Trend's Not Dead

(It's just moved to a trendier neighbourhood)

Dr Thomas Babbedge[1], J Scott Kerson[2]

[1] Chief Scientist & Deputy Head of Systematic Strategies
[2] Senior Managing Director & Head of Systematic Strategies

E-mail: investorrelations@greshamllc.com

Published May 2019. http://ssrn.com/abstract=3404669

Abstract

We explore the reduced performance of trend followers over the past decade but fail to find evidence that this is due to the commonly proffered reason of over-crowding of the strategy. Instead we find that the cause can be laid at the feet of the markets themselves – those markets commonly traded by trend followers have simply not trended as strongly in the past decade. By using a novel dataset of alternative commodity markets we show that the 'trendiness' of less mainstream markets, selected based on a set of simple criteria, is inherently higher and that trend following in these markets has continued to be significantly better.

Keywords: trend-following, momentum, crowding, alternative markets, CTA

1. Trend Followers

As is well known, classical trend following in liquid markets has struggled over most of the 10 years since the global financial crisis (GFC), and stands in sharp relief to the performance of similar systems prior and during the crisis. This is demonstrated by the performance of representative indices such as Société Générale's SG Trend Index and BarclayHedge's Barclay CTA Index.

Taking March 2009 as the start of the post-GFC period[1] we find that the Sharpe Ratio (SR) of the Barclay CTA Index has been essentially zero (0.1 +/- 0.3 s.e.) compared to a SR of 0.8 (+/-0.2 s.e.) before then. We can ask how significant this difference in SRs is via Opdyke 2007's [1] work on the asymptotic distribution of measured SRs. Whilst the probability that the pre-GFC SR is positive is 99%, it is only 60% for the post-GFC SR and the probability that the post-GFC observed SR is less than the pre-GFC period is 95%.

2. Why has the performance declined?

2.1 Is it over-crowding?

A common hypothesis is that the amount of capital deployed in trend following strategies has reached the scale where competitive saturation is now a significant concern. Competitive saturation refers to the degradation in performance caused by increased competition for the same source of alpha – i.e., the compression in returns caused by more people applying the same investment approach to the same markets. Indeed, from Figure 1. we can see that recent reduced CTA performance has been coincident with AUM in Managed Futures strategies being at historic highs.

[1] Exact date choice has minimal impact on conclusions

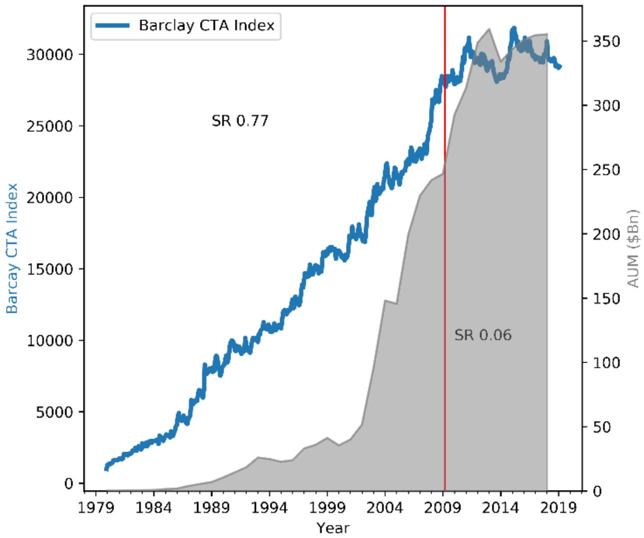

Figure 1 – Barclay CTA Index on left hand axis (blue) and Managed Futures AUM (shaded region using secondary axis). Source: BarclayHedge

Whilst the size and number of futures markets has also increased over time, this has been outstripped by the growth of CTAs, with the ratio of managed futures AUM to total $ ADV in futures markets doubling from pre- to post-GFC periods (0.16 to 0.27).

But correlation does not neccesarily mean causation. We here attempt to measure any impact on CTA performance arising from a general crowding of the strategy.

Direct observation is of course impossible, because one cannot evaluate market behavior on a counterfactual basis. We can however simulate the counterfactual; *what would have happened if one had traded behind everybody else?*

Implementation lag refers to the negative impact on performance of the inevitable delay between sample time (when the model 'sees' the price) and execution time (when the model 'fills' its desired holdings).

We use alpha decay—the deterioration of performance and Sharpe Ratio—as a proxy for both the impact of saturation and of implementation lag as both should manifest themselves in terms of the magnitude and speed of alpha decay. Thus if increasing competitive saturation is costly, we should observe a change in performance over time (i.e., an acceleration in alpha decay). Similarly, if delayed implementation is costly, we should see performance drop as a function of trade lag (i.e., a step function change in the level).

2.1.1 Quantifying saturation via alpha decay. The crux of this analysis is that if the recent growth in assets and players is cannibalizing alpha, then we should see an increasingly negative cost to 'trading late', because all those assets and players will have created a 'footprint' in the market, and the late entrant will buy after the competition has bought, or sold after they've sold. Given that we know that the number and size of assets and players has increased over the past few years, we would expect to observe an increasingly severe cost of delayed execution over the same period, if those assets and players have saturated liquid futures markets.

We backtest a trend following simulation on a set of over one hundred liquid futures markets from 2000-2019 (across bonds, rates, currencies, equities and commodities), comparing the resulting performance when we either assume the theoretical – but unachievable – case of simultaneous sampling and execution (Lag 0) to the case where we trade a *full* 24h later hours (Lag 1). The Lag 0 SR before fees is 0.75, dropping to 0.7 for Lag 1. At 10% annualized volatility, 0.05 Sharpe points equates to 50bps annualized loss in performance, or about 8% of net alpha (for a Lag 0 after fees SR of 0.66).

Clearly, the delay is not costless, but we should note two things: firstly, Lag 0 is an unachievable best case (one can never trade and sample at the same price simultaneously), and Lag 1 is a worst case (since one would normally sample and then trade some short time afterwards). Thus, one would expect the actual SR and cost to land somewhere in between the Lag 0 and Lag 1 scenarios. To address the possibility of crowding leading to increased alpha degradation we need to know if this cost has been accelerating. This would manifest itself as an increasing performance differential over time. The top panel of Figure 2 shows the cumulative differential between the Lag 0 and Lag 1 account curves. These differentials have been stable over time, and there is no obvious acceleration over the recent past. The gradients in the two periods are entirely consistent with being the same – that is, the rate of alpha decay with lag being the same in both periods. Thus, we see no footprint of increased trend follower AUM leading to competitive saturation and over-crowding.

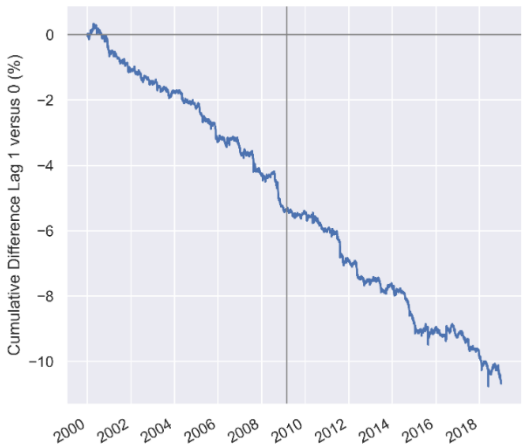

Figure 2 – cumulative lag 1 under-performance versus lag 0 backtest, showing the consistent and persistent gradient. Source: Gresham Investment Management (GIM), Bloomberg

2.2 Why *haven't assets swallowed alpha?*

One explanation is 'stock versus flow'. The natural concern is that any individual CTA's will overestimate available liquidity inasmuch as it fails to fully consider the combined assets of similar participants, who will also presumably be making their own assessment of available liquidity. However, this phrasing of the issue ignores a key differentiation between positions and trades – what we call the *stock* (the collective position across the space) and the *flow* (the incremental changes in that position by participant, for which the question of liquidity is highly relevant). Indeed, even for two hypothetical CTA's with identical market allocations, they may have substantial differences in their respective parameterisations (eg, speed) of their strategies.

Here we examine whether CTA's with similar trend following strategies, and hence similar 'stocks' of positions (e.g., generally being long or short a given market at the same time), will also exhibit correlated 'flows', or changes in those positions (trades).

2.2.1 *Toy model.* Two similar trend following strategies are run on a single market of arbitrary choice (WTI Crude oil). Here trend following has been defined as being an exponentially weighted moving average crossover (EWMAC). The two strategies have similar effective speeds in terms of information window, where the effective speed is defined as the number of days into the past that contain 50% of the EWMAC weight. For CTA A a single medium speed EWMAC has been used. For CTA B a mix of both a fast and slow EWMAC has been used. Both CTAs have an effective speed of around 45-50 days. In Figure 3 we first compare the trend signal from both CTAs (stock), and then their changes in signal (flow).

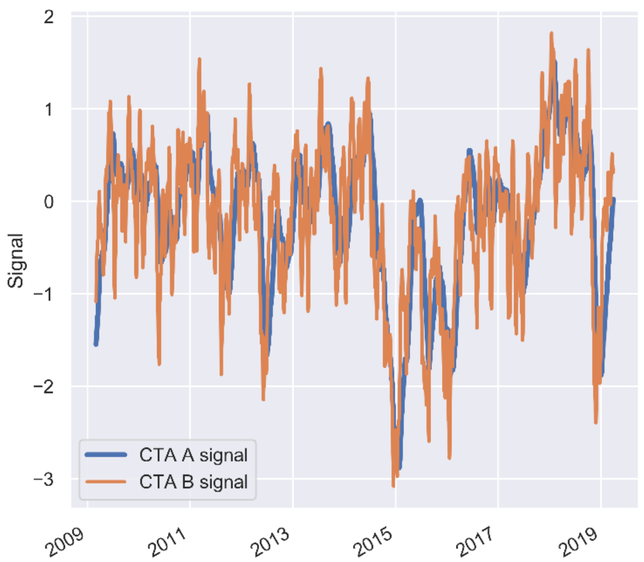

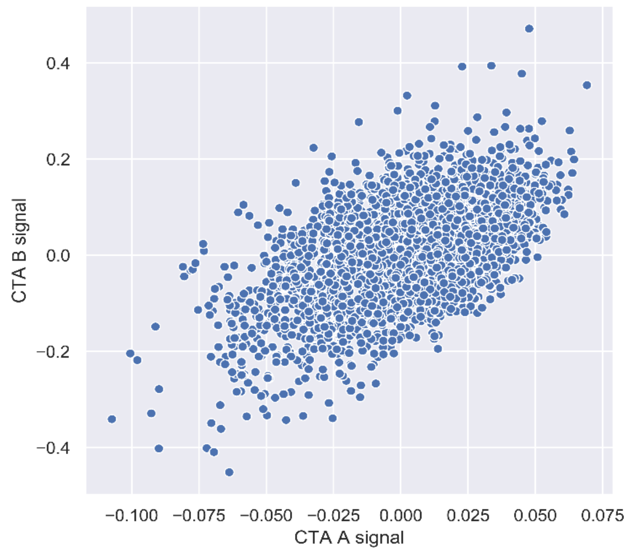

Figure 3 - Comparison of signal (first panel) and delta signal (second panel) for CTA A and B on WTI Crude. They are 0.78 and 0.56 correlated, respectively. Source: GIM, Bloomberg

To generalise the result, we extend this approach to over 100 liquid futures markets, finding that the mean signal correlation across these markets is 0.77 over the past decade, whilst for Δsignal the mean correlation is 0.58 – in other words, the stock (as represented by signal) between the two are about 80% correlated, but the flow of trades (as represented by Δsignal) between them are less than 60% correlated. Next, because signals are all normalised into the same units, we can aggregate all the data into a single relationship. This is displayed as a density plot in Figure 4 due to the large number of data points (260,000). For this super-sample, signal correlation is 0.79 and Δsignal correlation is 0.58 – very similar to the individual market analysis.

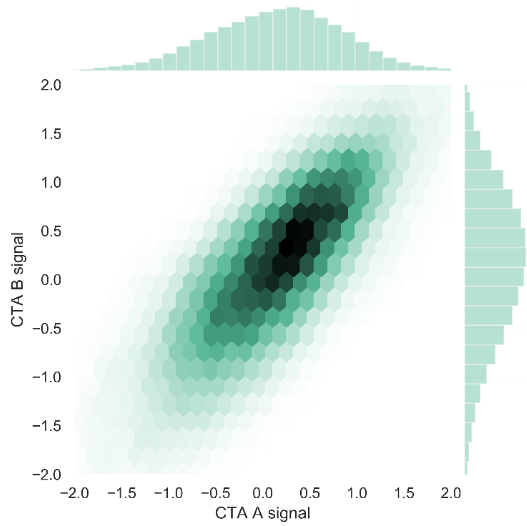

Figure 4 – Signal density for CTA A and B across liquid futures. Source: GIM, Bloomberg

Note that the Δsignal correlation is likely to represent an *upper* limit for the degree of overlapping trading behavior, because the only difference introduced was in terms of trend horizons and even then, they were 'effective speed' matched – we will relax and test this hypothesis next.

2.2.2 *A step closer to realism*. In the real world, different CTAs – even in the narrowly defined trend bucket - employ a wide range of different techniques to achieve their ends: there are different definitions of 'trend' (EWMA oscillator, break-out, etc), different 'splines' or response functions mapping raw signal to model conviction, different risk models for inverse-vol scaling, different portfolio risk controls, different smoothing, buffering and trade/position limits… the list is as potentially as long as there are lines of code in the strategy codebase.

We attempt to construct a more realistic comparison between two (somewhat arbitrary) trend-following CTAs. For CTA A we adopt a plain-vanilla 1 month realized volatility for inverse position sizing, for which we then simulate positions and trades. For CTA B an approach more similar to our own strategies has been adopted, including our proprietary robust volatility model, signal and position buffering, and a signal spline incorporating endogenous awareness of forecast uncertainty and trend exhaustion.

We can't meaningfully aggregate positions across all futures markets (as notional positions are not normalised) but we can find the correlation for each market in turn, and the average correlation of each pairwise position was 0.74, and the average trade correlation was 0.30 – again, not high, and substantially lower for the 'flow' than for the 'stock'. So, despite having very similar positions, two CTAs' trades can in fact be quite uncorrelated.

2.2.3 *Stress scenario liquidation*. The 'flow' property that we've established is all well and good, as it's certainly helpful to know the extent to which similarly spirited CTA's can exhibit markedly different trading behaviour in normal market conditions. However, the same analysis demonstrates that the 'stock' property of these CTA's is likely to be quite similar, which raises a question about liquidation risk, rather than normal trading patterns – i.e., suppose that two CTA's have a large and overlapping exposure to a given market, and they both want to reduce that exposure relatively quickly and simultaneously (this could be reaction to correlated redemption requests, a spike in volatility or other exogenous market event). To examine this scenario, we consider periods when both CTA A *and* B held large positions in a market (defined as a position ≥ 90th percentile of the distribution of all absolute positions in the simulation). We then select trades which were reducing for *either* CTA and then pool across all the markets considered, which yields an average conditional trade correlation of 0.00 +/- 0.03 at 95% C.I. – in other words, even lower correlation across trades than in the 'normal' case. The distribution of the individual market correlation measures is shown in Figure 5.

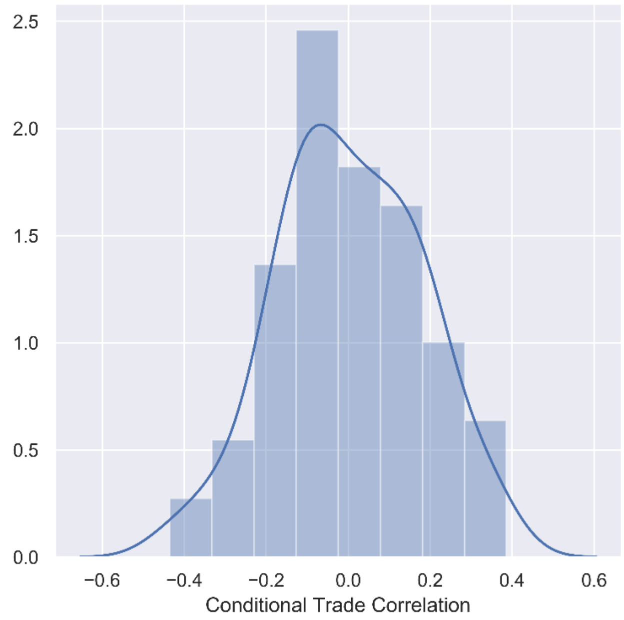

Figure 5 – Conditional trade correlations for each market, with a mean correlation of 0. Source: GIM, Bloomberg

2.3 Maybe it's the signal?

When we looked for evidence of over-crowding we failed to find its footprint in the lag-trading analysis. Furthermore, the notion that all trend followers' trading activity is similar was found to be less likely than is commonly believed. So, if we cannot convincingly blame over-crowding for poor trend performance post-GFC, perhaps we can instead blame the machinary of trend -following itself. Maybe EWMACs and their ilk no longer efficiently capture trends in markets?

Using the same trend following definition as used in §2.1, in Figure 6 we plot risk-adjusted quarterly returns[2] of futures markets[3] versus the resulting simulated quarterly return from trend following[4] on those individual markets, splitting the data into pre- and post-GFC. For both periods we overlay a loess line of best fit. The resulting convex 'CTA smile' is a well-known result and demonstrates how trend following is akin to a synthetic long straddle (e.g. Merton 1981 [2], Fung & Hsieh 1997 [3], Dao at el. 2016 [4]). It is perhaps remarkable that the pre- and post-GFC relationship is virtually identical. Crucially, therefore, the mechanism by which trend following translates market moves into trend returns has not altered.

[2] Chosen to be similar in timeframe to the horizon of medium-speed trend followers

[3] Risk-adjusted to an annualised risk of 10%

[4] Again, targeting 10% annualised risk

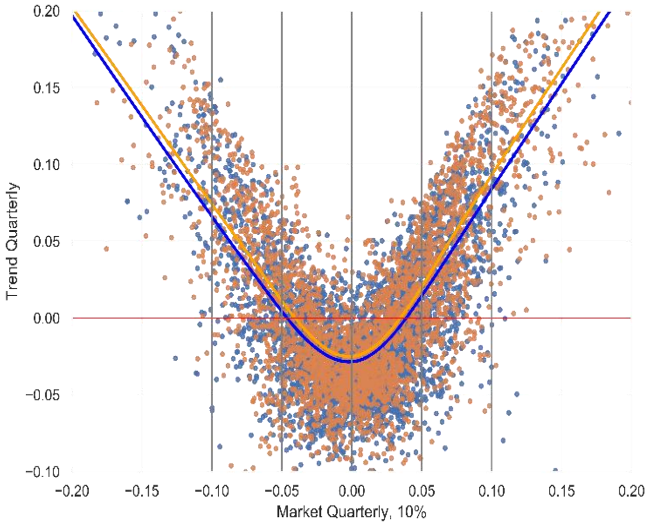

Fig 6. Quaterly return CTA smile for liquid futures in two periods (orange = pre-GC, blue = post-GFC). Loess fits indicated. Market quaterly returns are risk-adjusted to 10% annualised risk. Source: GIM, Bloomberg

2.3.1 So what changed? If we look at the density of data in different regions of the observed CTA smile we find that there is a difference between the two periods. Table 1 sets out the proportion of quarterly market returns that were 'small' (absolute returns < 5%) and 'large' (absolute returns > 10%). There has been a marked shift of occurrence away from *large* trends and into *small* trends. This is illustrated by Figure 7.

Table 1. Occurrence counts for small and large risk-adjusted market quarterly returns

	Small Trend (|Mkt Retn| < 5%)	Large Trend (|Mkt Retn| > 10%)
pre-GFC	59% of quarters	10% of quarters
post-GFC	68% of quarters	5% of quarters

Given that trend following, viewed as a straddle, can be characterised as bearing an options cost when markets are not trending (the central region) and a pay-off when markets are trending (the tails) this observation explains the weak performance of trend following in the post-GFC period – markets spent more of their time in small weak trends and the occurrence of larger trends was almost halved. It is beyond the scope of this paper to proffer a reason as to why markets have exhibited less trend in the past decade but the fact the cause lies with the markets rather than with trend following itself suggests that those same markets could exhibit larger trends again in the future, with a commensurate improvement in trend following performance. However, as we do not have a crystal ball we will instead look elsewhere for markets that have continued to exhibit larger trends.

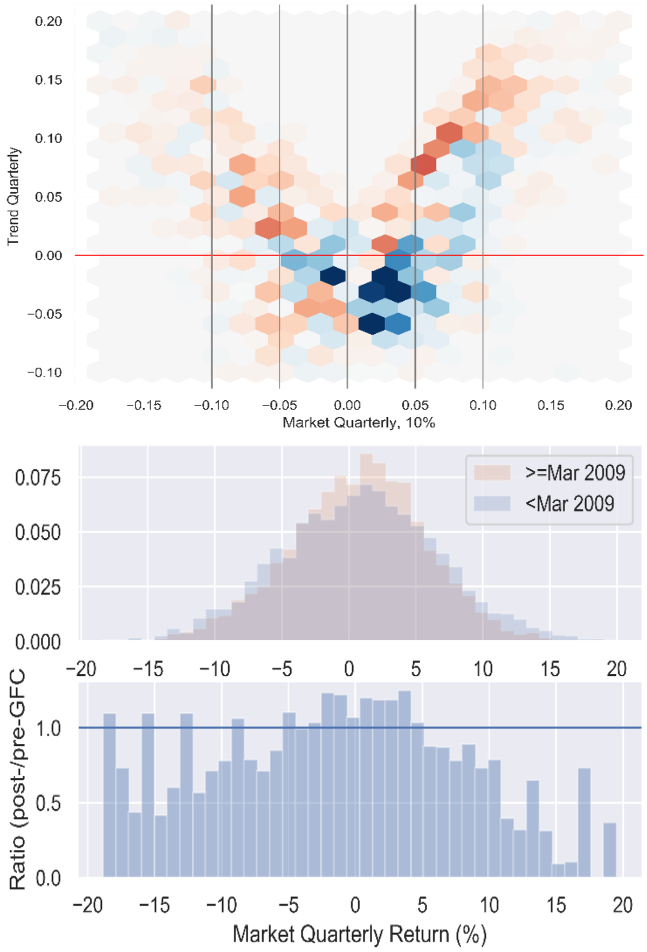

Figure 7 –* Top panel *shows increased occurences in blue and decreased in red when comparing post-GFC to pre-GFC period.* Middle panel *compares the distribution of risk-adjusted market quarterly returns in the two periods.* Bottom panel *displays the ratio of post-GFC histogram to pre-GFC. Source: GIM, Bloomberg

3. Alternative markets

Our hypothesis is that markets that exhibit certain characteristics should be inherently more 'trendy'. Namely:

- Are dominated by hedgers, not speculators – less competition, natural alpha transfer
- Are structurally insulated from risk on/off and typical macro factors – no policy driven capping/flooring of trends
- Exhibit fixed or inelastic supply/demand – forces prices to do all the work to clear markets
- Lack fungibility and temporal arbitrage – maintain diversification, inherit lots of carry

3.1 Alternative commodity markets: one such neighbourhood?

We believe that *Alternative Commodity* markets demonstrate these characterisitics, identifying 95 markets (that we currently trade). As an example, one can trade freight futures based on the Panamax[5] Timecharter Index. The availability of these ships is a classic case of inelastic supply and demand since it takes between 1 and 3 years to construct a new ship and that ship can then be in service for 25 to 30 years.

We choose to represent inherent trendiness via the cumulative auto-correlation term from Lo 2002 [5] since this provides a simple and intuitive measure of the extent to which a returns time series is auto-correlated over extended periods. See the second square-root term in Equation 1.

$$\sigma_{annual} = \sqrt{n}\sigma_{daily}\sqrt{1 + 2\sum_{i=1}^{n}\frac{n-i}{n}\rho_i} \qquad \text{(Eq 1)}$$

We measure this for both liquid futures markets pre-/post-GFC (Figure 8) and also compare to alternative commodities post-GFC (Figure 9), considering auto-correlation lags out to 1 year. Two observations can be made:

i) Just as with the smile trend densities in §2.2.1 we see a decline in auto-correlation 'trendiness' for liquid futures for the recent period

ii) We see that alternative commodity markets tend to have a larger auto-correlation trendiness term, as per the hypothesis

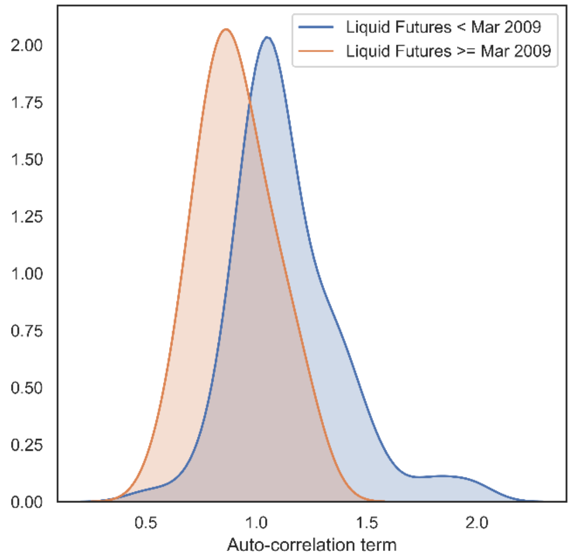

Figure 8 - Trendiness for liquid futures pre- and post-GFC, showing the reduction in the measure post-GFC. Source: GIM, Bloomberg

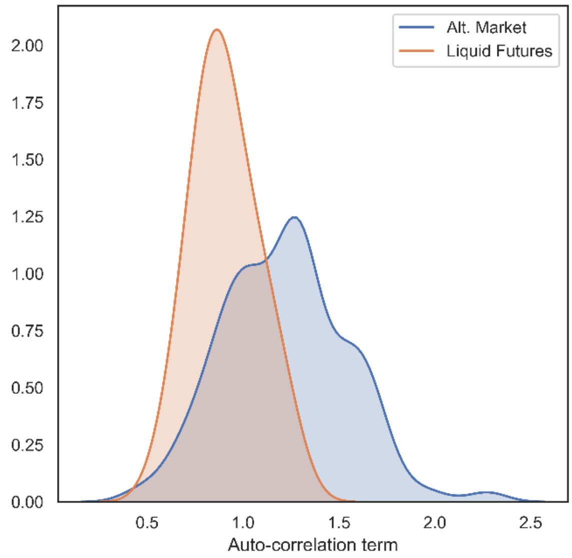

Figure 9 - Trendiness of liquid futures cf. alternative commodities for the post-GFC period, showing the higher level in alternative commodities. Source: GIM, Bloomberg

3.1.1 Trend following in alternative commodities. We run the same trend following backtest on this set of 95 alternative commodity markets[6] and construct the same CTA smile as before. As before, the loess fit is essentially identical to that seen for liquid futures markets in §2.2. Crucially, though, we now see an *increased* density of *large* quarterly market risk-adjusted returns and a *decreased* occurrence of *small* moves. In Figure 10. we present the differential density chart comparing Alternative Commodities to liquid futures post-GFC and provide fractions in Table 2.

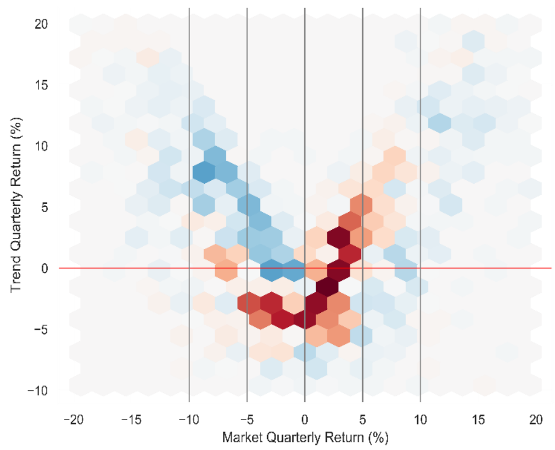

Figure 10 - shows increased occurences in blue and decreased in red when comparing alternative commodity markets to liquid futures markets. We see higher rates of large quarterly market returns and lower rates of small market moves for the alternatives. Source: GIM, Bloomberg

[5] The largest size of ship able to navigate the Panama Canal

[6] Being careful to apply realistic trading cost estimates based on our proprietary dataset of actual trading costs

Table 1. Occurrence counts for small and large market quarterly returns

	Small Trend (|Mkt Retn| < 5%)	Large Trend (|Mkt Retn| > 10%)
Liquid Futures pre-GFC	59% of quarters	10% of quarters
Liquid Futures post-GFC	68% of quarters	5% of quarters
Alt Comms post-GFC	*56% of quarters*	*14% of quarters*

3.1.2 Comparison to the mainstream. Finally, we construct a portfolio of alternative commodities and compare the simulated[7] cumulative performance after all fees and costs to that of the Barclay CTA Index in Figure 11. In Table 2. we provide correlations to major representative macro factors.

As per the original hypothesis we observe that simulated historical performance of the alternative commodities trend following has been far better than similar strategies applied to liquid futures markets in the post-GFC period, whilst exhibiting low correlation to more mainstream factors.

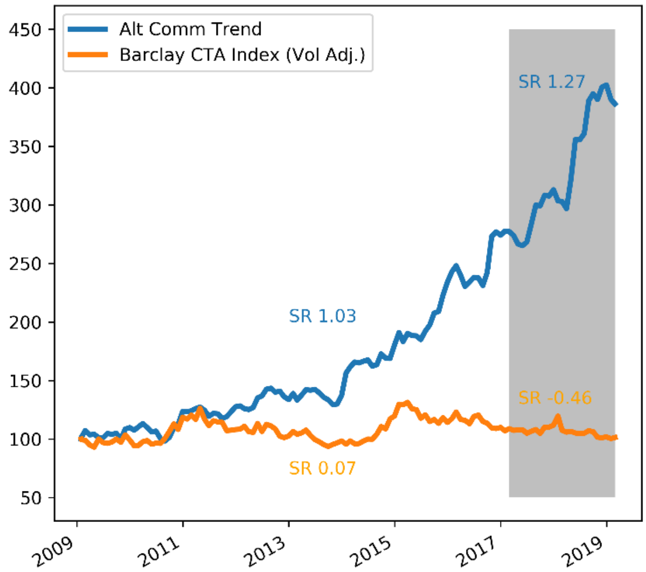

Figure 11 - Performance comparison for the Barclay CTA Index and the Alternative Commodity trend strategy. Shaded region indicates performance from live trading of the strategy. Source: GIM, Bloomberg

Table 2 - Correlations (monthly) between the Alternative Commodity Trend strategy and other macro factors

	Alt Comm Trend	Barclay CTA	S&P 500	BCOM
Barclay CTA	0.15			
S&P 500	0.06	0.43		
BCOM	0.17	0.26	0.67	
Barclays Ag. Bond	0.14	0.19	-0.07	-0.19

4. Concluding remarks

We were unable to find evidence that the poor performance of mainstream trend followers over the past decade (post-GFC) was due to over-crowding and found that even similar trend following approaches can result in lowly-correlated trading activity. Indeed, the 'mechanical' transformation of market moves into resulting trend following returns was shown to be the same pre-/post-GFC, implying that the act of trend following itself was not 'broken'. Rather, it appears that the cause lies with the behaviour of the markets themselves, with a marked reduction in the occurrence of large (quarterly) moves in markets. Therein lies some hope for mainstream trend followers since the cause appears to be exogenous and one might expect that the behaviour of markets could change again in the future.

Not content with waiting for this potential but uncertain future improvement, we instead looked to identify markets that should, in principle, exhibit stronger trending behaviours. We found that a novel dataset of *alternative commodity markets*, selected based on a set of simple criteria, had inherently higher trendiness and that, as a result, trend following in these alternative markets has continued to be significantly better than for the mainstream. Thus, it seems, Trend is not dead – it has just moved to a more trendy neighbourhood.

References

[1] Opdyke, JD. "Comparing Sharpe Ratios: So Where Are the p-Values?" *Journal of Asset Management*, vol. 8, no. 5, 2007, pp. 308–336.

[2] Merton RC. "On market timing and investment performance i: An equilibrium theory of value for market forecasts" *Journal of Business*, 54:363–407, 1981.

[3] Fung, W., Hseih D. "Emprical Characteristics of Dynamical Trading Systems: The Case of Hedge Funds" *The Review of Financial Studies*, 2, 275-302. 1997.

[4] Dao, TL. et al. "Tail protection for long investors: Trend convexity at work" http://ssrn.com/abstract=2777657, 2016.

[5] Lo, AW. "The Statistics of Sharpe Ratios" *Financial Analysts Journal*: Vol 58, No 4, 2002, pp 36-52.

[7] From March 2017 the returns are from the live track record of our alternative commodities strategy

Endnotes

Glossary

This material is not intended to be a recommendation or investment advice, does not constitute a solicitation to buy, sell or hold a security or an investment strategy, and is not provided in a fiduciary capacity. The information provided does not take into account the specific objectives or circumstances of any particular investor, or suggest any specific course of action. Investment decisions should be made based on an investor's objectives and circumstances and in consultation with his or her advisors. The views and opinions expressed are for informational and educational purposes only as of the date of production/writing and may change without notice at any time based on numerous factors, such as market or other conditions, legal and regulatory developments, additional risks and uncertainties and may not come to pass. This material may contain "forward-looking" information that is not purely historical in nature. Such information may include, among other things, projections, forecasts, estimates of market returns, and proposed or expected portfolio composition. Any changes to assumptions that may have been made in preparing this material could have a material impact on the information presented herein by way of example. Past performance is no guarantee of future results. Investing involves risk; principal loss is possible. All information has been obtained from sources believed to be reliable, but its accuracy is not guaranteed. There is no representation or warranty as to the current accuracy, reliability or completeness of, nor liability for, decisions based on such information and it should not be relied on as such.

All investments carry a certain degree of risk and there is no assurance that an investment will provide positive performance over any period of time. Commodity Trading Involves Substantial Risk of Loss. It is not possible to invest directly in an index.

Gresham Investment Management LLC is a U. S. registered investment adviser and an affilliate of Nuveen, LLC.
Nuveen provides investment advisory solutions through its investment specialists.

This information does not constitute investment research as defined under MiFID. In Europe this document is issued by the offices and branches of Nuveen Real Estate Management Limited (reg. no. 2137726) or Nuveen UK Limited (reg. no. 08921833); (incorporated and registered in England and Wales with registered office at 201 Bishopsgate, London EC2M 3BN), both of which entities are authorized and regulated by the Financial Conduct Authority to provide investment products and services. Please note that branches of Nuveen Real Estate Management Limited or Nuveen UK Limited are subject to limited regulatory supervision by the responsible financial regulator in the country of the branch.

870337-G-R-E-07/20