KENSINGTON Asset Management

Trend Following & Tail Risk

April 2020

"CUT SHORT YOUR LOSSES. LET YOUR PROFITS RUN ON." *– David Ricardo, late 18th century*

Key Topics

- Defining Trend Following
- Why Should Trend Following Work?
- Managing Tail Risk with Skew
- Positive Convexity Leads to Skew
- Behavior in Volatile Environments
- Timeframes and Trend Following

Across the globe and throughout history, merchants and traders have relied on trend following strategies. Some of the earliest forms of trend following can be traced back to ancient Athenian merchants, and it has been practiced for centuries even before the advent of the modern stock market (Hasansodzic & Lo, 2010). The earliest trend followers combined basic price information with superstition. As time progressed, methods became more advanced with ledgers and charts. Fortunately, for the purposes of this article, we have at our disposal modern tools and methods to provide the reader with an introductory understanding of how trend following can manage tail risk.

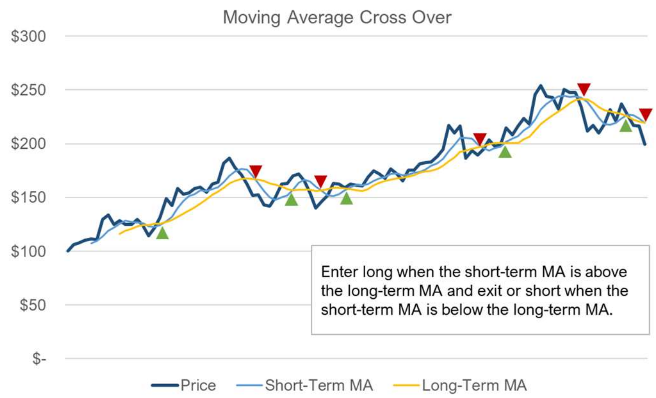

Figure 1. A visual example of a simple Moving Average Crossover rule.

What is Trend Following?

There are generally two ways to implement trend following:

1. Observing an asset's price performance in isolation to make allocation decisions
2. Observing relative price performance between assets

In this article we will focus on the former type of trend following, as relative momentum has been extensively documented (though there may be a deeper underlying relationship between the two). There is no one set way to systematically trend follow prices, but the basic building blocks of trend following strategies consist of trading rules using price/return metrics that have been constructed on short, medium, and/or long-term look-back periods. An example of a well-known, simple priced based rule is the Moving-Average Crossover, which utilizes two moving averages; one constructed on a shorter time frame than the other. Where the short-term moving average is in relation to the longer-term moving average determines the entry into a long or short/neutral position. Figure 1 provides a visual representation of the strategy. An example of a return-based rule is using the sign of past rolling returns to determine the position[1] (Moskowitz, Ooi, & Pedersen, 2012). Other systematic trend following rules may utilize volatility measures or price ranges to determine "breakouts," where the sustainability and direction of the trend is determined through the crossing of a threshold.

1. The Time Series Momentum factor as observed by Moskowitz et al. utilizes the past 12-month excess return as a predictor of future returns.

info@kensingtonassetmanagement.com | Barton Oaks Plaza Bldg II | 901 S Mopac Expy | Ste 225 | Austin, TX 78746 | 877.891.1222

There is no limit to the number of trading rules in existence, and generally one might find that many are mathematically tied to each other. In fact, with the two examples previously mentioned, the simple returns-based rule is in some sense a specific case of the Moving-Average Crossover[2] (Bruder, Dao, Richard, & Roncalli 2011). Despite the differences in trend following methodologies, they all strive to do the same thing: harvest persistence in price movements.

Why Should Trend Following Work?

It is difficult to discuss the benefits of trend following without at least lightly treading into a conversation of market efficiency. Proponents of any form of the Efficient Market Hypothesis will stress that patterns in market prices cannot consistently be exploited to outperform a naïve buy-and-hold strategy without assuming greater risk. However, with the growing body of literature showcasing the success of trend following in different settings, trend following has become a commonly accepted "anomaly." When it comes to anomalies related to trend and momentum, there are two schools of thought. On one hand, efficient market theorists will stress risk-based explanations for its success, often citing risk from economic factors or firm characteristics. On the other hand, behavioral finance makes a strong case for investor irrationality, and alongside structural limits to arbitrage, this can lead to trends arising in prices. Then there are those who posit a compromise, believing that markets are "micro-efficient" but "macro-inefficient," whereby securities within an asset class may be priced rationally, but entire asset classes themselves may be priced above or below fundamental value[3]. (Jung and Shiller 2005).

While this debate rages on and is important, there is another approach to which we can evaluate trend following outside of academic theories of irrationality and efficiency. We instead evaluate how trend following can mechanically manage tail risk due to positive skew arising from an option like payoff structure, which is not an anomalous feature.

Managing Tail Risk with Skew

Many studies cover multi-asset trend following strategies with volatility scaling of positions, a common methodology of many CTAs. It goes without saying that diversification is beneficial, and furthermore, while volatility scaling can contribute to risk management, it can complicate attribution of performance (Huang, Li, Wang, & Zhou, 2020). For a basic, introductory understanding of systematic trend following, we will instead examine the behavior of individual naïve trend following strategies over fast to slow time frames applied to a single asset class. In our back test we will be covering S&P 500 Index price returns over an 85-year period from 1935 through the end of 2019. Moving Average Crossovers are used as the trend signals and allocations are made in a long-neutral manner. Table 1 and Figure 2 record the performance statistics. A surface level glance at CAGR and Sharpe ratio might encourage one to simply dismiss trend following; three out of the four trend following strategies have lower Sharpe ratios compared to a simple "buy and hold" approach. But the Sharpe ratio represents an incomplete picture of what investors experience in their daily lives; volatility is only average deviation from mean returns and reveals nothing about symmetry of occurrences.

2. Bruder et al. show that time series momentum has a centered trend signal, whereas more generally moving average crossovers can either back or forward weight trend signals during an observation period.
3. Known as "Samuelson's Dictum", this statement comes from a quote of Paul Samuelson to John Cambell and Robert Shiller. Paul Samuelson independently developed similar ideas of market efficiency during the same period as Eugene Fama in the 1960s.

Strategy	5 Day MA X 20 Day MA	10 Day MA X 50 Day MA	20 Day MA X 100 Day MA	50 Day MA X 200 Day MA	S&P 500
CAGR	3.27%	4.46%	4.74%	6.80%	7.10%
Maximum Drawdown	-41.31%	-52.11%	-39.39%	-37.01%	-59.99%
Sharpe Ratio	0.03	0.14	0.17	0.35	0.30
Mean of Monthly Returns	0.00	0.00	0.00	0.01	0.01
Monthly Volatility	0.03	0.03	0.03	0.03	0.04
Average Up Month	2.73%	2.69%	2.87%	2.97%	3.38%
Average Down Month	-1.90%	-1.52%	-1.49%	-1.35%	-3.43%
Mean of Yearly Returns	0.03	0.04	0.05	0.07	0.07
Yearly Volatility	0.11	0.12	0.12	0.12	0.17
Skew of Monthly Returns	0.33	0.16	-0.43	-0.45	-0.52
Kurtosis of Monthly Returns	2.93	1.32	6.42	4.37	3.50
Yearly Skew	0.06	0.60	0.60	0.43	-0.46
Yearly Kurtosos	-0.07	1.20	0.11	-0.13	0.04
Average Up Year	10.46%	12.42%	13.95%	13.51%	17.48%
Average Down Year	-8.09%	-6.21%	-6.53%	-6.73%	-12.71%

Table 1. Performance statistics from 1935 through the end of 2019 using daily price returns. Calculations performed by Kensington Asset Management, LLC

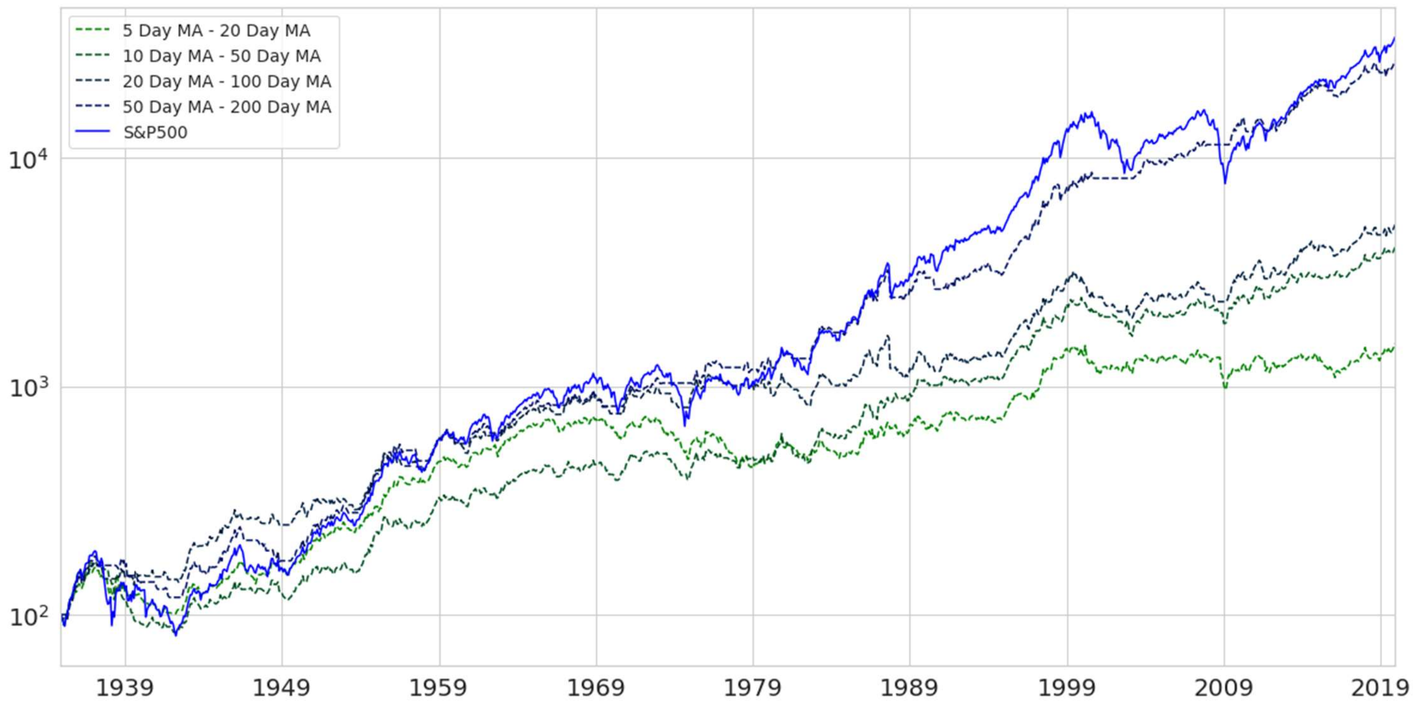

Figure 2. Log equity curve of strategies from 1935 through the end of 2019. Calculations performed by Kensington Asset Management, LLC

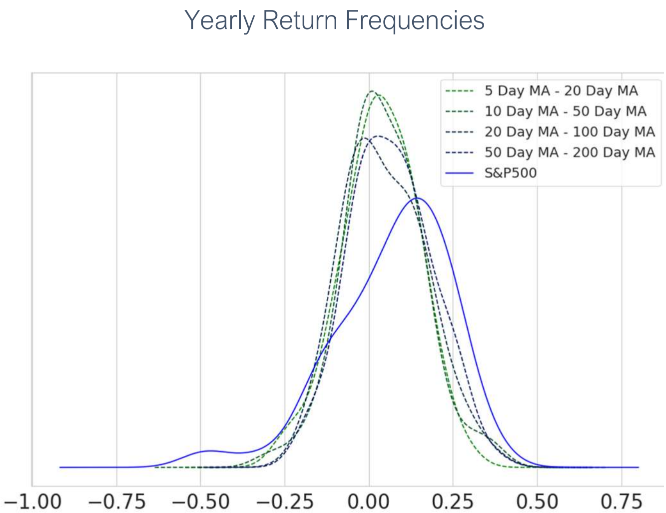

Figure 3. Frequencies of annual returns from 1935 through the end of 2019.

Fortunately, we have skew to address those concerns. As a reminder, for unimodal (single peaked) probability distributions, skewness compares the relative size of the left and right tails. In Table 1, when returns are measured on a yearly basis for trend following, they are positively skewed whereas the S&P 500's annual returns are negatively skewed. What this shows is that trend following strategies are likely to experience less large, negative returns relative to their own large positive returns, while the opposite holds true for the S&P 500. Figure 3 helps provide a visual representation of this effect: we can clearly see that for the S&P 500's annual return distribution, the left tail extends out further than the right tail. Moreover, its left tail falls above the left tails of the trend following return distributions. The key takeaway from the skewness measures is that a buy and hold strategy of the S&P 500 exposes investors to more asymmetric downside risk than does a trend following strategy. The tradeoff here is higher returns and a higher Sharpe ratio. The average up and down returns by year and month, alongside the maximum drawdowns recorded in Table 1 additionally reflect this fact.

Positive Convexity leads to Skew

For a further understanding of how positive skew is realized, we can look at how trend following strategies mechanically create nonlinear optionlike payoffs (Bruder & Gaussel, 2011). Figure 4 shows annual returns of each trend following strategy plotted against annual returns of the S&P 500, demonstrating a positive convex (upward curved) relationship. For reference, a strategy that performs linearly with the same exact return occurrences as the S&P 500 would have returns plotted on the 45-degree line. Points falling above that line would show a year of outperformance, and points falling below that line would indicate the opposite. In Figure 4, when the S&P 500 experiences extreme negative annual returns, trend following outperforms the S&P 500 the most.

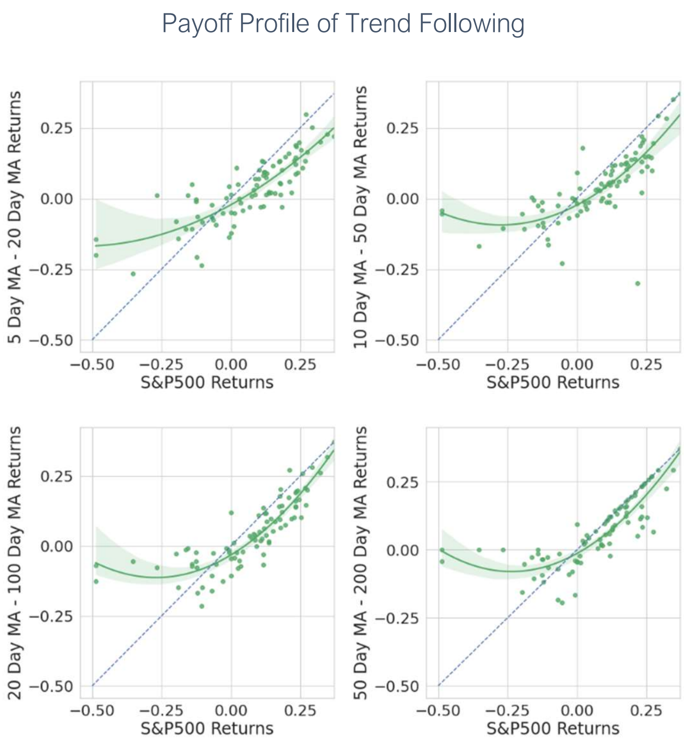

Figure 4. Annual returns of trend following strategies plotted against returns of the S&P500. Calculations performed by Kensington Asset Management, LLC

However, when the S&P 500 doesn't experience extreme negative returns, naïve trend following does not seem to compare as favorably. In Figure 4 there are instances when the S&P 500 is slightly negative but experiences smaller losses than trend following. Years of small or moderate losses in the underlying asset can either arise from constant steady price decreases or large short-term price swings from higher volatility. In the latter scenario, trend following strategies can be stopped out of positions too early, locking in losses when it would have otherwise recovered or profited. Clearly, shaping the payoff profile to be convex does not come free, as it is paid for by false positive signals from reversals that happen too quickly. This can reduce the annual expected returns of trend following, which is suggested by the lower mean of yearly returns in Table 1. It is fair then to interpret these lower returns as the cost for tail risk protection.

Behavior in Volatile Environments

Given that trend following depends on persistence in prices, it is useful to observe how it behaves when conditioned on different volatility regimes. Figure 5 shows the mean of annual returns divided into quintiles of volatility. The first thing to note is that trend following appears to perform well relative to its underlying asset at the extreme ends of daily volatility. When low volatility of the S&P 500 persists for an extended period, steady daily gains lead to large compounded returns. On the other end of the spectrum, the highest levels of volatility typically coincide during market sell offs. This would explain why the shortest of trend following strategies performs negatively in yearly time frames, while the longer-term strategies excel. Nevertheless, even the short-term trend following strategy on average outperforms the S&P 500 during crisis years, though at a higher cost. This cost can be mitigated by being selective with time frames.

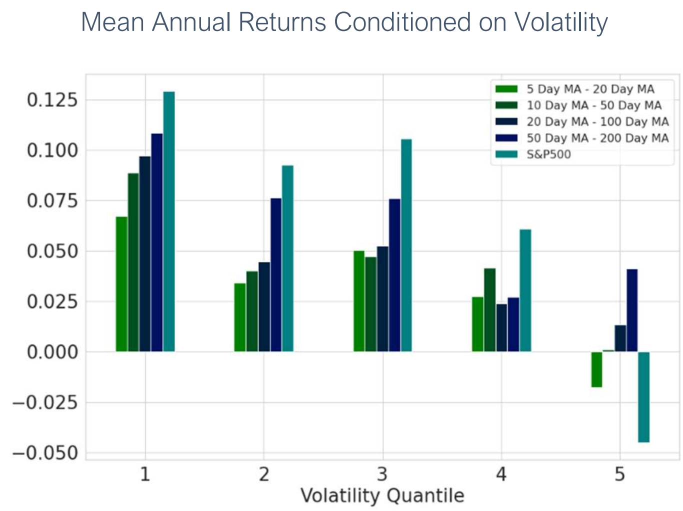

Figure 5. Mean of annual returns divided into quintiles of volatility.

Timeframes Matter

It is important to note that the convexity and skewness of trend following strategies depend on how well their time frames match the period of measurement for returns. In Table 1, we see that strategies with longer term look-back periods do not exhibit positive skew when the returns are measured monthly. This makes sense as they are constructed to capture long term trends that wouldn't appear within a shorter time frame. When measuring returns on an annual basis, all trend following strategies exhibit positive skew, with the intermediate-term strategies having the highest. The longest-term strategy doesn't have the highest skew, but it seems to have the best balance with total performance, and even has a better Sharpe ratio than the S&P 500. This highlights the importance of strategy specification in minimizing the cost of tail risk protection. The other two intermediate-term strategies may have higher skew, but their reduced returns from increased price sensitivity can be seen as overinsurance. For assets that follow longer economic cycles, longer trend following time frames and reduced turnover is warranted.

Summary

In this article, we provided an introductory overview on the key behaviors of trend following. We defined trend following and a couple of different simple systematic methods. Despite the different methods, many may in fact be mathematically related, but they all seek a similar goal. We briefly discuss the conflict between riskbased and behavioral-based explanations of trend following, but ultimately underscore the practicality of its use for tail risk management. Trend following's utility as a tail risk mitigator is shown through the positive skew of its return distribution, indicating that it experiences less extreme negative returns relative to its extreme positive returns. This positive skew is a result of mechanically shaping the payoff structure of the strategy to be positively convex through its rules, without having to rely on market inefficiencies or mispricing. The tradeoff for positive skew is the potentially lower expected returns as unfavorable volatility regimes can lead to quick price reversals. These price reversals lock in losses that would have otherwise recovered or profited and represent the "cost" of tail risk protection. Given the cost, it is important to consider the timeframes of the trend following strategies. Nevertheless, trend following strategies can be additive to a portfolio for those who want to smooth long term returns by reducing extreme losses from tail events.

References

- Bruder, B., Dao, T.-L., Richard, J.-C., & Roncalli, T. (2011). Trend Filtering Methods for Momentum Strategies. SSRN Electronic Journal. doi: 10.2139/ssrn.2289097
- Bruder, B., & Gaussel, N. (2011). Risk-Return Analysis of Dynamic Investment Strategies. SSRN Electronic Journal. doi: 10.2139/ssrn.2465623
- Jung, J., & Shiller, R. J. (2005). Samuelsons Dictum and The Stock Market. Economic Inquiry, 43(2), 221–228. doi: 10.1093/ei/cbi015
- Lo, A. W.-C., & Hasanhodzic, J. (2010). The Evolution of Technical Analysis: Financial Prediction from Babylonian Tablets to Bloomberg Terminals. Hoboken, NJ: John Wiley & Sons.
- Moskowitz, T., Ooi, Y. H., & Pedersen, L. H. (2012). Time Series Momentum. Journal of Financial Economics, 104(2), 228–250. doi: 10.1016/j.jfineco.2011.11.003.
- Huang, D., Li, J., Wang, L., & Zhou, G. (2020). Time series momentum: Is it there? Journal of Financial Economics, 135, 774–794. doi: 10.1016/j.jfineco.2019.08.004

Disclosure