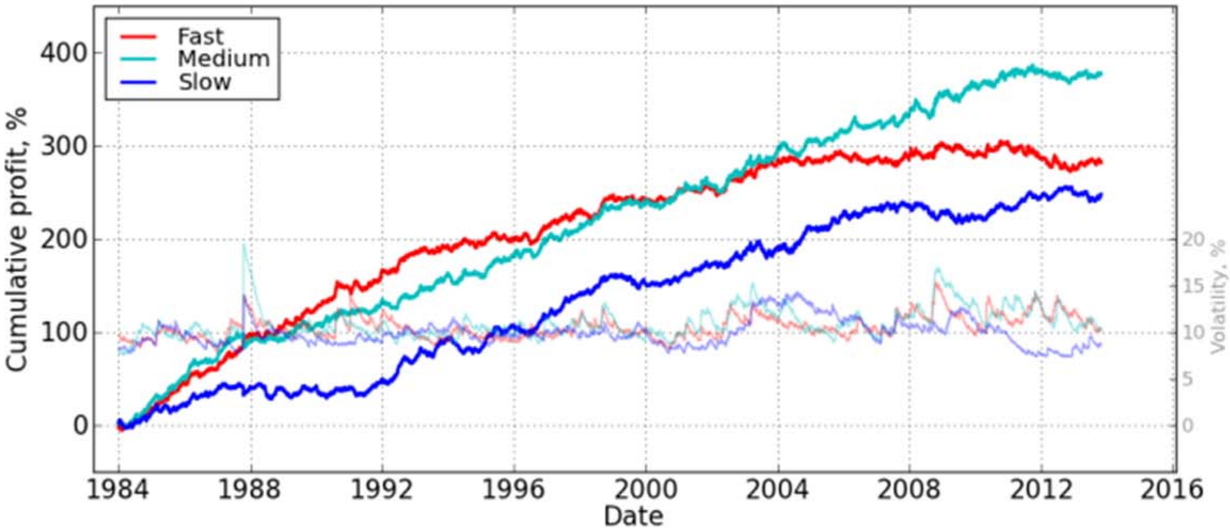

Figure 1: *The gross performance (not compounded) and the realised volatility (determined by a 100 day exponentially-weighted estimate) for our three trend following strategies. The Sharpe ratios (SR) are 0.87, 1.12, and 0.81 for the fast, medium, and slow strategies respectively for the above 30 year time period (the standard error on the SR is approximately 0.2 for 30 years of data). The realised volatility is around 10.5% for all three strategies.*

Fixed Inc	Indices	FX	Commodities
10y Tnote	S&P500	Euro	Gold
Eurodollar	Nikkei[3]	JYen	Live Cattle
Bunds[1]	Hang Seng[4]	BPound	Copper[7,8]
JGB2	ASX200[5]	SFranc	WTI Oil
Bonds	FTSE[6]	CDollar	Corn

Table 1: *The 20 futures in our trend following portfolio. These are (now) traded on the CME except: [1]Eurex; [2]TSE; [3]OSE; [4]HKFE; [5]SFE; [6]LIFFE; [7]LME; [8]we synthesise prices for a Copper futures contract with a fixed expiry date, rather than a fixed time to expiry as found on the LME.*

simulation. For the purpose of this paper we wish our results to be free from any uncertainty regarding the validity of such an extrapolation. Similarly, we restrict ourselves to results since Jan 1984, when most of our markets have started, in order to avoid the effects a changing number of assets can have on the portfolio-level performance.

Trading system

We use an exponentially weighted moving-average (ewma) cross-over system (the difference between a fast and a slow ewma of the price) to determine whether we go long or short in a market; we are long when the fast average is above the slow and *vice versa*. The positions are scaled by the inverse of a 20 day exponentially weighted estimate of volatility, which has the effect of keeping the volatility of the final 'profit and loss' time series (PL) approximately constant with time for a single asset, and equal between assets.

As well as the per-asset PL results, we consider the results of a portfolio of the 20 assets. The portfolio is simply a linear combination of the per-asset PL streams, with equal weights. Recall that each asset achieves approximately the same volatility due to the volatility scaling of the positions, and we achieve approximately equal sector risk because we have 5 markets per sector. We remain free to apply a final overall scaling to all the positions in order to achieve a desired volatility at the portfolio level.

Returning to the moving-average cross-over system, there are two important parameters: the short and long periods of the ewmas. These determine the effective look-back time and holding period, although this is not how the system is specified. We will consider three different trend following systems of different speeds, called 'fast', 'medium', and 'slow', which have a turnover of 1, 6, and 13 weeks respectively[1]. The fast system is correlated 26% and 5% to the medium and slow systems respectively, and these have a 34% correlation to each other. Note that our systems trade every day, regardless of speed.

[1] We define the turnover of a system as the ratio of the mean absolute position to the mean absolute 5 day change in position, roughly indicating how many weeks it takes to close a position but ignoring small trades back and forth on a daily level.

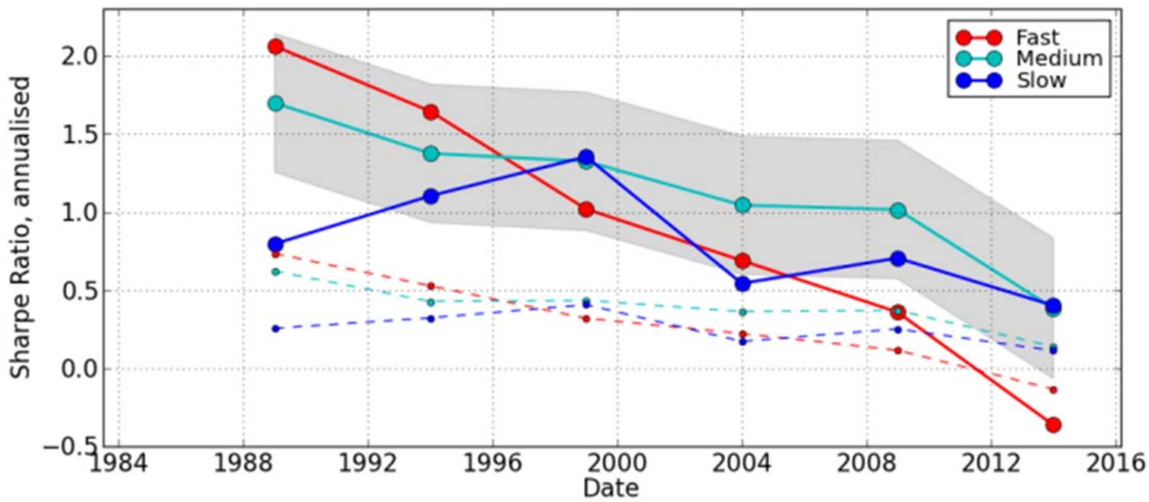

Figure 2: *The gross performance of the portfolio in six separate 5 year periods (84-88, 89-93, ..., 09-13), for our three trend following strategies. The standard error of these estimates is shown in grey around the results of the medium strategy. Dashed lines give the average SR for the individual asset PL; these appear to mimic the performance of the portfolio, indicating that the decline at the portfolio level is due to the decline of performance of underlying assets and not an increase in between-asset correlation.*

Another way of implementing a momentum strategy is to look at the past return for some look-back period and then hold a position (long or short according to the sign of the return) for some holding period, as implemented in [1]. In Section 3 we check our results against this alternative implementation, and generally find the results to be similar (at the same turnover level).

In this study we focus on gross performance, in order to simply document the historical properties of futures price series. We do not attempt to translate these results into Net PL. That is, we are not including transaction costs (broker commission and slippage), fees (performance and management), or any earned interest in our simulations.

3 Results

First we look at the results of the portfolio of 20 markets. In Figure 1, we see that all three strategies have performed well, with annualised Sharpe ratios[2] of 0.81 or higher for this 30 year period. However, we can see that the performance of the fast strategy has declined, having initially outperformed the others but remaining flat since 2004. This is even clearer once we compute the SR for different time periods, as can be seen in Figure 2, and further examined in Table 2. The fastest strategy has seen a significant decline in performance, and we can rule out with a probability <0.0001 that the SR in the first and last 10 year periods have come from a distribution with the same mean.

Speed	84-93	94-03	04-13	Δ	t	p
Fast	1.85	0.85	0.00	1.85	4.2	<0.01%
Medium	1.54	1.18	0.70	0.84	1.9	3%
Slow	0.95	0.94	0.57	0.38	0.9	20%

Table 2: *Ten year gross SR for our strategies, where Δ is the difference between the first and last 10y periods, and t and p are the corresponding t-statistic and p-value. The p-value estimates the probability of achieving the observed Δ value under the null hypothesis that the true performance did not change in these two ten year periods, and the difference is just due to sample error. All strategies have seen a decline in performance, but this is most significant for the fastest system.*

Second we look at the performance of the systems at the asset level. The decline in risk-adjusted performance at the portfolio level, as seen in Figure 2, could be due to increased correlations between the constituent asset PL, reducing the diversification and causing one to scale down one's positions in order to achieve the same annualised volatility at the portfolio level, or a decline in per-asset performance, or a mixture of both. In Figure 2 we see that the average per-asset performance has declined in an identical fashion to the portfolio performance, confirming that this is the main cause of the portfolio-level decline. To double-check this we also consider how the correlation of the per-asset PL has changed over time. We find the average correlation is ~4% (84-88) and rises to ~7% (09-13), and this is true for the 3 different trading speeds. The risk-adjusted

[2] The Sharpe ratio (SR) is a measure of risk-adjusted return, $E[\mathrm{PL} - R_f]/\sigma(\mathrm{PL} - R_f)$. It is generally measured with respect to a benchmark R_f, such as the risk free rate. However, in this study we do not use a benchmark in our SR calculations.

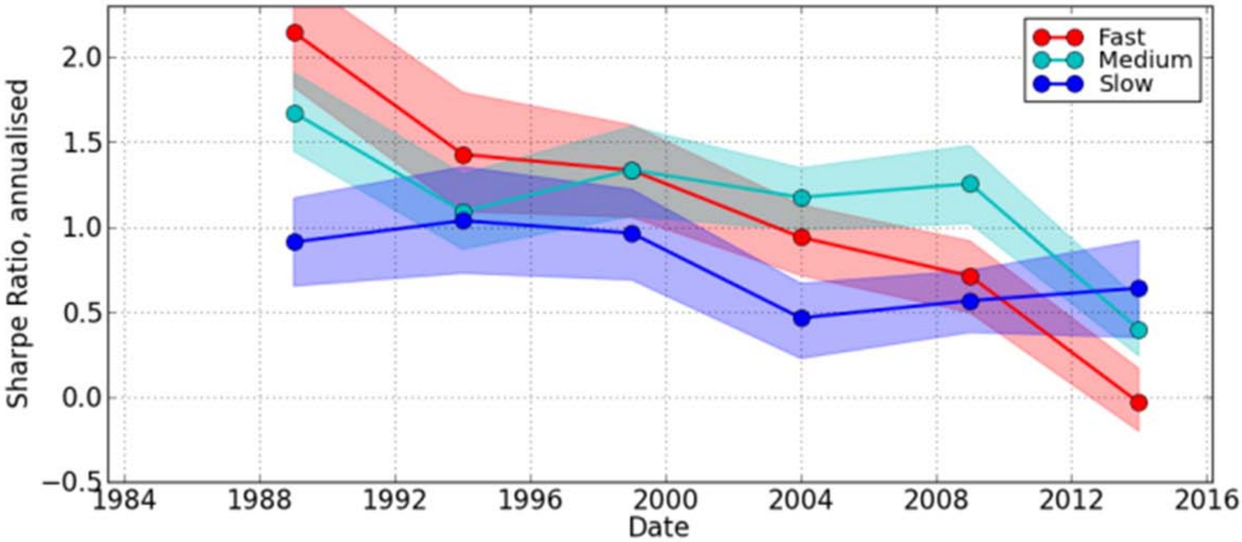

Figure 3: *Five year gross SR results for randomly selected portfolios of 20 markets (selected from a total list of 38). Plotted are the median results for each strategy speed, and the 15 - 85th percentile ranges from 1,000 realisations.*

performance of an equally weighted portfolio of N assets can be estimated as

$$R_p = \frac{\sqrt{N}}{\sqrt{(1+(N-1)\bar{\rho}}} R_a \qquad (1)$$

where R_a is the average risk-adjusted performance of the assets. With this we can estimate the effect the increase in correlation has on the portfolio level performance, and we find the multiplicative factor goes from 3.4 to 2.9. Therefore, had the per-asset results not changed over time, the increased correlations would alone cause a relative decline of about 15%. This is not much compared to the relative decline that we actually observe (100% for the fastest system).

Further analysis

We consider an alternative implementation of a trend following strategy, where instead of a price moving-average cross-over signal, we simply look at the past return for some look-back period and use this to determine the sign of our position. We then hold any position (which will also be inverse volatility scaled) for a certain holding period. However, we will have a number of positions open at any time if the holding period is greater than 1 day, and we sum over these at any time to evaluate our daily PL. We restrict ourselves to the parameter space where the holding period is half of the look-back period. We find look-back periods of 2 weeks, 1 quarter, and 1 year (and holding periods of 1 week, 33 days, and 6 months) approximately correspond to the fast, medium, and slow ewma cross-over strategies respectively; they have similar turnover levels (1, 6, and 13 weeks) and are correlated approximately 80% to the corresponding ewma strategy. The results for this implementation are as follows:

Speed	84-93	94-03	04-13
Fast	1.38	0.39	-0.01
Medium	1.19	1.07	0.59
Slow	1.26	0.91	0.60

We see the same feature - with performance declining overall, and most significantly for the fast strategy. We therefore think our results are robust to the exact implementation details of a momentum driven strategy.

We next consider the robustness of our results to the markets that we have used, as listed in Table 1. We extend our potential portfolio by another 18 markets (5 in fixed-income, commodities, indices, and 3 in currencies)[3] and we randomly choose a subset of 20 markets from these 38 and repeat the analysis for the fast, medium, and slow ewma cross-over systems. In Figure 3 we plot the median and 15-85 percentile range for the results from 1,000 random portfolio choices for each sped of system. We find the overall result robust to selection of the markets; the performance is seen to

[3] Lean Hogs, Silver, Natural Gas, Soy Beans, Aluminium, DJ Eurostoxx, Dax, Kospi, Nasdaq, the Hang Seng China Enterprises Index, Euribor, Short sterling, Bobl, 5Y T-notes, Schatz, Australian Dollar, Mexican Peso, New Zealand dollar

decline, particularly for the fastest strategy.

Although we do not want our analysis to rest heavily on synthesised data, we consider one very long time series: that of the Dow Jones Industrial Average stock index. This index was first calculated in 1896. However, futures on the index have only existed since 1997, and therefore we synthesise a rolled futures series using dividend yields and interest rates going back to 1900. We then simulate the hypothetical results of trading our original fast, medium, and slow ewma trend following strategies on this DJIA price series. We find results over this very long time period that support the hypothesis that performance has declined for the fastest strategy:

Speed	00-23	24-53	54-83	84-13
Fast	0.98	0.86	1.03	-0.10
Medium	0.55	0.51	0.43	0.08
Slow	0.01	0.43	0.15	0.43

The medium strategy has seen a less sharp decline than the fast strategy, and the slow performance is fairly flat. For reference, the standard error on a 30 year SR is approximately 0.2.

4 Discussion

Momentum is increasingly getting the stamp of approval from the academic community as one of several factors that can explain both the risk and return of equity prices [5] and more generally the returns of asset prices [1]. Its existence is controversial as illustrated by the words of Eugene Fama; “Of all the potential embarrassments to market efficiency, momentum is the primary one” [7].

At the same time as receiving academic attention, mainstream institutional investors have embraced momentum strategies. This is illustrated by the growth of assets in the CTA industry and the proliferation of products that seek to profit specifically from a momentum or trend following effect.

In this paper we have gone beyond just documenting the existence of a trend following effect in futures markets, and have explored how consistent it has been through time. In particular, we find evidence that the strength of fast trend following strategies has declined significantly over time. Our results are robust to the markets we use and the parameterisation of the system.

A glaring omission from this study is the issue of transaction costs which could be a major factor in the ultimate profitability of these strategies. Instead we have concentrated on the raw effect size, and have not attempted to make any claims as to the potential economic usefulness of it. In order to understand the economic significance it is necessary to consider brokerage fees, market impact (or slippage) and any other associated costs. Furthermore in the context of this study it would be necessary to understand how these costs have varied through time. Since this would require many more assumptions we decided that it would detract from the simplicity and clarity of this study, although we appreciate that costs could be substantial.

We have also not included interest in these simulations; investors of CTA Funds typically earn the risk-free rate on the cash that they do not require to margin positions. This can be around 90% of assets and therefore this would improve the performance of the total investment significantly. This is true especially in the past when the risk-free rate was higher than today. We have also not included management or performance fees in the simulations as they are not relevant for this study (having had a consistent effect over time).

We have not yet speculated about the causes of the decline in performance, but many would note that it coincides with a rise in assets under management in the CTA industry. Three possible explanations for a weakening of the momentum effect are:

1. *A growing volume of assets seeking to profit from the effect*. The exact mechanics of how this might weaken performance are beyond the scope of this discussion, but such an argument would be based on the premise that the more people who know about an ‘opportunity’ in financial markets, the less likely it is to persist. This explanation is consistent with the growth of the CTA industry.

2. *A reduction in transaction costs*. An efficient market perspective would argue that predictable effects can exist, provided they are not exploitable after costs. Viewed from this perspective a steady reduction in transaction costs over the last 40 years could explain the reduction in effect size. Whilst the experience of the authors as practitioners within the CTA industry is at odds with an efficient market perspective, nothing in

this paper discredits this being either a partial or complete explanation.

3. *A change in the behaviour of market participants*. The first two possible explanations are endogenous to those who seek to exploit this effect. It is possible that some exogenous change is responsible. Indeed the recent poor performance of trend following strategies has been attributed to the intervention of central banks curtailing trends in markets [6].

These three explanations are purely speculative and as such we do not seek within this paper to suggest which, if any, of them is the most plausible. It is also not possible to extrapolate our results into the future with any confidence. All we can say is that the strength of the momentum effect has not been historically stable, and so we have no reason to think it will be in future. Taking all the data at face value the outlook is unclear, except that our results suggest that fast is not always best, and perhaps patience is rewarded.

Acknowledgements

We are grateful to the Winton research department for help with the simulations and data, as well as many useful discussions on this topic over a number of years. With particular thanks to Matthew Beddall, and Tom Babbedge, for their valuable insights.

For correspondence please email
researchpapers@wintoncapital.com

References

[1] T.J.Moskowitz, Y.H.Ooi, L.H.Pedersen, *Time series momentum*, Journal of Financial Economics, vol 104, 228-250 (2012)

[2] N.Jegadeesh, S.Titman, *Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency*, Journal of Finance, vol 48 (1993)

[3] A.W.Lo, A.C.MacKinlay, *When Are Contrarian Profits Due to Stock Market Overreaction?*, Review of Financial Studies, vol 3, 175-205 (1989)

[4] M.M.Carhart, *On Persistence in Mutual Fund Performance*, Journal of Finance, vol 52, 57-82 (1997)

[5] E.F.Fama, K.R.French, *Size, Value, and Momentum in International Stock Returns*, Journal of Financial Economics, vol 105, 457-472 (2012).

[6] E.Cusworth, *CTA trend followers suffer in a market dominated by intervention*, Risk.net (2013)

[7] E.F.Fama, R.Litterman, *An Experienced View on Markets and Investing*, Financial Analysts Journal, vol 68 (2012)

LEGAL DISCLAIMER